



Cisco UCS Manager Server Management Guide, Release 4.1

First Published: 2020-02-20

Last Modified: 2024-08-29

Americas Headquarters

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



CONTENTS

PREFACE

Preface	xvii
Audience	xvii
Conventions	xvii
Related Cisco UCS Documentation	xix
Documentation Feedback	xix

CHAPTER 1

New and Changed Information	1
New and Changed Information	1

CHAPTER 2

Server Management Overview	3
Server Management Overview	3
Cisco UCS Manager User CLI Documentation	4
Cisco UCS Manager User Documentation	4

CHAPTER 3

Server License Management	7
Licenses	7
C-Direct Rack Licensing Support	8
Obtaining the Host ID for a Fabric Interconnect	9
Obtaining a License	10
Downloading Licenses to the Fabric Interconnect from the Local File System	11
Downloading Licenses to the Fabric Interconnect from a Remote Location	12
Installing a License	13
Viewing the Licenses Installed on a Fabric Interconnect	13
Determining the Grace Period Available for a Port or Feature	14
Determining the Expiry Date of a License	14
Uninstalling a License	14

CHAPTER 4	Registering Cisco UCS Domains with Cisco UCS Central	17
	Registration of Cisco UCS Domains	17
	Policy Resolution between Cisco UCS Manager and Cisco UCS Central	17
	Registering a Cisco UCS Domain with Cisco UCS Central	19
	Configuring Policy Resolutions between Cisco UCS Manager and Cisco UCS Central	20
	Setting Cisco UCS Central Registration Properties in Cisco UCS Manager	20
	Unregistering a Cisco UCS Domain from Cisco UCS Central	21

CHAPTER 5	Power Capping and Power Management in Cisco UCS	23
	Power Capping in Cisco UCS	23
	Power Policy Configuration	24
	Power Policy for Cisco UCS Servers	24
	Configuring the Power Policy	24
	Power Supply for Redundancy Method	25
	Configuring Policy Driven Chassis Group Power Capping	25
	Policy Driven Chassis Group Power Capping	25
	Power Control Policy	26
	Creating a Power Control Policy	27
	Deleting a Power Control Policy	29
	Acoustic Mode Fan Profile	30
	Acoustic Mode Fan Profile	30
	Configuring Acoustic Mode	30
	Power Groups in UCS Manager	32
	Creating a Power Group	34
	Adding a Chassis to a Power Group	35
	Removing a Chassis from a Power Group	36
	Deleting a Power Group	36
	Blade Level Power Capping	36
	Manual Blade Level Power Cap	36
	Setting the Blade-Level Power Cap for a Server	37
	Viewing the Blade-Level Power Cap	38
	Fan Control Policy Configuration	38
	Fan Control Policy	38

Creating a Fan Control Policy	38
Global Power Profiling Policy Configuration	39
Global Power Profiling Policy	39
Configuring the Global Power Profile Policy	39
Global Power Allocation Policy Configuration	40
Global Power Allocation Policy	40
Configuring the Global Power Allocation Policy	40
Power Management During Power-on Operations	41
Power Sync Policy Configuration	41
Power Sync Policy	41
Power Synchronization Behavior	42
Creating a Power Sync Policy	42
Changing a Power Sync Policy	44
Deleting a Power Sync Policy	44
Rack Server Power Management	45
UCS Mini Power Management	45

CHAPTER 6
Blade Server Hardware Management 47

Blade Server Management	47
Guidelines for Removing and Decommissioning Blade Servers	48
Recommendations for Avoiding Unexpected Server Power Changes	48
Booting a Blade Server	49
Booting a Rack-Mount Server from the Service Profile	50
Determining the Boot Order of a Blade Server	50
Shutting Down a Blade Server	51
Shutting Down a Server from the Service Profile	51
Resetting a Blade Server	52
Resetting a Blade Server to Factory Default Settings	52
Reacknowledging a Blade Server	53
Removing a Server from a Chassis	54
Deleting the Inband Configuration from a Blade Server	54
Decommissioning a Blade Server	55
Removing a Non-Existent Blade Server Entry	55
Recommissioning a Blade Server	56

Reacknowledging a Server Slot in a Chassis	56
Removing a Non-Existent Blade Server from the Configuration Database	57
Turning the Locator LED for a Blade Server On and Off	57
Turning the Local Disk Locator LED on a Blade Server On and Off	58
Resetting the CMOS for a Blade Server	58
Resetting the CIMC for a Blade Server	59
Clearing TPM for a Blade Server	59
Viewing the POST Results for a Blade Server	60
Issuing an NMI from a Blade Server	60
Viewing Health Events for a Blade Server	61
Health LED Alarms	63
Viewing Health LED Alarms	63
Smart SSD	63
Monitoring SSD Health	64
<hr/>	
CHAPTER 7	Rack-Mount Server Hardware Management 67
Rack-Mount Server Management	67
Rack-Enclosure Server Management	68
Guidelines for Removing and Decommissioning Rack-Mount Servers	69
Recommendations for Avoiding Unexpected Server Power Changes	69
Booting a Rack-Mount Server	70
Booting a Rack-Mount Server from the Service Profile	71
Determining the Boot Order of a Rack-Mount Server	71
Shutting Down a Rack-Mount Server	72
Shutting Down a Server from the Service Profile	72
Resetting a Rack-Mount Server	73
Resetting a Rack-Mount Server to Factory Default Settings	74
Persistent Memory Scrub	75
Reacknowledging a Rack-Mount Server	75
Deleting the Inband Configuration from a Rack-Mount Server	76
Decommissioning a Rack-Mount Server	76
Recommissioning a Rack-Mount Server	77
Renumbering a Rack-Mount Server	77
Removing a Non-Existent Rack-Mount Server from the Configuration Database	78

Turning the Locator LED for a Rack-Mount Server On and Off	79
Turning the Local Disk Locator LED on a Rack-Mount Server On and Off	79
Resetting the CMOS for a Rack-Mount Server	80
Resetting the CIMC for a Rack-Mount Server	80
Clearing TPM for a Rack-Mount Server	81
Issuing an NMI from a Rack-Mount Server	81
Viewing Health Events for a Rack-Mount Server	82
Viewing the POST Results for a Rack-Mount Server	84
Viewing the Power Transition Log	84
Viewing Cisco UCS C125 M5 Server Slot ID	85

CHAPTER 8**S3X60 Server Node Hardware Management 87**

Cisco UCS S3260 Server Node Management	87
Booting a Cisco UCS S3260 Server Node	88
Booting a Cisco UCS S3260 Server Node from the Service Profile	88
Determining the Boot Order of a Cisco UCS S3260 Server Node	89
Shutting Down a Cisco UCS S3260 Server Node	89
Shutting Down a Cisco UCS S3260 Server Node from the Service Profile	90
Resetting a Cisco UCS S3260 Server Node	90
Resetting a Cisco UCS S3260 Server Node to Factory Default Settings	91
Reacknowledging a Cisco UCS S3260 Server Node	92
Removing a Cisco UCS S3260 Server Node from a Chassis	93
Deleting the Inband Configuration from a Cisco UCS S3260 Server Node	93
Decommissioning a Cisco UCS S3260 Server Node	94
Recommissioning a Cisco UCS S3260 Server Node	94
Reacknowledging a Server Slot in a S3260 Chassis	95
Removing a Non-Existent Cisco UCS S3260 Server Node from the Configuration Database	95
Turning the Locator LED for a Cisco UCS S3260 Server Node On and Off	96
Turning the Local Disk Locator LED on a Cisco UCS S3260 Server Node On and Off	96
Resetting the CIMC for a Cisco UCS S3260 Server Node	97
Resetting the CMOS for a Cisco UCS S3260 Server Node	97
Resetting the BIOS Password for a S3X60 Server	98
Issuing an NMI from a Cisco UCS S3260 Server Node	98
Viewing the POST Results for a Cisco UCS S3260 Server Node	99

Viewing Health Events for a Cisco UCS S3260 Server Node 99

Health LED Alarms 101

Viewing Health LED Alarms 101

CHAPTER 9**Server Pools 103**

Configuring Server Pools 103

Server Pools 103

Creating a Server Pool 103

Deleting a Server Pool 104

Adding Servers to a Server Pool 104

Removing Servers from a Server Pool 105

Configuring UUID Suffix Pools 105

UUID Suffix Pools 105

Creating a UUID Suffix Pool 105

Deleting a UUID Suffix Pool 107

Configuring IP Pools 107

IP Pools 107

Creating an IP Pool 108

Adding a Block to an IP Pool 109

Deleting a Block from an IP Pool 111

Deleting an IP Pool 111

CHAPTER 10**Server Boot 113**

Boot Policy 113

UEFI Boot Mode 114

UEFI Secure Boot 115

CIMC Secure Boot 116

Determining the CIMC Secure Boot Status 117

Enabling CIMC Secure Boot on a Rack Server 117

Creating a Boot Policy 118

SAN Boot 119

Configuring a SAN Boot for a Boot Policy 119

iSCSI Boot 121

iSCSI Boot Process 121

iSCSI Boot Guidelines and Prerequisites	122
Initiator IQN Configuration	123
Enabling MPIO on Windows	124
Configuring iSCSI Boot	124
Creating an iSCSI Adapter Policy	126
Deleting an iSCSI Adapter Policy	127
Creating an iSCSI Authentication Profile	128
Deleting an iSCSI Authentication Profile	129
Creating an iSCSI Initiator IP Pool	129
Creating an iSCSI Boot Policy	130
Creating an iSCSI vNIC for a Service Profile	131
Deleting an iSCSI vNIC from a Service Profile	133
Setting the Initiator IQN at the Service Profile Level	133
Changing the Initiator IQN at the Service Profile Level	134
Setting iSCSI Boot Parameters	134
Modifying iSCSI Boot Parameters	138
IQN Pools	142
Creating an IQN Pool	142
Adding a Block to an IQN Pool	143
Deleting a Block from an IQN Pool	144
Deleting an IQN Pool	145
LAN Boot	145
Configuring a LAN Boot for a Boot Policy	145
Local Devices Boot	146
Configuring a Local Disk Boot for a Boot Policy	147
Configuring a Virtual Media Boot for a Boot Policy	148
Configuring a NVMe Boot for a Boot Policy	150
Adding a Boot Policy to a vMedia Service Profile	150
Deleting a Boot Policy	152
UEFI Boot Parameters	153
Guidelines and Limitations for UEFI Boot Parameters	153
Setting UEFI Boot Parameters	153
Modifying UEFI Boot Parameters	154

CHAPTER 11**Service Profiles 157**

- Service Profiles in UCS Manager 157
- Service Profiles that Override Server Identity 158
- Service Profiles that Inherit Server Identity 158
- Guidelines and Recommendations for Service Profiles 159
- Methods of Creating Service Profiles 159
 - Creating a Service Profile with the Expert Wizard 159
 - Creating a Service Profile that Inherits Server Identity 161
 - Creating a Hardware Based Service Profile for a Blade Server 161
 - Creating a Hardware Based Service Profile for a Rack-Mount Server 162
- Inband Service Profiles 163
 - Deleting the Inband Configuration from a Service Profile 163
- Service Profile Tasks 163
 - Renaming a Service Profile 163
 - Cloning a Service Profile 164
 - Changing the UUID in a Service Profile 165
 - Modifying the Boot Order in a Service Profile 166
 - Creating a vNIC for a Service Profile 168
 - Deleting a vNIC from a Service Profile 168
 - Creating a vHBA for a Service Profile 169
 - Changing the WWPN for a vHBA 169
 - Clearing Persistent Binding for a vHBA 170
 - Deleting a vHBA from a Service Profile 170
 - Adding a vHBA Initiator Group to a Service Profile 170
 - Deleting a Service Profile 172
- Service Profile Association 173
 - Associating a Service Profile with a Server or Server Pool 173
 - Disassociating a Service Profile from a Server or Server Pool 174
- Service Profile Templates 174
 - Initial and Existing Templates 174
 - Creating a Service Profile Template 175
 - Creating One or More Service Profiles from a Service Profile Template 176
 - Creating a Template Based Service Profile for a Blade Server 176

Creating a Template Based Service Profile for a Rack-Mount Server	177
Creating a Service Profile Template from a Service Profile	178
Setting an Asset Tag for a Service Profile	178
Service Profile Template Tasks	179
Binding a Service Profile to a Service Profile Template	179
Unbinding a Service Profile from a Service Profile Template	179
Changing the UUID in a Service Profile Template	180
Resetting the UUID Assigned to a Service Profile from a Pool in a Service Profile Template	180
Resetting the MAC Address Assigned to a vNIC from a Pool in a Service Profile	181
Resetting the WWPN Assigned to a vHBA from a Pool in a Service Profile Template	182
Deleting the Inband Configuration from a Service Profile Template	183
Service Profile Association	183
Associating a Service Profile with a Server or Server Pool	183
Associating a Service Profile Template with a Server Pool	184
Disassociating a Service Profile from a Server or Server Pool	185
Disassociating a Service Profile Template from its Server Pool	185

CHAPTER 12
Server-Related Policies 187

BIOS Settings	187
Server BIOS Settings	187
Server BIOS Settings	187
Intel Speed Select	261
BIOS Policy	263
Default BIOS Settings	263
Creating a BIOS Policy	264
Modifying the BIOS Defaults	265
Viewing the Actual BIOS Settings for a Server	266
Memory RAS Features	266
Post-Package Repair (PPR)	266
Enabling Post Package Repair	267
Limiting Presented Memory	267
Limiting Memory Size	267
Partial Memory Mirroring	268
Enabling Partial Memory Mirroring	268

Trusted Platform Module	269
Trusted Platform Module	269
Intel Trusted Execution Technology	270
Configuring Trusted Platform	270
Configuring Trusted Platform	271
Viewing TPM Properties	272
Consistent Device Naming	272
Guidelines and Limitations for Consistent Device Naming (CDN)	272
Configuring Consistent Device Naming in a BIOS Policy	274
Configuring a CDN Name for a vNIC	275
CIMC Security Policies	275
IPMI Access Profile	275
Creating an IPMI Access Profile	276
Deleting an IPMI Access Profile	277
KVM Management Policy	277
Creating a KVM Management Policy	278
Graphics Card Policies	278
Creating a Graphics Card Policy	279
Local Disk Policies	279
Local Disk Configuration Policy	279
Guidelines for all Local Disk Configuration Policies	280
Guidelines for Local Disk Configuration Policies Configured for RAID	281
Creating a Local Disk Configuration Policy	282
Changing a Local Disk Configuration Policy	285
Deleting a Local Disk Configuration Policy	285
FlexFlash Support	286
FlexFlash FX3S Support	288
Starting Up Blade Servers with FlexFlash SD Cards	289
Enabling FlexFlash SD Card Support	290
Enabling Auto-Sync	290
Formatting the SD Cards	291
Resetting the FlexFlash Controller	291
Persistent Memory Modules	291
Scrub Policy	292

Scrub Policy Settings	292
Creating a Scrub Policy	294
Deleting a Scrub Policy	295
DIMM Error Management	295
DIMM Correctable Error Handling	295
Resetting Memory Errors	295
DIMM Blacklisting	296
Enabling DIMM Blacklisting	296
Serial over LAN Policy Settings	297
Serial over LAN Policy Overview	297
Creating a Serial over LAN Policy	297
Deleting a Serial over LAN Policy	298
Server Autoconfiguration Policies	299
Server Autoconfiguration Policy Overview	299
Creating an Autoconfiguration Policy	299
Deleting an Autoconfiguration Policy	300
Server Discovery Policy Settings	301
Server Discovery Policy Overview	301
Creating a Server Discovery Policy	301
Deleting a Server Discovery Policy	302
Hardware Change Discovery Policy	302
Configuring Hardware Change Discovery Policy	303
Server Inheritance Policy Settings	303
Server Inheritance Policy Overview	303
Creating a Server Inheritance Policy	303
Deleting a Server Inheritance Policy	304
Server Pool Policy Settings	305
Server Pool Policy Overview	305
Creating a Server Pool Policy	305
Deleting a Server Pool Policy	306
Server Pool Policy Qualifications Settings	306
Server Pool Policy Qualification Overview	306
Creating Server Pool Policy Qualifications	307
Deleting Server Pool Policy Qualifications	311

- Deleting Qualifications from Server Pool Policy Qualifications 312
- vNIC/vHBA Placement Policy Settings 312
 - vNIC/vHBA Placement Policies 312
 - vCon to Adapter Placement 313
 - For N20-B6620-2 and N20-B6625-2 Blade Servers 314
 - vCon to Adapter Placement for All Other Supported Servers 314
 - vNIC/vHBA to vCon Assignment 315
 - Creating a vNIC/vHBA Placement Policy 317
 - Deleting a vNIC/vHBA Placement Policy 319
 - Explicitly Assigning a vNIC to a vCon 319
 - Explicitly Assigning a vHBA to a vCon 321
 - Placing Static vNICs Before Dynamic vNICs 322
 - vNIC/vHBA Host Port Placement 324
 - Configuring Host Port Placement 324
- CIMC Mounted vMedia 325
 - Creating a vMedia Policy 326
 - Adding a vMedia Policy to a Service Profile 330
 - Viewing CIMC vMedia Policy 332

CHAPTER 13

Firmware Upgrades 333

- Firmware Upgrades 333
- Verifying Firmware Versions on Components 333

CHAPTER 14

Diagnostics Configuration 335

- Overview of Cisco UCS Manager Diagnostics 335
- Creating a Diagnostics Policy 335
- Diagnostics Test on a Blade Server 336
 - Starting a Diagnostics Test on a Blade Server 336
 - Stopping a Diagnostics Test on a Blade Server 337
- Diagnostics Test on a Rack Server 337
 - Starting a Diagnostics Test on a Rack Server 337
 - Stopping a Diagnostics Test on a Rack Server 338
- Starting a Diagnostics Tests on All Servers 338
- Stopping a Diagnostics Tests on All Servers 339

Viewing the Server Diagnostics Status/Result	339
Diagnostics Troubleshooting	340



Preface

- [Audience, on page xvii](#)
- [Conventions, on page xvii](#)
- [Related Cisco UCS Documentation, on page xix](#)
- [Documentation Feedback, on page xix](#)

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 <i>this font</i> .
TUI elements	In a Text-based User Interface, text the system displays appears in <code>this font</code> .
System output	Terminal sessions and information that the system displays appear in <code>this font</code> .
CLI commands	CLI command keywords appear in this font . Variables in a CLI command appear in <i>this font</i> .
[]	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 Roadmapdoc 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](#).

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](#) , on page 1

New and Changed Information

This section provides information on new feature and changed behavior in Cisco UCS Manager, Release 4.1(2b).

Table 1: New Features and Changed Behavior in Cisco UCS Manager, Release 4.1(2b)

Feature	Description	Where Documented
New RAS Memory BIOS Settings	Cisco UCS Manager now has the following BIOS tokens to mitigate Rowhammer-style attacks: <ul style="list-style-type: none">• Memory Refresh Rate• Panic and High Watermark	RAS Memory BIOS Settings, on page 222

This section provides information on new feature and changed behavior in Cisco UCS Manager, Release 4.1(2a).

Table 2: New Features and Changed Behavior in Cisco UCS Manager, Release 4.1(2a)

Feature	Description	Where Documented
Cisco UCS C240 SD M5 Server	Cisco UCS Manager provides support for rack server power management, FlexFlash FX3S support, and scrub policy on Cisco UCS C240 SD M5 Server.	<ul style="list-style-type: none">• FlexFlash FX3S Support, on page 288• Rack Server Power Management, on page 45

Feature	Description	Where Documented
New Processor BIOS Options	<p>Cisco UCS Manager now has BIOS tokens for Configurable TDP Level, External SSC Enable, PCIe PLL SSC, UPI Link Frequency Select, and Uncore Frequency Scaling. It also adds the following DCPMM Memory BIOS tokens:</p> <ul style="list-style-type: none"> • CR QoS • NVM Performance Setting • CR FastGo Config • Snoopy mode for AD • Snoopy mode for 2LM 	<ul style="list-style-type: none"> • Processor BIOS Settings, on page 190 • LOM and PCIe Slots BIOS Settings, on page 239 • Intel® Optane™ DC Persistent Memory (DCPMM) BIOS Tokens, on page 229
Acoustic Mode	<p>Reduces the fan speed to reduce noise levels in acoustic-sensitive environments in Cisco UCS C220 M5 Server, C240 M5 and C240 SD M5 Rack servers.</p>	<ul style="list-style-type: none"> • Acoustic Mode Fan Profile, on page 30

This section provides information on new feature and changed behavior in Cisco UCS Manager, Release 4.1(1a).

Table 3: New Features and Changed Behavior in Cisco UCS Manager, Release 4.1(1a)

Feature	Description	Where Documented
RAS Memory Features	<p>Cisco UCS Manager expands on Intel® Xeon® processor support for RAS memory features via the BIOS. These features expand on the capabilities of the processor to increase the performance and reliability of memory DIMMs.</p> <p>These include limiting presented memory to a subset of overall memory, ability to create a mirrored region in memory map with the attributes of a partial mirror copy, rather than the complete mirror copy.</p>	<p>Memory RAS Features, on page 266</p>



CHAPTER 2

Server Management Overview

- [Server Management Overview, on page 3](#)
- [Cisco UCS Manager User CLI Documentation, on page 4](#)
- [Cisco UCS Manager User Documentation, on page 4](#)

Server Management Overview

Cisco UCS Manager enables you to manage general and complex server deployments. For example, you can manage a general deployment with a pair of Fabric Interconnects (FIs), which is the redundant server access layer that you get with the first chassis that can scale up to 20 chassis' and up to 160 physical servers. This can be a combination of blades and rack mount servers to support the workload in your environment. As you add more servers, you can continue to perform server provisioning, device discovery, inventory, configuration, diagnostics, monitoring, fault detection, and auditing.

- Cisco UCS C220 M6 Server
- Cisco UCS C240 M6 Server
- Cisco UCS B200 M6 Server
- Cisco UCS VIC 1467 (mLOM)
- Cisco UCS VIC 1477 (mLOM)

The ,Cisco UCS 6332 Fabric Interconnects, Cisco UCS Mini 6324 Fabric Interconnects, and Cisco UCS 6200 Series Fabric Interconnects include centralized management. You can manage the UCS Blade Servers and Rack-Mount Servers that are in the same domain from one console. You can also manage the UCS Mini from the Cisco UCS Manager.

To ensure the optimum server performance, you can configure the amount of power that you allocate to servers. You can also set the server boot policy, the location from which the server boots, and the order in which the boot devices are invoked. You can create service profiles and assign the service profiles to servers. In service profile, you can configure vNICs and vHBAs, enables BIOS settings, apply firmware policy, and other settings. When the service profile is associated to a server, the configured configurations, policies, and settings are pushed to the server.

Cisco UCS Manager User CLI Documentation

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

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

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.
Cisco UCS Manager Server Management Guide	Discusses the new licenses, registering Cisco UCS domains with Cisco UCS Central, power capping, server boot, server profiles and server-related policies.
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.



CHAPTER 3

Server License Management

- [Licenses, on page 7](#)
- [C-Direct Rack Licensing Support, on page 8](#)
- [Obtaining the Host ID for a Fabric Interconnect, on page 9](#)
- [Obtaining a License, on page 10](#)
- [Downloading Licenses to the Fabric Interconnect from the Local File System, on page 11](#)
- [Downloading Licenses to the Fabric Interconnect from a Remote Location, on page 12](#)
- [Installing a License, on page 13](#)
- [Viewing the Licenses Installed on a Fabric Interconnect, on page 13](#)
- [Determining the Grace Period Available for a Port or Feature, on page 14](#)
- [Determining the Expiry Date of a License, on page 14](#)
- [Uninstalling a License, on page 14](#)

Licenses

Each Cisco UCS 6324 Fabric Interconnect comes with a factory installed port license that is shipped with the hardware. The C-direct port license is factory installed with a grace period, measured from first use of the port, and can be used for Cisco UCS rack servers. If multiple ports are acting within grace periods, the license is moved to the port whose grace period is closest to expiring.

Grace Period

If you attempt to use a port that does not have an installed license, Cisco UCS initiates a 120 day grace period. The grace period is measured from the first use of the port without a license and is paused when a valid license file is installed. The amount of time used in the grace period is retained by the system.



Note Each physical port has its own grace period. Initiating the grace period on a single port does not initiate the grace period for all ports.

If a licensed port is unconfigured, that license is transferred to a port functioning within a grace period. If multiple ports are acting within grace periods, the license is moved to the port whose grace period is closest to expiring.

High Availability Configurations

To avoid inconsistencies during failover, we recommend that both Fabric Interconnects in the cluster have the same number of ports licensed. If symmetry is not maintained and failover occurs, Cisco UCS enables the missing licenses and initiates the grace period for each port being used on the failover node.

C-Direct Rack Licensing Support

Release 4.1(1a) and Higher

Beginning with release 4.1(1a), Cisco UCS 64108 Fabric Interconnects use the `ETH_C_PORT_ACTIVATION_PKG` feature pack for C-Direct port licenses for ports 1-96. There are no default `ETH_C_PORT_ACTIVATION_PKG` licenses shipped with the Fabric Interconnect. You may purchase them as required.

C-direct support is only applicable on ports that are connected to the rack servers. The `ETH_C_PORT_ACTIVATION_PKG` is added to the existing license package with all the same properties as the existing licensing feature. The Subordinate Quantity property is added to the `ETH_PORT_ACTIVATION_PKG` to track ports connected to rack servers.

The License Tab in the Cisco UCS Manager GUI displays the new license and the **Subordinate Quantity** for the license. You can also use the **show feature** and **show usage** commands under **scope license** to view the license feature, the vendor version type, and the grace period for each license.

Release 4.0(1a) and Higher

Beginning with release 4.0(1a), Cisco UCS 6454 Fabric Interconnects use the `ETH_C_PORT_ACTIVATION_PKG` feature pack for C-Direct port licenses for ports 1-48. There are no default `ETH_C_PORT_ACTIVATION_PKG` licenses shipped with the Fabric Interconnect. You may purchase them as required.

C-direct support is only applicable on ports that are connected to the rack servers. The `ETH_C_PORT_ACTIVATION_PKG` is added to the existing license package with all the same properties as the existing licensing feature. The Subordinate Quantity property is added to the `ETH_PORT_ACTIVATION_PKG` to track ports connected to rack servers.

The License Tab in the Cisco UCS Manager GUI displays the new license and the **Subordinate Quantity** for the license. You can also use the **show feature** and **show usage** commands under **scope license** to view the license feature, the vendor version type, and the grace period for each license.

Release 3.2(3o) and Earlier

Each Cisco UCS Fabric Interconnect is shipped with a default number of port licenses that are factory licensed and shipped with the hardware. C-direct support is only applicable on ports that are connected to the rack servers. The `10G_C_PORT_ACTIVATION_PKG` and the `40G_ETH_C_PORT_ACTIVATION_PKG` are added to the existing license package with all the same properties as the existing licensing feature. The **Subordinate Quantity** property is added to the `10G_PORT_ACTIVATION_PKG` and `40G_ETH_PORT_ACTIVATION_PKG` to track ports connected to rack servers.

The License Tab in the Cisco UCS Manager GUI displays the new license and the **Subordinate Quantity** for the license. You can also use the **show feature** and **show usage** commands under **scope license** to view the license feature, the vendor version type, and the grace period for each license.

Ports connected to rack servers can use existing 10G_PORT_ACTIVATION_PKG, 40G_ETH_PORT_ACTIVATION_PKG if the license is available or if the license is not in use. Otherwise, you must purchase a 10G_C_PORT_ACTIVATION_PKG, the 40G_ETH_C_PORT_ACTIVATION_PKG to avoid the license grace period.

There is no change in the 10 GB ports. The 10G_PORT_ACTIVATION_PKG and 10G_C_PORT_ACTIVATION_PKG license packages include all of the same properties as the existing the ETH_PORT_ACTIVATION_PKG and the ETH_PORT_C_ACTIVATION_PKG license features.

Configuration and Restrictions

- The C-Direct rack licensing feature accounts for the rack server ports that are directly connected to the FI, but not to a CIMC port. The default quantity for the 10G_C_PORT_ACTIVATION_PKG and the 40G_ETH_C_PORT_ACTIVATION_PKG is always 0.
- When a 40 GB port, or a breakout port under a 40 GB breakout port is enabled without any connections, this port is allotted a license under the 40G_ETH_PORT_ACTIVATION_PKG, if available. If this port is connected to a Direct-Connect rack server after a time lag, it triggers a complete re-allocation of licenses, then this port passes through one of the following license allocation scenarios occurs:

When you enable a breakout port under a 40 GB breakout port, if that port is connected to a Direct-Connect rack server, and the 40G_C_PORT_ACTIVATION_PKG license files are installed on the FI, the following license allocation occurs:

- If no other ports under the breakout port are enabled, the parent 40 GB port is allotted a license under the 40G_C_PORT_ACTIVATION_PKG, and the used quantity is incremented for this instance.
 - If other ports are enabled, and if at least one port is not connected to a Direct Connect rack server, even if the port is not being used, the parent 40 GB port is allotted a license under the 40G_ETH_PORT_ACTIVATION_PKG, and the used quantity is incremented for this instance.
- When you enable a breakout port under a 40 GB breakout port and that port is connected to a Direct-Connect rack server, and the 40G_C_PORT_ACTIVATION_PKG license files are not installed on the FI, the following license allocation occurs:
 - If no ports under the breakout port are enabled, the parent 40 GB port is allotted a license under the 40G_ETH_PORT_ACTIVATION_PKG. The subordinate quantity is increased if the licenses are available in the 40G_ETH_PORT_ACTIVATION_PKG. If the licenses are not available, the used quantity under this feature is increased and the entire port goes in to the grace period.
 - If other ports are enabled and at least one port is not connected to a Direct Connect rack server, even if the port is not being used, the parent 40 GB port is allotted a license under the 40G_ETH_PORT_ACTIVATION_PKG, and the used quantity is incremented for this instance.

Obtaining the Host ID for a Fabric Interconnect

The host ID is also known as the serial number.

Procedure

-
- Step 1** In the **Navigation** pane, click **Equipment**.
 - Step 2** Expand **Equipment > Fabric Interconnects**.
 - Step 3** Click the node for the fabric interconnect for which you want to obtain the host ID.
 - Step 4** In the **Work** pane, click the **General** tab.
 - Step 5** In the **Properties** area, the host ID is listed in the **Serial Number (SN)** field.
-

What to do next

Obtain the required licenses from Cisco.

Obtaining a License



Note This process may change after the release of this document. If one or more of these steps no longer applies, contact your Cisco representative for information on how to obtain a license file.

Before you begin

Obtain the following:

- Host ID or serial number for the fabric interconnect
- Claim certificate or other proof of purchase document for the fabric interconnect or expansion module

Procedure

-
- Step 1** Obtain the product authorization key (PAK) from the claim certificate or other proof of purchase document.
 - Step 2** Locate the website URL in the claim certificate or proof of purchase document.
 - Step 3** Access the website URL for the fabric interconnect and enter the serial number and the PAK.

Cisco sends you the license file by email. The license file is digitally signed to authorize use on only the requested fabric interconnect. The requested features are also enabled once Cisco UCS Manager accesses the license file.

What to do next

Install the license on the fabric interconnect.

Downloading Licenses to the Fabric Interconnect from the Local File System



Note In a cluster setup, Cisco recommends that you download and install licenses to both fabric interconnects in matching pairs. An individual license is only downloaded to the fabric interconnect that is used to initiate the download.

This procedure is specific to fabric interconnects using port-based licenses and does not apply to fabric interconnects with perpetual license configurations.

Before you begin

Obtain the required licenses from Cisco.

Procedure

- Step 1** In the **Navigation** pane, click **Admin**.
- Step 2** Expand **All > License Management**.
- Step 3** Click the node for the fabric interconnect to which you want to download the license.
- Step 4** In the **Work** pane, click the **Download Tasks** tab.
- Step 5** Click **Download License**.
- Step 6** In the **Download License** dialog box, click the **Local File System** radio button in the **Location of the Image File** field.
- Step 7** In the **Filename** field, type the full path and name of the license file.

You cannot have spaces anywhere in the path name or the file name. For example, `c:\Path\Folder_Name\License.lic` is a valid path, but `c:\Path\Folder Name\License.lic` is invalid due to the space in "Folder Name".

If you do not know the exact path to the folder where the license file is located, click **Browse** and navigate to the file.
- Step 8** Click **OK**.

Cisco UCS Manager GUI begins downloading the license to the fabric interconnect.
- Step 9** (Optional) Monitor the status of the download on the **Download Tasks** tab.

Note
If Cisco UCS Manager reports that the bootflash is out of space, delete obsolete bundles on the **Packages** tab to free up space. To view the available space in bootflash, navigate to the fabric interconnect, click **Equipment**, and expand the **Local Storage Information** area on the **General** tab.
- Step 10** Repeat this task until all the required licenses have been downloaded to the fabric interconnect.

What to do next

After all of the download tasks complete, install the licenses.

Downloading Licenses to the Fabric Interconnect from a Remote Location



Note In a cluster setup, Cisco recommends that you download and install licenses to both fabric interconnects in matching pairs. An individual license is only downloaded to the fabric interconnect that is used to initiate the download.

This procedure is specific to fabric interconnects using port-based licenses and does not apply to fabric interconnects with perpetual license configurations.

Before you begin

Obtain the required licenses from Cisco.

Procedure

-
- Step 1** In the **Navigation** pane, click **Admin**.
- Step 2** Expand **All > License Management**.
- Step 3** Click the node for the fabric interconnect to which you want to download the license.
- Step 4** In the **Work** pane, click the **Download Tasks** tab.
- Step 5** Click **Download License**.
- Step 6** In the **Download License** dialog box, click the **Remote File System** radio button in the **Location of the Image File** field.
- Step 7** Specify the protocol, and enter the required information.
- You cannot have spaces anywhere in the path name or the file name. For example, `c:\Path\Folder_Name\License.lic` is a valid path, but `c:\Path\Folder Name\License.lic` is invalid due to the space in "Folder Name".
- Note** If you use a hostname rather than an IPv4 or IPv6 address, you must configure a DNS server. If the Cisco UCS domain is not registered with Cisco UCS Central or DNS management is set to **local**, configure a DNS server in Cisco UCS Manager. If the Cisco UCS domain is registered with Cisco UCS Central and DNS management is set to **global**, configure a DNS server in Cisco UCS Central.
- Step 8** Click **OK**.
- Cisco UCS Manager GUI begins downloading the license to the fabric interconnect.
- Step 9** (Optional) Monitor the status of the download on the **Download Tasks** tab.

Note

If Cisco UCS Manager reports that the bootflash is out of space, delete obsolete bundles on the **Packages** tab to free up space. To view the available space in bootflash, navigate to the fabric interconnect, click **Equipment**, and expand the **Local Storage Information** area on the **General** tab.

Step 10 Repeat this task until all the required licenses have been downloaded to the fabric interconnect.

What to do next

After all of the download tasks complete, install the licenses.

Installing a License

Before you begin

Obtain the required licenses from Cisco.

Procedure

- Step 1** In the **Navigation** pane, click **Admin**.
- Step 2** Expand **All > License Management**.
- Step 3** In the **Work** pane, click the **Downloaded License Files** tab.
- Step 4** Choose the license you want to install from the table.

Note

There is no downtime required or impact to traffic when installing a new port license.

- Step 5** Click the **Install License** button.
 - Step 6** In the **Install License** dialog box, click **Yes**.
- Cisco UCS Manager GUI installs the license and activates the unlicensed port or feature.
-

Viewing the Licenses Installed on a Fabric Interconnect

Procedure

- Step 1** In the **Navigation** pane, click **Admin**.
- Step 2** Expand **All > License Management**.
- Step 3** In the **Work** pane, click the **Installed Licenses** tab to view the details of all licenses installed on the fabric interconnect.

- Step 4** Click a license in the table to view the details of that license in the **Contents** tab.
You may need to expand the license file to view the details of individual licenses in the file.
-

Determining the Grace Period Available for a Port or Feature

Procedure

- Step 1** In the **Navigation** pane, click **Admin**.
Step 2 Expand **All > License Management**.
Step 3 In the **Work** pane, click the **General** tab.
Step 4 Click a feature in the table to view details for that feature, including the operational state and used grace period.
-

Determining the Expiry Date of a License

Procedure

- Step 1** In the **Navigation** pane, click **Admin**.
Step 2 Expand **All > License Management**.
Step 3 In the **Work** pane, click the **Installed Licenses** tab.
Step 4 Click a license in the table to view the details of that license in the **Contents** tab below.
Step 5 In the **Contents** tab, expand the license file to view all licenses in the file.
Step 6 In the **Expiry** column, view the expiry date of the license.
-

Uninstalling a License



Note Permanent licenses cannot be uninstalled if they are in use. You can only uninstall a permanent license that is not in use. If you try to delete a permanent license that is being used, Cisco UCS Manager rejects the request and display an error message.

Before you begin

Back up the Cisco UCS Manager configuration.

Procedure

-
- Step 1** In the **Navigation** pane, click **Admin**.
 - Step 2** Expand **All > License Management**.
 - Step 3** In the **Work** pane, click the **Installed Licenses** tab.
 - Step 4** Choose the license you want to uninstall from the table.
 - Step 5** Click the **Clear License** button.
 - Step 6** If a confirmation dialog box displays, click **Yes**.
-

Cisco UCS Manager deactivates the license, removes the license from the list of licenses, and deletes the license from the fabric interconnect. The port is moved into unlicensed mode. In a cluster setup, you must uninstall the license from the other fabric interconnect.



CHAPTER 4

Registering Cisco UCS Domains with Cisco UCS Central

- [Registration of Cisco UCS Domains, on page 17](#)
- [Policy Resolution between Cisco UCS Manager and Cisco UCS Central, on page 17](#)
- [Registering a Cisco UCS Domain with Cisco UCS Central, on page 19](#)
- [Configuring Policy Resolutions between Cisco UCS Manager and Cisco UCS Central , on page 20](#)
- [Setting Cisco UCS Central Registration Properties in Cisco UCS Manager, on page 20](#)
- [Unregistering a Cisco UCS Domain from Cisco UCS Central, on page 21](#)

Registration of Cisco UCS Domains

You can have Cisco UCS Central manage some or all of the Cisco UCS domains in your data center.

If you want Cisco UCS Central to manage a Cisco UCS domain, you need to register that domain. When you register, you must choose which types of policies and other configurations will be managed by Cisco UCS Central and Cisco UCS Manager. Cisco UCS Central can manage the same types of policies and configurations for all registered Cisco UCS domains. You can also choose to have different settings for each registered Cisco UCS domain.

Perform the following before registering a Cisco UCS domain with Cisco UCS Central:

- Configure an NTP server and the correct time zone in both Cisco UCS Manager and Cisco UCS Central to ensure that they are in sync. If the time and date in the Cisco UCS domain and Cisco UCS Central are out of sync, the registration might fail.
- Obtain the hostname or IP address of Cisco UCS Central
- Obtain the shared secret that was configured when Cisco UCS Central was deployed.

Policy Resolution between Cisco UCS Manager and Cisco UCS Central

For each Cisco UCS domain that you register with Cisco UCS Central, you can choose which application will manage certain policies and configuration settings. This policy resolution does not have to be the same for every Cisco UCS domain that you register with the same Cisco UCS Central.



Note Unregistering a Cisco UCS domain with Cisco UCS Central will terminate all open sessions.

You have the following options for resolving these policies and configuration settings:

- **Local**—The policy or configuration is determined and managed by Cisco UCS Manager.
- **Global**—The policy or configuration is determined and managed by Cisco UCS Central.

The following table contains a list of the policies and configuration settings that you can choose to have managed by either Cisco UCS Manager or Cisco UCS Central:

Name	Description
Infrastructure & Catalog Firmware	Determines whether the Capability Catalog and infrastructure firmware policy are defined locally or come from Cisco UCS Central.
Time Zone Management	Determines whether the date and time is defined locally or comes from Cisco UCS Central.
Communication Services	Determines whether HTTP, CIM XML, Telnet, SNMP, web session limits, and Management Interfaces Monitoring Policy settings are defined locally or in Cisco UCS Central.
Global Fault Policy	Determines whether the Global Fault Policy is defined locally or in Cisco UCS Central.
User Management	Determines whether authentication and native domains, LDAP, RADIUS, TACACS+, trusted points, locales, and user roles are defined locally or in Cisco UCS Central.
DNS Management	Determines whether DNS servers are defined locally or in Cisco UCS Central.
Backup & Export Policies	Determines whether the Full State Backup Policy and All Configuration Export Policy are defined locally or in Cisco UCS Central.
Monitoring	Determines whether Call Home, Syslog, and TFTP Core Exporter settings are defined locally or in Cisco UCS Central.
SEL Policy	Determines whether managed endpoints are defined locally or in Cisco UCS Central.
Power Allocation Policy	Determines whether the Global Power Allocation Policy is defined locally or in Cisco UCS Central.
Powerextended Policy	Determines whether the Powerextended Policy is defined locally or in Cisco UCS Central.
Powersave Policy	Determines whether the Powersave Policy is defined locally or in Cisco UCS Central.
Modular Chassis Fan Control Policy	Determines whether the Modular Chassis Fan Control Policy is defined locally or in Cisco UCS Central.

Name	Description
Power Policy	Determines whether the Power Policy is defined locally or in Cisco UCS Central.
Equipment Policy	Determines whether Equipment Global, LAN Cloud, and SAN Cloud Policies are defined locally or in Cisco UCS Central.
Port Configuration	Determines whether port configuration is defined locally or in Cisco UCS Central.

Registering a Cisco UCS Domain with Cisco UCS Central

Before you begin

Configure an NTP server and the correct time zone in both Cisco UCS Manager and Cisco UCS Central to ensure that they are in sync. If the time and date in the Cisco UCS domain and Cisco UCS Central are out of sync, the registration might fail.

Procedure

-
- Step 1** In the **Navigation** pane, click **Admin**.
- Step 2** Expand **All > Communication Management**.
- Step 3** Click the **UCS Central** node.
- Step 4** In the **Actions** area, click **UCS Central**.
- Step 5** In the **Actions** area, click **Register With UCS Central**.
- Step 6** In the **Register with UCS Central** dialog box, do the following:
- Complete the following fields:

Name	Description
Hostname/IP Address field	The hostname or IP address of the virtual machine where Cisco UCS Central is deployed. Note If you use a hostname rather than an IPv4 or IPv6 address, you must configure a DNS server. If the Cisco UCS domain is not registered with Cisco UCS Central or DNS management is set to local , configure a DNS server in Cisco UCS Manager. If the Cisco UCS domain is registered with Cisco UCS Central and DNS management is set to global , configure a DNS server in Cisco UCS Central.
Shared Secret field	The shared secret (or password) that was configured when Cisco UCS Central was deployed.

- In the **Policy Resolution Control** area, click one of the following radio buttons for each of the fields:
 - Local**—The policy or configuration is determined and managed by Cisco UCS Manager.

- **Global**—The policy or configuration is determined and managed by Cisco UCS Central.

c) Click **OK**.

Configuring Policy Resolutions between Cisco UCS Manager and Cisco UCS Central

Procedure

- Step 1** In the **Navigation** pane, click **Admin**.
- Step 2** Expand **All > Communication Management**.
- Step 3** Click the **UCS Central** node.
- Step 4** In the **Actions** area, click **UCS Central**.
- Step 5** In the **Policy Resolution Control** area, click one of the following radio buttons for each of the fields:
- **Local**—The policy or configuration is determined and managed by Cisco UCS Manager.
 - **Global**—The policy or configuration is determined and managed by Cisco UCS Central.
- Step 6** Click **Save Changes**.
-

Setting Cisco UCS Central Registration Properties in Cisco UCS Manager

Procedure

- Step 1** In the **Navigation** pane, click **Admin**.
- Step 2** Expand **All > Communication Management**.
- Step 3** Click the **UCS Central** node.
- Step 4** In the **Actions** area, click **UCS Central**.
- Step 5** In the **Status** area, complete the following as appropriate:
- a) Click the radio button for the **Cleanup Mode** that you want to use.
- This can be one of the following:
- **Localize Global**—When a Cisco UCS domain is unregistered, all global policies in the Cisco UCS domain will be localized to Cisco UCS Manager. The policies remain in the Cisco UCS domain,

policy ownership is now local to Cisco UCS Manager, and Cisco UCS Manager admin users can make changes.

Note

If you reregister the Cisco UCS domain with Cisco UCS Central, there can be policy conflicts due to the policies existing both in Cisco UCS Central and in Cisco UCS Manager. Either delete the local policies, or set the local policies to global before you try to create and associate a global service profile.

- **Deep Remove Global**—This option should only be used after careful consideration. When a Cisco UCS domain is unregistered, all global policies in the Cisco UCS domain are removed. If there are global service profiles, they will now refer to Cisco UCS Manager local default policies, and one of the following occurs:

- If there are local default policies present, the server will reboot.
- If there are no local default policies, the service profile association fails with a configuration error.

Note

The deep remove global cleanup mode does not remove global VSANs and VLANs when you unregister from Cisco UCS Central. Those must be removed manually if desired.

- b) Optionally check the **Suspend State** check box.

If checked, the Cisco UCS domain is temporarily removed from Cisco UCS Central, and all global policies revert to their local counterparts. All service profiles maintain their current identities. However, global pools are no longer visible and cannot be accessible by new service profiles.

- c) Optionally check the **Acknowledge State** check box.

If the event ID stream that represents time and consistency between Cisco UCS Manager and Cisco UCS Central becomes skewed or inconsistent, Cisco UCS Manager places itself in a Suspended State and disconnects itself from Cisco UCS Central.

If you check this check box, you acknowledge that inconsistencies exist between Cisco UCS Manager and Cisco UCS Central and are still willing to reconnect the Cisco UCS domain with Cisco UCS Central.

Step 6 Click **Save Changes**.

Unregistering a Cisco UCS Domain from Cisco UCS Central

When you unregister a Cisco UCS domain from Cisco UCS Central, Cisco UCS Manager no longer receives updates to global policies.

Procedure

- Step 1** In the **Navigation** pane, click **Admin**.
- Step 2** Expand **All > Communication Management**.

- Step 3** Click the **UCS Central** node.
- Step 4** In the **Actions** area, click **UCS Central**.
- Step 5** In the **Actions** area, click **Unregister From UCS Central**.
- Step 6** If a confirmation dialog box displays, click **Yes**.
- Step 7** Click **OK**.

For more information on the impact of unregistering and registering a Cisco UCS Domain with Cisco UCS Central, see [Policy Resolution between Cisco UCS Manager and Cisco UCS Central](#).



CHAPTER 5

Power Capping and Power Management in Cisco UCS

- [Power Capping in Cisco UCS, on page 23](#)
- [Power Policy Configuration, on page 24](#)
- [Power Policy for Cisco UCS Servers, on page 24](#)
- [Configuring the Power Policy, on page 24](#)
- [Power Supply for Redundancy Method, on page 25](#)
- [Configuring Policy Driven Chassis Group Power Capping, on page 25](#)
- [Policy Driven Chassis Group Power Capping, on page 25](#)
- [Power Control Policy, on page 26](#)
- [Acoustic Mode Fan Profile, on page 30](#)
- [Power Groups in UCS Manager, on page 32](#)
- [Blade Level Power Capping, on page 36](#)
- [Fan Control Policy Configuration, on page 38](#)
- [Global Power Profiling Policy Configuration, on page 39](#)
- [Global Power Allocation Policy Configuration, on page 40](#)
- [Power Sync Policy Configuration, on page 41](#)
- [Rack Server Power Management, on page 45](#)
- [UCS Mini Power Management , on page 45](#)

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, rack servers, UCS Mini, and mixed UCS domains.

Cisco UCS Manager supports power capping on the following:

- UCS 6400 Series Fabric Interconnects
- UCS 6300 Series Fabric Interconnects
- UCS 6324 Series Fabric Interconnects (Cisco UCS Mini)
- UCS 6200 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 7500 Watts) 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 @ 240 V
Grid	5000 Watts
N+1	7500 Watts
Non-Redundant	8280 Watts



Note This table is valid if there are four PSUs installed in the chassis.

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.

Starting with Cisco UCS Manager 4.1(2), the power control policy is also used for regulating fans in Cisco UCS C220 M5 and C240 M5 rack servers in acoustically-sensitive environments. The Acoustic setting for these fans is only available on these servers. On C240 SD M5 rack servers, Acoustic mode is the default mode.



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).
Owner field	The options are: <ul style="list-style-type: none"> • Local—This policy is available only to service profiles and service profile templates in this Cisco UCS domain. • Pending Global—Control of this policy is being transferred to Cisco UCS Central. Once the transfer is complete, this policy will be available to all Cisco UCS domains registered with Cisco UCS Central. • Global—This policy is managed by Cisco UCS Central. Any changes to this policy must be made through Cisco UCS Central.

Name	Description
Fan Speed Policy drop-down	<p>Enables to control the speed at which the C-Series (rack) server fans operate, thereby managing the cooling efficiency and power consumption of the server. The options are:</p> <ul style="list-style-type: none"> • 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. <p>Note The Performance option is not supported on Cisco UCS C-Series M5 and M6 servers.</p> <ul style="list-style-type: none"> • 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 fan runs on optimal speed. • Acoustic—The fan speed is reduced to reduce noise levels in acoustic-sensitive environments. Rather than regulating energy consumption and preventing component throttling as in other modes, the Acoustic option could result in short-term throttling to achieve a lowered noise level. <p>Note Low Power mode is the default mode for Cisco UCS C220 M5 and C240 M5 servers. For all other servers, Acoustic mode is the default mode.</p> <p>Important For Cisco UCS C125 M5 Server, ensure that you select the same Fan Speed Policy for all the servers in an enclosure. Cisco UCS Manager applies the Fan Speed Policy of the server which gets associated last. Having the same Fan Speed Policy for the all the server ensures that the desired Fan Speed Policy is applied irrespective of which server is associated last.</p>

Name	Description
Power Capping field	<p>What happens to a server when the demand for power within a power group exceeds the power supply. The options are:</p> <ul style="list-style-type: none"> • No Cap—The server runs at full capacity regardless of the power requirements of the other servers in its power group. <p>Note If you choose No Cap in the power capping field, ensure that Performance is not selected for the fan speed policy. This combination can lead to failure in associating the service profile with the server.</p> <ul style="list-style-type: none"> • Cap—The server is allocated a minimum amount of power capacity based on the server's priority relative to the other servers in its server group. If more power becomes available, Cisco UCS Manager 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. <p>When you select Cap, Cisco UCS Manager displays the Priority field.</p>
Priority field	<p>The priority the server has within its power group when power capping is in effect.</p> <p>Enter an integer between 1 and 10, where 1 is the highest priority.</p>

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**.

Acoustic Mode Fan Profile

Acoustic Mode Fan Profile

The Acoustic Mode fan profile is available on Cisco UCS C-Series rack servers..

Setting up an Acoustic Mode fan policy lets you reduce the noise level on Cisco UCS C-Series rack servers. The higher capacity fans on Cisco UCS C-Series rack servers increase the cooling capacity, but also create more acoustic noise. The standard fan profiles for Cisco UCS C-Series rack servers (Low Power, Balanced, High Power, and Max Power) are designed to regulate the server to optimize energy consumption. The primary goal of these fan profiles is to prevent throttling of CPU's and peripherals.

The goal of Acoustic Mode is reducing fan speed to reduce noise levels in acoustic-sensitive environments. Power capping has no effect when Acoustic Mode is selected.

Configuring Acoustic Mode

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**. Although these steps use the Power Control menu, you are creating a Fan Policy, which is administered through these menus.
 - Step 5** In the **Create Power Control Policy** dialog box, complete the following fields:

Name	Description
Name field	<p>The name of the policy.</p> <p>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.</p>
Description field	<p>A description of the policy. Cisco recommends including information about where and when to use the policy.</p> <p>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).</p>

Name	Description
Fan Speed Policy drop-down	<p>Fan speed is for C-Series Rack servers only. Acoustic mode is a fan policy available only on Cisco UCS C220 M5, C240 M5, C240 SD M5 Rack Servers.</p> <p>Fan speed can be one of the following:</p> <ul style="list-style-type: none"> • Acoustic—The fan speed is reduced to reduce noise levels in acoustic-sensitive environments. The Acoustic option can result in short-term throttling to achieve a lowered noise level. <p>Note For Cisco UCS C-Series M5 servers using Acoustic Mode, cap in the Power Capping field is automatically selected. Acoustic Mode is the default fan speed policy for rack servers.</p> <ul style="list-style-type: none"> • 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. <p>Note Performance Mode is not available on M5 servers.</p>
Power Capping field	<p>Power Capping occurs when the demand for power within a power group exceeds the power supply. For Cisco UCS C-Series M5 servers using Acoustic Mode, cap in the Power Capping field is automatically selected.</p> <ul style="list-style-type: none"> • No Cap—Allows you to set a priority for the server power throttling when Acoustic mode is selected. • cap—The server is allocated an amount of power capacity based on Acoustic Mode's need for power throttling and the server's priority relative to the other servers in its server group. <p>When cap is selected, Cisco UCS Manager GUI displays the Priority field.</p>

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. For Acoustic Mode, the default is 5.

