



Cisco Intersight Managed Mode Configuration Guide

First Published: 2017-08-13

Last Modified: 2024-09-03

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

THE SPECIFICATIONS AND INFORMATION REGARDING THE PRODUCTS IN THIS MANUAL ARE SUBJECT TO CHANGE WITHOUT NOTICE. ALL STATEMENTS, INFORMATION, AND RECOMMENDATIONS IN THIS MANUAL ARE BELIEVED TO BE ACCURATE BUT ARE PRESENTED WITHOUT WARRANTY OF ANY KIND, EXPRESS OR IMPLIED. USERS MUST TAKE FULL RESPONSIBILITY FOR THEIR APPLICATION OF ANY PRODUCTS.

THE SOFTWARE LICENSE AND LIMITED WARRANTY FOR THE ACCOMPANYING PRODUCT ARE SET FORTH IN THE INFORMATION PACKET THAT SHIPPED WITH THE PRODUCT AND ARE INCORPORATED HEREIN BY THIS REFERENCE. IF YOU ARE UNABLE TO LOCATE THE SOFTWARE LICENSE OR LIMITED WARRANTY, CONTACT YOUR CISCO REPRESENTATIVE FOR A COPY.

The Cisco implementation of TCP header compression is an adaptation of a program developed by the University of California, Berkeley (UCB) as part of UCB's public domain version of the UNIX operating system. All rights reserved. Copyright © 1981, Regents of the University of California.

NOTWITHSTANDING ANY OTHER WARRANTY HEREIN, ALL DOCUMENT FILES AND SOFTWARE OF THESE SUPPLIERS ARE PROVIDED "AS IS" WITH ALL FAULTS. CISCO AND THE ABOVE-NAMED SUPPLIERS DISCLAIM ALL WARRANTIES, EXPRESSED OR IMPLIED, INCLUDING, WITHOUT LIMITATION, THOSE OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE AND NON-INFRINGEMENT OR ARISING FROM A COURSE OF DEALING, USAGE, OR TRADE PRACTICE.

IN NO EVENT SHALL CISCO OR ITS SUPPLIERS BE LIABLE FOR ANY INDIRECT, SPECIAL, CONSEQUENTIAL, OR INCIDENTAL DAMAGES, INCLUDING, WITHOUT LIMITATION, LOST PROFITS OR LOSS OR DAMAGE TO DATA ARISING OUT OF THE USE OR INABILITY TO USE THIS MANUAL, EVEN IF CISCO OR ITS SUPPLIERS HAVE BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGES.

Any Internet Protocol (IP) addresses and phone numbers used in this document are not intended to be actual addresses and phone numbers. Any examples, command display output, network topology diagrams, and other figures included in the document are shown for illustrative purposes only. Any use of actual IP addresses or phone numbers in illustrative content is unintentional and coincidental.

All printed copies and duplicate soft copies of this document are considered uncontrolled. See the current online version for the latest version.

Cisco has more than 200 offices worldwide. Addresses and phone numbers are listed on the Cisco website at www.cisco.com/go/offices.

Cisco and the Cisco logo are trademarks or registered trademarks of Cisco and/or its affiliates in the U.S. and other countries. To view a list of Cisco trademarks, go to this URL: <https://www.cisco.com/c/en/us/about/legal/trademarks.html>. Third-party trademarks mentioned are the property of their respective owners. The use of the word partner does not imply a partnership relationship between Cisco and any other company. (1721R)

© 2020–2024 Cisco Systems, Inc. All rights reserved.



CONTENTS

PREFACE

[Communications, Services, Bias-free Language, and Additional Information](#) ix

CHAPTER 1

[About Intersight Managed Mode](#) 1

[About Intersight Managed Mode](#) 1

[Supported Hardware](#) 1

CHAPTER 2

[Setting Up Fabric Interconnects](#) 23

[Initial Fabric Interconnect Configuration](#) 23

[Configuring Fabric Interconnect-A Using the Console](#) 24

[Configuring Fabric Interconnect-B Using the Console](#) 26

[Configuring Fabric Interconnect-A Using the GUI](#) 27

[Configuring Fabric Interconnect-B Using the GUI](#) 30

[Fabric Interconnect Password Guidelines](#) 31

[Migrating to Cisco UCS 6500 Series Fabric Interconnects](#) 32

[Migrating from UCSX-I-9108-25G to UCSX-I-9108-100G IFM on an existing domain with Cisco UCS 6500 Series Fabric Interconnect and UCSX-9508 Chassis](#) 34

[Fabric Interconnect Views](#) 35

[Fabric Interconnects Table View](#) 35

[Fabric Interconnects Details View](#) 36

[Fabric Interconnects Inventory View](#) 39

[Fabric Interconnects Connections View](#) 41

[Fabric Interconnects UCS Domain Profile View](#) 41

[Fabric Interconnect Topology View](#) 42

[Fabric Interconnect Metrics View](#) 42

[Fabric Interconnect Actions](#) 42

[Configuring Traffic Mirroring \(SPAN\)](#) 45

CHAPTER 3	Chassis and FEX Lifecycle	49
	Chassis and Fabric Extender Discovery and Actions	49
	Chassis Table View	51
	Chassis Details View	52
	Chassis Inventory View	53
	Chassis Connections View	54
	Chassis Topology View	54
	Chassis Metrics View	54
	Chassis Actions	54
	Fabric Extender Details View	55
	Fabric Extender Inventory View	56
	Fabric Extender Connection View	56

CHAPTER 4	Server Lifecycle	59
	Server Discovery and Actions	59
	Servers Table View	64
	Server Details View	67
	Server Inventory View	68
	Server Topology View	72
	Server Metrics View	72
	Server Connectivity View	73
	Compliance with Hardware Compatibility List (HCL)	73

CHAPTER 5	Configuring UCS Domain Profiles	75
	About UCS Domain Profile	75
	Creating Domain Profile Template and Deriving Profiles	75
	Creating a UCS Domain Profile	77
	UCS Domain Profile Details	77

CHAPTER 6	Configuring Server Profiles	79
	Server Profiles	79
	Creating a UCS Server Profile	88
	UCS Server Profile Details	89

CHAPTER 7	Configuring UCS Chassis Profiles	93
	About UCS Chassis Profile	93
	Creating Chassis Profile Template and Deriving Profiles	94
	Creating a Chassis Profile	95
	UCS Chassis Profile Details	95

CHAPTER 8	Configuring UCS Domain Policies	97
	Domain Policies	97
	Creating a Port Policy	100
	Creating an Ethernet Network Group Policy	108
	Creating an Ethernet Network Control Policy	110
	Creating a VLAN Policy	112
	Creating a VSAN Policy	114
	Creating an NTP Policy	115
	Creating a Network Connectivity Policy	116
	Creating an SNMP Policy	118
	Creating a System QoS Policy	120
	Creating a Syslog Policy	122
	Creating a Switch Control Policy	123
	Creating a Flow Control Policy	129
	Creating a Link Aggregation Policy	130
	Creating a Link Control Policy	131
	Creating a Multicast Policy	133

CHAPTER 9	Configuring Server Policies	135
	Server Policies	136
	Creating a Policy	143
	Supported UCS Server Policies	143
	Default Server Policy Configurations	147
	Creating a Certificate Management Policy	148
	Creating an Adapter Configuration Policy	149
	Creating a LAN Connectivity Policy	152
	Creating an Ethernet Adapter Policy	160

Creating an Ethernet QoS Policy	169
Creating an Ethernet Network Policy	170
Creating an Ethernet Network Group Policy	175
Creating an Ethernet Network Control Policy	176
Creating a SAN Connectivity Policy	178
Creating a Fibre Channel Adapter Policy	184
Creating a Fibre Channel Network Policy	187
Creating a Fibre Channel QoS Policy	188
Create FC Zone Policy	189
Creating a Firmware Policy	191
Creating a BIOS Policy	192
Creating a Boot Order Policy	193
Configuring an iSCSI Boot Policy	202
Creating an iSCSI Adapter Policy	205
Creating an iSCSI Static Target Policy	206
Creating a Device Connector Policy	207
Creating a Drive Security Policy	208
Creating a Disk Group Policy	209
Creating an IMC Access Policy	212
Creating an IPMI Over LAN Policy	214
Creating an LDAP Policy	216
Creating a Local User Policy	220
Creating an NTP Policy	223
Creating an SD Card Policy	224
Create a Serial Over LAN Policy	226
Create SSH Policy	228
Creating a Virtual KVM Policy	229
Creating a Virtual Media Policy	230
Creating a Network Connectivity Policy	234
Creating a SMTP Policy	236
Creating an SNMP Policy	237
Creating a Storage Policy	240
Creating a Syslog Policy	251
Creating a Power Policy for Server	252

Creating a Thermal Policy for Server 254

Creating vNIC or vHBA Templates 255

CHAPTER 10**Configuring UCS Chassis Policies 259**

Chassis Policies 259

Creating an IMC Access Policy 260

Creating an SNMP Policy 261

Creating a Power Policy for Chassis 263

Creating a Thermal Policy 266

CHAPTER 11**Configuring Pools 269**

Pools 269

Identity (ID) Pools 269

Pool Allocation 270

Deleting Pools 270

Identity Retention 271

IP Pools 272

 Creating an IP Pool 272

 IP Pool Details 274

MAC Pools 275

 Creating a MAC Pool 275

 MAC Pool Details 276

UUID Pools 277

 Creating a UUID Pool 278

 UUID Pool Details 278

WWN Pools 279

 Creating a WWNN Pool 280

 WWNN Pool Details 280

 Creating a WWPN Pool 281

 WWPN Pool Details 282

IQN Pools 282

 Creating an IQN Pool 283

 IQN Pool Details 283

Resource Pools 285

Creating a Resource Pool 285
 Resource Pool Details 286
 Virtual Routing and Forwarding 288
 Creating a VRF Instance 288

CHAPTER 12 **Managing the Device Console** 291
 Device Console 291

CHAPTER 13 **Managing Firmware** 293
 Firmware Upgrade in a Cisco UCS Domain through Intersight 293
 Upgrading Fabric Interconnect Firmware 295
 Upgrading Server Firmware 297
 Upgrading Server Firmware Using Firmware Policy 298
 Upgrades and Replacement of RMA Servers and Fabric Interconnects 298

CHAPTER 14 **Managing Technical Support** 303
 Integration with Cisco TAC 303
 Tech Support Diagnostic File Collection 304



Communications, Services, Bias-free Language, and Additional Information

- To receive timely, relevant information from Cisco, sign up at [Cisco Profile Manager](#).
- To get the business impact you're looking for with the technologies that matter, visit [Cisco Services](#).
- To submit a service request, visit [Cisco Support](#).
- To discover and browse secure, validated enterprise-class apps, products, solutions and services, visit [Cisco Marketplace](#).
- To obtain general networking, training, and certification titles, visit [Cisco Press](#).
- To find warranty information for a specific product or product family, access [Cisco Warranty Finder](#).

Documentation Feedback

To provide feedback about Cisco technical documentation, use the feedback form available in the right pane of every online document.

Cisco Bug Search Tool

[Cisco Bug Search Tool](#) (BST) is a web-based tool that acts as a gateway to the Cisco bug tracking system that maintains a comprehensive list of defects and vulnerabilities in Cisco products and software. BST provides you with detailed defect information about your products and software.

Bias-Free Language

The documentation set for this product strives to use bias-free language. For purposes of this documentation set, bias-free is defined as language that does not imply discrimination based on age, disability, gender, racial identity, ethnic identity, sexual orientation, socioeconomic status, and intersectionality. Exceptions may be present in the documentation due to language that is hardcoded in the user interfaces of the product software, language used based on standards documentation, or language that is used by a referenced third-party product.



CHAPTER 1

About Intersight Managed Mode

- [About Intersight Managed Mode, on page 1](#)
- [Supported Hardware, on page 1](#)

About Intersight Managed Mode

Cisco Intersight™ is a management platform delivered as a service with embedded analytics for your Cisco and 3rd party IT infrastructure. Intersight Managed Mode (IMM) is a new architecture that manages the UCS Fabric Interconnected systems through a Redfish-based standard model. Intersight Managed Mode unifies the capabilities of the UCS Systems and the cloud-based flexibility of Intersight, thus unifying the management experience for the standalone and Fabric Interconnect attached systems. Intersight Management Model standardizes policy and operation management for UCS-FI-6454, UCS-FI-64108, UCS-FI-6536, UCSX-S9108-100G Fabric Interconnects and Cisco UCS B-Series (M5, M6), Cisco UCS C-Series (M5, M6, M7, M8), and Cisco UCS X-Series (M6, M7) servers.

You can choose between the native UCSM Managed Mode (UMM) or Intersight Managed Mode (IMM) for the Fabric attached UCS Systems during initial setup of the Fabric Interconnects. If you choose to switch back between UMM and IMM, you must erase the present configuration and start from initial setup.



Note Before erasing the configuration, you must ensure to unclaim the device from Intersight and decommission all rack servers.

- Before you set up Intersight Managed Mode, please review the system requirements, supported hardware and software, and the steps required to migrate from UMM to IMM.
- For latest updates on Intersight features and functionality, see [Help Center](#).
- Servers in IMM mode require a minimum of Essentials license.

Supported Hardware

Supported Hardware for Intersight Managed Mode

This section includes the supported hardware for Intersight Managed Mode.

Table 1 lists the hardware components and the minimum required infrastructure firmware version.

Table 2 includes the supported hardware components along with the supported server and infrastructure firmware versions.

Table 3 shows the supported combination of components.

**Note**

- The Intersight Managed Mode (IMM) supports up to 20 chassis and 160 servers.
- Cisco UCS 6454 and 64108 Fabric Interconnects, require the port-based licensing in IMM but will not be enforced until further notice.
Beginning with UCS software release version 4.2(3), the Cisco UCS 6536 Fabric Interconnect supports a perpetual software license. This license activates all ports and software features of the Fabric Interconnect.
- In IMM, after discovery of a rack server, online swapping of cables on rack network adapters between Fabric Interconnects is not supported.
- The minimum supported Infrastructure firmware version for Intersight Managed Mode is 4.1(3).

Table 1: Supported Hardware Components with Required Minimum Infrastructure Versions

Fabric Components				
Model	Component	Sub-Component	Sub-Component Model	Minimum Infrastructure Firmware Versions
UCSX-S9108-100G	Fabric Interconnect			4.3(4.240078)
UCS-FI-6454	Fabric Interconnect			4.1(3b)
UCS-FI-64108	Fabric Interconnect			4.1(3b)
UCS-FI-6536	Fabric Interconnect			4.2(2b)
N20-C6508	Chassis			4.1(3b)
UCSB-5108-AC2		Input/Output Module (IOM)	UCS-IOM-2204XP	4.1(3b)
			UCS-IOM-2208XP	
			UCS-IOM-2408	
		UCS-IOM-2304	4.2(3c)	
		UCS-IOM-2304V2		

Fabric Components				
Model	Component	Sub-Component	Sub-Component Model	Minimum Infrastructure Firmware Versions
UCSX-9508	Chassis			4.2(1e)
		X-Fabric Modules (XFM)	UCSX-F-9416	4.2(2a)
		Intelligent Fabric Module (IFM)	UCSX-I-9108-25G	4.2(1e)
			UCSX-I-9108-100G	4.2(2a)
Cisco Nexus 2232PP	Fabric Extender (FEX)			4.1(3b)
N9K-C93108YC-FX3	Fabric Extender (FEX)			4.2(2a)

Table 2: Supported Hardware Components with Required Minimum Firmware Versions

Fabric Components					
Model	Component	Sub-Component	Sub-Component Model	Minimum Infrastructure Firmware Versions	Minimum Server Firmware Versions
UCSX-410C-M7	X-Series M7 Server			4.2(3e)	5.1(1.230052)
		CPU	4 th Generation Intel [®] Xeon [®] Scalable Processors Note For more information, see the Cisco UCS X410c M7 Compute Node Spec Sheet .	4.2(3e)	5.1(1.230052)
		DIMMs	UCSX-MRX16G1RE1 UCSX-MRX32G1RE1 UCSX-MRX64G2RE1 UCSX-MR128G4RE1 UCSX-MR256G8RE1	4.2(3e)	5.1(1.230052)
			UCSX-MRX96G2RF3	4.3(4.240066)	5.2(2.240053)
		Adapters	UCSX-ML-V5Q50G (Secure Boot) UCSX-ME-V5Q50G (Secure Boot) UCSX-ML-V5D200G	N/A	5.1(1.230052)
			UCSX-ML-V5D200GV2 (Secure Boot)	N/A	5.2(0.230061)
		Trusted Platform Module (TPM)	UCSX-TPM-002C	N/A	5.1(1.230052)
			UCS-TPM-002D	N/A	5.2(2.240053)
		Graphics processing unit (GPU)	UCSX-GPU-A16 UCSX-GPU-A40 UCSX-GPU-A100-80	N/A	5.1(1.230052)

Fabric Components					
Model	Component	Sub-Component	Sub-Component Model	Minimum Infrastructure Firmware Versions	Minimum Server Firmware Versions
			UCSX-GPU-H100-80 UCSX-GPU-L40 UCSX-GPU-L4 UCSX-GPU-FLEX140 UCSX-GPU-FLEX170	N/A	5.2(0.230127)
			UCSX-GPU-L40S	N/A	5.2(2.240053)
		Storage Controller	UCSX-M2-HWRAID UCSX-X10C-RAIDF	N/A	5.1(1.230052)
			UCSX-M2-PT-FPN	N/A	5.2(0.230127)

Fabric Components					
Model	Component	Sub-Component	Sub-Component Model	Minimum Infrastructure Firmware Versions	Minimum Server Firmware Versions
UCSX-210C-M7	X-Series M7 Server			4.2(3b)	5.1(0.230096)
		CPU	4 th Generation Intel [®] Xeon [®] Scalable Processors.	4.2(3b)	5.1(0.230096)
			5 th Generation Intel [®] Xeon [®] Scalable Processors Note For more information, see the Cisco UCS X210c M7 Compute Node Spec Sheet .	4.2(3b)	5.2(1.240010)
			Additional 5 th Generation Intel [®] Xeon [®] Scalable Processors: <ul style="list-style-type: none"> • UCSX-CPU-I4510T • UCSX-CPU-I4510 • UCSX-CPU-I4509Y • UCSX-CPU-I3508U 	4.2(3b)	5.2(2.240053)
		DIMMs	DDR5-4800MT/s: <ul style="list-style-type: none"> • UCSX-MRX16G1RE1 • UCSX-MRX32G1RE1 • UCSX-MRX64G2RE1 • UCSX-MR128G4RE1 • UCSX-MR256G8RE1 	4.2(3b)	5.2(1.240010)
	4.2(3b)		5.2(1.240010)		

Fabric Components					
Model	Component	Sub-Component	Sub-Component Model	Minimum Infrastructure Firmware Versions	Minimum Server Firmware Versions
			DDR5-5600MT/s: <ul style="list-style-type: none"> • UCSX-MRX16GIRE3 • UCSX-MRX32GIRE3 • UCSX-MRX48GIRF3 • UCSX-MRX64G2RE3 • UCSX-MRX96G2RF3 • UCSX-MR128G4RE3 • UCSX-MR256G8RE3 		
		Adapters	UCSX-ML-V5Q50G (Secure Boot) UCSX-ME-V5Q50G (Secure Boot) UCSX-ML-V5D200G	N/A	5.1(0.230096)
			UCSX-ML-V5D200GV2 (Secure Boot)	N/A	5.2(0.230061)
		Trusted Platform Module (TPM)	UCSX-TPM-002C	N/A	5.1(0.230096)
			UCS-TPM-002D	N/A	5.2(2.240053)
		Graphics processing unit (GPU)	UCSX-GPU-T4-MEZZ	N/A	5.1(0.230096)
			UCSX-GPU-A16		
			UCSX-GPU-A40		
			UCSX-GPU-A100-80	N/A	5.1(0.230096)
			UCSX-GPU-H100-80	N/A	5.2(0.230127)
			UCSX-GPU-L40		
			UCSX-GPU-L4		
		UCSX-GPU-FLEX140			
		UCSX-GPU-FLEX170			
			UCSX-GPU-L40S	N/A	5.2(2.240053)
				N/A	5.1(0.230096)

Fabric Components						
Model	Component	Sub-Component	Sub-Component Model	Minimum Infrastructure Firmware Versions	Minimum Server Firmware Versions	
		Storage Controller	UCSX-X10C-PT4F UCSX-X10C-RAIDF UCSX-M2-HWRAID			
			UCSX-M2-PT-FPN	N/A	5.2(0.230041)	
UCSX-210C-M6	X-Series M6 Server			4.2(1a)	5.0(1b)	
		Adapters	UCSX-V4-Q25GML UCSX-V4-Q25GME	N/A	5.0(1b)	
			UCSX-ML-V5Q50G (Secure Boot) UCSX-ME-V5Q50G (Secure Boot)	N/A	5.1(0.230054)	
			UCSX-ML-V5D200G	N/A	5.0(2b)	
			UCSX-ML-V5D200GV2 (Secure Boot)	N/A	5.2(0.230061)	
			Rear Mezzanine Adapters	UCSX-V4-PCIME	N/A	5.0(2d)
		Trusted Platform Module (TPM)	UCSX-TPM-002C	N/A	5.0(1b)	
			UCS-TPM-002D	N/A	4.3(4.240152)	
		Graphics processing unit (GPU)	UCSX-GPU-A100-80	N/A	5.0(2e)	
			UCSX-GPU-T4-MEZZ UCSX-GPU-T4-16 UCSX-GPU-A16 UCSX-GPU-A40	N/A	5.0(2d)	
			Storage Controller	UCSX-X10C-PT4F UCSX-X10C-RAIDF UCSX-M2-HWRAID	N/A	5.0(4b)

Fabric Components					
Model	Component	Sub-Component	Sub-Component Model	Minimum Infrastructure Firmware Versions	Minimum Server Firmware Versions
UCSB-B200-M6	B-Series M6 Server			4.1(3b)	4.2(1d)
		Adapters	UCSB-ML-V5Q10G	N/A	4.2(1d)
			UCSB-MLOM-40G-04 UCSB-VIC-M84-4P	4.1(3b)	4.2(1d)
		Trusted Platform Module (TPM)	UCSX-TPM-002C	N/A	4.2(1d)
			UCS-TPM-002D	N/A	4.3(4.240152)
		Storage Controller	UCS-M2-HWRAID UCSB-RAID12G-M6 UCSB-MSTOR-M6 UCSB-LSTOR-PT-M6	N/A	4.2(1d)
UCSB-B200-M5 UCSB-B480-M5	B-Series M5 Server			4.1(3b)	4.1(3b)
		Adapters	UCSB-MLOM-40G-03 UCSB-VIC-M83-8P	4.1(3b)	4.2(2e)
			UCSB-MLOM-40G-04 UCSB-VIC-M84-4P	4.1(3b)	4.1(3b)
			UCSB-MLOM-PT-01	4.1(3b)	N/A
		Trusted Platform Module (TPM)	UCSX-TPM2-001 UCSX-TPM2-002	N/A	4.1(3b)
		Storage Controller	UCS-M2-HWRAID UCSB-MRAID12G UCSB-MRAID12G-HE UCSB-LSTOR-PT	N/A	4.1(3c)
UCSC-C245-M8	C-Series M8 Server	For more information on the sub-components bundled with the server, see Cisco UCS C245 M8 Server Spec Sheet .		4.2(3b)	4.3(4.241014)

Fabric Components					
Model	Component	Sub-Component	Sub-Component Model	Minimum Infrastructure Firmware Versions	Minimum Server Firmware Versions
UCSC-C220-M7	C-Series M7 Server			4.2(3b)	4.3(1.230097)
		CPU	4 th Generation Intel [®] Xeon [®] Scalable Processors.	4.2(3b)	4.3(1.230097)
			5 th Generation Intel [®] Xeon [®] Scalable Processors Note For more information, see the Cisco UCSC C220 M7 SFF Rack Server Spec Sheet .	4.2(3b)	4.3(3.240022)
			Additional 5 th Generation Intel [®] Xeon [®] Scalable Processors: <ul style="list-style-type: none"> • UCS-CPU-I4510T • UCS-CPU-I4510 • UCS-CPU-I4509Y • UCS-CPU-I3508U 	4.2(3b)	4.3(4.240152)
		DIMMs	DDR5-4800MT/s: <ul style="list-style-type: none"> • UCS-MRX16G1RE1 • UCS-MRX32G1RE1 • UCS-MRX64G2RE1 • UCS-MR128G4RE1 	4.2(3b)	4.3(3.240022)
			4.2(3b)	4.3(3.240022)	

Fabric Components					
Model	Component	Sub-Component	Sub-Component Model	Minimum Infrastructure Firmware Versions	Minimum Server Firmware Versions
			DDR5-5600MT/s: <ul style="list-style-type: none"> • UCS-MRX16G1RE3 • UCS-MRX32G1RE3 • UCS-MRX48G1RF3 • UCS-MRX64G2RE3 • UCS-MRX96G2RF3 • UCS-MR128G4RE3 <p>Note 48Gb and 96GB memory DIMMs are not supported on UCS-CPU-I3508U, UCS-CPU-I4509Y, UCS-CPU-I4510, UCS-CPU-I4510T</p>		
		Adapters	UCSC-M-V5Q50G UCSC-M-V5D200G	N/A	4.3(1.230097)
			UCSC-P-V5D200G (Secure Boot) UCSC-P-V5Q50G (Secure Boot)	4.3(2.230117)	4.3(2.230184)
			UCSC-M-V5D200GV2 (Secure Boot) UCSC-M-V5Q50GV2 (Secure Boot)	N/A	4.3(2.230258)
		Trusted Platform Module (TPM)	UCSX-TPM-002C	N/A	4.3(1.230097)
			UCS-TPM-002D	N/A	4.3(4.240152)
		Graphics processing unit (GPU)	UCSC-GPU-A16 UCSC-GPU-A100-80	N/A	4.3(1.230097)
				N/A	4.3(2.230207)

Fabric Components					
Model	Component	Sub-Component	Sub-Component Model	Minimum Infrastructure Firmware Versions	Minimum Server Firmware Versions
			UCSC-GPU-L4 UCSC-GPU-FLEX140		
		Storage Controller	UCS-M2-NVRAID	4.3(2.230117)	4.3(2.230207)
		Virtual Drives	UCS-SD16TKA3X-EP UCS-SD32TKA3X-EP UCS-SD16TBKANK9 UCS-SD19TKA1X-EV UCS-SD38TKA1X-EV UCS-SD76TKA1X-EV UCS-SD15TKA1X-EV UCS-SD38TBKANK9 UCS-SD76TBKANK9	4.3(2.230117)	4.3(2.230207)

Fabric Components					
Model	Component	Sub-Component	Sub-Component Model	Minimum Infrastructure Firmware Versions	Minimum Server Firmware Versions
UCSC-C240-M7	C-Series M7 Server			4.2(3b)	4.3(1.230097)
		CPU	4 th Generation Intel [®] Xeon [®] Scalable Processors.	4.2(3b)	4.3(1.230097)
			5 th Generation Intel [®] Xeon [®] Scalable Processors Note For more information, see the Cisco UCS C240 M7 SFF Rack Server Spec Sheet .	4.2(3b)	4.3(3.240022)
			Additional 5 th Generation Intel [®] Xeon [®] Scalable Processors: <ul style="list-style-type: none"> • UCS-CPU-I4510T • UCS-CPU-I4510 • UCS-CPU-I4509Y • UCS-CPU-I3508U 	4.2(3b)	4.3(4.240152)
		DIMMs	DDR5-4800MT/s: <ul style="list-style-type: none"> • UCS-MRX16G1RE1 • UCS-MRX32G1RE1 • UCS-MRX64G2RE1 • UCS-MR128G4RE1 • UCS-MR256G8RE1 	4.2(3b)	4.3(3.240022)
	4.2(3b)		4.3(3.240022)		

Fabric Components					
Model	Component	Sub-Component	Sub-Component Model	Minimum Infrastructure Firmware Versions	Minimum Server Firmware Versions
			DDR5-5600MT/s: <ul style="list-style-type: none"> • UCS-MRX16G1RE3 • UCS-MRX32G1RE3 • UCS-MRX48G1RF3 • UCS-MRX64G2RE3 • UCS-MRX96G2RF3 • UCS-MR128G4RE3 • UCS-MR256G8RE3 		
		Trusted Platform Module (TPM)	UCSX-TPM-002C	N/A	4.3(1.230097)
			UCS-TPM-002D	N/A	4.3(4.240152)
		Adapters	UCSC-M-V5Q50G UCSC-M-V5D200G	N/A	4.3(1.230097)
			UCSC-P-V5D200G (Secure Boot) UCSC-P-V5Q50G (Secure Boot)	4.3(2.230117)	4.3(2.230184)
			UCSC-M-V5D200GV2 (Secure Boot) UCSC-M-V5Q50GV2 (Secure Boot)	N/A	4.3(2.230258)
			Graphics processing unit (GPU)	UCSC-GPU-A16 UCSC-GPU-A100-80	N/A
		UCSC-GPU-H100-80 UCSC-GPU-L40 UCSC-GPU-L4 UCSC-GPU-FLEX140 UCSC-GPU-FLEX170		N/A	4.3(2.230207)
		UCSC-GPU-L40S		N/A	4.3(4.240152)

Fabric Components					
Model	Component	Sub-Component	Sub-Component Model	Minimum Infrastructure Firmware Versions	Minimum Server Firmware Versions
		Storage Controller	UCS-M2-NVRAID	4.3(2.230117)	4.3(2.230207)
		Virtual Drives	UCS-SD16TKA3X-EP UCS-SD32TKA3X-EP UCS-SD16TBKANK9 UCS-SD19TKA1X-EV UCS-SD38TKA1X-EV UCS-SD76TKA1X-EV UCS-SD15TKA1X-EV UCS-SD38TBKANK9 UCS-SD76TBKANK9	4.3(2.230117)	4.3(2.230207)

Fabric Components					
Model	Component	Sub-Component	Sub-Component Model	Minimum Infrastructure Firmware Versions	Minimum Server Firmware Versions
UCSC-C220-M6 UCSC-C240-M6	C-Series M6 Server	Adapters		4.1(3b)	4.2(1e)
			UCSC-PCIE-C25Q-04	4.1(3b)	4.2(1e)
			UCSC-PCIE-C100-04		
			UCSC-M-V100-04	4.1(3b)	4.2(1e)
			UCSC-M-V25-04		
			UCSC-M-V5D200G	N/A	4.2(2f)
			UCSC-M-V5Q50G	N/A	4.2(2b)
			UCSC-P-V5D200G (Secure Boot)	4.3(2.230117)	4.3(2.230184)
			UCSC-P-V5Q50G (Secure Boot)		
			UCSC-M-V5D200GV2 (Secure Boot)	N/A	4.3(2.230258)
		UCSC-M-V5Q50GV2 (Secure Boot)			
		Trusted Platform Module (TPM)	UCSX-TPM-002C	N/A	4.2(1e)
			UCS-TPM-002D	N/A	4.3(4.240152)
		Graphics processing unit (GPU)	UCSC-GPU-A16	N/A	4.2(3b)
			UCSC-GPU-A100-80		
			UCSC-GPU-L4	N/A	4.3(4.240152)
Storage Controller	UCS-M2-HWRAID	N/A	4.2(1e)		
	UCSC-RAID-M6T				
	UCSC-RAID-M6SD				
	UCSC-RAID-M6HD				
	UCSC-SAS-M6HD				
	UCSC-SAS-M6T				

Fabric Components					
Model	Component	Sub-Component	Sub-Component Model	Minimum Infrastructure Firmware Versions	Minimum Server Firmware Versions
UCSC-C245-M6 UCSC-C225-M6	C-Series M6 Server			4.1(3b)	4.2(1e)
		Adapters	UCSC-PCIE-C25Q-04 UCSC-PCIE-C100-04	4.1(3b)	4.2(1e)
			UCSC-M-V100-04 UCSC-M-V25-04	4.1(3b)	4.2(1e)
			UCSC-M-V5D200G	N/A	4.2(2f)
			UCSC-M-V5Q50G	N/A	4.2(2b)
			UCSC-P-V5D200G (Secure Boot) UCSC-P-V5Q50G (Secure Boot)	4.3(2.230117)	4.3(2.230184)
			UCSC-M-V5D200GV2 (Secure Boot) UCSC-M-V5Q50GV2 (Secure Boot)	N/A	4.3(2.230258)
			Trusted Platform Module (TPM)	UCSX-TPM2-002B-C	N/A
		UCS-TPM2-002D		N/A	4.3(4.240152)
		Graphics processing unit (GPU)	UCSC-GPU-A16 UCSC-GPU-A100-80	N/A	4.2(3b)
			Storage Controller	UCS-M2-HWRAID UCSC-RAID-M6T UCSC-RAID-M6SD UCSC-RAID-M6HD UCSC-SAS-M6HD UCSC-SAS-M6T	N/A

Fabric Components					
Model	Component	Sub-Component	Sub-Component Model	Minimum Infrastructure Firmware Versions	Minimum Server Firmware Versions
UCSC-C220-M5	C-Series M5 Server			4.1(3b)	4.1(3b)
UCSC-C240-M5		Adapters	UCSC-MLOM-C40Q-03	4.1(3b)	4.2(2g)
UCSC-C480-M5			UCSC-PCIE-C40Q-03		
			UCSC-PCIE-C25Q-04	4.1(3b)	4.1(3b)
			UCSC-MLOM-C25Q-04		
			UCSC-PCIE-C100-04		
			UCSC-MLOM-C100-04		
		Trusted Platform Module (TPM)	UCSX-TPM2-001 UCSX-TPM2-002	N/A	4.1(3b)
		Graphics processing unit (GPU)	UCSC-GPU-A100-80	N/A	4.2(3b)
		Storage Controller	UCS-M2-HWRAID UCSC-RAID-M5HD UCSC-RAID-M5 UCSC-SAS-M5, UCSC-SAS-M5HD UCSC-SAS12GHBA UCSC-9400-8E	N/A	4.1(3b)



Note Post Infra Firmware release 4.2(3c), the Server Firmware bundle in Intersight Infrastructure Service (IIS) will bear the version number in a new format instead of the letter format.

With Infra Firmware release 4.3(2), the Infra Firmware bundle in IIS will bear the version number in a new format instead of the letter format.

For example: 4.3(2.230117) , where 23 represents year, 0117 shows the incremental build number.

For more information on Cisco Intersight Infrastructure Firmware release notes, Server Firmware release notes, and Release Bundle Content document see [Release Notes](#).

Table 3: Supported Combination of Hardware Components in IMM

Component	Supported Combination
Cisco UCS X-Series Direct	<p>Integrated FI: UCSX-S9108-100G</p> <p>Cisco UCS X-Series Direct supports all components that are compatible with Cisco UCS X-Series modular system.</p> <p>It currently does not support Fabric Extender, Rack servers, or adding additional chassis in Intersight Managed Mode.</p>
Topologies	<p>Direct-Attached Racks through 10G/25G/100G connections</p> <p>Break-out port configuration through 10G/25G connections</p> <p>FEX-Attached Racks through 10GE connections</p> <p>Chassis through 10G/25G/100G connections</p> <p>N9K-C93108YC-FX3 FEX through 10G/25G connections</p>
Fabric Interconnects	<p>UCS-FI-6536 and direct-attached rack server are supported at 40G and 100G on Cisco UCS 1400 and 15000 series VIC adapters.</p>
Input/Output Module (IOM)	<ul style="list-style-type: none"> • UCS-IOM-2204XP and UCS-IOM-2208XP are not supported on Cisco UCS 6500 Series Fabric Interconnects. • UCS-IOM-2304 and UCS-IOM-2304V2 are supported only with Cisco UCS 6500 series Fabric Interconnect. • When there is a mixed IOM configuration, Access Policy deployment can fail resulting in Server Profile deployment failure. It will recover once both the IOMs are replaced.
X-Fabric Modules (XFM)	<p>UCS 9416 X-Fabric module is supported only on UCSX-9508 chassis and required for Peripheral Component Interconnect Express (PCIe) node and GPU discovery or inventory support in IMM.</p>
Fabric Extender (FEX)	<p>Cisco Nexus 2232PP is not supported on Cisco UCS 6500 Series Fabric Interconnects.</p>
Rear Mezzanine Adapters	<ul style="list-style-type: none"> • UCS PCI mezz card for X-Fabric connectivity. • The UCSX-210C Compute Node must include a UCSX-V4-PCIME or a supported mezz card when paired with a X440p PCIe node.

Component	Supported Combination
Adapters	

Component	Supported Combination
	<ul style="list-style-type: none"> • UCSX-X10C-GPUFM is an adapter that supports the GPU, UCSX-GPU-T4-MEZZ. For more information, see Cisco UCS X10c Front Mezzanine GPU Module Installation and Service Guide. • UCSX-V4-Q25GME is a mezz card requires UCS VIC 14000 bridge connector (UCSX-V4-BRIDGE) and UCSX-V4-Q25GML mLOM support in the X210c Compute Node. For more information, see Cisco UCS X210c M6 Compute Node. • The UCSX-210C Compute Node must include a UCSX-V4-PCIME or a supported mezz card when paired with a X440p PCIe node. • UCSX-ML-V5D200G adapter is supported on Cisco UCS 6500 series Fabric Interconnect at 40G and 100G speed, as well as on Cisco UCS 6400 series Fabric Interconnect at 25G speed. • Cisco UCS C-Series and X-Series M7 servers support only Cisco UCS 15000 series VIC adapters. • UCSX-ME-V5Q50G is a mezz card that requires UCS VIC 15000 bridge connector (UCSX-V5-BRIDGE) and UCSX-ML-V5Q50G mLOM support in the X210c Compute Node. However, this mezz adapter is not supported with UCSX-ML-V5D200G mLOM. <ul style="list-style-type: none"> • On a B-series server, installing a combination of Cisco UCS 1400 and UCS 15000 series VIC adapters is not supported. • Cisco UCS VIC 1300 Series adapters are supported on B-Series and C-Series M5 servers with the following combination. <ul style="list-style-type: none"> • UCS-FI-6454 and UCS-IOM-2408 • UCS-FI-6536 and UCS-IOM-2408 • UCS-FI-6454 and UCS-IOM-2204XP • UCS-FI-6454 and UCS-IOM-2208XP • UCS-FI-6536 and direct-attached rack server at 40G • UCS-FI-6454 and rack server connected through FEX • UCS-FI-6454 and direct-attached rack server with 10G QSA • UCS-FI-6536 and UCS-IOM-2304 or

Component	Supported Combination
	<p>UCS-IOM-2304V2</p> <ul style="list-style-type: none"> • UCS-FI-64108 and UCS-IOM-2408 • UCS-FI-64108 and UCS-IOM-2204XP • UCS-FI-64108 and UCS-IOM-2208XP • UCS-FI-64108 and direct-attached rack server • UCS-FI-64108 and rack server connected through FEX <ul style="list-style-type: none"> • UCSC-M-V100-04, UCSC-PCIE-C100-04, UCSC-MLOM-C100-04 are supported only on Cisco UCS 6500 Series Fabric Interconnects. <ul style="list-style-type: none"> • The following combinations are not supported on UCS C series M6 servers: <ul style="list-style-type: none"> • 1400 Series MLOM adapters with 15000 Series PCIE adapters • UCSC-M-V5Q50GV2 and UCSC-M-V5D200GV2 are not supported with 14xx PCIE adapters • Ensure that you have upgraded servers to the VIC supported release versions before installing the VIC adapters into the server. If you install VIC adapters on servers running an earlier release and later decide to upgrade the servers to the supported version, you need to perform A/C power cycle for the servers to enable the adapters.
Graphics processing unit (GPU)	<ul style="list-style-type: none"> • All supported X-Series GPU are supported on UCS X440P with UCSX-210C-M6 and UCSX-210C-M7 Compute Nodes. • Mixing of GPU models are not supported in the server. For more information, see Cisco UCS X440p PCIe Node Installation and Service Guide. • Specific GPUs are also supported on the X210c Compute Nodes. They require the UCSX-X10C-GPUFM adapter to support a GPU in the Front Mezz. • The GPU supported in the X210c M7 Front Mezz includes UCSX-GPU-T4-MEZZ. For more information, see Cisco UCS X10c Front Mezzanine GPU Module Installation and Service Guide.



CHAPTER 2

Setting Up Fabric Interconnects

- [Initial Fabric Interconnect Configuration, on page 23](#)
- [Configuring Fabric Interconnect-A Using the Console, on page 24](#)
- [Configuring Fabric Interconnect-B Using the Console, on page 26](#)
- [Configuring Fabric Interconnect-A Using the GUI, on page 27](#)
- [Configuring Fabric Interconnect-B Using the GUI, on page 30](#)
- [Fabric Interconnect Password Guidelines, on page 31](#)
- [Migrating to Cisco UCS 6500 Series Fabric Interconnects, on page 32](#)
- [Migrating from UCSX-I-9108-25G to UCSX-I-9108-100G IFM on an existing domain with Cisco UCS 6500 Series Fabric Interconnect and UCSX-9508 Chassis, on page 34](#)
- [Fabric Interconnect Views, on page 35](#)
- [Fabric Interconnects Table View, on page 35](#)
- [Fabric Interconnects Details View, on page 36](#)
- [Fabric Interconnects Inventory View, on page 39](#)
- [Fabric Interconnects Connections View, on page 41](#)
- [Fabric Interconnects UCS Domain Profile View, on page 41](#)
- [Fabric Interconnect Topology View, on page 42](#)
- [Fabric Interconnect Metrics View, on page 42](#)
- [Fabric Interconnect Actions, on page 42](#)
- [Configuring Traffic Mirroring \(SPAN\), on page 45](#)

Initial Fabric Interconnect Configuration

The initial configuration for a Fabric Interconnect can be done by using the serial console when the Fabric Interconnect boots for the first time. This can happen either during factory install, or after the existing configuration is cleared. The configuration wizard enables you to select the management mode and other parameters such as the administrative subnet, gateway, and DNS IP addresses for each Fabric Interconnect. For the management mode, you can choose whether you want to manage the Fabric Interconnect through Cisco UCS Manager or Cisco Intersight.

You can change the management mode for the Fabric Interconnects between Cisco Intersight and Cisco UCS Manager. However, this is a disruptive process because it will cause all endpoint configurations to be reset, and will result in the loss of the current configuration.



Note All the discovered servers, chassis, and Fabric Extenders (FEX) must be decommissioned before changing the management mode.

The erase configuration option, which is available in both management modes, allows you to clear the existing configuration and reboot the Fabric Interconnects. After the Fabric Interconnects are rebooted, the initial configuration screen appears, and you can configure the Fabric Interconnects with the appropriate management mode.

This configuration process is valid for Cisco UCS 6400 Series Fabric Interconnects, Cisco UCS 6500 Series Fabric Interconnects, and Cisco UCS Fabric Interconnects 9108 100G in a cluster setup.



Note Cisco UCS 6500 Series Fabric Interconnects support UCSM Managed Mode (UMM) from firmware version 4.2(3b) onwards.

To configure the Fabric Interconnects in a cluster:

1. [Configuring Fabric Interconnect-A Using the Console](#)
2. [Configuring Fabric Interconnect-B Using the Console](#)

After completing the initial configuration of the Fabric Interconnects, you must claim them for use with the Cisco Intersight platform. For more information about claiming devices in Cisco Intersight, see [Target Claim in Intersight Managed Mode](#).

After you claim the Fabric Interconnects, they appear in the list of available devices. The device type for Fabric Interconnects managed through Cisco Intersight is **Intersight Managed Domain**. The **Device IP** field shows the IP addresses of both the Fabric Interconnects, and the **Device ID** field shows the serial numbers of both the Fabric Interconnects. The Fabric Interconnects now appear in the **Fabric Interconnects** table view.

After you claim the Fabric Interconnects, you must configure the ports on the Fabric Interconnect to discover the connected chassis and servers. For each Fabric Interconnect, you can view the properties, and an inventory of its components, including ports, fan modules, and power supply units (PSUs).

Configuring Fabric Interconnect-A Using the Console

- Step 1** Connect to the console port.
- Step 2** Power on the Fabric Interconnect.
You will see the power-on self-test messages as the Fabric Interconnect boots.
- Step 3** When the unconfigured system boots, it prompts you for the setup method to be used. Enter **console** to continue the initial setup using the console CLI.
- Step 4** Enter the management mode for the Fabric Interconnect:
- **intersight** to manage the Fabric Interconnect through Cisco Intersight
 - **ucsm** to manage the Fabric Interconnect through Cisco UCS Manager

NOTE:

- Standalone option is not supported in the Intersight Managed Mode.
- If the Fabric Interconnect is an Intersight Managed Mode only Fabric Interconnect, which is the default mode, then you can select No and choose the required one.

- Step 5** Enter **y** to confirm that you want to continue the initial setup.
- Step 6** To use a strong password, enter **y**
- Step 7** Enter the password for the admin account. For more details, see [Fabric Interconnect Password Guidelines](#).
- Step 8** To confirm, re-enter the password for the admin account.
- Step 9** Enter **yes** to continue the initial setup for a cluster configuration.
- Step 10** Enter the Fabric Interconnect fabric (either **A** or **B**).
- Step 11** Enter the system name.
- Step 12** Enter the IPv4 or IPv6 address for the management port of the Fabric Interconnect.
If you enter an IPv4 address, you will be prompted to enter an IPv4 subnet mask. If you enter an IPv6 address, you will be prompted to enter an IPv6 network prefix.
- Step 13** Enter the respective IPv4 subnet mask or IPv6 network prefix, then press **Enter**.
You are prompted for an IPv4 or IPv6 address for the default gateway, depending on the address type you entered for the management port of the Fabric Interconnect.
- Step 14** Enter either of the following:
- IPv4 address of the default gateway
 - IPv6 address of the default gateway
- Step 15** Enter the IPv4 or IPv6 address for the DNS server.
The address type must be the same as the address type of the management port of the Fabric Interconnect.
- Step 16** Enter **yes** if you want to specify the default Domain name, or **no** if you do not.
- Step 17** (Optional) Enter the default Domain name.
- Step 18** Review the setup summary and enter **yes** to save and apply the settings, or enter **no** to go through the setup again to change some of the settings.
If you choose to go through the setup again, it provides the values that you previously entered, and the values appear in brackets. To accept previously entered values, press **Enter**.

Example

Here is an example of how to configure Fabric Interconnect-A in Cisco Intersight management mode for a cluster configuration using the console and management addresses:

```
Enter the configuration method (console/gui)? console
Enter the management mode [ucsm/intersight]? intersight
You have chosen to setup a new Fabric Interconnect in "intersight" managed mode. Continue?
(y/n): y
Enforce strong password? (y/n) [y]:n

Enter the password for "admin":
Confirm the password for "admin":
```

```

Enter the switch fabric (A/B) []: A

Enter the system name: UCS

Physical Switch Mgmt0 IP address : 15.XX.XX.XX

Physical Switch Mgmt0 IPv4 netmask : 255.255.255.X

IPv4 address of the default gateway : 15.XX.XX.XX

DNS IP address : 15.XX.XX.XX

Configure the default domain name? (yes/no) [n]:

Following configurations will be applied:

Management Mode=intersight
Switch Fabric=A
System Name=UCS-A
Enforced Strong Password=no
Physical Switch Mgmt0 IP Address=15.XX.XX.XX
Physical Switch Mgmt0 IP Netmask=255.255.255.X
Default Gateway=15.XX.XX.XX
Ipv6 value=0
DNS Server=15.XX.XX.XX

Apply and save the configuration (select 'no' if you want to re-enter)? (yes/no): yes
Applying configuration. Please wait.

Configuration file - Ok

```

What to do next

Configure the Fabric Interconnect-B using the console.

Configuring Fabric Interconnect-B Using the Console

This procedure describes setting up Fabric Interconnect-B using IPv4 or IPv6 addresses for the management port.

-
- Step 1** Connect to the console port.
- Step 2** Power up the Fabric Interconnect.
You will see the power-on self-test messages as the Fabric Interconnect boots.
- Step 3** When the unconfigured system boots, it prompts you for the setup method to be used. Enter **console** to continue the initial setup using the console CLI.
- Note** Fabric Interconnect-A should detect Fabric Interconnect-B in the cluster. If it does not, check the physical connections between the L1 and L2 ports, and verify that Fabric Interconnect-B has been enabled for a cluster configuration.
- Step 4** Enter **y** to add Fabric Interconnect-B to the cluster.
- Step 5** Enter the admin password of the peer Fabric Interconnect.
- Step 6** Enter the IP address for the management port on Fabric Interconnect-B.

Step 7 Review the setup summary and enter **yes** to save and apply the settings, or enter **no** to go through the setup again to change some of the settings.

If you choose to go through the setup again, it provides the values you previously entered, and the values appear in brackets. To accept previously entered values, press **Enter**.

Example

Here is an example of how to configure Fabric Interconnect-B in Cisco Intersight management mode for a cluster configuration using the console and management addresses:

```
Enter the configuration method. (console/gui) ? console
```

```
Installer has detected the presence of a peer Fabric interconnect. This Fabric interconnect will be added to the cluster. Continue (y/n) ? y
```

```
Enter the admin password of the peer Fabric interconnect:
Connecting to peer Fabric interconnect... done
Retrieving config from peer Fabric interconnect... done
Peer Fabric interconnect management mode : intersight
Peer Fabric interconnect Mgmt0 IPv4 Address: 15.XX.XX.XX
Peer Fabric interconnect Mgmt0 IPv4 Netmask: 255.255.255.0
```

```
Peer FI is IPv4 Cluster enabled. Please Provide Local Fabric Interconnect Mgmt0 IPv4 Address
```

```
Physical Switch Mgmt0 IP address : 15.XX.XX.XX
```

```
Local fabric interconnect model(UCS-FI-6454)
Peer fabric interconnect is compatible with the local fabric interconnect. Continuing with the installer...
```

```
Apply and save the configuration (select 'no' if you want to re-enter)? (yes/no): yes
Applying configuration. Please wait.
```

```
Configuration file - Ok
```

What to do next

Claim the Intersight Managed Domain through Cisco Intersight. For more information, see [Target Claim in Intersight Managed Mode](#).

Configuring Fabric Interconnect-A Using the GUI

Configuring Fabric Interconnect-A Using GUI

This procedure describes setting up Fabric Interconnect-A using GUI.

For detailed information on how to configure Fabric Interconnect-A using GUI, see [Initial System Setup for a Cluster Configuration](#).

1. Power on the Fabric Interconnect.

You will see the power on self-test messages as the Fabric Interconnect boots.

2. If the system obtains a lease, go to step 6, otherwise, continue to the next step.
3. Connect to the console port.
4. At the installation method prompt, enter **gui**.
5. If the system cannot access a DHCP server, you are prompted to enter the following information:
 - IPv4 or IPv6 address for the management port on the Fabric Interconnect.
 - IPv4 subnet mask or IPv6 prefix for the management port on the Fabric Interconnect.
 - IPv4 or IPv6 address for the default gateway assigned to the Fabric Interconnect.



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

6. Copy the web link from the prompt into a web browser and go to the Cisco UCS Fabric Interconnect Setup GUI launch page.



Note You can choose between two modes: UCSM Managed and Intersight Managed Fabric Interconnects based on your preferences.

7. In the **Cisco UCS Fabric Interconnect Setup GUI** launch page, select **Express Setup**.
8. In the **Express Setup** page, enter the Fabric Interconnects configuration details.



Note From Cisco UCS Manager 4.2(2) onwards, you can choose GUI setup method to configure Fabric Interconnects. If the Fabric Interconnect defaults to Intersight managed mode, you can choose to change during confirmation and select required mode again in console setup method alone.

9. In the **Basic Settings** area:
 - For the **Fabric Setup** option, select **Fabric A**.
 - Select IPv4 or IPv6 address that Cisco Intersight Managed Mode will use.

and click **Submit**.

10. In the **System Setup** area, complete the following fields:

Field	Description
Enforce Strong Password	Choose Yes or No to enforce strong password.

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

**Note**

- For Intersight Managed Mode Fabric Interconnects, DNS is mandatory
- Standalone option is not supported in the Intersight Managed Mode.

11. Click **Submit**.

A page displays the results of your setup operation.

What to do next

Configure the Fabric Interconnect-B using the GUI.

Configuring Fabric Interconnect-B Using the GUI

Configuring Fabric Interconnect-B Using GUI

This procedure describes setting up Fabric Interconnect-B using GUI.

You can either follow the procedure below for configuring the Fabric Interconnect-B or watch [Cisco UCS Manager Initial Setup part 2](#).



Note When adding a new Fabric Interconnect to an existing High Availability cluster, for example, during a new install or when replacing a Fabric Interconnect, the new device will not be able to log into the cluster as long as the authentication method is set to remote. To successfully add a new Fabric Interconnect to the cluster, the authentication method must be temporarily set to local and the local admin credentials of the primary Fabric Interconnect must be used.

1. Power up the Fabric Interconnect.
You will see the power-up self-test message as the Fabric Interconnect boots.
2. If the system obtains a lease, go to step 6, otherwise, continue to the next step.
3. Connect to the console port.
4. At the installation method prompt, enter **gui**.
5. If the system cannot access a DHCP server, you are prompted to enter the following information:
 - IPv4 or IPv6 address for the management port on the Fabric Interconnect.
 - IPv4 subnet mask or IPv6 prefix for the management port on the Fabric Interconnect.
 - IPv4 or IPv6 address for the default gateway assigned to the Fabric Interconnect.



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

6. Copy the web link from the prompt into a web browser and go to the Cisco UCS Fabric Interconnect Setup GUI launch page.



Note You can choose between two modes: UCSM Managed and Intersight Managed Fabric Interconnects based on your preferences.

7. In the **Cisco UCS Fabric Interconnect Setup GUI** launch page, select **Express Setup**.
8. In the **Express Setup** page, enter the Fabric Interconnects configuration details.



Note From Cisco UCS Manager 4.2(2) onwards, you can choose GUI setup method to configure Fabric Interconnects. If the Fabric Interconnect defaults to Intersight managed mode, you can choose to change during confirmation and select required mode again in console setup method alone.

9. In the **Basic Settings** area:
 - For the **Fabric Setup** option, make sure **Fabric B** is selected.
10. In the **System Setup** area, enter the password for the Admin account into the **Admin Password of Master** field.
The **Manager Initial Setup** area is displayed.
11. In the **Manager Initial Setup** area, the field that is displayed depends on whether you configured the first Fabric Interconnect with an IPv4 or IPv6 management address. Complete the field that is appropriate for your configuration, as follows:

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

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

What to do next

Claim the Intersight Managed Domain through Cisco Intersight. For more information, go to [Target Claim in Intersight Managed Mode](#).

Fabric Interconnect Password Guidelines

Cisco recommends using a strong password; otherwise, the password strength check for the admin user of the Fabric Interconnect, Cisco Intersight rejects any password that does not meet the following requirements:

- Must contain a minimum of eight characters and a maximum of 80 characters.

- Must contain at least three of the following:
 - Lower case letters
 - Upper case letters
 - Digits
 - Special characters
- Must not contain a character that is repeated more than three times consecutively, such as aaabbb.
- Must not be identical to the username or the reverse of the username.
- Must pass a password dictionary check. For example, the password must not be based on a standard dictionary word.
- Must not contain the following symbols: \$ (dollar sign), ? (question mark), and = (equals sign).
- Should not be blank.

Migrating to Cisco UCS 6500 Series Fabric Interconnects

Replacing Cisco UCS 6400 Series Fabric Interconnect-B with Cisco UCS 6500 Series Fabric Interconnect-B

This section describes the process of migrating from Cisco UCS 6400 Series Fabric Interconnects to Cisco UCS 6500 Series Fabric Interconnects.

Procedure:

1. To have minimal traffic loss during migration, ensure that there are redundant paths from the chassis over fabrics A and B, and vNIC should either be redundant or that the fabric failover is enabled. Since in-flight packet loss is expected during the IFM migration, best to perform the next set of operations only during a maintenance window.
2. Power down the Cisco UCS 6400 Series Fabric Interconnect by unplugging it from the power source. If you are monitoring the migration using a KVM session, you may need to reconnect the KVM session when you power down the Fabric Interconnect.
3. Disconnect all connections such as the network management port, L1/L2 ports, rack server, Chassis IOM/IFM ports, fabric extenders, uplink ports, and fibre connections from the Cisco UCS 6400 Fabric Interconnect-B.
4. Replace Cisco UCS 6400 Series Fabric Interconnect-B with Cisco UCS 6500 Series Fabric Interconnect.
5. Connect all connections such as the network management port, L1/L2 ports, rack server, Chassis IOM/IFM ports, fabric extenders, uplink ports, and fibre connections onto the new Cisco UCS 6500 Series Fabric Interconnects.

Proper cables should be used to connect to the UCS 6500 Series Fabric Interconnect. For more information, refer to the [Cisco UCS 6500 Series Fabric Interconnect Hardware Installation Guide](#) and [Cisco UCS 6500 Series Fabric Interconnect Data Sheet](#).



Note You can choose to migrate IFM to UCSX-I-9108-100G on UCSX-9508 Chassis.

