



Cisco UCS Manager CLI Configuration Guide, Release 2.0

First Published: September 06, 2011

Last Modified: September 04, 2012

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

Text Part Number: OL-25721-05

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 NONINFRINGEMENT 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.

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: <http://www.cisco.com/go/trademarks>. 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. (1110R)

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

© 2011-2012 Cisco Systems, Inc. All rights reserved.



CONTENTS

Preface

Preface xxxi

Audience xxxi

Conventions xxxi

Related Cisco UCS Documentation xxxiii

Documentation Feedback xxxiii

Obtaining Documentation and Submitting a Service Request xxxiii

PART I

Introduction 1

CHAPTER 1

New and Changed Information 3

New and Changed Information for this Release 3

CHAPTER 2

Overview of Cisco Unified Computing System 9

About Cisco Unified Computing System 9

Unified Fabric 10

Fibre Channel over Ethernet 11

Link-Level Flow Control 11

Priority Flow Control 11

Server Architecture and Connectivity 12

Overview of Service Profiles 12

Network Connectivity through Service Profiles 12

Configuration through Service Profiles 12

Service Profiles that Override Server Identity 13

Service Profiles that Inherit Server Identity 14

Service Profile Templates 15

Policies 15

Configuration Policies 15

Boot Policy	15
Chassis Discovery Policy	16
Dynamic vNIC Connection Policy	19
Ethernet and Fibre Channel Adapter Policies	19
Global Cap Policy	20
Host Firmware Package	21
IPMI Access Profile	21
Local Disk Configuration Policy	22
Management Firmware Package	22
Management Interfaces Monitoring Policy	23
Network Control Policy	23
Power Control Policy	24
Power Policy	24
Quality of Service Policy	25
Rack Server Discovery Policy	25
Server Autoconfiguration Policy	25
Server Discovery Policy	25
Server Inheritance Policy	26
Server Pool Policy	26
Server Pool Policy Qualifications	26
vHBA Template	27
VM Lifecycle Policy	27
vNIC Template	27
vNIC/vHBA Placement Policies	28
Operational Policies	28
Fault Collection Policy	28
Flow Control Policy	29
Maintenance Policy	29
Scrub Policy	29
Serial over LAN Policy	30
Statistics Collection Policy	30
Statistics Threshold Policy	30
Pools	31
Server Pools	31
MAC Pools	31

UUID Suffix Pools	32
WWN Pools	32
Management IP Pool	33
Traffic Management	33
Oversubscription	33
Oversubscription Considerations	33
Guidelines for Estimating Oversubscription	34
Pinning	35
Pinning Server Traffic to Server Ports	35
Guidelines for Pinning	36
Quality of Service	37
System Classes	37
Quality of Service Policy	38
Flow Control Policy	38
Opt-In Features	38
Stateless Computing	38
Multi-Tenancy	39
Virtualization in Cisco UCS	40
Overview of Virtualization	40
Overview of Cisco Virtual Machine Fabric Extender	41
Virtualization with Network Interface Cards and Converged Network Adapters	41
Virtualization with a Virtual Interface Card Adapter	41

CHAPTER 3**Overview of Cisco UCS Manager 43**

About Cisco UCS Manager	43
Tasks You Can Perform in Cisco UCS Manager	44
Tasks You Cannot Perform in Cisco UCS Manager	46
Cisco UCS Manager in a High Availability Environment	46

CHAPTER 4**Overview of Cisco UCS Manager CLI 47**

Managed Objects	47
Command Modes	47
Object Commands	49
Complete a Command	50
Command History	50

Committing, Discarding, and Viewing Pending Commands	50
Online Help for the CLI	51
CLI Session Limits	51
Web Session Limits	51
Setting the Web Session Limit for Cisco UCS Manager from the CLI	51
Pre-Login Banner	52
Creating the Pre-Login Banner	52
Modifying the Pre-Login Banner	53
Deleting the Pre-Login Banner	54

PART II

System Configuration 55**CHAPTER 5****Configuring the Fabric Interconnects 57**

Initial System Setup	57
Setup Mode	58
System Configuration Type	58
Management Port IP Address	58
Performing an Initial System Setup for a Standalone Configuration	59
Initial System Setup for a Cluster Configuration	61
Performing an Initial System Setup for the First Fabric Interconnect	61
Performing an Initial System Setup for the Second Fabric Interconnect	63
Enabling a Standalone Fabric Interconnect for Cluster Configuration	64
Changing the System Name	65
Changing the Management Subnet of a Cluster	65
Ethernet Switching Mode	66
Configuring Ethernet Switching Mode	67
Fibre Channel Switching Mode	68
Configuring Fibre Channel Switching Mode	68