Step 6 Click **OK**.

What to do next

Include the policy in a service profile or service profile template.

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
<p>Insufficient budget for power group POWERGROUP_NAME</p> <p>and/or</p> <p>Chassis N cannot be capped as group cap is low. Please consider raising the cap.</p> <p>and/or</p> <p>Admin committed insufficient for power group GROUP_NAME, using previous value N</p> <p>and/or</p> <p>Power cap application failed for chassis N</p>	<p>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.</p>	<p>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.</p>
<p>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</p>	<p>Displays when the power budget requirement for the chassis is more than the PSU power that is available.</p>	<p>Check the PSU input power and redundancy policy to ensure that enough power is available for the chassis.</p> <p>If a PSU failed, replace the PSU.</p>
<p>Power cap application failed for server N</p>	<p>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.</p>	<p>Do not power on un-associated servers.</p>
<p>P-State lowered as consumption hit power cap for server</p>	<p>Displays when the server is capped to reduce the power consumption below the allocated power.</p>	<p>This is an information message.</p> <p>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.</p>
<p>Chassis N has a mix of high-line and low-line PSU input power sources.</p>	<p>This fault is raised when a chassis has a mix of high-line and low-line PSU input sources connected.</p>	<p>This is an unsupported configuration. All PSUs must be connected to similar power sources.</p>

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:
- 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.
 - Click **Next**.
- Step 7** On the **Add Chassis Members** page of the **Create Power Group** wizard, do the following:
- In the **Chassis** table, choose one or more chassis to include in the power group.
 - 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.
 - Click **Next**.
- Step 8** On the **Add Rack Members** page of the **Create Power Group** wizard, do the following:
- In the **Rack Unit** table, choose one or more rack units to include in the power group.
 - 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.
 - Click **Next**.
- Step 9** On the **Add FEX Members** page of the **Create Power Group** wizard, do the following:
- In the **FEX** table, choose one or more FEX to include in the power group.
 - 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.
 - Click **Next**.

- Step 10** On the **Add FI Members** page of the **Create Power Group** wizard, do the following:
- In the **FI** table, choose one or more FI to include in the power group.
 - 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.
 - Click **Next**.
- Step 11** On the **Power Group Attributes** page of the **Create Power Group** wizard, do the following:
- 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.

- 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:
 - In the **Chassis** table, choose one or more chassis to include in the power group.
 - 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.
 - 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: <ul style="list-style-type: none"> • Unbounded—The server is not power capped under any circumstances. • Enabled—The Cisco UCS Manager GUI displays the Watts field. <p>Note Manual blade level power capping will limit the power consumption of a single system, regardless of available power in the chassis.</p>
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.

Fan Control Policy Configuration

Fan Control Policy

Fan Control Policies enable you to control the fan speed to bring down server power consumption and noise levels. With the introduction of Fan Control policies, you can determine the right fan speed for the server, based on the components in the server.

Globally managing the fan speed can help in power management by applying a single policy for all B-series server fans in an enclosure, based on general cooling needs. Set the fan speed on a per-chassis basis in the Global Policies.

Fan Control policy options include:

- **Balanced**—The fan runs at a faster speed when needed, based on the heat generated by the server. When possible, the fan returns to the minimum required speed. This is the default option.
- **Low Power**—The fan runs at the minimum speed that is required to keep the server cool.

Creating a Fan Control Policy

You can create a Fan Control Policy and define the right fan control setting based on the server configuration and server components.

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** In the **Fan Control Policy** area, click one of the following radio buttons to define the fan control setting:
- **Balanced**—This setting can cool almost any server configuration. This is the default option.
 - **Low Power**—This setting is ideal for minimal configuration servers.
- Step 5** Click **Save Changes**.
-

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	<p>The name of the policy.</p> <p>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.</p>
Description field	<p>A description of the policy. Cisco recommends including information about where and when to use the policy.</p> <p>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).</p>
Sync-Option field	<p>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:</p> <ul style="list-style-type: none"> • 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 all the Cisco UCS C-Series servers except Cisco UCS C125 M5 Servers.

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 6

Blade Server Hardware Management

- [Blade Server Management](#), on page 47
- [Booting a Blade Server](#), on page 49
- [Booting a Rack-Mount Server from the Service Profile](#), on page 50
- [Determining the Boot Order of a Blade Server](#), on page 50
- [Shutting Down a Blade Server](#), on page 51
- [Shutting Down a Server from the Service Profile](#), on page 51
- [Resetting a Blade Server](#), on page 52
- [Resetting a Blade Server to Factory Default Settings](#), on page 52
- [Reacknowledging a Blade Server](#), on page 53
- [Removing a Server from a Chassis](#), on page 54
- [Deleting the Inband Configuration from a Blade Server](#), on page 54
- [Decommissioning a Blade Server](#), on page 55
- [Removing a Non-Existent Blade Server Entry](#), on page 55
- [Recommissioning a Blade Server](#), on page 56
- [Reacknowledging a Server Slot in a Chassis](#), on page 56
- [Removing a Non-Existent Blade Server from the Configuration Database](#), on page 57
- [Turning the Locator LED for a Blade Server On and Off](#), on page 57
- [Turning the Local Disk Locator LED on a Blade Server On and Off](#), on page 58
- [Resetting the CMOS for a Blade Server](#), on page 58
- [Resetting the CIMC for a Blade Server](#), on page 59
- [Clearing TPM for a Blade Server](#), on page 59
- [Viewing the POST Results for a Blade Server](#), on page 60
- [Issuing an NMI from a Blade Server](#), on page 60
- [Viewing Health Events for a Blade Server](#), on page 61
- [Health LED Alarms](#), on page 63
- [Smart SSD](#), on page 63

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:
- Click the **Power Cycle** option.
 - (Optional) Check the check box if you want Cisco UCS Manager to complete all management operations that are pending on this server.
 - 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:
- Click **Reset to Factory Default**.
 - 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 56](#).

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.

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.

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 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 54](#).

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:
- Click **Diagnostic Interrupt**.
 - 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 health events that are triggered for the component.

Name	Description
Health Severity field	<p>Highest severity of all the health events that are triggered for the component. This can be one of the following:</p> <ul style="list-style-type: none"> • critical • major • minor • warning • info • cleared <p>Note The severity levels listed here are from highest to lowest severity.</p>
Health Details area	
Severity column	<p>Severity of the health event. This can be one of the following:</p> <ul style="list-style-type: none"> • critical • major • minor • warning • info • cleared <p>Note The severity levels listed here are from highest to lowest severity.</p>
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: <ul style="list-style-type: none"> • Critical—The blade health LED is blinking amber. This is indicated with a red dot. • Minor—The blade health LED is amber. This is indicated with an orange dot.
Description column	A brief description of the alarm.
Sensor ID column	The ID of the sensor the triggered the alarm.
Sensor Name column	The name of the sensor that triggered the alarm.

Viewing Health LED Alarms

Procedure

-
- Step 1** In the **Navigation** pane, click **Equipment**.
- Step 2** Expand **Equipment** > **Chassis** > *Chassis Number* > **Servers**.
- Step 3** Click the server for which you want to view health LED alarms.
- Step 4** In the **Work** pane, click the **General** tab.
- Step 5** In the **Actions** area, click **View Health LED Alarms**.

The **View Health LED Alarms** dialog box lists the health LED alarms for the selected server.

- Step 6** Click **OK** to close the **View Health LED Alarms** dialog box.
-

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 7

Rack-Mount Server Hardware Management

- [Rack-Mount Server Management, on page 67](#)
- [Rack-Enclosure Server Management, on page 68](#)
- [Guidelines for Removing and Decommissioning Rack-Mount Servers, on page 69](#)
- [Recommendations for Avoiding Unexpected Server Power Changes, on page 69](#)
- [Booting a Rack-Mount Server, on page 70](#)
- [Booting a Rack-Mount Server from the Service Profile , on page 71](#)
- [Determining the Boot Order of a Rack-Mount Server, on page 71](#)
- [Shutting Down a Rack-Mount Server, on page 72](#)
- [Shutting Down a Server from the Service Profile , on page 72](#)
- [Resetting a Rack-Mount Server, on page 73](#)
- [Resetting a Rack-Mount Server to Factory Default Settings, on page 74](#)
- [Persistent Memory Scrub, on page 75](#)
- [Reacknowledging a Rack-Mount Server, on page 75](#)
- [Deleting the Inband Configuration from a Rack-Mount Server, on page 76](#)
- [Decommissioning a Rack-Mount Server, on page 76](#)
- [Recommissioning a Rack-Mount Server, on page 77](#)
- [Renumbering a Rack-Mount Server, on page 77](#)
- [Removing a Non-Existent Rack-Mount Server from the Configuration Database, on page 78](#)
- [Turning the Locator LED for a Rack-Mount Server On and Off, on page 79](#)
- [Turning the Local Disk Locator LED on a Rack-Mount Server On and Off, on page 79](#)
- [Resetting the CMOS for a Rack-Mount Server, on page 80](#)
- [Resetting the CIMC for a Rack-Mount Server, on page 80](#)
- [Clearing TPM for a Rack-Mount Server, on page 81](#)
- [Issuing an NMI from a Rack-Mount Server, on page 81](#)
- [Viewing Health Events for a Rack-Mount Server, on page 82](#)
- [Viewing the POST Results for a Rack-Mount Server, on page 84](#)
- [Viewing the Power Transition Log, on page 84](#)
- [Viewing Cisco UCS C125 M5 Server Slot ID, on page 85](#)

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.

Rack-Enclosure Server Management

Beginning with release 4.0(1a), Cisco UCS Manager extends support for all existing features on Cisco UCS C125 M5 Server unless specifically noted in this guide.

Cisco UCS C125 M5 Servers are housed in the Cisco UCS C4200 Series Rack Server Chassis. Each Cisco UCS C4200 Series Rack Server Chassis supports two to four Cisco UCS C125 M5 Server nodes. To manage the Cisco UCS C125 M5 Server nodes, Cisco UCS Manager supports the following:

- **Enclosures:**

Cisco UCS Manager GUI path - **Equipment** > **Rack-Mounts** > **Enclosures**

Displays a list of all the Cisco UCS C4200 Series Rack Server Chassis managed by Cisco UCS Manager.

- **Rack Enclosure *rack_enclosure_number*:**

Cisco UCS Manager GUI path - **Equipment** > **Rack-Mounts** > **Enclosures** > **Rack Enclosure *rack_enclosure_number***

Each **Rack Enclosure** is one Cisco UCS C4200 Series Rack Server Chassis, which can contain up to four Cisco UCS C125 M5 Server nodes, four fan units, and two PSUs. See [Viewing Cisco UCS C125 M5 Server Slot ID, on page 85](#) for the slot IDs of the server.

Cisco UCS C125 M5 Servers can be managed the same way as other rack servers from **Rack Enclosure *rack_enclosure_number***.



Note Cisco UCS C125 M5 Servers is supported on the following Fabric Interconnects:

- Cisco UCS 6400 Series Fabric Interconnect
 - Cisco UCS 6300 Series Fabric Interconnect
-

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**.

Note

For Cisco UCS C125 M5 Servers, expand **Equipment** > **Rack Mounts** > **Enclosures** > **Rack Enclosure** *rack_enclosure_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 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**.
Note
For Cisco UCS C125 M5 Servers, expand **Equipment** > **Rack Mounts** > **Enclosures** > **Rack Enclosure** *rack_enclosure_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 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**.
- Note**
For Cisco UCS C125 M5 Servers, expand **Equipment** > **Rack Mounts** > **Enclosures** > **Rack Enclosure *rack_enclosure_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 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**.
- Note**
For Cisco UCS C125 M5 Servers, expand **Equipment** > **Rack Mounts** > **Enclosures** > **Rack Enclosure** *rack_enclosure_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:
- Click the **Power Cycle** option.
 - (Optional) Check the check box if you want Cisco UCS Manager to complete all management operations that are pending on this server.
 - 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**.

Note

For Cisco UCS C125 M5 Servers, expand **Equipment > Rack Mounts > Enclosures > Rack Enclosure *rack_enclosure_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, 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

After persistent memory scrub is complete, the following happen:

- All persistent memory data is erased
- Persistent memory configuration is reset to factory default settings.

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.

- Persistent memory module security is disabled

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**.

Note

For Cisco UCS C125 M5 Servers, expand **Equipment** > **Rack Mounts** > **Enclosures** > **Rack Enclosure *rack_enclosure_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, 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**.

Note

For Cisco UCS C125 M5 Servers, expand **Equipment** > **Rack Mounts** > **Enclosures** > **Rack Enclosure** *rack_enclosure_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, click **Decommission**, then click **OK**.

The server is removed from the Cisco UCS configuration.

Note

When you decommission the last Cisco UCS C125 M5 Server from a **Rack Enclosure**, Cisco UCS Manager removes the complete **Rack Enclosure** *rack_enclosure_number* entry from the navigation pane.

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:
- In the **Recommission** column, check the check box.
 - 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**.

Note

For Cisco UCS C125 M5 Servers, expand **Equipment** > **Rack Mounts** > **Enclosures** > **Rack Enclosure** *rack_enclosure_number* > **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**.

Note

For Cisco UCS C125 M5 Servers, expand **Equipment** > **Rack Mounts** > **Enclosures** > **Rack Enclosure** *rack_enclosure_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 Rack-Mount Server On and Off

Procedure

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

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

Note

For Cisco UCS C125 M5 Servers, expand **Equipment > Rack Mounts > Enclosures > Rack Enclosure *rack_enclosure_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**
 - **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**.

Note

For Cisco UCS C125 M5 Servers, expand **Equipment > Rack Mounts > Enclosures > Rack Enclosure *rack_enclosure_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 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**.
- Note**
For Cisco UCS C125 M5 Servers, expand **Equipment** > **Rack Mounts** > **Enclosures** > **Rack Enclosure** *rack_enclosure_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 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**.
- Note**
For Cisco UCS C125 M5 Servers, expand **Equipment** > **Rack Mounts** > **Enclosures** > **Rack Enclosure** *rack_enclosure_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 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**.
- Note**
For Cisco UCS C125 M5 Servers, expand **Equipment** > **Rack Mounts** > **Enclosures** > **Rack Enclosure *rack_enclosure_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**.
-

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**.

Note

For Cisco UCS C125 M5 Servers, expand **Equipment** > **Rack Mounts** > **Enclosures** > **Rack Enclosure** *rack_enclosure_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, 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**.

Note

For Cisco UCS C125 M5 Servers, expand **Equipment** > **Rack Mounts** > **Enclosures** > **Rack Enclosure** *rack_enclosure_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 health events that are triggered for the component.

Name	Description
Health Severity field	<p>Highest severity of all the health events that are triggered for the component. This can be one of the following:</p> <ul style="list-style-type: none"> • critical • major • minor • warning • info • cleared <p>Note The severity levels listed here are from highest to lowest severity.</p>
Health Details area	
Severity column	<p>Severity of the health event. This can be one of the following:</p> <ul style="list-style-type: none"> • critical • major • minor • warning • info • cleared <p>Note The severity levels listed here are from highest to lowest severity.</p>
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**.

Note

For Cisco UCS C125 M5 Servers, expand **Equipment > Rack Mounts > Enclosures > Rack Enclosure *rack_enclosure_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 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.

Viewing Cisco UCS C125 M5 Server Slot ID

Procedure

- Step 1** In the **Navigation** pane, click **Equipment**.
- Step 2** Expand **Equipment** > **Rack Mounts** > **Enclosures** > **Rack Enclosure** *rack_enclosure_number*.
- Step 3** In the **Work** pane, click the **Slots** tab.
-



CHAPTER 8

S3X60 Server Node Hardware Management

- [Cisco UCS S3260 Server Node Management, on page 87](#)
- [Booting a Cisco UCS S3260 Server Node, on page 88](#)
- [Booting a Cisco UCS S3260 Server Node from the Service Profile , on page 88](#)
- [Determining the Boot Order of a Cisco UCS S3260 Server Node, on page 89](#)
- [Shutting Down a Cisco UCS S3260 Server Node, on page 89](#)
- [Shutting Down a Cisco UCS S3260 Server Node from the Service Profile , on page 90](#)
- [Resetting a Cisco UCS S3260 Server Node, on page 90](#)
- [Resetting a Cisco UCS S3260 Server Node to Factory Default Settings, on page 91](#)
- [Reacknowledging a Cisco UCS S3260 Server Node, on page 92](#)
- [Removing a Cisco UCS S3260 Server Node from a Chassis, on page 93](#)
- [Deleting the Inband Configuration from a Cisco UCS S3260 Server Node, on page 93](#)
- [Decommissioning a Cisco UCS S3260 Server Node, on page 94](#)
- [Recommissioning a Cisco UCS S3260 Server Node, on page 94](#)
- [Reacknowledging a Server Slot in a S3260 Chassis, on page 95](#)
- [Removing a Non-Existent Cisco UCS S3260 Server Node from the Configuration Database, on page 95](#)
- [Turning the Locator LED for a Cisco UCS S3260 Server Node On and Off, on page 96](#)
- [Turning the Local Disk Locator LED on a Cisco UCS S3260 Server Node On and Off, on page 96](#)
- [Resetting the CIMC for a Cisco UCS S3260 Server Node, on page 97](#)
- [Resetting the CMOS for a Cisco UCS S3260 Server Node, on page 97](#)
- [Resetting the BIOS Password for a S3X60 Server, on page 98](#)
- [Issuing an NMI from a Cisco UCS S3260 Server Node, on page 98](#)
- [Viewing the POST Results for a Cisco UCS S3260 Server Node, on page 99](#)
- [Viewing Health Events for a Cisco UCS S3260 Server Node, on page 99](#)
- [Health LED Alarms, on page 101](#)

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:
- Click the **Power Cycle** option.
 - (Optional) Check the check box if you want Cisco UCS Manager to complete all management operations that are pending on this server.
 - 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:
- Click **Reset to Factory Default**.
 - 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 re-install 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.
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 93](#).

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**.
-

Resetting the BIOS Password for a S3X60 Server

This option allows you to reset the BIOS password without using the F2 BIOS configuration prompt. Resetting the BIOS password is not part of the normal maintenance of a server. After BIOS password reset, the server is rebooted immediately and the new BIOS password gets updated.

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 BIOS password.
- 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 BIOS Password**, 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:
- Click **Diagnostic Interrupt**.
 - 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 health events that are triggered for the component.

Name	Description
Health Severity field	<p>Highest severity of all the health events that are triggered for the component. This can be one of the following:</p> <ul style="list-style-type: none"> • critical • major • minor • warning • info • cleared <p>Note The severity levels listed here are from highest to lowest severity.</p>
Health Details area	
Severity column	<p>Severity of the health event. This can be one of the following:</p> <ul style="list-style-type: none"> • critical • major • minor • warning • info • cleared <p>Note The severity levels listed here are from highest to lowest severity.</p>
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: <ul style="list-style-type: none"> • 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 9

Server Pools

- [Configuring Server Pools, on page 103](#)
- [Configuring UUID Suffix Pools, on page 105](#)
- [Configuring IP Pools, on page 107](#)

Configuring Server Pools

Server Pools

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

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

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

Creating a Server Pool

Procedure

- Step 1** In the **Navigation** pane, click **Servers**.
- Step 2** Expand **Servers > Pools**.
- Step 3** Expand the node for the organization where you want to create the pool.
If the system does not include multi tenancy, expand the **root** node.
- Step 4** Right-click the **Server Pools** node and select **Create Server Pool**.
- Step 5** On the **Set Name and Description** page of the **Create Server Pool** wizard, complete the following fields:

Name	Description
Name field	The name of the server pool. This name can be between 1 and 32 alphanumeric characters. You cannot use spaces or any special characters other than - (hyphen), _ (underscore), : (colon), and . (period), and you cannot change this name after the object is saved.
Description field	A user-defined description of the server pool. Enter up to 256 characters. You can use any characters or spaces except ` (accent mark), \ (backslash), ^ (carat), " (double quote), = (equal sign), > (greater than), < (less than), or ' (single quote).

Step 6 Click **Next**.

Step 7 On the **Add Servers** page of the **Create Server Pool** wizard:

- a) Select one or more servers from the **Available Servers** table.
- b) Click the >> button to add the servers to the server pool.
- c) When you have added all desired servers to the pool, click **Finish**.

Deleting a Server Pool

Procedure

- Step 1** In the **Navigation** pane, click **Servers**.
- Step 2** Expand **Servers > Pools > Organization_Name**.
- Step 3** Expand the **Server Pools** node.
- Step 4** Right-click the pool you want to delete and select **Delete**.
- Step 5** If a confirmation dialog box displays, click **Yes**.

Adding Servers to a Server Pool

Procedure

- Step 1** In the **Navigation** pane, click **Servers**.
- Step 2** Expand **Servers > Pools > Organization_Name**.
- Step 3** Right-click the pool to which you want to add one or more servers and select **Add Servers to Server Pool**.
- Step 4** In the **Add Servers to Server Pool** dialog box, do the following:
 - a) In the **Servers** table, select the servers that you want to add to the server pool.

You can use the **Shift** key or **Ctrl** key to select multiple entries.

- b) Click the >> button to move those servers to the **Pooled Servers** table and add them to the server pool.
 - c) Click **OK**.
-

Removing Servers from a Server Pool

Procedure

- Step 1** In the **Navigation** pane, click **Servers**.
 - Step 2** Expand **Servers > Pools > Organization_Name**.
 - Step 3** Right-click the pool from which you want to remove one or more servers and select **Add Servers to Server Pool**.
 - Step 4** In the **Add Servers to Server Pool** dialog box, do the following:
 - a) In the **Pooled Servers** table, select the servers that you want to remove from the server pool.

You can use the **Shift** key or **Ctrl** key to select multiple entries.
 - b) Click the << button to move those servers to the **Servers** table and remove them from the server pool.
 - c) Click **OK**.
-

Configuring UUID Suffix Pools

UUID Suffix Pools

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

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

Creating a UUID Suffix Pool

Procedure

- Step 1** In the **Navigation** pane, click **Servers**.
- Step 2** Expand **Servers > Pools**.

Step 3 Expand the node for the organization where you want to create the pool.

If the system does not include multi tenancy, expand the **root** node.

Step 4 Right-click **UUID Suffix Pools** and select **Create UUID Suffix Pool**.

Step 5 In the **Define Name and Description** page of the **Create UUID Suffix Pool** wizard, complete the following fields:

Name	Description
Name field	The name of the UUID pool. This name can be between 1 and 32 alphanumeric characters. You cannot use spaces or any special characters other than - (hyphen), _ (underscore), : (colon), and . (period), and you cannot change this name after the object is saved.
Description field	The user-defined description of the pool. Enter up to 256 characters. You can use any characters or spaces except ` (accent mark), \ (backslash), ^ (carat), " (double quote), = (equal sign), > (greater than), < (less than), or ' (single quote).
Prefix field	This can be one of the following: <ul style="list-style-type: none"> • Derived—The system creates the suffix. • other—You specify the desired suffix. If you select this option, Cisco UCS Manager GUI displays a text field where you can enter the desired suffix, in the format <code>XXXXXXXX-XXXX-XXXX</code>.
Assignment Order field	This can be one of the following: <ul style="list-style-type: none"> • Default—Cisco UCS Manager selects a random identity from the pool. • Sequential—Cisco UCS Manager selects the lowest available identity from the pool.

Step 6 Click **Next**.

Step 7 In the **Add UUID Blocks** page of the **Create UUID Suffix Pool** wizard, click **Add**.

Step 8 In the **Create a Block of UUID Suffixes** dialog box, complete the following fields:

Name	Description
From field	The first UUID in the block.
Size field	The number of UUIDs in the block.

Step 9 Click **OK**.

Step 10 Click **Finish** to complete the wizard.

What to do next

Include the UUID suffix pool in a service profile and/or template.

Deleting a UUID Suffix Pool

If you delete a pool, Cisco UCS Manager does not reallocate any addresses from that pool that were assigned to vNICs or vHBAs. All assigned addresses from a deleted pool remain with the vNIC or vHBA to which they are assigned until one of the following occurs:

- The associated service profiles are deleted.
- The vNIC or vHBA to which the address is assigned is deleted.
- The vNIC or vHBA is assigned to a different pool.

Procedure

-
- Step 1** In the **Navigation** pane, click **Servers**.
 - Step 2** Expand **Servers > Pools > *Organization_Name***.
 - Step 3** Expand the **UUID Suffix Pools** node.
 - Step 4** Right-click the pool you want to delete and select **Delete**.
 - Step 5** If a confirmation dialog box displays, click **Yes**.
-

Configuring IP Pools

IP Pools

IP pools are collections of IP addresses that do not have a default purpose. You can create IPv4 or IPv6 address pools in Cisco UCS Manager to do the following:

- Replace the default management IP pool **ext-mgmt** for servers that have an associated service profile. Cisco UCS Manager reserves each block of IP addresses in the IP pool for external access that terminates in the Cisco Integrated Management Controller (CIMC) on a server. If there is no associated service profile, you must use the **ext-mgmt** IP pool for the CIMC to get an IP address.
- Replace the management inband or out-of-band IP addresses for the CIMC.



Note You cannot create iSCSI boot IPv6 pools in Cisco UCS Manager.

You can create IPv4 address pools in Cisco UCS Manager to do the following:

- Replace the default iSCSI boot IP pool **iscsi-initiator-pool**. Cisco UCS Manager reserves each block of IP addresses in the IP pool that you specify.

- Replace both the management IP address and iSCSI boot IP addresses.



Note The IP pool must not contain any IP addresses that were assigned as static IP addresses for a server or service profile.

Creating an IP Pool

Procedure

- Step 1** In the **Navigation** pane, click **LAN**.
- Step 2** In the **LAN** tab, expand **LAN > Pools > Organization_Name**.
- Step 3** Right-click **IP Pools** and select **Create IP Pool**.
- Step 4** In the **Define Name and Description** page of the **Create IP Pool** wizard, complete the following fields:

Name	Description
Name field	The name of the IP address pool. This name can be between 1 and 32 alphanumeric characters. You cannot use spaces or any special characters other than - (hyphen), _ (underscore), : (colon), and . (period), and you cannot change this name after the object is saved.
Description field	The user-defined description of the IP address pool. Enter up to 256 characters. You can use any characters or spaces except ` (accent mark), \ (backslash), ^ (carat), " (double quote), = (equal sign), > (greater than), < (less than), or ' (single quote).
Assignment Order field	This can be one of the following: <ul style="list-style-type: none"> • Default—Cisco UCS Manager selects a random identity from the pool. • Sequential—Cisco UCS Manager selects the lowest available identity from the pool.

- Step 5** Click **Next**.
- Step 6** In the **Add IPv4 Blocks** page of the **Create IP Pool** wizard, click **Add**.
- Step 7** In the **Create a Block of IPv4 Addresses** dialog box, complete the following fields:

Name	Description
From field	The first IPv4 address in the block.
Size field	The number of IP addresses in the pool.

Name	Description
Subnet Mask field	The subnet mask associated with the IPv4 addresses in the block.
Default Gateway field	The default gateway associated with the IPv4 addresses in the block.
Primary DNS field	The primary DNS server that this block of IPv4 addresses should access.
Secondary DNS field	The secondary DNS server that this block of IPv4 addresses should access.

Step 8

Click **Next**.

Step 9

In the **Add IPv6 Blocks** page of the **Create IP Pool** wizard, click **Add**.

Step 10**Step 11**

In the **Create a Block of IPv6 Addresses** dialog box, complete the following fields:

Name	Description
From field	The first IPv6 address in the block.
Size field	The number of IP addresses in the pool.
Prefix	The network address prefix associated with the IPv6 addresses in the block.
Default Gateway field	The default gateway associated with the IPv6 addresses in the block.
Primary DNS field	The primary DNS server that this block of IPv6 addresses should access.
Secondary DNS field	The secondary DNS server that this block of IPv6 addresses should access.

Step 12

Click **OK**.

Step 13

Click **Finish** to complete the wizard.

What to do next

Include the IP pool in a service profile and template.

Adding a Block to an IP Pool

You can add blocks of IPv4 or IPv6 addresses to IP pools.

Procedure**Step 1**

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

Step 2 In the **LAN** tab, expand **LAN > Pools > Organization_Name** .

Step 3 Expand the **IP Pools** node.

Step 4 Right-click the desired IP pool and select one of:

- **Create Block of IPv4 Addresses**
- **Create Block of IPv6 Addresses**

Step 5 Complete the fields in the appropriate dialog box.

a) In the **Create a Block of IPv4 Addresses** dialog box, complete the following fields:

Name	Description
Name column	The range of IPv4 addresses assigned to the block.
From column	The first IPv4 address in the block.
To column	The last IPv4 address in the block.
Subnet column	The subnet mask associated with the IPv4 addresses in the block.
Default Gateway column	The default gateway associated with the IPv4 addresses in the block.
Primary DNS column	The primary DNS server that this block of IPv4 addresses should access.
Secondary DNS column	The secondary DNS server that this block of IPv4 addresses should access.

b) In the **Create a Block of IPv6 Addresses** dialog box, complete the following fields:

Name	Description
Name column	The range of IPv6 addresses assigned to the block.
From column	The first IPv6 address in the block.
To column	The last IPv6 address in the block.
Prefix column	The network address prefix associated with the IPv6 addresses in the block.
Default Gateway column	The default gateway associated with the IPv6 addresses in the block.
Primary DNS column	The primary DNS server that this block of IPv6 addresses should access.
Secondary DNS column	The secondary DNS server that this block of IPv6 addresses should access.

Step 6 Click **OK**.

Deleting a Block from an IP Pool

Procedure

- Step 1** In the **Navigation** pane, click **LAN**.
- Step 2** In the **LAN** tab, expand **LAN > Pools > Root** .
- Step 3** Expand the **IP Pools** node.
- Step 4** Expand the pool for which you want to delete a block of IP addresses.
- Step 5** Right-click the IP address block that you want to delete and select **Delete**.
- Step 6** If a confirmation dialog box displays, click **Yes**.
-

Deleting an IP Pool

If you delete a pool, Cisco UCS Manager does not reallocate any addresses from that pool that were assigned to vNICs or vHBAs. All assigned addresses from a deleted pool remain with the vNIC or vHBA to which they are assigned until one of the following occurs:

- The associated service profiles are deleted.
- The vNIC or vHBA to which the address is assigned is deleted.
- The vNIC or vHBA is assigned to a different pool.

Procedure

- Step 1** In the **Navigation** pane, click **LAN**.
- Step 2** In the **LAN** tab, expand **LAN > Pools > Organization_Name** .
- Step 3** Expand the **IP Pools** node.
- Step 4** Right-click the IP pool you want to delete and select **Delete**.

Note

You cannot delete the default pools **ext-mgmt** and **iscsi-initiator-pool**.

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



CHAPTER 10

Server Boot

- [Boot Policy, on page 113](#)
- [UEFI Boot Mode, on page 114](#)
- [UEFI Secure Boot, on page 115](#)
- [CIMC Secure Boot, on page 116](#)
- [Creating a Boot Policy, on page 118](#)
- [SAN Boot, on page 119](#)
- [iSCSI Boot, on page 121](#)
- [LAN Boot, on page 145](#)
- [Local Devices Boot, on page 146](#)
- [Deleting a Boot Policy, on page 152](#)
- [UEFI Boot Parameters, on page 153](#)

Boot Policy

The Cisco UCS Manager enables you to create a boot policy for blade servers and rack servers.

The Cisco UCS Manager boot policy overrides the boot order in the BIOS setup menu and determines the following:

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

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

You can either create a named boot policy to associate with one or more service profiles, or create a boot policy for a specific service profile. A boot policy must be included in a service profile, and that service profile must be associated with a server for it to take effect. If you do not include a boot policy in a service profile, Cisco UCS Manager applies the default boot policy.



Note Changes to a boot policy might be propagated to all servers created with an updating service profile template that includes that boot policy. Re-association of the service profile with the server to rewrite the boot order information in the BIOS is automatically triggered.

You can also specify the following for the boot policy:

- Local LUN name. The name specified is the logical name in the storage profile, not the deployed name. Specify only a primary name. Specifying a secondary name results in a configuration error.
- Specific JBOD disk number for booting from JBOD disks.
- Any LUN for backward compatibility; however, we do not recommend this. Other devices must not have bootable images to ensure a successful boot.

UEFI Boot Mode

Unified Extensible Firmware Interface (UEFI) is a specification that defines a software interface between an operating system and platform firmware. Cisco UCS Manager uses UEFI to replace the BIOS firmware interfaces. This allows the BIOS to run in UEFI mode while still providing legacy support.

You can choose either legacy or UEFI boot mode when you create a boot policy. Legacy boot mode is supported for all Cisco UCS servers except Cisco UCS C125 M5 Server. UEFI boot mode is supported only on M3 and higher servers, and allows you to enable UEFI secure boot mode. Cisco UCS C125 M5 Server supports only UEFI boot mode.

UEFI PXE boot is supported with all Cisco VIC adapters on Cisco UCS rack servers integrated with Cisco UCS Manager Release 2.2(4) and later releases. Beginning with Cisco UCS Manager Release 2.2(1), UEFI PXE boot is supported on all Cisco blade servers.

The following limitations apply to the UEFI boot mode:

- UEFI boot mode is not supported with the following combinations:
 - Gen-3 Emulex and QLogic adapters on Cisco UCS blade and rack servers integrated with Cisco UCS Manager.
 - iSCSI boot for all adapters on Cisco UCS rack servers integrated with Cisco UCS Manager.
- To use UEFI boot mode with dual iSCSI vNICs, you must specify a common iSCSI initiator name in the service profile—either manually or via an IQN pool—to ensure it is applied to both underlying iSCSI eNICs. Do not configure the IQN within the individual iSCSI boot parameters. Failure to provide a common initiator name will prevent the Cisco UCS server from detecting the secondary path iSCSI LUN.
- You cannot mix UEFI and legacy boot mode on the same server.
- The server will boot correctly in UEFI mode only if the boot devices configured in the boot policy have UEFI-aware operating systems installed. If a compatible OS is not present, the boot device is not displayed on the **Actual Boot Order** tab in the **Boot Order Details** area.
- In some corner cases, the UEFI boot may not succeed because the UEFI boot manager entry was not saved correctly in the BIOS NVRAM. You can use the UEFI shell to enter the UEFI boot manager entry manually. This situation could occur in the following situations:

- If a blade server with UEFI boot mode enabled is disassociated from the service profile, and the blade is manually powered on using the **Equipment** tab or the front panel.
- If a blade server with UEFI boot mode enabled is disassociated from the service profile, and a direct VIC firmware upgrade is attempted.
- If a blade or rack server with UEFI boot mode enabled is booted off SAN LUN, and the service profile is migrated.

You can create UEFI boot parameters in Cisco UCS Manager. [UEFI Boot Parameters, on page 153](#) provides more information.

UEFI Secure Boot

Cisco UCS Manager supports UEFI secure boot on Cisco UCS B-Series M4 and higher Blade servers, Cisco UCS C-Series M3 and higher Rack servers, and Cisco UCS S-Series M4 Rack servers, and Cisco UCS C125 M5 Servers. Linux secure boot is supported on SLES 15, SLES 13 SP4, Red Hat Linux 7.6 operating systems starting with Release 4.0(4a). When UEFI secure boot is enabled, all executables, such as boot loaders and adapter drivers, are authenticated by the BIOS before they can be loaded. To be authenticated, the images must be signed by either the Cisco Certificate Authority (CA) or a Microsoft CA.

The following limitations apply to UEFI secure boot:

- UEFI boot mode must be enabled in the boot policy.
- UEFI boot mode is available only for drives.
- The Cisco UCS Manager software and the BIOS firmware must be at Release 2.2 or greater.



Note UEFI boot mode is supported on Cisco UCS C-Series and S-Series rack servers beginning with Release 2.2(3a).

- User-generated encryption keys are not supported.
- UEFI secure boot can only be controlled by Cisco UCS Manager.
- If you want to downgrade to an earlier version of Cisco UCS Manager, and you have a server in secure boot mode, you must disassociate, then re-associate the server before downgrading. Otherwise, server discovery is not successful.
- In Cisco UCS Manager Release 4.0, UEFI secure boot is supported on the following Operating Systems:
 - In Cisco UCS Manager Release 4.0(1), UEFI secure boot is supported only on Windows 2016 and Windows 2012 R2.
 - In Cisco UCS Manager Release 4.0(2), UEFI secure boot is supported only on Windows 2016 and Windows 2019.
 - In Cisco UCS Manager Release 4.0(4), UEFI secure boot is supported on the following:

Table 4: Linux Operating Systems

Linux OS	eNIC/nNIC	fNIC
RHEL 7.5	3.2.210.18.738.12	1.6.0.50
RHEL 7.6	3.2.210.18.738.12	2.0.0.37
CentOS 7.5	3.2.210.18.738.12	1.6.0.50
CentOS 7.6	3.2.210.18.738.12	1.6.0.50
SLES 12.4	3.2.210.18.738.12	2.0.0.32
SLES 15	3.2.210.18.738.12	2.0.0.39-71.0
ESXi	Inbox works	Inbox works

**Note**

- For ESXi, inbox drivers are signed and work as such. Async drivers are not signed and do not work.
- Oracle OS does not support IPv6.
- XEN OS does not support IPv6.

Table 5: Windows Operating Systems

Windows OS	neNIC	nfNIC
Windows 2016	5.3.25.4	3.2.0.3
Windows 2019	5.3.25.4	3.2.0.3

CIMC Secure Boot

With CIMC secure boot, only Cisco signed firmware images can be installed and run on the servers. When the CIMC is updated, the image is certified before the firmware is flashed. If certification fails, the firmware is not flashed. This prevents unauthorized access to the CIMC firmware.

Guidelines and Limitations for CIMC Secure Boot

- CIMC secure boot is supported on Cisco UCS M3 M4, M5 rack servers.
- After CIMC secure boot is enabled, you cannot disable it.
- After CIMC secure boot is enabled on a server, you cannot downgrade to a CIMC firmware image prior to 2.1(3).

Determining the CIMC Secure Boot Status

Procedure

- Step 1** In the **Navigation** pane, click **Equipment**.
- Step 2** Expand **Equipment > Rack-Mounts > Servers > Server Name**.
- Step 3** In the **Work** area, click the **Inventory** tab.
- Step 4** Click the **CIMC** subtab.
- Step 5** In the **CIMC** area, note the **Secure Boot Operational State** field.

This can be one of the following:

- **Unsupported**—CIMC secure boot is not supported on the server.
 - **Disabled**—CIMC secure boot is supported, but is disabled on the server.
 - **Enabling**—CIMC secure boot was enabled, and the operation is in process.
 - **Enabled**—CIMC secure boot is enabled on the server.
-

Enabling CIMC Secure Boot on a Rack Server

Procedure

- Step 1** In the **Navigation** pane, click **Equipment**.
- Step 2** Expand **Equipment > Rack-Mounts > Servers > Server Name**.
- Step 3** In the **Work** area, click the **Inventory** tab.
- Step 4** Click the **CIMC** subtab.
- Step 5** In the **Actions** area, click **Enable Secure Boot**.

CIMC secure boot is only supported on Cisco UCS M3 rack servers. If CIMC secure boot is not supported or is already enabled, this action is greyed.

- Step 6** Click **Yes** in the **Enable Secure Boot** confirmation dialog box.

Note

After enabled, you cannot disable CIMC secure boot.

Creating a Boot Policy

You can also create a local boot policy that is restricted to a service profile or service profile template. However, Cisco recommends that you create a global boot policy that can be included in multiple service profiles or service profile templates.

Procedure

- Step 1** In the **Navigation** pane, click **Servers**.
- Step 2** Expand **Servers > Policies**.
- Step 3** Expand the node for the organization where you want to create the policy.
If the system does not include multi tenancy, expand the **root** node.
- Step 4** Right-click **Boot Policies** and select **Create Boot Policy**.
The **Create Boot Policy** wizard displays.
- Step 5** Enter a unique name and description for the policy.
This name can be between 1 and 16 alphanumeric characters. You cannot use spaces or any special characters other than - (hyphen), _ (underscore), : (colon), and . (period), and you cannot change this name after the object is saved.
- Step 6** (Optional) After you make changes to the boot order, check the **Reboot on Boot Order Change** check box to reboot all servers that use this boot policy.
For boot policies applied to a server with a non-Cisco VIC adapter, even if the **Reboot on Boot Order Change** check box is not checked, when SAN devices are added, deleted, or their order is changed, the server always reboots when boot policy changes are saved.
- Step 7** (Optional) If desired, check the **Enforce vNIC/vHBA/iSCSI Name** check box.
- If checked, Cisco UCS Manager displays a configuration error and reports whether one or more of the vNICs, vHBAs, or iSCSI vNICs listed in the **Boot Order** table match the server configuration in the service profile.
 - If not checked, Cisco UCS Manager uses the vNICs or vHBAs (as appropriate for the boot option) from the service profile.
- Step 8** In the **Boot Mode** field, choose the **Legacy** or **UEFI** radio button to specify the boot mode for servers using this boot policy.
- Note**
- Cisco UCS C125 M5 Server supports only UEFI boot mode.
 - Legacy boot mode is not supported on Cisco UCS M7 and later server versions. You must use UEFI boot mode for these systems.
- Step 9** (Optional) If **Uefi** boot mode is selected, check the **Boot Security** check box to enable UEFI Secure Boot.
- Step 10** Configure one or more of the following boot options for the boot policy and set their boot order:

- **Local Devices boot**—To boot from local devices, such as local disks on the server, virtual media, or remote virtual disks, continue with [Configuring a Local Disk Boot for a Boot Policy, on page 147](#).
- **SAN boot**—To boot from an operating system image on the SAN, continue with [Configuring a SAN Boot for a Boot Policy, on page 119](#).
You can specify a primary and a secondary SAN boot. If the primary boot fails, the server attempts to boot from the secondary.
- **LAN boot**—To boot from a centralized provisioning server, continue with [Configuring a LAN Boot for a Boot Policy, on page 145](#).
- **iSCSI boot**—To boot from an iSCSI LUN, continue with [Creating an iSCSI Boot Policy, on page 130](#).

What to do next

Include the boot policy in a service profile and template.

After a server is associated with a service profile that includes this boot policy, you can verify the boot order in the **Boot Order Details** area on the **General** tab for the server.

SAN Boot

You can configure a boot policy to boot one or more servers from an operating system image on the SAN. The boot policy can include a primary and a secondary SAN boot. If the primary boot fails, the server attempts to boot from the secondary.

Cisco recommends using a SAN boot, because it offers the most service profile mobility within the system. If you boot from the SAN when you move a service profile from one server to another, the new server boots from the same operating system image. Therefore, the new server appears as the same server to the network.

To use a SAN boot, ensure that the following is configured:

- The Cisco UCS domain must be able to communicate with the SAN storage device that hosts the operating system image.
- A boot target LUN (Logical Unit Number) on the device where the operating system image is located.



Note SAN boot is not supported on Gen-3 Emulex adapters on Cisco UCS blade and rack servers.

Configuring a SAN Boot for a Boot Policy

You can also create a local boot policy that is restricted to a service profile or service profile template. However, Cisco recommends that you create a global boot policy that can be included in multiple service profiles or service profile templates.



Tip If you configure a local disk and a SAN LUN for the boot order storage type and the operating system or logical volume manager (LVM) is configured incorrectly, the server might boot from the local disk rather than the SAN LUN.

For example, on a server with Red Hat Linux installed, where the LVM is configured with default LV names and the boot order is configured with a SAN LUN and a local disk, Linux reports that there are two LVs with the same name and boots from the LV with the lowest SCSI ID, which could be the local disk.

This procedure continues directly from [Creating a Boot Policy, on page 118](#).

Before you begin



Note If you are creating a boot policy that boots the server from a SAN LUN and you require reliable SAN boot operations, Cisco recommends that you first remove all local disks and other SAN LUNs from the boot policy in the server service profile.

This does not apply to the UCS Mini Series.

Procedure

Step 1 Click the down arrows to expand the **vHBAs** area.

Step 2 Click the **Add SAN Boot** link.

Step 3 In the **Add San Boot** dialog box, specify the vHBA and type, then click **OK**.

You can specify a **Primary** or a **Secondary** SAN boot. If the primary boot fails, the server attempts to boot from the secondary. The **Any** option is for unsupported adapters that connect directly to the SAN storage device and bypasses UCS Manager. Do not use the **Any** option with SAN boot for a supported set of adapters which are managed by UCSM. For unsupported adapters, use the instructions from the vendor to configure the adaptor for booting.

Step 4 If this vHBA points to a bootable SAN image, click the **Add SAN Boot Target** link and, in the **Add SAN Boot Target** dialog box, specify the boot target LUN, boot target WWPN, and type, then click **OK**:

Step 5 Do one of the following:

- Add another boot device to the **Boot Order** table.
- Click **OK** to finish.

What to do next

Include the boot policy in a service profile and template.

After a server is associated with a service profile that includes this boot policy, you can verify the actual boot order in the **Boot Order Details** area on the **General** tab for the server.

iSCSI Boot

iSCSI boot enables a server to boot its operating system from an iSCSI target machine located remotely over a network.

iSCSI boot is supported on the following Cisco UCS VIC adapters:

- Cisco UCS VIC 1200 series
- Cisco UCS VIC 1300 series
- Cisco UCS VIC 1400 series

There are prerequisites that must be met before you configure iSCSI boot. For a list of prerequisites, see [iSCSI Boot Guidelines and Prerequisites, on page 122](#).

For a high-level procedure for implementing iSCSI boot, see [Configuring iSCSI Boot, on page 124](#).

iSCSI Boot Process

Cisco UCS Manager uses the iSCSI vNIC and iSCSI boot information created for the service profile in the association process to program the adapter, located on the server. After the adapter is programmed, the server reboots with the latest service profile values. After the power on self-test (POST), the adapter attempts to initialize using these service profile values. If the adapter can use the values and log in to its specified target, the adapter initializes and posts an iSCSI Boot Firmware Table (iBFT) to the host memory and a valid bootable LUN to the system BIOS. The iBFT that is posted to the host memory contains the initiator and target configuration that is programmed on the primary iSCSI vNIC.



Note Previously, the host could only detect one boot path, determined by which path completed LUN discovery first, and would boot from that path. Now, with two iSCSI boot vNICs configured, the host detects both boot paths. For multipath configurations, a single IQN must be assigned to both boot vNICs. This can be achieved by creating an iSCSI pool within the service profile configuration and updating the initiator name under 'iSCSI vNICs.' If different IQNs are configured on the boot vNICs, the host will boot using the IQN assigned to the boot vNIC with the lower PCI order.

The next step, which is the installation of the operating system (OS), requires an OS that is iBFT capable. During installation of the OS, the OS installer scans the host memory for the iBFT table and uses the information in the iBFT to discover the boot device and create an iSCSI path to the target LUN. Some OSs require a NIC driver to complete this path. If this step is successful, the OS installer finds the iSCSI target LUN on which to install the OS.



Note The iBFT works at the OS installation software level and might not work with HBA mode (also known as TCP offload). Whether iBFT works with HBA mode depends on the OS capabilities during installation. Also, for a server that includes a Cisco UCS M51KR-B Broadcom BCM57711 adapter, the iBFT normally works at a maximum transmission unit (MTU) size of 1500, regardless of the MTU jumbo configuration. If the OS supports HBA mode, you might need to set HBA mode, dual-fabric support, and jumbo MTU size after the iSCSI installation process.

iSCSI Boot Guidelines and Prerequisites

These guidelines and prerequisites must be met before configuring iSCSI boot:

- After the iSCSI boot policies are created, a user with ls-compute privileges can include them in a service profile or service profile template. However, a user with only ls-compute privileges cannot create iSCSI boot policies.
- To set up iSCSI boot from a Windows 2008 server where the second vNIC (failover vNIC) must boot from an iSCSI LUN, consult Microsoft Knowledge Base Article 976042. Microsoft has a known issue where Windows might fail to boot from an iSCSI drive or cause a bugcheck error if the networking hardware is changed. To work around this issue, follow the resolution recommended by Microsoft.
- The storage array must be licensed for iSCSI boot and the array side LUN masking must be properly configured.
- Two IP addresses must be determined, one for each iSCSI initiator. If possible, the IP addresses should be on the same subnet as the storage array. The IP addresses are assigned statically or dynamically using the Dynamic Host Configuration Protocol (DHCP).
- You cannot configure boot parameters in the Global boot policy. Instead, after configuring boot parameters, include the boot policy in the appropriate service profile.
- The operating system (OS) must be iSCSI Boot Firmware Table (iBFT) compatible.
 - For RHEL 7.x, the kernel parameter "rd.iscsi.ibft=1" is required before the installation. If the parameter is not entered, the iSCSI boot may fail.
 - For SLES 12.x, the following guidelines must be followed:
 - Hit "e" on the install disk before loading the kernel, edit the linuxefi (if using EFI) or kernel (if using legacy), and add the kernel parameter "rd.iscsi.ibft=1 rd.iscsi.firmware=1 rd.neednet=1". If the parameter is not entered, the iSCSI boot may fail.
 - On an existing system that uses iSCSI, ensure that the /etc/iscsi/iscsid.conf has node.startup=automatic (not manual). Add this parameter to the /etc/default/grub/ and then run grub2-mkconfig -o /boot/grub2/grub.cfg to rebuild grub config.
- For Cisco UCS M51KR-B Broadcom BCM57711 network adapters:
 - Servers that use iSCSI boot must contain the Cisco UCS M51KR-B Broadcom BCM57711 network adapter. For information on installing or replacing an adapter card, see the *Cisco UCS B250 Extended Memory Blade Server Installation and Service Note*. The service note is accessible from the *Cisco UCS B-Series Servers Documentation Roadmap* at <http://www.cisco.com/go/unifiedcomputing/b-series-doc>.
 - Set the MAC addresses on the iSCSI device.
 - If you are using the DHCP Vendor ID (Option 43), configure the MAC address of an iSCSI device in /etc/dhcpd.conf.
 - HBA mode (also known as TCP offload) and the boot to target setting are supported. However, only Windows OS supports HBA mode during installation.
 - Before installing the OS, disable the boot to target setting in the iSCSI adapter policy, then after installing the OS, re-enable the boot to target setting.



Note Each time you change an adapter policy setting, the adapter reboots to apply the new setting.

- When installing the OS on the iSCSI target, the iSCSI target must be ordered *before* the device where the OS image resides. For example, if you are installing the OS on the iSCSI target from a CD, the boot order should be the iSCSI target and then the CD.
 - After the server is iSCSI booted, do not modify the Initiator Name, Target name, LUN, iSCSI device IP, or Netmask/gateway using the Broadcom tool.
 - Do not interrupt the POST (power on self-test) process or the Cisco UCS M51KR-B Broadcom BCM57711 network adapter will fail to initialize.
- For Cisco UCS M81KR Virtual Interface Card and Cisco UCS VIC-1240 Virtual Interface Card:
For Cisco UCS VIC-1240 Virtual Interface Card:
- Do not set MAC addresses on the iSCSI device.
 - HBA mode and the boot to target setting are *not* supported.
 - When installing the OS on the iSCSI target, the iSCSI target must be ordered *after* the device where the OS image resides. For example, if you are installing the OS on the iSCSI target from a CD, the boot order should be the CD and then the iSCSI target.
 - If you are using the DHCP Vendor ID (Option 43), the MAC address of the overlay vNIC must be configured in `/etc/dhcpd.conf`.
 - After the server is iSCSI booted, do not modify the IP details of the overlay vNIC.
- The VMware ESX/ESXi operating system does not support storing a core dump file to an iSCSI boot target LUN. Dump files must be written to a local disk.

Initiator IQN Configuration

Cisco UCS uses the following rules to determine the initiator IQN for an adapter iSCSI vNIC at the time a service profile is associated with a physical server:

- An initiator IQN at the service profile level *and* at the iSCSI vNIC level cannot be used together in a service profile.
- If an initiator IQN is specified at the service profile level, all of the adaptor iSCSI vNICs are configured to use the same initiator IQN, except in the case of DHCP Option 43, where the initiator IQN is set to empty on the adaptor iSCSI vNIC.
- When an initiator IQN is set at the iSCSI vNIC level, the initiator IQN at the service profile level is removed, if one is present.
- If there are two iSCSI vNIC in a service profile and only one of them has the initiator IQN set, the second one is configured with the default IQN pool. You can change this configuration later. The only exception is if DHCP Option 43 is configured. In this case, the initiator IQN on the second iSCSI vNIC is removed during service profile association.



Note If you change an iSCSI vNIC to use the DHCP Option 43 by setting the vendor ID, it does not remove the initiator IQN configured at the service profile level. The initiator IQN at the service profile level can still be used by another iSCSI vNIC which does not use the DHCP Option 43.

Enabling MPIO on Windows

You can enable (MPIO) to optimize connectivity with storage arrays.



Note If you change the networking hardware, Windows might fail to boot from an iSCSI drive. For more information, see [Microsoft support Article ID: 976042](#).