UCS-IOM-2204/2208XP are not supported with Cisco UCS 6500 Series Fabric Interconnect. You can migrate to UCS-IOM-2408 on the Cisco UCS 5100 series Chassis.

6. Connect the power to the new Cisco UCS 6500 Series Fabric Interconnect, it will automatically boot and run POST tests.
Important Directly connect the console port to a terminal and observe the boot sequence. You should at some point see the Basic System Configuration Dialog, where you will configure the switch as a peer interconnect. If the Cisco UCS 6500 Series Fabric Interconnect have been previously configured or been part of a cluster, it will need to have all configuration information wiped before it can be added to a cluster. Immediately disconnect the L1 and L2 connections and login to the Fabric Interconnect and run the erase configuration to clear out existing configuration.
7. Create new Port policies for Cisco UCS 6500 Series Fabric Interconnect-B to reflect connectivity with the Fabric Interconnect.
 - Configure Ethernet/FC breakout ports as required.
 - Configure Port roles/ Port channels as required.
8. Edit the domain profile for the cluster and change the Fabric Interconnect-B Port policy to refer to the new Cisco UCS 6500 Series Fabric Interconnect-B Port policy.



Note **Deploy domain profile** will be part of the replacement workflow. Therefore, no need to deploy the profile after modifying the port policy. Domain profile deployment will fail as the Fabric Interconnect model has not been updated.

9. Go to **Operate > Fabric Interconnects** to view the new Cisco UCS 6500 Series Fabric Interconnect and the old Cisco UCS 6400 Series Fabric Interconnect.
10. Click **Replace Fabric Interconnect** option for Cisco UCS 6400 Series Fabric Interconnect to start the Replacement workflow.
11. Verify that
 - The disconnected Fabric Interconnect cluster is removed from inventory.
 - The domain profile is reassigned to the new Fabric Interconnect cluster and deployed.
 - The servers, chassis, and FEX are inventoried and discovered under the new Fabric Interconnect cluster.
 - The server and chassis profiles are redeployed with Fabric Interconnect related policies.

**Note**

- If there is a mix of different IFM models, the chassis profile will not be pushed until both the IFM are the same.

For example, if you migrate the IFM to UCSX-I-9108-100G from UCSX-I-9108-25G, and have chassis profile deployed prior to the migration, the chassis profile will not be deployed when the chassis has a mix of different IFM models. The chassis profile will be automatically deployed after both the IFM in the chassis have been migrated to UCSX-I-9108-100G.

Replacing Cisco UCS 6400 Series Fabric Interconnect-A with Cisco UCS 6500 Series Fabric Interconnect-A

Repeat the above-mentioned procedure for Fabric Interconnect-A and complete the UCS 6400 Series Fabric Interconnect to UCS 6500 Series Fabric Interconnect migration.

Migrating from UCSX-I-9108-25G to UCSX-I-9108-100G IFM on an existing domain with Cisco UCS 6500 Series Fabric Interconnect and UCSX-9508 Chassis

Follow the below procedure to move to 100G IFM from UCSX-I-9108-25G IFM with Cisco UCS 6500 Series Fabric Interconnect and UCSX-9508 chassis.

Procedure:

1. To have minimal traffic loss during migration, ensure that there are redundant paths from the chassis over fabrics A and B, and vNIC should either be redundant or that the fabric failover is enabled. Since inflight packet loss is expected during the IFM migration, best to perform the next set of operations only during a maintenance window.
2. Proceed with replacing one IFM at a time.

Start from Fabric Interconnect-B Port policy, unconfigure the server ports toward the migrating UCSX-I-9108-25G IFM. Once the server ports are unconfigured, the peer fabric interconnect will take over traffic-forwarding for these migrating UCS chassis.

3. Deploy the domain profile.
4. Disconnect the cables connecting the peer Cisco UCS 6500 Series Fabric Interconnect-B and the corresponding UCSX-I-9108-25G IFM from each migrating chassis.
5. Remove and replace the migrating UCSX-I-9108-25G IFM with UCSX-I-9108-100G IFM. Connect the UCSX-I-9108-100G IFM to the peer Cisco UCS 6500 Series Fabric Interconnect-B with proper cables. For more information, refer to the [Cisco UCS 6500 Series Fabric Interconnect Data Sheet](#).

At this point, the migrating UCS chassis will have a mix of UCSX-I-9108-25G and UCSX-I-9108-100G IFM.

6. Configure the Fabric Interconnect-B Port policy, then deploy the domain profile. Verify that the 100GbE links come up between the IFMs and the fabric interconnects.

- UCSX-I-9108-100G IFM will be auto-upgraded if the firmware is not the same with the Fabric Interconnect.
 - Once IFM comes online, it will be discovered and inventoried automatically. Blades in the chassis will get discovered, Server profile will get deployed automatically. Server will not be rebooted or have any disruption.
 - The Chassis profile will be automatically deployed after both the IFM in the chassis have been migrated to UCSX-I-9108-100G.
7. After completing IFM migration towards Cisco UCS 6500 Series Fabric Interconnect-B, repeat steps 3 through 7 for replacing the other UCSX-I-9108-25G connected to the Cisco UCS 6500 Series Fabric Interconnect-A and complete the UCSX-I-9108-100G IFM migration for the UCS domain.

Fabric Interconnect Views

Fabric Interconnects Table View

From the **Service Selector** drop-down list, select **Infrastructure Service**. Navigate to **Operate > Fabric Interconnects**, to launch the Fabric Interconnects Table view. Click the **Settings** icon (the gear icon representation), and select the columns that you want in the Table view. You can add the specific columns, or sort the columns by tags.

You can view the following details in the Fabric Interconnects Table view:

- **Name**—Displays the name of the Fabric Interconnect.
- **Health**—Status of the health of the Fabric Interconnect corresponds to the alarms on the servers. For more details, see [Alarms](#).
- **Contract Status**—Displays the status of service contract for the Fabric Interconnect based on the current validity of their associated contracts. You can identify the SmartNet Contract ID details of the server, and cross launch the [Cisco Commerce Software Subscriptions and Service Portal](#).
- **Management IP**—Displays the IP address of the management interface on the Fabric Interconnect.
- **Model**—Displays the Server Model of the Cisco Fabric Interconnect.
- **Expansion Modules**—Displays the total number of expansion modules in the Fabric Interconnect available to expand Ethernet, FCoE, or Fibre Channel ports.
- **Bundle Version**—The firmware bundle version to which the Fabric Interconnect was upgraded.



Note For newly claimed Fabric Interconnects, bundle version will be available only on the subsequent firmware upgrade.

- **NX-OS Version**—The firmware version running on the Fabric Interconnect.
- **User Label**—Displays the assigned user label that helps in identification of the Fabric Interconnect.

- **UCS Domain Profile**—UCS Domain Profile to which the Fabric Interconnect belongs. For standalone servers, this column is not applicable.
- **Ports**—Displays the **Total** ports, **Used** number of ports, and the **Available** ports.
- **Serial**—Displays the host ID of the Fabric Interconnect.
- **Organizations**—Lists the organizations to which the Fabric Interconnect is assigned.
- **Admin Evacuation Mode**—Displays the status of the evacuation mode in Enabled or Disabled state.

Fabric Interconnects Table Summary Dashboard

The following widgets are available in the Fabric Interconnects table view:

- **Health**—The pie chart provides a visual representation of the health of the Fabric Interconnects.
- **Connection**—The badge displays the connection status of the Fabric Interconnects.
- **Bundle Version**—The pie chart displays the firmware bundle version to which the Fabric Interconnect was upgraded.
- **NX-OS Version**—The pie chart displays the firmware version running on the Fabric Interconnect.
- **Models**—The pie chart displays the total number of Fabric Interconnects distributed by the model of the Fabric Interconnect.
- **Contract Status**—The badges display the status of the service contract of the Fabric Interconnects based on the current validity of their associated contracts.

Fabric Interconnects Details View

When you select a Fabric Interconnect in the **Fabric Interconnects** table view, a Details page with information specific to the Fabric Interconnect is displayed. If a Fabric Interconnect is in **Not Connected** status, you can view the device details to resolve the issue. To view further recommendations for troubleshooting, see [Troubleshooting](#).

In addition to the Fabric Interconnect **Health** status, you can view the following information in the Fabric Interconnects Details page:

- **Name**—Displays the name of the Fabric Interconnect.
- **Peer Switch**—Name of the Primary or Subordinate Fabric Interconnect, depending on the device you choose to view. Click **Peer FI** to view the details of the other Fabric Interconnect.
- **User Label**—The assigned user label for the Fabric Interconnect.
- **Model**—The model number of the Fabric Interconnect.
- **Organizations**—Displays the organization to which the Fabric Interconnect is assigned.
- **Expansion Modules**—The number of expansion modules in the Fabric Interconnect.
- **Serial**—The serial number of the Fabric Interconnect.
- **Management IP**—The IP address of the management interface on the Fabric Interconnect.

- **Switch Profile**—The name of the switch profile created for the UCS Domain that belongs to the Fabric Interconnect.
- **Switch Profile Status**—The current status of the switch profile associated with the Fabric Interconnect.
- **Bundle Version**—The firmware bundle version to which the Fabric Interconnect was upgraded.
- **NX-OS Version**—The firmware version running on the Fabric Interconnect.
- **Ports**—The total number of the ports.
- **Used**—The number of used ports
- **Available**—The number of ports available for use.
- **Tags**—The existing tags for the Fabric Interconnects. You can add new tags, or modify the existing ones from **Manage** tags.

The **Properties** area displays a graphical view of the Fabric Interconnect.



Note For Cisco UCS Fabric Interconnects 9108 100G, only the front view with the layout of all ports is displayed.

The **Health Overlay** function enables you to monitor the health of the ports on the Fabric Interconnect. Additionally, this area provides the following information:

- **Mode**—Fabric Interconnects operate in two main switching modes: Ethernet or Fibre Channel. These modes are independent of each other. They determine how the Fabric Interconnect behaves as a device between the server and network/server and storage device.
 - **Ethernet Mode**—The Ethernet switching mode determines how the Fabric Interconnect behaves as a switching device between the servers and the network. The Fabric Interconnect operates in either of the following Ethernet switching modes:
 - **End-Host Mode**—Allows the Fabric Interconnect to act as an end host to the network, representing all servers (hosts) connected to it through virtual Network Interface Cards (vNICs).
 - **Switch Mode**—Allows the Fabric Interconnect to run STP to avoid loops. Broadcast and multicast packets are handled in the traditional way.
 - **FC Mode**—The Fibre Channel switching mode determines how the Fabric Interconnect behaves as a switching device between the servers and storage devices. The Fabric Interconnect operates in either of the following Fibre Channel switching modes:
 - **End-Host Mode**—Allows the Fabric Interconnect to act as an end host to the connected fibre channel networks, representing all servers (hosts) connected to it through virtual Host Bus Adapters (vHBAs).
 - **Switch Mode**—Allows the Fabric Interconnect to connect directly to a storage device.
- **Admin Evac State**—Specifies the evacuation state of Fabric Interconnect traffic. This can be one of the following options:
 - **Disabled**—Restarts traffic on the Fabric Interconnect.
 - **Enabled**—Stops traffic on the Fabric Interconnect.

- **Oper Evac State**—Specifies the operational evacuation state of Fabric Interconnect traffic.
- **FC Zone Count**
 - **FC Zone Limit**—The maximum number of Fibre Channel zones allowed on this Fabric Interconnect.
 - **FC User Zone Limit**—The maximum number of user-created Fibre Channel zones allowed on this Fabric Interconnect.
 - **FC Zone Count**—The number of Fibre Channel zones defined on this Fabric Interconnect.
 - **FC User Zone Count**—The number of user-created Fibre Channel zones defined on this Fabric Interconnect.
- **Access**
 - **IP Address**—The IP address to use when communicating with the Fabric Interconnect.
 - **Subnet Mask**—The subnet mask associated with the IP address.
 - **Default Gateway**—The gateway associated with the IP address.
 - **MAC**—The MAC address.
- **VLAN Details**
 - **VLAN Port Limit**—The maximum number of VLAN ports allowed on this Fabric Interconnect.
 - **Access VLAN Port Count**—The number of available VLAN access ports.
 - **Border VLAN Port Count**—The number of available VLAN border ports.
 - **Compressed Optimization Sets**—The number of VP optimization groups.
 - **Compressed VLAN Port Count**—The number of compressed VLAN ports.
 - **Uncompressed VLAN Port Count**—The number of uncompressed VLAN ports.
 - **Reserved VLAN Range**—The range of VLAN IDs reserved for system use.
- **Fabric Interconnect License**

This section is relevant only to Cisco UCS 6536 Fabric Interconnect and Cisco UCS Fabric Interconnect 9108 100G. These Fabric Interconnects (FIs) support a perpetual software license that activates all ports and software features.



Note This section is not applicable to Cisco UCS 6454 and 64108 Fabric Interconnects, which require port-based licensing. To activate unused ports on these Fabric Interconnects, you need to install a product activation key for each port.

Supported Fabric Interconnects Models

The Fabric Interconnects models supported in Intersight Managed Mode are:

UCS-FI-6454

UCS-FI-64108

UCS-FI-6536

UCSX-S9108-100G

The Fabric Interconnects models supported in UCSM Managed Mode are:

UCS-FI-6248UP, UCS-FI-6296UP

UCS-FI-6332, UCS-FI-6332-16UP

UCS-FI-M-6324

UCS-FI-6454

UCS-FI-64108

UCS-FI-6536

Alarms

Intersight provides fault monitoring capabilities to track and set up alarms for all managed UCS and HyperFlex systems. An alarm alerts you about a failure in the setup (a fault) or a threshold that has been raised. An alarm in Intersight includes information about the operational state of the affected object at the time the fault was raised. Click on a specific alarm to view the fault code, the source type and name, component on which the fault occurred, and a description of the fault.



Note Intersight managed devices must be running with firmware version of 4.1(3) or later releases to generate alarms.

Click on any of the categories to view more details about the alarms.

- **All(Info)**—Displays the total number of faults both Critical and Warning.
- **Critical**—Displays the total number of Critical faults. Raised when a service-affecting condition requires an immediate corrective action. For example, the severity could indicate that the managed object is out of service and its capability must be restored immediately.
- **Warning**—Displays the total number of Warning faults. Raised when a potential or impending service-affecting fault occurs.

This fault could have no significant or immediate effects on the system. A warning status indicates that you must take the appropriate action to diagnose the fault and correct the problem to prevent it from becoming a more serious service-affecting fault.

Fabric Interconnects Inventory View

From the Fabric Interconnects Table View, select a Fabric Interconnect to navigate to its Inventory View.

For the selected Fabric Interconnect, you can view details of each of the following components and perform Traffic Mirroring on the Inventory tab:

- **Ports & Port Channels**—You can see a summary of the Ethernet ports, FC ports, Ethernet Port Channels, and FC Port Channels on the Fabric Interconnect. When you click a specific port, you can view the properties and graphical view of that port.

You can **Enable** or **Disable** a port or a port channel from this view. Disabling a port may lead to traffic disruption. The device connected to a disabled port will also go offline. Disabling a port channel will lead to the member ports also getting disabled.

Reset option from the Fabric Interconnects inventory view allows you to reset a port or an Ethernet port which has server role configuration. The **Reset** action is also available for Backplane Ports on the FEX under the FEX inventory view.



Note This action should be attempted only when the port is not converged due to incorrect configurations. Resetting a port will lead to traffic disruption.

- **Fan Modules**—You can see a summary of the fan modules on the Fabric Interconnect. When you click a specific fan module, you can view the list of fans on the fan module, and the properties and graphical view of that fan module.



Note For Cisco UCS Fabric Interconnects 9108 100G, fan modules are not displayed in this view; they can be viewed under Chassis > Inventory > Fabric Interconnect Module in the Chassis View.

- **PSUs**—You can see a summary of the Power Supply Units (PSUs) on the Fabric Interconnect. When you click a specific PSU, you can view the properties and graphical view of that PSU.



Note For Cisco UCS Fabric Interconnects 9108 100G, PSU modules are not displayed in this view; they can be viewed under Chassis > Inventory > Power in the Chassis View.

- **Local Storage**—You can see a summary of the partitions on the Fabric Interconnect, including details such as their size and current usage. For more information, see *Upgrading Fabric Interconnect Firmware* section of the [Managing Firmware](#) topic.
- **Traffic Mirroring (SPAN)**—You can create a SPAN (Switched Port Analyzer) session that mirrors/copies network traffic from one or more source ports, port-channels, VLANs, and vNICs to a destination port. This traffic can be then analyzed by a Traffic Analyzer. This helps in troubleshooting, performance monitoring, and protocol analysis of network traffic.

You can create, view, edit, or delete a SPAN session from this page. You can view the following details of a SPAN session: Name, Destination, Operational Status, Operational Status Reason, and Sources. For more information, see [Configuring Traffic Mirroring \(SPAN\)](#).

Fabric Interconnects Connections View

The Connections view provides a list of all the components that are directly or indirectly connected to your Fabric Interconnect, such as servers, chassis, and Fabric Extenders (FEX).



Note Cisco UCS Fabric Interconnects 9108 100G currently does not support Fabric Extenders, Rack servers, or adding additional chassis in Intersight Managed Mode.

Depending on the information available for the selected Fabric Interconnect, the following is displayed:

- **Compute**

- **Servers**—The details of all the servers that are connected to the Fabric Interconnect. These details are Name, Health, User Label, Slot Id, Management IP, Model, and Serial.
- **Chassis**—The details of all the chassis that are connected to the Fabric Interconnect. These details are Name, Health, Model, and Serial.



Note For Cisco UCS Fabric Interconnects 9108 100G, chassis actions such as recommissioning, decommissioning, and removal are not permitted.

- **Network**

- **Fabric Extenders**—The details of the Fabric Extenders that are connected to the Fabric Interconnect. These details are Name, Health, Model, Vendor, and Serial.



Note For Cisco UCS Fabric Interconnects 9108 100G, support for Fabric Extender (FEX) is currently not available.

- **Decommissioned**

- **Devices**—The details of decommissioned devices. These details are Type, Model, Serial, Decommissioned Date.

Fabric Interconnects UCS Domain Profile View

The **UCS Domain Profile** view displays a graphic representation of the port configuration, VLAN and VSAN configuration, and the UCS Domain Configuration. Additionally, the following information is displayed:

- **Details**

- **Status**—Status of the UCS Domain profile deployment to the assigned Fabric Interconnect pair
- **Name**

- **Fabric Interconnect A**—Name of Fabric Interconnect A
- **Fabric Interconnect B**—Name of Fabric Interconnect B
- **Last Update**—Date and time that the UCS Domain profile was last updated
- **Description**—Optional description of the UCS Domain profile
- **Tags**—The existing tags for the Domain. You can add new tags, or modify the existing ones from **Manage** tags.
- **Policies**

View the **Policies** that are attached to the UCS Domain profile. The **Policies** pane displays details of the **Port, VLAN and VSAN, and UCS Domain Configuration**. A graphical representation of the ports configuration on the Fabric Interconnects, including port roles and port channels and a list of associated policies is displayed. The VLAN, VSAN, and UCS Domain Configuration lists the Domain policies associated with the selected Domain profile.



Note For Cisco UCS Fabric Interconnects 9108 100G, there are updates in the Port Configuration. To know more about the updates to the Port policy and Switch Control Policy, see [Creating a Port Policy](#) and [Creating a Switch Control Policy](#).

Fabric Interconnect Topology View

For more information, see [Topology](#).

Fabric Interconnect Metrics View

For more information, see [Fabric Interconnects Metrics](#).

Fabric Interconnect Actions

The Fabric Interconnect Actions allows you to perform specific management operations on that Fabric Interconnect. In Cisco Intersight, when you click a Fabric Interconnect, the Fabric Interconnects Table view is displayed. In this page, click the Ellipsis (...) icon to perform Fabric Interconnect actions.

Fabric Interconnect Actions: You can perform the following operations to manage a Fabric Interconnect:

- **Enable/Disable Evacuation Mode**—Fabric Evacuation refers to the ability that allows you to stop all active traffic flowing through the selected Fabric Interconnect. You can use the **Enable Evacuation Mode** to evacuate all the Ethernet and Fibre Channel traffic flowing through the selected Fabric Interconnect, from all blades and rack servers. This Evacuation Mode option provides more control so that you can perform maintenance operations on the Fabric Interconnect with Evacuation Mode enabled, and also test the traffic high availability behavior during the setup.

Use the **Disable Evacuation Mode** to restore the traffic to flow through both the Fabric Interconnects paths.

An alarm is raised on the Fabric Interconnect when the Evacuation Mode is enabled. You can monitor the progress of the traffic evacuation in the Requests view. The evacuation state of the Fabric Interconnect is displayed in the **Fabric Interconnects Table View** and the **Fabric Interconnects Details View > Inventory** tab. The alarm gets cleared when the Evacuation Mode is disabled. For more information, see [Cisco Intersight Alarms Reference Guide](#).

- If Traffic Evacuation Mode is **Enabled**, the IOM or FEX backplane ports or ports towards the servers are set to **admin-down**. In case of directly attached rack servers, the server ports on the Fabric Interconnect connected to these rack servers are set to **admin-down**.
- During Fabric Evacuation for the Intersight Managed Mode domains with direct attached rack servers, the server ports are shut down. In case of Intersight Managed Mode domains with IOM or IFM and rack servers behind the FEX, the HIF ports are shut down.
- You must explicitly configure Traffic Evacuation Mode as **Disabled** to move the backplane ports or server ports on Fabric Interconnect connected to the rack servers back to the **Up** state and resume the traffic flow.



Note

- If traffic evacuation mode is enabled on a Fabric Interconnect within a pair, it will be disabled on its peer Fabric Interconnect. By default, Traffic Evacuation Mode is disabled.
 - You can toggle the **Force Evacuation** option when one or more of the peer Fabric Interconnect IOMs/FEX are not operable or disconnected.
-

Traffic evacuation for a Fabric Interconnect is not allowed if it can cause a traffic outage for servers in the domain. For example, if a chassis has connectivity to only one Fabric Interconnect, and you evacuate this Fabric Interconnect, servers on this chassis loses the connectivity to the rest of the data center. Select **Force Evacuation** if you want to override this restriction and proceed.



Note

You cannot perform evacuation on Fabric Interconnect when:

- The Peer Fabric Interconnect is disconnected.
 - The Fabric Interconnect firmware upgrade is in progress.
 - The domain profile deployment is in progress.
-

Traffic evacuation on a Fabric Interconnect can take 7 to 10 minutes to complete, depending on the system scale. During this time, traffic evacuation to the other Fabric Interconnect can take 1 to 2 seconds for an individual vNIC or vHBA. Ensure to configure redundancy for vNIC or vHBA to minimize traffic disruption.

- **Reboot**— Allows you to restart a Fabric Interconnect for recovery and troubleshooting purposes. By default, this action evacuates the data traffic for the selected Fabric Interconnect. All the traffic in the Fabric Interconnect stops and the traffic on failover enabled vNICs fails over to the peer Fabric Interconnect. The request will complete after the Fabric Interconnect configuration is restored and all

evacuated paths have come back up. To bypass the traffic evacuation during the Fabric Interconnect Reboot request, deselect the Evacuate Fabric Interconnect traffic before reboot check box.



Note For Cisco UCS X-Series Direct, you can reboot the Cisco UCS Fabric Interconnects 9108 100G either from the Fabric Interconnect Module list view or the Chassis Inventory view.

A reboot action could cause a traffic outage for servers in the domain and will not be allowed in such cases. You can enable the Force Reboot option to override and proceed with the reboot in the following cases:

- The peer Fabric Interconnect has been manually evacuated or a manual evacuation request is in progress.
- One or more of the peer Fabric Interconnect IOMs/FEX are not operable or disconnected.
- A Domain Profile deployment is in progress.
- The peer Fabric Interconnect is unreachable.



Note

- The Reboot or Force Reboot action is not allowed when a firmware upgrade is in progress.
- The Reboot or Force Reboot action will not initiate a Fabric Interconnect Reboot action again on the same Fabric Interconnect while a prior reboot is in progress.

- **Turn On Locator**— Turns on the locator LED on the device.



Note This action is disabled for Cisco UCS Fabric Interconnects 9108 100G and moved to the corresponding Chassis View.

- **Launch UCS Manager**— Launch Cisco UCS Management interface from Cisco Intersight.
- **Launch CLI**—Launch command line interface from Cisco Intersight.



Note Launch Cisco UCS Manager and Launch CLI options are available only for UCSM Managed Mode Fabric Interconnects.

- **Create Traffic Mirroring (SPAN) Session**—Create a traffic mirroring or SPAN (Switched Port Analyzer) session that copies network traffic from one or more source ports, port-channels, VLANs, and vNICs to a destination port. This traffic can be then analyzed by a Traffic Analyzer. This helps in troubleshooting, performance monitoring, and protocol analysis of network traffic.

You can edit or delete a SPAN session from Fabric Interconnects Inventory View. For more information, see [Configuring Traffic Mirroring \(SPAN\)](#).

- **Open TAC Case**—Open a case to report an issue with the server.
- **Upgrade Firmware**—Perform a firmware upgrade. For more information, see the [Firmware Upgrade](#).
- **Set User Label**—Allows you to set, update, or delete user labels for each Fabric Interconnect. It must be between 1 and 64 alphanumeric characters, containing only the following special characters: ! # \$ % & * + , () [] { } | / . ? @ _ : ; ~
- **Replace Fabric Interconnect**—Remove the old Fabric Interconnect and connect the new Fabric Interconnect.
- **Replace UCS Domain**—Remove the old Fabric Interconnect cluster and connect the new Fabric Interconnect cluster.



Note

- Replace Fabric Interconnect and Replace UCS Domain options are available only for Intersight Managed Mode Fabric Interconnects.
 - To perform Replace UCS Domain action, you must ensure that the Domain Profile is in associated state with the Fabric Interconnect.
-

- **Collect Tech Support Bundle**—Collect the tech support bundle. An Account Administrator or a user with Support Services role can select the device and collect the tech support bundle file for the selected device. The downloaded file can be accessed by navigating to **Admin > Tech Support Bundles** section. This file can be shared with the TAC team for troubleshooting any issue.

Configuring Traffic Mirroring (SPAN)

A Traffic Mirroring or SPAN (Switched Port Analyzer) session copies the network traffic from one or more sources ports, ports channel, VLANs, and vNICs to a designated local Switched Port Analyzer (SPAN) destination. The Cisco Intersight supports an Ethernet Traffic Mirroring session, for which the destination port must be an unconfigured ethernet port. An unconfigured port is a port that does not have any role assigned to it. You can connect this port to a traffic analyzer, facilitating troubleshooting, performance monitoring, and protocol analysis of network traffic.

Supported SPAN sources include uplink Ethernet ports, uplink Ethernet port channels, VLANs, and vNICs. A SPAN session can be configured for ingress, egress, or both directions.

Guidelines and Recommendations for Traffic Mirroring

While configuring or activating Traffic Mirroring, consider the following guidelines:

- You can create up to 16 SPAN sessions per Fabric Interconnect, out of which, only four sessions can be active at a time.
- A maximum of 128 sources, including up to 32 VLANs, can be added to a single SPAN or mirroring session.
- You can configure a port as a destination port for only one mirroring session.

- Ports and port-channels may be configured to mirror traffic in receive, transmit, or both directions. However, VLANs and vNICs are restricted to mirroring traffic in the receive direction only.
- All SPAN-related actions are disabled for a disconnected Fabric Interconnect. However, SPAN session will continue to copy packets while Fabric Interconnect is operational even though the Data Center (DC) is disconnected.
- SPAN sessions will be removed from an Fabric Interconnect that is unclaimed or replaced under RMA (Return Material Authorization).
- A minimum privilege level of Domain Administrator is required. For more information on privileges, see [Roles and Privilege](#).
- To SPAN traffic on both the primary and backup paths for failover enabled vNICs, add the selected vNIC as a SPAN source on both Fabric Interconnects.
- To SPAN all traffic on a VLAN common to both Fabric Interconnects, add the selected VLAN as a SPAN source on both Fabric Interconnects. SPAN functionality for private VLANs is not supported.

Creating an Ethernet Traffic Mirroring Session

You can create an Ethernet Traffic Mirroring session from Fabric Interconnect Inventory View or from Fabric Interconnect Actions.

To create an Ethernet Traffic Mirroring session:

1. Log in to Cisco Intersight with your Cisco ID and select Domain or Account Administrator role.
2. From the **Service Selector** drop-down list, select **Infrastructure Service**.
3. Go to **Fabric Interconnect** and select the desired Fabric Interconnect from the Fabric Interconnect Table View.
4. Do one of the following
 - Click **Inventory > Traffic Mirroring (SPAN) > Create**.
 - From the Fabric Interconnects Table View, click the ellipsis (...) icon of the desired Fabric Interconnect, then select **Create Traffic Mirroring (SPAN) Session**.
5. On the **General** page, configure the following parameters:

Name	Description
Name	The name of the traffic mirroring session. 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.
Enable Session	Toggle to enable or disable the session. It is enabled by default.

Name	Description
Enable SPAN Control Packets	Toggle to enable mirroring of outgoing control packets that are sent from the Fabric Interconnect's CPU.
Destination	
Select Port	Select the destination port from the available ports. Available ports are represented by white dots. Note Destination port must be an unconfigured ethernet port. An unconfigured port is a port that does not have any role assigned to it.
Admin Speed	Select the data transfer rate of the port channel to be mirrored.
FEC	The Forward Error Correction (FEC) configuration for the port: <ul style="list-style-type: none"> • Auto • Cl91—Supported with 25Gbps and 100Gbps Admin speed • Cl74—Supported with 25Gbps Admin speed

6. Click **Next**. The **Source Configuration** page appears.
7. On the **Source Configuration** page, click **Add** to select and add the sources to the mirroring session.

**Note**

- The SPAN source ports, port-channels, and VLANs will be automatically deleted during Domain Profile deployment if you unconfigure the corresponding port, port-channel, or VLANs.
- The SPAN source vNIC will be automatically deleted if you delete the vNIC.
- The SPAN source vNIC will be automatically deleted during Server Profile deployment if the corresponding vNIC is no longer placed on that Fabric Interconnect.

8. Click **Next**. The **Summary** page appears.
9. On the **Summary** page, verify the configuration details and click **Create**.

Editing or Deleting an Ethernet Traffic Mirroring (SPAN) Session

You can edit or delete an Ethernet Traffic Mirroring session from Fabric Interconnect Inventory View.

To edit or delete a SPAN session:

1. Log in to Cisco Intersight with your Cisco ID and select Domain or Account Administrator role.
2. From the **Service Selector** drop-down list, select **Infrastructure Service**.

3. Go to **Fabric Interconnect** and select the desired Fabric Interconnect from the Fabric Interconnect Table View.
4. Click **Inventory > Traffic Mirroring (SPAN)**.
5. Select the session that you want to edit or delete and do one of the following:
 - Click the **Edit** icon, make the necessary edits, and then save your changes.
 - Click the **Delete** icon.



Note The SPAN session will be asynchronously deleted from Cisco Intersight only after it is successfully removed from the Fabric Interconnect.



CHAPTER 3

Chassis and FEX Lifecycle

- [Chassis and Fabric Extender Discovery and Actions](#), on page 49
- [Chassis Table View](#), on page 51
- [Chassis Details View](#), on page 52
- [Chassis Inventory View](#), on page 53
- [Chassis Connections View](#), on page 54
- [Chassis Topology View](#), on page 54
- [Chassis Metrics View](#), on page 54
- [Chassis Actions](#), on page 54
- [Fabric Extender Details View](#), on page 55
- [Fabric Extender Inventory View](#), on page 56
- [Fabric Extender Connection View](#), on page 56

Chassis and Fabric Extender Discovery and Actions

Chassis and Fabric Extender Discovery

Chassis and Fabric Extenders (FEX) that are connected to a Fabric Interconnect are automatically discovered in Cisco Intersight. To discover chassis and FEX connected to a Fabric Interconnect, ensure that the Fabric Interconnect is claimed in Cisco Intersight.

After the Fabric Interconnect is claimed, do the following:

1. Connect the server ports to both Fabric Interconnects. For example, ports 1 and 2 to FI-A and ports 3 and 4 to FI-B.
2. Configure the server ports on both Fabric Interconnects by using a UCS Domain profile. *Creating a UCS Domain Profile* provides detailed information about creating a UCS Domain profile and assigning it to a UCS Fabric Interconnect Domain.

After the server ports are configured and applied, all the chassis and FEX that are connected to the Fabric Interconnect are automatically discovered. During discovery, the chassis and FEX will auto sync firmware with the Fabric Interconnect if their firmware versions do not match the firmware version of the Fabric Interconnect. Because of this, it may take 25-30 minutes for the chassis and FEX to appear in the GUI. You can check the chassis and FEX status through the nxos CLI by using the show fex command.

Chassis Actions

From the left navigation panel, click **Chassis** for the **Chassis** table view. You can perform the following operations to manage one or more chassis.

Chassis Actions

- **Rediscover**—Initiates the chassis discovery process and then the chassis inventory process.
- **Decommission**—Removes the chassis and IOM inventories. A decommissioned chassis is likely to be eventually recommissioned, a part of the chassis information, including the chassis ID, is retained by Cisco Intersight. Decommissioning is performed when a chassis is physically present and connected, but you want to temporarily remove it from the Cisco Intersight configuration.
- **Remove**—Removes the configuration of a physically removed chassis from Cisco Intersight.

Before physically removing a chassis from the system, ensure that you unconfigure the server ports to which the chassis is connected.

If you need to add a chassis, which was earlier removed, back to the Cisco Intersight configuration, the chassis must be reconnected and then rediscovered. During rediscovery Cisco Intersight will assign the chassis a new ID that may be different from ID that it was assigned earlier.

- **Recommission**—Brings the chassis and IOM back online and initiates the chassis discovery process and then the chassis inventory process. After this action is complete, you can access the chassis and any servers in it.

A list of decommissioned chassis is available in the **Devices** area under **Fabric Interconnects > Fabric Interconnect Name > Connections > Decommissioned**.

When you recommission the chassis, you have the option to configure the chassis ID.

- **Turn On/Off Locator**—Turns on/off the LED Locator.



Note This option is available only for Intersight Managed Mode servers.

- **Power Cycle Chassis Slot**—Power cycling a chassis slot is an attempt to recover a non-responding device. This operation can be initiated from Chassis Table View and Chassis Connections View by clicking on the Ellipsis (...) icon.



Note Power cycling a server slot will bring down the server, potentially impacting the application services. Therefore, this option should be utilized with caution to debug any issues with a server in that chassis slot.

- **Upgrade Firmware**—This action is supported only for Cisco UCS S3260 chassis.
- **Collect Tech Support Bundle**—Collects the tech support bundle. An account administrator can select the device and collect the tech support bundle file for the selected device. The downloaded file can be accessed by navigating to **Admin > Tech Support Bundles** section. This file can be shared with the TAC team for troubleshooting any issue.

FEX Actions

From the left navigation panel, click **Fabric Interconnects** > *Fabric Interconnect Name* > **Connections** > **Fabric Extenders** for the **FEX** table view. You can perform the following operations to manage one or more FEX.

FEX Actions

- **Decommission**—Decommissioning is performed when a FEX is physically present and connected, but you want to temporarily remove it from the Cisco Intersight configuration. This action takes the FEX offline and removes the FEX inventory. Because it is expected that a decommissioned FEX will be eventually recommissioned, a portion of the FEX information is retained by Cisco Intersight.
- **Remove**—Removing a FEX involves physically removing a FEX from the system. After the physical removal of the FEX is completed, the configuration for that FEX is removed from Cisco Intersight.

To add a removed FEX back to the Cisco Intersight configuration, it must be reconnected to server ports that are configured on the Fabric Interconnect. The FEX is automatically discovered. During discovery Cisco Intersight will assign the FEX a new ID that may be different from ID that it was assigned earlier.

- **Recommission**—Recommissioning a FEX brings the FEX back online, initiates the FEX discovery process and then the FEX inventory process. After this action is complete, you can access the FEX.

A list of decommissioned FEX is available in the **Devices** area under **Fabric Interconnects** > *Fabric Interconnect Name* > **Connections** > **Decommissioned**.

- **Turn On Locator**—Turn on the LED Locator on the selected FEX. Locators are indicators that help direct administrators to specific nodes in large data center environments.
- **Turn Off Locator**—Turn off the LED Locator on the selected FEX. Locators are indicators that help direct administrators to specific nodes in large data center environments.

Chassis Table View

From the **Service Selector** drop-down list, select **Infrastructure Service**. Navigate to **Operate** > **Chassis**, to launch the Chassis Table view. From this page, you can perform chassis actions, and navigate to the chassis details page. Click the **Settings** icon (the gear icon), and select the columns that you want in the Table view.

You can add specific columns or custom tags to the Chassis Table view to sort and filter the information.

You can view the following details in the Chassis Table view:

- **Name**—Name of the chassis. For UCSM Managed Mode (UMM) and Intersight Managed Mode (IMM) servers, the name is a combination of UCS domain name and chassis ID. The chassis ID is auto-assigned depending on the order in which the chassis are discovered. To configure the chassis ID, decommission the chassis and recommission it. During recommission, you can assign the chassis an ID of your choice.
- **Health**—The chassis health corresponds to the alarms on the chassis. For more details, see [Alarms](#).
- **Contract Status**—Displays the status of the service contract for the chassis based on the current validity of the associated contracts. You can identify the SmartNet Contract ID details of the chassis, and cross launch the [Cisco Commerce Software Subscriptions and Service Portal](#).
- **Chassis ID**—Displays the chassis ID.

- **UCS Domain**—UCS Domain to which the chassis belongs. For standalone servers, this column is not applicable.
- **Model**—Displays the chassis model.
- **User Label**—Displays the assigned user label that helps in identification of the chassis.
- **Serial**—Displays the host ID/serial number of the chassis.
- **Chassis Profile**—Displays the chassis profile that is associated with the chassis.
- **Management Mode**—The management mode of the chassis.

Chassis Table Summary Dashboard

The following widget is available in the Chassis Table view:

- **Contract Status**—The badges display distribution of the service contract status of managed UCS and Hyperflex clusters based on the validity of their associated contracts.

Chassis Details View

When you select a chassis in the chassis table view, a Details page with information specific to the chassis is displayed. In addition to the chassis **Health** status, you can view the following information in the Chassis Details page:

- **Name**
- **Serial**—The serial number of the chassis
- **Model**—The model number of the chassis, for example, UCSB-5108-AC2
- **Revision**—The revision number of the chassis
- **Part Number**—The part number of the chassis
- **Management Mode**—The management mode of the chassis.
- **Contract Status**—The contract status of the managed chassis based on the current validity of their associated contracts.
- **UCS Domain**—The name of the UCS Domain of which the selected chassis is a part
- **Chassis Profiles**—Displays the associated Chassis Profile configuration status.
- **Tags**—The existing tags for the selected object are displayed by default. Click **Manage** to add new tags or modify the existing ones.

The **Properties** area provides a graphical representation of the front and rear view of the chassis, the health overlay for the chassis, and an overview of the hardware properties of the chassis and its components.

Note:

- The Chassis Details View is supported on the Cisco UCS S3260 Chassis and Intersight Managed Mode chassis.

- For Cisco UCS X-Series Direct, the rear view of the chassis (UCSX-9508) is displayed, which is same as the rear view of the Cisco UCS Fabric Interconnects 9108 100G.

The **Alarms** area in Cisco Intersight provides fault monitoring capabilities to track and set up alarms for all managed UCS systems. An alarm alerts you about a failure in the endpoint (a fault) or a threshold that has been raised.

Chassis Inventory View

After a chassis is discovered, an inventory of all its components is made available. When you select a chassis in the **Chassis** table view, you can view the inventory of its components in the **Inventory** tab.

For the selected chassis, you can view details of each of the following components:

- **IO Modules**—You can see the name, vendor, model number, management IP address, operational state, and firmware version of the IO modules in the chassis. When you click a specific IO module, you can view the General properties, details of the Backplane Ports and Fabric ports, a Graphic View, and Health Overlay of the IO module.

Action: You can reset an IO module or its peer IO module from Chassis Inventory View. Resetting the peer IO module through the corresponding IO module initiates a reboot of peer IO module. This helps to recover a peer IO module that is not directly reachable from Intersight.

- In the Cisco UCS X-Series chassis, each Intelligent Fabric Module (IFM) contains fan modules. When you click a fan module, you can view the properties and operational state of the fans.
- For Cisco UCS X-Series Direct chassis (UCSX-9508), the IFM section is labeled 'Fabric Interconnect Module (Intelligent Fabric Module)'. Within the IFM section, the 'Fabric Ports' subsection is not present.
- **XFM Modules**—You can see a summary of the X-Fabric Modules (XFM) in the chassis. Click a specific XFM to view the details of the fan modules. When you click a fan module, you can view the ID, model, and operational state of the fans.



Note The XFM (UCSX-F-9416) slots are to be present in the respective two slots of UCSX 9508 Chassis.

- **Thermal**—The **General** tab in the Thermal section displays the thermal configuration and statistics. The **Fan Modules** tab displays the name, number of fans, model number, and operational state of the fan modules. When you click a specific fan module, you can view the general details, fan details, graphic view, and health overlay of the fan module.
- **Power**—The **General** tab in the Power section displays the power configuration and statistics. The **PSUs** tab displays the name, model number, vendor name, serial number, and operational state of the PSUs (Power Supply Units). When you click a specific PSU, you can view the general details, graphic view, and health overlay of the PSU.
- **Servers**—You can see the name, slot ID, model number, and serial number of the servers for the selected chassis.

Chassis Connections View

The Connections view provides a list of all the components that are directly or indirectly connected to your chassis, such as Fabric Interconnects and servers.

Depending on the information available for the selected chassis, the following is displayed:

- **Network**
 - **Switches**—Displays the details of the Fabric Interconnects that are connected to the chassis. These details are Name, Health, Model, Vendor, and Serial.

Chassis Topology View

For more information, see [Topology](#).

Chassis Metrics View

For more information, see [Chassis Metrics](#).

Chassis Actions

From the left navigation panel, click **Chassis** for the **Chassis Table View**. In this page, click the Ellipsis (...) icon to perform the following operations to manage one or more chassis. These actions can also be performed for the selected chassis in the Chassis Details View.

- **Rediscover**—Initiates the chassis discovery process and then the chassis inventory process.
- **Decommission**—Removes the chassis and IOM inventories. A decommissioned chassis is likely to be eventually recommissioned, a part of the chassis information, including the chassis ID, is retained by Cisco Intersight. Decommissioning is performed when a chassis is physically present and connected, but you want to temporarily remove it from the Cisco Intersight configuration.



Note The action to recommission, decommission, and remove are not available for the chassis of Cisco UCS X-Series Direct.

- **Remove**—Removes the configuration of a physically removed chassis from Cisco Intersight.

Before physically removing a chassis from the system, ensure that you unconfigure the server ports to which the chassis is connected.

If you need to add a chassis, which was earlier removed, back to the Cisco Intersight configuration, the chassis must be reconnected and rediscovered. During rediscovery Cisco Intersight will assign the chassis a new ID that may be different from the ID that it was assigned earlier.

- **Turn On/Off Locator**—Turns on/off the LED Locator.



Note This option is available only for Intersight Managed Mode servers.

- **Set User Label**—Allows you to set, update, or delete user labels for each chassis. It must be between 1 and 64 alphanumeric characters, containing only the following special characters: ! # \$ % & * + , () [] { } | / . ? @ _ : ; ~
- **Power Cycle Chassis Slot**—Power cycling a chassis slot is an attempt to recover a non-responding device. This operation can be initiated from Chassis Table View and Chassis Connections View by clicking on the Ellipsis (...) icon.



Note Power cycling a server slot will bring down the server, potentially impacting the application services. Therefore, this option should be utilized with caution to debug any issues with a server in that chassis slot.

- **Upgrade Firmware**—This action is supported only for Cisco UCS S3260 chassis.
- **Recommission**—Brings the chassis and IOM back online and initiates the chassis discovery process and then the chassis inventory process. After this action is complete, you can access the chassis and any servers in it.

A list of decommissioned chassis is available in the **Devices** area under **Fabric Interconnects > Fabric Interconnect Name > Connections > Decommissioned**.

When you recommission the chassis, you have the option to configure the chassis ID.



Note

- The action to recommission, decommission, and remove are not available for the chassis of Cisco UCS X-Series Direct.
- Reset IOM and Reset Peer IOM actions are not available for Cisco UCS X-Series Direct.

- **Collect Tech Support Bundle**—Collect the tech support bundle. An Account Administrator or a user with Support Services role can select the device and collect the tech support bundle file for the selected device. The downloaded file can be accessed by navigating to **Admin > Tech Support Bundles** section. This file can be shared with the TAC team for troubleshooting any issue.

Fabric Extender Details View

When you select a Fabric Extender (FEX) in the FEX table view, a Details page with information specific to the chassis is displayed. In addition to the FEX **Health** status, you can view the following information in the FEX Details page:

- **Name**
- **Serial**—The serial number of the Fabric Extender

- **Model**—The model number of the Fabric Extender
- **Vendor**—The name of the manufacturer
- **Revision**—The revision number of the Fabric Extender
- **Part Number**—The part number of the Fabric Extender
- **Ports**—The total number of ports on the Fabric Extender, and their operational status. The status can be:
 - **Used**—Number of ports that are currently connected to the Fabric Interconnects and servers
 - **Available**—Number of ports available for use on the Fabric Extender
- **Tags**—The existing tags for the Fabric Extender. You can add new tags, or modify the existing ones from **Manage** tags.

Fabric Extender Inventory View

After a Fabric Extender (FEX) is discovered, an inventory of all its components is made available. When you select a FEX in the **FEX** table view, you can view the inventory of its components on the **Inventory** tab.

For the selected FEX, you can view details of each of the following components:

- **Ports**—The details of all the **Backplane Ports** and **Fabric Ports** on the FEX that is selected.

The **Backplane Ports** table shows the server ports, which are host ports. This includes information such as the port **Name**, **Status**, **Port Channel ID** to which it belongs, **Speed** of the port and the **Peer** server port.

The **Fabric Ports** table shows the network ports that are connected to the Fabric Interconnect. This includes information such as the port **Name**, **Status**, **Port Channel ID** to which it belongs, **Switch Slot ID** of the Fabric Interconnect to which it is connected, the **Peer** Fabric Interconnect, and the **Switch Port ID** of the Fabric Interconnect to which it is connected.

It also includes detailed hardware information and graphic view of each port.

- **Fan Modules**—The details of all the fan modules on the FEX, such as **Name**, **Fans**, **Model**, and **Status**.

It also includes detailed hardware information and graphic view of each fan module and the fans in it.

- **PSUs**—The details of the Power Supply Units (PSUs) on the FEX, such as **Name**, **ID**, **Model**, **Vendor**, **Serial** and **Status**.

It also includes detailed hardware information and graphic view of each PSU.

Fabric Extender Connection View

The Connections view provides a list of all the components that are directly or indirectly connected to your Fabric Extender (FEX), such as servers and Fabric Interconnects.

Depending on the information available for the selected FEX, the following is displayed:

- **Compute**

- **Servers**—The details of all the servers that are connected to the FEX. These details are Name, Health, User Label, Model, and Serial.
- **Network**
 - **Switches**—Displays the details of the Fabric Interconnects that are connected to the FEX. These details are Name, Health, Model, Vendor, and Serial.



CHAPTER 4

Server Lifecycle

- [Server Discovery and Actions, on page 59](#)
- [Servers Table View, on page 64](#)
- [Server Details View, on page 67](#)
- [Server Inventory View, on page 68](#)
- [Server Topology View, on page 72](#)
- [Server Metrics View, on page 72](#)
- [Server Connectivity View, on page 73](#)
- [Compliance with Hardware Compatibility List \(HCL\), on page 73](#)

Server Discovery and Actions

After a chassis or FEX is discovered, the blade servers connected to the chassis or the rack servers connected to the FEX are automatically claimed and discovered. *Chassis and FEX Discovery and Operations* provides information about this process. For servers to be claimed and discovered, they must be in the factory default state.

For rack servers that are directly attached to the Fabric Interconnect, do the following after the Fabric Interconnect is claimed:

1. Connect the server ports to both Fabric Interconnects. For example, ports 1 and 2 to FI-A and ports 3 and 4 to FI-B.
2. Configure the server ports on both Fabric Interconnects.

The servers that are discovered appear on the **Servers** page.

Server Actions

The server actions enable you to manage the server. In Cisco Intersight, when you click on Servers, the Servers Table view is displayed. In Servers Table view page, click the ellipsis (...) icon to perform server actions.

You can perform the following operations to manage a server:

- **Power**
 - **Power On/Off**—Turns on/off the power of the server.
 - **Power Cycle**—Turns off and on for the server.

- **Hard Reset**—Reboots the server.
- **Shut Down OS**—Shuts down the server if supported by an operating system.

- **System**

- **Turn On/Off Locator**—Turns on/off the LED Locator.
- **Reset CMOS**—Resets the BIOS configuration settings to the original state hence helps in recovery when the server is not in a healthy state. The option to reset CMOS appears only when the server is powered off. For the reset to complete, the server must be powered on. There is an additional option to power on the server using the toggle button present in the Reset CMOS confirmation window.



Note This option is available only for Intersight Managed Mode servers.

- **Reset vKVM**—Resets the vKVM on the server. Clicking this option triggers a workflow that restarts the vKVM server running on Cisco IMC. You can use this action when the launch of vKVM is failing or when a running vKVM becomes unresponsive.



Note Reset vKVM is supported only on Firmware version 4.2(1a) and later. For earlier versions, Reboot Cisco IMC option should be used.

- **Secure Erase**—Initiates a workflow to delete all data from the server. This includes data on the BIOS, BMC, NVRAM, DIMM, eMMC, VIC, and storage disks components (except remote LUNs). During this time, disruptive server actions, such as power actions, OS installation, firmware upgrades, decommissioning, and server profile deployment are disabled.

Depending on the disk capacity, Secure Erase may take several hours to more than a day to complete. For more information, see [Secure Data Deletion](#).



Warning

- This action will permanently delete all data from the server. This data will not be recoverable
 - If Secure Erase is terminated before it is complete, data or drive corruption may occur.
-

- **Lock Front Panel**—Locks the physical power button on the server. For a server that already has the front panel locked, this option appears as **Unlock Front Panel**.



Note This option is available only for Intersight Managed Mode servers.

- **Rediscover**—Rediscover rebuilds the server inventory and internal connectivity. The operation can be used following any module replacements, additions, or removals, as well as when a server is relocated to another slot, chassis, or replaced. This operation will not disrupt any ongoing traffic.

- **Decommission**—Decommissions the server and removes the server from the Cisco UCS configuration. However, the server hardware physically remains in the Cisco UCS instance.

You can view the decommissioned servers by selecting **Operate > Fabric Interconnects** and clicking on the Fabric Interconnect that has the decommissioned servers. Go to the **Connections** tab and in the **Decommissioned** section click **Servers**. Click the **Ellipses (...)** to view the Server Actions.

Click **Recommission** to recommission the server. When you recommission a rack server, you have the option to configure the server ID.

To trigger the Secure Erase action before decommissioning a server, select the **Decommission Server With Secure Erase** check box. Use this option only if you want to delete all data from the server before it is decommissioned. Depending on the disk capacity, Secure Erase may take several hours to more than a day to complete. For more information, see [Secure Data Deletion](#).



Warning

Selecting the **Decommission Server With Secure Erase** check box will permanently delete all data from the server. This data will not be recoverable.

- **Reboot Cisco IMC**—Reboots the Cisco IMC.
- **Certificate:**
 - **Set KMIP Client Certificate**—To Configure a KMIP client certificate to ensure secure communication between the KMIP server and Cisco IMC. For more information, see [Configuring KMIP Client Certificate](#).
 - **IMC certificates**—To configure the certificate and private key on the server from a third-party managed Certificate Authority(CA). This option is available only for Intersight Managed Mode servers. For more information, see the *Creating a Certificate Management Policy* section in [Supported UCS Server Policies](#).
- **Reset Memory Errors**—Allows you to reset memory error counters for all Dual In-line Memory Modules (DIMMs) in a server to zero. After clearing errors, any DIMMs disabled due to prior uncorrectable errors will be re-enabled on the next boot, provided they encounter no further errors during Power-On Self-Test (POST).
- **Set Asset Tag**—Enables to set the custom asset tag.
- **Set User Label**—Allows you to set, update, or delete user labels for each server. It must be between 1 and 64 alphanumeric characters, containing only the following special characters: ! # \$ % & * + , () [] { } | / . ? @ _ : ; ~
- **Download System Event Log**—Downloads the system event logs of a selected server. These logs record server-related events, such as over and under voltage, temperature, and fan events.
- **Clear System Event Log**—Clears the system event logs of a selected server.
- **Install Operating System**—Perform an unattended OS installation on one or more Cisco UCS C-Series Standalone servers and Cisco Intersight Managed Mode (IMM) servers (C-Series, B-Series, and X-Series) from your centralized data center through a simple process. For more information, see [Installing an Operating System](#).
- **Upgrade Firmware**—Perform a firmware upgrade. For more information, see the [Firmware Upgrade](#).

- **Launch IMC**—Cross-launch Cisco Integrated Management Controller (CIMC) UI from Intersight. This action is available only for C-Series Standalone servers.



Note *Generate Technical Support Data for Local Download* and *Download Hardware Inventory Data to Local Download* options are not supported in the cross-launched CIMC interface.

- **Launch Virtual KVM**—Launch the virtual keyboard, video, and mouse (KVM) console directly for Fabric Interconnect-attached and Standalone server. Local network connectivity to the endpoint/server is required. For more information, see [Launch Virtual KVM](#).
- **Launch Tunneled Virtual KVM**—Tunneled vKVM works by tunneling the KVM traffic through Intersight. You can launch Tunneled vKVM sessions for all servers in Intersight Managed Mode, Cisco UCS C-Series Standalone M4, M5, M6, M7, and M8 servers, UCS S-Series, and Hyperflex HX-Series Edge Standalone M4 and M5 servers. For more information, see [Launch Tunneled Virtual KVM](#)
- **Start/Stop Alarm Suppression**—Select the server and choose the *Start Alarm Suppression* option to silence the non-critical alarms, or select the *Stop Alarm Suppression* option to resume receiving all alarm notifications. For more information, see [Alarm Suppression](#).
- **Open TAC Case**—Open a case to report an issue with the server.
- **Set License Tier**—Update the server to a new license tier. Updating license tier is not allowed on server(s) with an associated server profile. To move the license to another tier, unassign the profile from one or more selected servers and try again. For more information, see [Multiple Licensing Tiers](#).
- **Collect Tech Support Bundle**—Collect the tech support bundle. An Account Administrator or a user with Support Services role can select the device and collect the tech support bundle file for the selected device. The downloaded file can be accessed by navigating to **Admin > Tech Support Bundles** section. This file can be shared with the TAC team for troubleshooting any issue.

Bulk Server Actions

On the **Servers** table page, you can perform the following operations to manage more than one server.

- **Power**
 - **Power On**—Power on for one or more servers.
 - **Power Off**—Turn power off for one or more servers.
 - **Power Cycle**—Turns power off and on for one or more servers.
 - **Hard Reset**—Reboot the server.
 - **Shut Down OS**—Shuts down the server if supported by an operating system.
- **System**
 - **Turn On Locator**—Turn on the LED Locator.
 - **Turn Off Locator**—Turn off the LED Locator.
 - **Reset CMOS**—Resets the BIOS configuration settings to the original state hence helps in recovery when the server is not in a healthy state. The option to reset CMOS appears only when the server

is powered off. For the reset to complete, the server must be powered on. There is an additional option to power on the server using the toggle button present in the Reset CMOS confirmation window.



Note This option is available only for Intersight Managed Mode servers.

- **Reset vKVM**—Resets the vKVM on the server. Clicking this option triggers a workflow that restarts the vKVM server running on Cisco IMC. You can use this action when the launch of vKVM is failing or when a running vKVM becomes unresponsive.



Note This action is supported only on Firmware version 4.2(1a) and above. For lower versions, Reboot Cisco IMC option should be used.

- **Lock Front Panel**—Locks the physical power button on the server. For a server that already has the front panel locked, this option appears as **Unlock Front Panel**.



Note This option is available only for Intersight Managed Mode servers.

- **Reboot Cisco IMC**—Reboots the Cisco IMC.
- **Certificate:**
 - **Set KMIP Client Certificate**—To Configure a KMIP client certificate to ensure secure communication between the KMIP server and Cisco IMC. For more information, see [Configuring KMIP Client Certificate](#).
 - **IMC certificates**—To configure the certificate and private key on the server from a third-party managed Certificate Authority(CA). This option is available only for Intersight Managed Mode servers. For more information, see the *Creating a Certificate Management Policy* section in [Supported UCS Server Policies](#).
- **Reset Memory Errors**—Allows you to reset memory error counters for all Dual In-line Memory Modules (DIMMs) in a server to zero. After clearing errors, any DIMMs disabled due to prior uncorrectable errors will be re-enabled on the next boot, provided they encounter no further errors during Power-On Self-Test (POST).
- **Install Operating System**—Perform an unattended OS installation on one or more Cisco UCS C-Series Standalone servers and Cisco Intersight Managed Mode (IMM) servers (C-Series, B-Series, and X-Series) from your centralized data center through a simple process. For more information, see [Installing an Operating System](#).
- **Upgrade Firmware**—Perform a firmware upgrade. For more information, see the [Firmware Upgrade](#).
- **Start/Stop Alarm Suppression**—Select the servers and choose the *Start Alarm Suppression* option to silence the non-critical alarms, or select the *Stop Alarm Suppression* option to resume receiving all alarm notifications. For more information, see [Alarm Suppression](#).