CHAPTER 6**Configuring Ports and Port Channels 71**

Server and Uplink Ports on the 6100 Series Fabric Interconnect	71
Unified Ports on the 6200 Series Fabric Interconnect	72
Port Modes	72
Port Types	73
Beacon LEDs for Unified Ports	74

Guidelines for Configuring Unified Ports	74
Effect of Port Mode Changes on Data Traffic	75
Configuring the Port Mode	76
Configuring the Beacon LEDs for Unified Ports	78
Server Ports	79
Configuring a Server Port	79
Unconfiguring a Server Port	79
Uplink Ethernet Ports	80
Configuring an Uplink Ethernet Port	80
Unconfiguring an Uplink Ethernet Port	81
Appliance Ports	81
Configuring an Appliance Port	81
Assigning a Target MAC Address to an Appliance Port or Appliance Port Channel	83
Unconfiguring an Appliance Port	84
FCoE and Fibre Channel Storage Ports	84
Configuring a Fibre Channel Storage or FCoE Port	84
Unconfiguring a Fibre Channel Storage or FCoE Port	85
Restoring a Fibre Channel Storage Port Back to an Uplink Fibre Channel Port	85
Default Zoning	86
Enabling Default Zoning	86
Disabling Default Zoning	87
Uplink Ethernet Port Channels	87
Configuring an Uplink Ethernet Port Channel	87
Unconfiguring an Uplink Ethernet Port Channel	88
Adding a Member Port to an Uplink Ethernet Port Channel	89
Deleting a Member Port from an Uplink Ethernet Port Channel	90
Appliance Port Channels	90
Configuring an Appliance Port Channel	90
Unconfiguring an Appliance Port Channel	92
Enabling or Disabling an Appliance Port Channel	93
Adding a Member Port to an Appliance Port Channel	93
Deleting a Member Port from an Appliance Port Channel	94
Fibre Channel Port Channels	94
Configuring a Fibre Channel Port Channel	95
Unconfiguring a Fibre Channel Port Channel	96

Enabling or Disabling a Fibre Channel Port Channel	96
Adding a Member Port to a Fibre Channel Port Channel	97
Deleting a Member Port from a Fibre Channel Port Channel	97
Adapter Port Channels	98
Viewing Adapter Port Channels	98
Fabric Port Channels	99
Cabling Considerations for Fabric Port Channels	99
Configuring a Fabric Port Channel	100
Viewing Fabric Port Channels	100
Enabling or Disabling a Fabric Port Channel Member Port	101

CHAPTER 7**Configuring Communication Services 103**

Communication Services	103
Configuring CIM XML	104
Configuring HTTP	105
Unconfiguring HTTP	106
Configuring HTTPS	106
Certificates, Key Rings, and Trusted Points	106
Creating a Key Ring	107
Regenerating the Default Key Ring	107
Creating a Certificate Request for a Key Ring	108
Creating a Certificate Request for a Key Ring with Basic Options	108
Creating a Certificate Request for a Key Ring with Advanced Options	109
Creating a Trusted Point	111
Importing a Certificate into a Key Ring	112
Configuring HTTPS	113
Deleting a Key Ring	114
Deleting a Trusted Point	115
Unconfiguring HTTPS	115
Enabling HTTP Redirection	116
Configuring SNMP	116
Information about SNMP	116
SNMP Functional Overview	116
SNMP Notifications	117
SNMP Security Levels and Privileges	117

Supported Combinations of SNMP Security Models and Levels	118
SNMPv3 Security Features	119
SNMP Support in Cisco UCS	119
Enabling SNMP and Configuring SNMP Properties	119
Creating an SNMP Trap	120
Deleting an SNMP Trap	121
Creating an SNMPv3 User	122
Deleting an SNMPv3 User	123
Enabling Telnet	123
Disabling Communication Services	124