Before you begin

The server on which you enable the Microsoft Multipath I/O (MPIO) must have a Cisco VIC driver.

If there are multiple paths configured to the boot LUN, only one path should be enabled when the LUN is installed.

Procedure

-
- Step 1** In the service profile associated with the server, configure the primary iSCSI vNIC.
For more information, see [Creating an iSCSI vNIC for a Service Profile, on page 131](#).
 - Step 2** Using the primary iSCSI vNIC, install the Windows operating system on the iSCSI target LUN.
 - Step 3** After Windows installation completes, enable MPIO on the host.
 - Step 4** In the service profile associated with the server, add the secondary iSCSI vNIC to the boot policy.
For more information, see [Creating an iSCSI Boot Policy, on page 130](#).
-

Configuring iSCSI Boot

When you configure an adapter or blade in Cisco UCS to iSCSI boot from a LUN target, complete all of the following steps.

Procedure

	Command or Action	Purpose
Step 1	(Optional) Configure the iSCSI boot adapter policy.	For more information, see Creating an iSCSI Boot Policy, on page 130

	Command or Action	Purpose
Step 2	(Optional) Configure the authentication profiles for the initiator and target.	For more information, see Creating an iSCSI Authentication Profile, on page 128
Step 3	(Optional) To configure the iSCSI initiator to use an IP address from a pool of IP addresses, add a block of IP addresses to the iSCSI initiator pool.	For more information, see Creating an iSCSI Initiator IP Pool, on page 129
Step 4	Create a boot policy that can be used in any service profile. Alternatively, you can create a local boot policy only for the specific service policy. However, Cisco recommends that you create a boot policy that can be shared with multiple service profiles.	For more information about creating a boot policy that can be used in any service profile, see Creating an iSCSI Boot Policy, on page 130 .
Step 5	If you created a boot policy that can be used in any service profile, assign it to the service profile. Otherwise, proceed to the next step.	You can assign the boot policy to the service profile while configuring the iSCSI boot and vNIC parameters in the service profile in step 7.
Step 6	Create an iSCSI vNIC in a service profile.	For more information, see Creating an iSCSI vNIC for a Service Profile, on page 131
Step 7	Configure the iSCSI boot parameters, including the iSCSI qualifier name (IQN), initiator, target interfaces, and iSCSI vNIC parameters in a service profile in expert mode or service profile template.	For more information, see Creating a Service Profile with the Expert Wizard, on page 159 or Creating a Service Profile Template, on page 175 , respectively.
Step 8	Verify the iSCSI boot operation.	For more information, see Verifying iSCSI Boot .
Step 9	Before installing the OS, ensure that the OS is iSCSI Boot Firmware Table (iBFT) compatible. <ul style="list-style-type: none"> • For RHEL 7.x, the kernel parameter "rd.iscsi.ibft=1" is required before installing the OS. • For SLES 12.x, hit "e" on the install disk before loading the kernel, edit the linuxefi (if using EFI) or kernel (if using legacy), and add the kernel parameter "rd.iscsi.ibft=1 rd.iscsi.firmware=1 rd.neednet=1". 	If the correct parameter is not entered, the iSCSI boot operation may fail.
Step 10	Install the OS on the server.	For more information, see one of the following guides: <ul style="list-style-type: none"> • <i>Cisco UCS B-Series Blade Servers VMware Installation Guide</i>

	Command or Action	Purpose
		<ul style="list-style-type: none"> • <i>Cisco UCS B-Series Blade Servers Linux Installation Guide</i> • <i>Cisco UCS B-Series Blade Servers Windows Installation Guide</i>
Step 11	Boot the server.	

Creating an iSCSI Adapter Policy

Procedure

-
- Step 1** In the **Navigation** pane, click **Servers**.
 - Step 2** Expand **Servers > Policies**.
 - Step 3** Expand the node for the organization where you want to create the policy.
If the system does not include multi tenancy, expand the **root** node.
 - Step 4** Right-click **Adapter Policies** and choose **Create iSCSI Adapter Policy**.
 - Step 5** In the **Create iSCSI Adapter Policy** dialog box, complete the following fields:

Name	Description
Name field	The name of the policy. This name can be between 1 and 16 alphanumeric characters. You cannot use spaces or any special characters other than - (hyphen), _ (underscore), : (colon), and . (period), and you cannot change this name after the object is saved.
Connection Timeout field	The number of seconds to wait until Cisco UCS assumes that the initial login has failed and the iSCSI adapter is unavailable. Enter an integer between 0 and 255. If you enter 0, Cisco UCS uses the value set in the adapter firmware.
LUN Busy Retry Count field	The number of times to retry the connection in case of a failure during iSCSI LUN discovery. Enter an integer between 0 and 60. If you enter 0, Cisco UCS uses the value set in the adapter firmware.
DHCP Timeout field	The number of seconds to wait before the initiator assumes that the DHCP server is unavailable. Enter an integer between 60 and 300 (default: 60 seconds).

Name	Description
Enable TCP Timestamp check box	<p>Check this box if you want to use a TCP Timestamp. With this setting, transmitted packets are given a time stamp of when the packet was sent so that the packet's round-trip time can be calculated, when needed.</p> <p>Note This option only applies to servers with the Cisco UCS NIC M51KR-B adapter.</p>
HBA Mode check box	<p>Check this box to enable HBA mode (also known as TCP offload).</p> <p>Important This option should only be enabled for servers with the Cisco UCS NIC M51KR-B adapter running the Windows operating system.</p>
Boot to Target check box	<p>Check this box to boot from the iSCSI target.</p> <p>Note This option only applies to servers with the Cisco UCS NIC M51KR-B adapter. It should be disabled until you have installed an operating system on the server.</p>
Owner field	<p>This can be one of the following:</p> <ul style="list-style-type: none"> • Local—This policy is available only to service profiles and service profile templates in this Cisco UCS domain. • Pending Global—Control of this policy is being transferred to Cisco UCS Central. Once the transfer is complete, this policy will be available to all Cisco UCS domains registered with Cisco UCS Central. • Global—This policy is managed by Cisco UCS Central. Any changes to this policy must be made through Cisco UCS Central.

Step 6 Click **OK**.

What to do next

Include the adapter policy in a service profile and template.

Deleting an iSCSI Adapter Policy

Procedure

- Step 1** In the **Navigation** pane, click **Servers**.
- Step 2** Expand **Servers > Policies**.

- Step 3** Expand the node for the organization where you want to create the policy.
If the system does not include multi tenancy, expand the **root** node.
- Step 4** Expand the **Adapter Policies** node.
- Step 5** Right-click the adapter policy and choose **Delete**.
- Step 6** If a confirmation dialog box displays, click **Yes**.

Creating an iSCSI Authentication Profile

For iSCSI boot, you need to create both an initiator and a target iSCSI authentication profile.

Procedure

- Step 1** In the **Navigation** pane, click **Servers**.
- Step 2** Expand **Servers > Policies**.
- Step 3** Expand the node for the organization where you want to create the policy.
If the system does not include multi tenancy, expand the **root** node.
- Step 4** Right-click **iSCSI Authentication Profiles** and choose **Create iSCSI Authentication Profile**.
- Step 5** In the **Create Authentication Profile** dialog box, complete the following fields:

Name	Description
Name field	The name of the authentication profile. This name can be between 1 and 16 alphanumeric characters. You cannot use spaces or any special characters other than - (hyphen), _ (underscore), : (colon), and . (period), and you cannot change this name after the object is saved.
User Id field	The user Id associated with this profile. Enter between 1 and 128 characters, spaces, or special characters.
Password field	The password associated with this profile. Enter between 12 and 16 characters, including special characters.
Confirm Password field	The password again for confirmation purposes.

- Step 6** Click **OK**.

What to do next

Include the authentication profile in a service profile and template.

Deleting an iSCSI Authentication Profile

Procedure

-
- Step 1** In the **Navigation** pane, click **Servers**.
 - Step 2** Expand **Servers > Policies**.
 - Step 3** Expand the node for the organization where you want to create the policy.
If the system does not include multi tenancy, expand the **root** node.
 - Step 4** Expand the **iSCSI Authentication Profiles** node.
 - Step 5** Right-click the IP pool you want to delete and choose **Delete**.
 - Step 6** If a confirmation dialog box displays, click **Yes**.
-

Creating an iSCSI Initiator IP Pool

You can create a group of IP addresses to be used for iSCSI boot. Cisco UCS Manager reserves the block of IPv4 addresses you specify.

The IP pool must not contain any IP addresses that were assigned as static IP addresses for a server or service profile.

Procedure

-
- Step 1** In the **Navigation** pane, click **LAN**.
 - Step 2** Expand **LAN > Pools**.
 - Step 3** Expand the node for the organization where you want to create the pool.
If the system does not include multi tenancy, expand the **root** node.
 - Step 4** Expand the **IP Pools** node.
 - Step 5** Right-click **IP Pool iscsi-initiator-pool** and choose **Create Block of IPv4 Addresses**.
 - Step 6** In the **Create a Block of IPv4 Addresses** dialog box, complete the following fields:

Name	Description
Name column	The range of IPv4 addresses assigned to the block.
From column	The first IPv4 address in the block.
To column	The last IPv4 address in the block.
Subnet column	The subnet mask associated with the IPv4 addresses in the block.
Default Gateway column	The default gateway associated with the IPv4 addresses in the block.

Name	Description
Primary DNS column	The primary DNS server that this block of IPv4 addresses should access.
Secondary DNS column	The secondary DNS server that this block of IPv4 addresses should access.

Step 7 Click **OK**.

What to do next

Configure one or more service profiles or service profile templates to obtain the iSCSI initiator IP address from the iSCSI initiator IP pool.

Creating an iSCSI Boot Policy

You can add up to two iSCSI vNICs per boot policy. One vNIC acts as the primary iSCSI boot source, and the other acts as the secondary iSCSI boot source.

Procedure

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

Step 2 Expand **Servers > Policies**.

Step 3 Expand the node for the organization where you want to create the policy.

If the system does not include multi tenancy, expand the **root** node.

Step 4 Right-click **Boot Policies** and choose **Create Boot Policy**.

The **Create Boot Policy** wizard displays.

Step 5 Enter a unique name and description for the policy.

This name can be between 1 and 16 alphanumeric characters. You cannot use spaces or any special characters other than - (hyphen), _ (underscore), : (colon), and . (period), and you cannot change this name after the object is saved.

Step 6 (Optional) To reboot a server that uses this boot policy after you make changes to the boot order, check the **Reboot on Boot Order Change** check box.

In the Cisco UCS Manager GUI, if the **Reboot on Boot Order Change** check box is checked for a boot policy, and if CD-ROM or Floppy is the last device in the boot order, deleting or adding the device does not directly affect the boot order and the server does not reboot.

Note

This applies only to servers using the standard boot order.

Step 7 (Optional) If desired, check the **Enforce vNIC/vHBA/iSCSI Name** check box.

- If checked, Cisco UCS Manager displays a configuration error and reports whether one or more of the vNICs, vHBAs, or iSCSI vNICs listed in the **Boot Order** table match the server configuration in the service profile.
- If not checked, Cisco UCS Manager uses the vNICs or vHBAs (as appropriate for the boot option) from the service profile.

- Step 8** To add a iSCSI boot to the boot policy, do the following:
- Click the down arrows to expand the iSCSI vNICs area.
 - Click the **Add iSCSI Boot** link.
 - In the **Add iSCSI Boot** dialog box, enter a name for the iSCSI vNIC, and click **OK**.
 - Repeat steps b and c to create another iSCSI vNIC.

What to do next

Include the boot policy in a service profile and template.

After a server is associated with a service profile that includes this boot policy, you can verify the actual boot order in the **Boot Order Details** area on the **General** tab for the server.

Creating an iSCSI vNIC for a Service Profile

Procedure

- Step 1** In the **Navigation** pane, click **Servers**.
- Step 2** Expand **Servers > Service Profiles**.
- Step 3** Expand the node for the organization that contains the service profile for which you want to create an iSCSI vNIC.
- Step 4** Expand the service profile for which you want to create a iSCSI vNIC.
- Step 5** Right-click the **iSCSI vNICs** node and choose **Create vNICs**.
- Step 6** In the **Create iSCSI vNIC** dialog box, complete the following fields:

Name	Description
Name field	The name of the iSCSI vNIC. This name can be between 1 and 16 alphanumeric characters. You cannot use spaces or any special characters other than - (hyphen), _ (underscore), : (colon), and . (period), and you cannot change this name after the object is saved.
Overlay vNIC drop-down list	The LAN vNIC associated with this iSCSI vNIC, if any. The selection of an Overlay vNIC may influence the available VLAN options, especially for specific Virtual Interface Cards.

Name	Description
iSCSI Adapter Policy drop-down list	The iSCSI adapter policy associated with this iSCSI vNIC, if any. This drop-down list displays available policies, including default policies such as 'default' (supports IPv4 address). These policies define the iSCSI adapter's behavior and settings.
Create iSCSI Adapter Policy link	Click this link to create a new iSCSI adapter policy that will be available to all iSCSI vNICs.
MAC Address field	The MAC address associated with this iSCSI vNIC, if any. If the MAC address is not set, the Cisco UCS Manager GUI displays Derived .
MAC Pool field	The MAC pool associated with this iSCSI vNIC, if any.
VLAN drop-down list	<p>The virtual LAN associated with this iSCSI vNIC. The default VLAN is default.</p> <p>Note For the Cisco UCS M81KR Virtual Interface Card and the Cisco UCS VIC-1240 Virtual Interface Card, the VLAN that you specify must be the same as the native VLAN on the overlay vNIC.</p> <p>For the Cisco UCS M51KR-B Broadcom BCM57711 Adapter, the VLAN that you specify can be any VLAN assigned to the overlay vNIC.</p>

Step 7 In the **MAC Address Assignment** drop-down list in the **iSCSI MAC Address** area, choose one of the following:

- Leave the MAC address unassigned, **Select (None used by default)**. Select this option if the server that will be associated with this service profile contains a Cisco UCS M81KR Virtual Interface Card adapter or a Cisco UCS VIC-1240 Virtual Interface Card.

Important

If the server that will be associated with this service profile contains a Cisco UCS NIC M51KR-B adapter, you must specify a MAC address.

- A specific MAC address, select **00:25:B5:XX:XX:XX** and enter the address in the **MAC Address** field. To verify that this address is available, click the corresponding link.
- A MAC address from a pool, select the pool name from the list. Each pool name is followed by a pair of numbers in parentheses. The first number is the number of available MAC addresses in the pool and the second is the total number of MAC addresses in the pool.

If this Cisco UCS domain is registered with Cisco UCS Central, there might be two pool categories.

Domain Pools are defined locally in the Cisco UCS domain and **Global Pools** are defined in Cisco UCS Central.

Step 8 (Optional) If you want to create a MAC pool that will be available to all service profiles, click **Create MAC Pool** and complete the fields in the **Create MAC Pool** wizard.

For more information, see the *Creating a MAC Pool* section in *Cisco UCS Manager Network Management Guide, Release 3.2*.

Step 9 Click **OK**.

- Step 10** (Optional) If you want to set or change the initiator name, from the **iSCSI vNICs** tab, click **Reset Initiator Name** or **Change Initiator Name** and complete the fields in the **Change Initiator Name** dialog box or click . For more information, see [Setting the Initiator IQN at the Service Profile Level, on page 133](#).
-

Deleting an iSCSI vNIC from a Service Profile

Before you begin

Consider the impact of deleting the iSCSI vNIC. If the service profile is currently associated with a server and the iSCSI vNIC is in use (e.g., for iSCSI boot or data access), deleting it will disrupt the server's connectivity to its iSCSI targets and may lead to service interruption or inability to boot.

It is highly recommended to disassociate the service profile from any server, or power off the associated server, before deleting an iSCSI vNIC to prevent unexpected behavior or data access issues.

Procedure

- Step 1** In the **Navigation** pane, click **Servers**.
 - Step 2** Expand **Servers > Service Profiles**.
 - Step 3** Expand the node for the organization that contains the service profile from which you want to delete an iSCSI vNIC.
 - Step 4** Expand the service profile from which you want to delete an iSCSI vNIC.
 - Step 5** Expand the **iSCSI vNICs** node.
 - Step 6** Right-click the iSCSI vNIC you want to delete and choose **Delete**.
 - Step 7** If a confirmation dialog box displays, click **Yes**.
-

What to do next

After deleting the iSCSI vNIC, if the service profile was associated with a server, you may need to re-associate the service profile or reboot the server for the changes to fully take effect. If the deleted iSCSI vNIC was part of an iSCSI boot policy, ensure that the boot policy is updated accordingly.

Setting the Initiator IQN at the Service Profile Level

Procedure

- Step 1** In the **Navigation** pane, click **Servers**.
- Step 2** Expand **Servers > Service Profiles**.
- Step 3** Expand the desired node for the organization.
- Step 4** Click the service profile with the iSCSI vNIC that you want to change.
- Step 5** In the **Work** pane, click the **iSCSI vNICs** tab.

- Step 6** Click **Reset Initiator Name**.
- Step 7** If a confirmation dialog box displays, click **Yes**.

Changing the Initiator IQN at the Service Profile Level

Procedure

- Step 1** In the **Navigation** pane, click **Servers**.
- Step 2** Expand **Servers > Service Profiles**.
- Step 3** Expand the desired node for the organization.
- Step 4** Click the service profile with the iSCSI vNIC that you want to change.
- Step 5** In the **Work** pane, click the **iSCSI vNICs** tab.
- Step 6** In the **Actions** area, click **Change Initiator Name**.
- Step 7** In the **Change Initiator Name** dialog box, change the values in the following fields

Name	Description
Initiator Name Assignment drop-down list	Choose the IQN initiator name that you want to use from the drop-down list.
Initiator Name field	If you selected a manual initiator name assignment, enter the initiator name.
Create IQN Suffix Pool link	Click to create a new IQN suffix pool.

- Step 8** Click **OK**.

Setting iSCSI Boot Parameters

You can set iSCSI boot parameters, including the boot order, boot policy, iSCSI authentication profile, initiator interface, and target interface for an iSCSI vNIC.

Procedure

- Step 1** In the **Navigation** pane, click **Servers**.
- Step 2** Expand **Servers > Service Profiles**.
- Step 3** Expand the node for the organization that contains the service profile for which you want to create iSCSI boot parameters. If the system does not include multi-tenancy, expand the root node.
- Step 4** Click the service profile for which you want to create iSCSI boot parameters.
- Step 5** Click the **Boot Order** tab.
- Step 6** In the **Specific Boot Policy** area, click the down arrows to expand the **iSCSI vNICs** area.

Step 7 In the **iSCSI vNICs** area, double-click the iSCSI vNICs from which you want to boot the server to add them to the **Boot Order** table.

Step 8 In the **iSCSI vNICs** area, click the **Set Boot Parameters** link.

If there are two iSCSI vNICs, choose the one for which you want to set boot parameters.

Step 9 In the **Set iSCSI Boot Parameters** dialog box, complete the following fields:

Name	Description
Name field	The name of the iSCSI vNIC for which you are setting the boot parameters.
Authentication Profile drop-down list	The name of the associated iSCSI authentication profile.
Create Authentication Profile link	Click this link to create a new iSCSI authentication profile that will be available to all iSCSI vNICs.

Step 10 In the **Initiator Name** area, complete the following fields:

Name	Description
Initiator Name Assignment drop-down list	<p>Select how the iSCSI boot initiator name is assigned. Choose one of the following methods:</p> <ul style="list-style-type: none"> • Manual—You will enter a name in the Initiator Name field. The initiator name can contain up to 223 characters. • Pools—Choose an IQN suffix pool from which the name will be assigned. <p>Note Setting the Initiator Name from the Set iSCSI Boot Parameters dialog box sets the initiator IQN at the iSCSI vNIC level and not at the service profile level. If more than one path is configured, you must set the initiator IQN from the iSCSI vNICs tab or when creating a service profile.</p> <p>If you need to, you can change or reset the initiator name. For more information, see Changing the Initiator IQN at the Service Profile Level, on page 134.</p>
Create IQN Suffix Pool link	Click this link to create a new IQN suffix pool that will be available to all iSCSI vNICs.
Initiator Name field	<p>A regular expression that defines the name of the iSCSI initiator.</p> <p>You can enter any alphanumeric string as well as the following special characters:</p> <ul style="list-style-type: none"> • . (period) • : (colon) • - (dash)

Step 11 From the **Initiator IP Address Policy** drop-down list, choose of the following:

Option	Description
Select (DHCP used by default)	The system selects an interface automatically using DHCP. Proceed to Step 13.
Static	A static IPv4 address is assigned to the iSCSI boot vNIC based on the information entered in this area. Proceed to Step 12.
Pool	An IPv4 address is assigned to the iSCSI boot vNIC from the management IP address pool. Proceed to Step 13.

Step 12 If you chose **Static** from the **Initiator IP Address Policy** drop-down list, complete the following fields:

Name	Description
IPv4 Address field	The IPv4 address assigned to the iSCSI boot vNIC. If you want to specify this address, you must select Static in the Initiator IP Address Policy drop-down list.
Subnet Mask field	The subnet mask associated with the IPv4 address.
Default Gateway field	The default gateway associated with the IPv4 address.
Primary DNS field	The primary DNS server address.
Secondary DNS field	The secondary DNS server address.

Note

Verify the IPv4 address, subnet mask, and default gateway are correctly configured to ensure network reachability to the iSCSI target.

Step 13 For the iSCSI target interface, choose one of the following radio buttons:

Option	Description
iSCSI Static Target Interface	The system creates a static target interface that you need to configure. Proceed to Step 14.
iSCSI Auto Target Interface	The system creates an auto target interface. You need to specify whether the auto target uses an initiator or a DHCP vendor ID. Proceed to Step 16.

Step 14 If you chose **iSCSI Static Target Interface**, in the **Static Target Interface** table, click **Add**.

Step 15 In the **Create iSCSI Static Target** dialog box, complete the following fields:

Name	Description
iSCSI Target Name field	<p>A regular expression that defines the iSCSI Qualified Name (IQN) or Extended Unique Identifier (EUI) name of the iSCSI target.</p> <p>You can enter any alphanumeric characters as well as the following special characters:</p> <ul style="list-style-type: none"> • . (period) • : (colon) • - (dash) <p>Important This name must be properly formatted using standard IQN or EUI guidelines.</p> <p>The following examples show properly formatted iSCSI target names:</p> <ul style="list-style-type: none"> • iqn.2001-04.com.example • iqn.2001-04.com.example:storage:diskarrays-sn-a8675309 • iqn.2001-04.com.example:storage.tape1.sys1.xyz • iqn.2001-04.com.example:storage.disk2.sys1.xyz • eui.02004567A425678D
Priority field	The system-assigned priority for the iSCSI target. This field value is automatically set by the system and cannot be modified by the user.
Port field	<p>The port associated with the iSCSI target.</p> <p>Enter an integer between 1 and 65535. The default value is 3260. This port value (3260) is supported on IPv4 connections.</p>
Authentication Profile drop-down list	The name of the associated iSCSI authentication profile.
Create iSCSI Authentication Profile link	Click this link to create a new iSCSI authentication profile that will be available to all iSCSI vNICs.
LUN Id field	Specifies the LUN (Logical Unit Number) identifier in the iSCSI target. LUN IDs can range from 0 to 65535.

Step 16 If you chose **iSCSI Auto Target Interface**, enter either the initiator name or the DHCP vendor ID in the **DHCP Vendor Id** field. The initiator must have already been configured. The vendor ID can be up to 32 alphanumeric characters.

Step 17 Click **OK**.

Modifying iSCSI Boot Parameters

You can modify iSCSI boot parameters, including the boot order, boot policy, iSCSI authentication profile, initiator interface, and target interface for an iSCSI vNIC.

Procedure

- Step 1** In the **Navigation** pane, click **Servers**.
- Step 2** Expand **Servers > Service Profiles**.
- Step 3** Expand the node for the organization that contains the service profile for which you want to modify iSCSI boot parameters. If the system does not include multi-tenancy, expand the root node.
- Step 4** Click the service profile for which you want to modify iSCSI boot parameters.
- Step 5** Click the **Boot Order** tab.
- Step 6** In the **Specific Boot Policy** area, click the down arrows to expand the **iSCSI vNICs** area.
- Step 7** To add or delete an iSCSI vNIC from the boot order or to change the boot order, do one of the following:
- To add an iSCSI vNIC, in the **iSCSI vNICs** area, double-click an iSCSI vNICs to add it to the **Boot Order** table.
 - To delete an iSCSI vNIC from the boot order, in the **Boot Order** table, select the iSCSI vNIC and click **Delete**.
 - To change the iSCSI vNIC boot order, in the **Boot Order** table, select the iSCSI vNIC and click either **Move Up** or **Move Down**.
- Step 8** To change the boot parameters, in the **iSCSI vNICs** area, click the **Set Boot Parameters** link. If there are two iSCSI vNICs, choose the one for which you want to change boot parameters.
- Step 9** In the **Set iSCSI Boot Parameters** dialog box, change the values in any of the following fields:
- | Name | Description |
|--|---|
| Name field | The name of the iSCSI vNIC for which you are setting the boot parameters. |
| Authentication Profile drop-down list | The name of the associated iSCSI authentication profile. |
| Create Authentication Profile link | Click this link to create a new iSCSI authentication profile that will be available to all iSCSI vNICs. |
- Step 10** In the **Initiator Name** area, complete the following fields:

Name	Description
Initiator Name Assignment drop-down list	<p>Select how the iSCSI boot initiator name is assigned. Choose one of the following methods:</p> <ul style="list-style-type: none"> • Manual—You will enter a name in the Initiator Name field. The initiator name can contain up to 223 characters. • Pools—Choose an IQN suffix pool from which the name will be assigned. <p>Note Setting the Initiator Name from the Set iSCSI Boot Parameters dialog box sets the initiator IQN at the iSCSI vNIC level and not at the service profile level. If more than one path is configured, you must set the initiator IQN from the iSCSI vNICs tab or when creating a service profile.</p> <p>If you need to, you can change or reset the initiator name. For more information, see Changing the Initiator IQN at the Service Profile Level, on page 134.</p>
Create IQN Suffix Pool link	Click this link to create a new IQN suffix pool that will be available to all iSCSI vNICs.
Initiator Name field	<p>A regular expression that defines the name of the iSCSI initiator.</p> <p>You can enter any alphanumeric string as well as the following special characters:</p> <ul style="list-style-type: none"> • . (period) • : (colon) • - (dash)

Step 11 From the **Initiator IP Address Policy** drop-down list, change the selection to one of the following:

Option	Description
Select (DHCP used by default)	The system selects an interface automatically using DHCP. Proceed to Step 13.
Static	A static IPv4 address is assigned to the iSCSI boot vNIC based on the information entered in this area. Proceed to Step 12.
Pool	An IPv4 address is assigned to the iSCSI boot vNIC from the management IP address pool. Proceed to Step 13.

Step 12 If you chose **Static** from the **Initiator IP Address Policy** drop-down list, complete or change the following fields

Name	Description
IPv4 Address field	The IPv4 address assigned to the iSCSI boot vNIC. If you want to specify this address, you must select Static in the Initiator IP Address Policy drop-down list.
Subnet Mask field	The subnet mask associated with the IPv4 address.
Default Gateway field	The default gateway associated with the IPv4 address.
Primary DNS field	The primary DNS server address.
Secondary DNS field	The secondary DNS server address.

Step 13 For the iSCSI target interface, choose one of the following radio buttons:

Option	Description
iSCSI Static Target Interface	The system creates a static target interface that you need to configure. Proceed to Step 14.
iSCSI Auto Target Interface	The system creates an auto target interface. You need to specify whether the auto target uses an initiator or a DHCP vendor ID. Proceed to Step 15.

Step 14 If you chose **iSCSI Static Target Interface**, do one of the following in the **Static Target Interface** table:

- To add an iSCSI static target interface, click **Add** or to modify an iSCSI target interface, select the iSCSI target interface that you want to change and click **Modify**. Then and complete or change the following fields in the **Create iSCSI Static Target** dialog box:

Name	Description
iSCSI Target Name field	<p>A regular expression that defines the iSCSI Qualified Name (IQN) or Extended Unique Identifier (EUI) name of the iSCSI target.</p> <p>You can enter any alphanumeric characters as well as the following special characters:</p> <ul style="list-style-type: none"> • . (period) • : (colon) • - (dash) <p>Important This name must be properly formatted using standard IQN or EUI guidelines.</p> <p>The following examples show properly formatted iSCSI target names:</p> <ul style="list-style-type: none"> • iqn.2001-04.com.example • iqn.2001-04.com.example:storage.diskarrays-sn-a8675309 • iqn.2001-04.com.example:storage.tape1.sys1.xyz • iqn.2001-04.com.example:storage.disk2.sys1.xyz • eui.02004567A425678D
Priority field	The system-assigned priority for the iSCSI target. This field value is automatically set by the system and cannot be modified by the user.
Port field	The port associated with the iSCSI target. Enter an integer between 1 and 65535. The default value is 3260 . This port value (3260) is supported on IPv4 connections.
Authentication Profile drop-down list	The name of the associated iSCSI authentication profile.
Create iSCSI Authentication Profile link	Click this link to create a new iSCSI authentication profile that will be available to all iSCSI vNICs.
LUN Id field	Specifies the LUN (Logical Unit Number) identifier in the iSCSI target. LUN IDs can range from 0 to 65535.

- To delete an iSCSI target interface, select the iSCSI target interface that you want to delete and click **Delete**.

Note

If you have two iSCSI static targets and you delete the first priority target, the second priority target becomes the first priority target, although Cisco UCS Manager still shows it as the second priority target.

- Step 15** If you chose **iSCSI Auto Target Interface**, change the entry to either the initiator name or the DHCP vendor ID in the **DHCP Vendor Id** field. The initiator must have already been configured. The vendor ID can be up to 32 alphanumeric characters.
- Step 16** Click **OK**.

IQN Pools

An IQN pool is a collection of iSCSI Qualified Names (IQNs) for use as initiator identifiers by iSCSI vNICs in a Cisco UCS domain.

IQN pool members are of the form *prefix:suffix:number*, where you can specify the prefix, suffix, and a block (range) of numbers.

An IQN pool can contain more than one IQN block, with different number ranges and different suffixes, but sharing the same prefix.

Creating an IQN Pool



Note In most cases, the maximum IQN size (prefix + suffix + additional characters) is 223 characters. When using the Cisco UCS NIC M51KR-B adapter, you must limit the IQN size to 128 characters.

Procedure

- Step 1** In the **Navigation** pane, click **SAN**.
- Step 2** Expand **SAN > Pools**.
- Step 3** Expand the node for the organization where you want to create the pool.
If the system does not include multi tenancy, expand the **root** node.
- Step 4** Right-click **IQN Pools** and select **Create IQN Suffix Pool**.
- Step 5** In the **Define Name and Description** page of the **Create IQN Suffix Pool** wizard, fill in the following fields:

Field	Description
Name	The name of the iSCSI Qualified Name (IQN) pool. This name can be between 1 and 32 alphanumeric characters. You cannot use spaces or any special characters other than - (hyphen), _ (underscore), : (colon), and . (period), and you cannot change this name after the object is saved.
Description	The user-defined description of the pool. Enter up to 256 characters. You can use any characters or spaces except ` (accent mark), \ (backslash), ^ (carat), " (double quote), = (equal sign), > (greater than), < (less than), or ' (single quote).

Field	Description
Prefix	The prefix for any IQN blocks created for this pool. Enter from 1 to 150 characters. You can use any letter or number, as well as the special characters . (period), : (colon), and - (hyphen). For example, you could use iqn1.alpha.com .
Assignment Order field	This can be one of the following: <ul style="list-style-type: none"> • Default—Cisco UCS Manager selects a random identity from the pool. • Sequential—Cisco UCS Manager selects the lowest available identity from the pool.

Step 6 Click **Next**.

Step 7 In the **Add IQN Blocks** page of the **Create IQN Suffix Pool** wizard, click **Add**.

Step 8 In the **Create a Block of IQN Suffixes** dialog box, fill in the following fields:

Name	Description
Suffix field	The suffix for this block of iSCSI Qualified Names (IQNs). Enter from 1 to 64 characters. You can use any letter or number, as well as the special characters . (period), : (colon), and - (hyphen). For example, you could use alphadc-1 .
From field	The first suffix number in the block.
Size field	The number of suffixes in the block.

Step 9 Click **OK**.

Step 10 Click **Finish** to complete the wizard.

What to do next

Include the IQN suffix pool in a service profile and template.

Adding a Block to an IQN Pool

Procedure

-
- Step 1** In the **Navigation** pane, click **SAN**.
- Step 2** Expand **SAN > Pools**.
- Step 3** Expand the node for the organization containing the pool.
If the system does not include multi tenancy, expand the **root** node.

- Step 4** Expand the **IQN Pools** node.
- Step 5** Right-click the desired IQN pool and select **Create a Block of IQN Suffixes**.
- Step 6** In the **Create a Block of IQN Suffixes** dialog box, fill in the following fields:

Name	Description
Suffix field	The suffix for this block of iSCSI Qualified Names (IQNs). Enter from 1 to 64 characters. You can use any letter or number, as well as the special characters . (period), : (colon), and - (hyphen). For example, you could use alphadc-1 .
From field	The first suffix number in the block.
Size field	The number of suffixes in the block.

- Step 7** Click **OK**.
-

Deleting a Block from an IQN Pool

If you delete an address block from a pool, Cisco UCS Manager does not reallocate any addresses in that block that were assigned to vNICs or vHBAs. All assigned addresses from a deleted block remain with the vNIC or vHBA to which they are assigned until one of the following occurs:

- The associated service profiles are deleted.
- The vNIC or vHBA to which the address is assigned is deleted.
- The vNIC or vHBA is assigned to a different pool.

Procedure

- Step 1** In the **Navigation** pane, click **SAN**.
- Step 2** Expand **SAN > Pools**.
- Step 3** Expand the node for the organization containing the pool.
If the system does not include multi tenancy, expand the **root** node.
- Step 4** Expand the **IQN Pools** node.
- Step 5** Choose the IQN pool for which you want to delete a block of IQN suffixes.
- Step 6** In the **Work pane**, click the **IQN Blocks** tab.
- Step 7** Right-click the block to be deleted and select **Delete**.
- Step 8** Click **Yes** to confirm the deletion.
- Step 9** Click **Save Changes**.
-

Deleting an IQN Pool

If you delete a pool, Cisco UCS Manager does not reallocate any addresses from that pool that were assigned to vNICs or vHBAs. All assigned addresses from a deleted pool remain with the vNIC or vHBA to which they are assigned until one of the following occurs:

- The associated service profiles are deleted.
- The vNIC or vHBA to which the address is assigned is deleted.
- The vNIC or vHBA is assigned to a different pool.

Procedure

- Step 1** In the **Navigation** pane, click **SAN**.
- Step 2** Expand **SAN > Pools**.
- Step 3** Expand the node for the organization containing the pool.
If the system does not include multi tenancy, expand the **root** node.
- Step 4** Expand the **IQN Pools** node.
- Step 5** Right-click the pool that you want to delete and select **Delete**.
- Step 6** If a confirmation dialog box displays, click **Yes**.
-

LAN Boot

You can configure a boot policy to boot one or more servers from a centralized provisioning server on the LAN. A LAN (or PXE) boot is frequently used to install operating systems on a server from that LAN server.

You can add more than one type of boot device to a LAN boot policy. For example, you could add a local disk or virtual media boot as a secondary boot device.

Configuring a LAN Boot for a Boot Policy

You can also create a local boot policy that is restricted to a service profile or service profile template. However, Cisco recommends that you create a global boot policy that can be included in multiple service profiles or service profile templates.

You can add more than one type of boot device to a boot policy. For example, you can add a local disk or virtual media boot as a secondary boot device.

This procedure continues directly from [Creating a Boot Policy, on page 118](#).

Procedure

- Step 1** Click the down arrows to expand the **vNICs** area.

- Step 2** Click the **Add LAN Boot** link.
- Step 3** In the **Add LAN Boot** dialog box, enter the **vNIC** name, in the **vNIC** field for LAN Boot, select the **IP address Type**—None, IPv4, or IPv6, and then click **OK**.
- Step 4** Do one of the following:
- Add another boot device to the **Boot Order** table.
 - Click **OK** to finish.
-

What to do next

Include the boot policy in a service profile and template.

After a server is associated with a service profile that includes this boot policy, you can verify the actual boot order in the **Boot Order Details** area on the **General** tab for the server.

Local Devices Boot

Cisco UCS Manager allows you to boot from different local devices.



Note For Cisco UCS M3 and higher blade and rack servers using enhanced boot order, you can select both top-level and second-level boot devices.



Note When there are more than one boot options provided under same Controller, the boot options is considered as follows instead of the boot order set in Cisco UCS Manager:

- When OS is installed or booted, for UEFI Boot, the installed OS will push its boot option to zero priority (Top Priority) irrespective of the set boot options in Cisco UCS Manager.
- The boot order will be based on the Boot Device enumeration set by BIOS and on how controller exposes the device to host (or as provided in Cisco UCS Manager).

Local Disk Boot

If a server has a local drive, you can configure a boot policy to boot the server from the top-level local disk device or from any of the following second-level devices:

- Local LUN—Enables boot from local disk or local LUN.
- Local JBOD—Enables boot from a bootable JBOD.
- SD card—Enables boot from SD card.
- Internal USB—Enables boot for internal USB.
- External USB—Enables boot from external USB.
- Embedded Local LUN—Enables boot from the embedded local LUN on all Cisco UCS M4, M5 servers.

- Embedded Local Disk—Enables boot from the embedded local disk on all Cisco UCS M4, M5 servers.



Note For Cisco UCS C125 M5 Servers, if there is no separate PCIe storage controller, then do not use this option. Instead, use **Add Local Disk** option.



Note Second-level devices are only available for Cisco UCS M3 and higher blade and rack servers using enhanced boot order.

Virtual Media Boot

You can configure a boot policy to boot one or more servers from a virtual media device that is accessible from the server. A virtual media device mimics the insertion of a physical CD/DVD disk (read-only) or floppy disk (read-write) into a server. This type of server boot is typically used to manually install operating systems on a server.



Note Second-level devices are only available for Cisco UCS M3 and higher blade and rack servers using enhanced boot order.

Remote Virtual Drive Boot

You can configure a boot policy to boot one or more servers from a remote virtual drive that is accessible from the server.

NVMe Boot

Beginning with release 3.2(1) Cisco UCS Manager provides the option of adding an NVMe device to the Boot policy for M5 blade and rack servers. BIOS enumerates the NVMe devices present and boots to the first NVMe device having UEFI capable OS installed on it.

Cisco Boot Optimized M.2 RAID Controller

Beginning with 4.0(4a), Cisco UCS Manager supports Cisco boot optimized M.2 RAID controller based off Marvell 88SE92xx PCIe to SATA 6Gb/s controller (UCS-M2-HWRAID). BIOS enumerates the M.2 SATA drives installed on this controller followed by the front panel SATA drives to boot from the first SATA device having UEFI capable OS installed on it

Configuring a Local Disk Boot for a Boot Policy

You can also create a local boot policy that is restricted to a service profile or service profile template. However, Cisco recommends that you create a global boot policy that can be included in multiple service profiles or service profile templates.

You can add more than one type of boot device to a boot policy. For example, you could add an SD card boot as a secondary boot device.

This procedure continues directly from [Creating a Boot Policy, on page 118](#).

Procedure

-
- Step 1** Expand the **Local Devices** area.
- Step 2** Click any of the following links to add the device to the **Boot Order** table:
- **Add Local Disk** or
 - **Add Local LUN**
 - **Add Local JBOD**
 - **Add SD Card**
 - **Add Internal USB**
 - **Add External USB**
 - **Add Embedded Local LUN**
 - **Add Embedded Local Disk**

Important

In a setup with the Cisco Boot Optimized M.2 RAID Controller (UCS-M2-HWRAID), in the **Add Embedded Local Disk** dialog box, select **Any** to add the disk. Do not select **Primary** or **Secondary**.

Note

For Cisco UCS M3 and higher blade and rack servers using enhanced boot order, you can select both top-level and second-level boot devices.

- Step 3** Do one of the following:
- Add another boot device to the **Boot Order** table.
 - Click **OK** to finish.

What to do next

Include the boot policy in a service profile and template.

After a server is associated with a service profile that includes this boot policy, you can verify the actual boot order in the **Boot Order Details** area on the **General** tab for the server.

Configuring a Virtual Media Boot for a Boot Policy

You can also create a local boot policy that is restricted to a service profile or service profile template. However, Cisco recommends that you create a global boot policy that can be included in multiple service profiles or service profile templates.

You can add more than one type of boot device to a boot policy. For example, you could add a local disk boot as a second boot device.



Note Virtual Media requires the USB to be enabled. If you modify the BIOS settings that affect the USB functionality, you also affect the Virtual Media. Therefore, Cisco recommends that you leave the following USB BIOS defaults for best performance:

- Make Device Non Bootable—set to **disabled**
- USB Idle Power Optimizing Setting—set to **high-performance**

This procedure continues directly from [Creating a Boot Policy, on page 118](#).

Procedure

Step 1 Click the down arrows to expand the **Local Devices** area.

Step 2 Click any of the following links to add the device to the **Boot Order** table:

- **Add CD/DVD** or
 - **Add Local CD/DVD**
 - **Add Remote CD/DVD** (For KVM CD/DVD in rack servers)
In a setup with M5 blade servers, if an ISO is mapped to the KVM console, use only **Add Remote CD/DVD** in **Boot Order**.
- **Add Floppy** or
 - **Add Local Floppy**
 - **Add Remote Floppy**
- **Add Remote Virtual Drive**

Note

For Cisco UCS M3 and higher blade and rack servers using enhanced boot order, you can select both top-level and second-level boot devices.

Step 3 Do one of the following:

- Add another boot device to the **Boot Order** table.
- Click **OK** to finish.

What to do next

Include the boot policy in a service profile and template.

After a server is associated with a service profile that includes this boot policy, you can verify the actual boot order in the **Boot Order Details** area on the **General** tab for the server.

Configuring a NVMe Boot for a Boot Policy

You can also create a local boot policy that is restricted to a service profile or service profile template. However, Cisco recommends that you create a global boot policy that can be included in multiple service profiles or service profile templates.

You can add more than one type of boot device to a boot policy. For example, you could add an SD card boot as a secondary boot device.

This procedure continues directly from [Creating a Boot Policy, on page 118](#).

Procedure

Step 1 Click the down arrows to expand the **Local Devices** area.

Step 2 Click **Add NVMe** to add the device to the Boot Order table.

Note

NVMe boot policy is available only with **Uefi** boot mode, with or without boot security.

Step 3 Do one of the following:

- Add another boot device to the **Boot Order** table.
 - Click **OK** to finish.
-

What to do next

Include the boot policy in a service profile and template.

After a server is associated with a service profile that includes this boot policy, you can verify the actual boot order in the **Boot Order Details** area on the **General** tab for the server.

Adding a Boot Policy to a vMedia Service Profile

This procedure describes how to set the boot policy options for vMedia on the **Server Boot Order** page of the **Create Service Profile (expert)** wizard.

Procedure

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

Step 2 Expand **Servers > Service Profiles**.

Step 3 Expand the node for the organization where you want to create the service profile.

If the system does not include multi tenancy, expand the **root** node.

Step 4 Right-click the organization and select **Create Service Profile (expert)**.
The **Unified Computing System Manager** pane displays.

Step 5 In the **Name** field, enter a unique name that you can use to identify the service profile.

This name can be between 2 and 32 alphanumeric characters. You cannot use spaces or any special characters other than - (hyphen), _ (underscore), : (colon), and . (period), and this name must be unique across all service profiles and service profile templates within the same organization.

This name must be unique within the organization or sub-organization in which you are creating the service profile.

Step 6 From the **UUID Assignment** drop-down list, do one of the following:

Option	Description
Select (pool default used by default)	Assigns a UUID from the default UUID Suffix pool. Continue with Step 8.
Hardware Default	Uses the UUID assigned to the server by the manufacturer. If you choose this option, the UUID remains unassigned until the service profile is associated with a server. At that point, the UUID is set to the UUID value assigned to the server by the manufacturer. If the service profile is later moved to a different server, the UUID is changed to match the new server. Continue with Step 8.
XXXXXXXX-XXXX-XXXX-XXXX-XXXXXXXXXXXX	Uses the UUID that you manually assign. Continue with Step 7.
Pools <i>Pool_Name</i>	Assigns a UUID from the UUID Suffix pool that you select from the list at the bottom of the drop-down list. Each pool name is followed by two numbers in parentheses that show the number of UUIDs available in the pool and the total number of UUIDs in the pool. If you do not want use any of the existing pools, but want to create a pool that all service profiles can access, continue with Step 4. Otherwise, continue with Step 8.

Step 7 (Optional) If you selected the **XXXXXXXX-XXXX-XXXX-XXXX-XXXXXXXXXXXX** option, do the following:

a) In the **UUID** field, enter the valid UUID that you want to assign to the server that uses this service profile.

Step 8 (Optional) If you want to create a new UUID Suffix pool to use in this service profile, click **Create UUID Suffix Pool** and complete the fields in the **Create UUID Suffix Pool** wizard.

For more information, see [Creating a UUID Suffix Pool, on page 105](#).

Step 9 (Optional) In the text box, enter a description of this service profile.

The user-defined description for this service profile.

Enter up to 256 characters. You can use any characters or spaces except ` (accent mark), \ (backslash), ^ (carat), " (double quote), = (equal sign), > (greater than), < (less than), or ' (single quote).

Step 10 Click **Next**.

Step 11 **Navigate** to **Create Service Profile (expert)** and click **Server Boot Order**.
The **Boot Policy** pane displays.

Step 12 From the **Boot Policy** drop-down list, choose one of the following:

Option	Description
Select Boot Policy to use	Assigns the default boot policy to this service profile. Continue with Step 13.
Create a Specific Boot Policy	Enables you to create a local boot policy that can only be accessed by this service profile.
Boot Policies <i>Policy_Name</i>	Assigns an existing boot policy to the service profile. If you choose this option, Cisco UCS Manager displays the details of the policy. If you do not want use any of the existing policies, but want to create a policy that all service profiles can access, click Create Boot Policy . Otherwise, choose a policy from the list and continue with Step 13.

Step 13 If you created a new boot policy accessible to all service profiles and template, choose that policy from the **Boot Policy** drop-down list .

Step 14 Click **Next**.

What to do next

Associate your Service Profile with a Cisco UCS server.

Deleting a Boot Policy

Procedure

-
- Step 1** In the **Navigation** pane, click **Servers**.
 - Step 2** Expand **Servers > Policies > *Organization_Name***.
 - Step 3** Expand the **Boot Policies** node.
 - Step 4** Right-click the policy you want to delete and choose **Delete**.
 - Step 5** If a confirmation dialog box displays, click **Yes**.
-

UEFI Boot Parameters

UEFI boot mode for servers is dependent on information that is stored on the platform hardware. The boot entry, which contains information about the UEFI OS boot loader, is stored in the BIOS flash of the server. In Cisco UCS Manager releases earlier than Release 2.2(4), when a service profile is migrated from one server to another server, the boot loader information is not available on the destination server. Hence, the BIOS cannot load the boot loader information for the server to boot in UEFI boot mode.

Cisco UCSM Release 2.2(4) introduces UEFI boot parameters to provide the BIOS with information about the location of the UEFI OS boot loader on the destination server from where the BIOS loads it. Now, the server can use the boot loader information and boot in UEFI boot mode.

Guidelines and Limitations for UEFI Boot Parameters

- You can configure UEFI boot parameters only if the boot mode is UEFI.
- When you upgrade Cisco UCS Manager to Release 2.2(4) and higher UEFI boot failure during service profile migration is not handled automatically. You must explicitly create the UEFI boot parameters in the target device to successfully boot to the UEFI-capable OS.
- UEFI boot parameters are supported on all M3 and higher servers that support second-level boot order.
- You can specify UEFI boot parameters for the following device types:
 - SAN LUN
 - ISCSI LUN
 - Local LUN
- UEFI boot parameters are specific to each operating system. You can specify UEFI boot parameters for the following operating systems:
 - VMware ESX
 - SuSE Linux
 - Microsoft Windows
 - Red Hat Enterprise Linux 7

Setting UEFI Boot Parameters

Before you begin

Ensure that the **Boot Mode** of the boot policy is **Uefi**.

Procedure

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

- Step 2** Expand **Servers > Policies**.
- Step 3** Expand **Boot Policies** and select the boot policy for which you want to configure UEFI boot parameters.
- Step 4** In the **Work** pane, click the **General** tab.
- Step 5** To set UEFI boot parameters for a LUN, select the LUN in the **Boot Order** area and click **Set Uefi Boot Parameters**.

Important

You can configure UEFI boot parameters only for local LUNs, SAN LUNs, and iSCSI LUNs.

- Step 6** In the **Set Uefi Boot Parameters** dialog box, enter the following information:

Field	Description
Boot Loader Name	Specifies the name of the boot loader. This is a mandatory field. Example—grub.efi
Boot Loader Path	Specifies the path where the boot loader is located. This is a mandatory field. The name of the boot loader must not be included in this field. Only the path must be specified. Example—\EFI\RedHat
Boot Loader Description	Describes the boot loader. This is the human readable name that appears in the F6 boot menu.

- Step 7** Click **OK**.
- Step 8** Click **Save Changes**.

Modifying UEFI Boot Parameters

Before you begin

Ensure that the **Boot Mode** of the boot policy is **Uefi**.

Procedure

- Step 1** In the **Navigation** pane, click **Servers**.
- Step 2** Expand **Servers > Policies**.
- Step 3** Expand **Boot Policies**, and select the boot policy for which you want to modify UEFI boot parameters.
- Step 4** In the **Work** pane, click the **General** tab.
- Step 5** To modify UEFI boot parameters for a LUN with UEFI boot parameters, select the LUN in the **Boot Order** area and click **Modify Uefi Boot Parameters**.
- Important**
You can configure UEFI boot parameters only for local LUNs, SAN LUNs, and iSCSI LUNs.
- Step 6** In the **Modify Uefi Boot Parameters** dialog box, enter the following information:

Field	Description
Boot Loader Name	Specifies the name of the boot loader. This is a mandatory field.
Boot Loader Path	Specifies the path where the boot loader is located. This is a mandatory field.
Boot Loader Description	Describes the boot loader.

Step 7 Click **OK**

Step 8 Click **Save Changes**.



CHAPTER 11

Service Profiles

- [Service Profiles in UCS Manager, on page 157](#)
- [Service Profiles that Override Server Identity, on page 158](#)
- [Service Profiles that Inherit Server Identity, on page 158](#)
- [Guidelines and Recommendations for Service Profiles, on page 159](#)
- [Methods of Creating Service Profiles, on page 159](#)
- [Inband Service Profiles, on page 163](#)
- [Service Profile Tasks, on page 163](#)
- [Service Profile Association, on page 173](#)
- [Service Profile Templates, on page 174](#)
- [Service Profile Template Tasks, on page 179](#)
- [Service Profile Association, on page 183](#)

Service Profiles in UCS Manager

A service profile defines a single server and its storage and networking characteristics. You can create a service profile for Cisco UCS Manager. When a service profile is deployed to a server, UCS Manager automatically configures the server, adapters, fabric extenders, and fabric interconnects to match the configuration specified in the service profile.

A service profile includes four types of information:

- **Server definition:** Defines the resources (e.g. a specific server or a blade inserted to a specific chassis) that are required to apply to the profile.
- **Identity information:** Includes the UUID, MAC address for each virtual NIC (vNIC), and WWN specifications for each HBA.
- **Firmware revision specifications:** Used when a certain tested firmware revision is required to be installed or for some other reason a specific firmware is used.
- **Connectivity definition:** Configures network adapters, fabric extenders, and parent interconnects, however this information is abstract as it does not include the details of how each network component is configured.

The UCS system provides two types of service profiles: Service profiles that inherit server identity and service profiles that override server identity.

Service Profiles that Override Server Identity

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

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

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

- Virtualized identity information, including pools of MAC addresses, WWN addresses, and UUIDs
- Ethernet and Fibre Channel adapter profile policies
- Firmware package policies
- Operating system boot order policies

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

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



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

Service Profiles that Inherit Server Identity

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

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

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

- MAC addresses for the two NICs
- For a converged network adapter or a virtual interface card, the WWN addresses for the two HBAs
- BIOS versions
- Server UUID



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

Guidelines and Recommendations for Service Profiles

In addition to any guidelines or recommendations that are specific to policies and pools included in service profiles and service profile templates, such as the local disk configuration policy, adhere to the following guidelines and recommendations that impact the ability to associate a service profile with a server:

Limit to the Number of vNICs that Can Be Configured on a Rack-Mount Server

You can configure up to 56 vNICs per supported adapter, such as the Cisco UCS P81E Virtual Interface Card (N2XX-ACPCI01), on any rack-mount server that is integrated with Cisco UCS Manager.

No Power Capping Support for Rack-Mount Servers

Power capping is not supported for rack servers. If you include a power control policy in a service profile that is associated with a rack-mount server, the policy is not implemented.