- **Set License Tier**—Update one or more servers to a new license tier. Updating license tier is not allowed on server(s) with an associated server profile. To move the license to another tier, unassign the profile from one or more selected servers and try again. For more details, see [Multiple Licensing Tiers](#).

Servers Table View

From the **Service Selector** drop-down list, select **Infrastructure Service**. Navigate to **Operate > Servers**, to launch the Server Table view. From this page, you can launch device endpoints, perform bulk server actions, and navigate to the server details page. Click the **Settings** icon (the gear icon representation), and select the columns that you want in the Table view.

You can add specific columns or custom tags to the Servers Table view to sort and filter.

Each column in the Servers Table view except **Organization** can be sorted with the **Sort** option and you can also add filter based on any columns using **Add Filter** option to view and explore server inventory.

Servers Table Summary Dashboard

The following widgets are available in the Servers table view:



Note Except **Server Profiles** all other widgets are dynamic based on the **Add Filter** option that you choose.

- **Health**—The pie chart provides a visual representation of the health of the servers.
- **Power**—The badges display the number of servers powered off or on.
- **HCL Status**—The badges display the HCL status for the servers.
- **Bundle Version**—The pie chart displays the total number of servers distributed by the bundle version.
- **Firmware Version**—The pie chart displays the total number of servers distributed by the firmware version.
- **Models**—The pie chart displays the total number of servers distributed by server models.
- **Contract Status**—The badges display the status of the service contract of the managed UCS and HyperFlex servers distributed by the current validity of their associated contracts.
- **Profile Status**—The pie chart displays the total number of servers distributed by the status of the server profile deployment.
- **Requests(last 24h)**—The pie chart displays the number of completed and failed tasks for the last 24 hours.
- **Alarm Suppression**—The badges display the number of servers categorized by their alarm suppression status: active (Yes) or inactive (No).

You can view the following details in the Servers Table view:

- **Name**—Displays the name of the server.



Important

- For standalone server, the name is a combination of server model and server serial number.
 - For UCSM Managed Mode (UMM) and Intersight Managed Mode (IMM) B-Series servers, the name is a combination of UCS domain name, chassis ID, and server ID. The server ID is auto-assigned depending on the order in which the servers are discovered.
 - For C-Series servers, the name is a combination of the UCS domain name, and server ID. To configure the server ID, decommission the server and recommission it. During recommission, you can assign the server an ID of your choice.
 - The power icon displays the server power status **ON/OFF**.
 - The connection icon displays the server connection status.
-

- **Health**—Displays the server health state, which corresponds to the server's alarm indicators. The server health status can be Healthy, Warning, or Critical. You can hover over the health status to view the top three active alarms. An icon next to the Health status indicates that the server's alarm notifications are currently suppressed. For more information, see [Alarm Suppression](#).
- **Contract Status**—Displays the status of service contract for the managed UCS and HyperFlex servers based on the current validity of their associated contracts. You can identify the SmartNet Contract ID details of the server, and cross launch the [Cisco Commerce Software Subscriptions and Service Portal](#).
- **Alarm Suppression**—Displays the alarm suppression status on the server as *Yes* for active or *No* for inactive. For more information, see [Alarm Suppression](#).
- **Management IP**—Each server in Cisco UCS instance must have a management IP address assigned to its Cisco Integrated Management Controller (CIMC) or to the profile associated with the server. Cisco UCS Manager uses this IP address for external access that terminates in the CIMC.
- **Model**—Displays the server model.
- **CPU Capacity (GHz)**—The aggregated speed of the CPUs on this server. CPU Capacity is calculated as the Number of CPU Sockets x Enabled Cores x Speed.
- **Memory Capacity (GB)**—The amount of RAM installed on the server in Gigabytes.
- **UCS Domain**—Displays the name of the UCS Domain the server belongs to. For standalone server, this column is not applicable.
- **HX Cluster**—Displays the name of the HyperFlex cluster the server belongs to.
- **HCL Status**—Displays the compliance status with the Hardware Compatibility List (HCL) after checking the compatibility of the server model, processor, firmware, adapters, operating system and drivers. For more information, see [Compliance with Hardware Compatibility List \(HCL\)](#).
- **Management Mode**—Displays the management mode of the server (Standalone, Intersight, UCSM)
- **Server Profile**—Displays the server profile that is associated with the server.

- **Utility Storage**—Displays the storage utility that is associated with the server and whether it is in the OK state.
- **Bundle Version**—Displays the firmware bundle version to which the server was upgraded.
- **Firmware Version**—Displays the running server firmware version at the endpoint.
- **Serial**—Displays the host ID/serial number of the server.
- **User Label**—Displays the assigned user label that helps in identification of the server.
- **License Tier**—Displays the current license on the server. You can update one or more servers to a new license tier. From the ellipsis (...) on the far right column for a server, you can choose a new license tier from the drop-down. For updating multiple servers at once, select the desired servers, click the ellipsis (...) at the top left of the table and select **Set License Tier**. For more information, see [Multiple Licensing Tiers](#).



Note **Set License Tier** action is not permitted if there is a server profile assigned to the server you wish to update. Once you have unassigned the server profile, then the **Set License Tier** action is available on that server to move to a new license tier.

- **Asset Tag**—A tag that identifies the server. The tag must have the serial number and other identifiers as required. The serial number is required to track the server in case of a service or replacement request.
- **CPU**—Displays the number of CPUs in the server.
- **CPU Cores**—Displays the number of CPU cores in the server.

Properties

The **Properties** area displays a graphical view of the Server. The **Health Overlay** function enables you to monitor the health of the ports on the Server.

- **CPUs**—The number of CPUs in the server.
- **CPU Cores**—The number of CPU cores in the server.
- **Memory Speed (MHz)**—The speed of memory in the server in MHz.
- **Organizations**—Lists the organizations to which the server is assigned.
- **Lightning Icon**—Click the ellipsis (...) icon for operations that include Firmware Upgrade, OS Install, launch Tunneled vKVM, and Set License Tier. You can also perform Bulk Server Actions including power cycle, hard reset, and reboot IMC. For more information, see the Bulk Server Actions section in [Server Discovery and Actions](#).

Alarms

Intersight provides fault monitoring capabilities to track alarms for all managed UCS and HyperFlex systems. For more information, see [Alarms](#).

Server Details View

- **General**—The server dashboard provides a centralized overview where users can assess the server's health, configuration, and properties. It allows for easy monitoring of the server's status and components. The dashboard also displays events associated with the server, including configuration changes, hardware events, and system messages. For detailed information on health status, properties, and a list of current alarms, see [Alarms](#).
- **Inventory**—Provides details such as server summary, server properties, and an inventory of subsystems on your server such as CPU, memory, power supplies, fans, IO devices, storage, BIOS, and Cisco IMC. Inventory is updated through events as and when they are received from endpoints. In addition, inventory is updated on a daily basis for claimed devices and on a weekly basis for unclaimed devices. For more information, see [Servers Inventory View](#).
- **Server Profile**—Displays the associated Server Profile configuration status, server health, last updated time, and server availability. This tab appears only when you activate the Intersight Essentials license. For more information, see [Server Profiles](#).
- **HCL**—Displays the hardware compliance status of your Cisco UCS and HyperFlex systems. This tab appears only when you activate the Intersight Essentials license. You can view recommendations for driver versions based on the server model, adapters, and the server firmware version for a selected operating system. For more information, see [Servers HCL View](#).
- **Statistics**—Displays the power state and power consumption telemetry of servers. This tab appears only when you activate the Intersight Essentials license. You can select the time period for which you need to view the telemetry data. The selected time period can be up to six months.



Note The feature is supported only in the cloud environment.

- **Power State**—The telemetry data is updated every five minutes. A green block indicates that the server is powered ON, a red block indicates that the server is powered OFF, and a blank block indicates that the server is unclaimed.
- **Power Consumption**—The power consumption in Watts for X-Series servers can be viewed numerically and graphically.
 - **Current Consumption** of power is updated every five minutes. An info button is present beside the current consumption value that shows the time stamp for the latest available data.
 - **Max Consumption** indicates the maximum power that has been consumed in the selected time period.
 - **Min Consumption** indicates the minimum power consumed in the selected time period.
 - **Average Consumption** indicates the average power consumed over the selected time period.
- **Topology**—Displays a detailed view of the connections from a single C-Series server, including visibility to internal ports. For more information, see [Viewing Server Topology](#).
- **Connectivity**—Displays detailed VIF path information of servers. It helps in mapping and troubleshooting the vNIC and vHBA connections. For more information, see [Server Connectivity View](#).

In addition to the **Actions** listed in the **Server Actions** and **Bulk Server Actions** section, you can perform the following operations from the **Server** details page:

- **Launch Tunneled vKVM**
- **Launch KVM**
- **Launch Cisco IMC**
- **Add/ Edit Asset Tags**
- **Add/Edit User Label**
- **Start/Stop Alarm Suppression**
- **Open TAC Case**
- **Download System Event Log**
- **Clear System Event Log**



Note Standalone HyperFlex M4 and M5 Servers support only Launch KVM.

Server Inventory View

After a server is discovered, an inventory of all its components is made available. When you select a server in the **Server** table view, you can view the inventory of its components on the **Inventory** tab.

For the selected server, you can view details of each of the following components:

- **Boot**—You can see the actual boot order of the devices configured on the server. The boot order displays the details that include device name, device type, configuration details such as Boot Mode (Legacy or UEFI), and Secure Boot Mode (Enabled or Disabled). A device configured in the server profile of Boot Order Policy may not appear in the actual boot order, if the server BIOS does not detect the device during server boot.
- **Management Controller**—You can view the firmware version, a summary of the out-band management access, hardware details, and server certificate details. Also, you can view or copy the latest server certificate from the Certificate section.



Note The Server Certificate operations for UCS B-series (M5, M6) and X-Series (M6, M7) servers in Intersight Managed Mode (IMM) are supported only on Server Firmware 4.2 and later versions. However, there are no limitations with Server Firmware version for UCS C-Series (M5, M6, M7, M8) servers.

- **CPUs**—You can see details about the processors, including the architecture, model, socket designation, and vendor. Expanding **CPUs** displays the state and a summary of the hardware and resource details of each processor.

- **Memory**—You can see a summary of the memory cards, including their location, ID, capacity, and Clock Speed. Expanding **Memory** displays the state and hardware details of each memory card.
- **Network Adapters**—You can see details about the network adapter cards, including the slot to which they are connected, model, serial, vendor, and the interfaces to which they are connected.

On selecting an adaptor, you can view the following details:

- **General**— You can view the firmware version, interface details (DCE/NIC/HBA), hardware details, and a list of alarms pertaining to each adapter.
- **Interfaces**—You can view the Name, MAC Address, VIF ID, Oper State, Passive VIF ID, Passive Oper State, and Rate Limit. Click the ellipses (...) to perform the following **Actions** for vNICs and vHBAs:
 - **vNICs**— Enable, Enable (Active), Enable (Passive), Disable, Disable (Active), Disable (Passive), Reset Connectivity, Reset Connectivity (Active), Reset Connectivity (Passive).
 - **vHBAs**— Enable, Disable, Reset Connectivity.



Note Enable or Disable vNICs and vHBAs is not applicable for Standalone mode.

- **GPUs**—You can see a list of GPUs. Expanding GPUs displays the GPU Inventory details that include General and GPU Controllers information for each GPU.
 - **General**
 - **Main**—You can view the Slot ID, Model, Serial Number, Vendor, Number of GPUs, and Firmware version.
 - **PCIe Enclosure**—You can view the Slot ID, Model, Serial Number, and Vendor information.
 - **GPU Controllers**—You can view the GPU Controller Name and PCI Address information.

Any change operation including insert, remove, or replace operation to the GPU requires you to trigger the rediscovery. Thus the rediscovery enables to discover the changes and update the server inventory.

- **PCIe Devices**—You can see a list of PCIe devices. Expanding PCIe Devices displays the configuration and hardware information of each device.
 - **Configuration**—You can view the firmware version of the device.
 - **Hardware**—You can view the Slot ID, Product Name, Serial Number, and Vendor information of the device.
- **Storage Controllers**—You can see a list of storage controller associated with the server. The Storage Controllers table view displays the storage controller Name, ID, Type, Firmware Version, Serial Number, Hybrid Slots Supported.

Expanding **Storage Controllers** displays the controller configuration details and enables you to perform the following operations:

You can perform the following operations to manage one or more storage controllers.

- **General**

- **Configuration**—You can view the firmware version of the storage controller.



Note Cisco UCS C-series (M7, M8) in Intersight Managed Mode and Standalone mode supports UCS-M2-NVRAID storage controller in M.2 RAID0 and RAID1 Controllers.

Cisco UCS C-series (M7) in UCSM Managed Mode supports Cisco UCS C220 M7 storage controller (Up to 10 SAS/SATA) and Cisco UCS C240 M7 (Up to 28 SAS/SATA) or NVMe disk drives.

Cisco Tri-Mode 24G SAS RAID Controller w/4GB Cache (UCSC-RAID-HP) to handle U.3 drives in RAID or Direct attached modes.

- **Hardware**—You can view the Controller ID, ID Type, RAID Support, Number of Disks, Serial Number, Model, and Vendor information of the storage controller.
- **Physical Drives**—Enables a single physical drive or multiple physical drives to switch between **Unconfigured Good** and **JBOD** drive states.
Physical Drives table view displays PID, Vendor, Protocol, Security Flags, Type, Drive State.
The following two models are supported: UCS-NVM2-400GB, UCS-NVM2-960GB
- **Virtual Drives**—Enables you to select and remove the unused virtual drive to reclaim the used space in the RAID controller. Removing the virtual drive destroys all information on the file systems and deletes the virtual drive from the RAID controller.



Note This is the only storage operation supported in Cisco Boot Optimized M.2 RAID Controller.

• Storage Controller and Physical Drive Operations

The following table describes the supported SED drive operations.

Storage Controller and Physical Drive Operations	Description
Secure Erase	Use this option to delete the Key Encryption Key and erase the stored data in an SED. The Actions menu next to a physical drive displays this option.
Import Foreign Configuration	Use this option to clear the user configuration on the physical drive and delete the Virtual Drives. The Actions menu next to a Controller displays this option.

Storage Controller and Physical Drive Operations	Description
Clear Foreign Configuration	Use this option to clear or erase all the data stored on the physical drives or the virtual drives. The Actions menu next to a Controller displays this option.
Clear Configuration	Use this option to delete the Virtual Drives, or to clear any user configurations on the storage controller and reuse the controller, when a server is not associated with a server profile. The Actions menu next to a Controller displays this option.
Disable Security	Use this action to disable the security on a controller. The Actions menu next to a Controller displays this option.
Modify Security	Use this option to modify the Key Encryption Key after security has already been enabled on a controller. The Actions menu next to a Controller displays this option.
Unlock Disks	Use this option to unlock a drive to access its data when an encrypted drive is inserted from another server. The Actions menu next to a Controller displays this option.

- **Hybrid Storage Slots**—Hybrid Slots indicate whether the RAID controller can handle U.3 drives in SAS/SATA mode or not. You can view the Slot ID, Requested Mode and Current Mode. The applicable values are RAID and Direct.
- **TPM**—Trusted Platform Module (TPM) enables protection to data and hardware components of the claimed server. TPM also enables you to view the state of the key identifiers and a summary of hardware details.

TPM configuration can also be cleared or reset by using the **Clear TPM** option from the **Actions** button in the right corner of the server inventory view.

Caution:

Clear TPM is meant for disaster recovery and data loss operation. Do not use it unless it is necessary.

Before using the **Clear TPM** action, you must ensure the following:

- Server Profile is configured.
- Operating System is installed.
- Server is in the power off state.



Note TPM Clear action is supported only for Cisco UCS B-Series and C-Series M5 and above servers and the Firmware Version 4.2(2a) and above.

You can view the following components of TPM:

- Key identifiers
 - **Activation Status**—Shows the TPM is in Activated/Deactivated state. When the TPM configuration is cleared/reset, the activation status shows Deactivated.
 - **Enabled State**—Shows the TPM is in Enabled/Disabled state. When the TPM configuration is cleared/reset, the Enabled state shows Disabled.
- Hardware
 - **Ownership**—Shows the ownership status as Owned/Unowned. When the TPM configuration is cleared/reset, the ownership status shows Unowned. To regain the ownership at anytime, you must switch on the Power Cycle Server.



Note These properties can be viewed only for TPM 1.2 version. For 2.0, the Activation Status, Enabled State, and the Ownership Status can be viewed in the operating system.

- Version
- Model
- Vendor
- Serial
- Firmware Version



Note This property can be viewed only for TPM 2.0 version.

Server Topology View

For more information, see [Viewing Server Topology](#).

Server Metrics View

For more information, see [Server Metrics](#).

Server Connectivity View

The Connectivity view provides detailed information of a server's Virtual Interface (VIF) paths, showing the links between its virtual Network Interface Cards (vNICs) and Host Bus Adapters (vHBAs) to the network's Fabric Interconnect (FI) or Fabric Extender (FEX). This view facilitates the mapping of these connections and helps in troubleshooting. You can sort and filter the displayed data across all columns for more efficient analysis. The system automatically updates the VIF path information at five-minute intervals to keep the data up to date.



Note You can export the displayed information in the Table view to a CSV file. To export and save it in a CSV file format, click the **Export** icon.

You can view the following details in the Connectivity Table view:

- **vNIC/vHBA Name**—Displays the name of the vNIC or vHBA.
- **Network Adapter**—Displays the connected network adapter card.
- **Fabric Side**—Displays the FI to which vNIC or vHBA is connected.
- **Link State**—Displays the current state of the Virtual interface.
- **State Reason**—Provides an explanation for a link being down, if applicable.
- **Virtual Circuit ID**—Displays the Virtual Circuit ID of the vNIC or vHBA.
- **Pin Group**—Displays the configured Pin Group information for the statically pinned vNIC.
- **FI Uplink Interface**—Displays the pinned FI uplink interface for the vNIC.
- **FEX Host Interface**—Displays the configured Host Interface information.
- **FI Server Interface**—Displays the configured Server Interface information.

Compliance with Hardware Compatibility List (HCL)

Cisco Intersight provides the capability to evaluate and mitigate the impact of service issues from running non-validated combinations of firmware, server model, processor, adapters, operating system, and driver versions. Intersight evaluates the compatibility of your Cisco UCS systems, HyperFlex systems, Intersight Managed Mode (IMM) servers, and Cisco UCS S-Series servers to check if the hardware and software have been tested and validated by Cisco or Cisco partners. Intersight reports validation issues after checking the compatibility of the server model, processor, firmware, adapters, operating system and drivers, and displays the compliance status with the Hardware Compatibility List (HCL). **This feature requires a Cisco Intersight Essentials or above license.**

You can use Cisco UCS Tools, a host utility vSphere Installation Bundle (VIB), or OS Discovery Tool, an open source script to collect OS and driver information to evaluate HCL compliance. For more information about Hardware Compatibility Status, a detailed description and instructions on how to download Cisco UCS Tools, and for instructions on how to use the OS Discovery Tool, see [Compliance with Hardware Compatibility List \(HCL\)](#) in Resources.



CHAPTER 5

Configuring UCS Domain Profiles

- [About UCS Domain Profile, on page 75](#)
- [Creating Domain Profile Template and Deriving Profiles, on page 75](#)
- [Creating a UCS Domain Profile, on page 77](#)
- [UCS Domain Profile Details, on page 77](#)

About UCS Domain Profile

Overview of a UCS Domain Profile

A UCS Domain Profile configures a Fabric Interconnect pair through reusable policies, allows for configuration of the ports and port channels, and configures the VLANs and VSANs in the network. It defines the characteristics of and configures ports on Fabric Interconnects. You can create a UCS Domain profile and associate it with a Fabric Interconnect Domain. The Domain-related policies can be attached to the profile either at the time of creation or later. One UCS Domain profile can be assigned to one Fabric Interconnect Domain.



Important

- Cisco Intersight supports attaching one port policy per UCS Domain profile.
- Policies that are attached to a UCS Domain profile can be created ahead of creating a profile or during the creation of the profile.
- Policies that are attached to a UCS Domain and the global policies of all UCS Domain Profiles associated with a specific UCS Domain are shared.

Creating Domain Profile Template and Deriving Profiles

Domain profile templates enable you to define a template from which domain profiles can be easily derived and deployed at scale. This method is particularly useful when you have a large number of domain profiles that use the same set of policies. Any property modification made in the template is inherited by the derived profiles. You can deploy these modified profiles individually. For more information, see [Domain Profile Templates](#).

To create a domain profile template and derive profiles:

1. Log in to Cisco Intersight.
2. From the **Service Selector** drop-down list, select **Infrastructure Service**.
3. Navigate to **Configure > Templates**, to launch the **Profile Template Table View**.
4. To create a domain template, click **Create UCS Domain Profile Template**.
5. On the **General** page:
 - a. Choose an Organization for the template from the list.

This field supports the capability of configuration sharing across Organizations.
 - b. Enter a name for the template. Enter a tag for the template.
 - c. Tags must be in the key:value format. For example, Org: IT or Site: APJ.
 - d. Enter a description to help identify the template.
 - e. Click **Next**.
6. On the **VLAN & VSAN Configuration** page, select a VLAN and VSAN configuration for Fabric Interconnect A and Fabric Interconnect B, and then click **Next**.
7. On the **Ports Configuration** page, create or select a port policy for Fabric Interconnect A and Fabric Interconnect B, and then click **Next**.
8. On the **UCS Domain Configuration** page, create or select the management and network policies, and then click **Next**.
9. On the **Summary** page:
 - a. Review the details of the template and the policies attached to it.
 - b. If you opt to derive the profiles later, click **Close**. Otherwise, to derive the profiles from this template, click **Derive Profiles**.
10. To derive a template:
 - a. In the **General** page:
 1. Under **Domain Assignment**, assign the template to multiple domains.

You can choose any of the following options for the domain assignment:

 - **Assign now**—Use this option to immediately assign a domain to the domain profile.
 - **Assign later**—Use this option to assign a domain to the domain profile at a later time.
 2. In the **Number of Profiles to derive** field, enter the number of profiles you want to derive.
 - b. On the **Details** page, you can view and edit the description, tag, and auto-generated names of the profiles, and then click **Next**.
 - c. On the **Summary** page, review the domain profile template configuration, and then click **Derive Profiles**.

The table view displays the list of domains templates and their details.

Creating a UCS Domain Profile

A UCS Domain Profile configures a Fabric Interconnect pair through reusable policies, streamlines the deployment of Fabric Interconnect pairs, allows for configuration of the ports and port channels, and configures the VLANs and VSANs in the network.

-
- Step 1** Log in to Cisco Intersight with your Cisco ID and select admin role.
- Step 2** Navigate to **Service Profiles > UCS Domain Profiles** tab, and click **Create UCS Domain Profile**.
- Step 3** On the **General** page, enter a name for your profile. Optionally, include a short description and tag information to help identify the profile. Tags must be in the key:value format. For example, Org: IT or Site: APJ.
- Step 4** On the **Domain Assignment** page, assign a switch pair to the Domain profile. You can also click **Assign Later** and assign a switch pair to the Domain profile at a later time.
- Step 5** Click **Next**.
- Step 6** On the **VLAN & VSAN Configuration** page, attach VLAN and VSAN policies for each switch to the **UCS Domain Profile** and click **Next**.
- Note** To configure System Reserved VLANs, the VLAN and VSAN policies must not conflict with the reserved VLAN range. When there is a conflict, the deployment fails.
- Step 7** On the **Ports Configuration** page, attach a port policy for each switch to the **UCS Domain Profile** and click **Next**.
- Step 8** On the **UCS Domain Configuration** page, attach the required compute and management policies to the **UCS Domain Profile** and click **Next**.
- Note:** In this step, it is mandatory to create and attach the Switch Control Policy to enable VLAN port count optimization.
- Step 9** Click **Next**.
- Step 10** On the **Summary** page, verify the details of the UCS Domain Profile and the policies attached to it.
- Step 11** Click **Deploy** to deploy the UCS Domain Profile to the assigned Fabric Interconnect Domain.
-

UCS Domain Profile Details

The UCS Domain Profile Details page displays a graphic representation of the Port Configuration, VLAN and VSAN Configuration, and the UCS Domain Configuration in addition to the status and the **Actions** menu. Navigate to the UCS Domain Details from the UCS Domain Profiles Table view. On this page, you can:

- Perform UCS Domain profile **Actions**:
 - **Deploy**—Deploy the UCS Domain profile on a Fabric Interconnect pair.
 - **Unassign**—Unassign the UCS Domain profile from the Fabric Interconnect pair.
 - **Edit**—Edit the properties of the UCS Domain Profile.
 - **Clone**—Clone the UCS Domain profile with properties similar to an existing UCS Domain profile. The clones are associated with the same policies as on the original UCS Domain profile.
 - **Set Tags**

- View UCS Domain profile **Details**:

Property	Essential Information
Status	The status of deploying the UCS Domain profile on a Fabric Interconnect pair. This could be: <ul style="list-style-type: none"> • OK • Failed • Not Deployed
Name	The UCS Domain profile name.
Fabric Interconnect A	The name of the associated Fabric Interconnect A in the UCS Domain.
Fabric Interconnect B	The name of the associated Fabric Interconnect B in the UCS Domain.
Last Update	The date and time that the UCS Domain Profile was last updated.
Tags	The existing tags for the selected object are displayed by default. Click Manage to add new tags or modify the existing ones.

- View the **Policies** that are attached to the UCS Domain profile. The **Policies** pane displays details of the Ports, VLAN and VSAN, and UCS Domain Configuration. A graphical representation of the ports configuration on the Fabric Interconnects, including port roles and port channels and a list of associated policies is displayed. The VLAN, VSAN, and UCS Domain Configuration lists the Domain policies associated with the selected Domain profile.



CHAPTER 6

Configuring Server Profiles

- [Server Profiles, on page 79](#)
- [Creating a UCS Server Profile, on page 88](#)
- [UCS Server Profile Details, on page 89](#)

Server Profiles

In Cisco Intersight, a Server Profile enables resource management by streamlining policy alignment, and server configuration. To view the Server Profiles table view, from the **Service Selector** drop-down list, choose **Infrastructure Service**. Navigate to **Configure > Profiles**. You can create Server Profiles using the Server Profile wizard or you can import the configuration details of C-series servers in standalone mode and FI-attached servers in Intersight Managed Mode (IMM), directly from Cisco IMC. You can create Server Profiles using the Server Profile wizard to provision servers, create policies to ensure smooth deployment of servers, and eliminate failures that are caused by inconsistent configuration. The Server Profiles wizard groups the server policies into the following four categories to provide a quick Summary View of the policies that are attached to a profile:

- **Compute Policies**—BIOS, Boot Order, and Virtual Media.
- **Network Policies**—Adapter Configuration, iSCSI Boot, LAN Connectivity, and SAN Connectivity policies.
 - The LAN Connectivity policy allows you to create Ethernet Network Policy, Ethernet Network Control Policy, Ethernet Network Group Policy, Ethernet Adapter Policy, or Ethernet QoS Policy. When you attach a LAN Connectivity policy to a server profile, the addresses of the MAC address Pool, or the static MAC address, are automatically assigned.



Note A LAN Connectivity policy that has a static MAC address can be attached to only one server profile.

- The SAN Connectivity policy requires you to create Fibre Channel Network Policy, Fibre Channel Adapter Policy, or Fiber Channel QoS Policy. When you attach a SAN Connectivity policy to a server profile, the addresses of the WWPN and WWNN Pools, or the static WWPN and WWNN addresses, are automatically assigned.



Note A SAN Connectivity policy that has a static WWPN, or a static WWNN can be attached to only one server profile.

- **Storage Policies**—SD Card and Storage policies
- **Management Policies**—Device Connector, IPMI Over LAN, LDAP, Local User, Network Connectivity, SMTP, SNMP, SSH, Serial over LAN, Syslog, NTP, Certificate Management, and Virtual KVM policies

For more information and descriptions of the policies, see the **Server Policies** section. For an example of the policy creation workflow, see [Creating Network Policies](#).

Server Profile List View

When you select **Profiles > UCS Server Profiles** in the Intersight UI, the UCS server profile list view is seen.

The list view shows the following details in a tabular format:

- **Name** – The name of the server profile.
- **Status** – The deployment status of the server profile.

The **Status** of the profiles can have any of the following values:

- **Not Assigned**—Policies are not assigned to the server profile.



Note

- Once you deploy policies to the server profile, the status changes automatically from *Not Assigned* to the new status depending on the outcome. You may need to refresh your screen to view the updated status.
- You must do the Power Cycle/Power ON after each profile deployment.

- **OK**—Policies deployed successfully on the server profile
- **In Progress**—Deployment of policies to the server profile is in progress
- **Failed**—Server profile validation, configuration, or deployment has failed.
- **Inconsistent**—Indicates that the policy configuration has changes that have not yet been deployed or activated. It may also indicate that the policy configuration at the endpoint is not in sync with the last deployed policy configuration in the server profile. If the endpoint settings are altered manually after a server profile is deployed, Intersight automatically detects the configuration changes and they will be shown on the server profile as **Inconsistent**. For more information, see the *Server Profile Drift* and the *Deploying and Activating a Server Profile* sections.
- **Inconsistency Reason** – The reason for the status being shown as *Inconsistent*. Example - Not Deployed, Not Activated, Out of Sync
- **Target Platform** – Indicates if the platform for which the profile is applicable is a Standalone UCS server or FI-attached UCS server.

- UCS Server Template – The template attached to the server profile or from which the profile has been derived.
- Server – The name of the server to which the profile is attached.
- Resource Pool – The pool to which the profile belongs.
- User Label – A user label is an identifier that helps in filtering the server profiles. It must be between 1 and 64 alphanumeric characters, containing only the following special characters: ! # \$ % & * + , () [] { } | / . ? @ _ : ; ~
- Last Update – The date on which the profile was last updated.
- Organization – The name of the organization.



Note Some of the columns are disabled by default, such as, **User Label**. To view such columns in the server profiles table view, you need to enable them while customizing the table view.

Server Profile Actions

After creating server profiles, actions that can be performed on a server profile are as follows:

- Deploy – Deploy the profile to the attached server
- Activate – Activate the profile on the attached server. The server gets power cycled on activation.
- Edit – Edit the profile
- Clone – Clone the profile
- Attach to Template – Attach the server profile to any of the available templates.



-
- Note**
- While template creation, if you toggle ON the **Attach UCS Server Profile to Profile Template** button, the selected profile gets attached to the template under creation.
 - If you keep the toggle button OFF, the selected profile's properties are carried to the template but the profile does not get attached to it.
-

- Create a Template – A server profile can be used to create a template. This template can then be used to create multiple profiles with same configurations and deployed on multiple servers.
- Detach from Template – Detach the profile from the template.

**Note**

- **Create a Template** and **Attach to Template** actions can be performed only if a server profile is not attached to any template.
- A server profile can be attached to an existing template. This attachment overrides the config properties of the profile and replaces them with the template properties.
- A server profile attached to a template cannot be modified. The modifications can be done in the associated template.
- A server profile can be detached from a template and modified as per the requirements.
- A detached server profile can always be reattached to a template.

- Unassign Server – Unassign the server from the profile.
- Set User Label – You can also set, update, or delete user labels for each server profile through the **Set User Label** action.

Server Profile Details View

Clicking on a profile redirects to the **Server Profile Details View** that displays the configuration details of the profile under **General**, **Server**, **Inventory**, and **Connectivity** tabs. For more information, see [Servers Details View](#).

Server Profile Drift

A server profile drift occurs when the policy configuration at the endpoint is not in sync with the last deployed policy configuration in the Server Profile.

Cisco Intersight supports Server Profile Drift detection for standalone servers and Intersight Managed Mode servers. For Intersight Managed Mode servers, the firmware versions required for drift detection are:

- For 4.2 release, the Cisco IMC version must be 4.2(1b) or above.
- For 4.1 release, the Cisco IMC versions must be:
 - For rack servers - 4.1(3d) or above
 - For blade servers - 4.1(33e) or above

The check to look up for any configuration change at the endpoint is performed every 30 min.

To see the policy configurations that have changed at the endpoint relative to the currently deployed policy configuration in Intersight, navigate to server profile details view and click **View Changes**. You can choose to view the **Changes Only** or **All** the policy configuration details.

Property	Essential Information
Saved Settings	Displays the policy settings in Intersight.

Property	Essential Information
Last Deployed Settings	Displays the latest policy settings deployed on the server profile.
Endpoint Settings	Displays the configuration at the endpoint.

To move the Server Profile status back to **OK**, you can either redeploy the profile or change the values at the endpoint. You can use the Device Connector Policy in Intersight to control configuration changes allowed from Cisco IMC. In the Device Connector Policy, choose **Configuration from Intersight only** to stop allowing configuration changes from Cisco IMC directly.

Limitations of Server Profile Drift - Standalone Servers

For standalone servers, configuration changes at the endpoint will not be detected for the following policies under the specified conditions:

Policy	Configuration at the endpoint
SD Card Policy	If an SD card is removed.
Storage Policy	<ul style="list-style-type: none"> • If Expand to Available is set for any of the virtual drives in the policy. • If the Power Cycle is not done after every deployment. • If there are additional drive groups that are not configured from Intersight
Boot Order Policy	<p>If the Power Cycle is not done after every deployment.</p> <p>In SAN boot devices, Intersight does not detect drift for Interface Name and Target WWPN</p> <p>Note Cisco recommends using a SAN boot, because it offers the server profile mobility within the system. If you boot from the SAN when you move a server 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.</p> <p>To use a SAN boot, ensure that the following is configured:</p> <ul style="list-style-type: none"> • 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.

Policy	Configuration at the endpoint
Local User, SNMP, LDAP, and IPMI over LAN Policy	If there are changes to the Password at the endpoint.
Virtual Media policy	If there are changes to the Password, Mount Options, or Authentication Protocols at the endpoint.
BIOS Policy	<ul style="list-style-type: none"> • BIOS token values configured as 'platform-default' are changed to the default value for that platform. Drift detection does not occur for such BIOS tokens. For more details, see Table 16 of the Creating a BIOS Policy section in Supported UCS Server Policies. • BIOS tokens whose values depend on other BIOS token values are not considered for drift detection. Drift may get reported for a BIOS token whose value is not supported by the server on which the policy is being deployed. For more details, see Cisco UCS Server BIOS Tokens.
IPMI over LAN policy	'Privilege Level' field will not be considered.
Network Connectivity Policy	'Preferred IPv6 DNS Server' and 'Alternate IPv6 DNS Server' fields in the policy will not be considered. Server Profile may move to Out of Sync status temporarily.
Adapter Configuration Policy	This policy will not be considered for drift calculation.
Ethernet Adapter Policy	<p>If a usNIC or VMMQ has a different Ethernet Adapter policy, then the configuration changes will not be calculated for usNIC or VMMQ attached Ethernet Adapter policy.</p> <p>Due to VMQ configuration restrictions, VMQ Number of Interrupts will override the value of Interrupts in Ethernet Adapter Policy, and VMQ Number of Virtual Machine Queues will override the value of Receive Queue Count, Transmit Queue Count, and Completion Queue Count (Receive+Transmit) of Ethernet Adapter Policy. Drift will not be detected for Number of Interrupts, Number of Virtual Machine Queues, Receive Queue Count, Transmit Queue Count, and Completion Queue Count.</p> <p>Intersight does not detect drift for 'Number of Interrupts', 'Number of Virtual Machine Queues', 'Receive Queue Count', 'Transmit Queue Count', and 'Completion Queue Count'.</p>
LAN Connectivity Policy	The CDN and PCI Order fields will not be considered.

Policy	Configuration at the endpoint
IMC Access Policy	If both In-Band IPv6 and IPv4 configurations are available, the IPv6 DNS configuration is prioritized.

Limitations of Server Profile Drift - Intersight Managed Mode Servers

For Intersight Managed Mode servers, server configuration changes at the endpoint will not be detected for the following policies under the specified conditions:



Note The Name field is not supported for any policy because Name is not an endpoint setting.



Note Drift detection is not supported for pools and IDs.

Policy	Configuration at the endpoint
SD Card Policy	Drift detection is not supported if an SD card is removed.
Storage Policy, Boot Order Policy, BIOS Policy, Virtual Media Policy	Drift detection is not supported for Storage policy, Boot Order Policy, BIOS Policy, and Virtual Media Policy on Intersight Managed Mode servers.
Local User Policy, SNMP Policy, Certificate Management Policy	Drift detection is not supported if there are changes to secure fields such as Password, Community Strings, and Private Key at the endpoint.

Policy	Configuration at the endpoint
LAN Connectivity Policy	Drift detection is not supported for: <ul style="list-style-type: none"> • VMQ connection <ul style="list-style-type: none"> • Number of interrupts • Number of Virtual Machine Queues • Consistent Device Naming (CDN) • Auto vNICs Placement IDs • PCI Order • Ethernet Adapter Policy <ul style="list-style-type: none"> • Interrupts Settings - Interrupts • Completion - Completion Queue Count, Completion Ring Size • VMMQ Adapter Policy • usNIC Adapter Policy <p>Note Drift detection is supported only when the servers are powered on.</p>
IMC Access Policy	Drift detection is not supported for Out-of-Band configuration.
SAN Connectivity Policy	Drift detection is not supported for Auto vNICs Placement IDs. <p>Note Drift detection is supported only when the servers are powered on.</p>
Power Policy	Drift detection is not supported for the Power Priority property.

Server Profile Import

Intersight provides the capability to import configuration details of C-series servers in standalone mode and FI-attached servers in Intersight Managed Mode (IMM), directly from Cisco IMC. The Server Profile import enables you to migrate the configuration of your existing servers to Intersight without having to create a profile and the policies manually. The Server Profile import operation creates a profile and the associated policies based on the server configuration. You can create a *golden* configuration profile and clone it and apply to another server already claimed in Intersight.

You can import a server profile configuration from the following locations in Intersight:

- **Servers** table view—Select a Cisco UCS C-Series Standalone server or any FI-attached server in Intersight Managed Mode (IMM) from the table view and click the ellipses (...) and select **Import Server Profile**.

- Click a C-series server in standalone mode or any FI-attached server in Intersight Managed Mode (IMM) in the Servers table view to access the Server details page. Click **Actions** on the top-right corner and select **Import Server Profile**. This option is enabled only when no server profile is associated with the server.



Note A partially imported server profile cannot be attached to a template or cannot be used for creating a template.

For more information on how to import a Server Profile Import and about the detection of manual configuration changes at the endpoint, see [Importing a Server Profile](#) in [Resources](#).

Estimate Impact

The Estimate Impact workflow, for standalone and Intersight Managed Mode servers, analyzes the disruptions that would be caused by the various policies attached to a server profile, when the server profile is deployed. The analyze impact workflow is triggered when a policy is attached, detached, or updated. The Disruption is indicated against each policy. The disruptions, which could be caused by the policies, are:

- Immediate reboot is required for standalone server policies such as Persistent Memory policy or Adapter policy. In such cases, the disruption indicated against the policy is **Immediate Reboot**.
- An Activate action on the server profile needs the server to reboot and activate the policy configuration on the server. In such cases, the disruption indicated against the policy is **Activate Requires Reboot**.
- Some policies, such as IMC Access policy, cause a brief outage of the server management network. In such cases, the disruption indicated against the policy is **Network Management Outage**.

Deploying and Activating a Server Profile

Deploy and **Activate** are two explicit actions that can be performed on server profiles. Policy configuration staging happens as a part of server profile deployment. Policy staging allows you to stage the policy configurations and get an idea of the pending actions for activating the policies. You can activate the policy by rebooting servers manually or using the **Activate** action of the Server Profile during a maintenance window. Policy activation failures are identified when the **Activate** action is triggered.

The **Status** widget in the Server Profiles table view shows the number of profiles in **Inconsistent** state. A server profile will be in the **Inconsistent** state when it has policy changes that have not yet been deployed or activated. The **Inconsistency Reason** widget shows the reason why a profile is in the **Inconsistent** state. A server profile could be in an **Inconsistent** state because:

- There are changes in the policies attached to the server profile assigned to the server.
- The policy configuration is out-of-sync with the configuration deployed in the endpoints.
- The policy is in **Not Activated** state.

You can use **Deploy** action to stage the configuration changes. During Deploy, you can choose to enable a toggle button to **Reboot Immediately**. If enabled, the server reboots and the server profile is activated immediately. If disabled, the policy configuration changes are activated at the next reboot.

The **Activate** action in the Server Profile details, reboots the server and activates the configuration on the server. You can trigger **Deploy** to stage the configuration changes and later trigger **Activate**, during the maintenance window, to activate the deployed configuration.

The **Update and Deploy** option in the policy edit page allows you to modify a policy configuration and deploy the changes on multiple server profiles to which the policy is attached.

Creating a UCS Server Profile

A server profile defines a server and its compute, storage, management, and network characteristics. When a server profile is deployed to a server, Cisco Intersight automatically configures the server and its connections to match the configuration specified in the server profile.



Note A Server profile can also be derived from Server Profile Templates. For more details, see [Server Profile Templates](#)

Step 1 Log in to Cisco Intersight with your Cisco ID and select admin role.

Step 2 Navigate to **Service Profiles > UCS Server Profiles** tab, and click **Create UCS Server Profile**.

Step 3 On the **General** page, enter the following information:

- a) **Name** of your server profile.
- b) **Target Platform** for which the profile is applicable. This can be **Standalone** servers or **FI Attached** servers.

A UCS server profile created for Standalone servers cannot be deployed on FI Attached servers. Similarly, a UCS server profile created for FI Attached servers cannot be deployed on Standalone servers.

- c) (Optional) **Tag** for the profile. Tags must be in the key:value format. For example, Org: IT or Site: APJ.
- d) (Optional) **Description** to help identify the profile.

Step 4 On the **Server Assignment** page, assign a server to the server profile. You can choose any of the following options for the server assignment:

- **Assign from a Specific Server**—Use this option for an immediate assignment of a server to the server profile.
- **Assign Server from a Resource Pool**—Use this option to assign a server from a resource pool to the server profile.
- **Assign by Chassis Slot Location**—Use this option to pre-assign a server to the server profile using the Domain Name, Chassis ID, and Slot ID.
- **Assign by Serial Number**—Use this option to pre-assign a server to the server profile using the Serial Number of the server.

Note

- Cisco UCS B-Series servers can be pre-assigned using **Chassis Slot Location** or **Serial Number**.
- Cisco Intersight Managed Mode C-Series servers and Cisco UCS C-Series Standalone servers can be pre-assigned only using **Serial Number**.

- **Assign Later**—Use this option to assign a server to the server profile at a later time.

The server assignment table displays list of servers or resource pools and their details. You can use any of the following options to view the details:

- **Show All** to view all the servers or resource pool currently present

- **Show Selected** to view the current server or resource pool selected
- **Unselect** to remove the selection.

Step 5 Click **Next**.

Step 6 On the **Compute Configuration** page, do the following:

- a) Choose the appropriate **UUID Assignment**:
 - **Pool**—Allows UUID Pool association to the server.
 - **Static**—Allows UUID association to the server using Static UUID address.
- b) Select the existing policies or create new policies.
- c) Click **Next**.

Step 7 On the **Management** page, attach the required policies to the **UCS Server Profile** and click **Next**.

Step 8 On the **Storage** page, attach the required policies to the **UCS Server Profile** and click **Next**.

Step 9 On the **Network Configuration** page, attach the required policies to the **UCS Server Profile** and click **Next**.

Step 10 On the **Summary** page, verify the details of the UCS Server Profile and the policies attached to it.

Step 11 Click **Deploy** to create the UCS Server Profile and deploy it to the assigned server.

Note • For the **Assign Server from a Resource Pool** assignment type, if a resource is not available in the resource pool, the status of the Server Profile changes to **Waiting for Resources**. Similar behavior is observed for the pre-assignment of the Server Profile. When a server is added to the resource pool at a later time, the server gets automatically added to the server profile from the **Waiting for Resources** status.

An alarm gets raised when the Server Profile is in the **Waiting** state. It gets auto-cleared when a server gets assigned to the Server Profile.

- Resource pool does not support dynamic selection of server. You can manually assign servers to a resource pool and can continue with the automated server profile assignment.
- The Server Profile pre-assignment is a one-time operation till the server is assigned. The Pre-assigned properties are lost once the server is assigned and continues to function as any other existing Server Profiles.
- For more information on creating a resource pool and viewing the resource pool details, see [Resource Pools](#).
- For more information on creating a UUID pool and viewing the UUID pool details, see [UUID Pools](#).

UCS Server Profile Details

The UCS Server Profile Details page displays details of the UCS Server profile and the server that it is assigned to. Navigate to the UCS Server Details from the UCS Server Profiles Table view. On this page, you can:

- Perform UCS Server profile **Actions**:
 - **Deploy**—Deploy the UCS Server profile on a Fabric Interconnect pair.



Note This action can be performed on a server profile that has servers assigned to it.

- **Unassign**—Unassign the UCS Server profile from the Fabric Interconnect pair.



Note This action can be performed on a server profile that has servers assigned to it.

- **Edit**—Edit the properties of the UCS Server Profile.
- **Clone**—Clone the UCS Server profile with properties similar to an existing UCS Server profile. The clones are associated with the same policies as on the original UCS Server profile.
- **Delete**—Delete the server profile.
- **Attach to template**—Attach the server profile to an existing server profile template.



Note This action can be performed on a server profile that is not attached to any template.

- **Create a template**—Create a new template using the properties of the server profile.



Note This action can be performed on a server profile that is not attached to any template.

- **Detach from template**—Detach the server profile from a template and modify its properties.



Note This action can be performed on a server profile that is attached to a server profile template.

- **Manage Tags**— Set a tag for a profile in the key:value format.

- View UCS Server profile **Details** in the **General** tab:

Property	Essential Information
Status	<p>The status of deploying the UCS Server profile on a Fabric Interconnect pair. This could be:</p> <ul style="list-style-type: none"> • OK • Failed • Not Assigned • Not Deployed

Property	Essential Information
Name	The UCS Server profile name.
Server	The name of the associated server.
Last Update	The date and time that the UCS Server profile was last updated.
Tags	The existing tags for the selected object are displayed by default. Click Manage to add new tags or modify the existing ones.

Displays the policies associated with the server profile. Click on the policy name to view details of the associated policy.

If you make changes to a policy attached to a Server Profile after it is deployed, or add a new policy to the profile, the Server Profile Table view displays Not Deployed Changes to reflect the edits to the profile or the referenced policies. The Server Profile Detail view highlights the referenced policies, and the View Changes window allows you to view the actual changes. You can also view the Configuration details from the Service Profiles table view.

- View the assigned server and its properties in the **Server** tab.
- View the inventory of the assigned server in the **Inventory** tab.



CHAPTER 7

Configuring UCS Chassis Profiles

- [About UCS Chassis Profile, on page 93](#)
- [Creating Chassis Profile Template and Deriving Profiles, on page 94](#)
- [Creating a Chassis Profile, on page 95](#)
- [UCS Chassis Profile Details, on page 95](#)

About UCS Chassis Profile

Overview of a UCS Chassis Profile

A UCS Chassis profile enables to create and associate chassis policy to an Intersight Managed Mode (IMM) claimed chassis. When a chassis profile is associated to a chassis, Cisco Intersight automatically configures the chassis to match the configuration specified in the policies of the chassis profile. The chassis-related policies can be attached to the profile either at the time of creation or later.



Important

- The chassis profile feature is available in Cisco Intersight only if:
 - You have installed the Cisco Intersight Essentials License.
 - You are either an Account Administrator or Server Administrator.
- Policies that are attached to a chassis profile can be created ahead of creating a profile or during the creation of the profile.
- If chassis policies are changed after deployment, the chassis profile will be set in **Pending Changes** state and you must manually re-associate the changed policies to chassis.
- Chassis policies will be applied to both the input/output modules (IOMs) in a chassis. The chassis policies association workflow will get failed even if the policy cannot be applied to one of the IOMs.
- The IMC Access policy and SNMP policy are not supported for Cisco UCS X-Series Direct chassis (UCSX-9508) currently.

Creating Chassis Profile Template and Deriving Profiles

Chassis profile templates enable you to define a template from which chassis profiles can be easily derived and deployed at scale. This method is particularly useful when you need a large number of chassis profiles that use the same set of policies. Any property modification made in the template is inherited by the derived profiles. You can deploy these modified profiles individually. For more information, see [Chassis Profile Templates](#).

To create a chassis profile template and derive profiles:

1. Log in to Cisco Intersight.
2. From the **Service Selector** drop-down list, select **Infrastructure Service**.
3. Navigate to **Configure > Templates**, to launch the **Chassis Profile Template Table View**.
4. To create a chassis template, click **Create UCS Chassis Profile Template**.
5. On the General page:
 - a. Choose an Organization for the template from the list. This field supports the capability of configuration sharing across Organizations.
 - b. Enter a name for the template.
 - c. Enter a tag for the template. Tags must be in the key:value format. For example, Org: IT or Site: APJ.
 - d. Enter a description to help identify the template.
 - e. Click **Next**.
6. On the **Chassis Configuration** page, create or select an IMC Access policy, Power policy, SNMP policy, and Thermal policy, and then click **Next**.
7. On the **Summary** page:
 - a. Review the details of the template and the policies attached to it.
 - b. If you opt to derive the profiles later, click **Close**. To derive the profiles from this template, click **Derive Profiles**.
8. To derive a template:
 - a. In the **General** page:
 1. Under **Chassis Assignment**, assign the template to multiple chassis.
You can choose any of the following options for the chassis assignment:
 - **Assign now**—Use this option to immediately assign a chassis to the chassis profile.
 - **Assign later**—Use this option to assign a chassis to the chassis profile at a later time.
 2. In the **Number of Profiles to derive** field, enter the number of profiles you want to derive.
 - b. On the **Details** page, you can view and edit the description, tag, and auto-generated names of the profiles, and then click **Next**.

- c. On the **Summary** page, review the chassis profile template configuration, and then click **Derive Profiles**.

The table view displays the list of chassis templates and their details.

Creating a Chassis Profile

A Chassis Profile configures a chassis through reusable policies.

-
- Step 1** Log in to Cisco Intersight with your Cisco ID and select admin role.
 - Step 2** Navigate to **Profiles > Chassis Profiles** tab and click **Create UCS Chassis Profile**.
 - Step 3** On the **General** page, select the organization and enter a name for your profile. Optionally, include a short description and tag information to help identify the profile. Tags must be in the key:value format. For example, Org: IT or Site: APJ.
 - Step 4** On the **Chassis Assignment** page, assign a chassis to the **Chassis Profile**. You can also click **Assign Later** and assign a chassis to the chassis profile later.
 - Step 5** Click **Next**.
 - Step 6** On the **Chassis Configuration** page, attach the required policies and click **Next**.
 - Step 7** On the **Summary** page, verify the details of the **UCS Chassis Profile** and the policies attached to it.
 - Step 8** Click **Deploy** to deploy the **UCS Chassis Profile** to the assigned Fabric Interconnect.
-

UCS Chassis Profile Details

On the **UCS Chassis Profile Details** page, you can:

- Perform chassis profile **Actions**:
 - **Deploy**—Deploy the chassis profile on a Fabric Interconnect pair.
 - **Edit**—Edit the properties of the chassis profile.
 - **Unassign Chassis**—Unassign the chassis profile from the Fabric Interconnect pair.
- View the UCS Chassis profile **Details**:
 - **Status**—The status of deploying the Chassis profile on a Fabric Interconnect pair, such as:
 - **OK**
 - **Not Assigned**
 - **Not Deployed**
 - **Failed**
 - **Not Deployed Changes**
 - **Name**—The chassis profile name.

- **Chassis**—The chassis details.
 - **Last Update**—The date and time that the chassis profile was last updated.
 - **Description**—The description of the chassis profile.
 - **Organization**—The selected organization is displayed. Click **default** to set a default organization.
 - **Tags**—The existing tags for the selected object are displayed by default. Click **Set** to add new tags or modify the existing ones.
-
- View the **Policies** that are attached to the chassis profile.



CHAPTER 8

Configuring UCS Domain Policies

- [Domain Policies](#), on page 97
- [Creating a Port Policy](#), on page 100
- [Creating an Ethernet Network Group Policy](#), on page 108
- [Creating an Ethernet Network Control Policy](#), on page 110
- [Creating a VLAN Policy](#), on page 112
- [Creating a VSAN Policy](#), on page 114
- [Creating an NTP Policy](#), on page 115
- [Creating a Network Connectivity Policy](#), on page 116
- [Creating an SNMP Policy](#), on page 118
- [Creating a System QoS Policy](#), on page 120
- [Creating a Syslog Policy](#), on page 122
- [Creating a Switch Control Policy](#), on page 123
- [Creating a Flow Control Policy](#), on page 129
- [Creating a Link Aggregation Policy](#), on page 130
- [Creating a Link Control Policy](#), on page 131
- [Creating a Multicast Policy](#), on page 133

Domain Policies

Domain policies in Cisco Intersight allow you to configure various parameters for UCS Fabric Interconnects, including port configuration, network control settings, and VLAN and VSAN settings. A domain policy can be assigned to any number of domain profiles to provide a configuration baseline. Domain policies in Cisco Intersight are a new feature, and native to the application. Policy-based configuration with Domain Profiles is a Cisco Intersight Essentials feature, and is supported on Cisco UCS B-Series M5 and M6 servers and Cisco UCS C-Series M5, M6, M7, and M8 servers, and Cisco UCS X-Series M6 and M7 servers that are in a UCS Domain.

The Domain Policy creation wizard in Cisco Intersight has two pages:

- **General**—The general page allows you to select the organization and enter a name for your policy. Optionally, include a short description and tag information to help identify the policy. Tags must be in the key:value format. For example, Org:IT or Site APJ
- **Policy Details**—The policy details page has properties that are applicable to UCS Domain Policies.

The following list describes the domain policies that you can configure in Cisco Intersight.

- **Port Policy**—Configures the ports and port roles for the Fabric Interconnect. Each Fabric Interconnect has a set of ports in a fixed port module that you can configure. You can enable or disable a port or a port channel.

The port policy is associated with a switch model. The network configuration limits also vary with the switch model.

- The maximum number of ports and port channels supported on Cisco UCS 6400 Series and 6500 Series Fabric Interconnects are:
 - Ethernet Uplink, Fibre Channel over Ethernet (FCoE) Uplink port channels, and Appliance port channels (combined) — 12
 - Ethernet Uplink ports per port channel — 16
 - FC uplink port channel — 4
 - FCoE Uplink ports per port channel —16
 - Ethernet Uplink and FCoE Uplink ports (combined) —31
 - Server ports—54 ports for Cisco UCS 6454 and 108 ports for Cisco UCS 64108 Fabric Interconnects

The maximum number of ports and port channels supported for Cisco UCS Fabric Interconnects 9108 100G are:

- Ethernet uplink, Fibre Channel over Ethernet (FCoE) uplink, Ethernet uplink port channels, FCoE uplink port channels, and Appliance port channels (combined) —8
 - Ethernet uplink ports per port channel —8
 - FC uplink port channel —4
 - FC uplink count —8
 - FC uplink per SAN port —8
 - FCoE Uplink ports per port channel —8
 - Ethernet Uplink and FCoE uplink ports (combined) —8
- **Ethernet Network Control Policy**—Configures the network control settings for appliance ports, appliance port channels, or vNICS.
 - **Ethernet Network Group Policy**—Configures the VLAN settings that include Native VLAN and QinQ VLAN for appliance ports, appliance port channels, or vNICs.



Note When an Ethernet Network Group is assigned to a Port Policy, the specified VLAN set must be either identical to or disjoint from the VLAN sets specified on other uplink interfaces. Ensure that the VLANs are defined in the VLAN Policy and that Auto Allow on Uplinks is disabled.

- **VLAN Configuration Policy**—Creates a connection to a specific external LAN.

- **VSAN Configuration Policy**—Partitions the Fibre Channel fabric into one or more zones. Each zone defines the set of Fibre Channel initiators and Fibre Channel targets that can communicate with each other in a VSAN.
- **NTP Policy**—Enables the NTP service to configure a UCS system that is managed by Cisco Intersight to synchronize the time with an NTP server. You must enable and configure the NTP service by specifying the IP/DNS address of at least one server or a maximum of four servers that function as NTP servers. When you enable the NTP service, Cisco Intersight configures the NTP details on the endpoint. For more information, see [Creating an NTP policy](#).
- **Network Connectivity Policy**—Specifies the DNS Domain settings that are used to add or update the resource records on the DNS server from the endpoints, and the DNS server settings for IPv4 and IPv6 on an endpoint.
- **System QoS Policy**—Implements network traffic prioritization based on the importance of the connected network by assigning system classes for individual vNICs. Intersight uses Data Center Ethernet (DCE) to handle all traffic inside a Cisco UCS domain. This industry standard enhancement to Ethernet divides the bandwidth of the Ethernet pipe into eight virtual lanes. Two virtual lanes are reserved for internal system and management traffic. You can configure quality of service (QoS) for the other six virtual lanes. System classes determine how the DCE bandwidth in these six virtual lanes is allocated across the entire Cisco UCS domain.