CHAPTER 8**Configuring Authentication 125**

Authentication Services	125
Guidelines and Recommendations for Remote Authentication Providers	125
User Attributes in Remote Authentication Providers	126
LDAP Group Rule	128
Configuring LDAP Providers	128
Configuring Properties for LDAP Providers	128
Creating an LDAP Provider	129
Changing the LDAP Group Rule for an LDAP Provider	132
Deleting an LDAP Provider	133
LDAP Group Mapping	133
Creating an LDAP Group Map	134
Deleting an LDAP Group Map	135
Configuring RADIUS Providers	135
Configuring Properties for RADIUS Providers	135
Creating a RADIUS Provider	136
Deleting a RADIUS Provider	138
Configuring TACACS+ Providers	138
Configuring Properties for TACACS+ Providers	138
Creating a TACACS+ Provider	139
Deleting a TACACS+ Provider	140
Configuring Multiple Authentication Systems	141
Multiple Authentication Systems	141
Provider Groups	141

Creating an LDAP Provider Group	141
Deleting an LDAP Provider Group	142
Creating a RADIUS Provider Group	143
Deleting a RADIUS Provider Group	144
Creating a TACACS Provider Group	145
Deleting a TACACS Provider Group	146
Authentication Domains	146
Creating an Authentication Domain	147
Selecting a Primary Authentication Service	148
Selecting the Console Authentication Service	148
Selecting the Default Authentication Service	149
Role Policy for Remote Users	150
Configuring the Role Policy for Remote Users	151

CHAPTER 9**Configuring Organizations 153**

Organizations in a Multi-Tenancy Environment	153
Hierarchical Name Resolution in a Multi-Tenancy Environment	154
Configuring an Organization Under the Root Organization	156
Configuring an Organization Under an Organization that is not Root	156
Deleting an Organization	157

CHAPTER 10**Configuring Role-Based Access Control 159**

Role-Based Access Control	159
User Accounts for Cisco UCS Manager	159
Guidelines for Cisco UCS Manager Usernames	160
Reserved Words: Locally Authenticated User Accounts	161
Guidelines for Cisco UCS Manager Passwords	162
Web Session Limits for User Accounts	162
User Roles	162
Default User Roles	163
Reserved Words: User Roles	164
Privileges	164
User Locales	166
Configuring User Roles	167
Creating a User Role	167

Adding Privileges to a User Role	167
Removing Privileges from a User Role	168
Deleting a User Role	169
Configuring Locales	169
Creating a Locale	169
Assigning an Organization to a Locale	170
Deleting an Organization from a Locale	170
Deleting a Locale	171
Configuring Locally Authenticated User Accounts	171
Creating a User Account	171
Enabling the Password Strength Check for Locally Authenticated Users	173
Setting Web Session Limits for User Accounts	174
Assigning a Role to a User Account	174
Assigning a Locale to a User Account	175
Removing a Role from a User Account	176
Removing a Locale from a User Account	176
Enabling or Disabling a User Account	177
Clearing the Password History for a Locally Authenticated User	178
Deleting a User Account	178
Password Profile for Locally Authenticated Users	179
Configuring the Maximum Number of Password Changes for a Change Interval	180
Configuring a No Change Interval for Passwords	181
Configuring the Password History Count	181
Monitoring User Sessions	182

CHAPTER 11

Managing Firmware	185
Overview of Firmware	185
Firmware Image Management	186
Firmware Image Headers	187
Firmware Image Catalog	187
Firmware Versions	188
Firmware Upgrades	189
Cautions, Guidelines, and Best Practices for Firmware Upgrades	189
Configuration Changes and Settings that Can Impact Upgrades	190
Hardware-Related Guidelines and Best Practices for Firmware Upgrades	191

Firmware- and Software-Related Best Practices for Upgrades	192
Required Order of Components for Firmware Activation	194
Required Order for Adding Support for Previously Unsupported Servers	195
Direct Firmware Upgrade at Endpoints	196
Stages of a Direct Firmware Upgrade	197
Outage Impacts of Direct Firmware Upgrades	198
Firmware Upgrades through Service Profiles	199
Host Firmware Package	199
Management Firmware Package	200
Stages of a Firmware Upgrade through Service Profiles	200
Firmware Downgrades	201
Completing the Prerequisites for Upgrading the Firmware	201
Prerequisites for Upgrading and Downgrading Firmware	201
Creating an All Configuration Backup File	202
Verifying the Operability of a Fabric Interconnect	203
Verifying the High Availability Status and Roles of a Cluster Configuration	203
Verifying the Status of an I/O Module	204
Verifying the Status of a Server	205
Verifying the Status of Adapters on Servers in a Chassis	205
Downloading and Managing Firmware Packages	206
Obtaining Software Bundles from Cisco	206
Downloading Firmware Images to the Fabric Interconnect from a Remote Location	207
Displaying the Firmware Package Download Status	209
Canceling an Image Download	209
Displaying All Available Software Images on the Fabric Interconnect	210
Displaying All Available Packages on the Fabric Interconnect	211
Determining the Contents of a Firmware Package	211
Checking the Available Space on a Fabric Interconnect	212
Deleting Firmware Packages from a Fabric Interconnect	213
Deleting Firmware Images from a Fabric Interconnect	213
Directly Upgrading Firmware at Endpoints	214
Updating and Activating the Firmware on an Adapter	214
Updating and Activating the BIOS Firmware on a Server	216
Updating and Activating the CIMC Firmware on a Server	218
Updating and Activating the Firmware on an IOM	220