QoS Policy Guidelines for vNICs

You can only assign a QoS policy to a vNIC if the priority setting for that policy is not set to **fc**, which represents the Fibre Channel system class. You can configure the priority for the QoS policy with any other system class.

QoS Policy Guidelines for vHBAs

You can only assign a QoS policy to a vHBA if the priority setting for that policy is set to **fc**, which represents the Fibre Channel system class.

The Host Control setting for a QoS policy applies to vNICs only. It has no effect on a vHBA.

Methods of Creating Service Profiles

Creating a Service Profile with the Expert Wizard

Procedure

- Step 1** In the **Navigation** pane, click **Servers**.
- Step 2** Expand **Servers** > **Service Profiles**.
- Step 3** Expand the node for the organization where you want to create the service profile.
If the system does not include multi tenancy, expand the **root** node.

- Step 4** Right-click the organization and select **Create Service Profile (expert)**.
- Step 5** In the **Identify Service Profile** panel, specify the service profile **Name**, UUID assignment and click **Next**.
You can provide an optional description for this service profile. If the UUID is not available, you can also create a UUID Suffix Pool from this panel.
- Note**
To create a service profile quickly, you can click **Finish** after specifying the name. Cisco UCS Manager creates a new service profile with the specified name and all system default values.
- Step 6** (Optional) In the **Networking** panel, specify the required information for the **Dynamic vNIC Connection Policy** and **LAN Connectivity** sections, then click **Next**.
You can create a dynamic vNIC connection policy and LAN connectivity policy from this panel.
- Note**
Dynamic vNICs, usNICs, and VMQs are not supported when creating a vNIC for a LAN Connectivity Policy. RoCE, VXLAN, NvGRE are not supported when configuring the Ethernet Adapter Policy.
- Step 7** (Optional) In the **Storage** panel, specify the SAN configuration information such as, **Local Storage Policy**, **SAN Connectivity**, **WWNN** and **VSAN**, then click **Next**.
You can create a local disk configuration policy and SAN connectivity policy from this panel.
- Note**
FC vNICs are not supported when configuring vHBA under the Storage Policy.
- Step 8** (Optional) In the **Zoning** panel, specify the required zoning information, then click **Next**.
You can create the vHBA initiator groups from this panel.
- Step 9** (Optional) In the **vNIC/vHBA Placement** panel, specify the placement method and PCI order, then click **Next**.
You can create a placement policy from this panel.
- Step 10** (Optional) In the **Server Boot Order** panel, specify the **Boot Policy** from the drop-down list, then click **Next**.
You can create a boot policy from this panel.
- Step 11** (Optional) In the **Maintenance Policy** panel, specify the maintenance policy, then click **Next**.
You can create a new maintenance policy and specify a maintenance schedule from this panel.
- Step 12** (Optional) In the **Server Assignment** panel, specify the **Server Assignment** from the drop down list and the power state to apply on assignment, then click **Next**.
You can create a server pool or a host firmware package from this panel.
- Step 13** (Optional) In the **Operational Policies** panel, specify the system operational information such as, **BIOS Configuration**, **External IPMI Management Configuration**, **Management IP Address**, **Monitoring Configuration (Thresholds)**, **Power Control Policy Configuration**, and **Scrub Policy**, then click **Finish**.
- Note**
To set up an Outband IPv4 address or an Inband IPv4 or IPv6 address, click the respective tabs and complete the required fields.

If you do not find the policies you need for each of these configurations, you can create them from this panel.

Creating a Service Profile that Inherits Server Identity

Procedure

- Step 1** In the **Navigation** pane, click **Servers**.
- Step 2** Expand **Servers > Service Profiles**.
- Step 3** Expand the node for the organization where you want to create the service profile.
If the system does not include multi tenancy, expand the **root** node.
- Step 4** Right-click the organization and select **Create Service Profile**.
- Step 5** In the **Naming** area of the **Create Service Profile** dialog box, complete the following fields:
- In the **Name** field, enter a unique name that you can use to identify the service profile.

This name can be between 2 and 32 alphanumeric characters. You cannot use spaces or any special characters other than - (hyphen), _ (underscore), : (colon), and . (period), and this name must be unique across all service profiles and service profile templates within the same organization.
 - In the **Description** field, enter a description of this service profile.
- Step 6** In the **vNICs** area of the **Create Service Profile** dialog box, choose the primary and secondary vNICs.
- Step 7** In the **vHBAs** area of the **Create Service Profile** dialog box, choose the primary and secondary vHBAs.
- Step 8** In the **Boot Order** area of the **Create Service Profile** dialog box, choose the primary and secondary boot devices.
- Step 9** (Optional) In the **Select** column of the **Server Association (optional)** area, click the radio button for a server to associate this service profile with that server.
- Step 10** Click **OK**.
-

Creating a Hardware Based Service Profile for a Blade Server

You cannot move a hardware based service profile to another server.

Procedure

- Step 1** In the **Navigation** pane, click **Equipment**.
- Step 2** Expand **Equipment > Chassis > Chassis Number > Servers**.
- Step 3** Choose the server for which you want to create a hardware based service profile.
- Step 4** In the **Work** pane, click the **General** tab.
- Step 5** In the **Actions** area, click **Create Service Profile**.

- Step 6** In the **Create Service Profile for Server** dialog box, do the following:
- From the **Create Service Profile in Organization** drop-down list, select the organization in which you want to create the service profile.
 - Click the **Hardware Based Service Profile** radio button.
 - In the **Name** field, enter a unique name for the service profile.
 This name can be between 2 and 32 alphanumeric characters. You cannot use spaces or any special characters other than - (hyphen), _ (underscore), : (colon), and . (period), and this name must be unique across all service profiles and service profile templates within the same organization.
 - If you want Cisco UCS Manager to create vNICs for the service profile, check the **Create Default vNICs** check box.
 - If you want Cisco UCS Manager to create vHBAs for the service profile, check the **Create Default vHBAs** check box.
 - Click **OK**.

Cisco UCS Manager inherits and automatically applies the identity and configuration information in the server, creates the service profile, and associates it with the server.

Creating a Hardware Based Service Profile for a Rack-Mount Server

You cannot move a hardware based service profile to another server.

Procedure

- Step 1** In the **Navigation** pane, click **Equipment**.
- Step 2** Expand **Equipment** > **Rack Mounts** > **Servers**.
- Note**
 For Cisco UCS C125 M5 Servers, expand **Equipment** > **Rack Mounts** > **Enclosures** > **Rack Enclosure *rack_enclosure_number*** > **Servers**.
- Step 3** Choose the server for which you want to create a hardware based service profile.
- Step 4** In the **Work** pane, click the **General** tab.
- Step 5** In the **Actions** area, click **Create Service Profile**.
- Step 6** In the **Create Service Profile for Server** dialog box, do the following:
- From the **Create Service Profile in Organization** drop-down list, select the organization in which you want to create the service profile.
 - Click the **Hardware Based Service Profile** radio button.
 - In the **Name** field, enter a unique name for the service profile.
 This name can be between 2 and 32 alphanumeric characters. You cannot use spaces or any special characters other than - (hyphen), _ (underscore), : (colon), and . (period), and this name must be unique across all service profiles and service profile templates within the same organization.
 - If you want Cisco UCS Manager to create vNICs for the service profile, check the **Create Default vNICs** check box.

- e) If you want Cisco UCS Manager to create vHBAs for the service profile, check the **Create Default vHBAs** check box.
- f) Click **OK**.

Cisco UCS Manager inherits and automatically applies the identity and configuration information in the server, creates the service profile, and associates it with the server.

Inband Service Profiles

Deleting the Inband Configuration from a Service Profile

This procedure removes the inband management IP address configuration from a service profile. If this action is greyed out, no inband configuration was configured.

Procedure

- Step 1** In the **Navigation** pane, click **Servers**.
- Step 2** Expand **Servers** > **Service Profiles** > *Service_Profile_Name*.
- Step 3** In the **Work** pane, click the **General** tab.
- Step 4** In the **Actions** area, click **Delete Inband Configuration**.
- Step 5** Click **Yes** in the **Delete** confirmation dialog box.

The inband management IP address configuration for the service profile is deleted.

Service Profile Tasks

Renaming a Service Profile

When you rename a service profile, the following occurs:

- Event logs and audit logs that reference the previous name for the service profile are retained under that name.
- A new audit record is created to log the rename operation.
- All records of faults against the service profile under its previous name are transferred to the new service profile name.



Note You cannot rename a service profile with pending changes.

Procedure

- Step 1** In the **Navigation** pane, click **Servers**.
- Step 2** Expand **Servers > Service Profiles**.
- Step 3** Expand the node for the organization that includes the service profile you want to rename.
If the system does not include multi tenancy, expand the **root** node.
- Step 4** Click the service profile you want to rename.
- Step 5** In the **Work** pane, click the **General** tab.
- Step 6** In the **Actions** area, click **Rename Service Profile**.
- Step 7** In the **Rename Service Profile** dialog box, enter the new name for the service profile in the **New Name** field.
This name can be between 2 and 32 alphanumeric characters. You cannot use spaces or any special characters other than - (hyphen), _ (underscore), : (colon), and . (period), and this name must be unique across all service profiles and service profile templates within the same organization.
- Step 8** Click **OK**.
-

Cloning a Service Profile

Procedure

- Step 1** In the **Navigation** pane, click **Servers**.
- Step 2** Expand **Servers > Service Profiles**.
- Step 3** Expand the node for the organization where you want to create the service profile.
If the system does not include multi tenancy, expand the **root** node.
- Step 4** Right-click the service profile you want to clone and select **Create a Clone**.
- Step 5** In the **Create Clone From Service Profile** dialog box:
- Enter the name you want to use for the new profile in the **Clone Name** field.
This name can be between 2 and 32 alphanumeric characters. You cannot use spaces or any special characters other than - (hyphen), _ (underscore), : (colon), and . (period), and this name must be unique across all service profiles and service profile templates within the same organization.
This name must be unique within the organization or sub-organization in which you are creating the service profile.
 - Click **OK**.
- Step 6** Navigate to the service profile you just created and make sure that all options are correct.
-

Changing the UUID in a Service Profile

Procedure

- Step 1** In the **Navigation** pane, click **Servers**.
- Step 2** Expand **Servers** > **Service Profiles**.
- Step 3** Expand the node for the organization that contains the service profile for which you want to change the UUID.
If the system does not include multi tenancy, expand the **root** node.
- Step 4** Choose the service profile that requires the UUID for the associated server to be changed.
- Step 5** In the **Work** pane, click the **General** tab.
- Step 6** In the **Actions** area, click **Change UUID**.
- Step 7** From the **UUID Assignment** drop-down list, do one of the following:

Option	Description
Select (pool default used by default)	Assigns a UUID from the default UUID Suffix pool. Continue with Step 9.
Hardware Default	Uses the UUID assigned to the server by the manufacturer. If you choose this option, the UUID remains unassigned until the service profile is associated with a server. At that point, the UUID is set to the UUID value assigned to the server by the manufacturer. If the service profile is later moved to a different server, the UUID is changed to match the new server. Continue with Step 9.
XXXXXXXX-XXXX-XXXX-XXXX-XXXXXXXXXXXX	Uses the UUID that you manually assign. Continue with Step 8.
Pools <i>Pool_Name</i>	Assigns a UUID from the UUID Suffix pool that you select from the list at the bottom of the drop-down list. Each pool name is followed by two numbers in parentheses that show the number of UUIDs still available in the pool and the total number of UUIDs in the pool. Continue with Step 9.

- Step 8** (Optional) If you selected the **XXXXXXXX-XXXX-XXXX-XXXX-XXXXXXXXXXXX** option, do the following:

- a) In the **UUID** field, enter the valid UUID that you want to assign to the server which uses this service profile.
- b) To verify that the selected UUID is available, click the **here** link.

Step 9 Click **OK**.

Modifying the Boot Order in a Service Profile

Procedure

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

Step 2 Expand **Servers > Service Profiles**.

Step 3 Expand the node for the organization that includes the service profile for which you want to change the boot order.

If the system does not include multi-tenancy, expand the **root** node.

Step 4 Click the service profile for which you want to change the boot order.

Step 5 In the **Work** pane, click the **Boot Order** tab.

Step 6 Click **Modify Boot Policy** to change the existing boot policy.

Step 7 In the **Modify Boot Policy** dialog box, choose one of the following from the **Boot Policy** drop-down list:

Option	Description
Select Boot Policy to use	Assigns the default boot policy to this service profile. Continue with Step 14.
Create a Specific Boot Policy	Enables you to create a local boot policy that can only be accessed by this service profile. Continue with Step 8.
Boot Policies <i>Policy_Name</i>	Assigns an existing boot policy to the service profile. If you choose this option, Cisco UCS Manager displays the details of the policy. If you do not want use any of the existing policies, but instead want to create a policy that all service profiles can access, click Create Boot Policy and continue with Step 2. Otherwise, continue with Step 14.

Step 8 If you chose to create a boot policy, in the **Create Boot Policy** dialog box, enter a unique name and description for the policy.

This name can be between 1 and 16 alphanumeric characters. You cannot use spaces or any special characters other than - (hyphen), _ (underscore), : (colon), and . (period), and you cannot change this name after the object is saved.

Step 9 (Optional) To reboot all servers that use this boot policy after you make changes to the boot order, check the **Reboot on Boot Order Change** check box.

In the Cisco UCS Manager GUI, if the **Reboot on Boot Order Change** check box is checked for a boot policy, and if CD-ROM or Floppy is the last device in the boot order, deleting or adding the device does not directly affect the boot order and the server does not reboot.

Step 10 (Optional) If desired, check the **Enforce vNIC/vHBA/iSCSI Name** check box.

- If checked, Cisco UCS Manager displays a configuration error and reports whether one or more of the vNICs, vHBAs, or iSCSI vNICs listed in the **Boot Order** table match the server configuration in the service profile.
- If not checked, Cisco UCS Manager uses the vNICs or vHBAs (as appropriate for the boot option) from the service profile.

Step 11 To add a local disk, virtual CD-ROM, or virtual floppy to the boot order, do the following:

- a) Click the down arrows to expand the **Local Devices** area.
- b) Click one of the following links to add the device to the **Boot Order** table:

- **Add Local Disk** or
 - **Add Local LUN**
 - **Add Local JBOD**
 - **Add SD Card**
 - **Add Internal USB**
 - **Add External USB**
 - **Add Embedded Local LUN**
 - **Add Embedded Local Disk**
- **Add CD/DVD** or
 - **Add Local CD/DVD**
 - **Add Local Remote CD/DVD**

In a setup with M5 blade servers, if an ISO is mapped to the KVM console, use only **Add Remote CD/DVD** in **Boot Order**.

- c) Add another boot device to the **Boot Order** table, or click **OK** to finish.

Step 12 To add a LAN boot to the boot order, do the following:

- a) Click the down arrows to expand the **vNICs** area.
- b) Click the **Add LAN Boot** link.
- c) In the **Add LAN Boot** dialog box, enter the **vNIC** name, in the vNIC field for LAN Boot, select the **IP address Type**—None, IPv4, or IPv6, and then click **OK**.
- d) Add another device to the **Boot Order** table, or click **OK** to finish.

Step 13 To add a SAN boot to the boot order, do the following:

- a) Click the down arrows to expand the **vHBAs** area.
- b) Click the **Add SAN Boot** link.
- c) In the **Add San Boot** dialog box, specify the vHBA and type, then click **OK**.

- d) If this vHBA points to a bootable SAN image, click the **Add SAN Boot Target** link and, in the **Add SAN Boot Target** dialog box, specify the boot target LUN, boot target WWPN, and type, then click **OK**.
- e) Add another boot device to the **Boot Order** table, or click **OK** to finish.

Step 14 Click **OK**.

Creating a vNIC for a Service Profile

Procedure

- Step 1** In the **Navigation** pane, click **Servers**.
 - Step 2** Expand **Servers > Service Profiles**.
 - Step 3** Expand the node for the organization that contains the service profile for which you want to create a vNIC.
 - Step 4** Expand the service profile for which you want to create a vNIC.
 - Step 5** Right-click the **vNICs** node and choose **Create vNICs**.
 - Step 6** In the **Create vNIC** dialog box, enter the name, select a **MAC Address Assignment**, and check the **Use vNIC Template** check box if you want to use an existing vNIC template.
You can also create a MAC pool from this area.
 - Step 7** Choose the **Fabric ID**, select the **VLANs** that you want to use, enter the **CDN Name** and **MTU**, and choose a **Pin Group**.
You can also create a VLAN and a LAN pin group from this area.
 - Step 8** In the **Operational Parameters** area, choose a **Stats Threshold Policy**.
 - Step 9** In the Adapter Performance Profile area, choose an **Adapter Policy**, **QoS Policy**, and a **Network Control Policy**.
You can also create an Ethernet adapter policy, QoS policy, and network control policy from this area.
 - Step 10** In the Connection Policies area, choose the **Dynamic vNIC**, **usNIC** or **VMQ** radio button, then choose the corresponding policy.
You can also create a dynamic vNIC, usNIC, or VMQ connection policy from this area.
 - Step 11** Click **OK**.
-

Deleting a vNIC from a Service Profile

Procedure

- Step 1** In the **Navigation** pane, click **Servers**.
- Step 2** Expand **Servers > Service Profiles**.

- Step 3** Expand the node for the organization that contains the service profile from which you want to delete a vNIC.
 - Step 4** Expand the service profile from which you want to delete a vNIC.
 - Step 5** Expand the vNICs node.
 - Step 6** Right-click the vNIC you want to delete and choose **Delete**.
 - Step 7** If a confirmation dialog box displays, click **Yes**.
-

Creating a vHBA for a Service Profile

Procedure

- Step 1** In the **Navigation** pane, click **Servers**.
 - Step 2** Expand **Servers > Service Profiles**.
 - Step 3** Expand the node for the organization that contains the service profile for which you want to create a vHBA.
 - Step 4** Expand the service profile for which you want to create a vHBA.
 - Step 5** Right-click the **vHBAs** node and choose **Create vHBAs**.
 - Step 6** In the **Create vHBAs** dialog box, enter the name and optional description.
 - Step 7** Choose the **Fabric ID**, **Select VSAN**, **Pin Group**, **Persistent Binding**, and **Max Data Field Size**.
You can also create a VSAN or SAN pin group from this area.
 - Step 8** In the **Operational Parameters** area, choose the **Stats Threshold Policy**.
 - Step 9** In the **Adapter Performance Profile** area, choose the **Adapter Policy** and **QoS Policy**.
You can also create a fibre channel adapter policy or QoS policy from this area.
 - Step 10** Click **OK**.
-

Changing the WWPN for a vHBA

Procedure

- Step 1** In the **Navigation** pane, click **Servers**.
- Step 2** Expand **Servers > Service Profiles**.
- Step 3** Expand the node for the organization that contains the service profile for which you want to change the WWPN.
- Step 4** Expand *Service_Profile_Name* > **vHBAs**.
- Step 5** Click the vHBA for which you want to change the WWPN.
- Step 6** In the **Work** pane, click the **General** tab.
- Step 7** In the **Actions** area, click **Change World Wide Name**.

- Step 8** In the **Change World Wide Port Name** dialog box, complete the required fields.
- Step 9** Click **OK**.
-

Clearing Persistent Binding for a vHBA

Procedure

- Step 1** In the **Navigation** pane, click **Servers**.
- Step 2** Expand **Servers > Service Profiles**.
- Step 3** Expand the node for the organization that contains the service profile for which you want to modify the vHBA.
- Step 4** Expand *Service_Profile_Name* > **vHBAs**.
- Step 5** Click the vHBA for which you want to clear the persistent binding.
- Step 6** In the **Work** pane, click the **General** tab.
- Step 7** In the **Actions** area, click **Clear Persistent Binding**.
- Step 8** If a confirmation dialog box displays, click **Yes**.
-

Deleting a vHBA from a Service Profile

Procedure

- Step 1** In the **Navigation** pane, click **Servers**.
- Step 2** Expand **Servers > Service Profiles**.
- Step 3** Expand the node for the organization that contains the service profile from which you want to delete a vHBA.
- Step 4** Expand the service profile from which you want to delete a vHBA.
- Step 5** Expand the **vHBAs** node.
- Step 6** Right-click the vHBA you want to delete and choose **Delete**.
- Step 7** If a confirmation dialog box displays, click **Yes**.
-

Adding a vHBA Initiator Group to a Service Profile

Procedure

- Step 1** Expand **Servers > Service Profiles**.

Step 2 Expand the node for the organization that contains the service profile to which you want to add a vHBA initiator group.

If the system does not include multi tenancy, expand the **root** node.

Step 3 Choose the service profile to which you want to add a vHBA initiator group.

Step 4 In the **Work** pane, click the **Storage > vHBA Initiator Groups**.

Step 5 On the icon bar at the right of the **Select vHBA Initiator Groups** table, click +.

Step 6 In the **Create vHBA Initiator Group** dialog box, complete the following fields to set the name and description:

Name	Description
Name field	The name of the vHBA initiator group. This name can be between 1 and 16 alphanumeric characters. You cannot use spaces or any special characters other than - (hyphen), _ (underscore), : (colon), and . (period), and you cannot change this name after the object is saved.
Description field	A description of the group. Enter up to 256 characters. You can use any characters or spaces except ` (accent mark), \ (backslash), ^ (carat), " (double quote), = (equal sign), > (greater than), < (less than), or ' (single quote).

Step 7 In the **Select vHBA Initiators** table, check the check box in the **Select** column for each vHBA you want to include in the vHBA initiator group.

Step 8 To add a storage connection policy to the initiator group, choose one of the following options:

- Choose an existing storage connection policy from the **Storage Connection Policy** drop-down list. Continue with Step 10.
- Click the **Create Storage Connection Policy** link if you want to create a new storage connection policy that will be available for use by other vHBA initiator groups within the Cisco UCS domain. For more information, see [Creating a Fibre Channel Storage Connection Policy](#). After you create the storage connection policy, continue with Step 10.
- Choose the **Specific Storage Connection Policy** option to create a storage connection policy that is only available to this vHBA initiator group. Continue with Step 9.

Step 9 In the **Specific Storage Connection Policy** area, complete the following fields to create a storage connection policy that is only available to this vHBA initiator group:

Name	Description
Description field	A description of the policy. Cisco recommends including information about where and when to use the policy. Enter up to 256 characters. You can use any characters or spaces except ` (accent mark), \ (backslash), ^ (carat), " (double quote), = (equal sign), > (greater than), < (less than), or ' (single quote).

Name	Description
Zoning Type field	<p>This can be one of the following:</p> <ul style="list-style-type: none"> • None—Cisco UCS Manager does not configure Fibre Channel zoning. • Single Initiator Single Target—Cisco UCS Manager automatically creates one zone for each vHBA and storage port pair. Each zone has two members. We recommend that you configure this type of zoning unless you expect the number of zones to exceed the maximum supported. • Single Initiator Multiple Targets—Cisco UCS Manager automatically creates one zone for each vHBA. We recommend that you configure this type of zoning if you expect the number of zones to reach or exceed the maximum supported.
FC Target Endpoints table	<p>The Fibre Channel target endpoints associated with this policy. This table contains the following columns and buttons:</p> <ul style="list-style-type: none"> • WWPN column—The World Wide Port Name associated with the endpoint. • Path column—The path to the endpoint. • VSAN column—The VSAN associated with the endpoint. • Add button—Creates a new FC target endpoint. • Delete button—Deletes the selected endpoint. • Properties button—Displays all properties for the selected endpoint.

Step 10 Click **OK**.

Step 11 If a confirmation dialog box displays, click **Yes**.

Deleting a Service Profile

Procedure

-
- Step 1** In the **Navigation** pane, click **Servers**.
- Step 2** Expand **Servers > Service Profiles > Organization Name**.
- Step 3** Right-click the service profile you want to delete and select **Delete**.
- Step 4** If a confirmation dialog box displays, click **Yes**.
- Step 5** Click **OK**.
-

Service Profile Association

Associating a Service Profile with a Server or Server Pool

Follow this procedure if you did not associate the service profile with a blade server or server pool when you created it, or to change the blade server or server pool with which a service profile is associated.

Procedure

- Step 1** In the **Navigation** pane, click **Servers**.
- Step 2** Expand **Servers** > **Service Profiles**.
- Step 3** Expand the node for the organization that contains the service profile that you want to associate with a new server or server pool.
- If the system does not include multi tenancy, expand the **root** node.
- Step 4** Right-click the service profile you want to associate with a server and select **Associate Service Profile**.
- Step 5** In the **Associate Service Profile** dialog box, select one of the following options:

Option	Description
Server Pool	Select a server pool from the drop-down list. Cisco UCS Manager assigns a server from this pool to the service profile. Continue with Step 7.
Server	Navigate to the desired available server in the navigation tree and select the server which will be assigned to the service profile. Continue with Step 7.
Custom Server	Specifies the chassis and slot that contains the server that will be assigned to the service profile. If the server is not in the slot or is otherwise unavailable, the service profile will be associated with the server when it becomes available. Continue with Step 6.

- Step 6** If you chose **Custom Server**, do the following:
- In the **Chassis Id** field, enter the number of the chassis where the selected server is located.
 - In the **Server Id** field, enter the number of the slot where the selected server is located.
- Step 7** If you want to restrict the migration of the service profile after it is associated with a server, check the **Restrict Migration** check box.

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

Step 8 Click **OK**.

Disassociating a Service Profile from a Server or Server Pool

When you disassociate a service profile, Cisco UCS Manager attempts to shutdown the operating system on the server. If the operating system does not shutdown within a reasonable length of time, Cisco UCS Manager forces the server to shutdown.

Procedure

- Step 1** In the **Navigation** pane, click **Servers**.
- Step 2** Expand **Servers > Service Profiles**.
- Step 3** Expand the node for the organization that contains the service profile that you want to disassociate from a server or server pool.
- If the system does not include multi tenancy, expand the **root** node.
- Step 4** Right-click the service profile you want to disassociate from a server and select **Disassociate Service Profile**.
- Step 5** In the **Disassociate Service Profile** dialog box, click **Yes** to confirm that you want to disassociate the service profile.
- Step 6** (Optional) Monitor the status and FSM for the server to confirm that the disassociation completed.
-

Service Profile Templates

Initial and Existing Templates

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



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

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

Cisco UCS supports the following types of service profile templates:

Initial template

Service profiles created from an initial template inherit all the properties of the template. Service profiles created from an initial service profile template are bound to the template. However, changes to the initial

template do not *automatically* propagate to the bound service profiles. If you want to propagate changes to bound service profiles, unbind and rebind the service profile to the initial template.

Updating template

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



Note Service profiles that are created from the initial template and normal service profiles fetch the lowest available IDs in the sequential pool when you press **Reset**.

Service profiles created from updating template might attempt to retain the same ID when you press **Reset** even when lower IDs of sequential pool are free.

Creating a Service Profile Template

Procedure

-
- Step 1** In the **Navigation** pane, click **Servers**.
- Step 2** Expand **Servers** > **Service Profile Templates**.
- Step 3** Expand the node for the organization where you want to create the service profile template.
If the system does not include multi tenancy, expand the **root** node.
- Step 4** Right-click the organization and choose **Create Service Profile Template**.
- Step 5** In the **Identify Service Profile Template** panel, specify the service profile **Name**, **Type**, and **UUID Assignment**, then click **Next**.
You can provide an optional description for this service profile template.
- Note**
To create a service profile template quickly, you can click **Finish** after specifying the name. Cisco UCS Manager creates a new service profile template with the specified name and all system default values.
- Step 6** (Optional) In the **Networking** panel, specify the required information for the **Dynamic vNIC Connection Policy** and **LAN Connectivity** sections, then click **Next**.
You can create a dynamic vNIC connection policy and LAN connectivity policy from this panel.
- Step 7** (Optional) In the **Storage** panel, specify the SAN configuration information such as, **Local Storage Policy**, **SAN Connectivity**, **WWNN**, and **vHBAs**, then click **Next**.
You can create a local disk configuration policy and SAN connectivity policy from this panel.
- Step 8** (Optional) In the **Zoning** panel, specify the required zoning information, then click **Next**.
You can create the vHBA initiator groups from this panel.
- Step 9** (Optional) In the **vNIC/vHBA Placement** panel, specify the placement method and PCI order, then click **Next**.

You can create a placement policy from this panel.

Step 10 (Optional) In the **Server Boot Order** panel, specify the **Boot Policy** from the drop-down list, then click **Next**.
You can create a boot policy from this panel.

Step 11 (Optional) In the **Maintenance Policy** panel, specify the maintenance policy, then click **Next**.
You can create a new maintenance policy and specify a maintenance schedule from this panel.

Step 12 (Optional) In the **Server Assignment** panel, specify the **Pool Assignment** from the drop down list and the power state to apply on assignment, then click **Next**.
You can create a server pool or a host firmware package from this panel.

Step 13 (Optional) In the **Operational Policies** panel, specify the system operational information such as, **BIOS Configuration**, **External IPMI Management Configuration**, **Management IP Address**, **Monitoring Configuration (Thresholds)**, **Power Control Policy Configuration**, and **Scrub Policy**, then click **Finish**.

Note

To set up an Outband IPv4 address or an Inband IPv4 or IPv6 address, click the respective tabs and complete the required fields.

If you do not find the policies you need for each of these configurations, you can create them from this panel.

Creating One or More Service Profiles from a Service Profile Template

Procedure

-
- Step 1** In the **Navigation** pane, click **Servers**.
 - Step 2** Expand **Servers > Service Profile Templates**.
 - Step 3** Expand the node for the organization that contains the service profile template that you want to use as the basis for your service profiles.
If the system does not include multi tenancy, expand the **root** node.
 - Step 4** Right-click the service profile template from which you want to create the profiles and select **Create Service Profiles From Template**.
 - Step 5** In the **Create Service Profiles From Template** dialog box, complete the required fields.
 - Step 6** Click **OK**.
-

Creating a Template Based Service Profile for a Blade Server

Before you begin

A qualified service profile template with the desired values must exist in Cisco UCS Manager.

Procedure

- Step 1** In the **Navigation** pane, click **Equipment**.
- Step 2** Expand **Equipment > Chassis > *Chassis Number* > Servers**.
- Step 3** Choose the server for which you want to create a template based service profile.
- Step 4** In the **Work** pane, click the **General** tab.
- Step 5** In the **Actions** area, click **Create Service Profile**.
- Step 6** In the **Create Service Profile for Server** dialog box, do the following:
- Click the **Template Based Service Profile** radio button.
 - In the **Name** field, enter a unique name for the service profile.

This name can be between 2 and 32 alphanumeric characters. You cannot use spaces or any special characters other than - (hyphen), _ (underscore), : (colon), and . (period), and this name must be unique across all service profiles and service profile templates within the same organization.
 - From the **Service Profile Template** drop-down list, select the template from which you want to create the service profile associated with this server.

Note
The drop-down list only lists service profile templates compatible with the selected blade server.
 - Click **OK**.
-

Creating a Template Based Service Profile for a Rack-Mount Server

Before you begin

A qualified service profile template with the desired values must exist in Cisco UCS Manager.

Procedure

- Step 1** In the **Navigation** pane, click **Equipment**.
- Step 2** Expand **Equipment > Rack Mounts > Servers**.
- Note**
For Cisco UCS C125 M5 Servers, expand **Equipment > Rack Mounts > Enclosures > Rack Enclosure *rack_enclosure_number* > Servers**.
- Step 3** Choose the server for which you want to create a template based service profile.
- Step 4** In the **Work** pane, click the **General** tab.
- Step 5** In the **Actions** area, click **Create Service Profile**.
- Step 6** In the **Create Service Profile for Server** dialog box, do the following:
- Click the **Template Based Service Profile** radio button.

- b) In the **Name** field, enter a unique name for the service profile.
This name can be between 2 and 32 alphanumeric characters. You cannot use spaces or any special characters other than - (hyphen), _ (underscore), : (colon), and . (period), and this name must be unique across all service profiles and service profile templates within the same organization.
 - c) From the **Service Profile Template** drop-down list, select the template from which you want to create the service profile associated with this server.
 - d) Click **OK**.
-

Creating a Service Profile Template from a Service Profile

Procedure

- Step 1** In the **Navigation** pane, click **Servers**.
 - Step 2** Expand **Servers > Service Profiles**.
 - Step 3** Expand the node for the organization that contains the service profile that you want to use as the basis for your template.
If the system does not include multi tenancy, expand the **root** node.
 - Step 4** Right-click the service profile from which you want to create the template and select **Create a Service Profile Template**.
 - Step 5** In the **Create Template From Service Profile** dialog box, complete the required fields.
 - Step 6** Click **OK**.
-

Setting an Asset Tag for a Service Profile

Procedure

- Step 1** Navigate to **Servers > Service Profiles**
- Step 2** Expand the node for which you want to create the asset tag.
If the system does not include multitenancy expand the **root** node
- Step 3** In the Work pane, click the **General** tab.
- Step 4** In the **Asset Tag** field, enter a name to identify the server.
The name can be between 2 to 32 alphanumeric characters. You can use space or any special characters other than ` (accent mark), \ (backslash), ^ (carat), " (double quote), = (equal sign), > (greater than), < (less than), or ' (single quote).
- Step 5** Click **Save Changes**.

- Step 6** Reboot the server manually for the changes to be in effect.
-

Service Profile Template Tasks

Binding a Service Profile to a Service Profile Template

You can bind a service profile to a service profile template. When you bind the service profile to a template, Cisco UCS Manager configures the service profile with the values defined in the service profile template. If the existing service profile configuration does not match the template, Cisco UCS Manager reconfigures the service profile. You can only change the configuration of a bound service profile through the associated template.

Procedure

- Step 1** In the **Navigation** pane, click **Servers**.
- Step 2** Expand **Servers > Service Profiles**.
- Step 3** Expand the node for the organization that includes the service profile you want to bind.
If the system does not include multi-tenancy, expand the **root** node.
- Step 4** Click the service profile you want to bind.
- Step 5** In the **Work** pane, click the **General** tab.
- Step 6** In the **Actions** area, click **Bind to a Template**.
- Step 7** In the **Bind to a Service Profile Template** dialog box, do the following:
- From the **Service Profile Template** drop-down list, choose the template to which you want to bind the service profile.
 - Click **OK**.
-

Unbinding a Service Profile from a Service Profile Template

Procedure

- Step 1** In the **Navigation** pane, click **Servers**.
- Step 2** Expand **Servers > Service Profiles**.
- Step 3** Expand the node for the organization that includes the service profile you want to unbind.
If the system does not include multi-tenancy, expand the **root** node.
- Step 4** Click the service profile you want to unbind.
- Step 5** In the **Work** pane, click the **General** tab.

Step 6 In the **Actions** area, click **Unbind from the Template**.

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

Changing the UUID in a Service Profile Template

Procedure

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

Step 2 Expand **Servers > Service Profile Templates**.

Step 3 Expand the node for the organization that contains the service profile template for which you want to change the UUID.

If the system does not include multi tenancy, expand the **root** node.

Step 4 Choose the service profile template whose UUID assignment you want to change.

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

Step 6 In the **Actions** area, click **Change UUID**.

Step 7 From the **UUID Assignment** drop-down list, choose one of the following:

Option	Description
Select (pool default used by default)	Assigns a UUID from the default UUID Suffix pool.
Hardware Default	Uses the UUID assigned to the server by the manufacturer. If you choose this option, the UUID remains unassigned until the service profile is associated with a server. At that point, the UUID is set to the UUID value assigned to the server by the manufacturer. If the service profile is later moved to a different server, the UUID is changed to match the new server.
Pools <i>Pool_Name</i>	Assigns a UUID from the UUID Suffix pool that you select from the list at the bottom of the drop-down list. Each pool name is followed by two numbers in parentheses that show the number of UUIDs still available in the pool and the total number of UUIDs in the pool.

Step 8 Click **OK**.

Resetting the UUID Assigned to a Service Profile from a Pool in a Service Profile Template

If you change the UUID suffix pool assigned to an updating service profile template, Cisco UCS Manager does not change the UUID assigned to a service profile created with that template. If you want Cisco UCS Manager to assign a UUID from the newly assigned pool to the service profile, and therefore to the associated

server, you must reset the UUID. You can only reset the UUID assigned to a service profile and its associated server under the following circumstances:

- The service profile was created from an updating service profile template and includes a UUID assigned from a UUID suffix pool.
- The UUID suffix pool name is specified in the service profile. For example, the pool name is not empty.
- The UUID value is not 0, and is therefore not derived from the server hardware.

Procedure

- Step 1** In the **Navigation** pane, click **Servers**.
- Step 2** Expand **Servers > Service Profiles**.
- Step 3** Expand the node for the organization that contains the service profile for which you want to reset the UUID. If the system does not include multi tenancy, expand the **root** node.
- Step 4** Choose the service profile that requires the UUID for the associated server to be reset to a different UUID suffix pool.
- Step 5** In the **Work** pane, click the **General** tab.
- Step 6** In the **Actions** area, click **Reset UUID**.
- If this action is not visible, then the UUID configuration in the service profile does not meet the requirements for resetting a UUID.
- Step 7** If a confirmation dialog box displays, click **Yes**.
- Step 8** Click **OK**
-

Resetting the MAC Address Assigned to a vNIC from a Pool in a Service Profile

If you change the MAC pool assigned to an updating service profile template or to a service profile using LAN connectivity policy/vNIC template, Cisco UCS Manager does not change the MAC address assigned to a service profile. If you want Cisco UCS Manager to assign a MAC address from the newly assigned pool to the service profile, and therefore to the associated server, you must reset the MAC address. This applies to all the Service Profiles, whether they are created from a template or configured manually. You can only reset the MAC address assigned to a service profile and its associated server under the following circumstances:

- Service Profile created from an updating Service Profile template and includes a MAC address assigned from a MAC pool.
- Service Profile created using LAN connectivity policy/vNic template and includes a MAC address assigned from a MAC pool.
- The MAC pool name is specified in the service profile. For example, the pool name is not empty.
- The MAC address value is not 0, and is therefore not derived from the server hardware.

Procedure

- Step 1** In the **Navigation** pane, click **Servers**.
- Step 2** Expand **Servers > Service Profiles**.
- Step 3** Expand the node for the organization that contains the service profile for which you want to reset the MAC address.
- If the system does not include multi tenancy, expand the **root** node.
- Step 4** Expand *Service_Profile_Name* > **vNICs**.
- Step 5** Click the vNIC for which you want to reset the MAC address.
- Step 6** In the **Work** pane, click the **General** tab.
- Step 7** In the **Actions** area, click **Reset MAC Address**.
- Step 8** If a confirmation dialog box displays, click **Yes**.
- Step 9** Click **OK**.
-

Resetting the WWPN Assigned to a vHBA from a Pool in a Service Profile Template

If you change the WWPN pool assigned to an updating service profile template, Cisco UCS Manager does not change the WWPN assigned to a service profile created with that template. If you want Cisco UCS Manager to assign a WWPN from the newly assigned pool to the service profile, and therefore to the associated server, you must reset the WWPN. You can only reset the WWPN assigned to a service profile and its associated server under the following circumstances:

- The service profile was created from an updating service profile template and includes a WWPN assigned from a WWPN pool.
- The WWPN pool name is specified in the service profile. For example, the pool name is not empty.
- The WWPN value is not 0, and is therefore not derived from the server hardware.

Procedure

- Step 1** In the **Navigation** pane, click **Servers**.
- Step 2** Expand **Servers > Service Profiles**.
- Step 3** Expand the node for the organization that contains the service profile for which you want to reset the WWPN.
- If the system does not include multi tenancy, expand the **root** node.
- Step 4** Expand *Service_Profile_Name* > **vHBAs**.
- Step 5** Click the vHBA for which you want to reset the WWPN.
- Step 6** In the **Work** pane, click the **General** tab.

- Step 7** In the **Actions** area, click **Reset WWPN**.
 - Step 8** If a confirmation dialog box displays, click **Yes**.
 - Step 9** Click **OK**.
-

Deleting the Inband Configuration from a Service Profile Template

This procedure removes the inband management IP address configuration from a service profile template. If this action is greyed out, no inband configuration was configured.

Procedure

- Step 1** In the **Navigation** pane, click **Servers**.
 - Step 2** Expand **Servers > Service Profile Template > Service_Profile_Template_Name**.
 - Step 3** In the **Work** pane, click the **General** tab.
 - Step 4** In the **Actions** area, click **Delete Inband Configuration**.
 - Step 5** Click **Yes** in the **Delete** confirmation dialog box.
- The inband management IP address configuration for the service profile template is deleted.
-

Service Profile Association

Associating a Service Profile with a Server or Server Pool

Follow this procedure if you did not associate the service profile with a blade server or server pool when you created it, or to change the blade server or server pool with which a service profile is associated.

Procedure

- Step 1** In the **Navigation** pane, click **Servers**.
- Step 2** Expand **Servers > Service Profiles**.
- Step 3** Expand the node for the organization that contains the service profile that you want to associate with a new server or server pool.

If the system does not include multi tenancy, expand the **root** node.
- Step 4** Right-click the service profile you want to associate with a server and select **Associate Service Profile**.
- Step 5** In the **Associate Service Profile** dialog box, select one of the following options:

Option	Description
Server Pool	Select a server pool from the drop-down list. Cisco UCS Manager assigns a server from this pool to the service profile. Continue with Step 7.
Server	Navigate to the desired available server in the navigation tree and select the server which will be assigned to the service profile. Continue with Step 7.
Custom Server	Specifies the chassis and slot that contains the server that will be assigned to the service profile. If the server is not in the slot or is otherwise unavailable, the service profile will be associated with the server when it becomes available. Continue with Step 6.

Step 6 If you chose **Custom Server**, do the following:

- a) In the **Chassis Id** field, enter the number of the chassis where the selected server is located.
- b) In the **Server Id** field, enter the number of the slot where the selected server is located.

Step 7 If you want to restrict the migration of the service profile after it is associated with a server, check the **Restrict Migration** check box.

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

Step 8 Click **OK**.

Associating a Service Profile Template with a Server Pool

Follow this procedure if you did not associate the service profile template with a server pool when you created it, or to change the server pool with which a service profile created from this template is associated.

Procedure

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

Step 2 Expand **Servers > Service Profile Templates**.

Step 3 Expand the node for the organization that contains the service profile that you want to associate with a server pool.

If the system does not include multi tenancy, expand the **root** node.

Step 4 Right-click the service profile template you want to associate with a server pool and select **Associate with Server Pool**.

The **Associate with Server Pool** dialog box opens.

Step 5 From the **Server Pool** section of the **Pool Assignment** drop-down list, select a server pool.

If you select **Assign Later**, the service profile template is not associated with a server pool.

- Step 6** (Optional) From the **Select Qualification** drop-down list, select the server pool policy qualifications you want to apply to a server that is associated with a service profile created from this template.
- Step 7** Click **OK**.
-

Disassociating a Service Profile from a Server or Server Pool

When you disassociate a service profile, Cisco UCS Manager attempts to shutdown the operating system on the server. If the operating system does not shutdown within a reasonable length of time, Cisco UCS Manager forces the server to shutdown.

Procedure

- Step 1** In the **Navigation** pane, click **Servers**.
- Step 2** Expand **Servers > Service Profiles**.
- Step 3** Expand the node for the organization that contains the service profile that you want to disassociate from a server or server pool.
- If the system does not include multi tenancy, expand the **root** node.
- Step 4** Right-click the service profile you want to disassociate from a server and select **Disassociate Service Profile**.
- Step 5** In the **Disassociate Service Profile** dialog box, click **Yes** to confirm that you want to disassociate the service profile.
- Step 6** (Optional) Monitor the status and FSM for the server to confirm that the disassociation completed.
-

Disassociating a Service Profile Template from its Server Pool

Procedure

- Step 1** In the **Navigation** pane, click **Servers**.
- Step 2** Expand **Servers > Service Profile Templates**.
- Step 3** Expand the node for the organization that contains the service profile that you want to disassociate from its server pool.
- If the system does not include multi tenancy, expand the **root** node.
- Step 4** Right-click the service profile template you want to disassociate from its server pool and select **Disassociate Template**.
- Step 5** If a confirmation dialog box displays, click **Yes**.
-



CHAPTER 12

Server-Related Policies

- [BIOS Settings, on page 187](#)
- [Trusted Platform Module, on page 269](#)
- [Consistent Device Naming, on page 272](#)
- [CIMC Security Policies, on page 275](#)
- [Graphics Card Policies, on page 278](#)
- [Local Disk Policies, on page 279](#)
- [Persistent Memory Modules, on page 291](#)
- [Scrub Policy, on page 292](#)
- [DIMM Error Management, on page 295](#)
- [Serial over LAN Policy Settings, on page 297](#)
- [Server Autoconfiguration Policies, on page 299](#)
- [Server Discovery Policy Settings, on page 301](#)
- [Server Inheritance Policy Settings, on page 303](#)
- [Server Pool Policy Settings, on page 305](#)
- [Server Pool Policy Qualifications Settings, on page 306](#)
- [vNIC/vHBA Placement Policy Settings, on page 312](#)
- [CIMC Mounted vMedia, on page 325](#)

BIOS Settings

Server BIOS Settings

Server BIOS Settings

Cisco UCS provides two methods for making global modifications to the BIOS settings on servers in an Cisco UCS domain. You can create one or more BIOS policies that include a specific grouping of BIOS settings that match the needs of a server or set of servers, or you can use the default BIOS settings for a specific server platform.

Both the BIOS policy and the default BIOS settings for a server platform enable you to fine tune the BIOS settings for a server managed by Cisco UCS Manager.

Depending upon the needs of the data center, you can configure BIOS policies for some service profiles and use the BIOS defaults in other service profiles in the same Cisco UCS domain, or you can use only one of

them. You can also use Cisco UCS Manager to view the actual BIOS settings on a server and determine whether they are meeting current needs.



Note Cisco UCS Manager pushes BIOS configuration changes through a BIOS policy or default BIOS settings to the Cisco Integrated Management Controller (CIMC) buffer. These changes remain in the buffer and do not take effect until the server is rebooted.

We recommend that you verify the support for BIOS settings in the server that you want to configure. Some settings, such as Mirroring Mode for RAS Memory, are not supported by all Cisco UCS servers.

Main BIOS Settings

The following table lists the main server BIOS settings that you can configure through a BIOS policy or the default BIOS settings:

Name	Description
Properties	
Reboot on BIOS Settings Change	<p>When the server is rebooted after you change one or more BIOS settings.</p> <p>If you enable this setting, the server is rebooted according to the maintenance policy in the server's service profile. For example, if the maintenance policy requires user acknowledgment, the server is not rebooted and the BIOS changes are not applied until a user acknowledges the pending activity.</p> <p>If you do not enable this setting, the BIOS changes are not applied until the next time the server is rebooted, whether as a result of another server configuration change or a manual reboot.</p>
BIOS Setting	
Quiet Boot	<p>What the BIOS displays during Power On Self-Test (POST). This can be one of the following:</p> <ul style="list-style-type: none"> • Disabled—The BIOS displays all messages and Option ROM information during boot. • Enabled—The BIOS displays the logo screen, but does not display any messages or Option ROM information during boot. • Platform Default—The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.

Name	Description
POST error pause	<p>What happens when the server encounters a critical error during POST. This can be one of the following:</p> <ul style="list-style-type: none"> • Disabled—The BIOS continues to attempt to boot the server. • Enabled—The BIOS pauses the attempt to boot the server and opens the Error Manager when a critical error occurs during POST. • Platform Default—The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.
Resume on AC power loss	<p>How the server behaves when power is restored after an unexpected power loss. This can be one of the following:</p> <ul style="list-style-type: none"> • Stay Off—The server remains off until manually powered on. • Last State—The server is powered on and the system attempts to restore its last state. • Reset—The server is powered on and automatically reset. • Platform Default—The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.
Front panel lockout	<p>Whether the power and reset buttons on the front panel are ignored by the server. This can be one of the following:</p> <ul style="list-style-type: none"> • Disabled—The power and reset buttons on the front panel are active and can be used to affect the server. • Enabled—The power and reset buttons are locked out. The server can only be reset or powered on or off from the CIMC GUI. • Platform Default—The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.

Name	Description
CDN Control	<p>Consistent Device Naming allows Ethernet interfaces to be named in a consistent manner. This makes Ethernet interface names more uniform, easy to identify, and persistent when adapter or other configuration changes are made.</p> <p>Whether consistent device naming is enabled or not. This can be one of the following:</p> <ul style="list-style-type: none"> • Disabled—Consistent device naming is disabled for the BIOS policy. • Enabled—Consistent device naming is enabled for the BIOS policy. This enables Ethernet interfaces to be named consistently. This is the default option. • Platform Default—The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.

Processor BIOS Settings

The following table lists the processor BIOS settings that you can configure through a BIOS policy or the default BIOS settings:

Name	Description
Intel Turbo Boost Tech	<p>Whether the processor uses Intel Turbo Boost Technology, which allows the processor to automatically increase its frequency if it is running below power, temperature, or voltage specifications. This can be one of the following:</p> <ul style="list-style-type: none"> • Disabled—The processor does not increase its frequency automatically. • Enabled—The processor uses Turbo Boost Technology if required. • Platform Default—The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.

Name	Description
Enhanced Intel SpeedStep Tech	<p>Whether the processor uses Enhanced Intel SpeedStep Technology, which allows the system to dynamically adjust processor voltage and core frequency. This technology can result in decreased average power consumption and decreased average heat production. This can be one of the following:</p> <ul style="list-style-type: none"> • Disabled—The processor never dynamically adjusts its voltage or frequency. • Enabled—The processor utilizes Enhanced Intel SpeedStep Technology and enables all supported processor sleep states to further conserve power. • Platform Default —The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor. <p>We recommend that you contact your operating system vendor to make sure your operating system supports this feature.</p>
Intel HyperThreading Tech	<p>Whether the processor uses Intel Hyper-Threading Technology, which allows multithreaded software applications to execute threads in parallel within each processor. This can be one of the following:</p> <ul style="list-style-type: none"> • Disabled—The processor does not permit hyperthreading. • Enabled—The processor allows for the parallel execution of multiple threads. • Platform Default —The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor. <p>We recommend that you contact your operating system vendor to make sure the operating system supports this feature.</p>

Name	Description
Intel Speed Select	<p>Allows improved CPU performance by using Intel Speed Select technology to tune the CPU to run at one of three operating profiles, based on number of logical processor cores, frequency, and TDP thread setting, to improve performance over the basic Platform Default setting. These profiles correspond to High, Medium, and Low Core settings and can be one of the following:</p> <ul style="list-style-type: none"> • Base—The processor uses Base. • Config 1—The processor uses Config 1. • Config 2—The processor uses Config 2. • Platform Default —The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor. <p>Low core profiles enable the CPU to run at maximum frequency. Refer to Intel Speed Select, on page 261 sample settings.</p>
Core Multi Processing	<p>Sets the state of logical processor cores per CPU in a package. If you disable this setting, Intel Hyper Threading technology is also disabled. This can be one of the following:</p> <ul style="list-style-type: none"> • All—Enables multiprocessing on all logical processor cores. • 1 through n—Specifies the number of logical processor cores per CPU that can run on the server. To disable multiprocessing and have only one logical processor core per CPU running on the server, choose 1. • Platform Default —The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor. <p>We recommend that you contact your operating system vendor to make sure your operating system supports this feature.</p>

Name	Description
Execute Disable Bit	<p>Classifies memory areas on the server to specify where the application code can execute. As a result of this classification, the processor disables code execution if a malicious worm attempts to insert code in the buffer. This setting helps to prevent damage, worm propagation, and certain classes of malicious buffer overflow attacks. This can be one of the following:</p> <ul style="list-style-type: none"> • Disabled—The processor does not classify memory areas. • Enabled—The processor classifies memory areas. • Platform Default —The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor. <p>We recommend that you contact your operating system vendor to make sure your operating system supports this feature.</p>
Intel Virtualization Technology	<p>Whether the processor uses Intel Virtualization Technology, which allows a platform to run multiple operating systems and applications in independent partitions. This can be one of the following:</p> <ul style="list-style-type: none"> • Disabled—The processor does not permit virtualization. • Enabled—The processor allows multiple operating systems in independent partitions. • Platform Default —The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor. <p>Note If you change this option, you must power cycle the server before the setting takes effect.</p>

Name	Description
Hardware Prefetcher	<p>Whether the processor allows the Intel hardware prefetcher to fetch streams of data and instruction from memory into the unified second-level cache when necessary. This can be one of the following:</p> <ul style="list-style-type: none"> • Disabled—The hardware prefetcher is not used. • Enabled—The processor uses the hardware prefetcher when cache issues are detected. • Platform Default —The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor. <p>Note CPU Performance must be set to Custom in order to specify this value. For any value other than Custom, this option is overridden by the setting in the selected CPU performance profile.</p>
Adjacent Cache Line Prefetcher	<p>Whether the processor fetches cache lines in even/odd pairs instead of fetching just the required line. This can be one of the following:</p> <ul style="list-style-type: none"> • Disabled—The processor only fetches the required line. • Enabled—The processor fetches both the required line and its paired line. • Platform Default —The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor. <p>Note CPU Performance must be set to Custom in order to specify this value. For any value other than Custom, this option is overridden by the setting in the selected CPU performance profile.</p>

Name	Description
DCU Streamer Prefetch	<p>Whether the processor uses the DCU IP Prefetch mechanism to analyze historical cache access patterns and preload the most relevant lines in the L1 cache. This can be one of the following:</p> <ul style="list-style-type: none"> • Disabled—The processor does not try to anticipate cache read requirements and only fetches explicitly requested lines. • Enabled—The DCU prefetcher analyzes the cache read pattern and prefetches the next line in the cache if it determines that it may be needed. • Platform Default —The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.
DCU IP Prefetcher	<p>Whether the processor uses the DCU IP Prefetch mechanism to analyze historical cache access patterns and preload the most relevant lines in the L1 cache. This can be one of the following:</p> <ul style="list-style-type: none"> • Disabled—The processor does not preload any cache data. • Enabled—The DCU IP prefetcher preloads the L1 cache with the data it determines to be the most relevant. • Platform Default —The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.
KTI Prefetch	<p>KTI prefetch is a mechanism to get the memory read started early on a DDR bus. This can be one of the following:</p> <ul style="list-style-type: none"> • Disabled—The processor does not preload any cache data. • Enabled—The KTI prefetcher preloads the L1 cache with the data it determines to be the most relevant. • Platform Default —The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.
LLC Prefetch	<p>Whether the processor uses the LLC Prefetch mechanism to fetch the data into the LLC. This can be one of the following:</p> <ul style="list-style-type: none"> • Disabled—The processor does not preload any cache data. • Enabled—The LLC prefetcher preloads the L1 cache with the data it determines to be the most relevant. • Platform Default —The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.