Each system class reserves a specific segment of the bandwidth for a specific type of traffic, which provides a level of traffic management, even in an oversubscribed system. For example, you can configure the **Fibre Channel Priority** system class to determine the percentage of DCE bandwidth allocated to FCoE traffic. The configuration setup validates each input on the system class to prevent duplicate or invalid entries.

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

- **Platinum, Gold, Silver, and Bronze**—A configurable set of system classes that you can include in the QoS policy for a service profile. Each system class manages one lane of traffic. All properties of these system classes are available for you to assign custom settings and policies.
- **Best Effort**—A system class that sets the quality of service for the lane reserved for basic Ethernet traffic. Some properties of this system class are preset and cannot be modified. For example, this class has a drop policy that allows it to drop data packets if required. You cannot disable this system class.
- **Fibre Channel**—A system class that sets the quality of service for the lane reserved for Fibre Channel over Ethernet traffic. Some properties of this system class are preset and cannot be modified. For example, this class has a no-drop policy that ensures it never drops data packets. You cannot disable this system class.
- **Multicast Policy**—Configures Internet Group Management Protocol (IGMP) snooping and IGMP querier. IGMP Snooping dynamically determines hosts in a VLAN that should be included in multicast transmissions.

You can create, modify, and delete a multicast policy that can be associated to one or more VLANs. When a multicast policy is modified, all VLANs associated with that multicast policy are re-processed to apply the changes. By default, IGMP snooping is enabled and IGMP querier is disabled. On enabling IGMP querier, you can configure the IPv4 addresses for the local and peer IGMP snooping querier interfaces.

- **Simple Network Management Protocol (SNMP) Policy**—Configures the SNMP settings for sending fault and alert information by SNMP traps from the managed devices. Any existing SNMP Users or SNMP Traps configured previously on the managed devices are removed and replaced with users or traps that you configure in this policy.
- **Syslog Policy**—Enables to configure the local logging and remote logging (minimum severity) for an endpoint. This policy also provides configuration support to store the syslog messages in the local file and the remote syslog server.
- **Switch Control Policy**—Enables to configure and manage multiple network operations on the Fabric Interconnects (FI) that include:
 - **Port Count Optimization**—If the VLAN port count optimization is enabled, the Virtual Port (VP) groups are configured on the Fabric Interconnect (FI) and if VLAN port count optimization is disabled, the configured VP groups are removed from the FI.
 - **MAC Aging Time**—Allows to set the MAC aging time for the MAC address table entries. The MAC aging time specifies the time before a MAC entry expires and discards the entry from the MAC address table.
 - **Link Control Global Settings**—Enables configurations of message interval time in seconds and allows to reset the recovery action of an error-disabled port.
- **Flow Control Policy**—Enables configurations for Priority Flow Control for ports and port channels.
- **Link Control Policy**—Enables configurations of Link Control administrative state and configuration (normal or aggressive) mode for ports.
- **Link Aggregation Policy**— Enables to configure Link Aggregation properties. Link Aggregation combines multiple network connections in parallel to increase throughput and to provide redundancy.

Creating a Port Policy

The port policy is used for configuring the port parameters such as unified ports that carry Ethernet or Fibre Channel traffic, port roles and speed.

1. Log in to Cisco Intersight with your Cisco ID and select admin role.
2. From the **Service Selector** drop-down list, select **Infrastructure Service**.
3. Navigate to **Configure > Policies**, and then click **Create Policy**.
4. Select **Port**, and then click **Start**.
5. On the **General** page, configure the following parameters:

Property	Essential Information
Organization	Select the Organization.
Name	Enter a name for your policy.

Property	Essential Information
Switch Model	<p>Select any one of the following switch models:</p> <ul style="list-style-type: none"> • Cisco UCS 64108 Fabric Interconnect • Cisco UCS 6454 Fabric Interconnect • Cisco UCS 6536 Fabric Interconnect • Cisco UCS Fabric Interconnect 9108 100G <p>Note The switch models provide different network configuration capabilities to the policy. The switch model cannot be changed once the policy is created.</p>
Description (Optional)	Provide a short description
Add Tag (Optional)	Enter a tag in the key:value format. For example, Org: IT or Site: APJ.

6. On the **Policy Details** page, configure the following parameters:

Property	Essential Information
Unified Ports	<p>By default, all the unconfigured ports are Ethernet ports. Use the blue slider to select a range of Fibre Channel ports. The selected Fibre Channel ports are highlighted in blue.</p>
Fibre Channel (FC)	<p>Displays the port range selected for Fibre Channel.</p> <p>Note</p> <ul style="list-style-type: none"> • Valid FC port range for Cisco UCS 6454 Fabric Interconnect: Port 1-16 • Valid FC port range for Cisco UCS 64108 Fabric Interconnect: Port 1-16 • Valid FC port range for Cisco UCS 6536 Fabric Interconnect: Port 33-36 • Valid FC port range for Cisco UCS Fabric Interconnects 9108 100G: Port 1-2
Ethernet	Displays the port range selected for Ethernet.

7. On the **Breakout Options** page, configure the breakout ports on Fibre Channel or Ethernet.



Note To configure breakout port, you must upgrade your Fabric Interconnect firmware to firmware version 4.2(2a) and above.

In Cisco UCS 6536 Fabric Interconnect, only FC breakout is supported.

- Select the ports for breakout either by clicking on the valid ports within the graphic image or by selecting the port number in the table present below the image.

Following are the breakout port range for different Cisco UCS Fabric Interconnects:

- Cisco UCS 64108 Fabric Interconnect, the valid breakout port range is 97—108
- Cisco UCS 6454 Fabric Interconnect, the valid breakout port range is 49—54
- Cisco UCS 6536 Fabric Interconnects. the valid breakout port range is 1—36
- Cisco UCS Fabric Interconnects 9108 100G, the valid breakout port range is 1—8

- Click **Configure**.

A pop-up window appears. It displays the admin speeds that can be set for the breakout ports.

Ethernet breakout ports can be configured with three options : no breakout, Admin speed of 4x10G, and Admin speed of 4x25G

FC breakout ports can be configured in three different **Admin Speed**: 4x8G, 4x16G, and 4x32G

- Select the desired speed.



Note You can configure Ethernet breakout and switch between breakout speeds without requiring a FI reboot.

Changing the FC breakout speeds does not require FI reboot.

Switching from the Ethernet breakout to the FC breakout and vice versa, or from the Ethernet port to the FC breakout port and vice versa, requires an FI reboot each time.

- Click **Set**.
- Click **Next**.

8. On the **Port Roles** page, select the ports that have to be configured for port roles either in the graphic image or by selecting in the table present below the graphic image.

Selected Ports	Indicates the port number(s) selected.
Name	The user determined port name.
Type	The type can be Ethernet or FC .

<p>Role</p>	<p>Select the port role type:</p> <p>The roles for an Ethernet port are:</p> <ul style="list-style-type: none"> • Unconfigured—Default • Server—All server traffic travels through the input or output (I/O) module to server ports on the fabric interconnect. <p>Note</p> <ul style="list-style-type: none"> • For Cisco UCS 6454 Fabric Interconnect, the maximum number of server ports allowed is 54. For Cisco UCS 64108 Fabric Interconnect, the maximum number of server ports allowed is 108. • For Cisco UCS 6536 Fabric Interconnect, server roles are not supported on 10G breakout ports. • Server port configuration is supported for discovering direct-attach Cisco UCS C-Series servers only after configuring breakout port on Ports 49-54 for Cisco UCS 6454 Fabric Interconnect and on Ports 97-108 for Cisco UCS 64108 Fabric Interconnect. • Discovering chassis, blade server connected to chassis, or rack servers connected to FEX are not supported after configuring breakout port on Ports 49-54 for Cisco UCS 6454 Fabric Interconnect and on Ports 97-108 for Cisco UCS 64108 Fabric Interconnect. • Cisco UCS Fabric Interconnects 9108 100G, Server role is not available for port role configuration. <ul style="list-style-type: none"> • Ethernet Uplink—Ethernet traffic passes through the unified uplink port <p>Note The maximum number of combined Ethernet Uplink ports and FCoE Uplink ports allowed is 31.</p> <ul style="list-style-type: none"> • Appliance—Allows the Network File System to connect directly with the Fabric Interconnects, without traffic having to pass through the uplink ports. <p>The roles for an FC port are:</p> <ul style="list-style-type: none"> • FC Uplink —FC traffic passes through the FC uplink port. To specify the role of an FC port as an FC Uplink port the VSAN scope of the port must have been created as Storage and Uplink, or as Uplink in the VSAN Configuration policy. • FC Storage—FC port acts as a storage port. To specify the role of an FC port as an FC Storage port the VSAN scope of the port must have been created as Storage and Uplink, or as Storage in the VSAN Configuration policy. Moreover, the FC has to be in the switching mode. • Unconfigured—Unconfigured is the default role of the port.
--------------------	---

Admin Speed	<p>The administrative port speed. The options are:</p> <ul style="list-style-type: none"> • 1GBPS • 10GBPS • 25GBPS • 40GBPS • 100GBPS <p>Note</p> <ul style="list-style-type: none"> • Admin Speed cannot be selected for any role on breakout ports. • For Cisco UCS 6536 Fabric Interconnect, only 25G/40G/100G connectivity is supported for server ports. • For Cisco UCS Fabric Interconnects 9108 100G, 1Gbps speed is available only for Ports 7 and 8. <p>Note When the 25GBPS admin speed is selected, Enable 25GBPS Copper Cable Negotiation is automatically enabled for any copper cable that is more than 3 meters.</p> <p>Enable 25GBPS Copper Cable Negotiation:</p> <ul style="list-style-type: none"> • Supports only on Appliance, Ethernet Uplink, FCoE Uplink port roles. • Does not support breakout ports. • Supports firmware versions 4.2(1a) or higher. • Supports only for the FEC configuration set to Auto.
VSAN ID	The VSAN ID of an FC port as specified in the VSAN Configuration policy.
FEC	<p>The forward error correction configuration for the port:</p> <ul style="list-style-type: none"> • Auto • CI91—Supported with 25GBPS and 100GBPS Admin speed <p>Note CI91 is not present for <i>Server</i> Port role.</p> <ul style="list-style-type: none"> • CI74—Supported with 25GBPS Admin speed
Priority	Select the priority of the port for routing traffic and ensuring QoS.
Mode	Select the port mode. Port mode can be Trunk or Access.

<p>Connected Device Type and Device Number</p>	<p>Select the device type and device number for each port or a set of ports.</p> <p>Note This option is applicable for Server Roles only.</p> <p>By default, this option is disabled.</p> <p>To enable:</p> <ul style="list-style-type: none"> • Select the ports and click Configure. • Turn the Manual Chassis/Server Numbering button ON. <p>A table is displayed where you can specify the Connected Device Type and Device Number for each port.</p> <p>Note Auto-Fill Numbering can be enabled to edit Connected Device Type, Starting Device Number, and Ports per Device for each port according to your preferences.</p> <ul style="list-style-type: none"> • Click Save to see the Connected Device Type and Device Number columns in the Port Roles list view. <p>Note If the selected Device Number is already allocated for any other server/chassis on any other port then the next available number will be allocated to the server that is discovered. This action will not result in failure of Port Policy deployment.</p> <p>Note The Port Policy changes are not applicable for FEX.</p>
<p>Ethernet Network Group</p>	<p>Select the Ethernet Network Group policy that is to be attached to the ethernet uplink or appliance port. The Ethernet Network Group policy specifies the Allowed VLANs and the Native VLAN.</p> <p>Note Ethernet Network Group policy applies only for ports with ethernet uplink and appliance roles.</p> <p>Note To create Ethernet Network Groups for configuring Disjoint VLANs, ensure that the groups are completely disjoint. Partial overlap of VLANs is not allowed.</p>
<p>Ethernet Network Control</p>	<p>Select the Ethernet Network Control policy that is to be attached to the appliance port. The Ethernet Network Control policy allows you to enable or disable CDP, specify the MAC Register Mode, the action to be taken on uplink fail, the MAC security details and LLDP details.</p> <p>Note Ethernet Network Control policy applies only for a port with an appliance role.</p>
<p>Port</p>	<p>Select the valid port range:</p> <ul style="list-style-type: none"> • Port 1-96—Auto, 10GBPS, and 25GBPS • Port 89-96—Auto, 1GBPS, 10GBPS, and 25GBPS • Port 97-108—Auto, 40GBPS, and 100GBPS

Port Channels

Click **Create Port Channel** where you can choose the role for the selected ports.

Select the ports for configuration either by clicking on the ports within the graphic image or in the box next to the desired port within the table.

<p>Role</p>	<p>The port channel role type. The role type can be:</p> <ul style="list-style-type: none"> • Ethernet Uplink Port Channel • FC Uplink Port Channel • FCoE Uplink Port Channel • Appliance Port Channel <p>Note</p> <ul style="list-style-type: none"> • The maximum number of ports allowed for: <ul style="list-style-type: none"> • Ethernet Uplink port channel, FCoE Uplink port channel, and Appliance port channel (combined) is 12 • FC uplink port channel is 4 • Ethernet ports per port channel is 16 • FCoE Uplink ports per port channel is 16 • You cannot combine normal ports and breakout ports for any port channel. For example, Uplink port channel ID 100 with members 1/96 and 1/97/1 are not allowed. • If a port with a speed of 100G in Cisco UCS 6536 Fabric Interconnect, is connected with N9K-C93180YC-FX3, then you must disable Auto Negotiation while assigning the port role. • For FC uplink Port Channel, port channel with different port speed is not allowed. For example, FC uplink port channel ID 101 with member 1/33 with port speed 8Gbps and 1/34 with port speed 16Gbps are not allowed.
<p>PC ID</p>	<p>Unique Identifier of the port channel, local to this switch.</p>

Admin Speed	<p>The administrative port channel speed options for Uplink, Uplink Port Channel, and FCoE Uplink Port Channel are:</p> <ul style="list-style-type: none"> • 1GBPS • 10GBPS • 25GBPS • 40GBPS • 100GBPS <p>The administrative port channel speed options for FC Uplink and FC Uplink Port Channel are:</p> <ul style="list-style-type: none"> • 8GBPS • 16GBPS • 32GBPS <p>Note You cannot select Admin Speed for any roles on breakout ports.</p>
Priority	Select the priority of the port channel for routing traffic and ensuring QoS.
Mode	Select the port channel mode. Port channel mode can be Trunk or Access.
Ethernet Network Group	<p>Select the Ethernet Network Group policy that is to be attached to the ethernet uplink or appliance port channel. The Ethernet Network Group policy specifies the Allowed VLANs and the Native VLAN.</p> <p>Note Ethernet Network Group policy applies to port channels with ethernet uplink and appliance roles.</p> <p>Note To create Ethernet Network Groups for configuring Disjoint VLANs, ensure that the groups are completely disjoint. Partial overlap of VLANs is not allowed.</p>
Ethernet Network Control	<p>Select the Ethernet Network Control policy that is to be attached to the appliance port channel. The Ethernet Network Control policy allows you to enable or disable CDP, specify the MAC Register Mode, the action to be taken on uplink fail, the MAC security details and LLDP details.</p> <p>Note Ethernet Network Control policy applies only for a port channel with an appliance role.</p>
Port Channel	Select the valid port channel range between 1 and 256.

Pin Groups	
<p>Pin Group is used to pin Ethernet/FC traffic from a vNIC/vHBA on a server to an uplink Ethernet/FC port or port channel on the Fabric Interconnect. You can use this pinning to manage the distribution of traffic from the servers. Static pinning is not supported when FI are in Switching Mode (Ethernet and FC).</p> <p>To configure pinning for a server, you must include the LAN/SAN pin group in the LAN/SAN connectivity policy.</p> <p>Click Create Pin Group to specify the ports/port channels in the FI through which the LAN and SAN data traffic can be made to flow.</p>	
Pin Group Type	<p>The type of the data traffic that needs to flow to the pinned ports/port channels. The type can be</p> <ul style="list-style-type: none"> • LAN • SAN
Pin Group Name	<p>The name of the Pin Group. This name will appear in LAN/SAN Connectivity policy creation page, once the Pin Group is created.</p>
Interface Type	<p>The type of the interface on the Fabric Interconnect.</p> <ul style="list-style-type: none"> • Port • Port Channels
Port Selection	<p>From the available table, you can select the ports and the breakout ports that should be pinned for data traffic flow.</p> <p>It is enabled by default.</p>

9. Click **Save**.

Creating an Ethernet Network Group Policy

An Ethernet Network Group policy enables you to manage settings for VLANs on a UCS Server. These settings include defining which VLANs are allowed, designating a Native VLAN, and specifying a QinQ VLAN.



Note When an Ethernet Network Group is assigned to a Port Policy, the specified VLAN set must be either identical to or disjoint from the VLAN sets specified on other uplink interfaces. Ensure that the VLANs are defined in the VLAN Policy and that **Auto Allow on Uplinks** is disabled.

This policy also supports VIC QinQ Tunneling. A QinQ (802.1Qin802.1Q) tunnel allows segregation and isolation of different VLANs within a network. To configure QinQ VLAN, you can specify the desired VLAN ID as part of the VLAN settings for the specific port, port channel, or vNIC. This enables the transmission of multiple VLANs over a single VLAN trunk.

1. Log in to Cisco Intersight with your Cisco ID and select admin role.
2. From the **Service Selector** drop-down list, select **Infrastructure Service**.
3. Navigate to **Configure > Policies**, and then click **Create Policy**.
4. Select **Ethernet Network Group**, and then click **Start**.
5. On the **General** page, configure the following parameters:

Property	Essential Information
Organization	Select the Organization.
Name	Enter a name for your policy.
Set Tag (Optional)	Enter a tag in the key:value format. For example, Org: IT or Site: APJ.
Description (Optional)	Provide a short description

6. On the **Policy Details** page, configure the following parameters:

Property	Essential Information
VLAN Settings	
Native VLANs	<p>This property allows you to specify the native VLAN ID for the virtual interface or its corresponding vethernet in a range of 1-4093.</p> <ul style="list-style-type: none"> • If the native VLAN is not already part of the allowed VLANs, it will be automatically added to the list of allowed VLANs. • If QinQ Tunneling is enabled, the native VLAN and Allowed VLAN properties are combined.
Enable QinQ Tunneling	Slide to enable VIC QinQ (802.1Qin802.1Q) Tunneling.
Allowed VLANs	<p>Refers to the VLANs that are permitted for the virtual interface. You can specify the allowed VLANs by providing a list of comma-separated VLAN IDs and VLAN ID ranges.</p> <p>For example, you can enter VLAN IDs 10, 20, 30-40 to allow VLANs 10, 20, and a range from 30 to 40.</p> <p>Note This property is displayed only when <i>Enable QinQ Tunneling</i> slider is disabled.</p>

Property	Essential Information
QinQ VLAN	<p>This property enables the configuration of QinQ Tunneling, that facilitates the encapsulation of multiple VLANs within a single VLAN. The supported VLAN IDs range from 2 to 4093 that allows you to effectively manage and segregate the network traffic.</p> <p>Note This property is available only when <i>Enable QinQ Tunneling</i> slider is enabled.</p>



Note To make the server an Isolated host or a Community host, specify the ID of an Isolated VLAN or a Community VLAN in both Allowed VLANs and Native VLAN

7. Click **Create**.

Creating an Ethernet Network Control Policy

Ethernet Network Control policies configure the network control settings for the UCS Domain. This policy is applicable only for the Appliance Ports defined in a Port Policy and for the vNICs defined in a LAN Connectivity Policy, on an FI-Attached UCS Servers.

1. Log in to Cisco Intersight with your Cisco ID and select admin role.
2. From the **Service Selector** drop-down list, select **Infrastructure Service**.
3. Navigate to **Configure > Policies**, and then click **Create Policy**.
4. Select **Ethernet Network Control**, and then click **Start**.
5. On the **General** page, configure the following parameters:

Property	Essential Information
Organization	Select the Organization.
Name	Enter a name for your policy.
Description (Optional)	Provide a short description
Add Tag (Optional)	Enter a tag in the key:value format. For example, Org: IT or Site: APJ.

6. On the **Policy Details** page, configure the following parameters:

Property	Essential Information
Enable CDP	Enables the Cisco Discovery Protocol (CDP) on an interface.

Property	Essential Information
MAC Register Mode	<p>Determines the MAC addresses to be registered with the switch. This can be:</p> <ul style="list-style-type: none"> • Only Native VLAN—MAC addresses are only added to the native VLAN. This option is the default, and it maximizes the port+VLAN count. • All Host VLANs—MAC addresses are added to all VLANs with which they are associated. Select this option if your VLANs are configured to use trunking but are not running in Promiscuous mode.
Action on Uplink Fail	<p>Determines how the interface behaves if no uplink port is available when the switch is in end-host mode.</p> <ul style="list-style-type: none"> • Link Down—Changes the operational state of a vNIC to down when uplink connectivity is lost on the switch, and enables fabric failover for vNICs. This is the default option. • Warning—Maintains server-to-server connectivity even when no uplink port is available, and disables fabric failover when uplink connectivity is lost on the switch.
MAC Security Forge	<p>Determines whether forged MAC addresses are allowed or denied when packets are sent from the server to the switch. This can be:</p> <ul style="list-style-type: none"> • Allow— All server packets are accepted by the switch, regardless of the MAC address associated with the packets. This is the default option. • Deny— After the first packet has been sent to the switch, all other packets must use the same MAC address or they will be silently rejected by the switch. In effect, this option enables port security for the associated vNIC.
LLDP	<p>Determines whether interfaces can transmit or receive LLDP packets.</p> <ul style="list-style-type: none"> • To enable or disable the transmission of LLDP packets on an interface, click Enable Transmit. • To enable or disable the receipt of LLDP packets on an interface, click Enable Receive.

- Click **Create**.

Creating a VLAN Policy

VLAN policies create a connection to a specific external LAN. The VLAN isolates traffic to that external LAN, including broadcast traffic. You can create VLANs and Private VLANs using the VLAN policy.



Note Ensure that each VLAN is associated with a multicast policy. You can edit the existing VLANs and associate them to a multicast policy. You cannot associate a Multicast policy to a Private VLAN.

- Log in to Cisco Intersight with your Cisco ID and select admin role.
- From the **Service Selector** drop-down list, select **Infrastructure Service**.
- Navigate to **Configure > Policies**, and then click **Create Policy**.
- Select **VLAN**, and then click **Start**.
- On the **General** page, configure the following parameters:

Property	Essential Information
Organization	Select the Organization.
Name	Enter a name for your policy.
Description (Optional)	Provide a short description
Add Tag (Optional)	Enter a tag in the key:value format. For example, Org: IT or Site: APJ.

- On the **Policy Details** page, click **Add VLAN** and configure the following policy details:



Note The maximum number of VLANs allowed per Ethernet Network Policy is 3000.

Property	Essential Information
Add VLANs	Click Add VLANs to add VLANs and Private VLANs
Name/Prefix	For a single VLAN, this is the VLAN name. For a range of VLANs, this is the prefix that the system uses for each VLAN name.

Property	Essential Information
VLAN IDs	<p>Enter the VLAN ID number or a number range between 2 and 4093. You can enter a range of IDs using a hyphen, and you can enter multiple IDs or ID ranges separated by commas. Examples of valid VLAN IDs or ID ranges are 50, 200, 2000-2100. You cannot use VLANs from 3915-4042, 4043-4047, 4094, and 4095 because these IDs are reserved for system use.</p> <p>The name that you assign to a VLAN ID adds a layer of abstraction that allows you to globally update all servers associated with service profiles that use the named VLAN.</p>
Auto Allow on Uplinks	<p>Used to determine whether this VLAN will be allowed on all uplink ports and port channels in this Fabric Interconnect.</p> <p>Enable to allow this VLAN on uplink ports and port channels.</p> <p>Disable to configure disjoint VLANs.</p>
Multicast Policy	<p>Click Select Policy and choose a Multicast policy that needs to be associated with VLAN.</p> <p>Click Create New to create a new Multicast policy that will be available to all VLANs.</p> <p>Note You cannot add Multicast policy for a Private VLAN.</p>
Enable VLAN Sharing	Enable to create Private VLANs.
Sharing Type	<p>The Sharing type can be:</p> <ul style="list-style-type: none"> • Primary: The Primary VLAN of a Private VLAN. Secondary VLANs are mapped to Primary VLANs. <p>Note You must create the Primary VLAN before creating the Isolated or Community VLANs.</p> <ul style="list-style-type: none"> • Isolated: One of the two Sharing Types of a Secondary VLAN. Only one Isolated VLAN can be mapped to a Primary VLAN. • Community: One of the Sharing Types of a Secondary VLAN. Multiple Community VLANs can be mapped to a Primary VLAN.

Property	Essential Information
Primary VLAN ID	The Primary VLAN to which a Community or Isolated VLAN is to be mapped. Note When a Secondary VLAN is mapped to a Primary VLAN, you cannot modify or delete the Primary VLAN.



Note If the VLAN configuration in the domain profile is modified, the corresponding changes in the server profile will take effect only after the server profile is redeployed.

- Click **Add**.

Creating a VSAN Policy

With the VSAN policy, you can create Virtual SANs (VSANs) to isolate devices physically connected to the same SAN fabric. VSANs improve security and stability in Fibre Channel fabrics and let you create several logical SANs over a common physical infrastructure.

- Log in to Cisco Intersight with your Cisco ID and select admin role.
- From the **Service Selector** drop-down list, select **Infrastructure Service**.
- Navigate to **Configure > Policies**, and then click **Create Policy**.
- Select **VSAN**, and then click **Start**.
- On the **General** page, configure the following parameters:

Property	Essential Information
Organization	Select the Organization.
Name	Enter a name for your policy.
Description (Optional)	Provide a short description
Add Tag (Optional)	Enter a tag in the key:value format. For example, Org: IT or Site: APJ.

- On the **Policy Details** page, do the following:
 - Click **Trunking Mode** to enable or disable Fibre Channel uplink trunking.

If you enable trunking for the named VSANs on a Fabric Interconnect, all named VSANs in the Cisco UCS domain are allowed on all Fibre Channel uplink ports on that Fabric Interconnect. If you configure Fabric Interconnects for Fibre Channel end-host mode, enabling Fibre Channel uplink trunking renders all VSANs with an ID in the range from 3840 to 4079 non-operational.
 - Click **Add VSAN** and configure the following policy details:

Property	Essential Information
Name	The user-defined name given to the VSAN configuration.
VSAN Scope	<p>The scope of the VSAN. Indicate if the VSAN is a storage and uplink VSAN, a storage VSAN, or an uplink VSAN</p> <p>VSAN Scope can be:</p> <ul style="list-style-type: none"> • Storage and Uplink • Storage • Uplink <p>Note If you want to create an FC Zone policy for a VSAN, then the VSAN scope must be Storage.</p>
VSAN ID	The unique identifier for the VSAN on the switch. The VSAN ID can be between 1 and 4093.
FCoE VLAN ID	<p>The unique identifier assigned to the VLAN used for Fibre Channel connections.</p> <p>IDs of FCOE VLANs associated with the VSAN configuration must be between 2 and 4093. VLAN IDs from 3915-4042, 4043-4047, 4094, and 4095 are reserved for system use.</p> <p>By default, VLAN 4048 is mapped to VSAN-1 on the switch. Attempting to use VLAN 4048 for FCoE in a VSAN Policy will result in an error. In this case, you need to explicitly configure VSAN-1 to use a different FCOE VLAN ID in the VSAN policy.</p>

7. Click **Create**.

Creating an NTP Policy

The NTP policy enables the NTP service to configure a UCS system that is managed by Cisco Intersight to synchronize the time with an NTP server. You must enable and configure the NTP service by specifying the IP/DNS address of at least one server or a maximum of four servers that function as NTP servers. When you enable the NTP service, Cisco Intersight configures the NTP details on the endpoint.

1. Log in to Cisco Intersight with your Cisco ID and select admin role.
2. From the **Service Selector** drop-down list, select **Infrastructure Service**.
3. Navigate to **Configure > Policies**, and then click **Create Policy**.

4. Select **NTP**, and then click **Start**.
5. On the **General** page, configure the following parameters:

Property	Essential Information
Organization	Select the Organization.
Name	Enter a name for your policy.
Description (Optional)	Provide a short description
Add Tag (Optional)	Enter a tag in the key:value format. For example, Org: IT or Site: APJ.

6. On the **Policy Details** page, configure the following parameters:

Property	Essential Information
Enable NTP	Enables NTP policy configuration.
NTP Servers	A collection of NTP Server IP addresses or hostnames.
Time Zone	A collection of time zones from which you can select a time zone for the endpoint. This property is applicable to switches and to Cisco IMC (standalone) servers.

When a hostname is used for NTP configuration, DNS server information must be configured in the Network Connectivity policy.

7. Click **Create**.

Creating a Network Connectivity Policy

The Network Connectivity policy enables you to configure and assign IPv4 and IPv6 addresses.

Dynamic DNS

Dynamic DNS (DDNS) is used to add or update the resource records on the DNS server. When you enable the DDNS option, the DDNS service records the current hostname, Domain name, and the management IP address and updates the resource records in the DNS server.

1. Log in to Cisco Intersight with your Cisco ID and select admin role.
2. From the **Service Selector** drop-down list, select **Infrastructure Service**.
3. Navigate to **Configure > Policies**, and then click **Create Policy**.
4. Select **Network Connectivity**, and then click **Start**.
5. On the **General** page, configure the following parameters:

Property	Essential Information
Organization	Select the Organization.
Name	Enter a name for your policy.
Description (Optional)	Provide a short description
Add Tag (Optional)	Enter a tag in the key:value format. For example, Org: IT or Site: APJ.

6. On the **Policy Details** page, configure the following properties:

Common Properties

Property	Essential Information
Enable Dynamic DNS	Enables Dynamic DNS. This property is not applicable to Fabric Interconnects.
Dynamic DNS Update Domain	Specify the dynamic DNS Domain. The Domain can be either a main Domain or a sub-Domain. This property is not applicable to Fabric Interconnects.

IPv4 Properties

Property	Essential Information
Obtain IPv4 DNS Server Addresses from DHCP	Whether the IPv4 addresses are obtained from Dynamic Host Configuration Protocol (DHCP) or from a specifically configured set of DNS servers. <ul style="list-style-type: none"> • Enabled—Intersight uses DHCP • Disabled—Intersight uses a configured set of IPv4 DNS servers. This property is not applicable to Fabric Interconnects.
Preferred IPv4 DNS Server	The IP address of the primary DNS server. This property is displayed only when Obtain IPv4 DNS Server Addresses from DHCP is disabled.
Alternate IPv4 DNS Server	The IP address of the secondary DNS server. This property is displayed only when Obtain IPv4 DNS Server Addresses from DHCP is disabled.

Property	Essential Information
Enable IPv6	Whether IPv6 is enabled. You can configure IPv6 properties only if this property is enabled.

IPv6 Properties

Property	Essential Information
Obtain IPv6 DNS Server Addresses from DHCP	<p>Whether the IPv6 addresses are obtained from Dynamic Host Configuration Protocol (DHCP) or from a specifically configured set of DNS servers.</p> <ul style="list-style-type: none"> • Enabled—Intersight uses DHCP • Disabled—Intersight uses a configured set of IPv6 DNS servers. <p>This property is not applicable to Fabric Interconnects.</p>
Preferred IPv6 DNS Server	The IP address of the primary DNS server. This property is displayed only when Obtain IPv6 DNS Server Addresses from DHCP is disabled.
Alternate IPv6 DNS Server	The IP address of the secondary DNS server. This property is displayed only when Obtain IPv6 DNS Server Addresses from DHCP is disabled.

7. Click **Create**.

Creating an SNMP Policy

The SNMP policy configures the SNMP settings for sending fault and alert information by SNMP traps from the managed devices. This policy supports SNMP versions such as SNMPv1, SNMPv2(includes v2c), and SNMPv3. Any existing SNMP Users or SNMP Traps configured previously on the managed devices are removed and replaced with users or traps that you configure in this policy.

Using the SNMP Policy you can enable or disable SNMP, specify the access and community strings, and provide the SNMP user details that is used to retrieve data.

1. Log in to Cisco Intersight with your Cisco ID and select admin role.
2. From the **Service Selector** drop-down list, select **Infrastructure Service**.
3. Navigate to **Configure > Policies**, and then click **Create Policy**.
4. Select **SNMP**, and then click **Start**.
5. In the **General** page, configure the following parameters:

Property	Essential Information
Organization	Select the organization.

Property	Essential Information
Name	Enter a name for your policy.
Tag (optional)	Enter a tag in the key value format. For example, Org: IT or Site: APJ.
Description (optional)	Enter a short description.

6. In the **Policy Details** page, configure the following parameters:

Property	Essential Information
Enable SNMP	Displays the state of the SNMP Policy on the endpoint. Enable this option for the endpoint to send SNMP traps to the designated host.
Access Community String	Enter the SNMPv1, SNMPv2 community string or the SNMPv3 username. This field allows maximum of 18 characters.
Trap Community String	Enter the SNMP community group name used for sending SNMP trap to other devices. Note This field is applicable only for SNMPv2c trap host or destination.
System Contact	The contact person responsible for the SNMP implementation. Enter a string up to 64 characters, such as an email address or a name and telephone number.
System Location	The location of host on which the SNMP agent (server) runs.
SNMP Users	
Name	Enter the SNMP username. This field must have a minimum of 1 and a maximum of 31 characters.
Security Level	Select the security mechanism for communication between the agent and the manager that include: <ul style="list-style-type: none"> • AuthPriv • AuthNoPriv
Auth Type	Select SHA as the authorization protocol for authenticating the user. Note The MD5 authorization protocol is not supported.
Auth Password	Enter the authorization password for the user.

Property	Essential Information
Auth Password Confirmation	Enter the authorization password confirmation for the user.
Privacy Type	Select AES as the privacy protocol for the user. Note The DES privacy type is deprecated to meet security standards.
Privacy Password	Enter the privacy password for the user.
Privacy Password Confirmation	Enter the privacy password confirmation for the user.
SNMP Trap Destinations	
Enable	Enable this option to use the SNMP policy.
SNMP Version	Select V2 or V3 as the SNMP version for the trap.
User	Select the SNMP user for the trap. You can define maximum of 15 trap users. Note This field is applicable only to SNMPv3.
Trap Type	Select the trap type to receive a notification when a trap is received at the destination: <ul style="list-style-type: none"> • Trap • Inform
Destination Address	Provide the address to which the SNMP trap information can be sent. You are allowed to define maximum of 10 trap destinations.
Port	Enter the port number for the server to communicate with trap destination. The range is from 1 to 65535. The default is 162.

7. Click **Create**.

Creating a System QoS Policy

A System Quality of Service (QoS) policy assigns a system class to the outgoing traffic. This system class determines the quality of service for the outgoing traffic.

1. Log in to Cisco Intersight with your Cisco ID and select admin role.
2. From the **Service Selector** drop-down list, select **Infrastructure Service**.
3. Navigate to **Configure > Policies**, and then click **Create Policy**.

4. Select **System QoS**, and then click **Start**.
5. On the **General** page, configure the following parameters:

Property	Essential Information
Organization	Select the Organization.
Name	Enter a name for your policy.
Description (Optional)	Provide a short description
Add Tag (Optional)	Enter a tag in the key:value format. For example, Org: IT or Site: APJ.

6. On the **Policy Details** page, configure the following parameters:

Property	Essential Information
Platinum Gold Silver Bronze	This option enables you to configure the associated QoS class on the fabric interconnect and assign the class to a QoS policy. Note The Best Effort or Fibre Channel system classes are enabled by default.
CoS	Set the class of service (CoS) by entering an integer value between 0 and 6, with 0 being the lowest priority and 6 being the highest priority. Set the value to 0 only when you require the system class to be the default system class for traffic if the QoS policy is deleted or the assigned system class is disabled.
Weight	An integer between 1 and 10. If you enter an integer, Cisco UCS determines the percentage of network bandwidth assigned to the priority level as described in the Weight field.
Allow Packet Drops	You can select to allow the packet drop for this system class during transmission. This field is always selected for the Best Effort class, which allows dropped packets, and always not selected for the Fibre Channel class, which never allows dropped packets.
MTU	The maximum transmission unit (MTU) for the channel. You can enter an integer between 1500 and 9216. This value corresponds to the maximum packet size.

7. Click **Create**.

Creating a Syslog Policy

The Syslog policy defines the minimum severity as logging level from an endpoint. The policy also defines the target destination to store the Syslog messages, and the Hostname or the IP Address, the port information, and the communication protocol for the Remote Logging Servers.

1. Log in to Cisco Intersight with your Cisco ID and select admin role.
2. From the **Service Selector** drop-down list, select **Infrastructure Service**.
3. Navigate to **Configure > Policies**, and then click **Create Policy**.
4. Select **Syslog**, and then click **Start**.
5. On the **General** page, configure the following parameters:

Property	Essential Information
Organization	Select the Organization.
Name	Enter a name for your policy.
Description (Optional)	Provide a short description
Add Tag (Optional)	Enter a tag in the key:value format. For example, Org: IT or Site: APJ.

6. On the **Policy Details** page, configure the following parameters:

Property	Essential Information
Local Logging	
Minimum Severity to Report	Select the lowest severity level to report in the remote log. The severity levels are: <ul style="list-style-type: none"> • 0 Emergency • 1 Alert • 2 Critical • 3 Error • 4 Warning • 5 Notice • 6 Informational • 7 Debug
Remote Logging - Syslog Server 1 and Syslog Server 2	
Enable	Select this option to enable or disable the Syslog policy.

Property	Essential Information
Hostname/IP Address	<p>Enter the hostname or IP address of the Syslog server to store the Cisco IMC log. You can set an IPv4 or IPv6 address or a domain name as the remote system address.</p> <p>Note If you have both IPv4 and IPv6 as the remote logging addresses, ensure to configure IPv4 and IPv6 in the Fabric Interconnect through the command-line interface (CLI).</p>
Minimum Severity To Report	<p>Select the lowest severity level to report in the remote log. The severity levels are:</p> <ul style="list-style-type: none"> • 0 Emergency • 1 Alert • 2 Critical • 3 Error • 4 Warning • 5 Notice • 6 Informational • 7 Debug

7. Click **Create**.

Creating a Switch Control Policy

The Switch Control policy supports VLAN port count optimization, configuring MAC address aging time, and configuring Link Control Global settings.

1. Log in to Cisco Intersight with your Cisco ID and select admin role.
2. From the **Service Selector** drop-down list, select **Infrastructure Service**.
3. Navigate to **Configure > Policies**, and then click **Create Policy**.
4. Select **Switch Control**, and then click **Start**.
5. In the **General** page, configure the following parameters:

Property	Essential Information
Organization	Select the organization.
Name	Enter a name for your policy.

Property	Essential Information
Tag (optional)	Enter a tag in the key value format. For example, Org: IT or Site: APJ.
Description (optional)	Enter a short description.

6. On the Policy Details page, configure the following parameters:

Property	Essential Information
Switching Mode	
Ethernet	<p>Specify the Ethernet switching mode. The switching mode can be End Host or Switch.</p> <p>In End Host mode, the Fabric Interconnects appear to the upstream devices as end hosts with multiple links. In this mode, the switch does not run Spanning Tree Protocol and avoids loops by following a set of rules for traffic forwarding.</p> <p>In Switch mode, the switch runs Spanning Tree Protocol to avoid loops, and broadcast and multicast packets are handled in the conventional way.</p>
FC	<p>Specify the FC switching mode. The switching mode can be End Host or Switch.</p> <p>End-host mode allows the Fabric Interconnect to act as an end host to the connected Fibre Channel networks, representing all servers (hosts) connected to it through vHBAs. The end-host mode is achieved by pinning (dynamically pinned or hard pinned) vHBAs to Fibre Channel uplink ports, which makes the Fibre Channel ports appear as server ports (N-ports) to the rest of the fabric. When in end-host mode, the Fabric Interconnect avoids loops by ensuring that uplink ports do not receive traffic from one another.</p> <p>Switch mode is the traditional Fibre Channel switching mode. Switch mode allows the Fabric Interconnect to connect directly to a storage device. Enabling Fibre Channel switch mode is useful in POD models where there is no SAN (for example, a single Cisco UCS system connected directly to storage), or where a SAN exists (with an upstream MDS).</p>
VLAN Port Count	

Property	Essential Information
Enable VLAN Port Count Optimization	<p>Select to enable the VLAN port count optimization. This option is disabled by default.</p> <p>Note</p> <ul style="list-style-type: none">• PV Count with VLAN Port Count Optimization Enabled on Cisco UCS 6400 Series and 6500 Series FI in IMM is 108000.• VLAN Port Count Optimization is always enabled for Cisco UCS Fabric Interconnects 9108 100G.
System Reserved VLANs	

Property	Essential Information
Reserved VLAN Start ID	<p>Select this option to specify the Start ID of the reserved VLAN range. By default, the Start ID is 3915. VLAN ID with Start ID + 127 cannot be used in configuring VLAN or VSAN policy. For example, if the VLAN Start ID is changed to 3912, the Reserved VLAN range is 3912-4039. The Reserved VLAN range cannot be used for user-defined VLAN or VSAN policy.</p> <p>Note Before you begin:</p> <ul style="list-style-type: none"> • Remove any existing VLANs in the new reserved VLAN range. • Ensure that there are no VLANs or FCoE VLANs in the reserved VLAN block being used in the VLAN or VSAN policy. In other words, ensure that the VLAN and VSAN policies in both Fabric Interconnect A and B do not conflict with the reserved VLAN range. • If the Reserved VLAN Start ID is changed, VLANs in the old range which are not included in the new range will be available for VLAN and VSAN policies after the new switch control policy is deployed. • The default reserved VLAN range is 3916–4095. This system reserved VLAN range can be changed but note that VLANs 1002-1005 are blocked for internal use and cannot be used as part of system reserved range. <p>Note</p> <ul style="list-style-type: none"> • Fabric Interconnect reboots for the changes to take effect. Reboot occurs only once even if multiple changes are made. • On a device unclaim, the previously configured reserved VLAN will not be removed. On a subsequent claim, users will have to configure reserved VLAN via the Switch Control Policy if they intend to use a new range.

Property	Essential Information
Reserved VLAN End ID	The End ID of the reserved VLAN range. The system blocks 128 reserved VLANs from the specified VLAN Start ID. By default, the End ID is 4042. This ID cannot be used in configuring VLAN policy.
MAC Address Table Aging Time	
Default	Select this option to set the default MAC address aging time to 14500 seconds for the End-Host mode.
Custom	Select this option to allow the user to configure the MAC address aging time on the switch. For the switch model UCS-FI-6454 or higher versions, the valid time range is 120 to 918000 seconds. After the time range is defined by the user, the switch resets the defined time to its lower multiple of 5.
Never	Select this option to disable the MAC address aging process. This option ensures the MAC entries never expire and are not discarded from the MAC address table.
Aging Time (Seconds)	Define the MAC address aging time in seconds. This field is valid only when the Custom option is selected.
Unidirectional Link Detection (UDLD) Global Settings	
Message Interval	Define the UDLD probe message interval (time in seconds) on ports that are in advertisement mode and are bidirectional. Note The valid message interval time ranges between 7 and 90 Seconds.
Recovery Action	Select Reset to recover an error-disabled port. Note The option None is selected by default.
Fabric port-channel vHBA	

Property	Essential Information
<p>Enable the fabric port-channel vHBA reset</p>	<p>A virtual host bus adapter (vHBA) logically connects a virtual machine to a virtual interface on the fabric interconnect and allows the virtual machine to send and receive traffic through that interface. This is currently accomplished by using the fibre channel modes (End Host mode/Switch mode).</p> <p>The port channel operations involve addition or removal of a member link between Fabric Interconnect and I/O Module (IOM). Such operations may result in a long I/O pause or connection drop from virtual machines to its targets and require a vHBA reset support.</p> <p>With the fabric port-channel vHBA reset set to enabled, when the Cisco UCS IOM port-channel membership changes, the Fabric Interconnect sends a Registered State Change Notification (RSCN) packet to each vHBA configured via that Cisco UCS IOM. The RSCN enables the virtual interface card (VIC) or VIC Driver to reset the Fabric port-channel vHBA and to restore the connectivity.</p> <p>By default, the Fabric port-channel vHBA reset is set to disabled.</p> <p>When disabled (default), vHBA reset is done only when all the members of a fabric port-channel are down.</p> <p>Note</p> <ul style="list-style-type: none"> • The feature is supported on Cisco Intersight Infrastructure firmware version 4.1(3e) and above. • ESX NFNIC driver version 5.0.0.37 and later or 4.0.0.87 and later process this RSCN. • Linux FNIC driver version 2.0.0.85 and later process this RSCN.

7. Click **Create**



Note

- On the Policy Details page, all the existing Switch Control policies show the value of Link Control Global Settings fields as blank. These policies display the correct values on policy edit/update.
- When you change the switching mode of a Fabric Interconnect, the Fabric Interconnect goes for a reboot.

Creating a Flow Control Policy

Configure the Priority Flow Control for each port, to enable the no-drop behavior for the CoS defined by the System QoS Policy and an Ethernet QoS policy. In Auto and On priorities, the Receive and Send link level flow control will be Off.

1. Log in to Cisco Intersight with your Cisco ID and select admin role.
2. From the **Service Selector** drop-down list, select **Infrastructure Service**.
3. Navigate to **Configure > Policies**, and then click **Create Policy**.
4. Select **Flow Control**, and then click **Start**.
5. On the **General** page, configure the following parameters:

Property	Essential Information
Organization	Select the Organization.
Name	Enter a name for your policy.
Add Tag (Optional)	Enter a tag in the key:value format. For example, Org: IT or Site: APJ.
Description (Optional)	Provide a short description

6. On the **Policy Details** page, configure the following parameters:

Property	Essential Information
Priority Flow Control Mode	
Auto	Auto receives and sends the priority flow. This field is enabled by default.
On	Enables priority control flow on the local port. Note You cannot enable Send and Receive direction at the same time.

Property	Essential Information
Off	Enables Link Level Flow Control on the local port. Note You can enable Send and Receive direction at the same time.
	Send When enabled, the link level flow control is configured in the send direction.
	Receive When enabled, the link level flow control is configured in the receive direction.



Note If Priority Flow Control is in **Auto/On** mode then the Flow Control cannot be enabled and the options are not listed. To enable Flow Control, you must set the Priority Flow Control in **Off** mode.



Note Flow Control should be enabled only on interfaces that are connected to Flow Control capable devices. The following interface types are supported:

- Ethernet uplink ports and port channels

7. Click **Create**.

Creating a Link Aggregation Policy

This policy can be used to configure Link Aggregation properties.

1. Log in to Cisco Intersight with your Cisco ID and select admin role.
2. From the **Service Selector** drop-down list, select **Infrastructure Service**.
3. Navigate to **Configure > Policies**, and then click **Create Policy**.
4. Select **Link Aggregation**, and then click **Start**.
5. On the **General** page, configure the following parameters:

Property	Essential Information
Organization	Select the Organization.
Name	Enter a name for your policy.
Add Tag (Optional)	Enter a tag in the key:value format. For example, Org: IT or Site: APJ.
Description (Optional)	Provide a short description

6. On the **Policy Details** page, configure the following parameters:

Property	Essential Information
Suspend Individual	
False	Select False to continue to receive PDUs from the peer port.
True	Select True to suspend a port that is not receiving the PDUs from the peer port.
LACP Rate	
Normal	The port is expected to receive 1 PDU every 30 seconds. The timeout for this is 90 seconds.
Fast	The port is expected to receive 1 PDU every 1 second from the peer port. The time out for this is 3 seconds.



Note Link Aggregation should be enabled only on interfaces that are connected to link aggregation capable devices. The following interface types are supported:

- Ethernet uplink port channel
- FCoE uplink port channel

7. Click **Create**.

Creating a Link Control Policy

This policy enables configuration of link control administrative state and configuration (normal or aggressive) mode for ports.

1. Log in to Cisco Intersight with your Cisco ID and select admin role.
2. From the **Service Selector** drop-down list, select **Infrastructure Service**.

3. Navigate to **Configure > Policies**, and then click **Create Policy**.
4. Select **Link Control**, and then click **Start**.
5. On the **General** page, configure the following parameters:

Property	Essential Information
Organization	Select the Organization.
Name	Enter a name for your policy.
Add Tag (Optional)	Enter a tag in the key:value format. For example, Org: IT or Site: APJ.
Description (Optional)	Provide a short description

6. On the **Policy Details** page, configure the following parameters:

Property	Essential Information
Link Control Administrative State	
The link control state of the port configured and managed by the administrator.	
Link Control Mode	
Normal	Detects unidirectional links caused by misconnected interfaces on fiber-optic connections.
Aggressive	<p>Detects unidirectional links caused by to one-way traffic on fiber-optic and twisted-pair links and by misconnected interfaces on fiber-optic links.</p> <ul style="list-style-type: none"> • When UDLD Administrative State is disabled, the policy cannot be set to Aggressive mode • When configuring the UDLD Mode (normal or aggressive), ensure the same mode is configured on both sides of the unidirectional link.



Note Link Control policy should be enabled only on interfaces that are connected to link control capable devices. The following interface types are supported:

- Ethernet uplink ports
- FCoE uplink ports
- Ethernet uplink port channels
- FCoE uplink port channels

- Click **Create**.

Creating a Multicast Policy

The multicast policy is used to configure Internet Group Management Protocol (IGMP) snooping and IGMP querier.



Note Ensure that each VLAN is associated with a multicast policy. You can edit the existing VLANs and associate them to a multicast policy.

- Log in to Cisco Intersight with your Cisco ID and select admin role.
- From the **Service Selector** drop-down list, select **Infrastructure Service**.
- Navigate to **Configure > Policies**, and then click **Create Policy**.
- Select **Multicast**, and then click **Start**.
- On the **General** page, configure the following parameters:

Property	Essential Information
Organization	Select the Organization.
Name	Enter a name for your policy.
Description (Optional)	Provide a short description
Add Tag (Optional)	Enter a tag in the key:value format. For example, Org: IT or Site: APJ.

- On the **Policy Details** page, configure the following parameters:

Property	Essential Information
Snooping State	<p>Determines whether IGMP snooping examines IGMP protocol messages within a VLAN to discover which interfaces are connected to hosts or other devices interested in receiving multicast traffic. This can be one of the following:</p> <ul style="list-style-type: none"> • Enabled—IGMP snooping is used for VLANs associated with this policy. • Disabled—IGMP snooping is not used for associated VLANs.

Property	Essential Information
Querier State	<p>Determines whether IGMP snooping querier sends out periodic IGMP queries that trigger IGMP report messages from hosts that want to receive IP multicast traffic. This can be one of the following:</p> <ul style="list-style-type: none"> • Enabled—Periodic IGMP queries are sent out. • Disabled—No IGMP queries are sent out. This is the default option.
Querier IP Address	<p>The IPv4 address for the IGMP snooping querier interface.</p> <p>This field appears only when Querier State is enabled.</p>
Querier IP Address Peer	<p>(Optional) The IPv4 address for the peer IGMP snooping querier interface. The peer IP address is assigned to FI-B.</p> <p>This field appears only when Querier State is enabled.</p>

7. Click **Create**.



CHAPTER 9

Configuring Server Policies

- [Server Policies, on page 136](#)
- [Creating a Policy, on page 143](#)
- [Supported UCS Server Policies, on page 143](#)
- [Default Server Policy Configurations, on page 147](#)
- [Creating a Certificate Management Policy, on page 148](#)
- [Creating an Adapter Configuration Policy, on page 149](#)
- [Creating a LAN Connectivity Policy, on page 152](#)
- [Creating an Ethernet Adapter Policy, on page 160](#)
- [Creating an Ethernet QoS Policy, on page 169](#)
- [Creating an Ethernet Network Policy, on page 170](#)
- [Creating an Ethernet Network Group Policy, on page 175](#)
- [Creating an Ethernet Network Control Policy, on page 176](#)
- [Creating a SAN Connectivity Policy, on page 178](#)
- [Creating a Fibre Channel Adapter Policy, on page 184](#)
- [Creating a Fibre Channel Network Policy, on page 187](#)
- [Creating a Fibre Channel QoS Policy, on page 188](#)
- [Create FC Zone Policy, on page 189](#)
- [Creating a Firmware Policy, on page 191](#)
- [Creating a BIOS Policy, on page 192](#)
- [Creating a Boot Order Policy, on page 193](#)
- [Configuring an iSCSI Boot Policy, on page 202](#)
- [Creating an iSCSI Adapter Policy, on page 205](#)
- [Creating an iSCSI Static Target Policy, on page 206](#)
- [Creating a Device Connector Policy, on page 207](#)
- [Creating a Drive Security Policy, on page 208](#)
- [Creating a Disk Group Policy, on page 209](#)
- [Creating an IMC Access Policy, on page 212](#)
- [Creating an IPMI Over LAN Policy, on page 214](#)
- [Creating an LDAP Policy, on page 216](#)
- [Creating a Local User Policy, on page 220](#)
- [Creating an NTP Policy, on page 223](#)
- [Creating an SD Card Policy, on page 224](#)
- [Create a Serial Over LAN Policy, on page 226](#)

- [Create SSH Policy, on page 228](#)
- [Creating a Virtual KVM Policy, on page 229](#)
- [Creating a Virtual Media Policy, on page 230](#)
- [Creating a Network Connectivity Policy, on page 234](#)
- [Creating a SMTP Policy, on page 236](#)
- [Creating an SNMP Policy, on page 237](#)
- [Creating a Storage Policy, on page 240](#)
- [Creating a Syslog Policy, on page 251](#)
- [Creating a Power Policy for Server, on page 252](#)
- [Creating a Thermal Policy for Server, on page 254](#)
- [Creating vNIC or vHBA Templates, on page 255](#)

Server Policies

Policies in Cisco Intersight provide different configurations for UCS servers, including BIOS settings, firmware versions, disk group creation, Simple Mail Transfer Protocol (SMTP), Intelligent Platform Management Interface (IPMI) settings, and more. A policy that is once configured can be assigned to any number of servers to provide a configuration baseline. Policies in Cisco Intersight are native to the application and are not directly imported from the UCS Systems. Policy-based configuration with Server Profiles is a Cisco Intersight Essentials license tier functionality.

The Server Policy creation wizard in Cisco Intersight has two pages:

- **General**—The general page allows you to select the organization and enter a name for your policy. Optionally, include a short description and tag information to help identify the policy. Tags must be in the key:value format. For example, Org: IT or Site: APJ.
- **Policy Details**—The policy details page has properties that are applicable to standalone UCS servers, FI-attached UCS servers, or both. You can view these properties separately for **All Platforms**, **UCS Servers (Standalone)**, and **UCS Servers (FI-Attached)** by clicking on these options.

Server Policies can be imported as part of importing configuration details (server profiles and policies) of a Cisco C-Series Standalone server from Cisco IMC. For more information, see [Importing a Server Profile](#).

The following list describes the server policies that you can configure in Cisco Intersight.

- **Adapter Configuration Policy**—Configures the Ethernet and Fibre-Channel settings for the VIC adapter.
- **BIOS Policy**—Automates the configuration of BIOS settings on the managed devices. You can create one or more BIOS policies which contain a specific grouping of BIOS settings. If you do not specify a BIOS policy for a server, the BIOS settings remain as they are. If a BIOS policy is specified, the values that are specified in the policy replace any previously configured values on a server (including bare metal server configuration settings). To apply the BIOS policy settings, you must reboot the server.

To simplify creating a BIOS policy, you can use a pre-defined Cisco Provided BIOS Configuration. These configurations are ready to use without requiring any modifications. For more information, see [Creating a BIOS Policy](#).



Note If you are using a Cisco Provided Configuration in a policy, you can still modify the configurations while you create the policy. However, you cannot modify the default Cisco Provided Configurations.

- **Boot Order Policy**—Configures the linear ordering of devices and enables you to change the boot order and boot mode. You can also add multiple devices under various device types, rearrange the boot order, and set parameters for each boot device type.

The inventory view enables you to view the actual boot order configured on a server. The boot order displays the details that include device name, device type, configuration details such as Boot Mode (Legacy or UEFI), and Secure Boot Mode (Enabled or Disabled).



Note A device configured in the server profile of Boot Order Policy may not appear in the actual boot order, if the server BIOS does not detect the device during server boot.

Intersight provides a One-Time Boot (OTB) option to set a boot device that temporarily overrides the Boot Order Policy and the existing boot order. To set a One-Time Boot Device, select **Power Cycle** or **Power On** from the **Servers Table view** or from the **Server Details** page and toggle ON the **Set One Time Boot Device** Option. This operation attempts to boot from the One Time Boot device as part of the power cycle or power on action. After power cycle or power on, OTB configuration will be cleared to enable the next reboot to follow the default Boot Order.