Activating the Board Controller Firmware on a Server	222
Activating the Cisco UCS Manager Software	223
Activating the Firmware on a Fabric Interconnect	224
Updating Firmware through Service Profiles	225
Host Firmware Package	225
Management Firmware Package	226
Effect of Updates to Host Firmware Packages and Management Firmware Packages	226
Creating or Updating a Host Firmware Package	229
Creating or Updating a Management Firmware Package	230
Managing the Capability Catalog	232
Capability Catalog	232
Contents of the Capability Catalog	232
Updates to the Capability Catalog	233
Activating a Capability Catalog Update	233
Verifying that the Capability Catalog is Current	234
Restarting a Capability Catalog Update	234
Viewing a Capability Catalog Provider	235
Downloading Individual Capability Catalog Updates	237
Obtaining Capability Catalog Updates from Cisco	237
Updating the Capability Catalog from a Remote Location	237
Updating Management Extensions	238
Management Extensions	238
Activating a Management Extension	239

CHAPTER 12
Configuring DNS Servers 241

DNS Servers in Cisco UCS	241
Configuring a DNS Server	241
Deleting a DNS Server	242

CHAPTER 13
Configuring System-Related Policies 243

Configuring the Chassis Discovery Policy	243
Chassis Discovery Policy	243
Configuring the Chassis Discovery Policy	246
Configuring the Chassis Connectivity Policy	247
Chassis Connectivity Policy	247

Configuring a Chassis Connectivity Policy	248
Configuring the Rack Server Discovery Policy	248
Rack Server Discovery Policy	248
Configuring the Rack Server Discovery Policy	249
Configuring the Aging Time for the MAC Address Table	250
Aging Time for the MAC Address Table	250
Configuring the Aging Time for the MAC Address Table	250

CHAPTER 14**Managing Licenses 251**

Licenses	251
Obtaining the Host ID for a Fabric Interconnect	252
Obtaining a License	253
Installing a License	254
Viewing the Licenses Installed on a Fabric Interconnect	254
Viewing License Usage for a Fabric Interconnect	255
Uninstalling a License	257

CHAPTER 15**Managing Virtual Interfaces 259**

Virtual Interfaces	259
Virtual Interface Subscription Management and Error Handling	259

PART III**Network Configuration 261****CHAPTER 16****Configuring VLANs 263**

Named VLANs	263
Private VLANs	264
VLAN Port Limitations	265
Configuring Named VLANs	266
Creating a Named VLAN Accessible to Both Fabric Interconnects (Uplink Ethernet Mode)	266
Creating a Named VLAN Accessible to Both Fabric Interconnects (Ethernet Storage Mode)	267
Creating a Named VLAN Accessible to One Fabric Interconnect (Uplink Ethernet Mode)	268

	Creating a Named VLAN Accessible to One Fabric Interconnect (Ethernet Storage Mode) 269
	Deleting a Named VLAN 270
	Configuring Private VLANs 271
	Creating a Primary VLAN for a Private VLAN (Accessible to Both Fabric Interconnects) 271
	Creating a Primary VLAN for a Private VLAN (Accessible to One Fabric Interconnect) 272
	Creating a Secondary VLAN for a Private VLAN (Accessible to Both Fabric Interconnects) 272
	Creating a Secondary VLAN for a Private VLAN (Accessible to One Fabric Interconnect) 273
	Viewing the VLAN Port Count 274
<hr/>	
CHAPTER 17	Configuring LAN Pin Groups 275
	LAN Pin Groups 275
	Configuring a LAN Pin Group 275
<hr/>	
CHAPTER 18	Configuring MAC Pools 277
	MAC Pools 277
	Creating a MAC Pool 277
	Deleting a MAC Pool 278
<hr/>	
CHAPTER 19	Configuring Quality of Service 281
	Quality of Service 281
	Configuring System Classes 281
	System Classes 281
	Configuring a System Class 282
	Disabling a System Class 283
	Configuring Quality of Service Policies 284
	Quality of Service Policy 284
	Configuring a QoS Policy 284
	Deleting a QoS Policy 286
	Configuring Flow Control Policies 287
	Flow Control Policy 287