Name	Description
XPT Prefetch	<p>Whether XPT prefetch is used to enable a read request that is sent to the last level cache to issue a copy of that request to the memory controller prefetcher. This can be one of the following:</p> <ul style="list-style-type: none"> • Disabled—The CPU does not use the XPT Prefetch option. • Enabled—The CPU enables the XPT prefetcher option. • Auto—The CPU auto enables the XPT prefetcher option. • Platform Default —The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.
Direct Cache Access	<p>Allows processors to increase I/O performance by placing data from I/O devices directly into the processor cache. This setting helps to reduce cache misses. This can be one of the following:</p> <ul style="list-style-type: none"> • Auto—The CPU determines how to place data from I/O devices into the processor cache. • Disabled—Data from I/O devices is not placed directly into the processor cache. • Enabled—Data from I/O devices is placed directly into the processor cache. • Platform Default —The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.
Processor C State	<p>Whether the system can enter a power savings mode during idle periods. This can be one of the following:</p> <ul style="list-style-type: none"> • Disabled—The system remains in a high-performance state even when idle. • Enabled—The system can reduce power to system components such as the DIMMs and CPUs. • Platform Default —The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor. <p>We recommend that you contact your operating system vendor to make sure your operating system supports this feature.</p>

Name	Description
Processor C1E	<p>Allows the processor to transition to its minimum frequency upon entering C1. This setting does not take effect until after you have rebooted the server. This can be one of the following:</p> <ul style="list-style-type: none"> • Disabled—The CPU continues to run at its maximum frequency in the C1 state. • Enabled—The CPU transitions to its minimum frequency. This option saves the maximum amount of power in the C1 state. • Platform Default —The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.
Processor C3 Report	<p>Whether the processor sends the C3 report to the operating system. This can be one of the following:</p> <ul style="list-style-type: none"> • Enabled—The processor sends the C3 report to the OS. • Disabled—The processor does not send the C3 report. • ACPI C2—The processor sends the C3 report using the advanced configuration and power interface (ACPI) C2 format. • ACPI C3—The processor sends the C3 report using the ACPI C3 format. • Platform Default —The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor. <p>On the Cisco UCS B440 Server, the BIOS Setup menu uses enabled and disabled for these options. If you specify acpi-c2 or acpi-c3, the server sets the BIOS value for that option to enabled.</p>
Processor C6 Report	<p>Whether the processor sends the C6 report to the operating system. This can be one of the following:</p> <ul style="list-style-type: none"> • Disabled—The processor does not send the C6 report. • Enabled—The processor sends the C6 report. • Platform Default —The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.

Name	Description
Processor C7 Report	<p>Whether the processor sends the C7 report to the operating system. This can be one of the following:</p> <ul style="list-style-type: none"> • C7—The processor sends the report using the C7 format. • C7s—The processor sends the report using the C7s format. • Disabled—The processor does not send the C7 report. • Enabled—The processor sends the C7 report. • Platform Default —The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.
Processor CMCI	<p>Enables CMCI generation. This can be one of the following:</p> <ul style="list-style-type: none"> • Disabled—The processor disables CMCI. • Enabled—The processor enables CMCI. • Platform Default —The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.
CPU Performance	<p>Sets the CPU performance profile for the server. This can be one of the following:</p> <ul style="list-style-type: none"> • Custom—All performance profile options can be configured from the BIOS setup on the server. In addition, the Hardware Prefetcher and Adjacent Cache-Line Prefetch options can be configured as well. • Enterprise—For M3 and higher servers, all prefetchers and data reuse are enabled. • High Throughput—Data reuse and the DCU IP prefetcher are enabled, and all other prefetchers are disabled. • HPC—All prefetchers are enabled and data reuse is disabled. This setting is also known as high-performance computing. • Platform Default —The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.

Name	Description
Max Variable MTRR Setting	<p>Allows you to select the number of mean time to repair (MTRR) variables. This can be one of the following:</p> <ul style="list-style-type: none"> • Auto Max—BIOS uses the default value for the processor. • 8—BIOS uses the number specified for the variable MTRR. • Platform Default —The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.
Local X2 APIC	<p>Allows you to set the type of Application Policy Infrastructure Controller (APIC) architecture. This can be one of the following:</p> <ul style="list-style-type: none"> • Disabled—Processor disables Local X2 APIC. • Enabled—Processor enables Local X2 APIC. • XAPIC—Uses the standard xAPIC architecture. • X2APIC—Uses the enhanced x2APIC architecture to support 32 bit addressability of processors. • AUTO—Automatically uses the xAPIC architecture that is detected. • Platform Default —The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.

Name	Description
Power Technology	<p>Enables you to configure the CPU power management settings for the following options:</p> <ul style="list-style-type: none"> • Enhanced Intel Speedstep Technology • Intel Turbo Boost Technology • Processor Power State C6 <p>Power Technology can be one of the following:</p> <ul style="list-style-type: none"> • Disabled—The server does not perform any CPU power management and any settings for the BIOS parameters mentioned above are ignored. • Energy Efficient—The server determines the best settings for the BIOS parameters mentioned above and ignores the individual settings for these parameters. • Performance—The server automatically optimizes the performance for the BIOS parameters mentioned above. • Custom—The server uses the individual settings for the BIOS parameters mentioned above. You must select this option if you want to change any of these BIOS parameters. • Platform Default —The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.

Name	Description
Energy Performance	<p>Allows you to determine whether system performance or energy efficiency is more important on this server. This can be one of the following:</p> <ul style="list-style-type: none"> • Performance — The server provides all server components with full power at all times. This option maintains the highest level of performance and requires the greatest amount of power. • Balanced Performance — The server provides all server components with enough power to keep a balance between performance and power. • Balanced Energy — The server provides all server components with enough power to keep a balance between performance and power. • Energy Efficient — The server provides all server components with less power to keep reduce power consumption. • Platform Default —The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor. <p>Note Power Technology must be set to Custom or the server ignores the setting for this parameter.</p>
Frequency Floor Override	<p>Whether the CPU is allowed to drop below the maximum non-turbo frequency when idle. This can be one of the following:</p> <ul style="list-style-type: none"> • Disabled— The CPU can drop below the maximum non-turbo frequency when idle. This option decreases power consumption but may reduce system performance. • Enabled— The CPU cannot drop below the maximum non-turbo frequency when idle. This option improves system performance but may increase power consumption. • Platform Default —The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.

Name	Description
P STATE Coordination	<p>Allows you to define how BIOS communicates the P-state support model to the operating system. There are 3 models as defined by the Advanced Configuration and Power Interface (ACPI) specification.</p> <ul style="list-style-type: none"> • HW ALL—The processor hardware is responsible for coordinating the P-state among logical processors with dependencies (all logical processors in a package). • SW ALL—The OS Power Manager (OSPM) is responsible for coordinating the P-state among logical processors with dependencies (all logical processors in a physical package), and must initiate the transition on all of the logical processors. • SW ANY—The OS Power Manager (OSPM) is responsible for coordinating the P-state among logical processors with dependencies (all logical processors in a package), and may initiate the transition on any of the logical processors in the domain. • Platform Default —The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor. <p>Note Power Technology must be set to Custom or the server ignores the setting for this parameter.</p>
DRAM Clock Throttling	<p>Allows you to tune the system settings between the memory bandwidth and power consumption. This can be one of the following:</p> <ul style="list-style-type: none"> • Auto — CPU determines the DRAM Clock Throttling settings. • Balanced— DRAM clock throttling is reduced, providing a balance between performance and power. • Performance—DRAM clock throttling is disabled, providing increased memory bandwidth at the cost of additional power. • Energy Efficient—DRAM clock throttling is increased to improve energy efficiency. • Platform Default —The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.

Name	Description
Channel Interleaving	<p>Whether the CPU divides memory blocks and spreads contiguous portions of data across interleaved channels to enable simultaneous read operations. This can be one of the following:</p> <ul style="list-style-type: none"> • Auto—The CPU determines what interleaving is done. • 1 Way— • 2 Way • 3 Way • 4-way—The maximum amount of channel interleaving is used. • Platform Default—The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.
Rank Interleaving	<p>Whether the CPU interleaves physical ranks of memory so that one rank can be accessed while another is being refreshed. This can be one of the following:</p> <ul style="list-style-type: none"> • Auto—The CPU determines what interleaving is done. • 1 Way— • 2 Way • 4-way • 8 Way—The maximum amount of rank interleaving is used. • Platform Default—The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.

Name	Description
Sub NUMA Clustering	<p>This setting determines if the CPU supports sub NUMA clustering, which keeps the tag directory and memory channel in the same region to improve memory performance for certain workloads. This can be one of the following:</p> <ul style="list-style-type: none"> • Auto—The BIOS determines the Sub-NUMA Clustering behavior automatically. • Disabled—Sub NUMA Clustering is turned off, and the CPU operates using a standard memory configuration without SNC optimization. This is the default option. • Enabled—Sub NUMA Clustering is activated, dividing the CPU into regions where the tag directory and memory channel remain in the same region for improved performance. • SNC 2—Sub NUMA Clustering divides the CPU into two NUMA regions, optimizing memory performance for workloads benefiting from moderate NUMA segmentation. • SNC 4—Sub NUMA Clustering divides the CPU into four NUMA regions, optimizing memory performance for workloads benefiting from moderate NUMA segmentation. • Platform Default— The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor. <p>Note The values SNC 2 and SNC 4 are not supported on Cisco UCS C220 M8, C240 M8, X210c M8 servers.</p>
IMC Interleaving	<p>This BIOS option controls the interleaving between the Integrated Memory Controllers (IMCs).</p> <ul style="list-style-type: none"> • 1-way Interleave—There is no interleaving. • 2-way Interleave—Addresses are interleaved between the two IMCs. • Auto —CPU determines the IMC Interleaving mode. • Platform Default —The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.

Name	Description
Memory Interleaving	<p>Whether the CPU interleaves the physical memory so that the memory can be accessed while another is being refreshed. This controls fabric level memory interleaving. Channel, die and socket have requirements based on memory populations and will be ignored if the memory does not support the selected option. This can be one of the following:</p> <ul style="list-style-type: none"> • None • Channel • Die • Socket • Auto—This is the default option. • Platform Default—The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.
Demand Scrub	<p>Whether the system corrects single bit memory errors encountered when the CPU or I/O makes a demand read. This can be one of the following:</p> <ul style="list-style-type: none"> • Disabled— Single bit memory errors are not corrected. • Enabled— Single bit memory errors are corrected in memory and the corrected data is set in response to the demand read. • Platform Default —The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.
Patrol Scrub	<p>Whether the system actively searches for, and corrects, single bit memory errors even in unused portions of the memory on the server. This can be one of the following:</p> <ul style="list-style-type: none"> • Disabled—The system checks for memory ECC errors only when the CPU reads or writes a memory address. • Enabled—The system periodically reads and writes memory searching for ECC errors. If any errors are found, the system attempts to fix them. This option may correct single bit errors before they become multi-bit errors, but it may adversely affect performance when the patrol scrub is running. • Platform Default —The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.

Name	Description
DCPMM Firmware Downgrade	<p>This can be one of the following:</p> <ul style="list-style-type: none"> • Disabled—Support is disabled. • Enabled—Support is enabled. • Platform Default —The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.
Configurable TDP Control	<p>Allows you to set customized value for Thermal Design Power (TDP). This can be one of the following:</p> <ul style="list-style-type: none"> • Auto— Uses the rated TDP value of the processor. • Manual—Allows you to customize the TDP value.
Altitude	<p>The approximate number of meters above sea level at which the physical server is installed. This can be one of the following:</p> <ul style="list-style-type: none"> • Auto—The CPU determines the physical elevation. • 300 M—The server is approximately 300 meters above sea level. • 900 M—The server is approximately 900 meters above sea level. • 1500 M—The server is approximately 1500 meters above sea level. • 3000 M—The server is approximately 3000 meters above sea level. • Platform Default —The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.

Name	Description
Package C State Limit	

Name	Description
	<p>The amount of power available to the server components when they are idle. This can be one of the following:</p> <p>Note Cisco recommends setting Package C State Limit as No Limit or Platform Default for C220 M4, C240 M4, C460 M4, and S3260 M4 servers.</p> <p>If you are changing the Package C State Limit token from any other value to No Limit, then ensure that the Power Technology is set to Custom.</p> <ul style="list-style-type: none"> • Auto —The CPU determines the physical elevation. • No Limit—The server may enter any available C state. • C0—The server provides all server components with full power at all times. This option maintains the highest level of performance and requires the greatest amount of power. • C1—When the CPU is idle, the system slightly reduces the power consumption. This option requires less power than C0 and allows the server to return quickly to high performance mode. • C3—When the CPU is idle, the system reduces the power consumption further than with the C1 option. This requires less power than C1 or C0, but it takes the server slightly longer to return to high performance mode. • C6—When the CPU is idle, the system reduces the power consumption further than with the C3 option. This option saves more power than C0, C1, or C3, but there may be performance issues until the server returns to full power. • C2—When the CPU is idle, the system reduces the power consumption further than with the C1 option. This requires less power than C1 or C0, but it takes the server slightly longer to return to high performance mode. • C7—When the CPU is idle, the server makes a minimal amount of power available to the components. This option saves the maximum amount of power but it also requires the longest time for the server to return to high performance mode. • C7s—When the CPU is idle, the server makes a minimal amount of power available to the components. This option saves more power than C7, but it also requires the longest time for the server to return to high performance mode. • c0 c1 state—This is the default option. • c2

Name	Description
	<ul style="list-style-type: none"> • c6 Non retention • Platform Default —The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.
CPU Hardware Power Management	<p>Enables processor Hardware Power Management (HWPM). This can be one of the following:</p> <ul style="list-style-type: none"> • Platform Default —The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor. • Disabled—HWPM is disabled. • HWPM Native Mode—HWPM native mode is enabled. • HWPM OOB Mode—HWPM Out-Of-Box mode is enabled. • Native Mode with no Legacy (only GUI)
Energy Performance Tuning	<p>Determines if the BIOS or Operating System can turn on the energy performance bias tuning. The options are BIOS and OS.</p> <ul style="list-style-type: none"> • BIOS— • OS— • Platform Default —The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.
Workload Configuration	<p>This feature allows for workload optimization. The options are Balanced and I/O Sensitive:</p> <ul style="list-style-type: none"> • Balanced • IO Sensitive—This is the default option. • NUMA — (Only GUI) • UMA — (Only GUI) • Platform Default —The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor. <p>Cisco recommends using Balanced.</p>

Name	Description
Core Performance Boost	<p>Whether the AMD processor increases its frequency on some cores when it is idle or not being used much. This can be one of the following:</p> <ul style="list-style-type: none"> • Auto—The CPU automatically determines how to boost performance. This is the default option • Disabled—Core performance boost is disabled. • Platform Default —The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.
Uncore Frequency Scaling	<p>Allows you configure the scaling of the uncore frequency of the processor. This can be one of the following:</p> <ul style="list-style-type: none"> • Mode 1—Uncore frequency of the processor scales up or down based on the load. (Default.) • Mode 0—Uncore frequency of the processor remains fixed. • Platform Default —The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor. <p>Refer to the Intel Dear Customer Letter (DCL) to know the fixed higher and lower values for Uncore Frequency Scaling.</p>
Configurable TDP Level	<p>Allows adjustments in processor thermal design power (TDP) values. By modifying the processor behavior and the performance levels, power consumption of a processor can be configured and TDP can be adjusted at the same time. Hence, a processor operates at higher or lower performance levels, depending on the available cooling capacities and desired power consumption.</p> <p>This can be one of the following:</p> <ul style="list-style-type: none"> • Normal—The CPU operates at its normal performance level. (Default.) • Level 1 • Level 2 <p>Note Refer to the Intel Dear Customer Letter (DCL) for the values for TDP level.</p>

Name	Description
UPI Link Speed	<p>Allows you to configure the Intel Ultra Path Interconnect (UPI) link speed between multiple sockets. This can be one of the following:</p> <ul style="list-style-type: none"> • Auto—Automatically configures the optimal link speed. (Default) • 9.6GT/s (gigatransfers per second)—Configures the optimal link speed at 9.6GT/s • 10.4GT/s—Configures the optimal link speed at 10.4GT/s • 11.2GT/s—Configures the optimal link speed at 11.2GT/s
Global C-state Control	<p>Whether the AMD processors control IO-based C-state generation and DF C-states. This can be one of the following:</p> <ul style="list-style-type: none"> • Auto—The CPU automatically determines how to control IO-based C-state generation. • Disabled—Global C-state control is disabled. This is the default option. • Enabled—Global C-state control is enabled. • Platform Default—The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.
L1 Stream HW Prefetcher	<p>Whether the processor allows the AMD hardware prefetcher to speculatively fetch streams of data and instruction from memory into the L1 cache when necessary. This can be one of the following:</p> <ul style="list-style-type: none"> • Auto—The CPU determines how to place data from I/O devices into the processor cache. This is the default option. • Disabled—The hardware prefetcher is not used. • Enabled—The processor uses the hardware prefetcher when cache issues are detected. • Platform Default—The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.

Name	Description
L2 Stream HW Prefetcher	<p>Whether the processor allows the AMD hardware prefetcher to speculatively fetch streams of data and instruction from memory into the L2 cache when necessary. This can be one of the following:</p> <ul style="list-style-type: none"> • Auto—The CPU determines how to place data from I/O devices into the processor cache. This is the default option. • Disabled—The hardware prefetcher is not used. • Enabled—The processor uses the hardware prefetcher when cache issues are detected. • Platform Default —The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.
AMD Memory Interleaving	<p>Whether the AMD CPU interleaves the physical memory so that the memory can be accessed while another is being refreshed. This controls fabric level memory interleaving. Channel, die and socket have requirements based on memory populations and will be ignored if the memory does not support the selected option. This can be one of the following:</p> <ul style="list-style-type: none"> • Auto—The CPU determines how to interleave memory. • Channel—Interleaves the physical address space over multiple channels, as opposed to each channel owning single consecutive address spaces. • Die—Interleaves the physical address space over multiple dies, as opposed to each die owning single consecutive address spaces. • None—Consecutive memory blocks are accessed from the same physical memory. • Socket—Interleaves the physical address space over multiple sockets, as opposed to each socket owning single consecutive address spaces. • Platform Default —The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.

Name	Description
AMD Memory Interleaving Size	<p>Determines the size of the memory blocks to be interleaved. It also determines the starting address of the interleave (bit 8,9,10 or 11). This can be one of the following:</p> <ul style="list-style-type: none"> • 1 KB • 2 KB • 256 Bytes • 512 Bytes • Auto—The CPU determines the size of the memory block. • Platform Default—The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.
Chipselect Interleaving	<p>Whether memory blocks across the DRAM chip selects for node 0 are interleaved. This can be one of the following:</p> <ul style="list-style-type: none"> • Auto—The CPU automatically determines how to interleave chip selects. This is the default option. • Disabled—Chip selects are not interleaved within the memory controller. • Enabled—Chip selects are interleaved within the memory controller. • Platform Default—The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.
Bank Group Swap	<p>Determines how physical addresses are assigned to applications. This can be one of the following:</p> <ul style="list-style-type: none"> • Auto—The CPU automatically determines how to assign physical addresses to applications. This is the default option. • Disabled—Bank group swap is not used. • Enabled—Bank group swap is used to improve the performance of applications. • Platform Default—The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.

Name	Description
Determinism Slider	<p>Allows AMD processors to determine how to operate. This can be one of the following:</p> <ul style="list-style-type: none"> • Auto—The CPU automatically uses default power determinism settings. This is the default option. • Performance—Processor operates at the best performance in a consistent manner. • Power—Processor operates at the maximum allowable performance on a per die basis. • Platform Default —The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.
IOMMU	<p>Input Output Memory Management Unit (IOMMU) allows AMD processors to map virtual addresses to physical addresses. This can be one of the following:</p> <ul style="list-style-type: none"> • Auto—The CPU determines how map these addresses. This is the default option. • Disabled—IOMMU is not used. • Enabled—Address mapping takes place through the IOMMU. • Platform Default —The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.
Downcore control	<p>Allows AMD processors to disable cores and, thus, select how many cores to enable. This can be one of the following:</p> <ul style="list-style-type: none"> • FOUR (2+2)—Two cores enabled on each CPU complex. • FOUR (4+0)—Four cores enabled on one CPU complex. • SIX (3+3)—Three cores enabled on each CPU complex. • THREE (3+0)—Three cores enabled on one CPU complex. • TWO (1+1)—Two cores enabled on each CPU complex. • TWO (2+0)—Two cores enabled on one CPU complex. • Auto—The CPU determines how many cores need to be enabled. • Platform Default —The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.

Name	Description
SVM Mode	<p>Whether the processor uses AMD Secure Virtual Machine Technology. This can be one of the following: This can be one of the following:</p> <ul style="list-style-type: none"> • Disabled—The processor does not use SVM Technology. • Enabled—The processor uses SVM Technology. This is the default option. • Platform Default —The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.
SMT Mode	<p>Whether the processor uses AMD Simultaneous MultiThreading Technology, which allows multithreaded software applications to execute threads in parallel within each processor. This can be one of the following:</p> <ul style="list-style-type: none"> • Disabled—The processor does not use SMT Technology. • Enabled—The processor uses SMT Technology. This is the default option. • Platform Default —The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.
SMEE	<p>Whether the processor uses the Secure Memory Encryption Enable (SMEE) function, which provides memory encryption support. This can be one of the following:</p> <ul style="list-style-type: none"> • Disabled—The processor does not use the SMEE function. • Enabled—The processor uses the SMEE function. • Platform Default—The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.
UPI Link Enablement	<p>Enables the number of Ultra Path Interconnect (UPI) links required by the processor. This can be one of the following</p> <ul style="list-style-type: none"> • Auto—This is the default option. • 1 • 2

Name	Description
UPI Power Management	<p>The UPI power management can be used for conserving power on the server. This can be one of the following:</p> <ul style="list-style-type: none"> • Enabled—Enables the processor to support this functionality. • Disabled—Disables the processor to support this functionality. This is the default option.
C1 Auto UnDemotion	<p>Select whether to enable processors to automatically undemote from C1. This can be one of the following:</p> <ul style="list-style-type: none"> • Auto— This is the default option. • Enabled—Enables the processor to support this functionality. • Disabled—Disables the processor to support this functionality.
C1 Auto Demotion	<p>If enabled, CPU automatically demotes to C1 based on un-core auto-demote information. This can be one of the following:</p> <ul style="list-style-type: none"> • Auto — This is the default option. • Enabled—Enables the processor to support this functionality. • Disabled—Disables the processor to support this functionality.
Memory Interleaving Size	<p>Determines the size of the memory blocks to be interleaved. It also determines the starting address of the interleave (bit 8, 9, 10 or 11). This can be one of the following:</p> <ul style="list-style-type: none"> • 1 KB • 2 KB • 256 Bytes • 512 Bytes • Auto—The CPU determines the size of the memory block. • Platform Default —The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.

Name	Description
UPI Link Frequency Select	<p>Allows you to enable or disable extended APIC support. This can be one of the following:</p> <ul style="list-style-type: none"> • Auto—This option configures the optimal link speed automatically. This is the default option. • 9.6GT/S—This option configures the optimal link speed at 9.6GT/s. • 10.4GT/S—This option configures the optimal link speed at 10.4GT/s. • 11.2GT/S—This option configures the optimal link speed at 10.4GT/s. • 12.8GT/S—This option configures the optimal link speed at 12.8GT/s. • 14.4GT/S—This option configures the optimal link speed at 14.4GT/s. • 16.0GT/S—This option configures the optimal link speed at 16.0GT/s. • 20.0GT/S—This option configures the optimal link speed at 20.0GT/s.
Burst and Postponed Refresh	<p>Allows the memory controller to defer the refresh cycles when the memory is active and accomplishes the refresh within a specified window. The deferred refresh cycles may run in a burst of several refresh cycles. This can be one of the following:</p> <ul style="list-style-type: none"> • Enabled • Disabled—This is the default option. <p>Note It is recommended to leave this setting in the default state of Disabled to mitigate Rowhammer-style attacks.</p>
Operation Mode	<p>Allows you to set the Operation Mode. This can be one of the following:</p> <ul style="list-style-type: none"> • Test Only—Support is enabled. • Test and Repair—Support is disabled.

Name	Description
xGMI Force Link Width	<p>This setting allows users to manually specify the number of lanes used for the xGMI (Inter-chip Global Memory Interconnect) link width to x4/x8/x16.</p> <p>This can be one of the following:</p> <ul style="list-style-type: none">• 0 - Force xGMI link width to x4.• 1 - Force xGMI link width to x8.• 2 - Force xGMI link width to x16.• Auto - Use the default xGMI link width controller settings. This is the default option.• Platform Default —The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.

Name	Description
CPU Downcore control F19 M10h-1Fh	<p>Enables manage the number of active cores on AMD® Family 19h processors. This token can be used to optimize performance, power consumption, or compatibility based on specific needs. F refers to the processor family and M refers to the model. The available options include:</p> <ul style="list-style-type: none"> • Auto —The system automatically selects the optimal number of active cores based on the current workload and system configuration. This is the default option. • ONE (1+_0) — Enables only one core per CPU. • TWO (2+_0) —Enables two cores per CPU. • THREE (3+_0) —Enables three cores per CPU. • FOUR (4+_0) —Enables four cores per CPU. • FIVE (5+_0) —Enables five cores per CPU. • SIX (6+_0) —Enables six cores per CPU. • SEVEN (7+_0) —Enables seven cores per CPU. • EIGHT (8+_0) —Enables eight cores per CPU. • NINE (9+_0) —Enables nine cores per CPU. • TEN (10+_0) —Enables ten cores per CPU. • ELEVEN (11+_0) —Enables eleven cores per CPU. • TWELVE (12+_0) —Enables twelve cores per CPU. • THIRTEEN (13+_0) —Enables thirteen cores per CPU. • FOURTEEN (14+_0) —Enables fourteen cores per CPU. • FIFTEEN (15+_0) —Enables fifteen cores per CPU. <p>Note The values from <i>Eight (8+_0)</i> to <i>Fifteen (15+_0)</i> are applicable only for 5th Generation AMD® EPYC® 9xx5 processors.</p>

I/O BIOS Settings for Intel

The following table lists the Intel Directed I/O BIOS settings that you can configure through a BIOS policy or the default BIOS settings:

Name	Description
Intel VT for directed IO	<p>Whether the processor uses Intel Virtualization Technology for Directed I/O (VT-d). This can be one of the following:</p> <ul style="list-style-type: none"> • Disabled—The processor does not use virtualization technology. • Enabled—The processor uses virtualization technology. • Platform Default—The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor. <p>Note This option must be enabled if you want to change any of the other Intel Directed I/O BIOS settings.</p>
Intel VTD interrupt Remapping	<p>Whether the processor supports Intel VT-d Interrupt Remapping. This can be one of the following:</p> <ul style="list-style-type: none"> • Disabled—The processor does not support remapping. • Enabled—The processor uses VT-d Interrupt Remapping as required. • Platform Default—The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.
Intel VTD coherency support	<p>Whether the processor supports Intel VT-d Coherency. This can be one of the following:</p> <ul style="list-style-type: none"> • Disabled—The processor does not support coherency. • Enabled—The processor uses VT-d Coherency as required. • Platform Default—The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.
Intel VTD ATS support	<p>Whether the processor supports Intel VT-d Address Translation Services (ATS). This can be one of the following:</p> <ul style="list-style-type: none"> • Disabled—The processor does not support ATS. • Enabled—The processor uses VT-d ATS as required. • Platform Default—The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.

Name	Description
Intel VTD pass through DMA support	<p>Whether the processor supports Intel VT-d Pass-through DMA. This can be one of the following:</p> <ul style="list-style-type: none"> • Disabled—The processor does not support pass-through DMA. • Enabled—The processor uses VT-d Pass-through DMA as required. • Platform Default—The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.

I/O BIOS Settings for AMD

The following table lists the Input/Output BIOS settings that you can configure through a BIOS policy for AMD:

Name	Description
PCIe ARI Support	<p>The PCIe Alternative Routing ID (ARI) Interpretation feature specification supports greater numbers of virtual functions through the implementation of ARI, which reinterprets the device number field in the PCIe header allowing for more than eight functions. This can be one of the following:</p> <ul style="list-style-type: none"> • Disabled—PCIe ARI Support is not available. • Enabled—PCIe ARI Support is available. • Auto—PCIe ARI Support is in auto mode. This is the default option.
IPv4 PXE Support	<p>Enables or disables IPv4 support for PXE. This can be one of the following:</p> <ul style="list-style-type: none"> • Disabled—IPv4 PXE support is not available. • Enabled—IPv4 PXE support is available. This is the default option.
IPv4 HTTP Support	<p>Enables or disables IPv4 support for HTTP. This can be one of the following:</p> <ul style="list-style-type: none"> • Disabled—IPv4 HTTP support is not available. • Enabled—IPv4 HTTP support is available. This is the default option.

Name	Description
IPv6 HTTP Support	<p>Enables or disables IPv6 support for HTTP. This can be one of the following:</p> <ul style="list-style-type: none"> • Disabled—IPv6 HTTP support is not available. • Enabled—IPv6 HTTP support is available. This is the default option.
Network Stack	<p>This option allows you to monitor IPv6 and IPv4. This can be one of the following</p> <ul style="list-style-type: none"> • Disabled—Network Stack support is not available. <p>Note When disabled, the value set for IPV4 PXE Support does not impact the system.</p> <ul style="list-style-type: none"> • Enabled—Network Stack support is available. This is the default option. <p>Note When Network Stack token value is Disabled, the below tokens and their values are also set</p> <ul style="list-style-type: none"> • IPV4PXE - Disabled • IPV4HTTP - Disabled • IPV6HTTP - Disabled

RAS Memory BIOS Settings

The following table lists the RAS memory BIOS settings that you can configure through a BIOS policy or the default BIOS settings:

Name	Description
ACPI SRAT Special Purpose Memory Flag	<p>This setting allows or disallows the ACPI SRAT special purpose memory flag to be set when the UEFI Memory Map special purpose flag is enabled. This can be one of the following:</p> <ul style="list-style-type: none"> • Disabled • Enabled—This is the default option.
UEFI Memory Map Special Purpose Memory Flag	<p>This setting determines the behavior of the UEFI memory map special knob, which impacts CXL cards under certain operating systems. This can be one of the following:</p> <ul style="list-style-type: none"> • Disabled • Enabled—This is the default option.

Name	Description
Memory Refresh Rate	<p>Controls the refresh rate of the memory controller and might affect the memory performance and power depending on memory configuration and workload. This can be one of the following:</p> <ul style="list-style-type: none">• 1x Refresh—3.9us. This is the default option.• 2x Refresh
Panic and High Watermark	<p>Controls the delayed refresh capability of the memory controller. This can be one of the following:</p> <ul style="list-style-type: none">• High—The memory controller is allowed to postpone up to a maximum of eight refresh commands. The memory controller executes all the postponed refreshes within the refresh interval. For the ninth refresh command, the refresh priority becomes Panic and the memory controller pauses the normal memory transactions until all the postponed refresh commands are executed.• Low—This is the default option. The memory controller is not allowed to postpone refresh commands. <p>Note It is recommended to leave this setting in the default state (Low) which will help to reduce susceptibility to Rowhammer-style attacks.</p>

Name	Description
Memory RAS configuration	<p>How the memory reliability, availability, and serviceability (RAS) is configured for the server. This can be one of the following:</p> <ul style="list-style-type: none"> • Maximum Performance—Optimizes the system performance and disables all the advanced RAS features. • Mirroring—System reliability is optimized by using half the system memory as backup. This mode is used for UCS M4 and lower blade servers. • Lockstep—If the DIMM pairs in the server have an identical type, size, and organization and are populated across the SMI channels, you can enable lockstep mode to minimize memory access latency and provide better performance. Lockstep is enabled by default for B440 servers. • Mirror Mode 1LM—Mirror Mode 1LM will set the entire 1LM memory in the system to be mirrored, consequently reducing the memory capacity by half. This mode is used for UCS M5 blade servers. • Partial Mirror Mode 1LM—Partial Mirror Mode 1LM will set a part of the 1LM memory in the system to be mirrored, consequently reducing the memory capacity by half. This mode is used for UCS M5 blade servers. • Sparing—System reliability is optimized by holding memory in reserve so that it can be used in case other DIMMs fail. This mode provides some memory redundancy, but does not provide as much redundancy as mirroring. • ADDDC Sparing—System reliability is optimized by holding memory in reserve so that it can be used in case other DIMMs fail. This mode provides some memory redundancy, but does not provide as much redundancy as mirroring. • Platform Default—The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.

Name	Description
NUMA optimized	<p>Whether the BIOS supports NUMA. This can be one of the following:</p> <ul style="list-style-type: none"> • Disabled—The BIOS does not support NUMA. • Enabled—The BIOS includes the ACPI tables that are required for NUMA-aware operating systems. If you enable this option, the system must disable Inter-Socket Memory interleaving on some platforms. • Platform Default—The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.
Select PPR Type Configuration	<p>Post Package Repair (PPR) provides the ability to repair faulty memory cells by replacing them with spare cells. This can be one of the following:</p> <ul style="list-style-type: none"> • Disabled—The BIOS does not support selecting PPR Type. • Hard PPR—This results in a permanent remapping of damaged storage cells. • Platform Default—The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.
Memory Size Limit in GB	<p>Limits the capacity in Partial Memory Mirror Mode up to 50 percent of the total memory capacity. The memory size can range from 0 GB to 65535 GB in increments of 1 GB.</p>
Mirroring Mode	<p>Memory mirroring enhances system reliability by keeping two identical data images in memory.</p> <p>This option is only available if you choose the mirroring option for Memory RAS Config. It can be one of the following:</p> <ul style="list-style-type: none"> • Inter-Socket—Memory is mirrored between two Integrated Memory Controllers (IMCs) across CPU sockets. • Intra-Socket—One IMC is mirrored with another IMC in the same socket. • Platform Default—The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.

Name	Description
Sparing Mode	<p>Sparing optimizes reliability by holding memory in reserve so that it can be used in case other DIMMs fail. This option provides some memory redundancy, but does not provide as much redundancy as mirroring. The available sparing modes depend on the current memory population.</p> <p>This option is only available if you choose sparing option for Memory RAS Config. It can be one of the following:</p> <ul style="list-style-type: none"> • DIMM Sparing—One DIMM is held in reserve. If a DIMM fails, the contents of a failing DIMM are transferred to the spare DIMM. • Rank Sparing—A spare rank of DIMMs is held in reserve. If a rank of DIMMs fails, the contents of the failing rank are transferred to the spare rank. • Platform Default—The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.
LV DDR Mode	<p>Whether the system prioritizes low voltage or high frequency memory operations. This can be one of the following:</p> <ul style="list-style-type: none"> • Auto—The CPU determines whether to prioritize low voltage or high frequency memory operations. • Power Saving Mode—The system prioritizes low voltage memory operations over high frequency memory operations. This mode may lower memory frequency in order to keep the voltage low. • Performance Mode—The system prioritizes high frequency operations over low voltage operations. • Platform Default—The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.
DRAM Refresh Rate	<p>The refresh interval rate for internal memory. This can be one of the following:</p> <ul style="list-style-type: none"> • 1x • 2x • 3x • 4x • Auto • Platform Default—The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.

Name	Description
DDR3 Voltage Selection	<p>The voltage to be used by the dual-voltage RAM. This can be one of the following:</p> <ul style="list-style-type: none"> • DDR3-1500mv • DDR3-1350mv • Platform Default—The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.
Partial Memory Mirror Mode	<p>Partial Memory Mirroring enables you to partially mirror by GB or by a percentage of the memory capacity. Depending on the option selected here, you can define either a partial mirror percentage or a partial mirror capacity in GB in available fields. You can partially mirror up to 50 percent of the memory capacity. It can be one of the following:</p> <ul style="list-style-type: none"> • Disabled—Partial Memory Mode is disabled. This is the default option. • Percentage—The amount of memory to be mirrored in the Partial Memory Mode is defined as a percentage of the total memory. • Value in GB—The amount of memory to be mirrored in the Partial Memory Mode is defined in GB. • Platform Default—The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor. <p>Note Partial Memory Mirror Mode is mutually exclusive to standard Mirroring Mode.</p> <p>Partial Mirrors 1-4 can be used in any number or configuration, provided they do not exceed the capacity limit set in GB or Percentage in the related options.</p>
Partial Mirror Percentage	Limits the amount of available memory to be mirrored as a percentage of the total memory. This can range from 0.00 % to 50.00 % in increments of 0.01 %.
Partial Mirror1 Size in GB	Limits the amount of memory in Partial Mirror1 in GB. This can range from 0 GB to 65535 GB in increments of 1 GB.
Partial Mirror2 Size in GB	Limits the amount of memory in Partial Mirror2 in GB. This can range from 0 GB to 65535 GB in increments of 1 GB.
Partial Mirror3 Size in GB	Limits the amount of memory in Partial Mirror3 in GB. This can range from 0 GB to 65535 GB in increments of 1 GB.

Name	Description
Partial Mirror4 Size in GB	Limits the amount of memory in Partial Mirror4 in GB. This can range from 0 GB to 65535 GB in increments of 1 GB.
CR FastGo Config	<p>CR FastGo Config improves DDRT non-temporal write bandwidth when FastGO is disabled. When FastGO is enabled, it gives faster flow of NT writes into the uncore, When FastGO is disabled, it lessens NT writes queueing up in the CPU uncore, thereby improving sequentially at DCPMM, resulting in improved bandwidth.</p> <p>Applies to all Cisco UCS M5 and Cisco UCS M6 servers.</p> <p>The values can be one of the following:</p> <ul style="list-style-type: none"> • Default—Enables FastGO. • Option 1—Disables FastGO. • Option 2, Option 3, Option 4, Option 5—Not applicable.
CR QoS	<p>Prevents DRAM and overall system BW drop in the presence of concurrent DCPMM BW saturating threads, with minimal impact to homogenous DDRT-only usages, Good for multi-tenant use cases, VMs, etc. Targeted for App Direct, but also improves memory mode. Targets the “worst-case” degradations.</p> <p>Applies to all M5 and M6 servers.</p> <p>The values can be one of the following:</p> <ul style="list-style-type: none"> • Disabled—Feature disabled. This is the default option. • Recipe 1—6 modules, 4 modules per socket optimized • Recipe 2—2 modules per socket optimized • Recipe 3—1 module per socket optimized
NVM Performance Setting	<p>NVM Performance Setting enables efficient major mode arbitration between DDR and DDRT transactions on the DDR channel to optimize channel BW and DRAM latency.</p> <p>Applies to all M5 and M6 servers.</p> <p>The values can be one of the following:</p> <ul style="list-style-type: none"> • BW Optimized—Optimized for DDR and DDRT BW. This is the default option. • Latency Optimized—Better DDR latency in the presence of DDRT BW. • Balanced Profile—Optimized for Memory mode.

Name	Description
Snoopy mode for 2LM	<p>Enables snoop-mode for DCPMM accesses while maintaining directory on all DRAM accesses. Snoops maintain cache coherence between sockets. Directory reduces snoops by keeping the remote node information locally (in memory). Directory lookups and updates add memory traffic.</p> <p>Directory is a good tradeoff for DRAM, but not necessarily for DCPMM. For non-NUMA workload, when the feature is enabled, directory updates to DCPMM are eliminated, thereby helping DDRT bandwidth bound workloads. Directory is disabled for far memory accesses and instead snoops remote sockets to check for ownership. Directory is used only for DRAM (near memory).</p> <ul style="list-style-type: none"> • Enabled • Disabled—This is the default option.
Snoopy mode for AD	<p>Enables snoop-mode for DCPMM accesses while maintaining directory on all DRAM accesses. Snoops maintain cache coherence between sockets. Directory reduces snoops by keeping the remote node information locally (in memory). Directory lookups and updates add memory traffic.</p> <p>Directory is a good tradeoff for DRAM, but not necessarily for DCPMM. For non-NUMA workload, when the feature is enabled, directory updates to DCPMM are eliminated, thereby helping DDRT bandwidth bound workloads. Directory is disabled for accesses to AD and instead snoops remote sockets to check for ownership. Directory is used only for DRAM accesses.</p> <ul style="list-style-type: none"> • Enabled • Disabled—This is the default option.

Intel® Optane™ DC Persistent Memory (DCPMM) BIOS Tokens

The following table lists the Intel® Optane™ DC Persistent Memory (DCPMM) BIOS settings that you can configure through a BIOS policy or the default BIOS settings:

Name	Description
NVM Performance Setting	<p>NVM Performance Setting enables efficient major mode arbitration between DDR and DDRT transactions on the DDR channel to optimize channel BW and DRAM latency.</p> <p>Applies to all M5 and M6 servers.</p> <p>The values can be one of the following:</p> <ul style="list-style-type: none"> • BW Optimized—Optimized for DDR and DDRT BW. This is the default option. • Latency Optimized—Better DDR latency in the presence of DDRT BW. • Balanced Profile—Optimized for Memory mode.
CR QoS	<p>Prevents DRAM and overall system BW drop in the presence of concurrent DCPMM BW saturating threads, with minimal impact to homogenous DDRT-only usages, Good for multi-tenant use cases, VMs, etc. Targeted for App Direct, but also improves memory mode. Targets the “worst-case” degradations.</p> <p>Applies to all M5 and M6 servers.</p> <p>The values can be one of the following:</p> <ul style="list-style-type: none"> • Disabled—Feature disabled. This is the default option. • Recipe 1—6 modules, 4 modules per socket optimized • Recipe 2—2 modules per socket optimized • Recipe 3—1 module per socket optimized
CR FastGo Config	<p>CR FastGo Config improves DDRT non-temporal write bandwidth when FastGO is disabled. When FastGO is enabled, it gives faster flow of NT writes into the uncore, When FastGO is disabled, it lessens NT writes queueing up in the CPU uncore, thereby improving sequentially at DCPMM, resulting in improved bandwidth.</p> <p>Applies to all Cisco UCS M5 and Cisco UCS M6 servers.</p> <p>The values can be one of the following:</p> <ul style="list-style-type: none"> • Default—Enables FastGO. • Option 1—Disables FastGO. • Option 2, Option 3, Option 4, Option 5—Not applicable.

Name	Description
Snoopy mode for AD	<p>Enables snoopy-mode for DCPMM accesses while maintaining directory on all DRAM accesses. Snoops maintain cache coherence between sockets. Directory reduces snoops by keeping the remote node information locally (in memory). Directory lookups and updates add memory traffic.</p> <p>Directory is a good tradeoff for DRAM, but not necessarily for DCPMM. For non-NUMA workload, when the feature is enabled, directory updates to DCPMM are eliminated, thereby helping DDRT bandwidth bound workloads. Directory is disabled for accesses to AD and instead snoops remote sockets to check for ownership. Directory is used only for DRAM accesses.</p> <ul style="list-style-type: none"> • Enabled • Disabled This is the default option.
Snoopy mode for 2LM	<p>Enables snoopy-mode for DCPMM accesses while maintaining directory on all DRAM accesses. Snoops maintain cache coherence between sockets. Directory reduces snoops by keeping the remote node information locally (in memory). Directory lookups and updates add memory traffic.</p> <p>Directory is a good tradeoff for DRAM, but not necessarily for DCPMM. For non-NUMA workload, when the feature is enabled, directory updates to DCPMM are eliminated, thereby helping DDRT bandwidth bound workloads. Directory is disabled for far memory accesses and instead snoops remote sockets to check for ownership. Directory is used only for DRAM (near memory).</p> <ul style="list-style-type: none"> • Enabled • Disabled This is the default option.

Serial Port BIOS Settings

The following table lists the serial port BIOS settings that you can configure through a BIOS policy or the default BIOS settings:

Name	Description
Serial port A enable	<p>Whether serial port A is enabled or disabled. This can be one of the following:</p> <ul style="list-style-type: none"> • Disabled—The serial port is disabled. • Enabled—The serial port is enabled. • Platform Default—The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.

USB BIOS Settings

The following table lists the USB BIOS settings that you can configure through a BIOS policy or the default BIOS settings:

Name	Description
Make Device Non Bootable	<p>Whether the server can boot from a USB device. This can be one of the following:</p> <ul style="list-style-type: none"> • Disabled—The server can boot from a USB device. • Enabled—The server cannot boot from a USB device. • Platform Default—The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.
Legacy USB Support	<p>Whether the system supports legacy USB devices. This can be one of the following:</p> <ul style="list-style-type: none"> • Disabled—USB devices are only available to EFI applications. • Enabled—Legacy USB support is always available. • Auto—Disables legacy USB support if no USB devices are connected. • Platform Default—The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.
USB Idle Power Optimizing Setting	<p>Whether the USB Idle Power Optimizing setting is used to reduce USB EHCI idle power consumption. Depending upon the value you choose, this setting can have an impact on performance. This can be one of the following:</p> <ul style="list-style-type: none"> • high-performanceHigh Performance—The USB System Idle Power Optimizing setting is disabled, because optimal performance is preferred over power savings. Selecting this option can significantly improve performance. We recommend you select this option unless your site has server power restrictions. • Lower Idle Power—The USB System Idle Power Optimizing setting is enabled, because power savings are preferred over optimal performance. • Platform Default—The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.

Name	Description
USB Front Panel Access Lock	<p>USB front panel access lock is configured to enable or disable the front panel access to USB ports. This can be one of the following:</p> <ul style="list-style-type: none"> • Disabled • Enabled • Platform Default—The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.
Port 60/64 Emulation	<p>Whether the system supports 60h/64h emulation for complete USB keyboard legacy support. This can be one of the following:</p> <ul style="list-style-type: none"> • Disabled—60h/64 emulation is not supported. • Enabled—60h/64 emulation is supported. <p>You should select this option if you are using a non-USB aware operating system on the server.</p> <ul style="list-style-type: none"> • Platform Default—The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.
USB Port Front	<p>Whether the front panel USB devices are enabled or disabled. This can be one of the following:</p> <ul style="list-style-type: none"> • Disabled—Disables the front panel USB ports. Devices connected to these ports are not detected by the BIOS and operating system. • Enabled—Enables the front panel USB ports. Devices connected to these ports are detected by the BIOS and operating system. • Platform Default—The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.

Name	Description
USB Port Internal	<p>Whether the internal USB devices are enabled or disabled. This can be one of the following:</p> <ul style="list-style-type: none"> • Disabled—Disables the internal USB ports. Devices connected to these ports are not detected by the BIOS and operating system. • Enabled—Enables the internal USB ports. Devices connected to these ports are detected by the BIOS and operating system. • Platform Default—The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.
USB Port KVM	<p>Whether the vKVM ports are enabled or disabled. This can be one of the following:</p> <ul style="list-style-type: none"> • Disabled—Disables the KVM keyboard and/or mouse devices. Keyboard and/or mouse will not work in the KVM window. • Enabled—Enables the KVM keyboard and/or mouse devices. • Platform Default—The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.
USB Port Rear	<p>Whether the rear panel USB devices are enabled or disabled. This can be one of the following:</p> <ul style="list-style-type: none"> • Disabled—Disables the rear panel USB ports. Devices connected to these ports are not detected by the BIOS and operating system. • Enabled—Enables the rear panel USB ports. Devices connected to these ports are detected by the BIOS and operating system. • Platform Default—The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.

Name	Description
USB Port SD Card	<p>Whether the SD card drives are enabled or disabled. This can be one of the following:</p> <ul style="list-style-type: none"> • Disabled—Disables the SD card drives. The SD card drives are not detected by the BIOS and operating system. • Enabled—Enables the SD card drives. • Platform Default—The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.
USB Port VMedia	<p>Whether the virtual media devices are enabled or disabled. This can be one of the following:</p> <ul style="list-style-type: none"> • Disabled—Disables the vMedia devices. • Enabled—Enables the vMedia devices. • Platform Default—The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.
All USB Devices	<p>Whether all physical and virtual USB devices are enabled or disabled. This can be one of the following:</p> <ul style="list-style-type: none"> • Disabled—All USB devices are disabled. • Enabled—All USB devices are enabled. • Platform Default—The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.
xHCI Mode	<p>Whether xHCI mode is enabled or disabled. This can be one of the following:</p> <ul style="list-style-type: none"> • Disabled—xHCI mode is disabled. • Enabled—xHCI mode is enabled. • Platform Default—The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.

PCI Configuration BIOS Settings

The following table lists the PCI configuration BIOS settings that you can configure through a BIOS policy or the default BIOS settings:

Name	Description
Maximum memory below 4GB	<p>Whether the BIOS maximizes memory usage below 4GB for an operating system without PAE support, depending on the system configuration. This can be one of the following:</p> <ul style="list-style-type: none"> • Disabled—Does not maximize memory usage. Choose this option for all operating systems with PAE support. • Enabled—Maximizes memory usage below 4GB for an operating system without PAE support. • Platform Default—The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.
Memory mapped IO above 4GB	<p>Whether to enable or disable memory mapped I/O of 64-bit PCI devices to 4GB or greater address space. Legacy option ROMs are not able to access addresses above 4GB. PCI devices that are 64-bit compliant but use a legacy option ROM may not function correctly with this setting enabled. This can be one of the following:</p> <ul style="list-style-type: none"> • Disabled—Does not map I/O of 64-bit PCI devices to 4GB or greater address space. • Enabled—Maps I/O of 64-bit PCI devices to 4GB or greater address space. • Platform Default—The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.

Name	Description
VGA Priority	<p>Allows you to set the priority for VGA graphics devices if multiple VGA devices are found in the system. This can be one of the following:</p> <ul style="list-style-type: none"> • Onboard—Priority is given to the onboard VGA device. BIOS post screen and OS boot are driven through the onboard VGA port. • Offboard—Priority is given to the PCIE Graphics adapter. BIOS post screen and OS boot are driven through the external graphics adapter port. • Onboard VGA Disabled—Priority is given to the PCIE Graphics adapter, and the onboard VGA device is disabled. <p>Note The vKVM does not function when the onboard VGA is disabled.</p> <ul style="list-style-type: none"> • Platform Default—The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor. <p>Note Only onboard VGA devices are supported with Cisco UCS B-Series servers.</p>
ASPM Support	<p>Allows you to set the level of ASPM (Active Power State Management) support in the BIOS. This can be one of the following:</p> <ul style="list-style-type: none"> • Disabled—ASPM support is disabled in the BIOS. • Auto—The CPU determines the power state. • ForceL0—Force all links to L0 standby (L0s) state. • Platform Default—The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.
BME DMA Mitigation Support	<p>Allows you to disable the PCI BME bit to mitigate the threat from an unauthorized external DMA. This can be one of the following:</p> <ul style="list-style-type: none"> • Disabled—PCI BME bit is disabled in the BIOS. • Enabled—PCI BME bit is enabled in the BIOS. • Platform Default—The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.