Note

- The OTB option is available for servers that have been configured with a Boot Order Policy that is associated with a server profile. For a successful OTB configuration, you must deploy a server profile with a Boot Order Policy in Intersight in advance.
- Any out-of-band- boot order change will not reflect on the Intersight UI for OTB device configuration.

In the case of **PXE Boot** configuration, importing the server policy will not create the PXE device under boot policy if either the MAC address or both the slot and port are not present for a given PXE device under the Boot policy on the server. However, if both slot and port are present, boot order is set to **ANY** for the bootable interface on a given slot on the server. For non-VIC adapters you can configure PXE Boot with the MAC address, or both the slot and port, or slot only.

In the case of **SAN Boot** device configuration in the legacy mode, provide the boot target Logical Unit Number (LUN), device slot ID, interface name, and target WWPN. For **SAN Boot** device configuration in the Unified Extensible Firmware Interface (UEFI) mode, provide the bootloader name, description, and path in addition to the fields listed in the legacy mode.

In the case of **iSCSI Boot** provide the target interface details, authentication mechanism, and initiator IP source.

- In the case of **Non-Volatile Memory Express (NVMe) Boot**, configure the NVMe drive as bootable in the UEFI mode. During the server profile deployment, this NVMe configuration setting enables selecting the BIOS in a defined order.
- **Certificate Management Policy**—Allows you to specify the certificate details for an external certificate and attach the policy to servers. Cisco Intersight currently supports the following certificates:
 - Root CA certificates
 - IMC certificates
- **Disk Group Policy**—Disk Group Policy is now a part of Storage Policy.
- **Device Connector Policy**—Lets you choose the **Configuration from Intersight only** option to control configuration changes allowed from Cisco IMC. The **Configuration from Intersight only** option is enabled by default. You will observe the following changes when you deploy the Device Connector policy in Intersight:
 - Validation tasks will fail:
 - If Intersight Read-only mode is enabled in the claimed device.
 - If the firmware version of the Cisco UCS Standalone C-Series Servers is lower than 4.0(1).
 - If Intersight Read-only mode is enabled, firmware upgrades will be successful only when performed from Intersight. Firmware upgrade performed locally from Cisco IMC will fail.
 - IPMI over LAN privileges will be reset to **read-only** level if **Configuration from Intersight only** is enabled through the Device Connector policy, or if the same configuration is enabled in the Device Connector in Cisco IMC.



Attention The Device Connector Policy will not be imported as part of the Server Profile Import.

- **Ethernet Adapter Policy**—Governs the host-side behavior of the adapter, including how the adapter handles traffic. For each VIC Virtual Ethernet Interface, you can configure various features such as VXLAN, NVGRE, ARFS, Interrupt settings, and TCP Offload settings.

To simplify creating an Ethernet Adapter policy, you can use a pre-defined Cisco Provided Ethernet Adapter Configuration. These configurations are ready to use without requiring any modifications. For more information, see [Creating an Ethernet Adapter Policy](#).



Note If you are using a Cisco Provided Configuration in a policy, you can still modify the configurations while you create the policy. However, you cannot modify the default Cisco Provided Configurations.

- **Ethernet Network Policy**—Allows to define the port to carry single VLAN(Access) or multiple VLANs(Trunk) traffic. You can configure the Default VLAN and QinQ VLAN settings for vNICs. You can specify the VLAN to be associated with an Ethernet packet if no tag is found.

- **Ethernet Network Control Policy**—Configures the network control settings for the appliance ports, appliance port channels, or vNICs.
- **Ethernet Network Group Policy**—Configures the VLAN settings that include Native VLAN and QinQ VLAN for appliance ports, appliance port channels, or vNICs.
- **Ethernet QoS Policy**—Assigns a system class to the outgoing traffic for a vNIC. This system class determines the quality of service for the outgoing traffic. For certain adapters, you can also specify additional controls like burst and rate on the outgoing traffic.
- **Fibre Channel Adapter Policy**—Governs the host-side behavior of the adapter, including how the adapter handles traffic. You can enable FCP Error Recovery, change the default settings of Queues, and Interrupt handling for performance enhancement.

To simplify creating an Fibre Channel Adapter policy, you can use a pre-defined Cisco Provided Fibre Channel Adapter Configuration. These configurations are ready to use without requiring any modifications. For more information, see [Creating an Fibre Channel Adaptor Policy](#).



Note If you are using a Cisco Provided Configuration in a policy, you can still modify the configurations while you create the policy. However, you cannot modify the default Cisco Provided Configurations.

- **Fibre Channel Network Policy**—Governs the VSAN configuration for the virtual interfaces.
- **Fibre Channel QoS Policy**—Assigns a system class to the outgoing traffic for a vHBA. This system class determines the quality of service for the outgoing traffic. For certain adapters, you can also specify additional controls like burst and rate on the outgoing traffic.
- **IPMI over LAN Policy**—Defines the protocols for interfacing with a service processor that is embedded in a server platform. The Intelligent Platform Management Interface (IPMI) enables an operating system to obtain information about the system health and control system hardware and directs the Cisco IMC to perform the required actions. You can create an IPMI Over LAN policy to manage the IPMI messages through Cisco Intersight. You can assign these user roles to an IPMI user per session:
 - **admin**—IPMI users can perform all available actions. If you select this option, IPMI users with the "Administrator" user role can create admin, user, and read-only sessions on this server.
 - **read-only**—Can view information but cannot make any changes. IPMI users with the "Administrator", "Operator", or "User" user roles can only create read-only IPMI sessions, regardless of their other IPMI privileges.
 - **user**—IPMI users can perform some functions but cannot perform administrative tasks. If you select this option, IPMI users with the "Administrator" or "Operator" user role can create user and read-only sessions on this server.



Important The encryption key to use for IPMI Communication. The key must have an even number of hexadecimal characters and not exceeding 40 characters. You can use "00" to disable the encryption key use. If the encryption key specified is less than 40 characters, then the IPMI commands must add zeroes to the encryption key to achieve a length of 40 characters.

- **LAN Connectivity Policy**—Determines the connections and the network communication resources between the server and the LAN on the network. You must create the Ethernet Adapter, Ethernet QoS, and Ethernet Network policies as part of the LAN connectivity policy. For IMM servers, use a MAC pool, or static MAC addresses, to assign MAC addresses to servers and to identify the vNICs that the servers use to communicate with the network. For more information about creating Network Policies, see [Creating Network Policies](#).
- **LDAP Policy**—Specifies the LDAP configuration settings and preferences for an endpoint. The endpoints support LDAP to store and maintain directory information in a network. The LDAP policy determines configuration settings for LDAP Servers, DNS parameters including options to obtain a domain name used for the DNS SRV request, Binding methods, Search parameters, and Group Authorization preferences. Through an LDAP policy, you can also create multiple LDAP groups and add them to the LDAP server database.
- **Local User Policy**—Automates the configuration of local user preferences. You can create one or more Local User policies which contain a list of local users that need to be configured.
- **Persistent Memory Policy**—Persistent Memory Modules (PMem Modules) are non-volatile memory modules that bring together the low latency of memory and the persistence of storage. PMem Modules provide faster access to data and retain across power cycles, based on the mode. Intersight supports the configuration of Intel® Optane™ PMem Module modules on the UCS M5 servers that are based on the Second Generation Intel® Xeon® Scalable processors. Intel® Optane™ PMem Modules can be used only with the Second-Generation Intel® Xeon® Scalable processors. The Persistent Memory Policy allows the configuration of security, Goals, and Namespaces of Persistent Memory Modules:
 - **Security**—Used to configure the secure passphrase for all the persistent memory modules.
 - **Goal**—Used to configure volatile memory and regions in all the PMem Modules connected to all the sockets of the server. Intersight supports only the creation and modification of a Goal as part of the Persistent Memory policy. Some data loss occurs when a Goal is modified during the creation or modification of a Persistent Memory Policy. For information on the data loss, see the Data Loss during Persistent Memory Policy Configuration and Deployment table in [Resources](#).
 - **Namespaces**—Used to partition a region mapped to a specific socket or a PMem Module on a socket. Intersight supports only the creation and deletion of Namespaces as part of the Persistent Memory Policy. Modifying a Namespace is not supported. Some data loss occurs when a Namespace is created or deleted during the creation of a Persistent Memory policy. For information on the data loss, see the Data Loss during Persistent Memory Policy Configuration and Deployment table in [Resources](#).

It is important to consider the memory performance guidelines and population rules of the Persistent Memory Modules before they are installed or replaced, and the policy is deployed. The population guidelines for the PMem Modules can be divided into the following categories, based on the number of CPU sockets:

- Dual CPU for UCS [C220 M6](#), [C240 M6](#), and [B200 M6](#) servers
- Dual CPU for UCS [C220 M5](#), [C240 M5](#), and [B200 M5](#) servers
- Quad CPU for UCS [C480 M5](#) and [B480 M5](#) servers
- Dual CPU for UCS [S3260 M5](#) servers

For more information about creating a Persistent Memory policy, exceptions to the policy, and other caveats regarding the policy, see Persistent Memory Policy in [Resources](#).

- **SAN Connectivity Policy**—Determines the network storage resources and the connections between the server and the SAN on the network. This policy enables you to configure vHBAs that the servers use to communicate with the Storage Area Network. You can use WWNN and WWPNN address pools, or static WWNN and WWPNN addresses to add vHBAs and to configure them. You must create the Fibre Channel Adapter, Fibre Channel QoS, and Fibre Channel Network policies as part of the SAN connectivity policy. For more information about creating Network policies, see [Creating Network Policies](#).
- **SD Card Policy**—Configures the Cisco FlexFlash and FlexUtil Secure Digital (SD) cards for the Cisco UCS C-Series Standalone M4 and M5 servers. This policy specifies details of virtual drives on the SD cards. You can configure the SD cards in the Operating System Only, Utility Only, or Operating System + Utility modes.

When two cards are present in the Cisco FlexFlash controller and Operating System is chosen in the SD card policy, the configured OS partition is mirrored. If only single card is available in the Cisco FlexFlash controller, the configured OS partition is non-RAID. The utility partitions are always set as non-RAID.



-
- Note**
1. This policy is currently not supported on Cisco UCS M6 servers.
 2. You can enable up to two utility virtual drives on Cisco UCS M5 servers, and any number of supported utility virtual drives on Cisco UCS M4 servers.
 3. Diagnostics is supported only for Cisco UCS M5 servers.
 4. User Partition drives can be renamed only on Cisco UCS M4 servers.
 5. FlexFlash configuration is not supported on Cisco UCS C460 M4 servers.
 6. For the Operating System+Utility mode, the Cisco UCS M4 servers require two FlexFlash cards, and the Cisco UCS M5 servers require at least 1 FlexFlash + 1 FlexUtil card.
-

- **SMTP Policy**—Sets the state of the SMTP client in the managed device. You can specify the preferred settings for outgoing communication and select the fault severity level to report and the mail recipients.
- **SOL Policy**—Enables the input and output of the serial port of a managed system to be redirected over IP. You can create one or more Serial over LAN policies which contain a specific grouping of Serial over LAN attributes that match the needs of a server or a set of servers.
- **SSH Policy**—Enables an SSH client to make a secure, encrypted connection. You can create one or more SSH policies that contain a specific grouping of SSH properties for a server or a set of servers.
- **Simple Network Management Protocol (SNMP) Policy**—Configures the SNMP settings for sending fault and alert information by SNMP traps from the managed devices. Any existing SNMP Users or SNMP Traps configured previously on the managed devices are removed and replaced with users or traps that you configure in this policy. If you have not added any users or traps in the policy, the existing users or traps on the server are removed but not replaced.
- **Storage Policy**—A Storage policy allows you to create drive groups, virtual drives, configure the storage capacity of a virtual drive, and configure the M.2 RAID controllers.

- **Syslog Policy**—Defines the logging level (minimum severity) to report for a log file collected from an endpoint, the target destination to store the Syslog messages, and the Hostname/IP Address, port information, and communication protocol for the Remote Logging Server(s).
- **Virtual Media Policy**—Enables you to install an Operating System on the server using the KVM console and virtual media, mount files to the host from a remote file share, and enable virtual media encryption. You can create one or more Virtual Media policies, which can contain virtual media mappings for different OS images, and configure up to two virtual media mappings, one for ISO files (through CDD), and the other for IMG files (through HDD).

For more information about the various mount options for the Virtual Media volumes, see [Virtual Media Mount options](#).

- **Virtual KVM Policy**—Enables specific grouping of virtual KVM properties. This policy allows you specify the number of allowed concurrent KVM sessions, port information, and video encryption options.
- **IMC Access Policy**—Enables to manage and configure your network through mapping of IP pools to the server profile. This policy allows you to configure a VLAN and associate it with an IP address through the IP pool address.

In-Band IP address, Out-of-Band IP address, or both In-Band and Out-of-Band IP addresses can be configured using IMC Access Policy and are supported on the following:

- Drive Security, SNMP, Syslog, and vMedia policies
- vKVM, IPMI, SOL, and vMedia policies using vKVM client

- **Power Policy**—Enables the management of power for FI-attached servers and chassis. This policy allows you to set the power profiling the power priority of the server, and the power restore state of the system. For more information, see [Creating a Power Policy for Server](#)
- **NTP Policy**—Allows you to enable the NTP service on an Intersight Managed Cisco IMC (Standalone) server. The NTP service synchronizes the time with an NTP server. You must enable and configure the NTP service by specifying the IP address or DNS of a minimum of one to a maximum of four NTP servers.

NTP policy also allows you to configure the timezone on Cisco IMC (Standalone) server. When you enable the NTP service and select Timezone, Cisco Intersight configures the NTP details and Timezone on the endpoint.

- **FC Zone Policy**—Allows you to set up access control between hosts and storage devices. You can create a Single Initiator Single Target, or Single Initiator Multiple Target Zone on a VSAN with the scope FC Storage, and attach the Zone policy to the SAN Connectivity policy using the vHBA.



Note You can configure zones only when the Fabric Interconnect is in FC switching mode

Configuration drift is not supported for the FC Zone policy

Creating a Policy

In Cisco Intersight, you can create a UCS Server or UCS Domain policy by using the policy wizard. To create and configure a new policy, do the following:

- Step 1** Log in to Cisco Intersight with your Cisco ID and select admin role.
- Step 2** From the **Service Selector** drop-down list, select **Infrastructure Service**.
- Step 3** Navigate to **Configure > Policies**, and then click **Create Policy**.
- Step 4** Select **UCS Server > <A UCS server policy>**.
- Step 5** Click **Start** to begin configuring the policy.
- Step 6** On the **General** page, enter the **Name** of the policy. Optionally, enter a **Description** and **Tags**.
- Step 7** On the **Policy Details** page, configure policy properties.

Some policy properties may be applicable to specific target platforms—Standalone UCS servers, FI-attached UCS servers, or both. You can view these properties separately for **All Platforms**, **UCS Servers (Standalone)**, and **UCS Servers (FI-Attached)** by clicking on these options. The properties that are applicable only to Standalone servers or FI-Attached servers are indicated by an icon alongside the property.

- Step 8** Click **Create**.

Supported UCS Server Policies

The following table provides a list of UCS server policies and the managed devices on which they are supported. All the server policies listed in this table are available with a Cisco Intersight Essentials license.

UCS Server Policy	Supported Servers												
	Cisco UCS C-Series									Cisco UCS B-Series		Cisco UCS X-Series	
	Standalone					IMM				IMM		IMM	
	M4	M5	M6	M7	M8	M5	M6	M7	M8	M5	M6	M6	M7
Certificate Management Policy	—	—	—	—	—	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Device Connector Policy	Yes	Yes	Yes	Yes	Yes	—	—	—	—	—	—	—	—
IPMI Over LAN Policy	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes

UCS Server Policy	Supported Servers												
	Cisco UCS C-Series									Cisco UCS B-Series		Cisco UCS X-Series	
	Standalone					IMM				IMM		IMM	
	M4	M5	M6	M7	M8	M5	M6	M7	M8	M5	M6	M6	M7
LDAP Policy	Yes	Yes	Yes	Yes	Yes	—	—	—	—	—	—	—	—
Local User Policy	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
NTP Policy	Yes	Yes	Yes	Yes	Yes	—	—	—	—	—	—	—	—
Network Connectivity Policy	Yes	Yes	Yes	Yes	Yes	—	—	—	—	—	—	—	—
Persistent Memory Policy	—	Yes	Yes	No	No	—	—	—	—	—	—	—	—
Power Policy	—	—	—	—	—	—	—	—	—	Yes	Yes	Yes	Yes
SD Card Policy	Yes	Yes	—	—	—	Yes	—	—	—	Yes	—	—	—
SMTP Policy	Yes	Yes	Yes	Yes	Yes	—	—	—	—	—	—	—	—
SNMP Policy	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
SSH Policy	Yes	Yes	Yes	Yes	Yes	—	—	—	—	—	—	—	—
Serial Over LAN (SoL) Policy	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Syslog Policy	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes

UCS Server Policy	Supported Servers												
	Cisco UCS C-Series									Cisco UCS B-Series		Cisco UCS X-Series	
	Standalone					IMM				IMM		IMM	
	M4	M5	M6	M7	M8	M5	M6	M7	M8	M5	M6	M6	M7
Virtual KVM Policy	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
BIOS Token Policy	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Virtual Media Policy	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
LAN Connectivity Policy	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
SAN Connectivity Policy	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Boot Order Policy	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Adapter Configuration Policy	Yes	Yes	Yes	Yes	Yes	—	—	—	—	—	—	—	—
Drive Security Policy	No	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Storage Policy	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
IMC Access Policy	—	—	—	—	—	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Ethernet Adapter Policy	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes

UCS Server Policy	Supported Servers												
	Cisco UCS C-Series									Cisco UCS B-Series		Cisco UCS X-Series	
	Standalone					IMM				IMM		IMM	
	M4	M5	M6	M7	M8	M5	M6	M7	M8	M5	M6	M6	M7
Ethernet Network Policy	Yes	Yes	Yes	Yes	Yes	—	—	—	—	—	—	—	—
Ethernet QoS Policy	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Ethernet Network Control Policy	—	—	—	—	—	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Ethernet Network Group Policy	—	—	—	—	—	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
FC Zone Policy	—	—	—	—	—	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Fibre Channel Adapter Policy	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Fibre Channel Network Policy	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Fibre Channel QoS Policy	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
iSCSI Boot Policy	—	—	—	—	—	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
iSCSI Adapter Policy	—	—	—	—	—	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes

UCS Server Policy	Supported Servers												
	Cisco UCS C-Series								Cisco UCS B-Series		Cisco UCS X-Series		
	Standalone					IMM				IMM		IMM	
	M4	M5	M6	M7	M8	M5	M6	M7	M8	M5	M6	M6	M7
iSCSI Static Target Policy	—	—	—	—	—	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Firmware Policy	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Thermal Policy	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	No	No	No	No

Default Server Policy Configurations

When a server is unassigned from a server profile, the following changes occur to the server configurations:

- **BIOS Policy**—All the tokens are reset to platform-default. For more information on the **platform-defaults**, see [Cisco UCS Server BIOS Token Guide](#).
- **Boot Order Policy**—Boot devices are not configured. **Boot mode** is set to **UEFI** and **Secure boot** is disabled. Server boots to **UEFI** shell on restart.
- **Certificate Management Policy**—Root CA certificates and IMC certificates are not configured.
- **Local User Policy**—No users are configured. Server Profile deployment is needed to set the Admin password.
- **IPMI over LAN Policy**—IPMI over LAN is disabled, and the **Encryption Key** is set to **0**.
- **Virtual KVM Policy**—Virtual KVM is disabled and **Remote Port** is set to **2068**.
- **Power Policy**—Power profiling is enabled, **Power Priority** is set to **Low**, and **Power Restore** is set to **Always Off**.
- **SNMP Policy**—SNMP is disabled, the **Port** is set to **161**, and all SNMP users and traps are deleted.
- **Serial Over LAN Policy**—Serial over LAN is disabled and the **Baud Rate** is set to **115200**.
- **Syslog Policy**—**Minimum Severity to Report** is set to **Debug** in **Local Logging**, and the remote logging servers are not configured.
- **Thermal Policy**—**FanControlMode** is set to **Acoustic mode**.
- **Virtual Media Policy**—Virtual Media is enabled, and the Virtual Media mounts are not configured.
- **LAN Connectivity Policy**—vNICs are not configured.
- **SAN Connectivity Policy**—vHBAs are not configured.

- **IMC Access Policy**—In-Band and Out-Of-Band settings are not configured.
- **UUID**—The UUID address is not configured.

The following server policy configurations will not be reset upon server profile actions:

- Firmware Policy
- Storage Policy
- SD Card Policy
- Drive Security Policy

Creating a Certificate Management Policy

In Intersight Managed Mode, the Certificate Management policy allows you to specify the certificate details for an external certificate and attach the policy to servers. Cisco Intersight currently supports the following certificates:

- **Root CA certificates:** A Root CA certificate is necessary for HTTPS boot authentication. You can deploy a maximum of 10 Root CA certificates using the Certificate Management Policy. For a successful boot, at least one valid and unexpired Root CA certificate is required. For more information, see [Creating a Boot Order Policy](#).



Note In Intersight Managed Mode servers, removing a server profile will delete the Root CA certificates from the CIMC.

However, for C-Series servers in Standalone mode, the Root CA certificates are not automatically removed. You must manually delete them from CIMC or perform a factory reset on the server. Additionally, when you export the profile of a C-Series server in Standalone mode, the certificate management policy will not be included.

- **IMC certificates:** This option is available only for Intersight Managed Mode servers.

1. Log in to Cisco Intersight with your Cisco ID and select admin role.
2. From the **Service Selector** drop-down list, select **Infrastructure Service**.
3. Navigate to **Configure > Policies**, and then click **Create Policy**.
4. Select **Certificate Management**, and then click **Start**.
5. On the **General** page, configure the following parameters:

Property	Essential Information
Organization	Select the Organization.
Name	Enter a name for your policy.

Property	Essential Information
Add Tag (Optional)	Enter a tag in the key:value format. For example, Org: IT or Site: APJ.

6. On the **Policy Details** page, add the certificate that you want to provide, and configure the following parameters:

Property	Essential Information
Root CA	<ul style="list-style-type: none"> • Certificate Name—Enter the name of the certificate. • Certificate—Enter the certificate details.
IMC	<ul style="list-style-type: none"> • Certificate—Enter the certificate details. • Private Key—Enter the private key details for the certificate.

7. Click **Create**.

Creating an Adapter Configuration Policy

An Adapter Configuration Policy configures the Ethernet and Fibre-Channel settings for the Virtual Interface Card (VIC) adapter.



Note This policy, if attached to a server profile that is assigned to an Intersight Managed Fabric Attached server, will be ignored.

1. Log in to Cisco Intersight with your Cisco ID and select admin role.
2. From the **Service Selector** drop-down list, select **Infrastructure Service**.
3. Navigate to **Configure > Policies**, and then click **Create Policy**.
4. Select **Adapter Configuration**, and then click **Start**.
5. On the **General** page, configure the following parameters:

Property	Essential Information
Organization	Select the Organization.
Name	Enter a name for your policy.
Description (Optional)	Provide a short description
Add Tag (Optional)	Enter a tag in the key:value format. For example, Org: IT or Site: APJ.

6. On the **Policy Details** page, click **Add VIC Adapter Configuration** and configure the following parameters:

Property	Essential Information
Add VIC Adapter Configuration	
PCI Slot	The PCI slot in which the adapter is installed. The range is from 1 to 15 and MLOM.
LLDP	The LLDP protocol status on the adapter interface. If checked, then Link Layer Discovery Protocol (LLDP) enables all the Data Center Bridging Capability Exchange protocol (DCBX) functionality, which includes FCoE, priority based flow control. Note LLDP is available only on some UCS C-Series servers. We recommend that you do not disable LLDP option, as it disables all the DCBX functionality.
FIP	The FIP protocol status on the adapter interface. If checked, then FCoE Initialization Protocol (FIP) mode is enabled. FIP mode ensures that the adapter is compatible with current FCoE standards. Note We recommend that you use FIP option only when explicitly directed to do so by a technical support representative.
Port Channel	The port channel status on the adapter interface. When Port Channel is enabled, two vNICs and two vHBAs are available for use on the adapter card. When disabled, four vNICs and four vHBAs are available for use on the adapter card. Disabling port channel reboots the server. Note Port Channel is supported only for Cisco VIC 1455/1457 adapters.

Property	Essential Information
<p>Enable Physical NIC Mode</p>	<p>When Physical NIC Mode is enabled, uplink ports of the VIC are set to pass-through mode. This allows the host to transmit packets without any modification. VIC ASIC does not rewrite the VLAN tag of the packets based on the VLAN and CoS settings for the vNIC.</p> <p>Note</p> <ul style="list-style-type: none"> • Enabling Physical NIC Mode reboots the server. • Physical NIC Mode supports UCS VIC 1400 Series and VIC 15000 Series adapters. • The minimum supported Cisco Server firmware version 4.2(2a) and later and Adapter firmware version 5.2(2a). • This feature is not supported for Cisco Intersight Managed FI Attached servers. • Only default vNICs will be added if the Physical NIC mode is enabled. • This option cannot be enabled on an adapter that has: <ul style="list-style-type: none"> • Port Channel mode enabled • VNTAG mode enabled • LLDP enabled • FIP mode enabled • Cisco IMC Management Enabled value set to Yes <p>When Physical NIC Mode is enabled, the following message is displayed in a pop-up window:</p> <p>After physical nic-mode mode switch, vNIC configurations will be lost and new default vNICs will be created.</p> <p>Click Ok.</p>

Property	Essential Information
DCE Interface	<p>The Forward Error Correction (FEC) mode setting for the DCE interfaces of the adapter.</p> <p>Note FEC mode setting is supported only for Cisco VIC 14xx adapters. FEC mode 'cl74' is unsupported for Cisco VIC 1495/1497. This setting will be ignored for unsupported adapters and for unavailable DCE interfaces</p>

7. Click **Add**.
8. Click **Create**.

Creating a LAN Connectivity Policy

A LAN Connectivity Policy determines the connections and the network communication resources between the server and the LAN on the network. You can specify MAC address pools, or static MAC addresses, to assign MAC addresses to servers and to identify the vNICs that the servers use to communicate with the network.

Prerequisites

Choose the following sub-policies or pool as per your requirement to create the LAN Connectivity policy

- **Ethernet Network Policy**—Specify if the port should carry single VLAN (Access) or multiple VLANs (Trunk) traffic. You can specify the VLAN to be associated with an Ethernet packet if no tag is found.
- **Ethernet QoS Policy**—Configure the maximum size for a Fibre Channel frame payload that the virtual interface supports, limit the data rate on the virtual interface, associate a Class of Service to the traffic on the virtual interface.
- **Ethernet Adapter Policy**—Configure features like VXLAN, NVGRE, ARFS, Interrupt settings, RoCE, and TCP Offload settings to govern the host side behavior of the adapter.
- **IQN Pool**—You can configure the Prefix and Suffix for the IQN block, the first suffix number in the block and the number of identifiers the block can hold .

1. Log in to Cisco Intersight with your Cisco ID and select admin role.
2. From the **Service Selector** drop-down list, select **Infrastructure Service**.
3. Navigate to **Configure > Policies**, and then click **Create Policy**.
4. Select **LAN Connectivity**, and then click **Start**.
5. On the **General** page, enter the following information:
 - **Name** of your policy.
 - **Target Platform** for which the policy is applicable. This can be **Standalone** servers or **FI Attached** servers.

A LAN Connectivity Policy created for Standalone servers cannot be deployed on FI Attached servers. Similarly, a LAN Connectivity Policy created for FI Attached servers cannot be deployed on Standalone servers.

- **Description** to help identify the policy.
- **Tag** for the policy. Tags must be in the key:value format. For example, Org: IT or Site: APJ.

6. On the **Policy Details** page, configure the following:

- For an FI-attached server, turn the **Enable Azure Stack Host QoS** button ON, to successfully deploy the Azure Stack QoS capability on the adapter with RDMA enabled.

Enabled—Enabling AzureStack-Host QoS on an adapter allows the user to carve out traffic classes for RDMA traffic and ensure a desired portion of the bandwidth is allocated to it.

Disabled—Disables the Azure Stack Host QoS feature on the adapter.

- Specify whether no IQN, an IQN pool, or a unique IQN identifier is to be associated with the policy by selecting **None**, **Pool**, or **Static**.
 - **None**—If you select this option, you do not have to specify any IQN details.
 - **Pool**—If you select this option, select the IQN pool that you want to associate with the LAN Connectivity policy.
 - **Static**—If you select this option, enter a static IQN for use as initiator identifiers by iSCSI vNICs in a Fabric Interconnect domain.
- Select the placement option for each vNIC—**Manual** or **Auto**
 - **Manual vNIC Placement**—If you select this option, you must manually specify the placement for each vNIC. You can also use the **Graphic vNICs Editor** to create and specify the placement for each vNIC manually by adding vNICs and slots, and defining the connection between them.



Note

- For manual placement, **PCI Link** is not supported on UCS VIC 1400 Series adapters.
- If a LAN Connectivity Policy has both Simple and Advanced placements, ensure the number provided in PCI Order is appropriate to prevent Server Profile deployment failure.
- **Auto vNIC Placement**—If you select this option, vNIC placement will be done automatically during profile deployment. This option is available only for Cisco Intersight Managed FI Attached servers.



- Note**
- Cisco UCS VIC 1300 Series adapters auto-upgrade is supported on B-Series server with Cisco Server firmware version 4.2(2e) and above.
 - Discovery of a C-Series server will not get triggered if the server with Cisco UCS VIC 1300 Series adapters has a Cisco Server firmware version lower than 4.2(2g). Upgrade the Cisco Server firmware to 4.2(2g) to enable server discovery.

7. To set up a vNIC without using a template, click **Add vNIC** and configure the following parameters:

Property	Essential Information
Add vNIC	
Ensure that you configure eth0 and eth1 interfaces for each VIC adapter you configure. You can add additional vNICs depending on your network requirements.	
Name	vNIC name.
Pin Group Name	<p>Name of the pin group that contains the specific port/port channels. All traffic from the vNIC is pinned to the specified uplink Ethernet ports or port channels.</p> <p>Note The pin group can be defined while creating a Port policy.</p> <p>If you do not assign a pin group to a vNIC, an uplink Ethernet port or port channel for traffic is chosen from that server interface dynamically. This choice is not permanent. A different uplink Ethernet port or port channel may be used for traffic from that server interface after an interface flap or a server reboot.</p>
MAC Address Pool	Click Select Pool and choose a MAC address pool for MAC address assignment.
Static	Click Static and enter a static MAC address for MAC address assignment. This option is available only for Cisco Intersight Managed FI Attached servers.
Placement	
Placement Settings for the virtual interface.	

Property	Essential Information
<p>Simple</p> <p>When you select Simple Placement, the Slot ID and PCI Link are determined automatically by the system. vNICs are deployed on the first VIC. The slot ID determines the first VIC. Slot ID numbering begins with MLOM, and thereafter it keeps incrementing by 1, starting from 1. The PCI link is always set to 0.</p>	
<p>Switch ID</p>	<p>Refers to the Fabric Interconnect that carries the vNIC traffic.</p>
<p>PCI Order</p>	<p>The order in which the virtual interface is brought up. The order assigned to an interface should be unique and in sequence starting with "0" for all the Ethernet and Fibre-Channel interfaces on each PCI link on a VIC adapter. The maximum value of PCI order is limited by the number of virtual interfaces (Ethernet and Fibre-Channel) on each PCI link on a VIC adapter.</p> <p>Note You cannot change the PCI order of two vNICs without deleting and recreating the vNICs.</p>
<p>Advanced</p>	
<p>Automatic Slot ID Assignment</p>	<p>When enabled, slot ID is determined automatically by the system.</p>
<p>Slot ID</p>	<p>When automatic slot ID assignment is disabled, the slot ID needs to be entered manually.</p> <p>Supported values are (1-15) and MLOM.</p>
<p>PCI link</p> <p>The PCI link used as transport for the virtual interface.</p> <p>PCI Link is only applicable for select Cisco UCS VIC 1300 Series models (UCSC-PCIE-C40Q-03, UCSB-MLOM-40G-03, UCSB-VIC-M83-8P) that support two PCI links. The value, if specified, for any other VIC model will be ignored.</p> <p>Note The host device order can get impacted when using both the PCI links and while adding or removing vNICs.</p>	

Property	Essential Information
Automatic PCI link Assignment	<p>When enabled, PCI link is determined automatically by the system.</p> <p>Note</p> <ul style="list-style-type: none"> • If Automatic assignment is enabled for both Slot ID and PCI link, then the behavior is same as Simple placement. All the vNICs are placed on the same PCI link (link 0). • If Automatic Slot ID assignment is disabled but automatic PCI link assignment is enabled, then you need to provide the slot ID and the vNIC will be placed on PCI link 0.
Load Balanced	<p>When Automatic PCI link assignment is disabled and Load Balanced is enabled, the system uniformly distributes the interfaces across the PCI Links.</p> <ul style="list-style-type: none"> • If automatic PCI link assignment is disabled and automatic Slot ID is enabled, you need to specify the PCI order to load balance the vNICs. • If both automatic PCI link assignment and automatic Slot ID are disabled, you need to specify the slot and the PCI order to load balance the vNICs. <p>Note You cannot change the PCI link mode of two vNICs from Load Balanced mode to Custom mode without deleting and recreating the vNICs.</p>
Custom	<ul style="list-style-type: none"> • If automatic PCI link assignment is disabled and automatic Slot ID is enabled, you need to provide the value of the PCI order, PCI link, and Switch ID. • If both automatic PCI link assignment and automatic Slot ID assignment are disabled, you need to provide the values of the Slot ID, PCI order and the PCI link. <p>Note You cannot change the PCI link mode of two vNICs from Custom mode to Load Balanced mode without deleting and recreating the vNICs.</p>

Property	Essential Information
Consistent Device Naming (CDN) Consistent Device Naming configuration for the virtual NIC.	
Source	Whether the source of the CDN name is the name of the vNIC instance or a user-defined name.
Failover Enabling failover ensures that traffic automatically fails over from one uplink to another in case of an uplink failure.	
Ethernet Network Policy	Select or create an Ethernet Network policy. Note This sub-policy is applicable only for the LAN Connectivity Policy on Standalone servers.
Ethernet Network Group Policy	Select or create an Ethernet Network Group policy. Note This sub-policy is applicable only for the LAN Connectivity Policy on FI-attached servers.
Ethernet Network Control Policy	Select or create an Ethernet Network Control policy. Note This sub-policy is applicable only for the LAN Connectivity Policy on FI-attached servers.
Ethernet QoS Policy	Select or create an Ethernet QoS policy.
Ethernet Adapter Policy	Select or create an Ethernet Adapter policy.
iSCSI Boot Policy	Select or create an iSCSI Boot policy. Note This sub-policy is applicable only for the LAN Connectivity Policy on FI-attached servers.
Connection: Disabled/usNIC/VMQ/SR-IOV	
Disabled	Does not configure a connection policy.
usNIC User Space NIC Settings that enable low-latency and higher throughput by bypassing the kernel layer when sending/receiving packets.	
Number of usNICs	Number of usNIC interfaces to be created.
usNIC Adapter Policy	Select the Ethernet Adapter policy to be associated with the usNICs.

Property	Essential Information
Class of Service	Class of service to be used for traffic on the usNIC.
VMQ Virtual Machine Queue Settings for the virtual interface that allow efficient transfer of network traffic to the guest operating system.	
Enable Multi Queue Support	Whether Virtual Machine Multi-Queue (VMMQ) is enabled in the policy. With VMMQ, multiple queues are allocated to a single VM.
Number of Sub vNICs	Number of sub vNICs that are available for Multi Queue.
Enable RoCE Settings	Whether Remote Direct Memory Access (RDMA) over Converged Ethernet (RoCE) is enabled over this virtual interface.
Memory Regions	The number of memory regions per adapter. Enter an integer between 1 and 524288. It is recommended that this number be an integer power of 2.
Queue Pairs	The number of queue pairs per adapter. Enter an integer between 1 and 8192. It is recommended that this number be an integer power of 2.
Resource Groups	The number of resource groups per adapter. Enter an integer between 1 and 128. It is recommended that this number be an integer power of 2 greater than or equal to the number of CPU cores on the system for optimum performance.
Version	Version of the RDMA protocol Version 1 is a link layer protocol. It allows communication between any two hosts in the same Ethernet broadcast domain. RoCEv2 is an internet layer protocol. RoCEv2 packets can be routed. This is possible because RoCEv2 packets now include an IP and UDP header.
SR-IOV Single Root Input/Output Virtualization (SR-IOV) allows multiple VMs running a variety of Linux guest operating systems to share a single PCIe network adapter within a host server. SR-IOV allows a VM to move data directly to and from the vNIC, bypassing the hypervisor for increased network throughput and lower server CPU overhead.	

Property	Essential Information
Number of VFs	Number of VFs to create. Enter a value between 1 and 64. Default value is 64.
Receive Queue Count Per VF	Number of Receive Queue resources to configure for each VF. Enter a value between 1 to 8. Default value is 4.
Transmit Queue Count Per VF	Number of Transmit Queue resources to configure for each VF. Enter a value between 1 to 8. Default value is 1.
Completion Queue Count Per VF	Number of Completion Queue resources to configure for each VF. Enter a value between 1 to 16. Default value is 5.
Interrupt Count Per VF	Number of Interrupt count to configure for each VF. Enter a value between 1 to 16. Default value is 8.

- To derive vNIC for FI-attached servers using a vNIC template, choose **vNIC from Template** from the **Add** drop-down list. For more information on creating vNIC templates, see *Creating vNIC or vHBA Templates*.



Note

- When deriving a vNIC from a template, the vNIC configuration is auto-populated from the template configuration. You can edit or delete parameters, which are enabled for configuration override through the vNIC template. For parameters that are not enabled for override, you can only view the configurations using the Eye icon.
- The parameters that have been overridden are indicated using an Overridden label. In the case of override-enabled parameters, the changes applied in the template are not reflected in the derived vNIC.
- Only those parameters can be modified in the derived vNIC instance which are not included in the template.
- If you attempt to derive a vNIC from a template while profile deployment is in progress, the task will be retried until the profile deployment is completed. You can find these details in the Requests tab.

- Click **Create**.

Configuration Feature Matrix for Supported Adapters in IMM

The following table shows the features supported by various adapters in Intersight Managed Mode.

Feature	Cisco UCS 1300 Series Adapter	Cisco UCS 1400/14000 Series Adapter	Cisco UCS 15000 Series Adapter
usNIC	Yes	Yes	Yes
VMQ	Yes	Yes	Yes
VMMQ	No	Yes	Yes

Feature	Cisco UCS 1300 Series Adapter	Cisco UCS 1400/14000 Series Adapter	Cisco UCS 15000 Series Adapter
SR-IOV	No	Yes	Yes
NetQueue	Yes	Yes	Yes
RoCEv1	Yes	No	No
RoCEv2	No	Yes	Yes
Geneve Offload	No	Yes	Yes
AzureQoS	No	Yes	Yes
RSS	Yes	Yes	Yes
RSSv2	No	No	Yes
NVGRE	Yes	Yes	Yes
ARFS	Yes	Yes	Yes
VIC QinQ Tunneling	No	Yes	Yes
VXLAN	Yes	Yes	Yes
Advance Filter	Yes	Yes	Yes
Interrupt Scaling/Group Interrupt	Yes	Yes	Yes
Host Port Configuration	Yes	No	No
vHBA Type	Yes	Yes	Yes
16K Ring Size	No	No	Yes
Precision Time Protocol	No	No	Yes
FC MQ	Yes	Yes	Yes
FC NVMe	Yes	Yes	Yes
ENS	No	Yes	Yes

Creating an Ethernet Adapter Policy

An Ethernet adapter policy governs the host-side behavior of the adapter, including how the adapter handles traffic. For each VIC Virtual Ethernet Interface, you can configure various features like Virtual Extensible LAN (VXLAN), Network Virtualization using Generic Routing Encapsulation (NVGRE), Accelerated Receive Flow Steering (ARFS), Interrupt settings, and TCP Offload settings.

The Ethernet Adapter policy include the recommended settings for the virtual Ethernet interface, for each supported server operating system. Operating systems are sensitive to the settings in these policies. In general, the storage vendors require non-default adapter settings. You can find the details of these required settings on the support list provided by those vendors.

GENEVE Offload

Cisco Intersight supports Generic Network Virtualization Encapsulation (GENEVE) Offload on the ESXi platform, which allows essentially any information to be encoded in a packet and passed between tunnel endpoints. GENEVE provides the overlay capability to create isolated, multi-tenant broadcast domains across data center fabrics on 1400 Series adapters. Using the GENEVE protocol allows you to create logical networks that span physical network boundaries.

GENEVE offload is present in all Ethernet adapter policies and is disabled by default. It is the recommended setting if using VMWare ESXi GENEVE.

For more information on how to implement GENEVE offload end-to-end configuration, see [Cisco UCS Manager Network Management](#) documentation.

Cisco recommends configuring the following values in the Ethernet adapter policy when GENEVE offload is enabled:

- Transmit Queues : 1
- TX Ring Size: 4096
- Receive Queues: 8
- RX Ring Size: 4096
- Completion Queues : 16
- Interrupts : 32

The following features are not supported when GENEVE offload is enabled on any interface:

- Azure Stack QoS
- RoCEv2 - you cannot have GENEVE enabled on one vNIC and RoCEv2 enabled on another.
- Advanced Filters
- VIC QinQ Tunneling

Support for usNIC and VIC QinQ Tunneling features on interfaces:



Note

- usNIC or VMQ is not compatible with GENEVE Offload on the same interface only for 1400 Series adapters.
 - usNIC or VMQ is compatible with GENEVE Offload on different interfaces for 1400 Series adapters.
 - usNIC and VMQ is compatible with GENEVE Offload on both the same and different interfaces for 1500 Series adapters.
-



Note On switching from GENEVE offload feature to Azure Stack QoS feature or vice versa, please do the following:

1. Disable the current feature
2. Reboot the server
3. Enable the required feature

Other limitations with GENEVE offload include:

- External outer IPV6 is NOT supported with GENEVE offload.
- GENEVE offload is supported with ESX 7.0 (NSX-T 3.0) and ESX 6.7U3(NSX-T 2.5).
- GENEVE offload is supported only with Cisco UCS VIC 1400/14000 and 15000 Series adapters. It is not supported on Cisco UCS VIC 1300 Series adapters or Cisco UCS VIC 1200 Series adapters.
- Cisco UCS VIC 1400/14000 and 15000 Series adapters.
- Minimum server firmware version for UCS C-Series Standalone: 4.1(2a)
- Minimum adapter firmware version: 5.1(2f)
- Cisco recommends that you remove the GENEVE offload configuration before downgrading to any non-supported release.

For details on supported features matrix with GENEVE offload, refer the table below.

Table 4: GENEVE Offload Supported Features Matrix for 1400 Series Adapters

	KVM VM - FEX	VXLAN	NVGRE	RoCEv2	usNIC	Netflow	Advanced Filters	VMQ/ VMMQ/ netqueue	arfs	Azure QoS
GENEVE offload enabled on the interface vnic1 and feature is enabled on vnic1	No	Yes	Yes	No	No	No	No	No	No	No
GENEVE offload that is enabled on the interface vnic1 and feature is enabled on vnic2	Yes	Yes	Yes	No	Yes	Yes	Yes	Yes	Yes	No

Table 5: GENEVE Offload Supported Features Matrix for 15000 Series Adapters

	VXLAN	NVGRE	RoCEv2	usNIC	Netflow	Advanced Filters	VMO/ VMMQ/ netqueue	arfs	quad port per adapter	physical nic node per adapter
GENEVE offload enabled on the same interface (vnic1) and feature is enabled on vnic1	Yes	Yes	No	Yes	Yes	Yes	Yes	No	Yes	Yes
GENEVE offload enabled on different interface (vnic1) and feature is enabled on vnic2	Yes	Yes	No	Yes	Yes	Yes	Yes	Yes	Yes	Yes

1. Log in to Cisco Intersight with your Cisco ID and select admin role.
2. From the **Service Selector** drop-down list, select **Infrastructure Service**.
3. Navigate to **Configure > Policies**, and then click **Create Policy**.
4. Select **Ethernet Adapter**, and then click **Start**.
5. In the **General** page, configure the following parameters:

Property	Essential Information
Organization	Select the organization.
Name	Enter a name for your policy.
Set Tags	Enter a tag in the key value format. For example, Org: IT or Site: APJ.
Description (optional)	Enter a short description.
Cisco Provided Ethernet Adapter Configuration	
Select Cisco Provided Configuration (Optional)	Click Select Cisco Provided Configuration , search, and choose from the available pre-defined configurations. Note After you choose a configuration, the policy is updated with the pre-defined values from the chosen configuration. You can modify the values in the Details page or skip Step 6 and proceed to create the policy using these pre-defined values.

6. On the **Policy Details** page, configure the following parameters:

Property	Essential Information
Enable Virtual Extensible LAN	Enables the Virtual Extensible LAN protocol on the virtual Ethernet interface.
Enable Network Virtualization using Generic Routing Encapsulation	Enables Network Virtualization using Generic Routing Encapsulation on the virtual Ethernet interface. Note The Transmit checksum offload and TSO must be enabled for the NVGRE offloading to be effective.
Enable Accelerated Receive Flow Steering	Enables Accelerated Receive Flow Steering (ARFS) on the virtual Ethernet interface. ARFS is hardware-assisted receive flow steering that can increase CPU data cache hit rate by steering kernel level processing of packets to the CPU where the application thread consuming the packet is running.
Enable Advanced Filter	Enables advanced filtering on the virtual Ethernet interface.
Enable Interrupt Scaling	Enables Interrupt Scaling of resources on the virtual Ethernet interface.
Enable Geneve Offload	Enables GENEVE overlay hardware offloads.
RoCE Settings	
Intersight supports RDMA over Converged Ethernet (RoCE) for Microsoft SMB Direct. It sends additional configuration information to the adapter while creating or modifying an Ethernet adapter policy.	
Enable RDMA over converged Ethernet	Enables RDMA over Converged Ethernet (RoCE) on the virtual Ethernet interface. RoCE allows direct memory access over an Ethernet network. RoCE is a link layer protocol, and hence, it allows communication between any two hosts in the same Ethernet broadcast domain. RoCE delivers superior performance compared to traditional network socket implementations because of lower latency, lower CPU utilization, and higher utilization of network bandwidth.

Property	Essential Information
Queue Pairs	<p>The number of queue pairs per adapter.</p> <p>Enter an integer between 0 and 8192. It is recommended that this number be an integer power of 2.</p> <p>Note This property is displayed only when Enable RDMA over converged Ethernet is enabled.</p>
Memory Regions	<p>The number of memory regions per adapter.</p> <p>Enter an integer between 0 and 524288. It is recommended that this number be an integer power of 2.</p> <p>Note This property is displayed only when Enable RDMA over converged Ethernet is enabled.</p>
Resource Groups	<p>The number of resource groups per adapter. It is recommended that this number be an integer power of 2 greater than or equal to the number of CPU cores on the system for optimum performance.</p> <p>Enter an integer between 0 and 128.</p> <p>Note This property is displayed only when Enable RDMA over converged Ethernet is enabled.</p>
Version	<p>Version of the RDMA protocol</p> <p>Version 1 is a link layer protocol. It allows communication between any two hosts in the same Ethernet broadcast domain.</p> <p>Note This property is displayed only when Enable RDMA over converged Ethernet is enabled.</p>
Interrupt Settings	
Interrupts	<p>Enter the number of interrupt resources to allocate. Typically this value is equal to the number of completion queue resources.</p> <p>Enter an integer between 1 and 1024.</p>

Property	Essential Information
Interrupt Mode	Select the preferred driver interrupt mode that include: <ul style="list-style-type: none"> • MSIx—Message Signaled Interrupts (MSI) with the optional extension. This is the recommended option. • MSI—Message Signaled Interrupts (MSI) only • INTx—PCI INTx interrupts
Interrupt Timer, us	The time to wait between interrupts or the idle period that must be encountered before an interrupt is sent. To turn off interrupt coalescing, enter 0 (zero) in this field. Enter an integer between 0 and 65535.
Interrupt Coalescing Type	Select the Interrupt Coalescing Type: <ul style="list-style-type: none"> • Min - The system waits for the time specified in the Coalescing Time field before sending another interrupt event. • Idle - The system does not send an interrupt until there is a period of no activity lasting as least the time specified in the Coalescing Time field.
Receive Receive Queue resource settings.	
Receive Queue Count	The number of queue resources to allocate. Enter an integer between 1 and 1000.
Receive Ring Size	The number of descriptors in each queue. Enter an integer between 64 and 4096.
Transmit Transmit Queue resource settings	
Transmit Queue Count	The number of queue resources to allocate. Enter an integer between 1 and 1000.
Transmit Ring Size	The number of descriptors in each queue. Enter an integer between 64 and 4096.
Completion Completion Queue resources settings	

Property	Essential Information
Completion Queue Count	<p>The number of completion queue resources to allocate. In general, the number of completion queue resources to allocate is equal to the number of transmit queue resources plus the number of receive queue resources.</p> <p>Enter an integer between 1 and 2000.</p>
Completion Ring Size	<p>The number of descriptors in each queue.</p> <p>Enter an integer between 1 and 256.</p> <p>Note This property is displayed only when Enable RDMA over converged Ethernet is enabled.</p>
Uplink Failback Timeout (seconds)	<p>Uplink Failback Timeout in seconds when uplink failover is enabled for a vNIC. After a vNIC has started using its secondary interface, this setting controls how long the primary interface must be available before the system resumes using the primary interface for the vNIC.</p> <p>Enter an integer between 0 and 600.</p>
<p>TCP Offload</p> <p>The TCP offload settings decide whether to offload the TCP related network functions from the CPU to the network hardware or not. These options help reduce the CPU overhead and increase the network throughput.</p>	
Enable Tx Checksum Offload	Enables the CPU to send all packets to the hardware so that the checksum can be calculated.
Enable Rx Checksum Offload	Enables the CPU to send all packet checksums to the hardware for validation.
Enable Large Send Offload	Enables the CPU to send large packets to the hardware for segmentation.
Enable Large Receive Offload	Enables the CPU to reassemble the segmented packets in hardware before sending them to the CPU.

Property	Essential Information
<p>Receive Side Scaling: Receive Side Scaling (RSS)/Receive Side Scaling Version 2 (RSSv2) supports multiple cores to process the incoming data traffic.</p> <p>RSSv2 is supported on Windows 2019 OS and later versions and it requires Windows NENIC driver. With RSS enabled Windows NENIC driver and Cisco UCS VIC adapter, you can configure multiple hardware receive queues on the Physical Function(PF). With VMMQ enabled on the VIC, you can configure multiple hardware receive queues per Virtual Machine(VM).</p> <p>Before using the RSSv2 functionality, ensure the NENIC driver supports RSSv2. In general, a NENIC driver supports 4 queues. With RSSv2, the NENIC driver has no upper limit on the number of hardware queues for PF or VM.</p>	
<p>Enable Receive Side Scaling</p>	<p>Enables receive side scaling and allows the incoming traffic to be spread across multiple CPU cores. This property supports both RSS and RSSv2.</p> <p>By default, RSS is enabled. RSSv2 is compatible with RSS. Based on the NENIC driver support on RSS or RSSv2, this property is supported accordingly.</p> <p>Note RSSv2 is supported on the following:</p> <ul style="list-style-type: none"> • Cisco UCS VIC 15000 Series adapters • Cisco UCS M6, M7, and M8 servers
<p>Enable IPv4 Hash</p>	<p>Enables the IPv4 address for traffic distribution.</p>
<p>Enable IPv6 Extension Hash</p>	<p>Enables the IPv6 extensions for traffic distribution.</p>
<p>Enable IPv6 Hash</p>	<p>Enables the IPv6 address for traffic distribution.</p>
<p>Enable TCP and IPv4 Hash</p>	<p>Enables both the IPv4 address and TCP port number for traffic distribution.</p>
<p>Enable TCP and IPv6 Extensions Hash</p>	<p>Enables both the IPv6 extensions and TCP port number for traffic distribution.</p>
<p>Enable TCP and IPv6 Hash</p>	<p>Enables both the IPv6 address and TCP port number for traffic distribution.</p>
<p>Enable UDP and IPv4 Hash</p>	<p>Enables both the IPv4 address and UDP port number for traffic distribution.</p>
<p>Enable UDP and IPv6 Hash</p>	<p>Enables both the IPv6 address and UDP port number for traffic distribution.</p>

7. Click **Create**.

Creating an Ethernet QoS Policy

An Ethernet Quality of Service (QoS) policy assigns a system class to the outgoing traffic for a vNIC. This system class determines the quality of service for the outgoing traffic. For certain adapters, you can also specify additional controls like burst and rate on the outgoing traffic.

1. Log in to Cisco Intersight with your Cisco ID and select admin role.
2. From the **Service Selector** drop-down list, select **Infrastructure Service**.
3. Navigate to **Configure > Policies**, and then click **Create Policy**.
4. Select **Ethernet QoS**, and then click **Start**.
5. In the **General** page, configure the following parameters:

Property	Essential Information
Organization	Select the organization.
Name	Enter a name for your policy.
Tag (optional)	Enter a tag in the key value format. For example, Org: IT or Site: APJ.
Description (optional)	Enter a short description.

6. On the **Policy Details** page, configure the following parameters:

Property	Essential Information
MTU, Bytes	The Maximum Transmission Unit (MTU) or packet size that the virtual interface accepts. The valid range is between 1500 and 9000. The default value is 1500.
Rate Limit, Mbps	The value in Mbps (0-100000) to use for limiting the data rate on the virtual interface. Setting this to zero will turn rate limiting off.
Class of Service	The Class of Service to be associated to the traffic on the virtual interface. The valid range is between 0 and 6. The default value is 3. Note This property is supported only on Standalone servers.

Property	Essential Information
Burst	<p>The burst traffic allowed on the vNIC in bytes.</p> <p>The valid range is between 1024 and 1000000. The default value is 1024.</p> <p>Note This property is supported only on FI-attached servers.</p>
Priority	<p>Select the priority matching the System QoS defined in the domain profile that include:</p> <ul style="list-style-type: none"> • Best-effort • Fibre Channel (FC) • Platinum • Gold • Silver • Bronze <p>Note</p> <ul style="list-style-type: none"> • The Best-effort system class is enabled by default. • This property is supported only on FI-attached servers.
Enable Trust Host CoS	<p>Select to enable the usage of the Class of Service to be associated to the traffic on the virtual interface.</p>

7. Click **Create**.

Creating an Ethernet Network Policy

An Ethernet Network policy sets the rules for the port to handle network traffic. This policy determines whether the port can carry single VLAN (Access) or multiple VLANs (Trunk) traffic.

This policy also supports VIC QinQ Tunneling. A QinQ (802.1Qin802.1Q) tunnel allows segregation and isolation of different VLANs within a network. To configure QinQ VLAN, you can specify the desired VLAN ID as part of the VLAN settings for the specific port, port channel, or vNIC. This enables the transmission of multiple VLANs over a single VLAN trunk.



Important This policy is supported only on C-Series Standalone servers.

An Ethernet Network policy determines if the port can carry single VLAN (Access) or multiple VLANs (Trunk) traffic. You can specify the VLAN to be associated with an Ethernet packet if no tag is found.

1. Log in to Cisco Intersight with your Cisco ID and select admin role.
2. From the **Service Selector** drop-down list, select **Infrastructure Service**.
3. Navigate to **Configure > Policies**, and then click **Create Policy**.
4. Select **Ethernet Network**, and then click **Start**.
5. On the **General** page, configure the following parameters:

Property	Essential Information
Organization	Select the Organization.
Name	Enter a name for your policy.
Description (Optional)	Provide a short description
Add Tag (Optional)	Enter a tag in the key:value format. For example, Org: IT or Site: APJ.