Configuring a Flow Control Policy 287

Deleting a Flow Control Policy 288

CHAPTER 20

Configuring Network-Related Policies 291

Configuring vNIC Templates 291

vNIC Template 291

Configuring a vNIC Template 292

Deleting a vNIC Template 294

Configuring Ethernet Adapter Policies 294

Ethernet and Fibre Channel Adapter Policies 294

Configuring an Ethernet Adapter Policy 295

Deleting an Ethernet Adapter Policy 297

Configuring Network Control Policies 297

Network Control Policy 297

Configuring a Network Control Policy 298

Deleting a Network Control Policy 299

CHAPTER 21

Configuring Upstream Disjoint Layer-2 Networks 301

Upstream Disjoint Layer-2 Networks 301

Guidelines for Configuring Upstream Disjoint L2 Networks 302

Pinning Considerations for Upstream Disjoint L2 Networks 303

Configuring Cisco UCS for Upstream Disjoint L2 Networks 304

Assigning Ports and Port Channels to VLANs 305

Removing Ports and Port Channels from VLANs 306

Viewing Ports and Port Channels Assigned to VLANs 307

PART IV

Storage Configuration 309

CHAPTER 22

Configuring Named VSANs 311

Named VSANs 311

Fibre Channel Uplink Trunking for Named VSANs 312

Guidelines and Recommendations for VSANs 312

Creating a Named VSAN Accessible to Both Fabric Interconnects (Fibre Channel Uplink Mode) 314

- Creating a Named VSAN Accessible to Both Fabric Interconnects (Fibre Channel Storage Mode) **315**
- Creating a Named VSAN Accessible to One Fabric Interconnect (Fibre Channel Uplink Mode) **316**
- Creating a Named VSAN Accessible to One Fabric Interconnect (Fibre Channel Storage Mode) **317**
- Deleting a Named VSAN **318**
- Changing the VLAN ID for the FCoE Native VLAN for a Named VSAN **318**
- Changing the VLAN ID for the FCoE Native VLAN for a Storage VSAN **319**
- Enabling or Disabling Fibre Channel Uplink Trunking **319**

CHAPTER 23**Configuring SAN Pin Groups 321**

- SAN Pin Groups **321**
- Configuring a SAN Pin Group **321**

CHAPTER 24**Configuring WWN Pools 323**

- WWN Pools **323**
- Creating a WWN Pool **324**
- Deleting a WWN Pool **325**

CHAPTER 25**Configuring Storage-Related Policies 327**

- Configuring vHBA Templates **327**
 - vHBA Template **327**
 - Configuring a vHBA Template **327**
 - Deleting a vHBA Template **329**
- Configuring Fibre Channel Adapter Policies **329**
 - Ethernet and Fibre Channel Adapter Policies **329**
 - Configuring a Fibre Channel Adapter Policy **330**
 - Deleting a Fibre Channel Adapter Policy **332**

PART V**Server Configuration 333**

CHAPTER 26**Configuring Server-Related Pools 335**

- Server Pool Configuration **335**
 - Server Pools **335**

Creating a Server Pool	335
Deleting a Server Pool	336
UUID Suffix Pool Configuration	337
UUID Suffix Pools	337
Creating a UUID Suffix Pool	337
Deleting a UUID Suffix Pool	338

CHAPTER 27**Setting the Management IP Address 339**

Management IP Address	339
Configuring the Management IP Address on a Blade Server	340
Configuring a Blade Server to Use a Static IP Address	340
Configuring a Blade Server to Use the Management IP Pool	341
Configuring the Management IP Address on a Rack Server	341
Configuring a Rack Server to Use a Static IP Address	341
Configuring a Rack Server to Use the Management IP Pool	342
Setting the Management IP Address on a Service Profile or Service Profile Template	343
Configuring the Management IP Pool	344
Management IP Pool	344
Configuring an IP Address Block for the Management IP Pool	344
Deleting an IP Address Block from the Management IP Pool	345