QPI BIOS Settings

The following table lists the QPI BIOS settings that you can configure through a BIOS policy or the default BIOS settings:

Name	Description
QPI Link Frequency Select	<p>The Intel QuickPath Interconnect (QPI) link frequency, in megatransfers per second (MT/s). This can be one of the following:</p> <ul style="list-style-type: none"> • 20.0GT/s • 16.0GT/s • 14.4GT/s • 12.8GT/s • 6.4 GT/s • 7.2 GT/s • 8.0 GT/s • 9.6 GT/s • Auto—The CPU determines the QPI link frequency. • Platform Default—The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.
QPI Snoop Mode	<p>This can be one of the following:</p> <ul style="list-style-type: none"> • Home Snoop—The snoop is always spawned by the home agent (centralized ring stop) for the memory controller. This mode has a higher local latency than early snoop, but it provides extra resources for a larger number of outstanding transactions. • Cluster On Die—This mode is available only for processors that have 10 or more cores. It is the best mode for highly NUMA optimized workloads. • Home Directory Snoop with OSB • Early Snoop—The distributed cache ring stops can send a snoop probe or a request to another caching agent directly. This mode has lower latency and it is best for workloads that have shared data sets across threads and can benefit from a cache-to-cache transfer, or for workloads that are not NUMA optimized. • Auto —The CPU determines the QPI Snoop mode. • Platform Default—The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.

Trusted Platform BIOS Settings

The following table lists the trusted platform BIOS settings that you can configure through a BIOS policy or the default BIOS settings:

Name	Description
Trusted Platform Module (TPM) Support	<p>Whether to enable or disable the Trusted Platform Module (TPM), which is a component that securely stores artifacts that are used to authenticate the server. This can be one of the following:</p> <ul style="list-style-type: none"> • Disabled—Disables TPM. • Enabled—Enables TPM. • Platform Default—Enables TPM.
Intel Trusted Execution Technology (TXT) Support	<p>Whether to enable or disable Intel Trusted Execution Technology (TXT), which provides greater protection for information that is used and stored on the business server. This can be one of the following:</p> <ul style="list-style-type: none"> • Disabled—Disables TXT. This is default option. • Enabled—Enables TXT. • Platform Default—Disables TXT. <p>When you only enable TXT, it implicitly enables TPM, VT, and VTd.</p>
SHA-256 PCR Bank	<p>The Platform Configuration Register (PCR) is a memory location in the TPM. Multiple PCRs are collectively referred to as a PCR bank. A Secure Hash Algorithm 256-bit or SHA-256 PCR Bank allows to enable or disable TPM security. This can be one of the following:</p> <ul style="list-style-type: none"> • Disabled—Disables SHA-256 PCR Bank. • Enabled—Enables SHA-256 PCR Bank. This is the default option.

LOM and PCIe Slots BIOS Settings

The following table lists the USB BIOS settings that you can configure through a BIOS policy or the default BIOS settings:

Name	Description
PCIe Slot SAS OptionROM	<p>Whether Option ROM is available on the SAS port. This can be one of the following:</p> <ul style="list-style-type: none"> • Disabled—The expansion slot is not available. • Enabled—The expansion slot is available. • UEFI Only—The expansion slot is available for UEFI only. • Legacy Only—The expansion slot is available for legacy only. • Platform Default—The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.

Name	Description
PCIe Slot <i>n</i> Link Speed	<p>This option allows you to restrict the maximum speed of an adapter card installed in PCIe slot <i>n</i>. This can be one of the following:</p> <ul style="list-style-type: none"> • GEN 1—2.5GT/s (gigatransfers per second) is the maximum speed allowed. • GEN 2—5GT/s is the maximum speed allowed. • GEN 3—8GT/s is the maximum speed allowed. • Auto—The maximum speed is set automatically. This is the default option. • Disabled—The maximum speed is not restricted. • Platform Default—The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.
PCIe Slot <i>n</i> OptionROM	<p>Whether Option ROM is available on the port. This can be one of the following:</p> <ul style="list-style-type: none"> • Disabled—The expansion slot is not available. • Enabled—The expansion slot is available. This is the default option. • Platform Default—The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.
PCIe Slot HBA OptionROM	<p>Whether Option ROM is available on the HBA port. This can be one of the following:</p> <ul style="list-style-type: none"> • Disabled—The expansion slot is not available. • Enabled—The expansion slot is available. • UEFI Only—The expansion slot is available for UEFI only. • Legacy Only—The expansion slot is available for legacy only. • Platform Default—The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.
PCIe Slot MLOM OptionROM	<p>Whether Option ROM is available on the MLOM port. This can be one of the following:</p> <ul style="list-style-type: none"> • Disabled—The expansion slot is not available. • Enabled—The expansion slot is available. • UEFI Only—The expansion slot is available for UEFI only. • Legacy Only—The expansion slot is available for legacy only. • Platform Default—The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.

Name	Description
PCIe Slot Nx OptionROM	<p>Whether Option ROM is available on the port. This can be one of the following:</p> <ul style="list-style-type: none"> • Disabled—The expansion slot is not available. • Enabled—The expansion slot is available. • UEFI Only—The expansion slot is available for UEFI only. • Legacy Only—The expansion slot is available for legacy only. • Platform Default—The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.
PCIe 10G LOM 2 Link	<p>Whether Option ROM is available on the 10G LOM port. This can be one of the following:</p> <ul style="list-style-type: none"> • Disabled—The expansion slot is not available. • Enabled—The expansion slot is available. • Platform Default—The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.
PCI ROM CLP	<p>PCI ROM Command Line Protocol (CLP) controls the execution of different Option ROMs such as PxE and iSCSI that are present in the card. By default, it is disabled.</p> <ul style="list-style-type: none"> • Disabled—The expansion slot is not available. • Enabled—The expansion slot is available. • Platform Default—The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.
SIOC1 Option ROM	<p>Whether the server can use Option ROM present in System IO Controller 1 (SIOC1). This can be one of the following:</p> <ul style="list-style-type: none"> • Disabled—The expansion slot is not available. • Enabled—The expansion slot is available. • UEFI Only—The expansion slot is available for UEFI only. • Legacy Only—The expansion slot is available for legacy only. • Platform Default—The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.

Name	Description
SIOC2 Option ROM	<p>Whether the server can use Option ROM present in System IO Controller 2 (SIOC2). This can be one of the following:</p> <ul style="list-style-type: none"> • Disabled—The expansion slot is not available. • Enabled—The expansion slot is available. • UEFI Only—The expansion slot is available for UEFI only. • Legacy Only—The expansion slot is available for legacy only. • Platform Default—The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.
SBMEZZ1 Option ROM	<p>Whether the server can use Option ROM present in SBMezz1 controller. This can be one of the following:</p> <ul style="list-style-type: none"> • Disabled—The expansion slot is not available. • Enabled—The expansion slot is available. • UEFI Only—The expansion slot is available for UEFI only. • Legacy Only—The expansion slot is available for legacy only. • Platform Default—The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.
SBMEZZ2 Option ROM	<p>Whether the server can use Option ROM present in SBMezz2 controller. This can be one of the following:</p> <ul style="list-style-type: none"> • Disabled—The expansion slot is not available. • Enabled—The expansion slot is available. • UEFI Only—The expansion slot is available for UEFI only. • Legacy Only—The expansion slot is available for legacy only. • Platform Default—The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.
IOESlot1 OptionROM	<p>Whether option ROM is enabled on the IOE slot 1. This can be one of the following:</p> <ul style="list-style-type: none"> • Disabled—The expansion slot is not available. • Enabled—The expansion slot is available. • UEFI Only—The expansion slot is available for UEFI only. • Legacy Only—The expansion slot is available for legacy only. • Platform Default—The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.

Name	Description
IOEMEZ1 OptionROM	<p>Whether option ROM is enabled on the IOE Mezz1. This can be one of the following:</p> <ul style="list-style-type: none"> • Disabled—The expansion slot is not available. • Enabled—The expansion slot is available. • UEFI Only—The expansion slot is available for UEFI only. • Legacy Only—The expansion slot is available for legacy only. • Platform Default—The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.
IOE Slot2 Option ROM	<p>Whether option ROM is enabled on the IOE slot 2. This can be one of the following:</p> <ul style="list-style-type: none"> • Disabled—The expansion slot is not available. • Enabled—The expansion slot is available. • UEFI Only—The expansion slot is available for UEFI only. • Legacy Only—The expansion slot is available for legacy only. • Platform Default—The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.
IO ENVME1 Option ROM	<p>Whether option ROM is enabled on the IOE NVMe1. This can be one of the following:</p> <ul style="list-style-type: none"> • Disabled—The expansion slot is not available. • Enabled—The expansion slot is available. • UEFI Only—The expansion slot is available for UEFI only. • Legacy Only—The expansion slot is available for legacy only. • Platform Default—The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.
IO ENVME2 Option ROM	<p>Whether option ROM is enabled on the IOE NVMe2. This can be one of the following:</p> <ul style="list-style-type: none"> • Disabled—The expansion slot is not available. • Enabled—The expansion slot is available. • UEFI Only—The expansion slot is available for UEFI only. • Legacy Only—The expansion slot is available for legacy only. • Platform Default—The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.

Name	Description
SBNVME1 Option ROM	<p>Whether the server can use Option ROM present in SBNVMe1 controller. This can be one of the following:</p> <ul style="list-style-type: none"> • Disabled—The expansion slot is not available. • Enabled—The expansion slot is available. • UEFI Only—The expansion slot is available for UEFI only. • Legacy Only—The expansion slot is available for legacy only. • Platform Default—The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.
PCIe Slot MRAID-<i>n</i> OptionROM	<p>Whether Option ROM is available on the MRAID port. This can be one of the following:</p> <ul style="list-style-type: none"> • Disabled—The expansion slot is not available. • Enabled—The expansion slot is available. • UEFI Only—The expansion slot is available for UEFI only. • Legacy Only—The expansion slot is available for legacy only. • Platform Default—The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.
PCIe Slot RAID OptionROM	<p>Whether Option ROM is available on the RAID port. This can be one of the following:</p> <ul style="list-style-type: none"> • Disabled—The expansion slot is not available. • Enabled—The expansion slot is available. • UEFI Only—The expansion slot is available for UEFI only. • Legacy Only—The expansion slot is available for legacy only. • Platform Default—The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.
PCIe Slot Rear Nvmen OptionRom	<p>Whether Option ROM is available on the Rear NVMe<i>n</i> port. This can be one of the following:</p> <ul style="list-style-type: none"> • Disabled—The expansion slot is not available. • Enabled—The expansion slot is available. • UEFI Only—The expansion slot is available for UEFI only. • Legacy Only—The expansion slot is available for legacy only. • Platform Default—The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.

Name	Description
Rear NVME <i>n</i> Link Speed	<p>This option allows you to restrict the maximum speed of an NVME card installed in the rear PCIe slot <i>n</i>. This can be one of the following:</p> <ul style="list-style-type: none"> • Gen 1—2.5GT/s (gigatransfers per second) is the maximum speed allowed. • Gen 2—5GT/s is the maximum speed allowed. • Gen 3—8GT/s is the maximum speed allowed. • Enabled—The maximum speed is restricted. • Auto—The maximum speed is set automatically. • Disabled—The maximum speed is not restricted. • Platform Default—The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.
Front NVME <i>n</i> Link Speed	<p>This option allows you to restrict the maximum speed of an NVME card installed in the front PCIe slot. This can be one of the following:</p> <ul style="list-style-type: none"> • Gen 1—2.5GT/s (gigatransfers per second) is the maximum speed allowed. • Gen 2—5GT/s is the maximum speed allowed. • Gen 3—8GT/s is the maximum speed allowed. • Auto—The maximum speed is set automatically. This is the default option. • Enabled—The maximum speed is restricted. • Disabled—The maximum speed is not restricted. • Platform Default—The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.
HBA Link Speed	<p>This option allows you to restrict the maximum speed of an HBA card. This can be one of the following:</p> <ul style="list-style-type: none"> • Gen 1—2.5GT/s (gigatransfers per second) is the maximum speed allowed. • Gen 2—5GT/s is the maximum speed allowed. • Gen 3—8GT/s is the maximum speed allowed. • Auto—The maximum speed is set automatically. • Disabled—The maximum speed is not restricted. • Platform Default—The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.

Name	Description
MLOM Link Speed	<p>This option allows you to restrict the maximum speed of an MLOM adapter. This can be one of the following:</p> <ul style="list-style-type: none"> • Gen 1—2.5GT/s (gigatransfers per second) is the maximum speed allowed. • Gen 2—5GT/s is the maximum speed allowed. • Gen 3—8GT/s is the maximum speed allowed. • Auto—The maximum speed is set automatically. This is the default option. • Disabled—The maximum speed is not restricted. • Enabled—The maximum speed is restricted. • Platform Default—The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.
MRAID Link Speed	<p>This option allows you to restrict the maximum speed of MRAID. This can be one of the following:</p> <ul style="list-style-type: none"> • Gen 1—2.5GT/s (gigatransfers per second) is the maximum speed allowed. • Gen 2—5GT/s is the maximum speed allowed. • Gen 3—8GT/s is the maximum speed allowed. • Auto—The maximum speed is set automatically. • Enabled—The maximum speed is not restricted. • Disabled—The maximum speed is not restricted. • Platform Default—The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.
RAID-<i>n</i> Link Speed	<p>This option allows you to restrict the maximum speed of RAID. This can be one of the following:</p> <ul style="list-style-type: none"> • Gen 1—2.5GT/s (gigatransfers per second) is the maximum speed allowed. • Gen 2—5GT/s is the maximum speed allowed. • Gen 3—8GT/s is the maximum speed allowed. • Auto—The maximum speed is set automatically. • Disabled—The maximum speed is not restricted. • Platform Default—The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.

Name	Description
All Onboard LOM	<p>Whether all onboard LOM ports are enabled or disabled. This can be one of the following:</p> <ul style="list-style-type: none"> • Enabled—All onboard LOM are enabled. • Disabled—All onboard LOM are disabled. • Platform Default—The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.
LOM Port 1 OptionRom	<p>Whether Option ROM is available on the LOM port 1. This can be one of the following:</p> <ul style="list-style-type: none"> • Disabled—The expansion slot is not available. • Enabled—The expansion slot is available. • UEFI Only—The expansion slot is available for UEFI only. • Legacy Only—The expansion slot is available for legacy only. • Platform Default—The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.
LOM Port 2 OptionRom	<p>Whether Option ROM is available on the LOM port 2. This can be one of the following:</p> <ul style="list-style-type: none"> • Disabled—The expansion slot is not available. • Enabled—The expansion slot is available. • UEFI Only—The expansion slot is available for UEFI only. • Legacy Only—The expansion slot is available for legacy only. • Platform Default—The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.
Slot <i>n</i> State	<p>The state of the adapter card installed in PCIe slot <i>n</i>. This can be one of the following:</p> <ul style="list-style-type: none"> • Disabled—The expansion slot is not available. • Enabled—The expansion slot is available. • UEFI Only—The expansion slot is available for UEFI only. • Legacy Only—The expansion slot is available for legacy only. • Platform Default—The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.

Name	Description
SBNVMe1 OptionROM	<p>Whether the server can use Option ROM present in SBNVMe1 controller. This can be one of the following:</p> <ul style="list-style-type: none"> • Disabled—The expansion slot is not available. • Enabled—The expansion slot is available. • Platform Default—The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.
SBNVMe2 OptionROM	<p>Whether the server can use Option ROM present in SBNVMe2 controller. This can be one of the following:</p> <ul style="list-style-type: none"> • Disabled—The expansion slot is not available. • Enabled—The expansion slot is available. • Platform Default—The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.
SIOCNVMe1 OptionROM	<p>Whether the server can use Option ROM present in SIOCNVMe1 controller. This can be one of the following:</p> <ul style="list-style-type: none"> • Disabled—The expansion slot is not available. • Enabled—The expansion slot is available. • Platform Default—The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.
SIOCNVMe2 OptionROM	<p>Whether the server can use Option ROM present in SIOCNVMe2 controller. This can be one of the following:</p> <ul style="list-style-type: none"> • Disabled—The expansion slot is not available. • Enabled—The expansion slot is available. • Platform Default—The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.
SBLom1 OptionROM	<p>Whether the server can use Option ROM present in the SBLom1 controller. This can be one of the following:</p> <ul style="list-style-type: none"> • Disabled—The expansion slot is not available. • Enabled—The expansion slot is available. • Platform Default—The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.

Name	Description
SBNVMen Link Speed	<p>Link speed for SBNVMe slot n. This can be one of the following:</p> <ul style="list-style-type: none"> • Gen 1—2.5GT/s (gigatransfers per second) is the maximum speed allowed. • Gen 2—5GT/s is the maximum speed allowed. • Gen 3—8GT/s is the maximum speed allowed. • Enabled—The maximum speed is restricted. • Disabled—The maximum speed is not restricted. • Platform Default—The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.
SIOCNVMen Link Speed	<p>Link speed for SIOCNVMe slot n. This can be one of the following:</p> <ul style="list-style-type: none"> • Gen 1—2.5GT/s (gigatransfers per second) is the maximum speed allowed. • Gen 2—5GT/s is the maximum speed allowed. • Gen 3—8GT/s is the maximum speed allowed. • Enabled—The maximum speed is restricted. • Disabled—The maximum speed is not restricted. • Platform Default—The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.
SIOCn Link Speed	<p>Link speed for SIOC slot n. This can be one of the following:</p> <ul style="list-style-type: none"> • Gen 1—2.5GT/s (gigatransfers per second) is the maximum speed allowed. • Gen 2—5GT/s is the maximum speed allowed. • Gen 3—8GT/s is the maximum speed allowed. • Enabled—The maximum speed is restricted. • Disabled—The maximum speed is not restricted. • Platform Default—The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.

Name	Description
SBMezzn Link Speed	<p>Link speed for SBMezz slot n. This can be one of the following:</p> <ul style="list-style-type: none"> • Gen 1—2.5GT/s (gigatransfers per second) is the maximum speed allowed. • Gen 2—5GT/s is the maximum speed allowed. • Gen 3—8GT/s is the maximum speed allowed. • Enabled—The maximum speed is restricted. • Disabled—The maximum speed is not restricted. • Platform Default—The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.
IOESlotn Link Speed	<p>Link speed for IOE slot n. This can be one of the following:</p> <ul style="list-style-type: none"> • Gen 1—2.5GT/s (gigatransfers per second) is the maximum speed allowed. • Gen 2—5GT/s is the maximum speed allowed. • Gen 3—8GT/s is the maximum speed allowed. • Enabled—The maximum speed is restricted. • Disabled—The maximum speed is not restricted. • Platform Default—The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.
IOEMezzn Link Speed	<p>Link speed for IOEMezz slot n. This can be one of the following:</p> <ul style="list-style-type: none"> • Gen 1—2.5GT/s (gigatransfers per second) is the maximum speed allowed. • Gen 2—5GT/s is the maximum speed allowed. • Gen 3—8GT/s is the maximum speed allowed. • Enabled—The maximum speed is restricted. • Disabled—The maximum speed is not restricted. • Platform Default—The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.

Name	Description
IOENVMen Link Speed	<p>Link speed for IOENVMe slot n. This can be one of the following:</p> <ul style="list-style-type: none"> • Gen 1—2.5GT/s (gigatransfers per second) is the maximum speed allowed. • Gen 2—5GT/s is the maximum speed allowed. • Gen 3—8GT/s is the maximum speed allowed. • Enabled—The maximum speed is restricted. • Disabled—The maximum speed is not restricted. • Platform Default—The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.
CDN Support for LOMs	<p>Whether the Ethernet Networking Identifier naming convention is according to Consistent Device Naming (CDN) or the traditional way of naming conventions. This can be one of the following:</p> <ul style="list-style-type: none"> • Enabled—OS Ethernet Network Identifier is named in a consistent device naming (CDN) convention according to the physical LAN on Motherboard (LOM) port numbering; LOM Port 0, LOM Port 1 and so on. • Disabled—OS Ethernet Networking Identifier is named in a default convention as ETH0, ETH1 and so on. By default, CDN option is disabled. • Platform Default—The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.
VMD Enable	<p>Whether NVMe SSDs that are connected to the PCIe bus can be hot swapped. It also standardizes the LED status light on these drives. LED status lights can be optionally programmed to display specific Failure indicator patterns.</p> <p>This can be one of the following:</p> <ul style="list-style-type: none"> • Enabled—Hot swap of NVMe SSDs that are connected to the PCIe bus is allowed. • Disabled—Hot swap of NVMe SSDs that are connected to the PCIe bus is not allowed. • Platform Default—The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.

Name	Description
ACS Control SLOT-<i>n</i> <i>n</i> = 11 to 14	<p>Access Control Services (ACS) allow the processor to enable or disable peer-to-peer communication between multiple devices for Control Slot <i>n</i>. This can be one of the following:</p> <ul style="list-style-type: none"> • Enabled— Enables peer-to-peer communication between multiple devices for Control Slot <i>n</i>. • Disabled— Disables peer-to-peer communication between multiple devices for Control Slot <i>n</i>. • Platform Default—The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.
PCIe Slot GPU<i>n</i> OptionROM Only for Cisco UCS C480 M5 ML Server	<p>Whether the Option ROM is enabled on GPU slot <i>n</i>. <i>n</i> is the slot number, which can be numbered 1 through 8. This can be one of the following:</p> <ul style="list-style-type: none"> • Disabled—The expansion slot is not available. • Enabled—The expansion slot is available. • Platform Default—The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.
ACS Control GPU-<i>n</i> <i>n</i> = 1 to 8	<p>Access Control Services (ACS) allow the processor to enable or disable peer-to-peer communication between multiple devices for GPUs. This can be one of the following:</p> <ul style="list-style-type: none"> • Disabled— Enables peer-to-peer communication between multiple devices for GPUs. • Enabled— Disables peer-to-peer communication between multiple devices for GPUs. • Platform Default—The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.
PCIe PLL SSC	<p>Reduces EMI interference by down-spreading the clock by 0.5%. Disable this feature to centralize the clock without spreading. For all Cisco UCS M5 servers, this option is Disabled by default.</p> <ul style="list-style-type: none"> • Disabled— Clock is centralized without spreading. • Auto— EMI interference is auto adjusted. • ZeroPointFive— EMI interference us reduced by down-spreading the clock by 0.5%. • Platform Default— The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.

Name	Description
NVME-<i>n</i> OptionROM	Whether Option ROM is available on the NVME port. This can be one of the following: <ul style="list-style-type: none"> • Disabled—The expansion slot is not available. • Enabled—The expansion slot is available. This is the default option.
PCIe Slot OCP Link Speed	This option allows you to restrict the maximum speed of OCP. This can be one of the following: <ul style="list-style-type: none"> • Gen 1—2.5GT/s (gigatransfers per second) is the maximum speed allowed. • Gen 2—5GT/s is the maximum speed allowed. • Gen 3—8GT/s is the maximum speed allowed. • Auto—The maximum speed is set automatically. This is the default option. • Disabled—The maximum speed is not restricted. • Platform Default—The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.
RAID<i>n</i> OptionROM	Whether Option ROM is available on the RAID port. This can be one of the following: <ul style="list-style-type: none"> • Disabled—The expansion slot is not available. • Enabled—The expansion slot is available. This is the default option.
IOENVMe<i>n</i> OptionROM	Whether Option ROM is available on the IOENVMe port. This can be one of the following: <ul style="list-style-type: none"> • Disabled—The expansion slot is not available. • Enabled—The expansion slot is available. This is the default option.
GPU<i>n</i> OptionRom	Whether Option ROM is available on the GPU port. This can be one of the following: <ul style="list-style-type: none"> • Disabled—The expansion slot is not available. • Enabled—The expansion slot is available. This is the default option.

Name	Description
RAID Link Speed	<p>This option allows you to restrict the maximum speed of RAID. This can be one of the following:</p> <ul style="list-style-type: none"> • Gen 1—2.5GT/s (gigatransfers per second) is the maximum speed allowed. • Gen 2—5GT/s is the maximum speed allowed. • Gen 3—8GT/s is the maximum speed allowed. • Auto—The maximum speed is set automatically. This is the default option. • Enabled—The maximum speed is not restricted. • Disabled—The maximum speed is not restricted. • Platform Default—The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.
External SSC enable	<p>This option allows you to Enable/Disable the Clock Spread Spectrum of the external clock generators.</p> <ul style="list-style-type: none"> • Disabled— Clock Spread Spectrum support is not available. • Enabled— Clock Spread Spectrum support is always available. • Platform Default — The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.

Graphics Configuration BIOS Settings

The following tables list the graphics configuration BIOS settings that you can configure through a BIOS policy or the default BIOS settings:

Name	Description
Integrated Graphics	<p>Enables integrated graphics. This can be one of the following:</p> <ul style="list-style-type: none"> • Platform Default—The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.
Integrated Graphics Aperture Size	<p>Allows you to set the size of mapped memory for the integrated graphics controller. This can be one of the following:</p> <ul style="list-style-type: none"> • Platform Default—The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.
Onboard Graphics	<p>Enables onboard graphics (KVM). This can be one of the following:</p> <ul style="list-style-type: none"> • Platform Default—The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.

Boot Options BIOS Settings

The following table lists the boot options BIOS settings that you can configure through a BIOS policy or the default BIOS settings:



Note BIOS parameter virtualization capability in Cisco UCS Manager maps a unified set of BIOS settings in a service profile to the actual BIOS supporting parameters. However, not all BIOS setting items are applicable to every server model/platform. When you create a custom BIOS policy and have the **Boot Option Retry** selected, and when there is no bootable option available, the reboot fails on the Cisco UCS B420 M3 or Cisco UCS B420 M4 servers and Cisco UCS Manager displays this message : *Reboot and Select proper Boot device or Insert Boot Media in selected Boot device and press a key*. You must manually set a boot option after the boot path is corrected, in order to enable the servers to reboot after a power outage. For more information about BIOS default server policies and the BIOS options and their default settings, see [BIOS Policy, on page 263](#) and [Server BIOS Settings, on page 187](#).

Server Management BIOS Settings

The following tables list the server management BIOS settings that you can configure through a BIOS policy or the default BIOS settings:

General Settings

Name	Description
Assert NMI on SERR	<p>Whether the BIOS generates a non-maskable interrupt (NMI) and logs an error when a system error (SERR) occurs. This can be one of the following:</p> <ul style="list-style-type: none"> • Disabled—The BIOS does not generate an NMI or log an error when a SERR occurs. • Enabled—The BIOS generates an NMI and logs an error when a SERR occurs. You must enable this setting if you want to enable Assert NMI on PERR. • Platform Default—The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.
Assert NMI on PERR	<p>Whether the BIOS generates a non-maskable interrupt (NMI) and logs an error when a processor bus parity error (PERR) occurs. This can be one of the following:</p> <ul style="list-style-type: none"> • Disabled—The BIOS does not generate an NMI or log an error when a PERR occurs. • Enabled—The BIOS generates an NMI and logs an error when a PERR occurs. You must enable Assert NMI on SERR to use this setting. • Platform Default—The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.

Name	Description
OS Boot Watchdog Timer	<p>Whether the BIOS programs the watchdog timer with a predefined timeout value. If the operating system does not complete booting before the timer expires, the CIMC resets the system and an error is logged. This can be one of the following:</p> <ul style="list-style-type: none"> • Disabled—The watchdog timer is not used to track how long the server takes to boot. • Enabled—The watchdog timer tracks how long the server takes to boot. If the server does not boot within the predefined length of time, the CIMC resets the system and logs an error. • Platform Default—The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor. <p>This feature requires either operating system support or Intel Management software.</p>
OS Boot Watchdog Timer Policy	<p>What action the system takes if the watchdog timer expires. This can be one of the following:</p> <ul style="list-style-type: none"> • Power Off—The server is powered off if the watchdog timer expires during OS boot. • Reset—The server is reset if the watchdog timer expires during OS boot. • Platform Default—The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor. <p>This option is only available if you enable the OS Boot Watchdog Timer.</p>

Name	Description
OS Boot Watchdog Timer Timeout	<p>What timeout value the BIOS uses to configure the watchdog timer. This can be one of the following:</p> <ul style="list-style-type: none"> • 5-minutes—The watchdog timer expires 5 minutes after the OS begins to boot. • 10-minutes—The watchdog timer expires 10 minutes after the OS begins to boot. • 15-minutes—The watchdog timer expires 15 minutes after the OS begins to boot. • 20-minutes—The watchdog timer expires 20 minutes after the OS begins to boot. • Platform Default—The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor. <p>This option is only available if you enable the OS Boot Watchdog Timer.</p>
FRB-2 Timer	<p>Whether the FRB-2 timer is used to recover the system if it hangs during POST. This can be one of the following:</p> <ul style="list-style-type: none"> • Disabled—The FRB-2 timer is not used. • Enabled—The FRB-2 timer is started during POST and used to recover the system if necessary. • Platform Default—The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.

Console Redirection Settings

Name	Description
Console redirection	<p>Allows a serial port to be used for console redirection during POST and BIOS booting. After the BIOS has booted and the operating system is responsible for the server, console redirection is irrelevant and has no effect. This can be one of the following:</p> <ul style="list-style-type: none"> • Disabled—No console redirection occurs during POST. • serial-port-a or COM 0 or Enabled—Enables serial port A for console redirection during POST. This option is valid for blade servers and rack-mount servers. • serial-port-b or COM 1—Enables serial port B for console redirection and allows it to perform server management tasks. This option is only valid for rack-mount servers. • Platform Default—The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor. <p>Note If you enable this option, you also disable the display of the Quiet Boot logo screen during POST.</p>
Flow Control	<p>Whether a handshake protocol is used for flow control. Request to Send / Clear to Send (RTS/CTS) helps to reduce frame collisions that can be introduced by a hidden terminal problem. This can be one of the following:</p> <ul style="list-style-type: none"> • None—No flow control is used. • RTS-CTS—RTS/CTS is used for flow control. • Platform Default—The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor. <p>Note This setting must match the setting on the remote terminal application.</p>

Name	Description
Baud rate	<p>What Baud rate is used for the serial port transmission speed. If you disable Console Redirection, this option is not available. This can be one of the following:</p> <ul style="list-style-type: none"> • 9.6k—A 9600 Baud rate is used. • 19.2k—A 19200 Baud rate is used. • 38.4k—A 38400 Baud rate is used. • 57.6k—A 57600 Baud rate is used. • 115.2k—A 115200 Baud rate is used. This is the default option. • Platform Default—The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor. <p>Note This setting must match the setting on the remote terminal application.</p>
Terminal type	<p>What type of character formatting is used for console redirection. This can be one of the following:</p> <ul style="list-style-type: none"> • PC-ANSI—The PC-ANSI terminal font is used. • VT100—A supported vt100 video terminal and its character set are used. • VT100-PLUS—A supported vt100-plus video terminal and its character set are used. • VT-UTF8—A video terminal with the UTF-8 character set is used. • Platform Default—The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor. <p>Note This setting must match the setting on the remote terminal application.</p>

Name	Description
Legacy OS redirection	<p>Whether redirection from a legacy operating system, such as DOS, is enabled on the serial port. This can be one of the following:</p> <ul style="list-style-type: none"> • Disabled—The serial port enabled for console redirection is hidden from the legacy operating system. • Enabled— The serial port enabled for console redirection is visible to the legacy operating system. • Platform Default—The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.
<p>Putty KeyPad set console-redir-config putty-function-keypad</p>	<p>Allows you to change the action of the PuTTY function keys and the top row of the numeric keypad. This can be one of the following:</p> <ul style="list-style-type: none"> • VT100—The function keys generate ESC OP through ESC O[. • LINUX—Mimics the Linux virtual console. Function keys F6 to F12 behave like the default mode, but F1 to F5 generate ESC [[A through ESC [[E. • XTERMR6—Function keys F5 to F12 behave like the default mode. Function keys F1 to F4 generate ESC OP through ESC OS, which are the sequences produced by the top row of the keypad on Digital terminals. • SCO—The function keys F1 to F12 generate ESC [M through ESC [X. The function and shift keys generate ESC [Y through ESC [j. The control and function keys generate ESC [k through ESC [v. The shift, control and function keys generate ESC [w through ESC [{. • ESCN—The default mode. The function keys match the general behavior of Digital terminals. The function keys generate sequences such as ESC [11~ and ESC [12~. • VT400—The function keys behave like the default mode. The top row of the numeric keypad generates ESC OP through ESC OS. • Platform Default—The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.

Name	Description
Out of Band Management	<p>Used for Windows Special Administration Control (SAC). This option allows you to configure the COM port 0 that can be used for Windows Emergency Management services. ACPI SPCR table is reported based on this setup option. This can be one of the following:</p> <ul style="list-style-type: none"> • Disabled—Configures the COM port 0 as a general purpose port for use with the Windows Operating System. • Enabled—Configures the COM port 0 as a remote management port for Windows Emergency Management services. • Platform Default—The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.
Redirection After BIOS POST	<p>Whether BIOS console redirection should be active after BIOS POST is complete and control given to the OS bootloader. This can be one of the following:</p> <ul style="list-style-type: none"> • Always Enable—BIOS Legacy console redirection is active during the OS boot and run time. • Bootloader—BIOS Legacy console redirection is disabled before giving control to the OS boot loader. • Platform Default—The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.

Intel Speed Select

Intel Speed Select lets you optimize CPU performance by selecting one of three operating profiles, based on number of logical processor cores, frequency, and TDP thread setting. Speed Select is configurable with a BIOS policy based on the SKU. At boot time, the BIOS discovers the supported CPU and configures it to one of three specified profiles.

Set profiles by configuring the Speed Select BIOS policy tokens. You create policies based on performance needs. Tuning can be based on workload or VM needs, need for improved server utilization in the data center, or improved per-core performance SLAs. Low core profiles let the CPU run at maximum frequency.

The Speed Select profiles are:

- **Base**—High Core Configuration. Suggested for general purpose VM instances.
- **Config 1**—Medium Core Configuration. Suggested for high performance VM instances or Storage SKUs.
- **Config 2**—Low Core Configuration. Suggested for very high frequency Bare Metal servers or VM instances.

The following tables show sample profile values based on the Intel® Xeon® Gold CPU 6240Y. Base is the default profile.

Table 6: Base Config

SKU	Cores	Base Frequency (GHz)	All Core Turbo	LLC/Core (Mb)
8260Y	24	2.4	3.1	1.49
6240Y	18	2.6	3.3	1.375
4214Y	12	2.2	2.7	1.375

SKU	TDP (W)	Tj Max (°C)
8260Y	160	104
6240Y	150	88
4214Y	85	88

Table 7: Config 1

SKU	Cores	Base Frequency (GHz)	All Core Turbo	LLC/Core (Mb)
8260Y	20	2.5	3.3	1.79
6240Y	14	2.8	3.4	1.77
4214Y	10	2.3	2.7	1.65

SKU	TDP (W)	Tj Max (°C)
8260Y	155	TControl Spec
6240Y	150	TControl Spec
4214Y	85	TControl Spec

Table 8: Config 2

SKU	Cores	Base Frequency (GHz)	All Core Turbo	LLC/Core (Mb)
8260Y	16	2.7	3.6	2.23
6240Y	8	3.1	3.6	3.09
4214Y	8	2.4	2.9	2.06

SKU	TDP (W)	Tj Max (°C)
8260Y	150	TControl Spec

SKU	TDP (W)	Tj Max (°C)
6240Y	150	TControl Spec
4214Y	85	TControl Spec



Note While changing the profiles can make the CPU run at a different frequency, the CIMC processor inventory page will show information for the processor Base frequency, core, and threads. It will not display running frequencies.

BIOS Policy

The BIOS policy is a policy that automates the configuration of BIOS settings for a server or group of servers. You can create global BIOS policies available to all servers in the root organization, or you can create BIOS policies in sub-organizations that are only available to that hierarchy.

To use a BIOS policy, do the following:

1. Create the BIOS policy in Cisco UCS Manager.
2. Assign the BIOS policy to one or more service profiles.
3. Associate the service profile with a server.

During service profile association, Cisco UCS Manager modifies the BIOS settings on the server to match the configuration in the BIOS policy. If you do not create and assign a BIOS policy to a service profile, the server uses the default BIOS settings for that server platform.

Default BIOS Settings

Cisco UCS Manager includes a set of default BIOS settings for each type of server supported by Cisco UCS. The default BIOS settings are available only in the root organization and are global. Only one set of default BIOS settings can exist for each server platform supported by Cisco UCS. You can modify the default BIOS settings, but you cannot create an additional set of default BIOS settings.

Each set of default BIOS settings are designed for a particular type of supported server and are applied to all servers of that specific type which do not have a BIOS policy included in their service profiles.

Unless a Cisco UCS implementation has specific needs that are not met by the server-specific settings, we recommend that you use the default BIOS settings that are designed for each type of server in the Cisco UCS domain.

Cisco UCS Manager applies these server platform-specific BIOS settings as follows:

- The service profile associated with a server does not include a BIOS policy.
- The BIOS policy is configured with the platform-default option for a specific setting.

You can modify the default BIOS settings provided by Cisco UCS Manager. However, any changes to the default BIOS settings apply to all servers of that particular type or platform. If you want to modify the BIOS settings for only certain servers, we recommend that you use a BIOS policy.

The BIOS tokens for M5 servers and later are read-only and cannot be modified. For a complete and up to date list of BIOS tokens, defaults, and values, refer [Cisco UCS M5 Server BIOS Tokens](#) .

Creating a BIOS Policy



Note Cisco UCS Manager pushes BIOS configuration changes through a BIOS policy or default BIOS settings to the Cisco Integrated Management Controller (CIMC) buffer. These changes remain in the buffer and do not take effect until the server is rebooted.

We recommend that you verify the support for BIOS settings in the server that you want to configure. Some settings, such as Mirroring Mode for RAS Memory, are not supported by all Cisco UCS servers.

Procedure

- Step 1** In the **Navigation** pane, click **Servers**.
- Step 2** Expand **Servers > Policies**.
- Step 3** Expand the node for the organization where you want to create the policy.
If the system does not include multi tenancy, expand the **root** node.
- Step 4** Right-click **BIOS Policies** and select **Create BIOS Policy**.
- Step 5** On the **Main** page of the **Create BIOS Policy** wizard, enter a name for the BIOS policy in the **Name** field.
This name can be between 1 and 16 alphanumeric characters. You cannot use spaces or any special characters other than - (hyphen), _ (underscore), : (colon), and . (period), and you cannot change this name after the object is saved.
- Step 6** In the **Create BIOS Policy** wizard, do the following to configure the BIOS settings:
 - a) If you want to change a BIOS setting, click the desired radio button or make the appropriate choice from the drop-down list.
For descriptions and information about the options for each BIOS setting, see the following topics:
 - **Main** page: [Main BIOS Settings, on page 188](#)
 - **Advanced** page: [Main BIOS Settings, on page 188](#)
 - **Processor** page: [Processor BIOS Settings, on page 190](#)
 - **IO BIOS for Intel** page: [I/O BIOS Settings for Intel, on page 219](#)
 - **IO BIOS for AMD** page: [I/O BIOS Settings for AMD, on page 221](#)
 - **RAS Memory** page: [RAS Memory BIOS Settings, on page 222](#)
 - **Serial Port** page: [Serial Port BIOS Settings, on page 231](#)
 - **USB** page: [USB BIOS Settings, on page 232](#)
 - **PCI Configuration** page: [PCI Configuration BIOS Settings, on page 235](#)

- **QPI** page: [QPI BIOS Settings, on page 238](#)
- **LOM and PCIe Slots** subtab: [LOM and PCIe Slots BIOS Settings, on page 239](#)
- **Trusted Platform** subtab: [Trusted Platform BIOS Settings, on page 239](#)
- **Graphics Configuration** subtab: [Graphics Configuration BIOS Settings, on page 254](#)
- **Boot Options** page: [Boot Options BIOS Settings, on page 255](#)
- **Server Management** page: [Server Management BIOS Settings, on page 255](#)

b) Click **Next** after each page.

Step 7 After you configure all of the BIOS settings for the policy, click **Finish**.

Modifying the BIOS Defaults

We recommend that you verify the support for BIOS settings in the server that you want to configure. Some settings, such as Mirroring Mode for RAS Memory, are not supported by all Cisco UCS servers.

Unless a Cisco UCS implementation has specific needs that are not met by the server-specific settings, we recommend that you use the default BIOS settings that are designed for each type of server in the Cisco UCS domain.

Procedure

- Step 1** In the **Navigation** pane, click **Servers**.
- Step 2** Expand **Servers > Policies**.
- Step 3** Expand the node for the organization where you want to create the policy.
If the system does not include multi tenancy, expand the **root** node.
- Step 4** Expand **BIOS Defaults** and select the server model number or desired policy for which you want to modify the default BIOS settings.
- Step 5** In the **Work** pane, click the appropriate tab and then click the desired radio button or make a choice from the drop-down list to modify the default BIOS settings:
For descriptions and information about the options for each BIOS setting, see the following topics. Not all BIOS settings are available for each type of server.
- **Main** tab: [Main BIOS Settings, on page 188](#)
 - **Advanced** tab:
 - **Processor** subtab: [Processor BIOS Settings, on page 190](#)
 - **IO BIOS for Intel** subtab: [I/O BIOS Settings for Intel, on page 219](#)
 - **IO BIOS for AMD** page: [I/O BIOS Settings for AMD, on page 221](#)
 - **RAS Memory** subtab: [RAS Memory BIOS Settings, on page 222](#)

- **Serial Port** subtab: [Serial Port BIOS Settings, on page 231](#)
- **USB** subtab: [USB BIOS Settings, on page 232](#)
- **PCI Configuration** subtab: [PCI Configuration BIOS Settings, on page 235](#)
- **QPI** subtab: [QPI BIOS Settings, on page 238](#)
- **LOM and PCIe Slots** subtab: [LOM and PCIe Slots BIOS Settings, on page 239](#)
- **Trusted Platform** subtab: [Trusted Platform BIOS Settings, on page 239](#)
- **Graphics Configuration** subtab: [Graphics Configuration BIOS Settings, on page 254](#)
- **Boot Options** tab: [Boot Options BIOS Settings, on page 255](#)
- **Server Management** tab: [Server Management BIOS Settings, on page 255](#)

Step 6 Click **Save Changes**.

Viewing the Actual BIOS Settings for a Server

Follow this procedure to see the actual BIOS settings on a server.

Procedure

- Step 1** In the **Navigation** pane, click **Equipment**.
- Step 2** Expand **Equipment** > **Chassis** > *Chassis Number* > **Servers**.
- Step 3** Choose the server for which you want to view the actual BIOS settings.
- Step 4** On the **Work** pane, click the **Inventory** tab.
- Step 5** Click the **Motherboard** subtab.
- Step 6** In the **BIOS Settings** area, click the **Expand** icon to the right of the heading to open that area.

Each tab in the **BIOS Settings** area displays the settings for that server platform. Some of the tabs contain subtabs with additional information.

Memory RAS Features

The Intel® Xeon® processor supports additional RAS memory features via the BIOS. These features expand on the capabilities of the processor to increase the performance and reliability of memory DIMMs.

Post-Package Repair (PPR)

Post Package Repair (PPR) allows you to use spare rows in the DRAM bank within the DDR4 DRAM to replace faulty rows detected during system boot time. Cisco UCS M5 platforms apply hard PPR. In hard PPR,

the repair is permanent. The remapping of a faulty row to a spare row cannot be reverted. The remapping persists even after removal of power. If a PPR event occurs, the platform firmware generates a customer visible fault to schedule for system reboot for the repair to take effect.

The number of spare rows in the DRAM bank varies based on DIMM manufactures and models. The spare rows that are available after executing the PPR event are not visible to the platform firmware. Thereby, when all the spare rows that are available in the platform firmware visibility are utilized, the repair will not take effect and the memory errors may reoccur on the same DIMM.

Enabling Post Package Repair

When enabled, the repair process is irrevocable.

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** In the Policies section, right-click the BIOS Policy section and select **Create BIOS Policy** from the popup. In the BIOS Policy form, enter a name and optional description. Click **OK** to create the policy.
 - Step 5** Go to **Policies > Root > BIOS Policies** and select the new policy.
 - Step 6** In the main work pane, select the **Advanced** tab, then select the **RAS Memory** tab.
 - Step 7** To enable automatic repair of faulty cell areas detected during system boot, in **Select PPR Type Configuration** select **Hard PPR**.
 - Step 8** Click **Save Changes**.
-

Limiting Presented Memory

The amount of memory presented to the user can be limited in the BIOS. When the system is fully populated with high capacity DIMM modules, it may be desirable to reduce the amount of memory actually presented for use.

The memory limit will be applied evenly across all installed and available DIMMs to the extent possible. The minimum amount of presented memory you can specify is 1 GB. The following parameters apply:

0 = No limit. Full amount of installed memory is presented.

1 to $2^{31} - 1$ = Size of presented memory in gigabytes (GB)

Actual presented memory size will always be equal or less than specified memory size.

Limiting Memory Size

Actual presented memory size will always be equal or less than specified memory size.

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** In the Policies section, right-click the BIOS Policy section and select **Create BIOS Policy** from the popup. In the BIOS Policy form, enter a name and optional description. Click **OK** to create the policy.
 - Step 5** Go to **Policies > Root > BIOS Policies** and select the new policy.
 - Step 6** In the main work pane, select the **Advanced** tab, then select the **RAS Memory** tab.
 - Step 7** To limit the amount of presented memory to be mirrored, go to **Memory Size Limit in GB** and enter a value (in GB) for the desired amount of memory to be presented to the user.
 - Step 8** Click **Save Changes**.
-

Partial Memory Mirroring

Partial Memory Mirroring of DIMMs is an advanced RAS feature. Only Gold and Platinum SKU CPUs support this feature.

Partial DIMM Mirroring creates a mirrored copy of a specific region of memory cells, rather than keeping the complete mirror copy. Partial Memory Mirroring can be performed from either BIOS policy setup menu or from the Linux Operating System. Partial Mirroring creates a mirrored region in memory map with the attributes of a partial mirror copy. Up to 50% of the total memory capacity can be mirrored, using up to 4 partial mirrors.

For mirroring, at least two DDR channels must be populated in each IMC. Partial mirroring supports one DDR4 mirror region per IMC, with a maximum of four mirror regions.

In a two-way channel interleave, two channels are populated in each IMC. In a three-way channel interleave, three channels are populated in each IMC.

Partial mirroring is incompatible with rank sparing and ADDDC.

The following rules apply to partial mirroring:

- The DIMM population must be identical for the mirrored channels.
- The mirror pair must be in the same M2M, within an IMC DDR channel.
- DDR4 partial mirror regions within one iMC must be either two-way channel interleaves or three-way channel interleave. Two and three-way channel interleaves cannot be mixed. When the mirror region spans across iMCs, the channel interleaves must be the same.

Enabling Partial Memory Mirroring

The amount of partial DIMM memory mirroring can be configured either in percentage of available memory resources or in gigabytes..

Before you begin



Note Partial Memory Mirror Mode is mutually exclusive to standard Mirror Mode.

Partial Mirroring is incompatible with rank sparing and ADDDC. Verify that these are not selected.

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** In the Policies section, right-click the BIOS Policy section and select **Create BIOS Policy** from the popup. In the BIOS Policy form, enter a name and optional description. Click **OK** to create the policy.
- Step 5** Go to **Policies > Root > BIOS Policies** and select the new policy.
- Step 6** In the main work pane, select the **Advanced** tab, then select the **RAS Memory** tab.
- Step 7** Go to **Memory RAS Configuration** and select **Partial Mirror Mode 1LM** from the dropdown list.
- Step 8** To configure the partial mirror in percentage, go to **Partial Memory Mirror Mode** and select **Percentage** from the dropdown.
- Step 9** Go to **Partial Mirror percentage** and type a value between 0.01 and 50.00, representing the desired percentage of memory to be mirrored.
- Step 10** To configure the partial mirror in gigabytes, go to **Partial Memory Mirror Mode** and select **Value in GB** from the dropdown.
- Step 11** Enter a value between 1 and GB of memory displayed in the limits field for **Partial Mirror 1**.
- Step 12** If desired, enter additional values into **Partial Mirror 2**, **Partial Mirror 3**, and **Partial Mirror 4**. The total of values entered into these mirrors cannot exceed the total memory available.
- Step 13** Click **Save Changes**.

What to do next

Reboot the system.

Trusted Platform Module

Trusted Platform Module

The Trusted Platform Module (TPM) is a component that can securely store artifacts that are used to authenticate the server. These artifacts can include passwords, certificates, or encryption keys. A TPM can also be used to store platform measurements that help ensure that the platform remains trustworthy. Authentication (ensuring that the platform can prove that it is what it claims to be) and attestation (a process helping to prove that a

platform is trustworthy and has not been breached) are necessary steps to ensure safer computing in all environments. It is a requirement for the Intel Trusted Execution Technology (TXT) security feature, which must be enabled in the BIOS settings for a server equipped with a TPM. Cisco UCS M4 and higher blade and rack-mount servers include support for TPM. TPM is enabled by default on these servers.



Important

- If you upgrade Cisco UCS Manager to Release 2.2(4) and higher, TPM is enabled.
- When TPM is enabled and you downgrade Cisco UCS Manager from Release 2.2(4), TPM is disabled.

Intel Trusted Execution Technology

Intel Trusted Execution Technology (TXT) provides greater protection for information that is used and stored on the business server. A key aspect of that protection is the provision of an isolated execution environment and associated sections of memory where operations can be conducted on sensitive data, invisible to the rest of the system. Intel TXT provides for a sealed portion of storage where sensitive data such as encryption keys can be kept, helping to shield them from being compromised during an attack by malicious code. Cisco UCS M4 and higher blade and rack-mount servers include support for TXT. TXT is disabled by default on these servers.

TXT can be enabled only after TPM, Intel Virtualization technology (VT) and Intel Virtualization Technology for Directed I/O (VT-d) are enabled. When you only enable TXT, it also implicitly enables TPM, VT, and VT-d.

Configuring Trusted Platform

Procedure

- Step 1** In the **Navigation** pane, click **Servers**.
- Step 2** Expand **Servers > Policies**.
- Step 3** Expand the node for the organization where you want to configure TPM.
- Step 4** Expand **BIOS Policies** and select the BIOS policy for which you want to configure TPM.
- Step 5** In the **Work** pane, click the **Advanced** tab.
- Step 6** Click the **Trusted Platform** subtab.
- Step 7** To configure TPM, click one of the following:

Option	Description
disabled	Disables TPM
enabled	Enables TPM
Platform Default	Enables TPM

- Step 8** To configure TXT, click one of the following:

Option	Description
disabled	Disables TXT
enabled	Enables TXT
Platform Default	Disables TXT

Step 9 Click **Save Changes**.

Configuring Trusted Platform

Procedure

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

Step 2 Expand **Servers > Policies**.

Step 3 Expand the node for the organization where you want to configure TPM.

Step 4 Expand **BIOS Policies** and select the BIOS policy for which you want to configure TPM.

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

Step 6 Click the **Trusted Platform** subtab.

Step 7 To configure TPM, click one of the following:

Option	Description
disabled	Disables TPM
enable	Enables TPM
Platform Default	Enables TPM

Step 8 To configure TXT, click one of the following:

Option	Description
disabled	Disables TXT
enable	Enables TXT
Platform Default	Disables TXT

Step 9 Click **Save Changes**.

Viewing TPM Properties

Procedure

- Step 1** In the **Navigation** pane, click **Equipment**.
 - Step 2** Expand **Equipment** > **Chassis** > **Chassis Number** > **Cartridges** > **Cartridge Number** > **Servers**
 - Step 3** Choose the server for which you want to view the TPM settings.
 - Step 4** On the **Work** pane, click the **Inventory** tab.
 - Step 5** Click the **Motherboard** subtab.
-

Consistent Device Naming

When there is no mechanism for the Operating System to label Ethernet interfaces in a consistent manner, it becomes difficult to manage network connections with server configuration changes. Consistent Device Naming (CDN), introduced in Cisco UCS Manager Release 2.2(4), allows Ethernet interfaces to be named in a consistent manner. This makes Ethernet interface names more persistent when adapter or other configuration changes are made.

To configure CDN for a vNIC, do the following:

- Enable consistent device naming in the BIOS policy.
- Associate the BIOS policy with a service profile.
- Configure consistent naming for a vNIC.

Guidelines and Limitations for Consistent Device Naming (CDN)

- CDN is supported on the following Operating Systems:
 - Windows 2016 and later Windows releases
 - Windows Server 2019
 - Red Hat Enterprise Linux (RHEL) 6.x and 7.x
 - SLES 12 SP3, SLES 12 SP4, and SLES 15 (for 4.0(4a) and later)
 - ESXi 6.7
- Consistent device naming (CDN) is supported on all M3 and higher blade and rack-mount servers.
- BIOS and adapter firmware must be part of the Release 2.2(4) or higher bundle to support CDN.
- If the RHEL Operating System is installed on the server, CDN will appear when running the command "**biosdevname -d**" as "**sysfs label**". CDN will not change the kernel name.
- CDN is supported for vNIC template.