6. On the **Policy Details** page, configure the following parameters:

Property	Essential Information
VLAN Mode	

Property	Essential Information
	<p>Assign traffic flow to the VLAN to determine if the port can carry single VLAN (Access) or multiple VLANs (Trunk) traffic.</p> <ul style="list-style-type: none"> <p>Access Mode—Traffic is received and sent in native formats with no VLAN tagging. Anything arriving on an access port is assumed to belong to the VLAN assigned to the port.</p> <p>You can configure a port in access mode and specify the VLAN to carry the traffic for that interface. If you do not configure the VLAN for a port in access mode, or an access port, the interface carries the traffic for the default VLAN, which is VLAN 1. You can change the access port membership in a VLAN by configuring the VLAN. You must create the VLAN before you can assign it as an access VLAN for an access port. If you change the access VLAN on an access port to a VLAN that is not yet created, the UCS Manager shuts down that access port.</p> <p>If an access port receives a packet with an 802.1Q tag in the header other than the access VLAN value, that port drops the packet without learning its MAC source address. If you assign an access VLAN that is also a primary VLAN for a private VLAN, all access ports with that access VLAN receives all the broadcast traffic for the primary VLAN in the private VLAN mode.</p> <p>Trunk Mode—Trunk ports allow multiple VLANs to transport between switches over that trunk link. A trunk port can carry untagged packets simultaneously with the 802.1Q tagged packets. When you assign a default port VLAN ID to the trunk port, all untagged traffic travels on the default port VLAN ID for the trunk port, and all untagged traffic is assumed to belong to this VLAN. This VLAN is referred to as the native VLAN ID for a trunk port. The native VLAN ID is the VLAN that carries untagged traffic on trunk ports.</p> <p>The trunk port sends an egressing packet with a VLAN that is equal to the default port VLAN ID as untagged; all the other egressing packets are tagged by the trunk port. If you do not configure a native VLAN ID, the trunk port uses the default VLAN.</p>

Property	Essential Information
	This property is applicable only to Standalone servers, and not to FI Attached servers. For FI Attached mode, VLAN Mode is configured as Trunk .
Access Mode	
Enable QinQ Tunneling	Slide to enable VIC QinQ (802.1Qin802.1Q) Tunneling.
Default VLAN	Refers to the VLAN ID assigned to the traffic on the virtual interface by default. The range for the Default VLAN ID is from 0 to 4094.
QinQ VLAN	<p>This property enables the configuration of QinQ Tunneling, that facilitates the encapsulation of multiple VLANs within a single VLAN. This supported VLAN ID range is from 2 to 4093, allowing you to effectively manage and segregate the network traffic.</p> <p>Note This property is displayed only when <i>Enable QinQ Tunneling</i> slider is enabled.</p>
Trunk Mode	
Enable QinQ Tunneling	Slide to enable VIC QinQ (802.1Qin802.1Q) Tunneling.
Default VLAN	Refers to the VLAN ID assigned to the traffic on the virtual interface by default. The range for the Default VLAN ID is from 0 to 4094.
QinQ VLAN	<p>This property enables the configuration of QinQ Tunneling, that facilitates the encapsulation of multiple VLANs within a single VLAN. This supported VLAN ID range is from 2 to 4093, allowing you to effectively manage and segregate the network traffic.</p> <p>Note This property is displayed only when <i>Enable QinQ Tunneling</i> slider is enabled.</p>

7. Click **Create**.

Creating an Ethernet Network Group Policy

An Ethernet Network Group policy enables you to manage settings for VLANs on a UCS Server. These settings include defining which VLANs are allowed, designating a Native VLAN, and specifying a QinQ VLAN.



Note When an Ethernet Network Group is assigned to a Port Policy, the specified VLAN set must be either identical to or disjoint from the VLAN sets specified on other uplink interfaces. Ensure that the VLANs are defined in the VLAN Policy and that **Auto Allow on Uplinks** is disabled.

This policy also supports VIC QinQ Tunneling. A QinQ (802.1Qin802.1Q) tunnel allows segregation and isolation of different VLANs within a network. To configure QinQ VLAN, you can specify the desired VLAN ID as part of the VLAN settings for the specific port, port channel, or vNIC. This enables the transmission of multiple VLANs over a single VLAN trunk.

1. Log in to Cisco Intersight with your Cisco ID and select admin role.
2. From the **Service Selector** drop-down list, select **Infrastructure Service**.
3. Navigate to **Configure > Policies**, and then click **Create Policy**.
4. Select **Ethernet Network Group**, and then click **Start**.
5. On the **General** page, configure the following parameters:

Property	Essential Information
Organization	Select the Organization.
Name	Enter a name for your policy.
Set Tag (Optional)	Enter a tag in the key:value format. For example, Org: IT or Site: APJ.
Description (Optional)	Provide a short description

6. On the **Policy Details** page, configure the following parameters:

Property	Essential Information
VLAN Settings	

Property	Essential Information
Native VLANs	<p>This property allows you to specify the native VLAN ID for the virtual interface or its corresponding vethernet in a range of 1-4093.</p> <ul style="list-style-type: none"> • If the native VLAN is not already part of the allowed VLANs, it will be automatically added to the list of allowed VLANs. • If QinQ Tunneling is enabled, the native VLAN and Allowed VLAN properties are combined.
Enable QinQ Tunneling	<p>Slide to enable VIC QinQ (802.1Qin802.1Q) Tunneling.</p>
Allowed VLANs	<p>Refers to the VLANs that are permitted for the virtual interface. You can specify the allowed VLANs by providing a list of comma-separated VLAN IDs and VLAN ID ranges.</p> <p>For example, you can enter VLAN IDs 10, 20, 30-40 to allow VLANs 10, 20, and a range from 30 to 40.</p> <p>Note This property is displayed only when <i>Enable QinQ Tunneling</i> slider is disabled.</p>
QinQ VLAN	<p>This property enables the configuration of QinQ Tunneling, that facilitates the encapsulation of multiple VLANs within a single VLAN. The supported VLAN IDs range from 2 to 4093 that allows you to effectively manage and segregate the network traffic.</p> <p>Note This property is available only when <i>Enable QinQ Tunneling</i> slider is enabled.</p>



Note To make the server an Isolated host or a Community host, specify the ID of an Isolated VLAN or a Community VLAN in both Allowed VLANs and Native VLAN

7. Click **Create**.

Creating an Ethernet Network Control Policy

Ethernet Network Control policies configure the network control settings for the UCS Domain. This policy is applicable only for the Appliance Ports defined in a Port Policy and for the vNICs defined in a LAN Connectivity Policy, on an FI-Attached UCS Servers.

1. Log in to Cisco Intersight with your Cisco ID and select admin role.
2. From the **Service Selector** drop-down list, select **Infrastructure Service**.
3. Navigate to **Configure > Policies**, and then click **Create Policy**.
4. Select **Ethernet Network Control**, and then click **Start**.
5. On the **General** page, configure the following parameters:

Property	Essential Information
Organization	Select the Organization.
Name	Enter a name for your policy.
Description (Optional)	Provide a short description
Add Tag (Optional)	Enter a tag in the key:value format. For example, Org: IT or Site: APJ.

6. On the **Policy Details** page, configure the following parameters:

Property	Essential Information
Enable CDP	Enables the Cisco Discovery Protocol (CDP) on an interface.
MAC Register Mode	<p>Determines the MAC addresses to be registered with the switch. This can be:</p> <ul style="list-style-type: none"> • Only Native VLAN—MAC addresses are only added to the native VLAN. This option is the default, and it maximizes the port+VLAN count. • All Host VLANs—MAC addresses are added to all VLANs with which they are associated. Select this option if your VLANs are configured to use trunking but are not running in Promiscuous mode.
Action on Uplink Fail	<p>Determines how the interface behaves if no uplink port is available when the switch is in end-host mode.</p> <ul style="list-style-type: none"> • Link Down—Changes the operational state of a vNIC to down when uplink connectivity is lost on the switch, and enables fabric failover for vNICs. This is the default option. • Warning—Maintains server-to-server connectivity even when no uplink port is available, and disables fabric failover when uplink connectivity is lost on the switch.

Property	Essential Information
MAC Security Forge	Determines whether forged MAC addresses are allowed or denied when packets are sent from the server to the switch. This can be: <ul style="list-style-type: none"> • Allow— All server packets are accepted by the switch, regardless of the MAC address associated with the packets. This is the default option. • Deny— After the first packet has been sent to the switch, all other packets must use the same MAC address or they will be silently rejected by the switch. In effect, this option enables port security for the associated vNIC.
LLDP	Determines whether interfaces can transmit or receive LLDP packets. <ul style="list-style-type: none"> • To enable or disable the transmission of LLDP packets on an interface, click Enable Transmit. • To enable or disable the receipt of LLDP packets on an interface, click Enable Receive.

7. Click **Create**.

Creating a SAN Connectivity Policy

A Storage Area Network (SAN) connectivity policy determines the network storage resources and the connections between the server and the storage device on the network. This policy enables you to specify WWPN address pools, or a static WWPN address to add a vHBA. Similarly, you can specify a WWNN pool, or a static WWNN address to configure vHBAs that the servers use to communicate with the SAN.

Prerequisites

The following sub-policies are required to create the SAN Connectivity policy:

- **Fibre Channel Network Policy**—Configure the VSAN ID on the virtual interfaces.
- **Fibre Channel QoS Policy**—Limit the data rate on the virtual interface, configure the maximum size for a Fibre Channel frame payload bytes that the virtual interface supports, associate a Class of Service to the traffic on the virtual interface.
- **Fibre Channel Adapter Policy**—Govern the host side behavior of the adapter. You can enable FCP Error Recovery, change the default settings of Queues, and change Interrupt handling for performance enhancement.
- **Fibre Channel Zone Policy**—Specify direct access storage path configurations in the FC Zone policy, to set up access control between hosts and storage devices. You can create a Single Initiator Single Target, or Single Initiator Multiple Target zone on a VSAN with FC Storage scope.

- **WWNN Pool**—A World Wide Name (WWN) pool that contains only WW node names for use by the Fibre Channel vHBAs in a Cisco UCS Domain. You can also assign a static WWNN to a Fibre Channel vHBA in a Cisco UCS Domain.
- **WWPN Pool**—A World Wide Name (WWN) pool that contains only WW port names for use by the Fibre Channel vHBAs in a Cisco UCS Domain. You can also assign a static WWPN to a Fibre Channel vHBA in a Cisco UCS Domain.

1. Log in to Cisco Intersight with your Cisco ID and select admin role.
2. From the **Service Selector** drop-down list, select **Infrastructure Service**.
3. Navigate to **Configure > Policies**, and then click **Create Policy**.
4. Select **SAN Connectivity**, and then click **Start**.
5. On the **General** page, enter the following information:
 - **Name** of your policy.
 - **Target Platform** for which the policy is applicable. This can be **Standalone** servers or **FI Attached** servers.
A SAN Connectivity Policy created for Standalone servers cannot be deployed on FI Attached servers. Similarly, a SAN Connectivity Policy created for FI Attached servers cannot be deployed on Standalone servers.
 - **Description** to help identify the policy.
 - **Tag** for the policy. Tags must be in the key:value format. For example, Org: IT or Site: APJ.
6. On the **Policy Details** page, configure the following:
 - Select the placement option—**Manual** or **Auto**
 - **Manual vHBAs Placement**—If you select this option, you must manually specify the PCI slot and PCI order for each vHBA. You can also use the **Graphic vHBAs Editor** to create and specify the placement for each vHBA manually by adding vHBAs and slots, and defining the connection between them.



Note

- For manual placement, **PCI Link** is not supported on UCS VIC 1400 Series adapters
 - If a SAN Connectivity Policy has both Simple and Advanced placements, ensure the number provided in PCI Order is appropriate to prevent Server Profile deployment failure.
-
- **Auto vHBAs Placement**—If you select this option, vHBA placement will be done automatically during profile deployment. This option is available only for Cisco Intersight Managed FI Attached servers.
- Create or select a **WWNN Address Pool**, or select **Static** and enter a WWNN address. The Static option is available only for Cisco Intersight Managed FI Attached servers.

7. To set up a vHBA without using a template, click **Add vHBA** and configure the following parameters:

Property	Essential Information
Add vHBA	
Name	Name of the virtual Fibre Channel interface.
vHBA Type	<p>Type of vHBA configuration for SAN Connectivity Policy.</p> <ul style="list-style-type: none"> • fc-initiator—The type of Fibre Channel zoning to be configured for the vHBA is of the initiator type. • fc-target—The type of Fibre Channel zoning to be configured for the vHBA is of the target type. • fc-nvme-initiator—The vHBA type is initiator and applies the NVMe interface to Fibre Channel. • fc-nvme-target—The vHBA type is target and applies the NVMe interface to Fibre Channel. <p>The NVM Express (NVMe) interface allows host software to communicate with a non-volatile memory subsystem. It is optimized for Enterprise non-volatile storage, which is typically attached as a register level interface to the PCI Express (PCIe) interface.</p> <p>Note</p> <ul style="list-style-type: none"> • This configuration is supported only on Cisco VIC 1400 series and higher series of adapters. • 1300 series adapters support only fc-initiator, and fc-nvme-initiator. • Prior to connection, association with adapter should be fine. • After connection with adapter, check vhma_type in the vnic.cfg file. <p>For fc-nvme-initiator type, vhma_type should read the name.</p> <p>For fc-initiator type, vhma_type should not be present.</p>

Property	Essential Information
Pin Group Name	<p>Name of the pin group that contains the specific port/port channels. All traffic from the vHBA is pinned to the specified FC/FCoE uplink ports or port channels.</p> <p>Note The pin group can be defined while creating a Port policy.</p> <p>If you do not assign a pin group to a vHBA, an uplink FC/FCoE uplink port or port channel for traffic is chosen from that server interface dynamically. This choice is not permanent. A different FC/FCoE uplink port or port channel may be used for traffic from that server interface after an interface flap or a server reboot.</p>
WWPN Address Pool	<p>Click Select Pool and choose a WWPN address pool.</p>
Static	<p>Click Static and enter a static WWPN address. This option is available only for Cisco Intersight Managed FI Attached servers.</p>
Placement Placement Settings for the virtual interface.	
Simple <p>When you select Simple Placement, the Slot ID and PCI Link are determined automatically by the system. vHBAs are deployed on the first VIC. The slot ID determines the first VIC. Slot ID numbering begins with MLOM, and thereafter it keeps incrementing by 1, starting from 1. The PCI link is always set to 0.</p>	
Switch ID	<p>Refers to the Fabric Interconnect that carries the vHBA traffic.</p>
PCI Order	<p>The order in which the virtual interface is brought up. The order assigned to an interface should be unique and in sequence starting with "0" for all the Ethernet and Fibre-Channel interfaces on each PCI link on a VIC adapter. The maximum value of PCI order is limited by the number of virtual interfaces (Ethernet and Fibre-Channel) on each PCI link on a VIC adapter.</p> <p>Note You cannot change the PCI order of two vHBAs without deleting and recreating the vHBAs.</p>
Advanced	

Property	Essential Information
Automatic Slot ID Assignment	When enabled, slot ID is determined automatically by the system.
Slot ID	When automatic slot ID assignment is disabled, the slot ID needs to be entered manually. Supported values are (1-15) and MLOM.
<p>PCI link</p> <p>The PCI link used as transport for the virtual interface.</p> <p>PCI Link is only applicable for select Cisco UCS VIC 1300 Series models (UCSC-PCIE-C40Q-03, UCSB-MLOM-40G-03, UCSB-VIC-M83-8P) that support two PCI links. The value, if specified, for any other VIC model will be ignored.</p> <p>Note The host device order can get impacted when using both the PCI links.</p>	
Automatic PCI link Assignment	<p>When enabled, PCI link is determined automatically by the system.</p> <p>Note</p> <ul style="list-style-type: none"> • If Automatic assignment is enabled for both Slot ID and PCI link, then the behavior is same as Simple placement. All the vHBAs are placed on the same PCI link (link 0). • If Automatic Slot ID assignment is disabled but automatic PCI link assignment is enabled, then you need to provide the slot ID and the vHBA will be placed on PCI link 0.
Load Balanced	<p>When Automatic PCI link assignment is disabled and Load Balanced is enabled, the system uniformly distributes the interfaces across the PCI Links.</p> <ul style="list-style-type: none"> • If automatic PCI link assignment is disabled and automatic Slot ID is enabled, you can specify the PCI order to load balance the vHBAs. • If both automatic PCI link assignment and automatic Slot ID are disabled, you can specify the slot and the PCI order to load balance the vHBAs. <p>Note You cannot change the PCI link mode of two vHBAs from Load Balanced mode to Custom mode without deleting and recreating the vHBAs.</p>

Property	Essential Information
Custom	<ul style="list-style-type: none"> • If automatic PCI link assignment is disabled and automatic Slot ID is enabled, you need to provide the value of the PCI order, PCI link, and Switch ID. • If both automatic PCI link assignment and automatic Slot ID assignment are disabled, you need to provide the values of the Slot ID, PCI order, and the PCI link. <p>Note You cannot change the PCI link mode of two vHBAs from Custom mode to Load Balanced mode without deleting and recreating the vHBAs.</p>
Persistent LUN Bindings	
Enable Persistent LUN Bindings	Enables retention of LUN ID associations in memory until they are manually cleared.
Fibre Channel Network	Select or create a Fibre Channel Network policy.
Fibre Channel QoS	Select or create a Fibre Channel QoS policy.
Fibre Channel Adapter	Select or create a Fibre Channel Adapter policy.
FC Zone	Select or create the FC Zone policy to be attached.

- To derive vHBA for FI-attached servers using a vHBA template, choose **vHBA from Template** from the **Add** drop-down list. For more information on creating vHBA templates, see *Creating vNIC or vHBA Templates*.

**Note**

- When deriving a vHBA from a template, the vHBA configuration is auto-populated from the template configuration. You can edit or delete parameters, which are enabled for configuration override through the vHBA template. For parameters that are not enabled for override, you can only view the configurations using the **Eye** icon.
- The parameters that have been overridden are indicated using an **Overridden** label. In the case of override-enabled parameters, the changes applied in the template are not reflected in the derived vHBA.
- Only those parameters can be modified in the derived vHBA instance which are not included in the template.
- If you attempt to derive a vHBA from a template while profile deployment is in progress, the task will be retried until the profile deployment is completed. You can find these details in the **Requests** tab.

- Click **Create**.

Creating a Fibre Channel Adapter Policy

A Fibre Channel adapter policy governs the host-side behavior of the adapter, including how the adapter handles traffic. You can enable FCP Error Recovery, change the default settings of Queues, and Interrupt handling for performance enhancement.

1. Log in to Cisco Intersight with your Cisco ID and select admin role.
2. From the **Service Selector** drop-down list, select **Infrastructure Service**.
3. Navigate to **Configure > Policies**, and then click **Create Policy**.
4. Select **Fibre Channel Adapter**, and then click **Start**.
5. In the **General** page, configure the following parameters:

Property	Essential Information
Organization	Select the organization.
Name	Enter a name for your policy.
Set Tags (optional)	Enter a tag in the key value format. For example, Org: IT or Site: APJ.
Description (optional)	Enter a short description.
Cisco Provided Fibre Channel Adapter Configuration	
Select Cisco Provided Configuration (Optional)	Click Select Cisco Provided Configuration , search, and choose from the available pre-defined configurations. Note After you choose a configuration, the policy is updated with the pre-defined values from the chosen configuration. You can modify the values in the Details page or skip Step 6 and proceed to create the policy using these pre-defined values.

6. On the **Policy Details** page, configure the following parameters:

Property	Essential Information
Error Recovery	
FCP Error Recovery	Enables the use of FCP Sequence Level Error Recovery protocol (FC-TAPE) on the virtual interface.

Property	Essential Information
Port Down Timeout, ms	The number of milliseconds a remote Fibre Channel port should be offline before informing the SCSI upper layer that the port is unavailable. Enter an integer between 0 and 240000.
I/O Retry Timeout, Seconds	The number of seconds the adapter waits before aborting the pending command and resending the same I/O request. Enter an integer between 1 and 59.
Link Down Timeout, ms	The number of milliseconds the uplink port should be offline before it informs the system that the uplink port is down and fabric connectivity has been lost. Enter an integer between 0 and 240000.
Port Down IO Retry, ms	The number of times an IO request to a port is returned because the port is busy before the system decides the port is unavailable. Enter an integer between 0 and 255.
Error Detection	
Error Detection Timeout	Error Detection Timeout, also referred to as EDTOV, is the number of milliseconds to wait before the system assumes that an error has occurred. Enter an integer between 1000 and 10000.
Resource Allocation	
Resource Allocation Timeout	The number of milliseconds to wait before the system assumes that a resource cannot be properly allocated. Enter an integer between 5000 and 100000.
Flogi	
Flogi Retries	The number of times that the system tries to log in to the fabric after the first failure.
Flogi Timeout, ms	The number of milliseconds that the system waits before it tries to log in again. Enter an integer between 1000 and 255000.
Plogi	

Property	Essential Information
Plogi Retries	The number of times that the system tries to log into a port after the first failure. Enter an integer between 0 and 255.
Plogi Timeout, ms	The number of milliseconds that the system waits before it tries to log in again. Enter an integer between 1000 and 255000
Interrupt	
Mode	Select the preferred driver interrupt mode: <ul style="list-style-type: none"> • MSIx—Message Signaled Interrupts (MSI) with the optional extension. This is the recommended option. • MSI—Message Signaled Interrupts (MSI) only • INTx—PCI INTx interrupts
IO Throttle	
I/O Throttle Count	The number of I/O operations that can be pending in the vHBA at one time. Enter an integer between 1 and 1024.
LUN	
Maximum LUNs Per Target	The maximum number of LUNs that the driver will export. This is usually an operating system platform limitation. Enter an integer between 1 and 1024. For fc-initiator vHBA type, enter an integer between 1 and 4096. Note The fc-initiator vHBA maximum LUN configuration requires the minimum server firmware version 4.2(3d). For more information on the supported firmware for adapters, see Supported Hardware .
LUN Queue Depth	The number of commands that the HBA can send and receive in a single transmission per LUN. Enter an integer between 1 and 254.
Receive	
Receive Ring Size	The number of descriptors in each queue. Enter an integer between 64 and 2048.

Property	Essential Information
Transmit	
Transmit Ring Size	The number of descriptors in each queue. Enter an integer between 64 and 2048.
SCSI I/O	
SCSI I/O Queues	The number of SCSI I/O queue resources the system should allocate. Enter an integer between 1 and 245.
SCSI I/O Ring Size	The number of descriptors in each SCSI I/O queue. Enter an integer between 64 and 512.

- Click **Create**.

Creating a Fibre Channel Network Policy

A Fibre Channel Network policy governs the Virtual Storage Area Network (VSAN) configuration for the virtual interfaces.

- Log in to Cisco Intersight with your Cisco ID and select admin role.
- From the **Service Selector** drop-down list, select **Infrastructure Service**.
- Navigate to **Configure > Policies**, and then click **Create Policy**.
- Select **Fibre Channel Network**, and then click **Start**.
- On the **General** page, configure the following parameters:

Property	Essential Information
Organization	Select the Organization.
Name	Enter a name for your policy.
Description (Optional)	Provide a short description
Add Tag (Optional)	Enter a tag in the key:value format. For example, Org: IT or Site: APJ.

- On the **Policy Details** page, configure the following parameters:

Property	Essential Information
Default VLAN	Default VLAN of the virtual interface in Standalone Rack server. Setting the value to 0 is equivalent to None and will not associate any default VLAN to the traffic on the virtual interface. Valid values are 0 to 4094.
VSAN ID	Default VSAN ID of the virtual interface. Setting the ID to 0 will not associate any default VSAN to the traffic on the virtual interface.

7. Click **Create**.

Creating a Fibre Channel QoS Policy

The Fibre Channel QoS policy assigns a system class to the outgoing traffic for a vHBA. This system class determines the quality of service for the outgoing traffic. For certain adapters, you can also specify additional controls like burst and rate on the outgoing traffic.

1. Log in to Cisco Intersight with your Cisco ID and select admin role.
2. From the **Service Selector** drop-down list, select **Infrastructure Service**.
3. Navigate to **Configure > Policies**, and then click **Create Policy**.
4. Select **Fibre Channel QoS**, and then click **Start**.
5. In the **General** page, configure the following parameters:

Property	Essential Information
Organization	Select the organization.
Name	Enter a name for your policy.
Tag (optional)	Enter a tag in the key value format. For example, Org: IT or Site: APJ.
Description (optional)	Enter a short description.

6. On the **Policy Details** page, configure the following parameters:

Property	Essential Information
Rate Limit, Mbps	Used for limiting the data rate on the virtual interface. The valid range is between 0 and 100000. The default value is Zero.

Property	Essential Information
Maximum Data Field Size, Bytes	The maximum size of the Fibre Channel frame payload bytes that the virtual interface supports. The valid range is between 256 and 2112. The default value is 2112.
Class of Service	The Class of Service to be associated to the traffic on the virtual interface. The valid range is between 0 and 6. The default value is 3. Note <ul style="list-style-type: none"> • FCoE traffic has a reserved QoS system class that should not be used by any other type of traffic. If any other type of traffic has a CoS value that is used by FCoE, the value is remarked to 0. • This property is supported only on Standalone servers.
Burst	The burst traffic allowed on the vNIC in bytes. The valid range is between 1024 and 1000000. The default value is 1024. Note This property is supported only on FI-attached servers.
Priority	The priority matching the System QoS defined in the domain profile. The Fibre Channel (FC) is enabled by default. Note This property is supported only on FI-attached servers.

7. Click **Create**.

Create FC Zone Policy

This policy allows you to set up access control between hosts and storage devices.

Certain points to be noted when creating the FC Zone policy:

- Deploying a storage VSAN using a domain profile, for the first time, clears all the unmanaged zones from the Fabric Interconnect.
- SAN boot targets with a storage VSAN have a zone entry in the Fabric Interconnect.
- A one-time SAN boot with a storage VSAN has a zone entry in the Fabric Interconnect.

- Editing the FC Zone policy causes the server profile status to be changed to Pending Changes.
 - When the Fabric Interconnect is rebooted, there is a replay of zones in the configuration.
 - Detection of configuration drift is not supported for FC Zone policy.
1. Log in to Cisco Intersight with your Cisco ID and select admin role.
 2. From the **Service Selector** drop-down list, select **Infrastructure Service**.
 3. Navigate to **Configure > Policies**, and then click **Create Policy**.
 4. Select **FC Zone**, and then click **Start**.
 5. On the **General** page, configure the following parameters:

Property	Essential Information
Organization	Select the Organization.
Name	Enter a name for your policy.
Set Tags (Optional)	Enter a tag in the key:value format. For example, Org: IT or Site: APJ.
Description (Optional)	Provide a short description

6. On the **Policy Details** page, configure the following parameters:

Property	Essential Information
FC Target Zoning Type	Type of FC Zoning. FC Zoning can be of type: <ul style="list-style-type: none"> • Single Initiator Single Target • Single Initiator Multiple Target • None <p>Note If you select FC Zoning Type as None, you cannot add targets nor view the table of added FC Zone sets.</p>
Add Target	Click to add target details of the FC Zone policy.
Name	Name of the FC Zone policy.
WWPN	WWPN that is a member of the FC Zone.
Switch ID	Unique identifier of the Fabric object. The Switch ID can be A or B.

Property	Essential Information
VSAN ID	<p>Unique identifier of the VSAN on which the FC Zone is to be created. Valid values for the VSAN ID are 1 to 4093.</p> <p>Note The VSAN ID scope should be Storage in the VSAN policy specified for the domain.</p>

7. Click **Create**.

Creating a Firmware Policy

This policy allows you to see the firmware present in your systems, as against the firmware baseline. Firmware policy also enables you to bring the firmware of your systems in line with the desired version and thereby enables the drive to compliance.

1. On the **General** page, configure the following parameters:

Property	Essential Information
Organization	Select the Organization.
Name	Enter a name for your policy.
Set Tags (Optional)	Enter a tag in the key:value format. For example, Org: IT or Site: APJ.
Description (Optional)	Provide a short description

2. On the **Policy Details** page, onfigure the following parameters:

Property	Essential Information
Advanced Mode	Enable Advanced Mode to exclude components during firmware upgrade.
Exclude Drives	Enable Advanced Mode and select the Exclude Drives checkbox to exclude drives from the firmware upgrade.
Exclude Storage Controllers	Enable Advanced Mode and select the Exclude Storage Controllers checkbox to exclude storage controllers from the firmware upgrade.
Server Model	Select the server family for the firmware upgrade. Click + to add more server models.
Firmware Version	Select the bundle version to which the server is to be upgraded.

3. Click **Create**.

Creating a BIOS Policy

A BIOS policy automates the configuration of BIOS settings on servers. You can create one or more BIOS policies that contain a specific grouping of BIOS settings, matching the needs of a server or a set of servers. If you do not specify a BIOS policy for a server, the BIOS settings will default to set of values for a brand new baremetal server or to a set of values previously configured using Cisco IMC. If a BIOS policy is specified, its values replace any previously configured values on the server.

All BIOS tokens are not applicable to all servers. If unsupported tokens are pushed to a server, those tokens are ignored.

1. Log in to Cisco Intersight with your Cisco ID and select admin role.
2. From the **Service Selector** drop-down list, select **Infrastructure Service**.
3. Navigate to **Configure > Policies**, and then click **Create Policy**.
4. Select **BIOS**, and then click **Start**.
5. On the **General** page, configure the following parameters:

Property	Essential Information
Organization	Select the Organization.
Name	Enter a name for your policy.
Set Tags (Optional)	Enter a tag in the key:value format. For example, Org: IT or Site: APJ.
Description (Optional)	Enter a short description.
Cisco Provided BIOS Configuration	
Select Cisco Provided Configuration (optional)	<p>Click Select Cisco Provided Configuration to search and choose one of the pre-defined BIOS configuration settings.</p> <p>Note After you choose a configuration, the policy is updated with the pre-defined values from the chosen configuration. You can modify the values in the Details page or skip Step 6 and proceed to create the policy using these pre-defined values.</p>

6. On the **Policy Details** page, configure the following BIOS policy options. For more information, see [Cisco UCS Server BIOS Tokens in Intersight Managed Mode](#).
7. Click **Create**.

Creating a Boot Order Policy

The Boot Order policy configures the linear ordering of devices and enables you to change the boot order and boot mode. You can also add multiple devices under various device types, rearrange the boot order, and set parameters for each boot device type.

1. Log in to Cisco Intersight with your Cisco ID and select admin role.
2. From the **Service Selector** drop-down list, select **Infrastructure Service**.
3. Navigate to **Configure > Policies**, and then click **Create Policy**.
4. Select **Boot Order**, and then click **Start**.
5. In the **General** page, configure the following parameters:

Property	Essential Information
Organization	Select the organization.
Name	Enter a name for your policy.
Tag (optional)	Enter a tag in the key value format. For example, Org: IT or Site: APJ.
Description (optional)	Enter a short description.

6. On the **Policy Details** page, configure the following parameters:

Property	Essential Information
BootMode	<p>The type of boot mode that is enabled. This can be one of the following:</p> <ul style="list-style-type: none"> • Legacy—Uses the Master Boot Record (MBR) partitioning scheme. Select if the system is not UEFI-enabled. • UEFI—Uses the GUID Partition Table (GPT). Select Unified Extensible Firmware Interface (UEFI) if the system is UEFI-enabled. <p>Note The Legacy boot mode is currently not supported on Cisco UCS C225 M6, C245 M6, C220 M7, C240 M7 and C245 M8 servers.</p>
Enable Secure Boot Mode	<p>This option is available only when UEFI Boot Mode is enabled.</p> <p>Secure boot mode enforces that a device boots using the software that is trusted by the Original Equipment Manufacturer (OEM).</p>

Property	Essential Information
Add Boot Device	

Property	Essential Information
	<p>Select to add and configure a boot device. The configuration options vary with boot device types. The supported boot devices and its configuration options for UCS standalone and FI-attached servers are listed below:</p> <ul style="list-style-type: none"> • FlexMMC Boot <ul style="list-style-type: none"> Note <ul style="list-style-type: none"> • FlexMMC boot is supported only with UEFI Boot Mode for C-series standalone servers. • Secure Boot option is supported for FlexMMC. <p>For more information on the firmware requirements for FlexMMC Boot, see Firmware Requirements for FlexMMC Boot Option.</p> <p>Configuration options:</p> <ul style="list-style-type: none"> • Device Name—Name of the boot device. • Sub-Type—The sub-type for the selected device <ul style="list-style-type: none"> • None • FlexMMC Mapped DVD • FlexMMC Mapped HDD <ul style="list-style-type: none"> • HTTP Boot <ul style="list-style-type: none"> Note HTTP/HTTPS boot is supported only with UEFI Boot Mode for both IMM servers and C-series standalone servers. <p>For more information on the firmware requirements for HTTP Boot, see Firmware Requirements for HTTP Boot Option.</p> <p>Configuration options:</p> <ul style="list-style-type: none"> • Device Name—Name of the boot device. • IP Type—The IP address family type to use during the HTTP boot process. • IP Config Type—The IP config type to use during the HTTP Boot process.

Property	Essential Information
	<ul style="list-style-type: none"> • DHCP <ul style="list-style-type: none"> • [Optional] URI—The boot resource location in URI format. Note If you do not enter a URI, ensure that DHCP is configured with client extensions. • Interface Name (Only for UCS Server (FI-Attached))—The name of the underlying vNIC that will be used by the HTTP boot device. You can select a vNIC that was configured using the LAN Connectivity Policy. For more information, see the LAN Connectivity Policy section.

Property	Essential Information
	<ul style="list-style-type: none"> • Static <p><i>When IP Config Type is Static and IP Type is IPv4:</i></p> <ul style="list-style-type: none"> • DNS IP—The IP address of DNS server. • Gateway IP—The IP address of default gateway. • Static IP—IPv4 or IPv6 static Internet Protocol address. • Network Mask—Network mask of the IPv4 address. • URI—The boot resource location in URI format. • Interface Name—The name of the underlying vNIC that will be used by the HTTP boot device. You can select a vNIC that was configured using the LAN Connectivity Policy. <p><i>When IP Config Type is Static and IP Type is IPv6:</i></p> <ul style="list-style-type: none"> • DNS IP—The IP address of DNS server. • Gateway IP—The IP address of default gateway. • Static IP—IPv4 or IPv6 static Internet Protocol address. • Prefix Length—A prefix length which masks the IP address and divides the IP address into network address and host address. • URI—The boot resource location in URI format. • Interface Name—The name of the underlying vNIC that will be used by the HTTP boot device. You can select a vNIC that was configured using the LAN Connectivity Policy.

Property	Essential Information
	<ul style="list-style-type: none"> <li data-bbox="1036 289 1481 342">• Protocol—The protocol used for HTTP Boot. <p data-bbox="1052 363 1481 583">To use the HTTPS protocol, you must have a valid Root CA Certificate for authentication. You can deploy Root CA certificates using the Certificate Management Policy. For more information, see the <i>Creating a Certificate Management Policy</i> section.</p> <p data-bbox="1052 604 1481 919">Note Certificate Management Policy does not support addition, deletion, and modification of a single certificate. Even if one of the certificates is added, deleted or modified in policy, the Server Profile will need to be redeployed or Server Action must be performed, for certificate changes to take effect.</p>

Property	Essential Information
	<ul style="list-style-type: none"> • Interface Source (Only for C-series standalone servers)—Lists the supported Interface Source for HTTP device. • Interface Name (Only for VIC Adapters) <ul style="list-style-type: none"> • Slot—The slot ID of the adapter on which the underlying virtual ethernet interface is present. • Interface Name—The name of the underlying virtual ethernet interface used by the HTTP boot device. • Port (Only for VIC Adapters) <ul style="list-style-type: none"> • Slot—The slot ID of the adapter on which the underlying virtual ethernet interface is present. • Port—The Port ID of the adapter on which the underlying virtual ethernet interface is present. If no port is specified, the default value is -1. Supported values are 0 to 255. • MAC Address <ul style="list-style-type: none"> • Slot—The slot ID of the adapter on which the underlying virtual ethernet interface is present. • MAC—The MAC address of the underlying virtual ethernet interface used by the HTTP boot device. • iSCSI Boot <ul style="list-style-type: none"> • Device Name—Name of the boot device. • Slot—The slot id of the boot device. • Port—The port id of the boot device.

Property	Essential Information
	<ul style="list-style-type: none"> • Local CDD <ul style="list-style-type: none"> • Device Name—Name of the boot device. • Local Disk <ul style="list-style-type: none"> Note This device allows the host to use the virtual drive as a bootable device. • Device Name—Name of the boot device. • Slot—The slot id of the boot device. • NVMe <ul style="list-style-type: none"> • Device Name—Name of the boot device. • Bootloader Name—Name of the bootloader image. • Bootloader Description—Description of the bootloader. • Bootloader Path—Path to the boatloader image. Note The NVMe device can be configured only on UEFI mode. • PCH Storage <ul style="list-style-type: none"> • Device Name—Name of the boot device. • LUN—The Logical Unit Number (LUN) of the boot device (0-255). Note Only UEFI boot mode is supported with software RAID configuration. • PXE Boot <ul style="list-style-type: none"> • Device Name—Name of the boot device. • IP Type—The IP address family type to use during the PXE boot process. • Slot—The slot ID of the adapter on which the virtual ethernet interface is present. • Interface Name/Port/ MAC Address—The name or address of the underlying virtual ethernet interface used by the PXE boot device.

Property	Essential Information
	<ul style="list-style-type: none"> • SAN Boot <ul style="list-style-type: none"> • Device Name—Name of the boot device. • LUN—The Logical Unit Number (LUN) of the boot device (0-255). • Slot—The slot id of the boot device. This field is applicable only for Standalone servers. • Interface Name—The name of the underlying vHBA interface. • Target WWPN—The WWPN Address of the underlying fibre channel interface • Bootloader Name — The name of the bootloader image. This field is available only in UEFI Mode. • Bootloader Description— The details of the bootloader image. This field is available only in UEFI Mode. • Bootloader Path— The path of the bootloader image. This field is available only in UEFI Mode. • SD Card <ul style="list-style-type: none"> • Device Name—Name of the boot device. • LUN—The Logical Unit Number (LUN) of the boot device (0-255). • Sub-Type— The sub-type for the selected device: <ul style="list-style-type: none"> • FlexUtil • FlexFlash • SDCard • UEFI Shell <ul style="list-style-type: none"> • Device Name—Name of the boot device.

Property	Essential Information
	<ul style="list-style-type: none"> • USB <ul style="list-style-type: none"> • Device Name—Name of the boot device. • Sub-Type— The sub-type for the selected device: <ul style="list-style-type: none"> • CD • FDD • HDD • Virtual Media <ul style="list-style-type: none"> • Device Name—Name of the boot device. • Sub-Type— The sub-type for the selected device: <ul style="list-style-type: none"> • None <ul style="list-style-type: none"> Note This option is not supported on UCS FI-attached servers. • CIMC Mapped DVD • CIMC Mapped HDD • KVM Mapped DVD • KVM Mapped HDD • KVM Mapped FDD <p>Note The device name of the boot devices can be any string that adheres to the following constraints. It should start and end with an alphanumeric character. It can have underscores and hyphens. It cannot be more than 30 characters.</p>

7. Click **Create**.

Configuring an iSCSI Boot Policy

iSCSI boot support allows you to initialize the Operating System on FI-attached blade and rack servers from a remote disk across a Storage Area Network. The remote disk, known as the target, is accessed using TCP/IP and iSCSI boot firmware.

Prerequisites

The following are required to configure the iSCSI boot device:

- **iSCSI Static Target Policy**—When you select **Static** as the mode for configuring the iSCSI boot policy, you can use the iSCSI Static Target policy to specify the primary target details. You can also specify the details of a secondary target, if required.
- **iSCSI Adapter Policy**—Using this policy you can specify the TCP and DHCP Connection Timeout and the retry count when the logical unit number of the boot device is busy.
- **Creating an IQN Pool**—Using this policy you can specify the TCP and DHCP Connection Timeout and the retry count when the logical unit number of the boot device is busy.

1. Log in to Cisco Intersight with your Cisco ID and select admin role.
2. From the **Service Selector** drop-down list, select **Infrastructure Service**.
3. Navigate to **Configure > Policies**, and then click **Create Policy**.
4. Select **iSCSI Boot**, and then click **Start**.
5. In the **General** page, configure the following parameters:

Property	Essential Information
Organization	Select the organization.
Name	Enter a name for your policy.
Tag (optional)	Enter a tag in the key value format. For example, Org: IT or Site: APJ.
Description (optional)	Enter a short description.

6. On the **Policy Details** page, configure the following parameters:

Property	Essential Information
Target Interface	Target interface can be Auto or Static.
DHCP Vendor ID/IQN	If you select Auto for the target interface, specify the Initiator name, or the DHCP vendor ID. The vendor ID can be up to 32 alphanumeric characters.
Static	If the target interface is Static specify the following parameters.
Primary Target	Select the Primary Target policy. iSCSI target is the remote disk in the storage area network from which the operating system is initialized. This policy specifies the Target Name, the IP Address of the target, the Port, and the LUN ID.
Secondary Target	Select the Secondary Target policy. Secondary Target is optional

Property	Essential Information
Adapter Policy	Select the Adapter Policy for the iSCSI boot device. The Adapter Policy specifies the TCP and DHCP Timeouts, and the Retry Count if the LUN ID is busy.
<p>Authentication</p> <p>You can select CHAP or Mutual CHAP as the authentication method and specify the parameters. If you have selected CHAP, specify the CHAP authentication parameters for iSCSI Target. Mutual CHAP is a two-way DHCP mechanism and is more secure.</p>	
CHAP	<p>For CHAP authentication, enter:</p> <ul style="list-style-type: none"> • Username: The user Id of the Initiator/Target Interface. Enter between 1 and 128 characters, spaces, or special characters. • Password: Password of Initiator or Target Interface. Enter between 12 and 16 characters, including special characters except spaces, tabs, line breaks. • Password Confirmation: Re-enter the password that you entered. Both the password and password confirmation have to match.
Mutual CHAP	<p>Mutual CHAP is a two-way CHAP mechanism. For Mutual CHAP authentication, enter:</p> <ul style="list-style-type: none"> • Username: The user Id of the Initiator or Target Interface. Enter between 1 and 128 characters, spaces, or special characters. • Password: Password of Initiator or Target Interface. Enter between 12 and 16 characters, including special characters except spaces, tabs, line breaks. • Password Confirmation: Re-enter the password that you entered. Both the password and password confirmation have to match.

Property	Essential Information
Initiator IP Source	<p>Select the method that determines the Initiator IP Source. The methods to determine the Initiator IP Source are:</p> <ul style="list-style-type: none"> • Pool: You can select an IP pool • Auto: The IP is automatically determined • Static: You can specify a static IP address as the Initiator IP. Select Static and specify: <ul style="list-style-type: none"> • IP Address: Enter the Static IP address provided for iSCSI Initiator. • Subnet Mask: Enter the 32-bit number that masks an IP address and divides the IP address into network address and host address.. • Default Gateway: Enter the IP address of the default IPv4 gateway. • Primary DNS: Enter the IP address of the primary Domain Name System server. • Secondary DNS: Enter the IP address of the secondary Domain Name System server.

7. Click **Create**.

Creating an iSCSI Adapter Policy

The iSCSI Adapter policy allows you to configure values for TCP Connection Timeout, DHCP Timeout, and the Retry Count if the specified LUN ID is busy.

1. Log in to Cisco Intersight with your Cisco ID and select admin role.
2. From the **Service Selector** drop-down list, select **Infrastructure Service**.
3. Navigate to **Configure > Policies**, and then click **Create Policy**.
4. Select **iSCSI Adapter**, and then click **Start**.
5. In the **General** page, configure the following parameters:

Property	Essential Information
Organization	Select the organization.
Name	Enter a name for your policy.

Property	Essential Information
Tag (optional)	Enter a tag in the key value format. For example, Org: IT or Site: APJ.
Description (optional)	Enter a short description.

6. On the **Policy Details** page, configure the following parameters:

Property	Essential Information
TCP Connection Timeout	Enter the number of seconds after which the TCP connection times out.
DHCP Timeout	Enter the number of seconds after which the DHCP times out.
LUN Busy Retry Count	Enter the number of times connection is to be attempted when the LUN ID is busy.

7. Click **Create**.

Creating an iSCSI Static Target Policy

The iSCSI Static Target policy allows you to specify the name, IP address, port, and logical unit number of the primary target for iSCSI boot. You can optionally specify these details for a secondary target as well.

1. Log in to Cisco Intersight with your Cisco ID and select admin role.
2. From the **Service Selector** drop-down list, select **Infrastructure Service**.
3. Navigate to **Configure > Policies**, and then click **Create Policy**.
4. Select **iSCSI Static Target**, and then click **Start**.
5. In the **General** page, configure the following parameters:

Property	Essential Information
Organization	Select the organization.
Name	Enter a name for your policy.
Tag (optional)	Enter a tag in the key value format. For example, Org: IT or Site: APJ.
Description (optional)	Enter a short description.

6. On the **Policy Details** page, configure the following parameters:

Property	Essential Information
Target Name	Enter the name of the target.

Property	Essential Information
IP Address	Enter the target IP address.
Port	Enter the port number of the target.
LUN ID	Enter the ID of the boot logical unit number.

- Click **Create**.

Creating a Device Connector Policy

Device Connector Policy lets you choose the **Configuration from Intersight only** option to control configuration changes allowed from Cisco IMC. The **Configuration from Intersight only** option is enabled by default. You will observe the following changes when you deploy the Device Connector policy in Intersight:

- Validation tasks will fail:
 - If Intersight Read-only mode is enabled in the claimed device.
 - If the firmware version of the Cisco UCS Standalone C-Series Servers is lower than 4.0(1).
- If Intersight Read-only mode is enabled, firmware upgrades will be successful only when performed from Intersight. Firmware upgrade performed locally from Cisco IMC will fail.
- IPMI over LAN privileges will be reset to read-only level if Configuration from Intersight only is enabled through the Device Connector policy, or if the same configuration is enabled in the Device Connector in Cisco IMC.

- Log in to Cisco Intersight with your Cisco ID and select admin role.
- From the **Service Selector** drop-down list, select **Infrastructure Service**.
- Navigate to **Configure > Policies**, and then click **Create Policy**.
- Select **Device Connector**, and then click **Start**.
- On the **General** page, configure the following parameters:

Property	Essential Information
Organization	Select the Organization.
Name	Enter a name for your policy.
Description (Optional)	Provide a short description
Add Tag (Optional)	Enter a tag in the key:value format. For example, Org: IT or Site: APJ.

- In the **Policy Details** page, enable or disable **Configuration from Intersight only**. This option is enabled by default.
- Click **Create**.

Creating a Drive Security Policy

Drive Security Policy enables you to configure security keys either locally or remotely using a KMIP server.

1. Log in to Cisco Intersight with your Cisco ID and select **Admin** role.
2. From the **Service Selector** drop-down list, select **Infrastructure Service**.
3. Navigate to **Configure > Policies**, and then click **Create Policy**.
4. Select **Drive Security**, and then click **Start**.
5. On the **General** page, configure the following parameters:

Property	Essential Information
Organization	Select the Organization.
Name	Enter a name for your policy.
Set Tag (Optional)	Enter a tag in the key:value format. For example, Org: IT or Site: APJ.
Description	Enter a description for your policy.

6. Do one of the following on the **Policy Details** page:
 - To configure a manual key locally:
 - a. Under **Drive Security Configuration Type**, click **Manual Key**, and then do one of the following:
 - If you are configuring the security key for the first time, in the **New Security Key Passphrase** field, enter a new alphanumeric passphrase for the security key. The passphrase should be at least 8 to 32 characters long and should include at least one uppercase letter, one lowercase letter, one number, and one special character.



Caution To avoid losing the manual security key passphrase and ensure its recovery, it is recommended to save the key externally.

- If security is already enabled with a manual key and you want to reset the key:
 1. In the **New Security Key Passphrase** field, enter an alphanumeric passphrase. The passphrase should be at least 8 to 32 characters long and should include at least one uppercase letter, one lowercase letter, one number, and one special character.



Caution To avoid losing the manual security key passphrase and ensure its recovery, it is recommended to save the key externally.

2. Select the **Drive Security is Already Enabled with a Manual Key** checkbox, and then enter the existing key in the **Current Security Key Passphrase** field.



Note The key entered in the **Current Security Key Passphrase** field is not validated immediately. The validation happens only while you deploy the profile. If there is a key mismatch, the policy deployment will fail.

- To configure a KMIP server for remotely managing the key:
 1. Under **Drive Security Configuration** Type, click **Remote Key Management**.
 2. Configure the following parameters:

Property	Essential Information
Hostname/IP Address	Enter the IP address of the KMIP server that you want to use.
Port	Enter the port number for the KMIP server. The default port is 5696.
Timeout	Enter the time that will be allowed to elapse within which the KMIP client should connect. The recommended timeout interval is up to 65 seconds.

3. To configure a fallback KMIP server, add the details of an additional KMIP server under the **Secondary KMIP Server**.
4. In the **Server Public Root CA Certificate** field, copy-paste the root certificate from the KMIP server.
5. If your KMIP server supports authentication, click the **Enable Authentication** option for additional security and enter your username and password.



Note You can use authentication only if the KMIP server supports it.

7. Click **Create**.

The newly created policy is displayed in the table view on the **Policy Details** page.

See also, [Secure Self-Encrypting Drives](#).

Creating a Disk Group Policy

The Disk Group policy defines how a disk group (a group of physical disks that are used for creating virtual drives) is created and configured, and specifies the RAID level to be used for the disk group. With this policy,

you can select the physical disks that have to be part of a disk group. When a Disk Group policy is associated with multiple virtual drives in a Storage policy, the virtual drives share the same disk group space.



Note This policy is not applicable for virtual drives for a Cisco Boot Optimised M.2 RAID Controller.

1. Log in to Cisco Intersight with your Cisco ID and select admin role.
2. From the **Service Selector** drop-down list, select **Infrastructure Service**.
3. Navigate to **Configure > Policies**, and then click **Create Policy**.
4. Select **Disk Group**, and then click **Start**.
5. In the **General** page, configure the following parameters:

Property	Essential Information
Organization	Select the organization.
Name	Enter a name for your policy.
Tag (optional)	Enter a tag in the key value format. For example, Org: IT or Site: APJ.
Description (optional)	Enter a short description.

6. On the **Policy Details** page, configure the following parameters:

Property	Essential Information
Virtual Drive Configuration	

Property	Essential Information
<p>RAID Level</p>	<p>Set the Redundant Array of Inexpensive Disks (RAID) level to ensure availability and redundancy of data, and I/O performance.</p> <p>The supported RAID levels for the disk group are:</p> <ul style="list-style-type: none"> • RAID0—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. • RAID1—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. • RAID5—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. • RAID6—Data is striped across all disks in the array and two sets of parity data are used to provide protection against failure of up to two physical disks. In each row of data blocks, two sets of parity data are stored. • RAID10—This RAID uses mirrored pairs of disks to provide complete data redundancy and high throughput rates through block-level striping. RAID 10 is mirroring without parity and block-level striping. A minimum of four disks are required for RAID 10. • RAID50—Data is striped across multiple striped parity disk sets to provide high throughput and multiple disk failure tolerance. • RAID60—Data is striped across multiple striped dual parity disk sets to provide high throughput and greater disk failure tolerance.
<p>Local Disk Configuration - Disk Group (Span 0)</p>	
<p>Drive Number</p>	<p>Specify the drive number for the disk group associated with the RAID controller.</p>
<p>Dedicated Hot Spares</p>	

Property	Essential Information
Dedicated Hot Spares	Select Enable to use a hot spare drive in the case of disk failure in the disk group.
Drive Number	Specify the identified drive number to act as a dedicated hot spare for the disk group.
Set Disks in JBOD state to Unconfigured good	Select to allow users to convert any disks in JBOD to be un-configured good disks so that they can be used in the RAID group.



Attention All virtual drives in a disk group should be managed by using the same disk group policy.

- Click **Create**.

Creating an IMC Access Policy

The IMC Access policy allows you to configure your network and associate an IP address from an IP Pool with a server. In-Band IP address, Out-Of-Band IP address, or both In-Band and Out-Of-Band IP addresses can be configured using IMC Access Policy and is supported on Drive Security, SNMP, Syslog, and vMedia policies.



- Note**
- The Out-of-Band IP address support for SNMP policy is available only for the Fabric Interconnects running on Infrastructure Firmware 4.3(2.230129) or later versions.
 - The Out-of-Band IP address support for SNMP policy is not available for servers in Cisco UCS X-Series Direct system.
 - Disjoint Layer 2 rules apply to the Inband VLANs configured on the blades within the same chassis. You can configure multiple Inband VLANs for blades within the same chassis, provided the VLANs belong to the same Disjoint Group and are allowed on the same uplink.

- Log in to Cisco Intersight with your Cisco ID and select admin role.
- From the **Service Selector** drop-down list, select **Infrastructure Service**.
- Navigate to **Configure > Policies**, and then click **Create Policy**.
- Select **IMC Access**, and then click **Start**.
- On the **General** page, configure the following parameters:

Property	Essential Information
Organization	Select the Organization.
Name	Enter a name for your policy.

Property	Essential Information
Description (Optional)	Provide a short description
Add Tag (Optional)	Enter a tag in the key:value format. For example, Org: IT or Site: APJ.

6. On the **Policy Details** page, configure the following parameters:

Property	Essential Information
In-Band Configuration	Enable, to have the server management services made available using the uplink port.
VLAN ID	Enter the VLAN ID to be used for server access over the inband network. The field value can be between 4 and 4093.
IPv4 address configuration	Select to determine the type of network for this policy. Note You can select only IPv4 address configuration or both IPv4 and IPv6 configurations.
IPv6 address Configuration	Select to determine the type of network for this policy. Note You can select only IPv6 address configuration or both IPv4 and IPv6 configurations.
IP Pool	
Select IP Pool	Click to view the list of IP Pools available and select an IP pool for In-Band configuration. Note Ensure that the default gateway specified in the IP Pool used for IMC Access Policy has connectivity to Cisco IMC. For more information, see the <i>Creating an IP Pool</i> section.

Property	Essential Information
Out-Of-Band Configuration	Enable, to have the server management services made available using the management port.
	IP Pool
	Select IP Pool Click to view the list of IP Pools available and select an IP pool for the Out-Of-Band configuration. Note Only IPv4 addresses are supported for Out-Of-Band configuration.

Creating an IPMI Over LAN Policy

The IPMI over LAN policy defines the protocols for interfacing with a service processor that is embedded in a server platform. The Intelligent Platform Management Interface (IPMI) enables an operating system to obtain information about the system health and control system hardware and directs the Cisco IMC to perform the required actions. You can create an IPMI Over LAN policy to manage the IPMI messages through Cisco Intersight.

1. Log in to Cisco Intersight with your Cisco ID and select admin role.
2. From the **Service Selector** drop-down list, select **Infrastructure Service**.
3. Navigate to **Configure > Policies**, and then click **Create Policy**.
4. Select **IPMI Over LAN**, and then click **Start**.
5. On the **General** page, configure the following parameters:

Property	Essential Information
Organization	Select the Organization.
Name	Enter a name for your policy.
Description (Optional)	Provide a short description
Add Tag (Optional)	Enter a tag in the key:value format. For example, Org: IT or Site: APJ.

6. On the **Policy Details** page, configure the following parameters:

Property	Essential Information
Enable IPMI Over LAN	The state of the IPMI Over LAN service on the endpoint.

Property	Essential Information
<p>Privilege Level</p>	<p>You can assign these privileges to the IPMI sessions on the server:</p> <ul style="list-style-type: none"> • admin—You can create admin, user, and read-only sessions on servers with the "Administrator" user role. • read-only—You can only create read-only IPMI sessions on servers with the "Read-only" user role. • user—You can create user and read-only sessions, but not admin sessions on servers with the "User" role. <p>Note</p> <ul style="list-style-type: none"> • This configuration is supported only on Cisco UCS C-Series Standalone and C-Series Intersight Managed Mode Servers. • The value of the Privilege field must match exactly the role assigned to the user attempting to log in. For example, if this field is set to read-only and a user with the admin role attempts to log in through IPMI, that login attempt will fail.
<p>Encryption Key</p>	<p>The encryption key to use for IPMI Communication. The key must have an even number of hexadecimal characters and not exceeding 40 characters. You can use "00" to disable the encryption key use. If the encryption key specified is less than 40 characters, then the IPMI commands must add zeroes to the encryption key to achieve a length of 40 characters.</p> <p>Note This encryption key configuration is supported only on Cisco UCS C-Series Standalone and C-Series Intersight Managed Mode servers. To support this configuration on Intersight Managed Mode servers, a minimum firmware version 4.2(3a) is required.</p>

7. Click **Create**.

Creating an LDAP Policy

Lightweight Directory Access Protocol (LDAP) stores and maintains directory information in a network. When LDAP is enabled in the Cisco IMC, user authentication and role authorization is performed by the LDAP server for user accounts not found in the local user database. You can enable and configure LDAP, and configure LDAP servers and LDAP groups.



Note This policy, if attached to a server profile that is assigned to an Intersight Managed FI-attached UCS server, will be ignored.

1. Log in to Cisco Intersight with your Cisco ID and select admin role.
2. From the **Service Selector** drop-down list, select **Infrastructure Service**.
3. Navigate to **Configure > Policies**, and then click **Create Policy**.
4. Select **LDAP**, and then click **Start**.
5. In the **General** page, configure the following parameters:

Property	Essential Information
Organization	Select the organization.
Name	Enter a name for your policy.
Tag (optional)	Enter a tag in the key value format. For example, Org: IT or Site: APJ.
Description (optional)	Enter a short description.

6. On the **Policy Details** page, configure the following parameters:

Property	Essential Information
Enable LDAP	The state of the LDAP service on the endpoint.
Base Settings	
Base DN	Base Distinguished Name. This field describes where to load users and groups from. It must be in the dc=domain,dc=com format for Active Directory servers.
Domain	The IPv4 domain that all users must be in. This field is required unless you specify at least one Global Catalog server address.