CHAPTER 28**Configuring Server-Related Policies 347**

Configuring BIOS Settings	347
Server BIOS Settings	347
Main BIOS Settings	348
Processor BIOS Settings	350
Intel Directed I/O BIOS Settings	356
RAS Memory BIOS Settings	358
Serial Port BIOS Settings	360
USB BIOS Settings	360
PCI Configuration BIOS Settings	361
Boot Options BIOS Settings	362
Server Management BIOS Settings	363
BIOS Policy	368
Default BIOS Settings	368

Creating a BIOS Policy	369
Modifying BIOS Defaults	370
Viewing the Actual BIOS Settings for a Server	371
Configuring IPMI Access Profiles	372
IPMI Access Profile	372
Configuring an IPMI Access Profile	372
Deleting an IPMI Access Profile	374
Adding an Endpoint User to an IPMI Access Profile	374
Deleting an Endpoint User from an IPMI Access Profile	375
Configuring Local Disk Configuration Policies	376
Local Disk Configuration Policy	376
Guidelines for all Local Disk Configuration Policies	376
Guidelines for Local Disk Configuration Policies Configured for RAID	377
Creating a Local Disk Configuration Policy	379
Viewing a Local Disk Configuration Policy	380
Deleting a Local Disk Configuration Policy	380
Configuring Scrub Policies	381
Scrub Policy	381
Creating a Scrub Policy	381
Deleting a Scrub Policy	382
Configuring Serial over LAN Policies	383
Serial over LAN Policy	383
Configuring a Serial over LAN Policy	383
Viewing a Serial over LAN Policy	384
Deleting a Serial over LAN Policy	385
Configuring Server Autoconfiguration Policies	385
Server Autoconfiguration Policy	385
Configuring a Server Autoconfiguration Policy	385
Deleting a Server Autoconfiguration Policy	387
Configuring Server Discovery Policies	387
Server Discovery Policy	387
Configuring a Server Discovery Policy	387
Deleting a Server Discovery Policy	389
Configuring Server Inheritance Policies	389
Server Inheritance Policy	389

Configuring a Server Inheritance Policy	389
Deleting a Server Inheritance Policy	391
Configuring Server Pool Policies	391
Server Pool Policy	391
Configuring a Server Pool Policy	391
Deleting a Server Pool Policy	392
Configuring Server Pool Policy Qualifications	393
Server Pool Policy Qualifications	393
Creating a Server Pool Policy Qualification	393
Deleting a Server Pool Policy Qualification	394
Creating an Adapter Qualification	395
Deleting an Adapter Qualification	396
Configuring a Chassis Qualification	396
Deleting a Chassis Qualification	397
Creating a CPU Qualification	398
Deleting a CPU Qualification	399
Creating a Power Group Qualification	400
Deleting a Power Group Qualification	400
Creating a Memory Qualification	401
Deleting a Memory Qualification	402
Creating a Physical Qualification	402
Deleting a Physical Qualification	403
Creating a Storage Qualification	403
Deleting a Storage Qualification	405
Configuring vNIC/vHBA Placement Policies	405
vNIC/vHBA Placement Policies	405
vCon to Adapter Placement	406
vNIC/vHBA to vCon Assignment	407
Configuring a vNIC/vHBA Placement Policy	409
Deleting a vNIC/vHBA Placement Policy	410
Explicitly Assigning a vNIC to a vCon	410
Explicitly Assigning a vHBA to a vCon	411