- Multiple vNICs within the same service profile cannot have the same CDN name.
- When a CDN name is not specified for a vNIC, the vNIC name is used as the CDN name.
- The CDN name that you configure for a vNIC appears as **Admin CDN Name**. The CDN name that is finally applied to the vNIC appears as **Oper CDN Name**. For example, if the **Admin CDN Name** for a vNIC called "vnic0" is cdn0, then the **Oper CDN Name** for this vNIC will be cdn0, but if the **Admin CDN Name** for the same vNIC is not specified, the **Oper CDN Name** will be vnic0.
- In Cisco UCS Manager Release 3.1 and older releases, downgrade of the adapter firmware is prevented if a CDN-enabled BIOS policy is assigned to a server.
- In Cisco UCS Manager Release 2.2(4), downgrade of Cisco UCS Manager or BIOS is prevented, if CDN enabled BIOS policy is assigned on the associated server profile.
- When the applied BIOS policy is changed from CDN-disabled to CDN-enabled or from CDN-enabled to CDN-disabled, the host reboots with a warning, irrespective of whether reboot on BIOS update is enabled or not.
- It is recommended that you enable CDN in the BIOS policy and add CDN names to the vNICs before the Windows Operating System is installed.
- If the Windows Operating System is already installed on the server and CDN is then enabled in the BIOS policy, do the following:
 1. Uninstall the network drivers.
 2. Scan the system for hidden devices and uninstall them.
 3. Rescan the system for new hardware and install the network drivers again.



Note If this is not done, the vNICs will not come up with the configured CDN names.

- When the applied BIOS policy is changed from CDN-disabled to CDN-enabled or from CDN-enabled to CDN-disabled on a service profile, do the following:
 1. Uninstall the network drivers.
 2. Scan the system for hidden devices and delete them.
 3. Re-scan the system for new hardware and install the network drivers again.



Note When the BIOS policy is changed from CDN-enabled to CDN-disabled, ensure that the CDN names are removed from all the vNICs on the system.

- If any change is made to the vNICs, the BDF of all the devices on the system also changes. Following are some of the scenarios that trigger a change in the BDF of all the vNICs present on the system:
 - When a vNIC is added or deleted
 - When a vNIC is moved from one adapter on the system to another adapter on the system

When these changes are made to the system, do the following:

1. Uninstall the network driver from all the present network interfaces.
2. Scan the system for hidden devices and uninstall them.
3. Re-scan the system for new hardware and install the network driver on the network controllers again.

If the hidden devices are not deleted, the CDN names of the network adapters will not appear as configured on Cisco UCS Manager.

CDN with a Mixed Set of Adapters

When a CDN name is configured for a vNIC in a system with a mixed set of CDN-supported adapters and CDN-unsupported adapters, then system placement may not place CDN-configured vNICs on adapters that support CDN.

If CDN is enabled in the BIOS policy, and system placement places a CDN-configured vNIC (Admin CDN configured) on an adapter that does not support CDN, an info fault will be raised, but the configuration issue for the service profile will be ignored.

If CDN is enabled in the BIOS policy, and system placement places a vNIC (Admin CDN not configured) on an adapter that does not support CDN, an info fault will be raised, but the configuration issue for the service profile will be ignored. The **Oper CDN Name** in this case will be empty and will not be derived from the vNIC name.

If you want to deploy the CDN name as the host network interface name for a server, you must manually place a vNIC on a supported adapter.

Configuring Consistent Device Naming in a BIOS Policy

Procedure

- Step 1** In the **Navigation** pane, click **Servers**.
- Step 2** Expand **Servers > Policies**.
- Step 3** Expand **root**.
- Step 4** Expand **BIOS Policies**.
- Step 5** Select the BIOS policy for which you want to configure CDN.
- Step 6** Under the **Main** tab, click one of the following in the **Consistent Device Naming** field to configure CDN:

Option	Description
disabled	Disables CDN in the BIOS policy
enabled	Enables CDN in the BIOS policy
Platform Default	The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.

Step 7 Click **Save Changes**.

Configuring a CDN Name for a vNIC

When a CDN name is not specified for a vNIC, the vNIC name is used as the CDN name.

Procedure

- Step 1** In the **Navigation** pane, click **Servers**.
- Step 2** Expand **Servers > Service Profiles**.
- Step 3** Expand the node for the organization that contains the vNIC for which you want to configure a CDN name.
- Step 4** Expand the service profile and **vNICs** node that contain the vNIC for which you want to configure a CDN name.
- Step 5** Select the vNIC.
- Step 6** Click on the **General** tab.
- Step 7** In the **Properties** area, choose **User Defined** as the **CDN Source**.
- Step 8** Enter the CDN name for the vNIC in the **CDN Name** field.

Note

The CDN name that you configure for a vNIC appears as **CDN Name**. The CDN name that is finally applied to the vNIC appears as **Oper CDN Name**. For example, if the **CDN Name** for a vNIC called "vnic0" is cdn0, then the **Oper CDN Name** for this vNIC will be cdn0, but if the **CDN Name** for the same vNIC is not specified, the **Oper CDN Name** will be vnic0.

Step 9 Click **Save Changes**.

CIMC Security Policies

Cisco UCS Manager provides the following policies to increase security:

- KVM Management Policy
- IPMI Access Profile

IPMI Access Profile

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

You can also restrict remote connectivity by disabling or enabling IPMI over LAN in the IPMI access profile. IPMI over LAN is disabled by default on all unassociated servers, and on all servers without an IPMI access

policy. When an IPMI access policy is created, the IPMI over LAN is set to enabled by default. If you do not change the value to disabled, IPMI over LAN will be enabled on all associated servers.

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

Creating an IPMI Access Profile

Before you begin

An IPMI profile requires that one or more of the following resources already exist in the system:

- Username with appropriate permissions that can be authenticated by the operating system of the server
- Password for the username
- Permissions associated with the username

Procedure

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

Step 2 Expand **Servers > Policies**.

Step 3 Expand the node for the organization where you want to create the policy.

If the system does not include multi tenancy, expand the **root** node.

Step 4 Right-click **IPMI Access Profiles** and select **Create IPMI Access Profile**.

Step 5 In the **Create IPMI Access Profile** dialog box:

- a) Enter a unique name and description for the profile.
- b) In the **IPMI Over LAN** field, choose whether to allow or restrict remote connectivity.
- c) Click **OK**.

Step 6 In the **IPMI Users** area of the navigator, click +.

Step 7 In the **Create IPMI User** dialog box:

- a) Complete the following fields:

Name	Description
Name field	The username to associate with this IPMI or Redfish profile. Enter 1 to 16 alphanumeric characters. You can also use @ (at sign), _ (underscore), and - (hyphen). You cannot change this name once the profile has been saved.
Password field	The password associated with this username. Enter 1 to 20 standard ASCII characters, except for = (equal sign), \$ (dollar sign), and (vertical bar).
Confirm Password field	The password a second time for confirmation purposes.

Name	Description
Role field	The user role. This can be one of the following: <ul style="list-style-type: none"> • Admin • Read Only
Description field	User-defined description of the IPMI or Redfish user.

b) Click **OK**.

Step 8 Repeat Steps 6 and 7 to add another user.

Step 9 Click **OK** to return to the IPMI profiles in the **Work** pane.

What to do next

Include the IPMI profile in a service profile and/or template.

Deleting an IPMI Access Profile

Procedure

- Step 1** In the **Navigation** pane, click **Servers**.
- Step 2** Expand **Servers > Policies > Organization_Name**.
- Step 3** Expand the **IPMI Profiles** node.
- Step 4** Right-click the profile you want to delete and select **Delete**.
- Step 5** If a confirmation dialog box displays, click **Yes**.

KVM Management Policy

The KVM Management policy allows you to determine whether vMedia encryption is enabled when you access a server via KVM.

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



Note After a KVM vMedia session is mapped, if you change the KVM management policy, it will result in a loss of the vMedia session. You must re-map the KVM vMedia session again.

Before Cisco UCS Manager Release 4.0(4), port 2068 was the only KVM port. Beginning with Release 4.0(4), you can configure a port number between 1024 and 49151 as the KVM port. Port 2068 continues to be the default KVM port number.

Creating a KVM Management Policy

Procedure

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

Step 2 Expand **Servers > Policies**.

Step 3 Expand the node for the organization where you want to create the policy.

If the system does not include multi tenancy, expand the **root** node.

Step 4 Right-click **KVM Management Policies** and select **Create KVM Management Policy**.

Step 5 In the **Create KVM Management Policy** dialog box:

- a) Enter a unique name and description for the policy.
- b) In the **vMedia Encryption** field, choose whether to enable vMedia encryption.

Note

Starting with UCS Manager 4.2, vMedia Encryption is always enabled for security purposes. It cannot be modified by the user.

- c) In the **KVM Port** field, enter a port number between 1024 and 49151 for KVM.

The default KVM port number is 2068.

- d) Click **OK**.

Note

After a KVM vMedia session is mapped, if you change the KVM management policy, it will result in a loss of the vMedia session. You must re-map the KVM vMedia session again.

Graphics Card Policies

Cisco UCS Manager Release 3.1(3) extends graphics card support to include the ability to change the graphics card mode. You can now configure graphics card modes by using a graphics card policy. The graphics card modes are:

- Compute
- Graphics
- Any Configuration

Creating a Graphics Card Policy



Note Cisco UCS Manager pushes the configuration changes to the GPU through the Graphics Card policy to the Processor Node Utility Operating System (PNUOS). These changes do not take effect until the server is rebooted.

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 **Graphics Card Policies** and select **Create Graphics Card Policy**.
- Step 5** On the **Main** page of the **Create Graphics Card Policy** dialog box:
- Enter a unique name for the policy.
 - (Optional) Enter a description for the policy.
 - In the **Graphics Card Mode** field, choose one of the options:
 - **Compute**
 - **Graphics**
 - **Any Configuration**
 - Click **OK**.
-

Local Disk Policies

Local Disk Configuration Policy

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

The local disk modes include the following:

- **No Local Storage**—For a diskless server or a SAN only configuration. If you select this option, you cannot associate any service profile which uses this policy with a server that has a local disk.
- **RAID 0 Striped**—Data is striped across all disks in the array, providing fast throughput. There is no data redundancy, and all data is lost if any disk fails.

- **RAID 1 Mirrored**—Data is written to two disks, providing complete data redundancy if one disk fails. The maximum array size is equal to the available space on the smaller of the two drives.
- **Any Configuration**—For a server configuration that carries forward the local disk configuration without any changes.
- **No RAID**—For a server configuration that removes the RAID and leaves the disk MBR and payload unaltered.

If you choose **No RAID** and you apply this policy to a server that already has an operating system with RAID storage configured, the system does not remove the disk contents. Therefore, there may be no visible differences on the server after you apply the **No RAID** mode. This can lead to a mismatch between the RAID configuration in the policy and the actual disk configuration shown in the **Inventory > Storage** tab for the server.

To make sure that any previous RAID configuration information is removed from a disk, apply a scrub policy that removes all disk information after you apply the **No RAID** configuration mode.

- **RAID 5 Striped Parity**—Data is striped across all disks in the array. Part of the capacity of each disk stores parity information that can be used to reconstruct data if a disk fails. RAID 5 provides good data throughput for applications with high read request rates.
- **RAID 6 Striped Dual Parity**—Data is striped across all disks in the array and two parity disks are used to provide protection against the failure of up to two physical disks. In each row of data blocks, two sets of parity data are stored.
- **RAID 10 Mirrored and Striped**—RAID 10 uses mirrored pairs of disks to provide complete data redundancy and high throughput rates.
- **RAID 50 Striped Parity and Striped**—Data is striped across multiple striped parity disk sets to provide high throughput and multiple disk failure tolerance.
- **RAID 60 Striped Dual Parity and Striped**—Data is striped across multiple striped dual parity disk sets to provide high throughput and greater disk failure tolerance.

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



Note For a Cisco UCS C-Series server integrated with Cisco UCS Manager, with an embedded on-board RAID controller, the local disk mode should always be **Any Configuration**, and the RAID must be configured directly on the controller.

Guidelines for all Local Disk Configuration Policies

Before you create a local disk configuration policy, consider the following guidelines:

No Mixed HDDs and SSDs

Do not include HDDs and SSDs in a single server or RAID configuration.

JBOD Mode Support

The B200 M3 server supports JBOD mode for local disks.



Note Only B200 M3 and B22 M3 blade servers support the JBOD mode for local disks.

Guidelines for Local Disk Configuration Policies Configured for RAID

Configure RAID Settings in Local Disk Configuration Policy for Servers with MegaRAID Storage Controllers

If a blade server or integrated rack-mount server has a MegaRAID controller, you must configure RAID settings for the drives in the Local Disk Configuration policy included in the service profile for that server. You can do this either by configuring the local disk configuration policy in the service profile using one of the defined RAID modes for that server, or you can use the **Any Configuration** mode with the LSI Utilities toolset to create the RAID volumes.

If you do not configure your RAID LUNs before installing the OS, disk discovery failures might occur during the installation and you might see error messages such as “No Device Found.”

Server May Not Boot After RAID1 Cluster Migration if Any Configuration Mode Specified in Service Profile

After RAID1 clusters are migrated, you need to associate a service profile with the server. If the local disk configuration policy in the service profile is configured with **Any Configuration** mode rather than **RAID1**, the RAID LUN remains in "inactive" state during and after association. As a result, the server cannot boot.

To avoid this issue, ensure that the service profile you associate with the server contains the identical local disk configuration policy as the original service profile before the migration and does not include the **Any Configuration** mode.

Do Not Use JBOD Mode on Servers with MegaRAID Storage Controllers

Do not configure or use JBOD mode or JBOD operations on any blade server or integrated rack-mount server with a MegaRAID storage controllers. JBOD mode and operations are not intended for nor are they fully functional on these servers.

Maximum of One RAID Volume and One RAID Controller in Integrated Rack-Mount Servers

A rack-mount server that has been integrated with Cisco UCS Manager can have a maximum of one RAID volume irrespective of how many hard drives are present on the server.

All the local hard drives in an integrated rack-mount server must be connected to only one RAID Controller. Integration with Cisco UCS Manager does not support the connection of local hard drives to multiple RAID Controllers in a single rack-mount server. We therefore recommend that you request a single RAID Controller configuration when you order rack-mount servers to be integrated with Cisco UCS Manager.

In addition, do not use third party tools to create multiple RAID LUNs on rack-mount servers. Cisco UCS Manager does not support that configuration.

Maximum of One RAID Volume and One RAID Controller in Blade Servers

A blade server can have a maximum of one RAID volume irrespective of how many drives are present in the server. All the local hard drives must be connected to only one RAID controller. For example, a B200 M3

server has an LSI controller and an Intel Patsburg controller, but only the LSI controller can be used as a RAID controller.

In addition, do not use third party tools to create multiple RAID LUNs on blade servers. Cisco UCS Manager does not support that configuration.

License Required for Certain RAID Configuration Options on Some Servers

Some Cisco UCS servers require a license for certain RAID configuration options. When Cisco UCS Manager associates a service profile containing this local disk policy with a server, Cisco UCS Manager verifies that the selected RAID option is properly licensed. If there are issues, Cisco UCS Manager displays a configuration error during the service profile association.

For RAID license information for a specific Cisco UCS server, see the *Hardware Installation Guide* for that server.

B420 M3 Server Does Not Support All Configuration Modes

The B420 M3 server does not support the following configuration modes in a local disk configuration policy:

- No RAID
- RAID 6 Striped Dual Parity

In addition, the B420 M3 does not support JBOD modes or operations.

Creating a Local Disk Configuration Policy

Procedure

-
- Step 1** In the **Navigation** pane, click **Servers**.
- Step 2** Expand **Servers > Policies**.
- Step 3** Expand the node for the organization where you want to create the policy.
If the system does not include multi tenancy, expand the **root** node.
- Step 4** Right-click **Local Disk Config Policies** and choose **Create Local Disk Configuration Policy**.
- Step 5** In the **Create Local Disk Configuration Policy** dialog box, complete the following fields:

Name	Description
Name field	The name of the policy. This name can be between 1 and 16 alphanumeric characters. You cannot use spaces or any special characters other than - (hyphen), _ (underscore), : (colon), and . (period), and you cannot change this name after the object is saved.

Name	Description
Description field	<p>A description of the policy. Cisco recommends including information about where and when to use the policy.</p> <p>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).</p>
Mode drop-down list	<p>This can be one of the following local disk policy modes:</p> <ul style="list-style-type: none"> • No Local Storage • RAID 0 Striped • RAID 1 Mirrored • Any Configuration • No RAID <p>If you choose No RAID and you apply this policy to a server that already has an operating system with RAID storage configured, the system does not remove the disk contents. Therefore, there may be no visible differences on the server after you apply the No RAID mode. This can lead to a mismatch between the RAID configuration in the policy and the actual disk configuration shown in the Inventory > Storage tab for the server.</p> <p>To make sure that any previous RAID configuration information is removed from a disk, apply a scrub policy that removes all disk information after you apply the No RAID configuration mode.</p> <ul style="list-style-type: none"> • RAID 5 Striped Parity • RAID 6 Striped Dual Parity • RAID 10 Mirrored and Striped • RAID 50 Striped Parity and Striped • RAID 60 Striped Dual Parity and Striped <p>Note Some Cisco UCS servers require a license for certain RAID configuration options. When Cisco UCS Manager associates a service profile containing this local disk policy with a server, Cisco UCS Manager verifies that the selected RAID option is properly licensed. If there are issues, Cisco UCS Manager displays a configuration error during the service profile association.</p> <p>For RAID license information for a specific Cisco UCS server, see the <i>Hardware Installation Guide</i> for that server.</p>

Name	Description
Protect Configuration check box	<p>If checked, the server retains the configuration in the local disk configuration policy even if the server is disassociated from the service profile.</p> <p>Caution Protect Configuration becomes non-functional if one or more disks in the server are defective or faulty.</p> <p>This property is checked by default.</p> <p>When a service profile is disassociated from a server and a new service profile associated, the setting for the Protect Configuration property in the new service profile takes precedence and overwrites the setting in the previous service profile.</p> <p>With this option enabled, the data on the disk is protected even after the server is decommissioned and then recommissioned. Hence, reassociation of the server with a service profile fails.</p> <p>Note If you disassociate the server from a service profile with this option enabled and then associate it with a new service profile that includes a local disk configuration policy with different properties, the server returns a configuration mismatch error and the association fails.</p>
FlexFlash State radio button	<p>To enable or disable the FlexFlash controller on the SD card, click the appropriate button.</p> <p>Note This parameter only applies to a server with an SD card module.</p>
FlexFlash RAID Reporting State radio button	<p>To enable or disable RAID reporting, click the appropriate button. When RAID reporting is enabled, the RAID status is monitored and faults are enabled.</p> <p>Note If only one SD card is installed, the RAID state will be displayed as Disabled and the RAID health as NA even if RAID reporting is enabled.</p>
FlexFlash Removable State radio button	<p>To select the removable state of the FlexFlash SD card, click the appropriate button.</p> <ul style="list-style-type: none"> • Yes—Use this option to define the SD card as removable. • No—Use this option to define the SD card as fixed or non-removable. • No Change—Use this option if the hypervisor does not require a preset state for the SD card.

Step 6 Click **OK**.

Changing a Local Disk Configuration Policy

This procedure describes how to change a local disk configuration policy from an associated service profile. You can also change a local disk configuration policy from the **Policies** node from **Servers**.

Procedure

-
- Step 1** In the **Navigation** pane, click **Servers**.
- Step 2** Expand **Servers** > **Service Profiles**.
- Step 3** Expand the organization that includes the service profile with the local disk configuration policy you want to change.
- If the system does not include multi tenancy, expand the **root** node.
- Step 4** Click the service profile that contains the local disk configuration policy you want to change.
- Step 5** In the **Work** pane, click the **Storage** tab.
- Step 6** In the **Actions** area, click **Change Local Disk Configuration Policy**.
- Step 7** In the **Change Local Disk Configuration Policy** dialog box, choose one of the following options from the **Select the Local Disk Configuration Policy** drop-down list.

Option	Description
Use a Disk Policy	Select an existing local disk configuration policy from the list below this option. Cisco UCS Manager assigns this policy to the service profile.
Create a Local Disk Policy	Enables you to create a local disk configuration policy that can only be accessed by the selected service profile.
No Disk Policy	Selects the default local disk policy.

Note
If a UCS server is attached to the Cisco UCS Manager, selecting the No Disk Policy can erase and replace the RAID with individual RAID 0 disks if the default RAID configuration is not supported on the attached server.

- Step 8** Click **OK**.
- Step 9** (Optional) Expand the **Local Disk Configuration Policy** area to confirm that the change has been made.
-

Deleting a Local Disk Configuration Policy

Procedure

-
- Step 1** In the **Navigation** pane, click **Servers**.
- Step 2** Expand **Servers** > **Policies** > *Organization_Name*.
- Step 3** Expand the **Local Disk Config Policies** node.

- Step 4** Right-click the policy you want to delete and select **Delete**.
- Step 5** If a confirmation dialog box displays, click **Yes**.

FlexFlash Support

Overview

Cisco UCS B-Series, C-Series M3 and higher, and S-Series M4 servers support internal Secure Digital (SD) memory cards. The SD cards are hosted by the Cisco Flexible Flash storage controller, a PCI-based controller which has two slots for SD cards. The cards contain a single partition called HV. When FlexFlash is enabled, Cisco UCS Manager displays the HV partition as a USB drive to both the BIOS and the host operating system.


You can populate one or both the SD card slots that are provided. If two SD cards are populated, you can use them in a mirrored mode.




Note Do not mix different capacity cards in the same server.

The SD cards can be used to store operating system boot images or other information. The following figure illustrates the SD card slots.

Figure 1: SD Card Slots





FlexFlash is disabled by default. You can enable FlexFlash in a local disk policy used in a service profile. When FlexFlash is enabled in a local disk policy, and the server is capable of supporting SD cards, the FlexFlash controller is enabled during service profile association. If a server is not capable of supporting SD cards or has an older CIMC version, a config failure message is displayed.

If you disable FlexFlash in a supported server, the Hypervisor or HV partition is immediately disconnected from the host. The FlexFlash controller will also be disabled as part of a related service profile disassociation.

The FlexFlash controller supports RAID-1 for dual SD cards. The FlexFlash scrub policy erases the HV partition in both cards, and brings the cards to a healthy RAID state.

You can configure new SD cards in a RAID pair and format them using one of the following methods:

- Format the SD cards. [Formatting the SD Cards, on page 291](#) provides detailed information.
- Disassociate the service profile from the server. Re-acknowledge the server after modifying the default scrub policy and then associate the server profile back to the server.

The *Scrub Policy Settings* section in the *Cisco UCS Manager Server Management Guide* provides more details about the usage of the scrub policy.



Note Disable the scrub policy as soon as the pairing is complete.

To boot from the HV partition, the SD card must be present in the boot policy used in the service profile.

FlexFlash Firmware Management

The FlexFlash controller firmware is bundled as part of the CIMC image. When you upgrade the CIMC, if a newer firmware version is available for the FlexFlash controller, the controller can no longer be managed, and the FlexFlash inventory displays the **Controller State** as **Waiting For User Action** and the **Controller Health** as **Old Firmware Running**. To upgrade the FlexFlash controller firmware, you need to perform a board controller update. For more information, see the appropriate *Cisco UCS B-Series Firmware Management Guide*, available at the following URL:

http://www.cisco.com/en/US/products/ps10281/products_installation_and_configuration_guides_list.html.

Limitations for the Cisco Flexible Flash Storage Controller:

- The Cisco Flexible Flash storage controller only supports 16 GB, 32 GB, and 64 GB SD cards.



Note The 16 GB and 32 GB cards are supported only on the M3 blade servers, and the 64 GB SD cards are supported only on the M4M5 blade servers.

- We do not recommend using an SD card from a rack server in a blade server, or using an SD card from a blade server in a rack server. Switching SD cards between server types might result in data loss from the SD card.
- Some Cisco UCS C-Series rack-mount servers have SD cards with four partitions: HV, HUU, SCU, and Drivers. Only the HV partition is visible in Cisco UCS Manager. You can migrate a four-partition SD card to a single HV partition card with a FlexFlash scrub policy but there may be data loss.
- The FlexFlash controller does not support RAID-1 sync (mirror rebuild). If the SD cards are in a degraded RAID state, or if any metadata errors are reported by the controller, you must run the FlexFlash scrub policy to pair the cards for RAID. For more information about the FlexFlash scrub policy, see [Server-Related Policies](#). The following conditions might result in degraded RAID or metadata errors:
 - Inserting a new or used SD card in one slot, when the server already has an SD card populated in the second slot.
 - Inserting two SD cards from different servers.
- The server firmware version must be at 2.2(1a) or higher.

FlexFlash FX3S Support

Beginning with Release 2.2(3), Cisco UCS Manager allows additional FlexFlash support with the FX3S controller. The FX3S controller is present on the following servers:

- Cisco UCS B200 M4 and M5 blade server
- Cisco UCS C220 M4 and M5 rack server
- Cisco UCS C240 M4 and M5 rack server
- C480 M5 rack server
- C480 M5 ML blade server
- B480 M5 blade server
- Cisco UCS C125 M5 Server

FlexFlash operations with the FX3S control are similar to those with the Cisco Flexible Flash storage controller. FlexFlash is disabled by default, and is enabled using a local disk policy. You can also reset the controller, format the SD cards, and enable automatic synchronization of your paired SD cards.

The SD cards for the FX3S controller contain a single partition called Hypervisor.

Limitations for the Cisco FX3S Controller:

- The FX3S controller supports only 32 GB and 64 GB SD cards. 16 GB cards are not supported.
- The FX3S controller supports 128 GB cards on M5 blades and above.
- We do not recommend using an SD card from a rack server in a blade server, or using an SD card from a blade server in a rack server. Switching SD cards between server types might result in data loss from the SD card.
- The server firmware version must be at 2.2(3a) or higher.

Starting Up Blade Servers with FlexFlash SD Cards

Use this procedure to start up blade servers using FlexFlash cards 16 GB and larger. This procedure requires that you know how to setup the blade server, software, and the associated infrastructure, and ensure that they are working. This Cisco UCS Manager controlled procedure is applicable to all blade servers, running any version of firmware. This procedure does not apply to rack servers. Follow this procedure before you enable FlexFlash cards in a working environment.



Caution If you use the following procedure with FlexFlash cards already in use, you will lose all data from the cards.



Note This procedure does not cover FlexFlash card usage or other functions of the FlexFlash system.

Procedure

-
- Step 1** Expand **Equipment** > **Chassis** > *Chassis Number* > **Servers**.
- Step 2** In the **Work** pane, check the details of the FlexFlash cards in the **FlexFlash Controller** window.
- Step 3** Expand **Servers** > **Service Profiles**.
- Step 4** Expand the node for the organization containing the pool.
If the system does not include multi tenancy, expand the **root** node.
- Step 5** Expand the node for the organization containing the service profile and click **Storage**.
- Step 6** In the **Work** pane, click **Change Local Disk Configuration Policy** in the **Actions** area and expand **Create Local Disk Configuration Policy** link. Follow the procedure in [Creating a Local Disk Configuration Policy, on page 282](#) to create a Local Disk Configuration Policy.
The FlexFlash policy name must not contain empty spaces or special characters.
- Step 7** Expand **Change Disk Local Configuration Policy**, and select the policy you just created and click **OK**.
- Step 8** Expand **Servers** > **Policies**.
- Step 9** Follow the procedure in [Creating a Scrub Policy, on page 294](#) and create a policy with a name such as *Scrub-FF-name* and click **OK**.
The Scrub policy name must not contain empty spaces or special characters.

- Step 10** Select the policy you created from the drop-down box.
 - Step 11** Expand **Equipment** > **Chassis** > *Chassis Number* > **Servers**.
 - Step 12** In the **Work** pane, click the **General** tab and select **Server Maintenance** from the **Actions** area.
 - Step 13** In the **Maintenance Server** dialogue box, click on the **Re-acknowledge** radio button, and then click **OK**.
 - Step 14** Click **Server Maintenance** in the **Action** area and click on the **Re-acknowledge** radio button again.
 - Step 15** From the **Inventory** tab, select the **Storage** sub-tab.
You can verify details of the enabled FlexFlash cards from the **FlexFlash Controller** window in the **Work** area.
 - Step 16** Launch KVM Manager and log on to the operating system. Verify details of the Hypervisor partition from the Devices and drives folder. Depending on the card size, the HV partition displays details of 32GB, 64GB, or 128 GB.
The FlexFlash cards are now synced and ready to use.
-

Enabling FlexFlash SD Card Support

Procedure

- Step 1** In the **Navigation** pane, click **Servers**.
 - Step 2** Expand **Servers** > **Policies**.
 - Step 3** Expand the node for the organization where you want to create the policy.
If the system does not include multi tenancy, expand the **root** node.
 - Step 4** Expand **Local Disk Config Policies** and choose the local disk config policy for which you want to enable FlexFlash support.
 - Step 5** In the **Work** pane, click the **General** tab.
 - Step 6** In the **FlexFlash State** field, click the **Enable** radio button.
 - Step 7** In the **FlexFlash RAID Reporting State** field, click the **Enable** radio button.
 - Step 8** Click **Save Changes**.
-

Enabling Auto-Sync

Procedure

- Step 1** In the **Navigation** pane, click **Equipment**.
- Step 2** Expand **Equipment** > **Chassis** > *Chassis Number* > **Servers**.
- Step 3** Click the server for which you want to enable auto-sync.
- Step 4** In the **Work** pane, click the **Inventory** tab.
- Step 5** Click the **Storage** subtab.
- Step 6** In the **Actions** area, click **Enable Auto-sync**.

- Step 7** In the **Enable Auto-sync** dialog box, choose the **Admin Slot Number** for the SD card that you want to use as the primary.
- Step 8** Click **OK**.
-

Formatting the SD Cards

Procedure

- Step 1** In the **Navigation** pane, click **Equipment**.
- Step 2** Expand **Equipment > Chassis > Chassis Number > Servers**.
- Step 3** Click the server for which you want to format the SD cards.
- Step 4** In the **Work** pane, click the **Inventory** tab.
- Step 5** Click the **Storage** subtab.
- Step 6** In the **Actions** area, click **Format SD Cards**.
- Step 7** Click **Yes** to format the SD cards.
-

Resetting the FlexFlash Controller

Procedure

- Step 1** In the **Navigation** pane, click **Equipment**.
- Step 2** Expand **Equipment > Chassis > Chassis Number > Servers**.
- Step 3** Click the server for which you want to reset the FlexFlash controller.
- Step 4** In the **Work** pane, click the **Inventory** tab.
- Step 5** Click the **Storage** subtab.
- Step 6** In the **Actions** area, click **Reset FlexFlash Controller**.
- Step 7** Click **Yes** to reset the FlexFlash controller.
-

Persistent Memory Modules

Cisco UCS Manager Release 4.0(4) introduces support for the Intel[®] Optane[™] Data Center persistent memory modules on the UCS M5 servers that are based on the Second Generation Intel[®] Xeon[®] Scalable processors. These persistent memory modules can be used only with the Second Generation Intel[®] Xeon[®] Scalable processors.

Persistent memory modules are non-volatile memory modules that bring together the low latency of memory and the persistence of storage. Data stored in persistent memory modules can be accessed quickly compared to other storage devices, and is retained across power cycles.

For detailed information about configuring persistent memory modules, see the *Cisco UCS: Configuring and Managing Intel® Optane™ Data Center Persistent Memory Modules Guide*.

Scrub Policy

Scrub Policy Settings

This policy determines what happens to local data and to the BIOS settings on a server during the discovery process, when the server is re-acknowledged, or when the server is disassociated from a service profile.



Note Local disk scrub policies only apply to hard drives that are managed by Cisco UCS Manager and do not apply to other devices such as USB drives.

Depending upon how you configure a scrub policy, the following can occur at those times:

Disk scrub

One of the following occurs to the data on any local drives on disassociation:

- If enabled, deletes initial 200MB of data from master boot record or the boot sectors. Thus, preventing the system to boot from an already installed OS if any. For secure deletion of data on drives, refer [UCS Secure Data Deletion For Commission Regulation \(EU\) 2019 /424 Users Guide](#).



Note Though the disk scrub policy is not intended to delete the user data that exceeds 200MB, Cisco UCS Manager cannot guarantee against data loss.

- If disabled (default), preserves all data on any local drives, including local storage configuration.

For a server associated with a service profile, disk scrub occurs during disassociation, based on the scrub policy used in the service profile. For an un-associated server, disk scrub occurs during the server discovery process, based on the default scrub policy.

Scrub policies are supported on all B-Series platforms and only on the following C-Series platforms:



Note You must re-acknowledge the server to see the changes related to LUN deletion if:

- you are scrubbing boot drives which have LUNs under the SAS controller in a set up with Cisco UCS S3260 M4 or Cisco UCS S3260 M5 Storage Server.
- you are scrubbing the LUNs on Cisco boot optimized M.2 RAID controller.

BIOS Settings Scrub

One of the following occurs to the BIOS settings when a service profile containing the scrub policy is disassociated from a server:

- If enabled, erases all BIOS settings for the server and resets them to the BIOS defaults for that server type and vendor.
- If disabled (default), preserves the existing BIOS settings on the server.

FlexFlash Scrub

FlexFlash Scrub enables you to pair new or degraded SD cards, resolve FlexFlash metadata configuration failures, and migrate older SD cards with 4 partitions to single partition SD cards. One of the following occurs to the SD card when a service profile containing the scrub policy is disassociated from a server, or when the server is reacknowledged:

- If enabled, the HV partition on the SD card is formatted using the PNUOS formatting utility. If two SD cards are present, the cards are RAID-1 paired, and the HV partitions in both cards are marked as valid. The card in slot 1 is marked as primary, and the card in slot 2 is marked as secondary.
- If disabled (default), preserves the existing SD card settings.



Note

- For a server associated with a service profile, FlexFlash scrub occurs during disassociation, based on the scrub policy used in the service profile. For an un-associated server, FlexFlash scrub occurs during the server discovery process, based on the default scrub policy.
 - Because the FlexFlash scrub erases the HV partition on the SD cards, we recommend that you take a full backup of the SD card(s) using your preferred host operating system utilities before performing the FlexFlash scrub.
 - To resolve metadata config failures in a service profile, you need to disable FlexFlash in the local disk config policy before you run the FlexFlash scrub, then enable FlexFlash after the server is reacknowledged.
 - Disable the scrub policy as soon as the pairing is complete or the metadata failures are resolved.
 - FlexFlash scrub is not supported for Cisco UCS S3260 Storage Server.
-

Persistent Memory Scrub

Persistent memory scrub enables you to preserve or remove the persistent memory configuration and data on a server.

- If enabled:
 - erases all the persistent memory data
 - resets the configuration to factory default
 - disables DIMM security
- If disabled (default), preserves the existing persistent memory configuration and data on the server. It does not change the DIMM lock state.

Creating a Scrub Policy

Procedure

- Step 1** In the **Navigation** pane, click **Servers**.
- Step 2** Expand **Servers > Policies**.
- Step 3** Expand the node for the organization where you want to create the policy.
If the system does not include multi tenancy, expand the **root** node.

- Step 4** Right-click **Scrub Policies** and select **Create Scrub Policy**.

Note

Cisco UCS Manager does not support NVME local disk scrub.

- Step 5** In the **Create Scrub Policy** wizard, complete the following fields:

Name	Description
Name field	The name of the policy. This name can be between 1 and 16 alphanumeric characters. You cannot use spaces or any special characters other than - (hyphen), _ (underscore), : (colon), and . (period), and you cannot change this name after the object is saved.
Description field	A description of the policy. Cisco recommends including information about where and when to use the policy. Enter up to 256 characters. You can use any characters or spaces except ` (accent mark), \ (backslash), ^ (carat), " (double quote), = (equal sign), > (greater than), < (less than), or ' (single quote).
Disk Scrub field	If this field is set to Yes , when a service profile containing this scrub policy is disassociated from a server, the initial 200MB of data is deleted from master boot record or the boot sectors. Thus, preventing the system to boot from an already installed OS if any. For secure deletion of data on drives, refer UCS Secure Data Deletion For Commission Regulation (EU) 2019 /424 Users Guide . If this field is set to No , the data on the local drives is preserved, including all local storage configuration. Note Though the disk scrub policy is not intended to delete the user data that exceeds 200MB, Cisco UCS Manager cannot guarantee against data loss.
BIOS Settings Scrub field	If the field is set to Yes , when a service profile containing this scrub policy is disassociated from a server, the BIOS settings for that server are erased and reset to the defaults for that server type and vendor. If this field is set to No , the BIOS settings are preserved.

Name	Description
FlexFlash Scrub field	If the field is set to Yes , the HV partition on the SD card is formatted using the PNUOS formatting utility when the server is reacknowledged. If this field is set to No , the SD card is preserved.
Persistent Memory Scrub field	If the field is set to Yes , when a service profile containing this scrub policy is disassociated from a server, all persistent memory modules for that server are erased and reset to the defaults for that server type and vendor. If this field is set to No , the persistent memory modules are preserved.

Step 6 Click **OK**.

Note

Disk Scrub option will scrub only boot-lun/boot-disk for Cisco UCS S3260 Storage Server and it will not scrub data-lun's/data-disks. The FlexFlash Scrub option is not supported for Cisco UCS S3260 Storage Server.

Deleting a Scrub Policy

Procedure

- Step 1** In the **Navigation** pane, click **Servers**.
- Step 2** Expand **Servers** > **Policies** > *Organization_Name*.
- Step 3** Expand the **Scrub Policies** node.
- Step 4** Right-click the policy you want to delete and select **Delete**.
- Step 5** If a confirmation dialog box displays, click **Yes**.

DIMM Error Management

DIMM Correctable Error Handling

In Cisco UCS Manager, when a DIMM encounters a significant correctable error in a given predefined window, it is stated as degraded and considered as a non-functional device.

The DIMM correctable error handling feature enables you to reset all the correctable and uncorrectable memory errors on all the DIMMs in a server. When you reset the error configuration, the error count of a given DIMM is cleared, the status changes to operable, and it resets the sensor state of the given DIMM.

Resetting Memory Errors

Use this procedure to reset all correctable and uncorrectable memory errors encountered by Cisco UCS Manager and the baseboard management controller (BMC).

Procedure

-
- Step 1** In the **Navigation** pane, click **Equipment**.
- Step 2** Expand **Equipment** > **Chassis** > *Chassis Number* > **Servers**.
- Step 3** Right-click on the server for which you want to reset the error configuration, and select **Reset All Memory Errors**. You can also select **Reset All Memory Errors** from the **Actions** area.
- Step 4** If the Cisco UCS Manager GUI displays a confirmation dialog box, click **Yes**.
-

DIMM Blacklisting

In Cisco UCS Manager, the state of the Dual In-line Memory Module (DIMM) is based on SEL event records. When the BIOS encounters a noncorrectable memory error during memory test execution, the DIMM is marked as faulty. A faulty DIMM is considered a nonfunctional device.

If you enable DIMM blacklisting, Cisco UCS Manager monitors the memory test execution messages and blacklists any DIMMs that encounter memory errors in the DIMM SPD data. To allow the host to map out any DIMMs that encounter uncorrectable ECC errors.

Enabling DIMM Blacklisting

The memory policy is a global policy that you can apply to existing servers on a Cisco UCS domain and also to the servers that are added after you set the memory policy.



-
- Note**
- This feature is supported both on the Cisco UCS B-Series blade servers and UCS C-Series rack servers.



-
- Note** Cisco UCS C-Series 420 M3 rack servers do not support this feature.
-

- This global policy cannot be added to a service profile.
-

Before you begin

- For Cisco B-Series blade server, the server firmware must be at Release 2.2(1) or a later release.
- For Cisco C-Series and S-Series rack server, the server firmware must be at Release 2.2(3).
- You must be logged in with one of the following privileges:
 - Admin
 - Server policy
 - Server profile server policy

Procedure

-
- Step 1** In the **Navigation** pane, click **Servers**.
 - Step 2** Expand **Servers > Policies**.
 - Step 3** Expand the node for the organization where you want to enable the blacklisting.
If the system does not include multitenancy, expand the **root** node.
 - Step 4** Expand **Memory Policy** and choose **default**.
 - Step 5** In the **Blacklisting** area, click the **Enabled** radio button.

The DIMM blacklisting is enabled for the domain level policy and these changes apply to all the servers on that particular domain.



Note If the Cisco IMC of a server does not support DIMM blacklisting, an information level fault is generated.

Serial over LAN Policy Settings

Serial over LAN Policy Overview

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

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

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

Creating a Serial over LAN Policy

Procedure

-
- Step 1** In the **Navigation** pane, click **Servers**.
 - Step 2** Expand **Servers > Policies**.
 - Step 3** Expand the node for the organization where you want to create the policy.
If the system does not include multi tenancy, expand the **root** node.
 - Step 4** Right-click **Serial over LAN Policies** and select **Create Serial over LAN Policy**.
 - Step 5** In the **Create Serial over LAN Policy** wizard, complete the following fields:

Name	Description
Name field	<p>The name of the policy.</p> <p>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.</p>
Description field	<p>A description of the policy. Cisco recommends including information about where and when to use the policy.</p> <p>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).</p>
Serial over LAN State field	<p>This can be one of the following:</p> <ul style="list-style-type: none"> • Disable—Serial over LAN access is blocked. • Enable—Serial over LAN access is permitted.
Speed drop-down list	<p>This can be one of the following:</p> <ul style="list-style-type: none"> • 9600 • 19200 • 38400 • 57600 • 115200

Step 6 Click **OK**.

Deleting a Serial over LAN Policy

Procedure

-
- Step 1** In the **Navigation** pane, click **Servers**.
- Step 2** Expand **Servers > Policies > *Organization_Name***.
- Step 3** Expand the **Serial over LAN Policies** node.
- Step 4** Right-click the policy you want to delete and select **Delete**.
- Step 5** If a confirmation dialog box displays, click **Yes**.
-

Server Autoconfiguration Policies

Server Autoconfiguration Policy Overview

Cisco UCS Manager uses this policy to determine how to configure a new server. If you create a server autoconfiguration policy, the following occurs when a new server starts:

1. The qualification in the server autoconfiguration policy is executed against the server.
2. If the server meets the required qualifications, the server is associated with a service profile created from the service profile template configured in the server autoconfiguration policy. The name of that service profile is based on the name given to the server by Cisco UCS Manager.
3. The service profile is assigned to the organization configured in the server autoconfiguration policy.

Creating an Autoconfiguration Policy

Before you begin

This policy requires that one or more of the following resources already exist in the system:

- Server pool policy qualifications
- Service profile template
- Organizations, if a system implements multitenancy

Procedure

- Step 1** In the **Navigation** pane, click **Equipment**.
- Step 2** Click the **Equipment** node.
- Step 3** In the **Work** pane, click the **Policies** tab.
- Step 4** Click the **Autoconfig Policies** subtab.
- Step 5** On the icon bar to the right of the table, click +.
- If the + icon is disabled, click an entry in the table to enable it.
- Step 6** In the **Create Autoconfiguration Policy** dialog box, complete the following fields:

Name	Description
Name field	The name of the policy. This name can be between 1 and 16 alphanumeric characters. You cannot use spaces or any special characters other than - (hyphen), _ (underscore), : (colon), and . (period), and you cannot change this name after the object is saved.

Name	Description
Description field	A description of the policy. Cisco recommends including information about where and when to use the policy. Enter up to 256 characters. You can use any characters or spaces except ` (accent mark), \ (backslash), ^ (carat), " (double quote), = (equal sign), > (greater than), < (less than), or ' (single quote).
Qualification drop-down list	The server pool policy qualification associated with this auto-configuration policy. If a new server is discovered that matches the criteria specified in the server pool policy qualification, Cisco UCS automatically creates a service profile based on the service profile template selected in the Service Profile Template Name drop-down list and associates the newly created service profile with the server.
Org drop-down list	The organization associated with this autoconfiguration policy. If Cisco UCS automatically creates a service profile to associate with a server, it places the service profile under the organization selected in this field.
Service Profile Template Name drop-down list	The service profile template associated with this policy.

Step 7 Click **OK**.

Deleting an Autoconfiguration Policy

Procedure

- Step 1** In the **Navigation** pane, click **Equipment**.
 - Step 2** Click the **Equipment** node.
 - Step 3** In the **Work** pane, click the **Policies** tab.
 - Step 4** Click the **Autoconfig Policies** subtab.
 - Step 5** Right-click the autoconfiguration policy that you want to delete and choose **Delete**.
 - Step 6** If a confirmation dialog box displays, click **Yes**.
-

Server Discovery Policy Settings

Server Discovery Policy Overview

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

If you create a server discovery policy, the following occurs when a new server starts:

1. The server discovery policy qualification is executed against the server.
2. If the server meets the required qualifications, Cisco UCS Manager applies the following to the server:
 - Depending on the option that you select for the action, UCS Manager discovers the new server immediately, or waits for a user acknowledgment of the new server
 - Applies the scrub policy to the server

If automatic deep discovery is triggered by any hardware insertion, removal, or replacement, the following occurs:

1. The server is moved to a “pending activities” list.
2. A critical hardware mismatch fault is raised on the server, indicating that UCSM has detected a hardware mismatch.
3. User must explicitly acknowledge the server to trigger the deep discovery.



Important

In Cisco UCS Manager Release 2.2 (4), blade servers do not support drives with a block size of 4K, but rack-mount servers support such drives. If a drive with a block size of 4K is inserted into a blade server, discovery fails and the following error message appears:

```
Unable to get Scsi Device Information from the system
```

If this error occurs, do the following:

1. Remove the 4K drive.
2. Reacknowledge the server.

Reacknowledging the server causes the server to reboot and results in loss of service.

Creating a Server Discovery Policy

Before you begin

If you plan to associate this policy with a server pool, create server pool policy qualifications.

Procedure

- Step 1** In the **Navigation** pane, click **Equipment**.
- Step 2** In the **Work** pane, click the **Policies** tab.
- Step 3** Click the **Server Discovery Policies** subtab.
- Step 4** Click the + icon on the table icon bar to open the **Create Server Discovery Policy** dialog box.
- Step 5** In the **Description** field, enter a description for the discovery policy.
- Step 6** In the **Action** field, select one of the following options:
- **Immediate**—Cisco UCS Manager attempts to discover new servers automatically
 - **User Acknowledged**—Cisco UCS Manager waits until the user tells it to search for new servers
- Step 7** (Optional) To associate this policy with a server pool, select server pool policy qualifications from the **Qualification** drop-down list.
- Step 8** (Optional) To include a scrub policy, select a policy from the **Scrub Policy** drop-down list.
- Step 9** Click **OK**.
-

What to do next

Include the server discovery policy in a service profile and/or template.

Deleting a Server Discovery Policy

Procedure

- Step 1** In the **Navigation** pane, click **Equipment**.
- Step 2** In the **Work** pane, click the **Policies** tab.
- Step 3** Click the **Server Discovery Policies** subtab.
- Step 4** Right-click the server discover policy that you want to delete and choose **Delete**.
- Step 5** If a confirmation dialog box displays, click **Yes**.
-

Hardware Change Discovery Policy

The Hardware Change Discovery is a global policy used to set the how Cisco UCS Manager behaves when there is a hardware component change. The policy has two values:

- **User Acknowledged**: You must acknowledge the server to clear all the hardware inventory mismatch faults.
- **Auto Acknowledged**: Triggers automatic deep discovery when a hardware component change is detected.

When UCSM detects any change in the server hardware component, a critical hardware inventory mismatch fault is raised on the server. You must manually acknowledge the server to clear the fault and complete the hardware inventory. Once you have acknowledged the server, deep discovery and deep association is triggered.

For rack servers, you must decommission and recommission the server to clear the fault and complete the hardware inventory.

You cannot make changes to the policy if there is a hardware inventory mismatch fault.

Configuring Hardware Change Discovery Policy

Procedure

Step 1 Navigate to **Equipment > Policies > Global Policies**

Step 2 Under **Hardware Change Discovery Policy**, choose one of the following:

- **User Acknowledged:** You must acknowledge the server to clear all the hardware inventory mismatch faults.
- **Auto Acknowledged:** Triggers automatic deep discovery when a hardware component change is detected.

Step 3 Click **Save Changes**.

Server Inheritance Policy Settings

Server Inheritance Policy Overview

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

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

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

Creating a Server Inheritance Policy

A blade server or rack-mount server with a VIC adapter, such as the Cisco UCS M81KR Virtual Interface Card, does not have server identity values burned into the server hardware at manufacture. As a result, the identity of the adapter must be derived from default pools. If the default pools do not include sufficient entries for one to be assigned to the server, service profile association fails with a configuration error.

Procedure

-
- Step 1** In the **Navigation** pane, click **Equipment**.
- Step 2** In the **Work** pane, click the **Policies** tab.
- Step 3** Click the **Server Inheritance Policies** subtab.
- Step 4** On the icon bar at the bottom of the table, click + **Add**.
If + **Add** is disabled, click an entry in the table to enable it.
- Step 5** In the **Create Server Inheritance Policy** dialog box, complete the following fields:

Name	Description
Name field	The name of the policy. This name can be between 1 and 16 alphanumeric characters. You cannot use spaces or any special characters other than - (hyphen), _ (underscore), : (colon), and . (period), and you cannot change this name after the object is saved.
Description field	A description of the policy. Cisco recommends including information about where and when to use the policy. Enter up to 256 characters. You can use any characters or spaces except ` (accent mark), \ (backslash), ^ (carat), " (double quote), = (equal sign), > (greater than), < (less than), or ' (single quote).
Qualification drop-down list	To associate this policy with one or more specific server pools, choose the server pool qualification policy that identifies these pools.
Org drop-down list	If you want to associate an organization with this policy, or if you want to change the current association, choose the organization from the drop-down list.

- Step 6** Click **OK**.
-

Deleting a Server Inheritance Policy

Procedure

-
- Step 1** In the **Navigation** pane, click **Equipment**.
- Step 2** In the **Work** pane, click the **Policies** tab.
- Step 3** Click the **Server Inheritance Policies** subtab.
- Step 4** Right-click the server inheritance policy that you want to delete and choose **Delete**.

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

Server Pool Policy Settings

Server Pool Policy Overview

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

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

Creating a Server Pool Policy

Before you begin

This policy requires that one or more of the following resources already exist in the system:

- A minimum of one server pool
- Server pool policy qualifications, if you choose to have servers automatically added to pools

Procedure

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

Step 2 Expand **Servers > Policies**.

Step 3 Expand the node for the organization where you want to create the policy.

If the system does not include multi tenancy, expand the **root** node.

Step 4 Right-click **Server Pool Policies** and select **Create Server Pool Policy**.

Step 5 In the **Create Server Pool Policy** dialog box, complete the following fields:

Name	Description
Name field	<p>The name of the policy.</p> <p>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.</p>

Name	Description
Description field	A description of the policy. Cisco recommends including information about where and when to use the policy. Enter up to 256 characters. You can use any characters or spaces except ` (accent mark), \ (backslash), ^ (carat), " (double quote), = (equal sign), > (greater than), < (less than), or ' (single quote).
Target Pool drop-down list	If you want to associate this policy with a server pool, select that pool from the drop-down list.
Qualification drop-down list	To associate this policy with one or more specific server pools, choose the server pool qualification policy that identifies these pools.

Step 6 Click **OK**.

Deleting a Server Pool Policy

Procedure

- Step 1** In the **Navigation** pane, click **Servers**.
- Step 2** Expand **Servers > Policies > Organization_Name**.
- Step 3** Expand the **Server Pool Policies** node.
- Step 4** Right-click the policy you want to delete and select **Delete**.
- Step 5** If a confirmation dialog box displays, click **Yes**.

Server Pool Policy Qualifications Settings

Server Pool Policy Qualification Overview

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

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

You can use the server pool policy qualifications to qualify servers according to the following criteria:

- Adapter type

- Chassis location
- Memory type and configuration
- Power group
- CPU cores, type, and configuration
- Storage configuration and capacity
- Server model

Depending upon the implementation, you might need to configure several policies with server pool policy qualifications including the following:

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

Creating Server Pool Policy Qualifications

Procedure

-
- Step 1** In the **Navigation** pane, click **Servers**.
- Step 2** Expand **Servers > Policies**.
- Step 3** Expand the node for the organization where you want to create the policy.
If the system does not include multi tenancy, expand the **root** node.
- Step 4** Right-click the **Server Pool Policy Qualifications** node and select **Create Server Pool Policy Qualification**.
- Step 5** In the **Create Server Pool Policy Qualification** dialog box, enter a unique name and description for the policy.
- Step 6** (Optional) To use this policy to qualify servers according to their adapter configuration, do the following:
- Click **Create Adapter Qualifications**.
 - In the **Create Adapter Qualifications** dialog box, complete the following fields:

Name	Description
Type drop-down list	The adapter type. Once you save the adapter qualification, this type cannot be changed.
PID field	A regular expression that the adapter PID must match.

Name	Description
Maximum Capacity field	<p>The maximum capacity for the selected type.</p> <p>To specify a capacity, choose select and enter the desired maximum capacity. You can enter an integer between 1 and 65535.</p>

- c) Click **OK**.

Step 7

(Optional) To use this policy to qualify servers according to the chassis in which they physically reside, do the following:

- a) Click **Create Chassis/Server Qualifications**.
- b) In the **Chassis Qualifications** area of the **Create Chassis and Server Qualifications** dialog box, complete the following fields to specify the range of chassis you want to use:
 - **First Chassis ID** field—The first chassis ID from which server pools associated with this policy can draw.
 - **Number of Chassis** field—The total number of chassis to include in the pool, starting with the chassis identified in the **First Chassis ID** field.