Property	Essential Information
Timeout	<p>The number of seconds that Intersight waits until the LDAP search operation times out.</p> <p>If the search operation times out, Intersight tries to connect to the next server listed on this tab, if one is available.</p> <p>Note The value you specify for this field could impact the overall time.</p>
Enable Encryption	If enabled, the server encrypts all information it sends to the LDAP server.
Binding Parameters	
Bind Method	<p>It can be one of the following:</p> <p>Anonymous—requires NULL username and password. If this option is selected and the LDAP server is configured for Anonymous logins, then the user can gain access.</p> <p>Configured Credentials—requires a known set of credentials to be specified for the initial bind process. If the initial bind process succeeds, then the distinguished name (DN) of the user name is queried and re-used for the re-binding process. If the re-binding process fails, then the user is denied access.</p> <p>Login Credentials—requires the user credentials. If the bind process fails, the user is denied access. By default, the Login Credentials option is selected.</p>
Bind DN	The distinguished name (DN) of the user. This field is editable only if you have selected Configured Credentials option as the binding method.
Bind Password	The password of the user. This field is editable only if you have selected Configured Credentials option as the binding method.
Search Parameters	
Filter	<p>This field must match the configured attribute in the schema on the LDAP server.</p> <p>By default, this field displays sAMAccountName.</p>
Group Attribute	<p>This field must match the configured attribute in the schema on the LDAP server.</p> <p>By default, this field displays memberOf.</p>

Property	Essential Information
<p>Attribute</p>	<p>An LDAP attribute that contains the role and locale information for the user. This property is always a name-value pair. The system queries the user record for the value that matches this attribute name.</p> <p>The LDAP attribute can use an existing LDAP attribute that is mapped to the Cisco IMC user roles and locales, or can modify the schema such that a new LDAP attribute can be created. For example, CiscoAvPair.</p> <p>Note If you do not specify this property, the user cannot login. Although the object is located on the LDAP server, it should be an exact match of the attribute that is specified in this field.</p>
<p>Group Authorization</p>	
<p>Group Authorization</p>	<p>If enabled, user authentication is also done on the group level for LDAP users that are not found in the local user database.</p>
<p>Nested Group Search Depth</p>	<p>Parameter to search for an LDAP group nested within another defined group in an LDAP group map. The parameter defines the depth of a nested group search.</p>
<p>Configure LDAP Servers</p>	
<p>Enable DNS</p>	<p>If enabled, you can use DNS to configure access to the LDAP servers.</p>
<p>Source</p>	<p>Specifies how to obtain the domain name used for the DNS SRV request. It can be one of the following:</p> <ul style="list-style-type: none"> • Extracted—specifies using domain name extracted-domain from the login ID • Configured—specifies using the configured-search domain. • Configured-Extracted—specifies using the domain name extracted from the login ID than the configured-search domain.
<p>Server</p>	<p>The IP address or host name of the LDAP server.</p>
<p>Port</p>	<p>The LDAP server port numbers.</p>

Property	Essential Information
User Search Precedence	The order of search between the local user database and LDAP user database. This can be one of the following: <ul style="list-style-type: none"> • Local User Database (Default setting) • LDAP User Database
Add New LDAP Group	
Name	The name of the group in the LDAP server database that is authorized to access the server.
Domain	The LDAP server domain the group must reside in.
Role	The role assigned to all users in this LDAP server group. This can be one of the following: <ul style="list-style-type: none"> • read-only—A user with this role can view information but cannot make any changes. • user—A user with this role can perform the following tasks: <ul style="list-style-type: none"> • View all information • Manage the power control options such as power on, power cycle, and power off • Launch the KVM console and virtual media • Clear all logs • Toggle the locator LED • Set time zone • Ping • admin—A user with this role can perform all actions available through the GUI, CLI, and IPMI.
Port	The LDAP server port numbers.
User Search Precedence	The order of search between the local user database and LDAP user database. This can be one of the following: <ul style="list-style-type: none"> • Local User Database (Default setting) • LDAP User Database

7. Click **Create**.

Creating a Local User Policy

The Local User policy automates the configuration of local user preferences. You can create one or more Local User policies which contain a list of local users that need to be configured.



Note By default, IPMI support is enabled for all users

1. Log in to Cisco Intersight with your Cisco ID and select admin role.
2. From the **Service Selector** drop-down list, select **Infrastructure Service**.
3. Navigate to **Configure > Policies**, and then click **Create Policy**.
4. Select **Local User**, and then click **Start**.
5. On the **General** page, configure the following parameters:

Property	Essential Information
Organization	Select the Organization.
Name	Enter a name for your policy.
Description (Optional)	Provide a short description
Add Tag (Optional)	Enter a tag in the key:value format. For example, Org: IT or Site: APJ.

6. On the **Policy Details** page, configure the following parameters:

Property	Essential Information
Password Properties	Password properties apply only to Rack servers and not to Blade Servers.
Enforce Strong Password	Enables strong password policy.
Change Password	Enables changing the existing password.
Enable Password Expiry	Enables password expiry on the endpoint. Note Password expiry once set by the admin is applicable for all users that are subsequently created. The valid Password Expiry Duration must be greater than the Notification Period and the Grace Period . If otherwise, you will see an User Password Expiry Policy configuration error .

Property	Essential Information
Password Expiry Duration	The time period that you can set for the existing password to expire (from the time you set a new password or modify an existing one). The range is between 1 to 3650 days.
Notification Period	Notifies the time by when the password expires. Enter a value between 0 to 15 days. Entering 0 disables this field.
Grace Period	Time period till when the existing password can still be used, after it expires. Enter a value between 0 to 5 days. Entering 0 disables this field.
Password History	The number of occurrences when a password was entered. When this is enabled, you cannot repeat a password. Enter a value between 0 to 5. Entering 0 disables this field.
Always Send User Password	When enabled, the user password is always sent to the endpoint device. When not enabled, the user password is sent to the endpoint device for new users and when the password is changed for existing users.
Add New User	
Enable	Enables the user account on the endpoint.
New User	Enables new user configuration.
Username	The username for the user. Enter between 1 and 16 characters.

Property	Essential Information
Role	<p>The role associated with the user on the endpoint.</p> <ul style="list-style-type: none"> • read-only—A user with this role can view information but cannot make any changes. • user—The user role type is supported only in racks. A user with this role can perform the following tasks: <ul style="list-style-type: none"> • View all information • Manage the power control options such as power on, power cycle, and power off • Launch the KVM console and virtual media • Clear all logs • Ping • admin—A user with this role can perform all actions available through the GUI, CLI, and IPMI.

Property	Essential Information
Password	<p>The password for this user name. When you move the mouse over the help icon beside the field, the following guidelines to set the password are displayed:</p> <ul style="list-style-type: none"> • The password must have a minimum of 8 and a maximum of 20 characters. This is an Intersight platform limitation. • The password must not contain the User Name. • The password must contain characters from three of the following four categories: <ul style="list-style-type: none"> • English uppercase characters (A through Z). • English lowercase characters (a through z). • Base 10 digits (0 through 9). • Non-alphabetic characters (!, @, #, \$, %, ^, &, *, -, _ , =, "). <p>These rules are meant to define a strong password for the user, for security reasons. However, if you want to set a password of your choice ignoring these guidelines, click the Disable Strong Password button on the Local Users tab. While setting a password when the strong password option is disabled, you can use between 1- 20 characters.</p> <p>Note You can change the password of a Local User policy by editing the policy. However, the Change Password option is disabled once the policy is deployed.</p>
Password Confirmation	The password repeated for confirmation purposes.

7. Click **Create**.

Creating an NTP Policy

The NTP policy enables the NTP service to configure a UCS system that is managed by Cisco Intersight to synchronize the time with an NTP server. You must enable and configure the NTP service by specifying the IP/DNS address of at least one server or a maximum of four servers that function as NTP servers. When you enable the NTP service, Cisco Intersight configures the NTP details on the endpoint.

1. Log in to Cisco Intersight with your Cisco ID and select admin role.

2. From the **Service Selector** drop-down list, select **Infrastructure Service**.
3. Navigate to **Configure > Policies**, and then click **Create Policy**.
4. Select **NTP**, and then click **Start**.
5. On the **General** page, configure the following parameters:

Property	Essential Information
Organization	Select the Organization.
Name	Enter a name for your policy.
Description (Optional)	Provide a short description
Add Tag (Optional)	Enter a tag in the key:value format. For example, Org: IT or Site: APJ.

6. On the **Policy Details** page, configure the following parameters:

Property	Essential Information
Enable NTP	Enables NTP policy configuration.
NTP Servers	A collection of NTP Server IP addresses or hostnames.
Time Zone	A collection of time zones from which you can select a time zone for the endpoint. This property is applicable to switches and to Cisco IMC (standalone) servers.

When a hostname is used for NTP configuration, DNS server information must be configured in the Network Connectivity policy.

7. Click **Create**.

Creating an SD Card Policy

The SD Card policy in Cisco Intersight configures the Cisco FlexFlash and FlexUtil Secure Digital (SD) cards for the Cisco UCS C-Series Standalone M4, M5 servers, and Cisco UCS C-Series M5 servers in a Cisco Intersight-Managed Fabric Interconnect Domain. This policy specifies details of virtual drives on the SD cards. You can configure the SD cards in the Operating System Only, Utility Only, or Operating System + Utility modes.

When two cards are present in the Cisco FlexFlash controller and Operating System is chosen in the SD card policy, the configured OS partition is mirrored. If only single card is available in the Cisco FlexFlash controller, the configured OS partition is non-RAID. The utility partitions are always set as non-RAID.

1. Log in to Cisco Intersight with your Cisco ID and select admin role.
2. From the **Service Selector** drop-down list, select **Infrastructure Service**.

3. Navigate to **Configure > Policies**, and then click **Create Policy**.
4. Select **SD Card**, and then click **Start**.
5. On the **General** page, configure the following parameters:

Property	Essential Information
Organization	Select the Organization.
Name	Enter a name for your policy.
Description (Optional)	Provide a short description
Add Tag (Optional)	Enter a tag in the key:value format. For example, Org: IT or Site: APJ.

6. On the **Policy Details** page, configure the following parameters:

Property	Essential Information
Operating System Only	
Operating System	Enables the Operating System partition.
Operating System Partition Name	The name for the Operating System partition.
Utility Only	
Diagnostics	Enables the Operating System health diagnostics utility.
Drivers	Enables virtual driver utility.
Host Upgrade Utility	Enables Host Upgrade Utility (HUU).
Server Configuration Utility	Enables Server Configuration Utility (SCU).
User Partition	Enables user partition.
User Partition Name	The user partition name.
Operating System + Utility	
Diagnostics	Enables the operating system health diagnostics utility.
Drivers	Enables virtual driver utility.
Host Upgrade Utility	Enables Host Upgrade Utility (HUU).
Server Configuration Utility	Enables Server Configuration Utility (SCU).
User Partition	Enables user partition.
User Partition Name	The user partition name.

Property	Essential Information
Operating System Partition	Enables the Operating System partition.
Operating System Partition Name	The name for the Operating System partition.

- Click **Create**.

Exceptions

- **SD Card Policy is not supported on M6 servers.**
- SD Card Policy is not imported with a Server Profile when the SD Cards are not present in the server.
- Diagnostics is applicable for M5 Series only.
- For the Operating System+Utility mode the M5 servers require at least 1 FlexFlash + 1 FlexUtil card.

Create a Serial Over LAN Policy

The Serial Over LAN policy enables the input and output of the serial port of a managed system to be redirected over IP. You can create one or more Serial over LAN policies which contain a specific grouping of Serial over LAN attributes that match the needs of a server or a set of servers.

- Log in to Cisco Intersight with your Cisco ID and select admin role.
- From the **Service Selector** drop-down list, select **Infrastructure Service**.
- Navigate to **Configure > Policies**, and then click **Create Policy**.
- Select **Serial Over LAN**, and then click **Start**.
- On the **General** page, configure the following parameters:

Property	Essential Information
Organization	Select the Organization.
Name	Enter a name for your policy.
Description (Optional)	Provide a short description
Add Tag (Optional)	Enter a tag in the key:value format. For example, Org: IT or Site: APJ.

- On the **Policy Details** page, configure the following parameters:

Property	Essential Information
Enable Serial Over LAN	The state of Serial Over LAN service on the endpoint.

Property	Essential Information
COM Port	<p>The serial port through which the system routes Serial Over LAN communication.</p> <ul style="list-style-type: none"> • com0—SoL communication is routed through COM port 0, an externally accessible serial port that supports either a physical RJ45 connection to an external device or a virtual SoL connection to a network device. <p>If you select this option, the system enables SoL and disables the RJ45 connection, which means that the server can no longer support an external serial device.</p> <ul style="list-style-type: none"> • com1—SoL communication is routed through COM port 1, an internal port accessible only through SoL. <p>If you select this option, you can use SoL on COM port 1 and the physical RJ45 connection on COM port 0.</p> <p>Note</p> <ul style="list-style-type: none"> • This is applicable to Cisco UCS C-Series Standalone servers only. • Serial Port is available only on some Cisco UCS C-Series servers. If it is unavailable, the server uses COM port 0 by default. Changing the Com Port setting disconnects any existing SoL sessions.
Baud Rate	<p>The Baud Rate used for Serial Over LAN communication. The rate can be:</p> <ul style="list-style-type: none"> • 9600 bps • 19.2 kbps • 38.4 kbps • 57.6 kbps • 115.2 kbps <p>Note</p> <ul style="list-style-type: none"> • The baud rate must match the baud rate configured in the server serial console. • The baud rate for Cisco UCS X-Series Direct is 115.2 kbps.

Property	Essential Information
SSH Port	<p>The SSH port used to access Serial Over LAN directly. Enables bypassing Cisco IMC shell to provide direct access to Serial Over LAN.</p> <p>The valid range is 1024 to 65535. The default value is 2400.</p> <p>Note</p> <ul style="list-style-type: none"> • This is applicable to Cisco UCS C-Series Standalone servers only. • Changing the SSH Port setting disconnects any existing SSH sessions.

7. Click **Create**.

Create SSH Policy

The SSH policy enables an SSH client to make a secure, encrypted connection. You can create one or more SSH policies that contain a specific grouping of SSH properties for a server or a set of servers.

1. Log in to Cisco Intersight with your Cisco ID and select admin role.
2. From the **Service Selector** drop-down list, select **Infrastructure Service**.
3. Navigate to **Configure > Policies**, and then click **Create Policy**.
4. Select **SSH**, and then click **Start**.
5. On the **General** page, configure the following parameters:

Property	Essential Information
Organization	Select the Organization.
Name	Enter a name for your policy.
Description (Optional)	Provide a short description
Add Tag (Optional)	Enter a tag in the key:value format. For example, Org: IT or Site: APJ.

6. On the **Policy Details** page, configure the following parameters:

Property	Essential Information
Enable SSH Policy	Enables SSH.
SSH Port	The port used for secure shell access.

Property	Essential Information
SSH Timeout (seconds)	The number of seconds to wait before the system considers a SSH request to have timed out. Enter an integer between 60 and 10,800. The default is 1,800 seconds.

- Click **Create**.

Creating a Virtual KVM Policy

The KVM console is an interface that emulates a direct keyboard, video, and mouse (KVM) connection to the server. It allows you to control the server from a remote location and to map physical locations to virtual drives that can be accessed by the server during this KVM session.

Enables specific grouping of virtual KVM properties. This policy lets you specify the number of allowed concurrent KVM sessions, port information, and video encryption options.

- Log in to Cisco Intersight with your Cisco ID and select admin role.
- From the **Service Selector** drop-down list, select **Infrastructure Service**.
- Navigate to **Configure > Policies**, and then click **Create Policy**.
- Select **Virtual KVM**, and then click **Start**.
- On the **General** page, configure the following parameters:

Property	Essential Information
Organization	Select the Organization.
Name	Enter a name for your policy.
Description (Optional)	Provide a short description
Add Tag (Optional)	Enter a tag in the key:value format. For example, Org: IT or Site: APJ.

- On the **Policy Details** page, configure the following parameters:

Property	Essential Information
Enable Virtual KVM	The state of the vKVM service on the endpoint.
Max Sessions	The maximum number of concurrent KVM sessions allowed.
Remote Port	The port for remote KVM communication. The port range is from 1024 to 49151. The default is 2068.

Property	Essential Information
Enable Video Encryption	Enables encryption on all video information sent through KVM. The Video Encryption is enabled by default. Note For firmware versions 4.2(1a) or higher, this encryption parameter is deprecated and disabling the encryption will further result in validation failure during the server profile deployment.
Enable Local Server Video	Enables KVM session displays on any monitor attached to the server. Note This is applicable to Cisco UCS C-Series Standalone servers only.
Allow Tunneled vKVM	Enable to allow tunneled vKVM on the endpoint. Note Applies only to Device Connectors that support Tunneled vKVM.

7. Click **Create**.

Exceptions

- The virtual media viewer is accessed through the KVM. If you disable the KVM console, Cisco IMC also disables access to all virtual media devices attached to the host.
- After a KVM vMedia session is mapped, if you change the KVM management policy, it will result in a loss of the vMedia session. You must re-map the KVM vMedia session again.

Creating a Virtual Media Policy

The Virtual Media policy enables you to install an operating system on the server using the KVM console and virtual media, mount files to the host from a remote file share, and enable virtual media encryption. You can create one or more virtual media policies, which could contain virtual media mappings for different OS images, and configure up to two virtual media mappings, one for ISO files through CDD and the other for IMG files through HDD.

1. Log in to Cisco Intersight with your Cisco ID and select admin role.
2. From the **Service Selector** drop-down list, select **Infrastructure Service**.
3. Navigate to **Configure > Policies**, and then click **Create Policy**.
4. Select **Virtual Media**, and then click **Start**.
5. In the **General** page, configure the following parameters:

Property	Essential Information
Organization	Select the organization.
Name	Enter a name for your policy.
Tag (optional)	Enter a tag in the key value format. For example, Org: IT or Site: APJ.
Description (optional)	Enter a short description.

6. On the **Policy Details** page, configure the following parameters:

Property	Essential Information
Enable Virtual Media	Select this option to enable the virtual media policy. This property is enabled by default.
Enable Virtual Media Encryption	Select this option to enable encryption of the virtual media communications. This property is enabled by default. Note For firmware versions 4.2(1a) or higher, this encryption parameter is deprecated and disabling the encryption will further result in validation failure during the server profile deployment.
Enable Low Power USB	Select this option to enable the appearance of virtual drives on the boot selection menu after mapping the image and rebooting the host. This property is enabled by default.
Add Virtual Media	
Virtual Media Type	Select the remote virtual media type: <ul style="list-style-type: none"> • CDD • HDD
NFS/CIFS/HTTP/HTTPS	
The properties below vary depending on the tab that is selected.	
Name	The identity of the image for virtual media mapping.

Property	Essential Information
File Location	<p>Provide the remote file location path: Host Name or IP address/file path/file name</p> <ul style="list-style-type: none"> • IP Address—The IP address or the hostname of the remote server. • File Path—The path to the location of the image on the remote server. • File Name—The name of the remote file in .iso or .img format. <p>The remote file location path for virtual media mapping, the options include:</p> <ul style="list-style-type: none"> • HDD Virtual Media: hostname or IP address /filePath/fileName.img • CDD Virtual Media: hostname or IP address /filePath/fileName.iso • HDD Virtual media for HTTP: http://server-hostname-or-ip/filePath/fileName.img • CDD Virtual media for HTTP: http://server-hostname-or-ip/filePath/fileName.iso • HDD Virtual media for HTTPS: https://server-hostname-or-ip/filePath/fileName.img • CDD Virtual media for HTTPS: https://server-hostname-or-ip/filePath/fileName.iso
Username	The username to log in to the remote server. This field is displayed on selecting CIFS, HTTP, or HTTPS.
Password	The password associated with the username. This field is displayed on selecting CIFS, HTTP, or HTTPS.

Property	Essential Information
Mount Options	<p>The mount options for the virtual media mapping. The field can be left blank or filled in a comma separated list using the following options:</p> <ul style="list-style-type: none"> • For NFS, supported options are ro, rw, noexec, noexec, soft, port=VALUE, timeo=VALUE, retry=VALUE. • For CIFS, supported options are soft, nounix, noserverino, guest, ver=3.0, or ver=2.0. <p>Note If the firmware version is 4.1 or higher, and the CIFS version is lower than 3.0, the mount option field must be entered with the version value (vers=VALUE). For example, vers=2.0.</p> <ul style="list-style-type: none"> • For HTTP and HTTPS, the only supported option is noauto.
Authentication Protocol	<p>Select the authentication protocol when CIFS is used for communication with the remote server. This field is displayed on selecting CIFS.</p> <ul style="list-style-type: none"> • 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. • ntlmv2—NTLMv2 security protocol. Use this option only with Samba Linux. • ntlmv2i—NTLMv2i security protocol. Use this option only with Samba Linux. • ntlmssp—NT LAN Manager Security Support Provider (NTLMSSP) protocol. Use this option only with Windows 2008 R2 and Windows 2012 R2. • ntlmsspi—NT LAN Manager Security Support Provider (NTLMSSPI) protocol. Use this option only when you enable Digital Signing in the CIFS Windows server.
Add	Click Add to confirm adding the virtual media.

- Click **Create**.

Exceptions

- When an answer file is embedded in the OS ISO, it fails to boot from vMedia when the bootmode is set to UEFI, and the OS installation fails on Cisco UCS C-Series Standalone M4 servers.
- vMedia mapping of the OS image for HTTPS based share fails to mount.

Creating a Network Connectivity Policy

The Network Connectivity policy enables you to configure and assign IPv4 and IPv6 addresses.

Dynamic DNS

Dynamic DNS (DDNS) is used to add or update the resource records on the DNS server. When you enable the DDNS option, the DDNS service records the current hostname, Domain name, and the management IP address and updates the resource records in the DNS server.

- Log in to Cisco Intersight with your Cisco ID and select admin role.
- From the **Service Selector** drop-down list, select **Infrastructure Service**.
- Navigate to **Configure > Policies**, and then click **Create Policy**.
- Select **Network Connectivity**, and then click **Start**.
- On the **General** page, configure the following parameters:

Property	Essential Information
Organization	Select the Organization.
Name	Enter a name for your policy.
Description (Optional)	Provide a short description
Add Tag (Optional)	Enter a tag in the key:value format. For example, Org: IT or Site: APJ.

- On the **Policy Details** page, configure the following properties:

Common Properties

Property	Essential Information
Enable Dynamic DNS	Enables Dynamic DNS. This property is not applicable to Fabric Interconnects.
Dynamic DNS Update Domain	Specify the dynamic DNS Domain. The Domain can be either a main Domain or a sub-Domain. This property is not applicable to Fabric Interconnects.

IPv4 Properties

Property	Essential Information
Obtain IPv4 DNS Server Addresses from DHCP	<p>Whether the IPv4 addresses are obtained from Dynamic Host Configuration Protocol (DHCP) or from a specifically configured set of DNS servers.</p> <ul style="list-style-type: none"> • Enabled—Intersight uses DHCP • Disabled—Intersight uses a configured set of IPv4 DNS servers. <p>This property is not applicable to Fabric Interconnects.</p>
Preferred IPv4 DNS Server	The IP address of the primary DNS server. This property is displayed only when Obtain IPv4 DNS Server Addresses from DHCP is disabled.
Alternate IPv4 DNS Server	The IP address of the secondary DNS server. This property is displayed only when Obtain IPv4 DNS Server Addresses from DHCP is disabled.

Property	Essential Information
Enable IPv6	Whether IPv6 is enabled. You can configure IPv6 properties only if this property is enabled.

IPv6 Properties

Property	Essential Information
Obtain IPv6 DNS Server Addresses from DHCP	<p>Whether the IPv6 addresses are obtained from Dynamic Host Configuration Protocol (DHCP) or from a specifically configured set of DNS servers.</p> <ul style="list-style-type: none"> • Enabled—Intersight uses DHCP • Disabled—Intersight uses a configured set of IPv6 DNS servers. <p>This property is not applicable to Fabric Interconnects.</p>
Preferred IPv6 DNS Server	The IP address of the primary DNS server. This property is displayed only when Obtain IPv6 DNS Server Addresses from DHCP is disabled.
Alternate IPv6 DNS Server	The IP address of the secondary DNS server. This property is displayed only when Obtain IPv6 DNS Server Addresses from DHCP is disabled.

7. Click **Create**.

Creating a SMTP Policy

Simple Mail Transfer Protocol (SMTP) sends server faults as email alerts to the configured SMTP server.

Sets the state of the SMTP client in the managed device. You can specify the preferred settings for outgoing communication and select the fault severity level to report and the mail recipients.



Note This policy, if attached to a server profile that is assigned to an Intersight Managed FI-attached UCS server, will be ignored.

1. Log in to Cisco Intersight with your Cisco ID and select admin role.
2. From the **Service Selector** drop-down list, select **Infrastructure Service**.
3. Navigate to **Configure > Policies**, and then click **Create Policy**.
4. Select **SMTP**, and then click **Start**.
5. On the **General** page, configure the following parameters:

Property	Essential Information
Organization	Select the Organization.
Name	Enter a name for your policy.
Description (Optional)	Provide a short description
Add Tag (Optional)	Enter a tag in the key:value format. For example, Org: IT or Site: APJ.

6. In the **Policy Details** page, configure the following parameters:

Property	Essential Information
Enable SMTP	Enables or disables the SMTP policy.
SMTP Server Address	The IP address or host name of the SMTP server.
SMTP Port	The port number used by the SMTP server for outgoing SMTP communication. The range is from 1 to 65535. The default is 25.
Minimum Severity	The minimum fault severity level to receive email notifications. Email notifications are sent for all faults whose severity is equal to or greater than the chosen level.
SMTP Alert Sender Address	The sender IP address or hostname of all the SMTP mail alerts.

Property	Essential Information
Mail Alert Recipients	A list of email addresses that will receive notifications for faults.

- Click **Create**.

Creating an SNMP Policy

The SNMP policy configures the SNMP settings for sending fault and alert information by SNMP traps from the managed devices. This policy supports SNMP versions such as SNMPv1, SNMPv2(includes v2c), and SNMPv3. Any existing SNMP Users or SNMP Traps configured previously on the managed devices are removed and replaced with users or traps that you configure in this policy. If you have not added any users or traps in the policy, the existing users or traps on the server are removed.

Using the SNMP Policy you can enable or disable SNMP, specify the access and community strings, and provide the SNMP user details that is used to retrieve data.



Note

- The Out-of-Band IP address support for SNMP policy is available only for the Fabric Interconnects running on Infrastructure Firmware 4.3(2.230129) or later versions.
- The Out-of-Band IP address support for SNMP policy is not available for servers in Cisco UCS X-Series Direct system.

- Log in to Cisco Intersight with your Cisco ID and select admin role.
- From the **Service Selector** drop-down list, select **Infrastructure Service**.
- Navigate to **Configure > Policies**, and then click **Create Policy**.
- Select **SNMP**, and then click **Start**.
- In the **General** page, configure the following parameters:

Property	Essential Information
Organization	Select the organization.
Name	Enter a name for your policy.
Tag (optional)	Enter a tag in the key value format.
Description (optional)	Enter a short description.

- In the **Policy Details** page, configure the following parameters:

Property	Essential Information
Enable SNMP	Displays the state of the SNMP Policy on the endpoint. Enable this option for the endpoint to send SNMP traps to the designated host.
SNMP Port	The port on which Cisco IMC SNMP agent runs.
Access Community String	Enter the SNMPv1, SNMPv2 community string or the SNMPv3 username. This field allows maximum of 18 characters. Note If the field is empty, it indicates that the SNMPv1 and SNMPv2c users are disabled.
SNMP Community Access	The controls access to the information in the inventory tables. Applicable only for SNMPv1 and SNMPv2c users. Note This property is supported only in UCS Standalone C-Series servers.
Trap Community String	Enter the SNMP community group name used for sending SNMP trap to other devices. Note This field is applicable only for SNMPv2c trap host or destination.
System Contact	The contact person responsible for the SNMP implementation. Enter a string up to 64 characters, such as an email address or a name and telephone number. Note This property is supported only in UCS Standalone C-Series servers.
System Location	The location of host on which the SNMP agent (server) runs. Note This property is supported only in UCS Standalone C-Series servers.
SNMP Engine Input ID	The user-defined unique identification of the static engine. Note This property is supported only in UCS Standalone C-Series servers.
SNMP Users	
Name	Enter the SNMP username. This field must have a minimum of 1 and a maximum of 31 characters.

Property	Essential Information
Security Level	Select the security mechanism for communication between the agent and the manager that include: <ul style="list-style-type: none"> • AuthPriv • AuthNoPriv
Auth Type	Select SHA as the authorization protocol for authenticating the user. Note The MD5 authorization protocol is not supported.
Auth Password	Enter the authorization password for the user.
Auth Password Confirmation	Enter the authorization password confirmation for the user.
Privacy Type	Select AES as the privacy protocol for the user. Note The DES privacy type is deprecated to meet security standards.
Privacy Password	Enter the privacy password for the user.
Privacy Password Confirmation	Enter the privacy password confirmation for the user.
SNMP Trap Destinations	
Enable	Enable this option to use the SNMP policy.
SNMP Version	Select v2 or v3 as the SNMP version for the trap.
User	Select the SNMP user for the trap. You can define maximum of 15 trap users. Note This field is applicable only to SNMPv3.
Trap Type	Select the trap type to receive a notification when a trap is received at the destination: <ul style="list-style-type: none"> • Trap • Inform
Destination Address	Provide the address to which the SNMP trap information can be sent. You are allowed to define maximum of 15 trap destinations.
Port	Enter the port number for the server to communicate with trap destination. The range is from 1 to 65535. The default is 162.

7. Click **Create**.

Creating a Storage Policy

The Storage policy allows you to create drive groups, virtual drives, configure the storage capacity of a virtual drive, and configure the M.2 RAID controllers.

1. Log in to Cisco Intersight with your Cisco ID and select admin role.
2. From the **Service Selector** drop-down list, select **Infrastructure Service**.
3. Navigate to **Configure > Policies**, and then click **Create Policy**.
4. Select **Storage**, and then click **Start**.
5. On the **General** page, configure the following parameters:

Property	Essential Information
Organization	Select the Organization.
Name	Enter a name for your policy.
Description (Optional)	Provide a short description
Add Tag (Optional)	Enter a tag in the key:value format. For example, Org: IT or Site: APJ.

6. On the **Policy Details** page, configure the following parameters:

Property	Essential Information
General Configuration	
Use JBOD Drives for Virtual Drive creation	Enable this option to use disks in JBOD state for creating virtual drives.
Unused Disks State	Select the state to which unused disks in this policy are to be moved. The state can be any one of UnconfiguredGood , or JBOD . Selecting No Change leaves the state unchanged.

Property	Essential Information
Default Drive Mode	<p>Select the default disk state that should be set on supported storage controller for newly inserted drives or on reboot. The state can be any one of UnconfiguredGood, JBOD, or RAID0.</p> <p>Unused Disks State should be No Change if Default Drive Mode is set to JBOD or RAID0.</p> <p>Note The default drive mode is supported only on M6 servers and for the following storage controllers.</p> <ul style="list-style-type: none"> • UCSC-RAID-M6T • UCSC-RAID-M6HD • UCSC-RAID-M6SD • UCSX-X10C-RAIDF <p>Configuration Limitations:</p> <ul style="list-style-type: none"> • When Default Drive State is JBOD or RAID0, then Unused Disks State should be No Change. • Use JBOD for VD creation cannot be enabled if Default Drive Mode is JBOD. • When Default Drive State is UnconfiguredGood, the drive state does not change on reboot. <p>Refer the table Default Drive Mode Scenarios for different Default Drive Mode scenarios.</p>
Secure JBOD Disk Slots	<p>Specify the JBOD drive slots that you want to encrypt. You may enter a comma or hyphen separated number range. For example: 1, 3 or 4-6, 8.</p>
M.2 RAID Configuration	<p>Enable this option to specify the Virtual Drive Name and Slot of the M.2 RAID controller for virtual drive creation.</p> <p>The disk slots used by the M.2 controller are automatically added.</p>

Property	Essential Information
Virtual Drive Name	<p>This field comes pre-filled with a default name. You can change it to your preferred name. A suffix will be added to your preferred name based on the selected controller slot.</p> <p>The name must be between 1 and 15 characters in length and can include letters, numbers, and the special characters hyphen (-), underscore (_), colon (:), and period (.).</p>
Slot of the M.2 RAID Controller for Virtual Drive Creation	<p>Select the slot of the M.2 RAID controller for virtual drive creation. The slots that can be selected are:</p> <ul style="list-style-type: none"> • MSTOR-RAID-1 — Select this option if there is only one M.2 RAID controller slot, or if there are two slots for the M.2 RAID controller and the virtual drive has to be created on the controller in the first slot. • MSTOR-RAID-2 — Select this option if there are two slots for the M.2 RAID controller and the virtual drive has to be created on the controller in the second slot. • MSTOR-RAID-1,MSTOR-RAID-2 — Select this option to create virtual drives on controllers in either or both slots.
Drive Group Configuration	<p>Enable to add RAID drive groups that can be used to create virtual drives. You can also specify the Global Hot Spares information.</p> <p>This configuration is not applicable for M.2 RAID controllers.</p>
Global Hot Spares	<p>Specify the disks that are to be used as hot spares, globally for all the RAID groups.</p> <p>The allowed value is a number range separated by a comma or a hyphen.</p>
Add Drive Group	Click to add a drive group.
Drive Group Name	<p>Enter the name of the drive group.</p> <p>The name must be between 1 and 15 characters in length and can include letters, numbers, and the special characters hyphen (-), underscore (_), colon (:), and period (.).</p>

Property	Essential Information
RAID Level	<p>The RAID level of a disk group describes how the data is organized on the disk group for the purpose of ensuring availability, redundancy of data, and I/O performance. The levels are:</p> <ul style="list-style-type: none"> • RAID0—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. • RAID1—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. • RAID5—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. • RAID6—Data is striped across all disks in the array and two sets of parity data are used to provide protection against failure of up to two physical disks. In each row of data blocks, two sets of parity data are stored. • RAID10—RAID 10 uses mirrored pairs of disks to provide complete data redundancy and high throughput rates through block-level striping. RAID 10 is mirroring without parity and block-level striping. A minimum of four disks are required for RAID 10. • RAID50—Data is striped across multiple striped parity disk sets to provide high throughput and multiple disk failure tolerance. • RAID60—Data is striped across multiple striped dual parity disk sets to provide high throughput and greater disk failure tolerance.
Secure Drive Group	Enable this option to configure encryption for drives that are part of the Virtual Drive.

Property	Essential Information
Number of Spans	<p>Number of span groups to be created for the RAID group. RAID levels with no nesting have a single span.</p> <p>Note Number of spans appears only when a RAID level with spans is selected.</p>
Drive Selection	
Drive Array Span 0	<p>Enter the drive array span. RAID levels RAID0, RAID1, RAID5, and RAID6 that do not have spans have only one disk group. RAID levels with spans have multiple disk groups with each disk group representing a span.</p> <p>RAID levels without spans have one span group and RAID levels with spans have two to eight span groups.</p> <p>Note If you have selected a RAID level without spans, then the field Drive Array Span 0 alone appears. If you have selected a RAID level with spans, you would have had to specify the number of spans. In this scenario, as many Drive Array Span fields as there are spans appear for you to specify the details.</p>
Dedicated Hot Spares	<p>Specify the collection of drives to be used as hot spares for this drive group.</p> <p>The allowed value is a number range separated by a comma or a hyphen.</p>
Add	Click Add to add the drive group.
Add Virtual Drive	
Drive Groups	Select the drive groups on which the virtual drive is to be created.
Number of Copies	Enter the number of copies of the virtual drive that is to be created. You can create a maximum of 10 copies.
Virtual Drive Configuration	
Virtual Drive Name	<p>Enter the name of the virtual drive.</p> <p>The name can be 1 to 15 characters long and can contain alphanumeric characters, and special characters '-' (hyphen), '_' (underscore), ':' (colon), and '.' (period).</p>

Property	Essential Information
Size (MiB)	Virtual drive size in MebiBytes. Size is mandatory except when the Expand to Available option is enabled.
Secured	Set this to enable encryption for the virtual drive. Note This option is not supported for UCS-M2-NVRAID (M.2 NVMe controller) as there are no SED drives that are supported on this controller.
RAID Type	Select the RAID type.
Expand to Available	Enable for the virtual drive to use all the space available in the disk group. When this flag is enabled, the size property is ignored.
Set as Boot Drive	Select to use this virtual drive as a boot drive. Note For standalone racks, you cannot set a drive, with a native block size of 4K, as the boot drive.
Strip Size	Select the strip size required. Allowed values are 64KiB, 128KiB, 256KiB, 512KiB, 1 MiB.
Access Policy	Select the type of access the host has to this virtual drive: <ul style="list-style-type: none"> • Read Write—Enables host to perform read-write on the virtual drive • Read Only—Host can only read from the virtual drive. • Blocked—Host can neither read nor write to the virtual drive.
Read Policy	Select the read ahead mode for this virtual drive: <ul style="list-style-type: none"> • Always Read Ahead • No Read Ahead

Property	Essential Information
Write Policy	Select the mode to be used to write to this virtual drive: <ul style="list-style-type: none"> • Write Through—Data is written through the cache and to the physical drives. Performance is improved, because subsequent reads of that data can be satisfied from the cache. • Write Back Good BBU—With this policy, write caching remains Write Back even if the battery backup unit is in good condition. • Always Write Back—Data is stored in the cache, and is only written to the physical drives when space in the cache is needed.
Disk Cache	Select the disk cache policy for this virtual drive. The values are: <ul style="list-style-type: none"> • Unchanged • Enabled • Disabled
Add	Click Add to add the virtual drive.
Single Drive RAID Configuration	Enable to create RAID0 virtual drives on each physical drive.
Drive Slots	Specify the set of drive slots where RAID0 virtual drives are to be created. <p>Note Single drive RAID allows you to add slots only where disks are planned to be inserted in future.</p>
Strip Size	Select the strip size required. Allowed values are 64KiB, 128KiB, 256KiB, 512KiB, 1MiB.
Access Policy	Select the type of access the host has to this virtual drive: <ul style="list-style-type: none"> • Read Write—Enables host to perform read-write on the virtual drive • Read Only—Host can only read from the virtual drive. • Blocked—Host can neither read nor write to the virtual drive.

Property	Essential Information
Read Policy	Select the read ahead mode for this virtual drive: <ul style="list-style-type: none"> • Always Read Ahead • No Read Ahead
Write Policy	Select the mode to be used to write to this virtual drive: <ul style="list-style-type: none"> • Write Through—Data is written through the cache and to the physical drives. Performance is improved, because subsequent reads of that data can be satisfied from the cache. • Write Back Good BBU—With this policy, write caching remains Write Back even if the battery backup unit is in good condition. • Always Write Back—Data is stored in the cache, and is only written to the physical drives when space in the cache is needed.
Disk Cache	Select the disk cache policy for this virtual drive. The values are: <ul style="list-style-type: none"> • Unchanged • Enabled • Disabled

Property	Essential Information
<p>Hybrid Slot Configuration</p>	<p>Select the following modes for server that supports Hybrid Drive Slots configuration:</p> <ul style="list-style-type: none"> • Direct Attached NVMe Slots—NVMe drives specified in the slot range will be moved to direct attached mode. • RAID Attached NVMe Slots—NVMe drives specified in the slot range will be moved to RAID attached mode. <p>Note</p> <ul style="list-style-type: none"> • NVMe Hybrid slots are supported only for UCSC-C240-M7, UCSC-C220-M7, and UCSC-C245-M8 servers in Standalone mode and Intersight Managed Mode. • Hybrid slots support is available for Slots 1–4 and Slots 101–104. • If an endpoint has Trimode 24G SAS RAID controller with PID UCSC-RAID-HP and Micron 7450 4GC cache Drive then the RAID attached NVMe slots can be used to create RAID configuration. • Combination of U.2 and U.3 drive PIDs are not recommended in the hybrid slots.

7. Click **Create**.



Note The Delete Virtual Drives option is not available in Storage Policy. Use the Storage Controllers page to delete virtual drives



Note Decommissioning or recommissioning operation will not delete the RAIDs or data on the disks.

The following table shows the behavior of Default Drive State in different scenario.

Table 6: Default Drive Mode Scenarios

Default Drive State	Host Reboot/ Host Boot	Hotplug	User Action (Service Profile deployment with Default Drive State)
UnconfiguredGood (OFF)	<ul style="list-style-type: none"> • All UnconfiguredGood drives remain UnconfiguredGood. • All previously converted JBOD continue to be JBOD. 	<ul style="list-style-type: none"> • Inserted drive remains UnconfiguredGood • JBOD from a different server remains UnconfiguredGood on this controller. 	<ul style="list-style-type: none"> • Setting UnconfiguredGood has no impact on the existing configuration. • Any JBOD device will remain as JBOD across controller boot. • Any UnconfiguredGood will remain UnconfiguredGood across controller boot.
JBOD	All unconfigured drives (non-user configured) are converted to JBOD.	Newly inserted unconfigured drive is converted to JBOD.	<ul style="list-style-type: none"> • All unconfigured drives (non-user configured drives) on the controller will be converted to JBOD. • User created UnconfiguredGood drive will remain UnconfiguredGood.

Default Drive State	Host Reboot/ Host Boot	Hotplug	User Action (Service Profile deployment with Default Drive State)
RAID0(RAID0 WriteBack)	All unconfigured drives will be converted to RAID0 WriteBack (WB). Note Unconfigured drives are the drives whose state remains unchanged by any user action.	Newly inserted unconfigured drive is converted to RAID0 WB.	<ul style="list-style-type: none"> All unconfigured drives (non-user created UnconfiguredGood) on the controller will be converted to RAID0 WriteBack (WB). User created UnconfiguredGood will remain UnconfiguredGood across controller reboot. Any RAID0 WriteBack device will remain as RAID0 WB across controller boot/reboot.



Note The Virtual Drives created by the system due to default drive state being **RAID0** will have **Server Profile Derived** as **No**.

The following table shows sample use cases for different Default Drive State scenarios.

Table 7: Various Drive Mode Use Cases

Use Case Scenario	Default Drive State
Using the server for JBOD Only (for example: Hyper converged, Hadoop data node and so on)	JBOD
Using the server for RAID volume (for example: SAP HANA database)	UnconfiguredGood
Using the server for Mixed JBOD and RAID volume	UnconfiguredGood
Using the server for per drive ROWB (for example: Hadoop data node)	RAID0 WriteBack

Creating a Syslog Policy

The Syslog policy defines the logging level (minimum severity) to report for a log file collected from an endpoint, the target destination to store the Syslog messages, and the Hostname/IP Address, port information, and communication protocol for the Remote Logging Server(s).

1. Log in to Cisco Intersight with your Cisco ID and select admin role.
2. From the **Service Selector** drop-down list, select **Infrastructure Service**.
3. Navigate to **Configure > Policies**, and then click **Create Policy**.
4. Select **Syslog**, and then click **Start**.
5. On the **General** page, configure the following parameters:

Property	Essential Information
Organization	Select the Organization.
Name	Enter a name for your policy.
Description (Optional)	Provide a short description
Add Tag (Optional)	Enter a tag in the key:value format. For example, Org: IT or Site: APJ.

6. On the **Policy Details** page, configure the following parameters:

Property	Essential Information
Local Logging	
Minimum Severity to Report	Select the lowest severity level to report in the remote log. The severity levels are: <ul style="list-style-type: none"> • 0 Emergency • 1 Alert • 2 Critical • 3 Error • 4 Warning • 5 Notice • 6 Informational • 7 Debug
Remote Logging - Syslog Server 1 and Syslog Server 2	

Property	Essential Information
Enable	Select this option to enable or disable the Syslog policy. Note When the Syslog Policy is created with Syslog Server 1 disabled and Syslog Server 2 enabled, it is observed that the Syslog server 1 always gets enabled first in the end point server.
Hostname/IP Address	Enter the hostname or IP address of the Syslog server to store the Cisco IMC log. You can set an IPv4 or IPv6 address or a domain name as the remote system address.
Port	Enter the destination port number of the Syslog server between 1 and 65535. The default port number is 514.
Protocol	Select the transport layer protocol for transmission of log messages to the syslog server. The options are: <ul style="list-style-type: none"> • TCP • UDP
Minimum Severity To Report	Select the lowest severity level to report in the remote log. The severity levels are: <ul style="list-style-type: none"> • 0 Emergency • 1 Alert • 2 Critical • 3 Error • 4 Warning • 5 Notice • 6 Informational • 7 Debug

7. Click **Create**.

Creating a Power Policy for Server

This policy enables the configuration of power redundancy, power profiling, and power restoration for servers.

1. Log in to Cisco Intersight with your Cisco ID and select admin role.
2. From the **Service Selector** drop-down list, select **Infrastructure Service**.
3. Navigate to **Configure > Policies**, and then click **Create Policy**.
4. Select **Power**, and then click **Start**.
5. On the **General** page, configure the following parameters:

Property	Essential Information
Organization	Select the Organization.
Name	Enter a name for your policy.
Set Tags (Optional)	Enter a tag in the key:value format. For example, Org: IT or Site: APJ.
Description (Optional)	Provide a short description

6. On the **Policy Details** page, navigate to **All Platforms** tab.
7. Configure the following parameters:

Property	Essential Information
Power Profiling	<p>Enables/disables the power profiling of the system</p> <p>Enabled—When enabled, it allows the CIMC to run power profiling utility during BIOS boot to determine the power needs of the server.</p> <p>Disabled—When disabled, power profiling is not run.</p> <p>Note This property is supported only on Cisco UCS X-Series servers.</p>
Power Priority	<p>Each server is assigned a power priority, which can be High, Medium, or Low. The power budgeted for the server depends on the power priority of the server. A server with higher priority gets a higher power budget. The default power priority of a server is Low.</p> <p>For more information on the firmware requirements for the Power Policy, see Firmware Requirements for Power Policy.</p>

Property	Essential Information
Power Restore	<p>Allows the user to configure the power restore state of the server on the CIMC. In the absence of IMM connectivity, the CIMC will use this policy to recover the host power after a power loss event.</p> <p>The power restoration states available are:</p> <ul style="list-style-type: none"> • Last State—Sets the host power to whatever state it was in before the power loss event. • Always On—Always power on the host after a power loss event. • Always Off—Always keep the host power off after a power loss event. <p>Firmware Requirements for Power Policy</p>

8. Click **Create**.

Creating a Thermal Policy for Server

This policy enables controlling the speed of the chassis fan.

1. Log in to Cisco Intersight with your Cisco ID and select admin role.
2. From the **Service Selector** drop-down list, select **Infrastructure Service**.
3. Navigate to **Configure > Policies**, and then click **Create Policy**.
4. Select **Thermal**, and then click **Start**.



Note Thermal Policy is not supported for Cisco UCS Standalone M4 servers, Cisco UCS B-Series servers, and Cisco UCS X-Series servers.

5. On the **General** page, configure the following parameters:

Property	Essential Information
Organization	Select the Organization.
Name	Enter a name for your policy.
Set Tags (Optional)	Enter a tag in the key:value format. For example, Org: IT or Site: APJ.
Description (Optional)	Provide a short description.

6. On the **Policy Details** page, configure the following parameters:

Property	Essential Information
Fan Control Mode controls the fan speed of the chassis.	
Balanced	The fans run faster when needed based on the heat generated by the server. When possible, the fans return to the minimum required speed.
Low Power	The fans run at slightly lower minimum speeds than the Balanced mode, to consume less power when possible.
High Power	The fans are kept at higher speed to emphasize performance over power consumption. Note This mode is supported for all Cisco UCS C-Series servers.
Maximum Power	The fan is always kept at the maximum speed. This option provides the most cooling and consumes most power. Note This mode is supported for all Cisco UCS C-Series servers.
Acoustic	The fan speed is reduced to reduce noise levels in acoustic-sensitive environments. Note This mode is supported for all Cisco UCS C-Series servers.

7. Click **Create**.

Creating vNIC or vHBA Templates

Creating vNIC or vHBA Templates

A vNIC or vHBA template consists of common configurations that you can reuse across multiple vNICs or vHBAs, used in various Server Profiles. This approach simplifies network configuration across multiple servers. You can create vNICs or vHBAs from the template using the **Derive** operation while creating the policy. Additionally, you can attach an existing vNIC or vHBA to a template to utilize the configurations set in the template. These templates can be created with or without override options. The override option allows the configuration of the derived vNIC or vHBA to override the template configuration.

To create a vNIC or vHBA template:

1. Log in to Cisco Intersight.
2. Navigate to the **Templates** tab, and click one of the following:
 - To create a vNIC template, click **Create vNIC Template**.

- To create a vHBA template, click **Create vHBA Template**.
3. On the **General** page:
 - a. Choose an Organization for the template from the list. This field supports the capability of configuration sharing across Organizations.
 - b. Enter a name for the template.
 - c. Enter a Tag for the template. Tags must be in the key:value format. For example, Org: IT or Site: APJ.
 - d. Enter a description to help identify the template.
 - e. Click **Next**.
 4. On the **Configuration** page:
 - a. If you want to allow the configuration of the derived vNIC or vHBA to override the template configuration, select the **Allow Override** checkbox.



Note Parameters that can be overridden are indicated by an **Override Allowed** label.

- b. Configure the template properties as required.
- c. Click **Create** to create the template.



-
- Note**
- You cannot modify a template to remove mandatory configurations when there are active derived vNICs or vHBAs with override enabled.
 - If there are active overridden properties, you cannot disable Override option in the template. It is required to detach the derived vNIC or vHBA from the template usage page and then attempt to disable the Override option in the template.
-

Deriving a vNIC from a template is done as part of creating the LAN Connectivity Policy. Similarly, deriving a vHBA from a template is done as part of creating the SAN Connectivity Policy. For more information, see *Creating a LAN Connectivity Policy* and *Creating a SAN Connectivity Policy*.

Tutorial 1: Working with a vNIC template with active derived vNICs, when Override is enabled

Consider the following scenario after a vNIC has been derived from a template and attached to a LAN Connectivity policy:

1. Do the following in the vNIC template:
 - a. Enable the **Override** option.
 - b. Modify **Failover** to **Enabled**. Note that the **Override** option is not available for **Failover**; hence, this property change will be propagated to the vNICs derived from this template.
 - c. Retain **MAC Pool** as is with no change. Note that the **Override** option is available for **MAC Pool**.

- d. Create an **Ethernet Adapter** policy and attach it to the template. Note that the **Override** option is available for **Ethernet Adapter**. The **Ethernet Adapter** policy is indicated with the **Overridden** label, which means that the configuration of the derived vNICs will override the configuration propagated from the template.
2. Review the modifications in the LAN Connectivity policy that utilizes the vNIC derived from the template:
 - a. **Failover** is propagated from the template and modified to **Enabled**.
 - b. **MAC Pool** is propagated from the template. Even though **Override** is allowed for MAC Pool, it does not display the **Overridden** label since the configuration remains consistent with the template. If you modify **MAC Pool** now, as the **Override** option is available for this property, the new configuration is applied to the vNIC instance and the **Overridden** label is displayed.
 - c. Create or attach a different **Ethernet Adapter** policy to the derived vNIC of the LAN connectivity policy. The **Ethernet Adapter** policy is marked with the **Overridden** label. The **Ethernet Adapter** policy is not propagated from the template to the derived vNIC in the LAN Connectivity policy.

Working with a vHBA template with active derived vHBA, when Override is disabled

Consider the following scenario after a vHBA has been derived from a template and attached to a SAN Connectivity policy:

1. Maintain the **Override** option as **Disabled** in the vHBA template.
2. Review the SAN Connectivity policy that utilizes the vHBA derived from the template:
 - Since override is not enabled, the parameters inherited from the template can only be viewed and not be modified.
 - Only those parameters can be modified that are excluded from the template. For example, under **Placement**, **Switch ID** is part of the template, it can only be viewed. The other parameters, which are not included in the template, can be modified.
3. To make modifications to the parameters included in the template, you must enable **Override** in the template, and then retry overriding the configuration in the SAN Connectivity policy.



CHAPTER 10

Configuring UCS Chassis Policies

- [Chassis Policies](#), on page 259
- [Creating an IMC Access Policy](#), on page 260
- [Creating an SNMP Policy](#), on page 261
- [Creating a Power Policy for Chassis](#), on page 263
- [Creating a Thermal Policy](#), on page 266

Chassis Policies

Chassis policies in Cisco Intersight allow you to configure various parameters of the chassis, including IP pool configuration, VLAN settings, SNMP authentication, and SNMP trap settings. A chassis policy can be assigned to any number of chassis profiles to provide a configuration baseline for a chassis.

To view the Chassis Policies table view, from the **Service Selector** drop-down list, choose **Infrastructure Service**. Navigate to **Configure > Policies**.

The Chassis Policy creation wizard in Cisco Intersight has two pages:

- **General**—The general page allows you to select the organization and enter a name for your policy. Optionally, include a short description and tag information to help identify the policy. Tags must be in the key:value format. For example, Org:IT or Site APJ
- **Policy Details**—The policy details page has properties that are applicable to UCS Chassis Policies.

Chassis Policies can also be cloned by using the **Policy Clone** wizard with properties that are similar to the existing policies. The clone policy action is available on both the policies list and detailed views. For more information, see [Cloning a Policy](#).

The following list describes the chassis policies that you can configure in Cisco Intersight.

- **IMC Access Policy**—Enables you to configure and manage your network by mapping the IP pools to the chassis profile. This policy allows you to configure a VLAN and associate it with an IP address using the IP pool.



Note Only In-Band configuration is supported for Chassis IMC Access Policy.

- **SNMP Policy**—Configures the SNMP settings for sending fault and alert information by SNMP traps from the managed devices. SNMP Users or SNMP Traps configured previously on the managed devices

are removed and replaced with users or traps that you configure in this policy. If you have not added any users or traps in the policy, the existing users or traps on the input/output module (IOM) are removed.

- **Power Policy**—Enables the management of power usage for the chassis. This policy allows you to configure the redundancy mode of the Chassis Power Supply Units (PSUs) and allocate power to the chassis. You can view the redundancy health, redundancy mode, input power health, and output power health of the chassis in the properties section of the **General** tab on the Chassis details view page. For Cisco UCS X9508 Chassis, you can configure Power Save Mode and Dynamic Power Reallocation.
- **Thermal Policy**—Allows the user to set the value of the Fan Control Mode for the chassis. The Fan Control Mode controls the speed of the chassis fan to maintain optimal server cooling.

Creating an IMC Access Policy

IMC Access policy allows to provide a VLAN ID and enables to associate it with an IP address from the selected IP pool.

1. Log in to Cisco Intersight with your Cisco ID and select admin role.
2. From the **Service Selector** drop-down list, select **Infrastructure Service**.
3. Navigate to **Configure > Policies**, and then click **Create Policy**.
4. Select **IMC Access**, and then click **Start**.
5. In the **General** page, configure the following parameters:

Property	Essential Information
Organization	Select the organization.
Name	Enter a name for your policy.
Tag (optional)	Enter a tag in the key value format.
Description (optional)	Enter a short description.

6. On the **Policy Details** page, configure the following parameters:

Property	Essential Information
VLAN ID	Enter the VLAN ID to be used for server access over the inband network. The field value can be between 4 and 4093.
IPv4 address configuration	Select to determine the type of network for this policy. Note You can select only IPv4 address configuration or both IPv4 and IPv6 configurations.
IPv6 address Configuration	Select to determine the type of network for this policy. You can select only IPv6 address configuration or both IPv4 and IPv6 configurations. Important IPv6 is supported only on UCS-IOM-2408

Property	Essential Information	
IP Pool	Select IP Pool	Click to view and select the IP pool list on the right pane.

- Click **Create**.



Note The IMC Access policy is not supported for Cisco UCS X-Series Direct chassis (UCSX-9508) currently.

Creating an SNMP Policy

The SNMP policy configures the SNMP settings for sending fault and alert information by SNMP traps from the managed devices. This policy supports SNMP versions such as SNMPv1, SNMPv2(includes v2c), and SNMPv3. Any existing SNMP Users or SNMP Traps configured previously on the managed devices are removed and replaced with users or traps that you configure in this policy. If you have not added any users or traps in the policy, the existing users or traps on the input/output module (IOM) are removed.

Using the SNMP Policy you can enable or disable SNMP, specify the access and community strings, and provide the SNMP user details that is used to retrieve data.

- Log in to Cisco Intersight with your Cisco ID and select admin role.
- From the **Service Selector** drop-down list, select **Infrastructure Service**.
- Navigate to **Configure > Policies**, and then click **Create Policy**.
- Select **SNMP**, and then click **Start**.
- In the **General** page, configure the following parameters:

Property	Essential Information
Organization	Select the organization.
Name	Enter a name for your policy.
Tag (optional)	Enter a tag in the key value format.
Description (optional)	Enter a short description.

- In the **Policy Details** page, configure the following parameters:

Property	Essential Information
Enable SNMP	Displays the state of the SNMP Policy on the endpoint. Enable this option for the endpoint to send SNMP traps to the designated host.

Property	Essential Information
Access Community String	Enter the SNMPv1, SNMPv2 community string or the SNMPv3 username. This field allows maximum of 18 characters. Note If the field is empty, it indicates that the SNMPv1 and SNMPv2c users are disabled.
Trap Community String	Enter the SNMP community group name used for sending SNMP trap to other devices. Note This field is applicable only for SNMPv2c trap host or destination.
SNMP Users	
Name	Enter the SNMP username. This field must have a minimum of 1 and a maximum of 31 characters.
Security Level	Select the security mechanism for communication between the agent and the manager that include: <ul style="list-style-type: none"> • AuthPriv • AuthNoPriv
Auth Type	Select SHA as the authorization protocol for authenticating the user Note The MD5 authorization protocol is not supported.
Auth Password	Enter the authorization password for the user.
Auth Password Confirmation	Enter the authorization password confirmation for the user.
Privacy Type	Select AES as the privacy protocol for the user.
Privacy Password	Enter the privacy password for the user.
Privacy Password Confirmation	Enter the privacy password confirmation for the user.
SNMP Trap Destinations	
Enable	Enable this option to allow and deploy the SNMP policy.
SNMP Version	Select v2 or v3 as the SNMP version for the trap.

Property	Essential Information
User	Select the SNMP user for the trap. You can define maximum of 15 trap users. Note This field is applicable only to SNMPv3.
Trap Type	Select the trap type to receive a notification when a trap is received at the destination: <ul style="list-style-type: none"> • Trap • Inform
Destination Address	Provide the address to which the SNMP trap information can be sent. You are allowed to define maximum of 15 trap destinations.
Port	Enter the port number for the server to communicate with trap destination. The range is from 1 to 65535. The default is 162.

7. Click **Create**.



Note SNMP policy is not supported for Cisco UCS X-Series Direct chassis (UCSX-9508) currently.

Creating a Power Policy for Chassis

This policy enables configuration of power redundancy and power allocation for chassis.

1. Log in to Cisco Intersight with your Cisco ID and select admin role.
2. From the **Service Selector** drop-down list, select **Infrastructure Service**.
3. Navigate to **Configure > Policies**, and then click **Create Policy**.
4. Select **Power**, and then click **Start**.
5. On the **General** page, configure the following parameters:

Property	Essential Information
Organization	Select the Organization.
Name	Enter a name for your policy.
Set Tags (Optional)	Enter a tag in the key:value format. For example, Org: IT or Site: APJ.
Description (Optional)	Provide a short description

6. On the **Policy Details** page, navigate to **UCS Chassis** tab.
7. Configure the following parameters:

Property	Essential Information
Power Redundancy sets the redundancy mode of the chassis power supplies.	
Grid	Grid mode requires two power sources. If one source fails, the surviving power supplies on the other source provides power to the chassis.
Not Redundant	Power Manager turns on the minimum number of PSUs required to support chassis power requirement. No redundant PSUs are maintained.
N+1	Power Manager turns on the minimum number of PSUs required to support chassis power requirements and one additional PSU for redundancy.
N+2	Power Manager turns on the minimum number of PSUs required to support chassis power requirements and two additional PSUs for redundancy. Note This mode is supported only for Cisco-UCSX-9508 chassis.
Power Save Mode	Enable to place additional PSU capacity in Power Save mode, when the requested power is less than the available power. Note This property is supported on: <ul style="list-style-type: none"> • Cisco-UCSX-9508 chassis with the minimum Cisco IMC firmware version of 4.2(1d). • Cisco-UCSB-5108 chassis with the minimum Cisco IMC firmware version of 4.3(2a).

Property	Essential Information
Dynamic Power Rebalancing	<p>Enable for dynamically reallocating power for the servers.</p> <p>When enabled, the power will be rebalanced across various chassis components including blades, Fans, IOMs/IFMs, and XFMs.</p> <p>Note This property is supported on:</p> <ul style="list-style-type: none"> • Cisco-UCSX-9508 chassis with the minimum Cisco IMC firmware version of 4.2(1d). • Cisco-UCSB-5108 chassis with the minimum Cisco IMC firmware version of 4.3(2a).
Extended Power Capacity	<p>Sets the Extended Power Capacity of the Chassis. When this mode is enabled, power is borrowed from the redundant power supplies which increases the power available to the chassis.</p> <p>Note This property is supported only on Cisco-UCSX-9508 chassis with the minimum Cisco IMC firmware version of 4.2(1d).</p>
Power Allocation (Watts)	<p>Allows the user to set the maximum power a chassis can consume.</p> <p>The value can range from minimum of system requirement to maximum of available power.</p> <p>Deploying a policy with a Power Allocation of 0 will uncap the chassis budget, that is, the chassis will be able to consume all of the available power.</p> <p>Note This property is supported on:</p> <ul style="list-style-type: none"> • Cisco-UCSX-9508 chassis with the minimum Cisco IMC firmware version of 4.2(1d). • Cisco-UCSB-5108 chassis with the minimum Cisco IMC firmware version of 4.3(2a).

8. Click **Create**.

Creating a Thermal Policy

This policy enables controlling the speed of the chassis fan.

1. Log in to Cisco Intersight with your Cisco ID and select admin role.
2. From the **Service Selector** drop-down list, select **Infrastructure Service**.
3. Navigate to **Configure > Policies**, and then click **Create Policy**.
4. Select **Thermal**, and then click **Start**.
5. On the **General** page, configure the following parameters:

Property	Essential Information
Organization	Select the Organization.
Name	Enter a name for your policy.
Set Tags (Optional)	Enter a tag in the key:value format. For example, Org: IT or Site: APJ.
Description (Optional)	Provide a short description

6. On the **Policy Details** page, configure the following parameters:

Property	Essential Information
Fan Control Mode	controls the fan speed of the chassis.
Balanced	The fans run faster when needed based on the heat generated by the server. When possible, the fans return to the minimum required speed.
Low Power	The fans run at slightly lower minimum speeds than the Balanced mode, to consume less power when possible.
High Power	The fans are kept at higher speed to emphasize performance over power consumption. Note This mode is supported only for UCS X-Series chassis.
Maximum Power	The fan are always kept at the maximum speed. This option provides the most cooling and consumes most power. Note This mode is supported only for UCS X-Series chassis.

Property	Essential Information
Acoustic	The fan speed is reduced to reduce noise levels in acoustic-sensitive environments. Note This mode is supported only for UCS X-Series chassis.

7. Click **Create**.



CHAPTER 11

Configuring Pools

- [Pools, on page 269](#)
- [Identity \(ID\) Pools, on page 269](#)
- [Pool Allocation, on page 270](#)
- [Deleting Pools, on page 270](#)
- [Identity Retention, on page 271](#)
- [IP Pools, on page 272](#)
- [MAC Pools, on page 275](#)
- [UUID Pools, on page 277](#)
- [WWN Pools, on page 279](#)
- [IQN Pools, on page 282](#)
- [Resource Pools, on page 285](#)
- [Virtual Routing and Forwarding, on page 288](#)

Pools

Pools are the basic building blocks for uniquely identifying hardware resources. They form the foundation of the UCS management model, enabling the association of Server Profiles with blade servers while maintaining the same ID and presentation to the upstream LAN or SAN.

Pools are classified into Resource Pools and Identity (ID) Pools.

Using the Pools Table View, you can monitor the utilization of Server Pools, track *Available* and *Used* identifiers, and make informed decisions regarding pool capacity and allocation.

Identity (ID) Pools

ID Pools are further classified into the following categories:

- **IP pool:** Provides the flexibility of dynamically assigning IP addresses to services running on a network element.
- **MAC address pool:** Provides unique IDs for network interface ports.
- **UUID pool:** Provides unique IDs for each server associated with the server profile.

- **WWNN and WWPN pool:** Provides unique IDs for Fibre Channel resources on a server (Fibre Channel nodes and ports).
- **IQN pool:** Provides a collection of iSCSI Qualified Names (IQNs) for use as initiator identifiers by iSCSI vNICs.

Pool Allocation

To ensure consistent and repeatable ID allocation across multiple allocation iterations, the smallest available ID is allocated sequentially during pool allocation iterations.

For example, consider a pool with a range of 1 to 20 IDs. The following table describes the ID allocation and reallocation iterations.

Use Case	Behavior
5 IDs are requested from the pool	IDs 1 to 5 are allocated
1-5 IDs are released	IDs 1 to 5 are released
5 IDs are requested from the pool	IDs 1 to 5 are allocated

You can have overlapping pools, where some IDs are shared between different pools across organizations. Since IDs are unique across the account, whether they are used in different organizations or not, if an ID is consumed in one pool, it will also be marked as **Used** in all pools where it is used. The Pools Table View provides a summary of allocated IDs, providing the capability to track ID usage in Pools across Organizations. The Source column in the Pools Table View indicates the pool from which the ID is being allocated. If the ID is being allocated from the current pool, the Source column will be marked as **Self**. If the ID is allocated from another pool, the Source column will be marked as **Other**.

Deleting Pools

You can delete a pool or address block in an Organization where none of the IDs are allocated. It will not impact allocations in any pool in another Organization.

-
- Step 1** Check if the ID is not currently associated with a Server Profile:
- In the **Pools Table View**, review the **Source** column to analyze the pool usage.
 - If **Source** is marked as **Other**, you can proceed to delete the ID as it is allocated from some other overlapping pool.
 - Do not attempt to delete pools marked as **Source = Self** as it is not allowed since they are in use in one of the profiles.
- Step 2** Click **Delete**.
-

Identity Retention

IP Addresses, MAC addresses, IQNs, UUIDs, WWNNs, and WWPNs are the typical identifiers that a physical server gets from a Server Profile.

Intersight uses best efforts to retain the LAN Connectivity Policy (LCP) identifiers, such as MAC, IQN, and iSCSI IP, as well as the SAN Connectivity Policy (SCP) identifiers, such as WWPN and WWNN, when modifying policies, profiles or templates.

If you change the LCP or SCP to a new policy that accesses different Pools with non-overlapping IDs, then expect all the IDs to change. Furthermore, expect changes if the new Pool does not have the exact ID available.

In the following scenarios, you can expect ID retention during edits or changes:

- When adding a vNIC or vHBA to an LCP or SCP.
- When changing policy LCP1/SCP1 to LCP2/SCP2 that uses the same Pool Reference.
- When changing policy LCP1/SCP1 to LCP2/SCP2 that uses a different Pool Reference, but with the same IDs available.
- When changing policy LCP1/SCP1 that uses Static Identifiers to LCP2/SCP2 that uses a Pool Reference with the same IDs available.
- When detaching a Server Profile from Template T1 and attaching the Server Profile to Template T2 with the same IDs available.
- When editing a Server Profile Template and changing LCP1/SCP1 to LCP2/SCP2 as above.
- When editing an existing LCP/SCP Policy and changing the Identifier Reference from Static to a Pool with the same IDs available.

Identity Reservation

Identities can be reserved prior to allocation to allow selecting a specific value from a pool, for purposes such as migration across environments. For example, Cisco UCSM to IMM.

Reserved Identifiers Guidelines

- The identifiers can be reserved only via the [IMM Transition Tool](#) or by using available pool Reservation APIs, such as <https://intersight.com/apidocs/apirefs/macpool/Reservations/model/>.
- The reservation of the identifiers can be done only for Fabric Interconnect-attached servers.
- Reserved identifiers are meant for one-time use and are removed from the reservation pool once consumed.

The reserved ID gets consumed either when a policy (with the reserved identifier) is attached to a server profile, or when the server profile is deployed.

The **Reserved Identifiers** tab shows the list of reserved identifier values, their type, and the corresponding Pool membership. The Pool membership appears blank when the allocation type is Static. You can select and delete any reserved identifier.

IP Pools

An IP Pool can contain one or more blocks of IPs that will get consumed in sequential order, beginning with the lowest block. IP pools support both IPv4 and IPv6 addresses.

Subnet Configuration in an IP Pool

You can create IP pools with either common subnet configurations for all IP blocks (pool-level) or different subnet configurations for each IP block (block-level).

After defining a pool with subnet configuration at the block level, you can migrate it to pool-level subnet configuration and vice versa. When migrating from a pool-level to a block-level subnet configuration, the subnet configuration will be replicated across each existing IP block. When migrating from block-level to pool-level subnet configuration, you need to reconfigure the common subnet settings at the pool-level.

If an IP block already has existing leases, migration is allowed only in the following scenarios:

- **Migrating from Pool-Level to Block-Level Configuration with Existing Leases:**

When you migrate from pool-level to block-level configuration with existing leases, the subnet configuration is moved to block-level without any changes. This means that the same subnet configuration previously set at the pool level is copied to each block. In such cases, the migration is allowed even if there are existing leases. After migration, if you observe that you cannot modify the subnet configuration of any block, it could be because it has existing active leases. Note that you cannot change the subnet configuration of any block if it already has active leases.

- **Migrating from Block-Level to Pool-Level Configuration with Existing Leases:**

When you migrate from block-level to pool-level configuration with existing leases, you must specify the subnet configuration at the pool level. If all the previous block-level subnet configurations are the same as the new pool-level subnet configuration, the migration is allowed. In this scenario, the migration is permitted even if there are existing leases.

Creating an IP Pool

IP Pools represent a collection of IP addresses that can be allocated to configuration entities such as server profiles. You can create IPv4 pool or IPv6 pool or both.

Step 1 From the left navigation panel, click **Create Pools > IP > Start**.

The **IP Pool** wizard appears.

Step 2 Add the following information on the **General** page:

- **Organization**—The organization of the IP pool.
- **Name**—Name of the IP pool.
- **Add Tag**—The tag to identify and search for the IP pool.
- **Description**—The description the IP pool.
- **Configure Subnet at Block Level**—Select the checkbox to enable subnet configurations for each IP block within IPv4 and IPv6 pools.

Step 3 Click **Next**.

Step 4 [Optional] Configure IPv4 pools:

- a) Use the **Configure IPv4 Pool** toggle button to enable IPv4 pool configuration. By default, it is enabled. You can opt to configure the IPv4 pool later.
- b) If you have opted to configure the **Netmask**, **Gateway**, **Primary DNS**, and **Secondary DNS** fields at the pool-level, enter these details under **Configuration**. If you have opted to configure these fields at the block-level, enter these details while configuring the IP block.
- c) Under **IP Blocks**, configure one or more IP blocks:
 1. Click **Add IP Blocks** to add an IP block.
 2. Enter the following parameters for the IP block:

Note You can configure the Netmask, Gateway, Primary DNS, and Secondary DNS fields at the pool-level or the block-level.

- **From**—Starting IP address of the IP pool.
- **Size**—Number of IP addresses allocated for the IP pool.
- **Netmask**—the netmask associated with the IP pool.
- **Gateway**—The IP address of the gateway for the IP pool.

Note If the IP Pool is to be used for an IMC Access policy, ensure that the gateway IP address specified in the IP Pool has connectivity to Cisco IMC.

- **Primary DNS**—the primary DNS server that this block of IP addresses should access.
- **Secondary DNS**—the secondary DNS server that this block of IP addresses should access.

Step 5 [Optional] Configure IPv6 pools:

- a) Use the **Configure IPv6 Pool** toggle button to enable IPv6 pool configuration. By default, it is enabled. You can also opt to configure an IPv6 pool later.
- b) If you have opted to configure the **Prefix**, **Gateway**, **Primary DNS**, and **Secondary DNS** fields at the pool level, enter these details under **Configuration**. If you have opted to configure these fields at the block level, enter these details while configuring the IP block.
- c) Under **IP Blocks**, configure one or more IP blocks:
 1. Click **Add IP Blocks** to add an IP block.
 2. Enter the following parameters for the IP block:

Note You can configure the Prefix, Gateway, Primary DNS, and Secondary DNS fields at the pool-level or the block-level.

- **From**—Starting IP address of the IP pool.
- **Size**—Number of IP addresses allocated for the IP pool.
- **Prefix**: The prefix associated with the IP pool.
- **Gateway**: The IP address of the gateway for the IP pool.

Note If the IP Pool is to be used for an IMC Access policy, ensure that the gateway IP address specified in the IP Pool has connectivity to Cisco IMC.

- **Primary DNS**—the primary DNS server that this block of IP addresses should access.
- **Secondary DNS**—the secondary DNS server that this block of IP addresses should access.

Step 6 Click **Create**.

The newly created IP pool appears in the list of IP pools.

IP Pool Details

Details

Displays the list of IP pools.

Property	Essential Information
Details	
Name	Displays the name of the IP pool
Type	Displays the type of the pool.
Size	Displays the total number of identifiers the IP pool contains.
Used	Displays the total number of identifiers in the IP pool that are in use and are no longer available.
Reserved	Displays the total number of identifiers in the IP pool that are reserved for later use.
Available	Displays the total number of identifiers in the IP pool that are available for use.
Description	A description of the IP pool.
Last Update	The date and time when the IP pool was last updated.
Organization	Users in a Default Organization automatically has access to all the resources available for the user account.
Configuration	
IPv4	Displays the IPv4 configuration of the pool, such as subnet mask, default gateway, primary DNS, and secondary DNS, when the subnet is configured at the pool level.
IPv6	Displays the IPv6 configuration of the pool, such as prefix, default gateway, primary DNS, and secondary DNS, when the subnet is configured at the pool level.

Property	Essential Information
From	Displays the starting IP of the pool. Note Cisco Intersight selects the identity in sequential manner, that is, the lowest available identity from the pool.
To	Displays the range of the block size. Note This value is dependent on the IP pool size property.
Size	Displays the IP pool size
Eye symbol	Displays the configuration of the pool, such as subnet mask, prefix, default gateway, primary DNS, and secondary DNS, when the subnet is configured at the block level.
Usage	
IP, VRFs, Status, Server Profile, and Source	Displays the IP address, VRF instances, status of usage (Reserved or Used) and associated server profiles. Source can be Self or Other , where Self is ID used or reserved by this pool and Other is ID used or reserved statically or by another pool.
Actions	
Edit	Allows to add or modify the configuration details of the IP pool.
Delete	Allows to delete the IP pool.

MAC Pools

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

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

Creating a MAC Pool

MAC Pools represent a collection of MAC addresses that can be allocated to vNICs of a server profile.

-
- Step 1** From the left navigation panel, click **Pools > MAC > Create MAC Pool**.
The **MAC Pool** wizard appears.
- Step 2** Add the following information on the **General** page:
- **Name**—Name of the MAC pool
 - **Description**—An optional description of the MAC pool.
 - **Organization**—The organization to which the MAC pool belongs.
 - **Add Tag**—An optional tag to identify and search for the MAC pool.
- Step 3** Click **Next**. The **Pool Details** page appears.
- Step 4** Add the following configuration information in the **MAC Blocks** area:
- **From**—Indicates the first MAC address in the block.
 - **Size**—Indicates the number of MAC addresses in the block.
- Step 5** To add more blocks, click + and then add the starting MAC address and total number of MAC addresses in the new block.
- Step 6** Click **Create**.
-

The newly created MAC pool appears in the list of MAC pools.

MAC Pool Details

Details

Displays the list of MAC pools.

Property	Essential Information
Name	The name of the MAC pool.
Size	The number of MAC addresses in the pool.
Used	The number of MAC addresses in the pool that have been used, and are no longer available.
Reserved	Displays the total number of identifiers in the MAC pool that have been reserved to be used later.
Available	Displays the total number of identifiers in the MAC pool that are available to be used.
Description	A description of the MAC pool.
Last Update	When the MAC pool was last updated.
Configuration	

Property	Essential Information
From	Displays the MAC prefix value of the pool. Note Cisco Intersight selects the identity in sequential manner, that is, the lowest available identity from the pool.
To	Displays the MAC suffix value of the pool.
Size	Displays the MAC pool size
Usage	
MAC address, Status, Server Profile, and Source	Displays the MAC address, status of usage (Reserved or Used) and associated server profiles. Source can be Self or Other , where Self is ID used or reserved by this pool and Other is ID used or reserved statically or by another pool.
Actions	
Edit	Allows to add or modify the configuration details of the MAC pool.
Delete	Allows to delete the MAC pool.

UUID Pools

A Universally Unique Identifier (UUID) pool is a collection of UUIDs that are assigned to servers. The prefix and suffix of the UUID are variable values. A UUID pool ensures that these variable values are unique for each server associated with a server profile that uses a particular pool to avoid conflicts.



- Note**
- The supported servers and its minimum firmware or Cisco IMC versions required for UUID pool are mentioned below:

Servers	Minimum firmware versions
Cisco UCS-B200-M5, UCS-B480-M5, Cisco UCS UCS-B200-M6	4.2(1b)
Cisco UCS-C220-M6, UCS-C240-M6	4.2(1b)
Cisco UCS-C225-M6, UCS-C245-M6	4.2(1i)
Cisco UCSX-210C-M6	5.0(1a)

- For more information on the server profile association using UUID pool, see [Configuring Server Profiles](#).

Creating a UUID Pool

UUID Pools represent a collection of UUID items that can be allocated to server profiles.

-
- Step 1** From the left navigation panel, click **Pools > UUID > Create UUID Pool**.
The **UUID Pool** wizard appears.
- Step 2** Add the following information on the **General** page:
- **Organization**—An organization to which the UUID pool belongs.
 - **Name**—Name of the UUID pool.
 - **Set Tags**—An optional tag to identify and search for the UUID pool.
 - **Description**—An optional description of the UUID pool.
- Step 3** Click **Next**. The **Pool Details** page appears.
- Step 4** In the **Configuration** section, add the UUID Prefix number in hexadecimal format. Example, 1728E8C7-7B40-47E8
- Step 5** In the **UUID Blocks** section, add the following configuration details:
- **From**—Indicates the UUID suffix number of the block in hexadecimal format. Example, 9EDE-0E52924AC87A
 - **Size**—Indicates the number of UUID identifiers in the block. The size ranges from 1 to 1000.
- Step 6** To add more blocks, click + and then add the starting UUID suffix and total number of UUID identifiers in the new block.
- Step 7** Click **Create**.
-

The newly created UUID pool appears in the list of UUID pools.

UUID Pool Details

Details

Displays the list of UUID pools.

Property	Essential Information
Details	
Name	Displays the name of the UUID pool
Type	Displays the type of the pool.
Size	Displays the total number of identifiers the UUID pool contains.
Used	Displays the number of UUID already in use from the pool.
Reserved	Displays the total number of UUID that have been reserved to be used later.
Available	Displays the number of UUID available for usage.

Property	Essential Information
Last Update	The date and time when the when the UUID pool was last updated.
Description	Description of the UUID Pool.
Organization	Displays the organization under which the UUID Pool is created.
Configuration	
UUID Prefix	Displays the UUID prefix value of the pool.
From	Displays the UUID suffix value of the pool. Note Cisco Intersight selects the identity in sequential manner, that is, the lowest available identity from the pool.
To	Displays the range of the block size. Note This value is dependent on the UUID pool size property.
Size	Displays the UUID pool size
Usage	
UUID, Status, Server Profile, and Source	Displays the UUID assigned to the server profile, status of usage (Reserved or Used) and the associated server profile. Source can be Self or Other , where Self is ID used or reserved by this pool and Other is ID used or reserved statically or by another pool.

WWN Pools

A World Wide Name (WWN) pool is a collection of WWNs for use by the Fibre Channel vHBAs in a Cisco UCS Domain. You create separate pools for the following:

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



Note A WWN ID can not be reused across WWPN and WWNN pools. To ensure the uniqueness of the Cisco UCS WWNNs and WWPNS in the SAN fabric, Cisco Intersight uses the following WWN prefix for all blocks in a pool: 20:00:00:25:B5:xx:xx:xx.

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

WWNN Pools

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

WWPN Pools

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

Creating a WWNN Pool

To ensure the uniqueness of the Cisco UCS WWNNs and WWPNs in the SAN fabric, Cisco Intersight uses the following WWN prefix for all blocks in a pool: 20:00:00:25:B5:xx:xx:xx.

-
- Step 1** From the left navigation panel, click **Pools > WWNN > Create WWNN Pool**.
The **WWNN Pool** wizard appears.
- Step 2** Add the following information on the **General** page:
- **Name**—Name of the WWNN pool
 - **Description**—An optional description of the WWNN pool.
 - **Organization**—An optional entry of the organization to which the WWNN pool belongs.
 - **Add Tag**—An optional tag to identify and search for the WWNN pool.
- Step 3** Click **Next**. The **Pool Details** page appears.
- Step 4** Add the following configuration information in the **Initiator Blocks** area:
- **From**—Indicates the first WWN identifier of the block.
 - **Size**—Indicates the maximum number of identifiers that the block can contain.
- Step 5** To add more blocks, click + and then add the starting WWN identifier and maximum number of identifiers that the block can contain.
- Step 6** Click **Create**.
-

The newly created WWNN pool appears in the list of WWNN pools.

WWNN Pool Details

Details

Displays the list of WWNN pools. To ensure the uniqueness of the Cisco UCS WWNNs in the SAN fabric, Cisco recommends using the following WWN prefix 20:00:00:25:b5:00:00:01

Property	Essential Information
Name	The name of the WWNN pool.
Size	The total number of WWNNs in the pool.
Used	The number of WWNNs in the pool that have been used, and are no longer available.
Reserved	Displays the total number of WWNNs in the pool that have been reserved to be used later.
Available	Displays the total number of WWNNs in the pool that are available to be used.
Description	A description of the WWNN pool.
Last Update	When the WWNN pool was last updated.
Configuration	
From	Displays the WWNN prefix value of the pool. Note Cisco Intersight selects the identity in sequential manner, that is, the lowest available identity from the pool.
To	Displays the WWNN suffix value of the pool.
Size	Displays the WWNN pool size
Usage	
Identifier, Status, Server Profile, and Source	Displays the WWNN, status of usage (Reserved or Used) and associated server profiles. Source can be Self or Other , where Self is ID used or reserved by this pool and Other is ID used or reserved statically or by another pool.
Actions	
Edit	Allows to add or modify the configuration details of the WWNN pool.
Delete	Allows to delete the WWNN pool.

Creating a WWPN Pool

To ensure the uniqueness of the Cisco UCS WWNNs and WWPNS in the SAN fabric, Cisco Intersight uses the following WWN prefix for all blocks in a pool: 20:00:00:25:B5:xx:xx:xx.

Step 1 From the left navigation panel, click **Pools > WWNN > Create WWPN Pool**.

The **WWPN Pool** wizard appears.

Step 2 Add the following information on the **General** page:

- **Name**—Name of the WWPN pool
- **Description**—An optional description of the WWPN pool.
- **Organization**—An optional entry of the organization to which the WWPN pool belongs.
- **Add Tag**—An optional tag to identify and search for the WWPN pool.

Step 3 Click **Next**. The **Pool Details** page appears.

Step 4 Add the following configuration information in the **Initiator Blocks** area:

- **From**—Indicates the first WWN identifier of the block.
- **Size**—Indicates the maximum number of identifiers that the block can contain.

Step 5 To add more blocks, click + and then add the starting WWN identifier and maximum number of identifiers that the block can contain.

Step 6 Click **Create**.

The newly created WWPN pool appears in the list of WWPN pools.

WWPN Pool Details

Details

Displays the list of WWPN pools. To ensure the uniqueness of the Cisco UCS WWPNs in the SAN fabric, Cisco recommends using the following WWN prefix 20:00:00:25:b5:00:00:01

Property	Essential Information
Name	The name of the World Wide Port Name pool.
Size	The total number of WWPNs in the pool.
Used	The number of WWPNs in the pool that have been used, and are no longer available.
Description	A description of the WWPN pool.
Last Update	When the WWPN pool was last updated.

IQN Pools

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

An IQN pool is a collection of iSCSI Qualified Names (IQNs) for use as initiator identifiers. The IQN pool details are used for configuring blocks of IQN identifiers.

Procedure

	Command or Action	Purpose
Step 1	From the left navigation panel, click Create Pools > IQN > Start .	The IQN Pool wizard appears.
Step 2	Add the following information on the General page: <ul style="list-style-type: none"> • Organization—The organization of the IQN pool. • Name—Name of the IQN pool. • Add Tag—The tag to identify and search for the IQN pool. • Description—The description the IQN pool. 	
Step 3	Click Next . The Pool Details page appears.	
Step 4	Add the following configuration information for IQN pools in the Configuration area: <ul style="list-style-type: none"> • Prefix—The prefix for any IQN blocks created for this pool. IQN prefix must have the following format "iqn-yyyy-mm.naming-authority", where the naming authority is usually the reverse syntax of the internet domain of the naming authority. Example, iqn1.2021-01.alpha.com • Suffix—The suffix for this block of IQNs. Enter from 1 to 64 characters. You can use any letter or number, as well as the special characters . (period), : (colon), and - (hyphen). • From—The first iSCSI Qualified Name (IQN) suffix in the block. • Size—The number of identifiers this block can hold. 	

The newly created IQN pool appears in the list of IQN pools.

IQN Pool Details

Details

Displays the list of IQN pools.

Property	Essential Information
Details	

Property	Essential Information
Name	Displays the name of the IQN pool.
Type	Displays the type of the pool.
Size	Displays the total number of identifiers the IQN pool contains.
Used	Displays the number of identifiers already in use from the pool.
Reserved	Displays the total number of identifiers that have been reserved to be used later.
Available	Displays the number of identifiers available for usage.
Description	A description of the IQN pool.
Last Update	The date and time when the IQN pool was last updated.
Organization	Users in a Default Organization automatically has access to all the resources available for the user account.
Tags	Displays the tags for the pools.
Configuration	
Prefix	Displays the prefix for IQN blocks created for this pool.
Suffix	Displays the suffix for this block of IQNs.
From	The first suffix number in the block. Note Cisco Intersight selects the identity in sequential manner, that is, the lowest available identity from the pool.
To	The number of identifiers that the block can hold.
Usage	
IQN Address, Status, Server Profile, and Source	Displays the IQN address, status of usage (Reserved or Used) and associated server profiles. Source can be Self or Other , where Self is ID used or reserved by this pool and Other is ID used or reserved statically or by another pool.
Actions	
Edit	Allows to add or modify the configuration details of the IQN pool.

Property	Essential Information
Delete	Allows to delete the IQN pool.

Resource Pools

Pools enable you to logically group and manage resources (servers and other endpoints) more efficiently. You can assign servers to a resource pool and can continue with the automated server profile assignment. For more information on the server profile association using resource pools, see [Configuring Server Profiles](#).

Persistent Resource Pool Assignment

When a server that is a part of a pool is decommissioned, it appears in the **Decommissioned Resources** section of the **Resource Pool Details View** and **Server Details View**. When the server is recommissioned, it gets assigned to the same pool again. The same behavior occurs when a server is decommissioned, moved to a different chassis or slot, and then recommissioned. Thus, you do not need to manage the pool reassignment for that server when physical changes occur in the deployment environment.



Note To convert an existing resource pool into a persistent Resource Pool, edit the Resource Pool.

Change in Behavior for API or Terraform Users

API or Terraform users can use APIs to create new Resource Pools using either Managed Object IDs (MOIDs) or Serial selectors.

However, if an API user edits a Resource Pool that uses a MOID from the UI to enable persistent Resource Pool assignment, the system will internally convert these MOID selectors to Serial selectors, and the MOIDs will no longer be accessible through APIs. For more information on the payload for creating Resource Pools, see [API documentation](#).



Note Using the edit resource pool option, a resource with an active lease cannot be removed from the resource pool.

Creating a Resource Pool

A resource pool represents a collection of resources that can be associated to the configuration entities such as server profiles.

Step 1 From the left navigation panel, click **Create Pools > Resource > Start**.

The **Resource Pool** wizard is displayed.

Step 2 Add the following information on the **General** page:

- **Organization**—The organization of the Resource pool.
- **Name**—Name of the Resource pool.
- **Target Platform**—The target platform type as UCS Standalone server or UCS FI-Attached server.

- **Set Tags**—The tag to identify and search for the Resource pool.
- **Description**—The description of the Resource pool.

Step 3 Click **Next**. The **Resource Pool Details** page is displayed with the list of discovered servers based on the target platform type.

Step 4 Select the servers from the **Resource Selection** table.

Step 5 Click **Create**.

The newly created Resource pool appears in the list of Resource pools.

Resource Pool Details

Details - Displays the details of the resource pools.

Property	Essential Information
Details	
Name	Displays the name of the resource pool.
Type	Displays the type of the pool.
Size	Displays the total number of resources that the Resource pool contains.
Used	Displays the number of resources that are already used, and are unavailable for use.
Available	Displays the number resource pool available for usage.
Last Update	The date and time of the resource pool that was last updated.
Resource	
Type	Displays the resource pool type. Note Currently, Intersight supports only server type as a resource for the resource pool.
Selection	Displays the resource pool selection type. Currently, only manual (Static) selection is supported.
Target Platform	Displays the target platform. This could any of the following: <ul style="list-style-type: none"> • Standalone • FI-Attached
Description	Description of the resource pool.

Property	Essential Information
Organization	Displays the organization under which the Resource Pool is created
Configuration Note The configuration properties of the resource pool differs with the resource type associated.	
Status	Displays the status of the resource. This can be any of the following: <ul style="list-style-type: none"> • Available—Indicates the resource is available for use. • Used—Indicates the resource is already used in a resource pool.
Decommissioned Resources: This section displays the details of decommissioned servers. Note The section is displayed only when a server has been decommissioned and is already a part of the resource pool.	
Name	Displays the name of the decommissioned server.
Type	Displays whether the server is a Cisco UCS C-Series server or a Cisco UCS B-Series server.
ID	Displays the unique ID assigned to the decommissioned server. This field applies only to Cisco UCS C-Series servers.
Model	Displays the model of the server.
Serial Number	Displays the serial number of the server.
Decommissioned Date	Displays the time stamp at which the server was decommissioned.
Usage	
Resource Name	Displays the resource name.
Leasing Entity	Displays the configuration entity. Note A resource can be part of different pools but are allowed to be associated to only one leasing entity.
Use Case	Displays the consumer of the resource. Example, Server Profile.

Property	Essential Information
Resource Usage	<p>Displays the resource consumption types. The types can be:</p> <ul style="list-style-type: none"> • Current—The resource is associated and used in the current resource pool. • Other Pool—The resource is associated and used in an other pool. • Direct—The resource is directly associated with the server profile without using resource pool.



Note Using an edit resource pool option, a resource with an active lease cannot be removed from the resource pool.

Virtual Routing and Forwarding

Virtual Routing and Forwarding (VRF) is an IP technology that allows multiple instances of a routing table to coexist on the same router at the same time. Because the routing instances are independent, the same or overlapping IP addresses can be used without conflict. A VRF creates a namespace for IP address management. IP pools are VRF-aware in Cisco Intersight.

VRF Guidelines

The following guidelines and limitations apply for VRF instances:

- Intersight creates a default VRF for an account, and manages IP address allocation within the context of this default VRF.
- Within a single VRF instance, IP addresses must be unique. Between different VRF instances, you can have overlapping IP addresses.
- If IP Pools are shared between VRF instances, ensure that there are no overlapping IP addresses.

Creating a VRF Instance

Virtual Routing and Forwarding (VRF) is a networking technology that creates multiple virtual networks within a single network entity.

Step 1 From the left navigation panel, click **Virtual Routing And Forwarding > VRFs > Create VRF**.

The **VRF** wizard appears.

Step 2 Add the following information on the **General** page:

- **Name**—Name of the VRF instance
- **Description**—An optional description of the VRF instance.

- **Organization**—An optional entry of the organization to which the VRF instance belongs.
- **Add Tag**—An optional tag to identify and search for the VRF instance.

Step 3 Click **Create**.

The newly created VRF instance appears in the list of VRFs.



CHAPTER 12

Managing the Device Console

- [Device Console](#), on page 291

Device Console

The Device Console, which is installed on the Fabric Interconnect, allows you to monitor the health of your devices, and the status of their connection to Intersight. You can use the Device Console GUI or CLI interface if you want to troubleshoot your devices, or if your devices are not connecting to Cisco Intersight.

To access the Device Console user interface, log in to the Fabric Interconnect using a management IP address. You must have administrator privileges to access Device Console UI. For more information, see the [Cisco Intersight Managed Mode Fabric Interconnect Admin Guide](#) .



CHAPTER 13

Managing Firmware

- [Firmware Upgrade in a Cisco UCS Domain through Intersight, on page 293](#)
- [Upgrading Fabric Interconnect Firmware, on page 295](#)
- [Upgrading Server Firmware, on page 297](#)
- [Upgrading Server Firmware Using Firmware Policy, on page 298](#)
- [Upgrades and Replacement of RMA Servers and Fabric Interconnects, on page 298](#)

Firmware Upgrade in a Cisco UCS Domain through Intersight

You can upgrade the firmware for various components in a Cisco UCS Domain through Cisco Intersight by choosing one of the following upgrade options:

Fabric Firmware Upgrade

Through this process, you can upgrade all the fabric components in a Cisco UCS Domain, including the two Fabric Interconnects and I/O modules. These components are upgraded to the firmware version included in the selected fabric firmware bundle. Fabric firmware upgrade does not support a partial upgrade to only some components in a Cisco UCS Domain. The fabric firmware upgrade process is valid only for Cisco UCS 6400 Series Fabric Interconnects.

Fabric firmware bundles are available in the Cisco Intersight repository and have two component images:

- NXOS image
- CMC image

The following workflow illustrates the fabric firmware upgrade process:

1. **Fabric Selection:** You can initiate the fabric firmware upgrade process by selecting a Fabric Interconnect and performing an **Upgrade Firmware** action on it. Fabric Interconnects are always upgraded as a pair, in which Fabric Interconnect-B is upgraded before Fabric Interconnect-A.
2. **Bundle Selection:** After you select the Fabric Interconnect pair to be upgraded, you must select the fabric firmware bundle to which the Fabric Interconnects need to be upgraded. The firmware selection screen displays a list of available firmware bundles and information about their firmware version, size, release date, and description. The selected firmware bundle is downloaded from the Cisco Intersight repository.
3. **Impact Estimation:** The Summary screen shows a summary of the selected switches, the firmware version running on them, and the firmware version to which they will be upgraded. You can choose to upgrade by clicking **Upgrade**, or change the configuration by clicking **Back**.

4. Upgrade Request Submission: After you click **Upgrade**, confirm the upgrade request.

The following workflow illustrates the tasks that occur automatically after you submit an upgrade request:

1. The system validates whether there is enough storage space for the firmware bundle. If the space on the Fabric Interconnect is insufficient, the upgrade fails.
2. The system checks whether the selected firmware bundle is already in the Fabric Interconnect cache. If the firmware bundle is not present, it is downloaded to the Fabric Interconnect cache.
3. Both the IO modules are updated and activated on all the connected chassis. IO module upgrade is completed when the IO modules are rebooted.
4. Click **Continue** to acknowledge and begin firmware upgrade on Fabric Interconnect-B. After Fabric Interconnect-B upgrade is complete, the Fabric Interconnect reboots and comes up with the new image. IOM-B is rebooted along with the Fabric Interconnect-B, and comes up with the upgraded image.
5. Click **Continue** to acknowledge and begin firmware upgrade on Fabric Interconnect-A. After Fabric Interconnect-A upgrade is complete, the Fabric Interconnect reboots and comes up with the new image. IOM-A is rebooted along with the Fabric Interconnect-A, and comes up with the upgraded image.

Host Firmware Upgrade

Through this process, you can upgrade all the server components for Cisco UCS B-Series and C-Series FI-Attached servers that are in Intersight Managed Mode. These components are upgraded to the firmware version included in the selected host firmware bundle.

Server firmware bundles are available in the Cisco Intersight repository, and have the following component images:

- CIMC image
- BIOS image
- Network adapter image



Note Only UCS VIC 1400 Series adapters are supported.

- Storage controller image
- Board controller image
- Disk image
- GPU image
- Memory card image
- M-Switch and PLX images

The following workflow illustrates the host firmware upgrade process.

1. **Server Selection:** You can initiate the host firmware upgrade process by selecting a server and performing an **Upgrade Firmware** action on it.

2. **Bundle Selection:** After you confirm the server to be upgraded, you must select the host firmware bundle to which the server needs to be upgraded. The firmware selection screen displays a list of available firmware bundles and information about their firmware version, size, release date, and description. The selected firmware bundle will be downloaded from the Cisco Intersight repository.
3. **Impact Estimation:** The Summary screen shows a summary of the selected server, the firmware version running on it, and the firmware version to which it will be upgraded. You can choose to upgrade by clicking **Upgrade**, or change the firmware version by clicking **Back**.
4. **Upgrade Request Submission:** After you click **Upgrade**, select whether you want the firmware to be installed immediately or when the device reboots. Confirm the upgrade request.

By default, firmware will be installed on next boot of the device.

The following workflow illustrates the tasks that occur automatically after you submit an upgrade request:

1. The system validates whether there is enough storage space for the firmware bundle. If the space on the Fabric Interconnect is insufficient, the upgrade fails.
2. The system checks whether the selected firmware bundle is already in the Fabric Interconnect cache. If the firmware bundle is not present, it is downloaded to the Fabric Interconnect cache.
3. Server firmware is upgraded as follows:
 - For B-Series servers:
 - a. Adapter firmware is updated and activated. Adapter upgrade is completed when the server is rebooted.
 - b. The Host Service Utility (HSU) is upgraded immediately or when the server reboots.
 - c. All server components are upgraded.
 - For C-Series servers:
 - a. The HSU is upgraded immediately or when the server reboots.
 - b. All server components are upgraded.
4. Click **Continue** to acknowledge and begin firmware upgrade.

Upgrading Fabric Interconnect Firmware

You can upgrade Intersight managed Fabric Interconnect using Cisco Intersight.

Before you begin

Before you upgrade your Intersight managed Fabric Interconnect firmware, consider the following prerequisites:

- Only Cisco UCS 6400 Series Fabric Interconnects in a Cisco UCS Domain may be upgraded.
- You must have at least the following available storage in the Fabric Interconnect partitions for the firmware bundle to be downloaded:
 - 90 percent free space in /var/tmp

- 20 percent free space in /var/sysmgr
- 30 percent free space in /mnt/pss
- 18 percent free space in /bootflash
- Only Cisco UCS Domains that are claimed through Intersight may be upgraded.
- All servers in the Cisco UCS Domain must be at license tier Essentials or above.

-
- Step 1** From the left navigation pane, click **Fabric Interconnects**, select a Fabric Interconnect, and perform an **Upgrade Firmware** action on it.
- Step 2** On the **Upgrade Firmware** page, click **Start**.
- Step 3** On the **General** page, confirm selection of the switch Domain and click **Next**.
- Step 4** On the **Version** page, select the fabric firmware bundle to which the Fabric Interconnects need to be upgraded, and click **Next**.

This page displays a list of available firmware bundles and information about their firmware version, size, release date, and description. The selected firmware bundle will be downloaded from the Intersight repository.

During upgrade of Intersight Managed Fabric Interconnect, the Fabric Interconnect traffic evacuation is enabled by default. Fabric Interconnect traffic evacuation evacuates all traffic that flows through the Fabric Interconnect from all servers attached to it, while upgrading the system. The traffic will fail over to the peer Fabric Interconnect for fail over vNICs. Before the traffic evacuation on a Fabric Interconnect, the user must acknowledge that replay on peer Fabric Interconnect is completed and all vEths are up. Use the *show interface virtual status* command to check the vEth status for respective vEth from NXOS.

Before the traffic evacuation, you can check the traffic flowing through the Fabric Interconnect by viewing the Transmit (Tx) and Receive (Rx) stats of Host Interfaces (HIFs). After the traffic evacuation, you can check the traffic flowing through the Fabric Interconnect (FI) by viewing the Transmit (Tx) and Receive (Rx) stats of Network Interfaces (NIFs).

Note For Fabric Interconnect traffic evacuation to be functional, vNIC failover must be enabled in the LAN Connectivity Policy

Select **Advanced Mode** to disable the Fabric Interconnect traffic evacuation.

- Step 5** On the **Summary** screen, verify the summary of the selected switches, the firmware version running on them, and the firmware version to which they will be upgraded, and click **Upgrade**.

You can choose to change the firmware version by clicking **Back**.

- Step 6** Confirm the upgrade request.

The firmware upgrade workflow begins. You can check the status of the upgrade workflow in the **Execution Flow** pane. Acknowledge any messages in the **Execution Flow** pane and click **Continue** to proceed with the upgrade.

Upgrading Server Firmware

Before you begin

Before you upgrade your server, consider the following prerequisites:

- Only Cisco UCS B-Series M5, M6, C-Series M5, M6, M7, M8 and X-Series M6 and M7 servers that are claimed through Intersight may be upgraded.
- Servers may be upgraded from a minimum of Cisco UCS HSU bundle release version 4.1(2a).
- All servers in the Cisco UCS Domain must be at license tier Essentials or higher.

Step 1 From the left navigation pane, click **Servers**, select a server, and perform an **Upgrade Firmware** action on it.

Note To upgrade more than one server :

- Ensure that the selected servers are of the same model and management mode. Following are examples of valid selections:
 - One or more B200 M5 servers
 - One or more C220 M5 servers

Following are examples of invalid selections:

- C220 M5 and C240 M5 servers
 - C220 M5 and B200 M5 servers
- Select all the servers to be upgraded, click the three-dot menu at the bottom of the list view, and perform the **Upgrade Firmware** action on them.

Step 2 On the **Upgrade Firmware** page, click **Start**.

Step 3 On the **General** page, confirm selection of the server and click **Next**.

Step 4 On the **Version** page, select the Cisco UCS HSU bundle to which the server must be upgraded, and click **Next**.

This page displays a list of available firmware bundles and information about their firmware version, size, release date, and description. The selected firmware bundle will be downloaded from the Cisco repository. By default, all the server components will be upgraded, including drives and storage controllers.

Select **Advanced Mode** to exclude drives and storage controllers from the upgrade.

Step 5 On the **Summary** screen, verify the summary of the selected servers, the firmware version running on them, and the firmware version to which they will be upgraded.

You can choose to change the configuration by clicking **Back**.

Step 6 Click **Upgrade**.

Step 7 In the **Upgrade Firmware** dialog box, choose one the following options:

- a) **Reboot Immediately To Begin Upgrade**—By default, server firmware is upgraded on next boot. Enable this option if you choose to reboot immediately to begin firmware upgrade.

b) Click **Upgrade** to confirm the upgrade request.

The firmware upgrade workflow begins. You can check the status of the upgrade workflow in the **Execution Flow** pane. Acknowledge any messages in the **Execution Flow** pane and click **Continue** to proceed with the upgrade.

Upgrading Server Firmware Using Firmware Policy

Perform the following steps to upgrade the server firmware using Firmware policy:

1. Create a Firmware policy as described in [Creating a Firmware Policy](#).
2. Attach the policy to the server profile of the server that needs to be upgraded.
The profile status will change to 'Pending Changes'.
3. Deploy the profile.
The firmware will be downloaded and prepared for upgrade upon the next reboot.
The profile status will change to 'Pending Reboot'.
4. Reboot the server.
The Activation workflow is triggered, and the firmware upgrade is monitored.
After the successful upgrade, the profile status changes to 'OK'.

Upgrades and Replacement of RMA Servers and Fabric Interconnects

RMA Upgrades

RMA is a Return Material Authorization process that enhances customer experience.

Upgrade of RMA Server

The RMA process triggers an automatic discovery workflow when you insert a new blade server, or when you replace an old blade server. The discovery workflow raises an alarm if the firmware of the blade server is outdated, and you will be asked to trigger an upgrade workflow.

1. Go to **Chassis > Inventory > Servers**
2. Select the server that you want to upgrade
3. Click Upgrade.
4. Select the firmware version to which you want to upgrade the server.

Relevant endpoints like Cisco IMC and Adaptor are upgraded to ensure that the server comes up in the Intersight Managed Mode, is available in the server list page, and is ready for use. You can upgrade the rest of the endpoints using the standard firmware upgrade method



Note The CMC version must be 4.1(3b) or later.

RMA support is not available for FI-attached C-Series servers in Intersight Managed Mode. You first need to convert the C-Series server in IMM to Standalone mode, verify the firmware, and then upgrade using HUU.

To convert the server from IMM to Standalone, See [Converting a Server in Intersight Managed Mode to Standalone Mode](#).

To upgrade the firmware of C-Series Standalone server, See [Upgrading UCS C-Series Standalone Servers Firmware](#).

Replacement of RMA Fabric Interconnect

When a single Fabric Interconnect, or a Fabric Interconnect cluster is faulty, and the Fabric Interconnects have been replaced of the same model, you can use the Replace option for migrating the configuration of the old Fabric Interconnects to the new ones. The workflows for replacing both a single Fabric Interconnect and a Fabric Interconnect cluster are detailed in the subsequent paragraphs.

Replacement of Single Fabric Interconnect in Cisco UCS 6400 Series and 6500 Series Fabric Interconnects

1. Remove the old Fabric Interconnect and connect the new Fabric Interconnect of the same model.
2. Move all cable connections, including servers, FEX fabrics, and blade chassis, from the old Fabric Interconnect to the new Fabric Interconnect.
3. Configure the new module to join the cluster, before triggering the 'Replace' action.
4. Go to **Operate > Fabric Interconnects** to view the list of Fabric Interconnects that have been replaced and to identify those for which the 'Replace' option is enabled.
5. Select the Replace Fabric Interconnect option.
6. Click Replace on the confirmation page to trigger the replacement workflow.

As part of the workflow:

- The disconnected Fabric Interconnect is removed from inventory.
- The domain profile is reassigned to the new Fabric Interconnect and deployed.
- The servers, chassis, and FEX are inventoried and discovered under the new Fabric Interconnect.
- The server and chassis profiles are redeployed with Fabric Interconnect related policies.

Replacement of Single Fabric Interconnect in Cisco UCS Fabric Interconnects 9108 100G

1. Remove the old Fabric Interconnect and insert the new Fabric Interconnect of the same model into the slot.

If the replacement Fabric Interconnect has an existing configuration, ensure that you perform an erase-configuration on the device before using it for replacement.
2. Reconnect all the cable connections to the new Fabric Interconnect.
3. Configure the new module to join the cluster, before triggering the 'Replace' action.
4. Go to **Operate > Fabric Interconnects** to view the list of Fabric Interconnects that have been replaced and to identify those for which the 'Replace' option is enabled.

5. Select the Replace Fabric Interconnect option.
6. Click Replace on the confirmation page to trigger the replacement workflow.

As part of the workflow:

- The disconnected Fabric Interconnect is removed from inventory.
- The domain profile is reassigned to the new Fabric Interconnect and deployed.
- The servers and chassis are inventoried and discovered under the new Fabric Interconnect.
- The server and chassis profiles are redeployed with Fabric Interconnect related policies.

Replacement of Fabric Interconnect Cluster in Cisco UCS 6400 Series and 6500 Series Fabric Interconnects

1. Remove the old Fabric Interconnect cluster and connect the new Fabric Interconnect cluster of the same model.
2. Move all the cable connections, including servers, FEX fabrics, and blade chassis, from the old Fabric Interconnects to the new Fabric Interconnects.
3. Claim the new Fabric Interconnects in Intersight.
4. Select the 'Replace UCS Domain' option displayed next to the old cluster on the Fabric Interconnects page.
5. Choose the new Fabric Interconnect cluster to replace the old one.

As part of the workflow:

- The old device registration is merged with the new device registration.
- The disconnected Fabric Interconnect cluster is removed from inventory.
- The domain profile is reassigned to the new Fabric Interconnect cluster and deployed.
- The servers, chassis, and FEX are inventoried and discovered under the new Fabric Interconnect cluster.
- The server and chassis profiles are redeployed with Fabric Interconnect related policies.

Replacement of Fabric Interconnect Cluster in Cisco UCS Fabric Interconnects 9108 100G

1. Remove the old Fabric Interconnect cluster and insert the new Fabric Interconnect cluster of the same model into the slots.

If Fabric Interconnects with existing configurations are used for replacement, ensure to perform an erase-configuration on the devices before using them for replacement.

2. Reconnect all the cable connections to the new Fabric Interconnects.
3. Claim the new Fabric Interconnects in Intersight.

The Primary Chassis Discovery workflow will run and fail, displaying a warning to run the 'Replace UCS Domain' action, as the chassis ID is found to be the same.

4. Select the 'Replace UCS Domain' option displayed next to the old cluster on the Fabric Interconnects page.

5. Choose the new Fabric Interconnect cluster to replace the old one.

As part of the workflow:

- The old device registration is merged with the new device registration.
- The disconnected Fabric Interconnect cluster is removed from inventory.
- The domain profile is reassigned to the new Fabric Interconnect cluster and deployed.
- The servers, chassis, and FEX are inventoried and discovered under the new Fabric Interconnect cluster.
- The server and chassis profiles are redeployed with Fabric Interconnect related policies.



Note Swapping of Fabric Interconnects between slots is not supported.

Replacement of Chassis in Cisco UCS X-Series Direct

Cisco UCS X-Series Direct also supports chassis RMA for UCSX-9508. To do this:

1. Insert the modules into the new chassis.
2. Power on the chassis.

As part of the workflow:

- Discovery of the chassis is automatically initiated.
- The servers and chassis are inventoried and discovered.
- The new chassis serial number is updated in the Intersight inventory.

Cisco Intersight Support for Auto Upgrade of IOM

You do not have to manually update the firmware of IOMs that have CMC lower than 4.1(3b). When the chassis is connected to the Fabric Interconnect, the firmware is automatically updated, the server port is configured in the Port Policy, the port policy is associated with the domain profile, and the domain profile is deployed.



CHAPTER 14

Managing Technical Support

- [Integration with Cisco TAC, on page 303](#)
- [Tech Support Diagnostic File Collection, on page 304](#)

Integration with Cisco TAC



Important

- Tech Support diagnostic files are generated locally at the endpoints and you cannot access them at any point. Intersight does not currently send any notifications about the Tech Support files or other case-related activities.
- Connected TAC is available only for cases opened directly with Cisco TAC.
- For partner support cases Connected TAC works as expected only if:
 - The partner opens a case on behalf of the Intersight user.(Or)
 - The partner has authorized Intersight users to open a case directly with Cisco TAC.

You can create a Cisco TAC Service Request (SR) directly from Intersight by launching **Cisco Support Case Manager** for the following:

- **HyperFlex Clusters** from the table view and details view.
- **IWE Clusters** from the table view and details view.
- **Servers** from the table view and details view.
- **Fabric Interconnects** from the table view.

You can also open a Cisco TAC case from the Intersight Mobile App.

Before you open a case, please ensure that the following requirements are met:

- A valid service contract (entitlement) exists for the hardware.
- Your Cisco ID is associated with the service contract.

To open a Cisco TAC case:

1. Select a **HyperFlex Cluster**, or a **IWE Cluster**, or a **Server**, or a **Fabric Interconnect** from the corresponding table view and click the ellipsis (...) in the actions column on the right. You can also Open a TAC Case from the **Actions** menu on the **HyperFlex Cluster**, or **IWE Cluster**, or the **Server Details** page.
2. Select **Open TAC Case**. The Open a TAC Case window displays with the name and serial number of the selected HyperFlex cluster or server or Fabric Interconnect.
3. Click **Continue** to launch **Cisco Support Case Manager**. On the **Cisco Support Case Manager** UI, verify the auto-populated details of your case, add a description and a title for your TAC Case, and click **Submit**.

For detailed information about the Proactive Support workflow, configuring the advanced options, and opting out of proactive RMA, see [Proactive RMA for Intersight Connected Devices](#).

For the requirements and benefits of proactive RMA, see [Proactive Support Enable Through Intersight](#).

Tech Support Diagnostic File Collection

When you open a case with Cisco TAC, Intersight collects Tech Support diagnostic files to assist with an open support case. The data collected could include (but is not limited to) hardware telemetry, system configuration, and any other details that aid in active troubleshooting of the TAC case. Tech Support collection is allowed to occur regardless of data collection options you specify. However, this information is not collected arbitrarily, but only when you open a case against a system, requiring assistance with the system support.



Note The Tech Support diagnostic file collection is not supported for unclaimed Intersight managed device.

Account admin users can also submit tech-support collection request from the Tech Support Bundles page by clicking Add Tech Support Bundle and providing device's PID, serial number, and platform type.

To initiate the Tech Support diagnostic file collection for Intersight Managed FI attached devices, enter PID and serial number of the device, and then choose **Intersight Managed Domain** as the platform type in the **Add Tech Support Bundle** window.

- For IMM devices, the tech-support collection follows a best-effort strategy, where the collection includes all possible endpoint logs in a bundle.
- The final collection status appears **Completed** when at least one endpoint's logs get collected in the bundle.
- *The tech_support.log and peer_tech_support.log* files in the final tech-support bundle contain information about the missing endpoint logs and any collection failures.

The following table provides the combination of input that are required to initiate the Tech Support diagnostic file collection.

Tech Support Bundle Type	PID & Serial Number
Chassis	IOM-1, or IOM-2, or chassis
Fabric Interconnect (FI)	FI-A or FI-B

Tech Support Bundle Type	PID & Serial Number
Blade Server	Blade or adapters connected to blade

For Intersight Managed FI Attached devices, Tech Support diagnostic file collection is supported on the following endpoints:

- Blade BMC
- Blade adapter
- Blade chassis
- Fabric Interconnect
- IO modules
- Rack servers
- Rack server adapters
- Server Bundle
- Fabric Extender