Example:

For example, if you want to use chassis 5, 6, 7, and 8, enter **5** in the **First Chassis ID** field and **4** in the **Number of Chassis** field. If you want to use only chassis 3, enter **3** in the **First Chassis ID** field and **1** in the **Number of Chassis** field.

Tip

If you want to use chassis 5, 6, and 9, create a chassis/server qualification for the range 5-6 and another qualification for chassis 9. You can add as many chassis/server qualifications as needed.

- c) Click **Finish**.

Step 8

(Optional) To use this policy to qualify servers according to both the chassis and slot in which they physically reside, do the following:

- a) Click **Create Chassis/Server Qualifications**.
- b) In the **Chassis Qualifications** area of the **Create Chassis and Server Qualifications** dialog box, complete the following fields to specify the range of chassis you want to use:
 - **First Chassis ID** field—The first chassis ID from which server pools associated with this policy can draw.
 - **Number of Chassis** field—The total number of chassis to include in the pool, starting with the chassis identified in the **First Chassis ID** field.
- c) In the **Server Qualifications** table, click **Add**.
- d) In the **Create Server Qualifications** dialog box, complete the following fields to specify the range of server locations you want to use:
 - **First Slot ID** field—The first slot ID from which server pools associated with this policy can draw.
 - **Number of Slots** field—The total number of slots from which server pools associated with this policy can draw.
- e) Click **Finish Stage**.

- f) To add another range of slots, click **Add** and repeat steps d and e.
- g) When you have finished specifying the slot ranges, click **Finish**.

Step 9

(Optional) To use this policy to qualify servers according to their memory configuration, do the following:

- a) Click **Create Memory Qualifications**.
- b) In the **Create Memory Qualifications** dialog box, complete the following fields:

Name	Description
Clock field	The minimum clock speed required, in megahertz.
Latency field	The maximum latency allowed, in nanoseconds.
Min Cap field	The minimum memory capacity required, in megabytes.
Max Cap field	The maximum memory capacity allowed, in megabytes.
Width field	The minimum width of the data bus.
Units field	The unit of measure to associate with the value in the Width field.

- c) Click **OK**.

Step 10

(Optional) To use this policy to qualify servers according to their CPU/Cores configuration, do the following:

- a) Click **Create CPU/Cores Qualifications**.
- b) In the **Create CPU/Cores Qualifications** dialog box, complete the following fields:

Name	Description
Processor Architecture drop-down list	The CPU architecture to which this policy applies.
PID field	A regular expression that the processor PID must match.
Min Number of Cores field	The minimum number of CPU cores required. To specify a capacity, choose select and enter an integer between 1 and 65535 in the associated text field.
Max Number of Cores field	The maximum number of CPU cores allowed. To specify a capacity, choose select and enter an integer between 1 and 65535 in the associated text field.
Min Number of Threads field	The minimum number of CPU threads required. To specify a capacity, choose select and enter an integer between 1 and 65535 in the associated text field.
Max Number of Threads field	The maximum number of CPU threads allowed. To specify a capacity, choose select and enter an integer between 1 and 65535 in the associated text field.

Name	Description
CPU Speed field	The minimum CPU speed required. To specify a capacity, choose select and enter the minimum CPU speed.
CPU Stepping field	The minimum CPU version required. To specify a capacity, choose select and enter the maximum CPU speed.

c) Click **OK**.

Step 11

(Optional) To use this policy to qualify servers according to their storage configuration and capacity, do the following:

- a) Click **Create Storage Qualifications**.
- b) In the **Create Storage Qualifications** dialog box, complete the following fields:

Name	Description
Diskless field	Whether the available storage must be diskless. This can be one of the following: <ul style="list-style-type: none"> • Unspecified—Either storage type is acceptable. • Yes—The storage must be diskless. • No—The storage cannot be diskless.
Number of Blocks field	The minimum number of blocks required. To specify a capacity, choose select and enter the number of blocks.
Block Size field	The minimum block size required, in bytes. To specify a capacity, choose select and enter the block size.
Min Cap field	The minimum storage capacity across all disks in the server, in megabytes. To specify a capacity, choose select and enter the minimum storage capacity.
Max Cap field	The maximum storage capacity allowed, in megabytes. To specify a capacity, choose select and enter the maximum storage capacity.
Per Disk Cap field	The minimum storage capacity per disk required, in gigabytes. To specify a capacity, choose select and enter the minimum capacity on each disk.
Units field	The number of units. To specify a capacity, choose select and enter the desired units.

Name	Description
Number of Flex Flash Cards field	The number of FlexFlash Cards. To specify a capacity, choose select and enter the desired units.
Disk Type field	The disk type. This can be one of the following: <ul style="list-style-type: none"> • Unspecified—Either disk type is acceptable. • HDD—The disk must be HDD. • SSD—The disk must be SSD (SATA or SAS).

c) Click **OK**.

Step 12 (Optional) To use this policy to qualify servers according to the model of the server, do the following:

- Click **Create Server Model Qualifications**.
- In the **Create Server Model Qualifications** dialog box, enter a regular expression that the server model must match in the **Model** field.
- Click **OK**.

Step 13 (Optional) To use this policy to qualify servers according to power group, do the following:

- Click **Create Power Group Qualifications**.
- In the **Create Power Group Qualifications** dialog box, choose a power group from the **Power Group** drop-down list.
- Click **OK**.

Step 14 (Optional) To use this policy to qualify the rack-mount servers that can be added to the associated server pool, do the following:

- Click **Create Rack Qualifications**.
- In the **Create Rack Qualifications** dialog box, complete the following fields:

Name	Description
First Slot ID field	The first rack-mount server slot ID from which server pools associated with this policy can draw.
Number of Slots field	The total number of rack-mount server slots from which server pools associated with this policy can draw.

Step 15 Verify the qualifications in the table and correct if necessary.

Step 16 Click **OK**.

Deleting Server Pool Policy Qualifications

Procedure

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

- Step 2** Expand **Servers > Policies > *Organization_Name***.
 - Step 3** Expand the **Server Pool Policy Qualifications** node.
 - Step 4** Right-click the policy qualifications you want to delete and select **Delete**.
 - Step 5** If a confirmation dialog box displays, click **Yes**.
-

Deleting Qualifications from Server Pool Policy Qualifications

Use this procedure to modify Server Pool Policy Qualifications by deleting one or more sets of qualifications.

Procedure

-
- Step 1** In the **Navigation** pane, click **Servers**.
 - Step 2** Expand **Servers > Policies > *Organization_Name***.
 - Step 3** Expand the **Server Pool Policy Qualifications** node.
 - Step 4** Choose the policy you want to modify.
 - Step 5** In the **Work** pane, choose the **Qualifications** tab.
 - Step 6** To delete a set of qualifications:
 - a) In the table, choose the row that represents the set of qualifications.
 - b) Right-click the row and select **Delete**.
 - Step 7** Click **Save Changes**.
-

vNIC/vHBA Placement Policy Settings

vNIC/vHBA Placement Policies

vNIC/vHBA placement policies are used to determine the following:

- How the virtual network interface connections (vCons) are mapped to the physical adapters on a server.
- What types of vNICs or vHBAs can be assigned to each vCon.

Each vNIC/vHBA placement policy contains four vCons that are virtual representations of the physical adapters. When a vNIC/vHBA placement policy is assigned to a service profile, and the service profile is associated with a server, the vCons in the vNIC/vHBA placement policy are assigned to the physical adapters and the vNICs and vHBAs are assigned to those vCons.

For blade or rack servers that contain one adapter, Cisco UCS assigns all vCons to that adapter. For servers that contain four adapters, Cisco UCS assigns vCon1 to Adapter1, vCon2 to Adapter2, vCon3 to Adapter3, and vCon4 to Adapter4.

For blade or rack servers that contain two or three adapters, Cisco UCS assigns the vCons based on the type of server and the selected virtual slot mapping scheme, which can be **Round Robin** or **Linear Ordered**. For details about the available mapping schemes, see [vCon to Adapter Placement, on page 313](#).

After Cisco UCS assigns the vCons, it assigns the vNICs and vHBAs based on the **Selection Preference** for each vCon. This can be one of the following:



Note You can specify the PCI order for the vHBA; however, the desired order works within a class of devices, such as vNICs or vHBAs and not across them. Within an adapter, vNICs are always placed ahead of the vHBAs.

- **All**—All configured vNICs and vHBAs can be assigned to the vCon, whether they are explicitly assigned to it, unassigned, or dynamic. This is the default.
- **Assigned Only**—vNICs and vHBAs must be explicitly assigned to the vCon. You can assign them explicitly through the service profile or the properties of the vNIC or vHBA.
- **Exclude Dynamic**—Dynamic vNICs and vHBAs cannot be assigned to the vCon. The vCon can be used for all static vNICs and vHBAs, whether they are unassigned or explicitly assigned to it.
- **Exclude Unassigned**—Unassigned vNICs and vHBAs cannot be assigned to the vCon. The vCon can be used for dynamic vNICs and vHBAs and for static vNICs and vHBAs that are explicitly assigned to it.
- **Exclude usNIC**—Cisco usNICs cannot be assigned to the vCon. The vCon can be used for all other configured vNICs and vHBAs, whether they are explicitly assigned to it, unassigned, or dynamic.



Note An SRIOV usNIC that is explicitly assigned to a vCon set to **Exclude usNIC** will remain assigned to that vCon.

If you do not include a vNIC/vHBA placement policy in the service profile, Cisco UCS Manager defaults to the **Round Robin** vCon mapping scheme and the **All** vNIC/vHBA selection preference, distributing the vNICs and vHBAs between the adapters based on the capabilities and relative capacities of each adapter.

vCon to Adapter Placement

Cisco UCS maps every vCon in a service profile to a physical adapter on the server. How that mapping occurs and how the vCons are assigned to a specific adapter in a server depends on the following:

- The type of server. N20-B6620-2 and N20-B6625-2 blade servers with two adapter cards use a different mapping scheme than other supported rack or blade servers.
- The number of adapters in the server.
- The setting of the virtual slot mapping scheme in the vNIC/vHBA placement policy, if applicable.

You must consider this placement when you configure the vNIC/vHBA selection preference to assign vNICs and vHBAs to vCons.



Note vCon to adapter placement is not dependent upon the PCIE slot number of the adapter. The adapter numbers used for the purpose of vCon placement are not the PCIE slot numbers of the adapters, but the ID assigned to them during server discovery.

vCon to Adapter Placement for Blade Servers

If the blade two adapters are numbered left to right while vCons are numbered right to left. If one of these blade servers has a single adapter, Cisco UCS assigns all vCons to that adapter. If the server has two adapters, the vCon assignment depends upon the virtual slot mapping scheme:

- **Round Robin**—Cisco UCS assigns vCon2 and vCon4 to Adapter1 and vCon1 and vCon3 to Adapter2. This is the default.
- **Linear Ordered**—Cisco UCS assigns vCon3 and vCon4 to Adapter1 and vCon1 and vCon2 to Adapter2.

vCon to Adapter Placement for All Other Supported Servers

For all supported servers, the vCon assignment depends on the number of adapters in the server and the virtual slot mapping scheme.

For blade or rack servers that contain one adapter, Cisco UCS assigns all vCons to that adapter. For servers that contain four adapters, Cisco UCS assigns vCon1 to Adapter1, vCon2 to Adapter2, vCon3 to Adapter3, and vCon4 to Adapter4.

For blade or rack servers that contain two or three adapters, Cisco UCS assigns the vCons based on the selected virtual slot mapping scheme: Round Robin or Linear Ordered.

Table 9: vCon to Adapter Placement Using the Round - Robin Mapping Scheme

Number of Adapters	vCon1 Assignment	vCon2 Assignment	vCon3 Assignment	vCon4 Assignment
1	Adapter1	Adapter1	Adapter1	Adapter1
2	Adapter1	Adapter2	Adapter1	Adapter2
3	Adapter1	Adapter2	Adapter3	Adapter2
4	Adapter1	Adapter2	Adapter3	Adapter4

Round Robin is the default mapping scheme.

Table 10: vCon to Adapter Placement Using the Linear Ordered Mapping Scheme

Number of Adapters	vCon1 Assignment	vCon2 Assignment	vCon3 Assignment	vCon4 Assignment
1	Adapter1	Adapter1	Adapter1	Adapter1
2	Adapter1	Adapter1	Adapter2	Adapter2
3	Adapter1	Adapter2	Adapter3	Adapter3
4	Adapter1	Adapter2	Adapter3	Adapter4

vNIC/vHBA to vCon Assignment

Cisco UCS Manager provides two options for assigning vNICs and vHBAs to vCons through the vNIC/vHBA placement policy: explicit assignment and implicit assignment.

Explicit Assignment of vNICs and vHBAs

With explicit assignment, you specify the vCon and, therefore, the adapter to which a vNIC or vHBA is assigned. Use this assignment option when you need to determine how the vNICs and vHBAs are distributed between the adapters on a server.

To configure a vCon and the associated vNICs and vHBAs for explicit assignment, do the following:

- Set the vCon configuration to any of the available options. You can configure the vCons through a vNIC/vHBA placement policy or in the service profile associated with the server. If a vCon is configured for **All**, you can still explicitly assign a vNIC or vHBA to that vCon.
- Assign the vNICs and vHBAs to a vCon. You can make this assignment through the virtual host interface placement properties of the vNIC or vHBA or in the service profile associated with the server.

If you attempt to assign a vNIC or vHBA to a vCon that is not configured for that type of vNIC or vHBA, Cisco UCS Manager displays a message advising you of the configuration error.

During service profile association, Cisco UCS Manager validates the configured placement of the vNICs and vHBAs against the number and capabilities of the physical adapters in the server before assigning the vNICs and vHBAs according to the configuration in the policy. Load distribution is based upon the explicit assignments to the vCons and adapters configured in this policy.

If the adapters do not support the assignment of one or more vNICs or vHBAs, Cisco UCS Manager raises a fault against the service profile.



Note You can specify the PCI order for the vHBA; however, the desired order works within a class of devices, such as vNICs or vHBAs and not across them. Within an adapter, vNICs are always placed ahead of the vHBAs.

Implicit Assignment of vNICs and vHBAs

With implicit assignment, Cisco UCS Manager determines the vCon and, therefore, the adapter to which a vNIC or vHBA is assigned according to the capability of the adapters and their relative capacity. Use this assignment option if the adapter to which a vNIC or vHBA is assigned is not important to your system configuration.

To configure a vCon for implicit assignment, do the following:

- Set the vCon configuration to **All**, **Exclude Dynamic**, or **Exclude Unassigned**. You can configure the vCons through a vNIC/vHBA placement policy or in the service profile associated with the server.
- Do not set the vCon configuration to **Assigned Only**. Implicit assignment cannot be performed with this setting.
- Do not assign any vNICs or vHBAs to a vCon.

During service profile association, Cisco UCS Manager verifies the number and capabilities of the physical adapters in the server and assigns the vNICs and vHBAs accordingly. Load distribution is based upon the

capabilities of the adapters, and placement of the vNICs and vHBAs is performed according to the actual order determined by the system. For example, if one adapter can accommodate more vNICs than another, that adapter is assigned more vNICs.

If the adapters cannot support the number of vNICs and vHBAs configured for that server, Cisco UCS Manager raises a fault against the service profile.

Implicit Assignment of vNICs in a Dual Adapter Environment

When you use implicit vNIC assignment for a dual slot server with an adapter card in each slot, Cisco UCS Manager typically assigns the vNICs/vHBAs as follows:

- If the server has the same adapter in both slots, Cisco UCS Manager assigns half the vNICs and half the vHBAs to each adapter.
- If the server has one non-VIC adapter and one VIC adapter, Cisco UCS Manager assigns two vNICs and two vHBAs to the non-VIC adapter and the remaining vNICs and vHBAs to the VIC adapter.
- If the server has two different VIC adapters, Cisco UCS Manager assigns the vNICs and vHBAs proportionally, based on the relative capabilities of the two adapters.

The following examples show how Cisco UCS Manager would typically assign the vNICs and vHBAs with different combinations of supported adapter cards:

- If you want to configure four vNICs and the server contains two Cisco UCS M51KR-B Broadcom BCM57711 adapters (with two vNICs each), Cisco UCS Manager assigns two vNICs to each adapter.
- If you want to configure 50 vNICs and the server contains a Cisco UCS CNA M72KR-E adapter (2 vNICs) and a Cisco UCS M81KR Virtual Interface Card adapter (128 vNICs), Cisco UCS Manager assigns two vNICs to the Cisco UCS CNA M72KR-E adapter and 48 vNICs to the Cisco UCS M81KR Virtual Interface Card adapter.
- If you want to configure 150 vNICs and the server contains a Cisco UCS M81KR Virtual Interface Card adapter (128 vNICs) and a Cisco UCS VIC-1240 Virtual Interface Card adapter (256 vNICs), Cisco UCS Manager assigns 50 vNICs to the Cisco UCS M81KR Virtual Interface Card adapter and 100 vNICs to the Cisco UCS VIC-1240 Virtual Interface Card adapter.



Note Exceptions to this implicit assignment occur if you configure the vNICs for fabric failover and if you configure dynamic vNICs for the server.

For a configuration that includes vNIC fabric failover where one adapter does not support vNIC failover, Cisco UCS Manager implicitly assigns all vNICs that have fabric failover enabled to the adapter that supports them. If the configuration includes only vNICs that are configured for fabric failover, no vNICs are implicitly assigned to the adapter that does not support them. If some vNICs are configured for fabric failover and some are not, Cisco UCS Manager assigns all failover vNICs to the adapter that supports them and a minimum of one nonfailover vNIC to the adapter that does not support them, according to the ratio above.

For a configuration that includes dynamic vNICs, the same implicit assignment would occur. Cisco UCS Manager assigns all dynamic vNICs to the adapter that supports them. However, with a combination of dynamic vNICs and static vNICs, at least one static vNIC is assigned to the adapter that does not support dynamic vNICs.

Creating a vNIC/vHBA Placement Policy

Procedure

-
- Step 1** In the **Navigation** pane, click **Servers**.
- Step 2** Expand **Servers > Policies**.
- Step 3** Expand the node for the organization where you want to create the policy.
If the system does not include multi tenancy, expand the **root** node.
- Step 4** Right-click **vNIC/vHBA Placement Policies** and choose **Create Placement Policy**.
- Step 5** In the **Create Placement Policy** dialog box, do the following:
- Complete the following fields:

Name	Description
Name field	The name for this placement policy. This name can be between 1 and 16 alphanumeric characters. You cannot use spaces or any special characters other than - (hyphen), _ (underscore), : (colon), and . (period), and you cannot change this name after the object is saved.

Name	Description
Virtual Slot Mapping Scheme field	<p>Cisco UCS assigns virtual network interface connections (vCons) to the PCIe adapter cards in the server. Each vCon is a virtual representation of a physical adapter that can be assigned vNICs and vHBAs.</p> <p>For blade or rack servers that contain one adapter, Cisco UCS assigns all vCons to that adapter. For servers that contain four adapters, Cisco UCS assigns vCon1 to Adapter1, vCon2 to Adapter2, vCon3 to Adapter3, and vCon4 to Adapter4.</p> <p>For blade or rack servers that contain two or three adapters, Cisco UCS assigns the vCons based on the selected virtual slot mapping scheme. This can be one of the following:</p> <ul style="list-style-type: none"> • Round Robin— In a server with two adapter cards, Cisco UCS assigns vCon1 and vCon3 to Adapter1, then assigns vCon2 and vCon4 to Adapter2. <p>In a server with three adapter cards, Cisco UCS assigns vCon1 to Adapter1, vCon2 and vCon4 to Adapter2, and vCon3 to Adapter3.</p> <p>This is the default scheme.</p> <ul style="list-style-type: none"> • Linear Ordered— In a server with two adapter cards, Cisco UCS assigns vCon1 and vCon2 to Adapter1, then assigns vCon3 and vCon4 to Adapter2. <p>In a server with three adapter cards, Cisco UCS assigns vCon1 to Adapter1 and vCon2 to Adapter2, then assigns vCon3 and vCon4 to Adapter3.</p> <p>Note</p> <p>If the blade two adapters are numbered left to right while vCons are numbered right to left. If one of these blade servers has a single adapter, Cisco UCS assigns all vCons to that adapter. If the server has two adapters, the vCon assignment depends upon the virtual slot mapping scheme:</p> <ul style="list-style-type: none"> • Round Robin—Cisco UCS assigns vCon2 and vCon4 to Adapter1 and vCon1 and vCon3 to Adapter2. This is the default. • Linear Ordered—Cisco UCS assigns vCon3 and vCon4 to Adapter1 and vCon1 and vCon2 to Adapter2. <p>After Cisco UCS assigns the vCons, it assigns the vNICs and vHBAs based on the Selection Preference for each vCon.</p>

- b) In the **Selection Preference** column for each **Virtual Slot**, choose one of the following from the drop-down list:

- **All**—All configured vNICs and vHBAs can be assigned to the vCon, whether they are explicitly assigned to it, unassigned, or dynamic. This is the default.

- **Assigned Only**—vNICs and vHBAs must be explicitly assigned to the vCon. You can assign them explicitly through the service profile or the properties of the vNIC or vHBA.
- **Exclude Dynamic**—Dynamic vNICs and vHBAs cannot be assigned to the vCon. The vCon can be used for all static vNICs and vHBAs, whether they are unassigned or explicitly assigned to it.
- **Exclude Unassigned**—Unassigned vNICs and vHBAs cannot be assigned to the vCon. The vCon can be used for dynamic vNICs and vHBAs and for static vNICs and vHBAs that are explicitly assigned to it.
- **Exclude usNIC**—Cisco usNICs cannot be assigned to the vCon. The vCon can be used for all other configured vNICs and vHBAs, whether they are explicitly assigned to it, unassigned, or dynamic.

Note

An SRIOV usNIC that is explicitly assigned to a vCon set to **Exclude usNIC** will remain assigned to that vCon.

- c) Click **OK**.

Deleting a vNIC/vHBA Placement Policy

Procedure

-
- Step 1** In the **Navigation** pane, click **Servers**.
 - Step 2** Expand **Servers > Policies > Organization_Name**.
 - Step 3** Expand the **vNIC/vHBA Placement Policies** node.
 - Step 4** Right-click the policy you want to delete and choose **Delete**.
 - Step 5** If a confirmation dialog box displays, click **Yes**.
-

Explicitly Assigning a vNIC to a vCon

Before you begin

Configure the vCons through a vNIC/vHBA placement policy or in the service profile with one of the following values:

- **Assigned Only**
- **Exclude Dynamic**
- **Exclude Unassigned**

If a vCon is configured for **All**, you can explicitly assign a vNIC or vHBA to that vCon. However, there is less control with this configuration.

Procedure

- Step 1** In the **Navigation** pane, click **Servers**.
- Step 2** Expand **Servers > Service Profiles**.
- Step 3** Expand the node for the organization which contains the service profile whose vNICs you want to explicitly assign to a vCon.

If the system does not include multi tenancy, expand the **root** node.

- Step 4** Expand *Service_Profile_Name* > **vNICs**.
- Step 5** Click on the vNIC that you want to explicitly assign to a vCon.
- Step 6** In the **Work** pane, click the **General** tab.
- Step 7** In the **Virtual Host Interface Placement** section, complete the following fields:

Name	Description
Desired Placement drop-down list	The user-specified virtual network interface connection (vCon) placement for the vNIC. This can be one of the following: <ul style="list-style-type: none"> • Any—Allows Cisco UCS Manager to determine the vCon to which the vNIC is assigned. • 1—Explicitly assigns the vNIC to vCon1. • 2—Explicitly assigns the vNIC to vCon2. • 3—Explicitly assigns the vNIC to vCon3. • 4—Explicitly assigns the vNIC to vCon4.
Actual Assignment field	The actual vCon assignment of the vNIC on the server.

If you attempt to assign a vNIC to a vCon that is not configured for that type of vNIC, Cisco UCS Manager displays a message box to advise you of the configuration error. You must either assign the vNIC to another vCon or change the vCon configuration in the service profile.

- Step 8** In the **Order** section, complete the following fields:

Name	Description
Desired Order field	The user-specified PCI order for the vNIC. Enter an integer between 0 and 128. You cannot create more than 128 vNICs for a server.
Actual Order field	The actual PCI order of the vNIC on the server.

- Step 9** Click **Save Changes**.

Explicitly Assigning a vHBA to a vCon

Before you begin

Configure the vCons through a vNIC/vHBA placement policy or in the service profile with one of the following values:

- **Assigned Only**
- **Exclude Dynamic**
- **Exclude Unassigned**

If a vCon is configured for **All**, you can explicitly assign a vNIC or vHBA to that vCon. However, there is less control with this configuration.

Procedure

-
- Step 1** In the **Navigation** pane, click **Servers**.
- Step 2** Expand **Servers > Service Profiles**.
- Step 3** Expand the node for the organization which contains the service profile whose vHBAs you want to explicitly assign to a vCon.
- If the system does not include multi tenancy, expand the **root** node.
- Step 4** Expand *Service_Profile_Name* > **vHBAs**.
- Step 5** Click on the vHBA that you want to explicitly assign to a vCon.
- Step 6** In the **Work** pane, click the **General** tab.
- Step 7** In the **Virtual Host Interface Placement** section, complete the following fields:

Name	Description
Desired Placement field	The user-specified virtual network interface connection (vCon) placement for the vHBA. This can be one of the following: <ul style="list-style-type: none"> • Any—Allows Cisco UCS Manager to determine the vCon to which the vHBA is assigned. • 1—Explicitly assigns the vHBA to vCon1. • 2—Explicitly assigns the vHBA to vCon2. • 3—Explicitly assigns the vHBA to vCon3. • 4—Explicitly assigns the vHBA to vCon4.
Actual Assignment field	The actual vCon assignment of the vHBA on the server.

If you attempt to assign a vHBA to a vCon that is not configured for that type of vHBA, Cisco UCS Manager displays a message box to advise you of the configuration error. You must either assign the vHBA to another vCon or change the vCon configuration in the service profile.

Step 8 In the **Order** section, complete the following fields:

Name	Description
Desired Order field	The user-specified PCI order for the vHBA. Enter an integer between 0 and 128. You cannot create more than 128 vHBAs for a server.
Actual Order field	The actual PCI order of the vHBA on the server.

Step 9 Click **Save Changes**.

Placing Static vNICs Before Dynamic vNICs

For optimal performance, static vNICs and vHBAs should be placed before dynamic vNICs on the PCIe bus. Static vNICs refer to both static vNICs and vHBAs. Cisco UCS Manager Release 2.1 provides the following functionality regarding the order of static and dynamic vNICs:

- After upgrading to Cisco UCS Manager Release 2.1, if no change is made to existing service profiles (profiles that are defined in releases prior to Cisco UCS Manager Release 2.1), the vNIC order does not change.
- After an upgrade to Cisco UCS Manager Release 2.1, any vNIC-related change would reorder the vNIC map. As a result, all dynamic vNICs would be placed after the static vNICs.
- For newly created service profiles in Cisco UCS Manager Release 2.1, static vNICs are always ordered before dynamic vNICs.
- The above behavior is independent of the sequence of creating or deleting static or dynamic vNICs.
- For SRIOV-enabled service profiles, UCSM places the vNIC Physical Function(PF) before the corresponding Virtual Functions (VFs). This scheme guarantees that the VFs are placed close to the parent PF vNIC on the PCIe bus and BDFs are in successive incremental order for the VFs.

Example

Beginning Device Order in Cisco UCS Manager Release 2.0:

```
dyn-vNIC-1 1
dyn-vNIC-2 2
```

New Device Order in Cisco UCS Manager Release 2.0 (Add 2 static vNICs):

```
dyn-vNIC-1 1
dyn-vNIC-2 2
eth-vNIC-1 3
eth-vNIC-2 4
```

After upgrading to Cisco UCS Manager Release 2.1, (Before any vNIC-related change is made to the service profile.)

```

dyn-vNIC-1 1
dyn-vNIC-2 2
eth-vNIC-1 3
eth-vNIC-2 4

```

New Device Order in Cisco UCS Manager Release 2.1 (Add 2 dynamic vNICs by changing the policy count from 2 to 4.)

```

dyn-vNIC-1 3
dyn-vNIC-2 4
eth-vNIC-1 1
eth-vNIC-2 2
dyn-vNIC-3 5
dyn-vNIC-4 6

```

Dynamic vNICs as Multifunction PCIe Devices

Cisco UCS Manager Version 2.1 provisions static vNICs as 0-function devices (new BUS for every static vNIC). Multifunction dynamic vNICs are placed from the new Bus-slot after the last static vNIC/vHBA.



Note Cisco UCS Manager Version 2.1 supports the new StaticZero mode.

Table 11: Version Compatibility

Cisco UCS Manager		
Version 1.4	Version 2.0	Version 2.1
Scheme: ZeroFunction	Scheme: ZeroFunction / MultiFunction	Scheme: ZeroFunction / MultiFunction / StaticZero
Static and Dynamic vNICs are all on Bus [0-57], Function [0] < ZeroFunction Mode >	Static vNICs and Dynamic vNICs are on Bus [0-57], Function [0-7]. Bus 0, Function 0 Bus 0, Function 7 Bus 1, Function 0 < MultiFunction Mode >	Static vNICs or PFs will be on Bus [0-57], Function [0]. SRIOV: Corresponding VFs will be on the same Bus and Functions [1-255] No-SRIOV: Dynamic vNICs are on Bus [0-57], Function [0-7] < StaticZero Mode >
	Upgrade from Balboa will not renumber BDFs (remain in ZeroFunction mode) until Bus <= 57. Once devices exceed 58, switch to MultiFunction mode.	Upgrade from Balboa will not renumber BDFs (remain in ZeroFunction mode) until Bus <=57. Once devices exceed 58 or Platform specific maximum PCIe Bus number or change to SRIOV configuration, switch to StaticZero mode.

Cisco UCS Manager		
Version 1.4	Version 2.0	Version 2.1
Scheme: ZeroFunction	Scheme: ZeroFunction / MultiFunction	Scheme: ZeroFunction / MultiFunction / StaticZero
		Upgrade from Cisco UCS Manager Version 2.0 will not renumber BDFs (remain in ZeroFunction / MultiFunction mode). Once devices exceed 58 or Platform specific maximum PCIe Bus number OR Change to SRIOV configuration, switch to StaticZero mode.

vNIC/vHBA Host Port Placement

After a vNIC/vHBA is assigned to a vCON, it can be placed on one of the host ports of a specific adapter. You can either explicitly specify the host port for placement, or allow Cisco UCS Manager to automatically assign vNICs/vHBAs to host ports.



Note You can perform vNIC/vHBA host port placement on servers that support Cisco UCS VIC 1340/1380 adapters.

Cisco UCS 1300 Series adapters have 2x8 PCIe third generation host ports. Each PCIe host port is capable of a maximum of 64 Gbps bandwidth.

The host port placement of the vNIC/vHBA determines the order of the vNIC/vHBA on the adapter. The vNICs/vHBAs placed on the first host port will be enumerated first, followed by the vNICs/vHBAs on the second host port.



Note The 64 Gbps maximum is a theoretical maximum, actual data transfer is limited to around 40 Gbps.

All the vNICs sharing the same PCIe Host Port will share this bandwidth. To make the optimal use of PCIe host port bandwidth, vNICs should be distributed across the two host ports.

Configuring Host Port Placement

You can configure host port placement for vNICs on servers that support Cisco UCS VIC 1340 and VIC 1380 adapters.

Procedure

- Step 1** In the **Navigation** pane, click **Servers**.
- Step 2** Expand **Servers** > **Service Profiles**.

- Step 3** Select the service profile which is associated with the vNIC that you want to place on a host port.
- Step 4** Expand *Service_Profile_Name* > vNICs
- Step 5** Under the **Network** tab, in the vNICs summary table, double-click the **Admin Host Port** value of the vNIC which you want to configure and select one of the following:
- **Any**—Allows Cisco UCS Manager to determine the host port to which the vNIC is assigned.
 - **1**—Explicitly assigns the vNIC to host port 1.
 - **2**—Explicitly assigns the vNIC to host port 2.
- Actual Host Port** displays the actual assignment of the vNIC on a host port. When this feature is not supported, this will appear as **None**.
- Step 6** Click **Save Changes**.
-

CIMC Mounted vMedia

Using Scriptable vMedia

Cisco UCS Manager allows provisioning of vMedia devices iso images for remote UCS servers. Using Scriptable vMedia, you can programmatically mount an IMG or an ISO image on a remote server. CIMC mounted vMedia provide communications between other mounted media inside your datacenter with no additional requirements media connection. Scriptable vMedia allows you to control virtual media devices without using a browser to manually map each UCS server individually.

Scriptable vMedia supports multiple share types including NFS, CIFS, HTTP, and HTTPS shares. **Scriptable vMedia** is enabled through BIOS configuration and configured through a Web GUI and CLI interface.

Cisco UCS Manager Scriptable vMedia supports the following functionality:

- Booting from a specific vMedia device
- Copying files from a mounted share to a local disk
- Installation and updating OS drivers



Note Cisco UCS Manager support for Scriptable vMedia is applicable for CIMC mapped devices only. Existing KVM based vMedia devices are not supported.

vMedia mount fails when the following conditions are met:

1. The remote vMedia image filename in the vMedia policy is set to **Service-Profile-Name**.
2. The service profile is renamed.

This is because the change in the name of the service profile does not change the remote vMedia image filename in the vMedia policy. The image filename still points to the older image on the remote device, which cannot be found.

Creating a vMedia Policy

A vMedia policy is used to configure the mapping information for remote vMedia devices. Two vMedia devices and mappings for CD and HDD are allowed in a vMedia policy. You can configure one ISO and one IMG at a time. ISO configurations maps to a CD drive and IMG configurations maps to a HDD device.



Note If you want to map a device to a remote folder, you must create an IMG and map it as a HDD device.

Before you begin

Make sure that you have access to the following:

- Remote vMedia Server
- vMedia Devices

Procedure

-
- Step 1** In the **Navigation** pane, click **Servers**.
- Step 2** Expand **Servers > Policies**.
- Step 3** Expand the node for the organization where you want to create the policy.
If the system does not include multi tenancy, expand the **root** node.
- Step 4** Right-click **vMedia Policies** and select **Create vMedia Policy**.
- Step 5** In the **Create vMedia Policy** dialog box, complete the following fields:

Name	Description
Name	The name of the vMedia policy. This name can be between 1 and 16 alphanumeric characters. You cannot use spaces or any special characters other than - (hyphen), _ (underscore), : (colon), and . (period), and you cannot change this name after the object is saved.
Description	A description of the policy. We recommend including information about where and when the policy should be used. Maximum 115 characters.

Name	Description
Retry on Mount Failure	<p>Designates if the vMedia will continue mounting when a mount failure occurs. This can be:</p> <ul style="list-style-type: none"> • Yes • No <p>Note The default setting is Yes. When Yes is selected the remote server will continue to try to mount the vMedia mount process until it is successful or you disable this option. If you select No, a warning message will appear indicating retry on mount failure will not work in case of mount failure.</p>

Step 6 On the icon bar to the right of the table, click +.

Step 7 In the **Create vMedia Mount** dialog box, complete the following fields:

Name	Description
Name	<p>Name of the vMedia Mount policy.</p> <p>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.</p>
Device Type	<p>The type of remote vMedia you plan to mount. This can be:</p> <ul style="list-style-type: none"> • CDD—Scriptable vMedia CD. • HDD—Scriptable vMedia HDD.
Protocol	<p>The protocol to use when communicating with the remote server. Click one of the following radio buttons to indicate the protocol you want to use to communicate with the mounted remote server. This can be:</p> <ul style="list-style-type: none"> • NFS - Network Files System. • CIFS - Common Internet File System. • HTTP - Hypertext Transfer Protocol. • HTTPS - Hypertext Transfer Protocol over Secure.

Name	Description
Authentication Protocol	<p>The protocol to use for authentication when you use CIFS as the protocol for communicating with the remote server. When you use any protocol other than CIFS, this field is not available. Select one of the following from the drop-down list to specify the authentication protocol.</p> <ul style="list-style-type: none"> • Default—NT LAN Manager Security Support Provider (NTLMSSP) protocol. Use this option only with Windows 2008 R2 and Windows 2012 R2. • None—No authentication is used • Ntlm—NT LAN Manager (NTLM) security protocol. Use this option only with Windows 2008 R2 and Windows 2012 R2. • Ntlmi—NTLMI security protocol. Use this option only when you enable Digital Signing in the CIFS Windows server. • Ntlmssp—NT LAN Manager Security Support Provider (NTLMSSP) protocol. Use this option only with Windows 2008 R2 and Windows 2012 R2. • Ntlmsspi—NTLMSSPi protocol. Use this option only when you enable Digital Signing in the CIFS Windows server. • Ntlmv2—NTLMv2 security protocol. Use this option only with Samba Linux. • Ntlmv2i—NTLMv2i security protocol. Use this option only with Samba Linux. <p>Note The authentication protocol options are available only when you select CIFS as the protocol. For all other protocols, the Authentication Protocol field is disabled.</p>
Hostname/IPAddress	<p>Enter the IP address or hostname of the location where the backup file is to be stored. This can be a server, storage array, local drive, or any read/write media that the fabric interconnect can access through the network.</p> <p>If you use a hostname, you must configure Cisco UCS Manager to use a DNS server. The hostname (DNS) can be used when Inband network is configured for that server.</p>

Name	Description
Image Name Variable	<p>The name to be used for the image. This can be:</p> <ul style="list-style-type: none"> • None—Filename must be entered in the Remote File field. • Service Profile Name—Filename automatically becomes the name of the service profile that the vMedia Policy is associated with. <p>Note</p> <ul style="list-style-type: none"> • If you select Service Profile Name as the Image Name variable, the Remote File field is disabled. • If you select Service Profile Name as the Image Name variable, do not rename the service profile. Renaming the service profile can result in vMedia mount failure.
Remote File	<p>Enter the full path to the ISO or other image file.</p> <p>Note</p> <p>Ensure that the full path to the file begins with “/” after the share name. This field can contain the filename [with the file extension] only.</p>
Remote Path	<p>Enter the share name on the remote server, for example “Share”.</p>
Username	<p>Enter the username that Cisco UCS Manager should use to log in to the remote server.</p> <p>This field does not apply if the protocol is NFS. This field is optional if the protocol is HTTP.</p>
Password	<p>Enter the password associated with the username.</p> <p>This field does not apply if the protocol is NFS. This field is optional if the protocol is HTTP.</p>
Remap on Eject	<p>Click this checkbox to remap mounted vMedia after it is ejected.</p>
Writable	<p>Click this checkbox to configure the vMedia mount as writable. If this checkbox is cleared, the vMedia mount remains read-only.</p> <p>vMedia mounts are read-only by default.</p> <p>You can configure a vMedia mount as writable only when both the following conditions are met:</p> <ul style="list-style-type: none"> • Device Type is HDD • Protocol is NFS or CIFS

Step 8

Click **OK**.

The remote server details are listed in the **vMedia Mounts** area of the **Create vMedia Mount** dialog box.

What to do next

Create a vMedia boot policy.

Adding a vMedia Policy to a Service Profile

Before you can use Scriptable vMedia, you must add the vMedia and Boot Policies to a Service Profile. After the vMedia and Boot Policies are added to a service profile you can associate the service profile with a Cisco UCS server. The following procedure describes how to add a vMedia policy to a Service Profile.

Before you begin

Configure the vMedia Policy you want to add to a service profile.

Procedure

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

Step 2 Expand **Servers > Service Profiles**.

Step 3 Expand the node for the organization where you want to create the service profile.

If the system does not include multi tenancy, expand the **root** node.

Step 4 Right-click the organization and select **Create Service Profile (expert)**.
The Unified Computing System Manager pane displays.

Step 5 In the **Name** field, enter a unique name that you can use to identify the service profile.

This name can be between 2 and 32 alphanumeric characters. You cannot use spaces or any special characters other than - (hyphen), _ (underscore), : (colon), and . (period), and this name must be unique across all service profiles and service profile templates within the same organization.

This name must be unique within the organization or sub-organization in which you are creating the service profile.

Step 6 From the **UUID Assignment** drop-down list, do one of the following:

Option	Description
Select (pool default used by default)	Assigns a UUID from the default UUID Suffix pool. Continue with Step 8.
Hardware Default	Uses the UUID assigned to the server by the manufacturer. If you choose this option, the UUID remains unassigned until the service profile is associated with a server. At that point, the UUID is set to the UUID value assigned to the server by the manufacturer. If the service profile is later moved to a different server, the UUID is changed to match the new server. Continue with Step 8.

Option	Description
XXXXXXXX-XXXX-XXXX-XXXX-XXXXXXXXXXXX	Uses the UUID that you manually assign. Continue with Step 7.
Pools <i>Pool_Name</i>	Assigns a UUID from the UUID Suffix pool that you select from the list at the bottom of the drop-down list. Each pool name is followed by two numbers in parentheses that show the number of UUIDs still available in the pool and the total number of UUIDs in the pool. If you do not want use any of the existing pools, but instead want to create a pool that all service profiles can access, continue with Step 4. Otherwise, continue with Step 8.

Step 7 (Optional) If you selected the **XXXXXXXX-XXXX-XXXX-XXXX-XXXXXXXXXXXX** option, do the following:

- a) In the **UUID** field, enter the valid UUID that you want to assign to the server which uses this service profile.
- b) To verify that the selected UUID is available, click the **here** link.

Step 8 (Optional) If you want to create a new UUID Suffix pool to use to use in this service profile, click **Create UUID Suffix Pool** and complete the fields in the **Create UUID Suffix Pool** wizard.

Step 9 (Optional) In the text box, enter a description of this service profile.

The user-defined description for this service profile.

Enter up to 256 characters. You can use any characters or spaces except ` (accent mark), \ (backslash), ^ (carat), " (double quote), = (equal sign), > (greater than), < (less than), or ' (single quote).

Step 10 Click **Next**.

Step 11 From the **vMedia** drop down list, choose one of the following:

Option	Description
Select vMedia Policy to use	Enables you to assign a vMedia policy to this service profile. Continue with Step 12.
Create a Specific vMedia Policy	Enables you to create a local vMedia policy that can only be accessed by this service profile.
vMedia Policies <i>Policy_Name</i>	Assigns an existing vMedia policy to the service profile. If you choose this option, Cisco UCS Manager displays the details of the policy. If you do not want use any of the existing policies but instead want to create a policy that all service profiles can access, click Create vMedia Policy . Otherwise, choose a policy from the list and continue with Step 13.

- Step 12** If you created a new vmedia policy accessible to all service profiles and template, choose that policy from the **vMedia** drop down list .
- Step 13** Click **Next**.
-

Viewing CIMC vMedia Policy

Before you begin

vMedia Policies are configured.

Procedure

- Step 1** In the **Navigation** pane, click **Servers**.
- Step 2** Expand **Policies > vMedia Policies**.
- Step 3** Expand the **vMedia Policies** node to view the list of **vMedia Policies**.
- Step 4** Double-click the name of a vMedia policy to view the properties for the selected **vMedia Mount**.
On the **Properties** page, you can modify the properties used for the **vMedia Mounts**.
-



CHAPTER 13

Firmware Upgrades

- [Firmware Upgrades, on page 333](#)
- [Verifying Firmware Versions on Components, on page 333](#)

Firmware Upgrades

Beginning with Cisco UCS Manager Release 4.1(1), Cisco is releasing unified Cisco UCS Manager software and firmware upgrades for the following platforms with every release of Cisco UCS Manager:

- Cisco UCS 6400 Series Fabric Interconnect with Cisco UCS B-Series, C-Series, and S-Series servers.
- Cisco UCS 6300 Series Fabric Interconnect with Cisco UCS B-Series, C-Series, and S-Series servers.
- Cisco UCS 6200 Series Fabric Interconnect with Cisco UCS B-Series, C-Series, and S-Series servers.
- Cisco UCS 6324 Fabric Interconnect with Cisco UCS B-Series Servers and C-Series Servers, which is also known as UCS Mini.

You can upgrade the firmware through Auto Install, packages in service profiles, using the firmware automatic synchronization server policy, and directly at endpoints. For more information on guidelines and installing firmware, see the *Cisco UCS Firmware Management Guide*.

Verifying Firmware Versions on Components

Procedure

- Step 1** In the **Navigation** pane, click **Equipment**.
 - Step 2** In the **Work** pane, click the **Firmware Management** tab.
 - Step 3** On the **Installed Firmware** tab, review the firmware versions listed for each component.
For more information on guidelines and installing firmware, see the *Cisco UCS Firmware Management Guide*.
-



CHAPTER 14

Diagnostics Configuration

- [Overview of Cisco UCS Manager Diagnostics, on page 335](#)
- [Creating a Diagnostics Policy, on page 335](#)
- [Diagnostics Test on a Blade Server, on page 336](#)
- [Diagnostics Test on a Rack Server, on page 337](#)
- [Starting a Diagnostics Tests on All Servers, on page 338](#)
- [Stopping a Diagnostics Tests on All Servers, on page 339](#)
- [Viewing the Server Diagnostics Status/Result, on page 339](#)
- [Diagnostics Troubleshooting, on page 340](#)

Overview of Cisco UCS Manager Diagnostics

The Cisco UCS Manager diagnostics tool enables you to verify the health of the hardware components on your servers. The diagnostics tool provides a variety of tests to exercise and stress the various hardware subsystems on the servers, such as memory and CPU. You can use the tool to run a sanity check on the state of your servers after you fix or replace a hardware component. You can also use this tool to run comprehensive burn-in tests before you deploy a new server in your production environment.

When a system is new, a default diagnostics policy is created in org scope. This default policy is named default and it cannot be deleted. The user will receive an error message if they try to delete it. The default diagnostic policy is the preferred way to execute the same set of tests across all servers. Any diagnostic policy, including the default can be customized.

The default policy only has one memory test. The default parameters of the memory test can be modified. In addition, the memory test within the default diagnostics policy can be deleted. If it does not have a memory test, the diagnostic policy will not run.

Creating a Diagnostics Policy

Before you begin

You must have admin privileges to perform this task.

Procedure

Step 1 Navigate to **Servers > Policies > Diagnostics Policies**.

Step 2 Click **Add**.

Step 3 Complete the following fields:

Field	Description
Name	Name of the diagnostics policy. The character limit is 16.
Description	Description of the diagnostics policy. This is optional.

Step 4 Click **Next**.

Step 5 Click **Add**.

Step 6 Complete the following fields:

Name	Description
Order	The order in which the tests will be executed.
CPU Filter	Sets the CPU filter to all CPUs or to a specified CPU.
Loop Count	Sets the loop count to the specified iterations. The range is from 1-1000.
Memory Chunk Size	Sets the memory chunk to 5mb-chunk or big-chunk.
Memory Size	Sets the memory size to a specific value.
Pattern	Sets the memory test to butterfly, killer, prbs, prbs-addr, or prbs-killer.

Step 7 Click **OK**.

Step 8 Click **Finish**.

Diagnostics Test on a Blade Server

Starting a Diagnostics Test on a Blade Server

Before you begin

You must have admin privileges to perform this task.

Procedure

- Step 1** Navigate to **Equipment > Chassis > Server**.
 - Step 2** Choose the server for which you want to start the diagnostics test.
 - Step 3** Click on the **Diagnostics** tab.
 - Step 4** Click **Start**. Once the diagnostics test has started, the button will be grayed out.
-

Stopping a Diagnostics Test on a Blade Server

Procedure

- Step 1** Navigate to **Equipment > Chassis > Server**.
 - Step 2** Choose the server for which you want to stop the diagnostics test.
 - Step 3** Click on the **Diagnostics** tab.
 - Step 4** Click **Stop**. Once the diagnostic text has stopped, the button will be grayed out.
-

Diagnostics Test on a Rack Server

Starting a Diagnostics Test on a Rack Server

Diagnostics Test is available for C220 M5, C240 M5, and C480 M5/C480 M5 ML rack servers.

Before you begin

You must have admin privileges to perform this task.

Procedure

- Step 1** Navigate to **Equipment > Rack Mounts > Server**.
 - Step 2** Choose the server for which you want to start the diagnostics test.
 - Step 3** Click on the **Diagnostics** tab.
 - Step 4** Click **Start**. Once the diagnostics test has started, the button will be grayed out.
-

Stopping a Diagnostics Test on a Rack Server

Procedure

-
- Step 1** Navigate to **Equipment > Rack Mounts > Server**.
 - Step 2** Choose the server for which you want to stop the diagnostics test.
 - Step 3** Click on the **Diagnostics** tab.
 - Step 4** Click **Stop**. Once the diagnostic text has stopped, the button will be grayed out.
-

Starting a Diagnostics Tests on All Servers



Note Starting diagnostics testing all servers will cause a reboot of each individual server.

Before you begin

You must have admin privileges to perform this task.

Procedure

-
- Step 1** Navigate to **Equipment > Diagnostics**.
 - Step 2** Click **Start**. Once the diagnostics test has started, the link will be grayed out. In the **Diagnostic Result** table, you can view the following information:

Field	Description
Name	The system-defined server name.
Chassis ID	The unique identifier for the chassis. This numeric identifier is assigned based on the location of the chassis within the system. Note Not applicable for rack servers.
PID	The server model PID.
Overall Progress Percentage	A description of the overall progress percentage of the diagnostics test on the server.
Operation Status	A description of the diagnostics operation status of the server.

Note

If a server fails to run the diagnostic test, click on the server link and to view the error description under the **Diagnostics** tab. You can also view the faults generated in the **Faults** tab.

Stopping a Diagnostics Tests on All Servers

Before you begin

You must have admin privileges to perform this task.

Procedure

-
- Step 1** Navigate to **Equipment > Diagnostics**.
- Step 2** Click **Stop**. Once the diagnostics test has stopped, the link will be grayed out.
-

Viewing the Server Diagnostics Status/Result

Before you begin

You can run the diagnostic test on individual servers through CLI and view the status on this page.

Procedure

-
- Step 1** In the **Navigation** pane, click **Equipment**.
- Step 2** Expand **Equipment > Chassis > Servers**.
or for rack servers, Expand **Equipment > Rack Mounts > Server**
- Step 3** Choose the server for which you want to view the diagnostic status and then click the **Diagnostics** tab.
You can view the following information:

Name	Description
Diagnostic Policies	Enables the user to select a diagnostics policy and apply it to a specific server.
Start/Stop	Enables the user to start or stop a diagnostics test on a specific server
Operation State	The server's diagnostics operation status. Possible values are Idle, In-Progress, Completed, Failed, and Cancelled.

Name	Description
FSM Status Descr	A brief description of the current task in the server's diagnostics operation.
FSM Progress	The overall progress of the diagnostics operation being executed on the server.
Test Overall Progress	The overall progress of the diagnostics test.
Error Description	A description of the error returned from the diagnostics operation.

Table 12: Diagnostic Result

Name	Description
ID	The unique identifier associated with the test.
Test Type	The type of diagnostics test.
Status	The status of the test execution. Values are: Idle, In Progress, Completed, or Failed.
Description	The description of the diagnostics test run. Once the test is complete, it provides detailed descriptions of the result.
Result	The result of the diagnostics test. Values are Pass, Fail, or NA.
Progress Percentage	The progress percentage of the diagnostics test.

Diagnostics Troubleshooting

Issue	Steps to Debug
If the BIOS detects a bad DIMM, the DIMM is disabled and is not visible to the Diagnostics operation.	Refer to memory-related faults in addition to the diagnostics operation results.

If the DIMM blacklisting feature is enabled and a DIMM is blacklisted, it is not visible to the Diagnostics operation.	Refer to memory-related faults in addition to the diagnostics operation results.
The Diagnostics operation may not execute successfully, if the server has bad DIMMs which prevent the server from booting.	NA
The Diagnostics operation can fail, if an uncorrectable error causes a server reboot.	NA
A Diagnostics operation failure can occur if there are memory errors that cause the Diagnostics operation to hang.	NA

<p>The Diagnostics operation can be interrupted by external events, such as a managed endpoint failover or a critical UCSM process restart. In these cases, the Diagnostics operation is cancelled and the Memory Tests are marked as failed.</p>	<p>The failure is triggered by external events. Retry the Diagnostics operation.</p>
<p>A Memory test fails with the error: Uncorrectable errors detected.</p>	<p>Check for server faults under the Chassis/Server/Faults tab. See the SEL logs for the DIMM errors under the Chassis/Server/SEL Logs tab.</p>
<p>A Memory test failure needs further analysis.</p>	<p>See the diagnostics operation logs in following log file archive on the primary FI in the /workspace partition: diagnostics/diag_log_<system-name>_<timestamp>_<chassis-id>_<blade-id>.tar.gz</p> <p>See the analysis file: tmp/ServerDiags/MemoryPmem2.<id>/MemoryPmem2.analysis in the previously mentioned log file archive.</p> <p>Use the following command to find the diagnostics logs with the analysis files:</p> <pre># for file in `ls /workspace/diagnostics/*diag*`; do tar -tzvf \$file grep analysis && echo "IN " \$file; done</pre>