CHAPTER 29
Configuring Server Boot 413

Boot Policy	413
-------------	-----

Creating a Boot Policy	414
SAN Boot	415
Configuring a SAN Boot for a Boot Policy	416
iSCSI Boot	418
iSCSI Boot Process	418
iSCSI Boot Guidelines and Prerequisites	418
Enabling MPIO on Windows	420
Configuring iSCSI Boot	420
Creating an iSCSI Adapter Policy	422
Deleting an iSCSI Adapter Policy	424
Creating an Authentication Profile	424
Deleting an Authentication Profile	425
Adding a Block of IP Addresses to the Initiator Pool	425
Deleting a Block of IP Addresses from the Initiator Pool	426
Creating an iSCSI Boot Policy	427
Deleting iSCSI Devices from a Boot Policy	429
Creating an iSCSI vNIC in a Service Profile	429
Deleting an iSCSI vNIC from a Service Profile	431
Creating an iSCSI Initiator that Boots Using a Static IP Address	432
Deleting the Static IP Address Boot Parameters from an iSCSI Initiator	433
Creating an iSCSI Initiator that Boots Using an IP Address from an IP Pool	434
Deleting the IP Pool Boot Parameter from an iSCSI Initiator	435
Creating an iSCSI Initiator that Boots Using DHCP	436
Deleting the DHCP Boot Parameter from an iSCSI Initiator	437
IQN Pools	438
Creating an IQN Pool	438
Adding a Block to an IQN Pool	440
Deleting a Block from an IQN Pool	440
Deleting an IQN Pool	441
Viewing IQN Pool Usage	442
Creating an iSCSI Static Target	442
Deleting an iSCSI Static Target	445
Creating an iSCSI Auto Target	445
Deleting an iSCSI Auto Target	447
Verifying iSCSI Boot	447

LAN Boot	448
Configuring a LAN Boot for a Boot Policy	448
Local Disk Boot	449
Configuring a Local Disk Boot for a Boot Policy	449
Virtual Media Boot	450
Configuring a Virtual Media Boot for a Boot Policy	450
Deleting a Boot Policy	451

CHAPTER 30

Deferring Deployment of Service Profile Updates	453
Deferred Deployment of Service Profiles	453
Deferred Deployment Schedules	454
Maintenance Policy	454
Pending Activities	455
Guidelines and Limitations for Deferred Deployment	455
Configuring Schedules	456
Creating a Schedule	456
Creating a One Time Occurrence for a Schedule	456
Creating a Recurring Occurrence for a Schedule	457
Deleting a One Time Occurrence from a Schedule	459
Deleting a Recurring Occurrence from a Schedule	459
Deleting a Schedule	460
Configuring Maintenance Policies	461
Creating a Maintenance Policy	461
Deleting a Maintenance Policy	462
Managing Pending Activities	462
Viewing Pending Activities	462
Deploying a Service Profile Change Waiting for User Acknowledgement	463
Deploying a Scheduled Service Profile Change Immediately	463

CHAPTER 31

Configuring Service Profiles	465
Service Profiles that Override Server Identity	465
Service Profiles that Inherit Server Identity	466
Guidelines and Recommendations for Service Profiles	466
Service Profile Templates	467
Creating a Service Profile Template	467

Creating a Service Profile Instance from a Service Profile Template	470
Creating a Hardware-Based Service Profile	471
Configuring a vNIC for a Service Profile	474
Configuring a vHBA for a Service Profile	476
Configuring a Local Disk for a Service Profile	478
Configuring Serial over LAN for a Service Profile	479
Service Profile Boot Definition Configuration	480
Configuring a Boot Definition for a Service Profile	480
Configuring a LAN Boot for a Service Profile Boot Definition	481
Configuring a Storage Boot for a Service Profile Boot Definition	482
Configuring a Virtual Media Boot for a Service Profile Boot Definition	483
Deleting a Boot Definition for a Service Profile	484
Service Profiles and Service Profile Template Management	485
Associating a Service Profile with a Blade Server or Server Pool	485
Associating a Service Profile with a Rack Server	486
Disassociating a Service Profile from a Server or Server Pool	486
Resetting the UUID Assigned to a Service Profile from a Pool in a Service Profile Template	487
Resetting the MAC Address Assigned to a vNIC from a Pool in a Service Profile Template	488
Resetting the WWPN Assigned to a vHBA from a Pool in a Service Profile Template	489

CHAPTER 32
Managing Power in Cisco UCS 491

Power Management in Cisco UCS	491
Rack Server Power Management	491
Power Management Precautions	491
Configuring the Power Policy	492
Power Policy	492
Configuring the Power Policy	492
Configuring the Global Cap Policy	493
Global Cap Policy	493
Configuring the Global Cap Policy	493
Configuring Policy-Driven Chassis Group Power Capping	493
Policy-Driven Chassis Group Power Capping	493
Power Groups	494

Creating a Power Group	494
Deleting a Power Group	495
Power Control Policy	496
Creating a Power Control Policy	496
Deleting a Power Control Policy	497
Configuring Manual Blade-Level Power Capping	497
Manual Blade-Level Power Capping	497
Setting the Blade-Level Power Cap for a Server	498
Viewing the Blade-Level Power Cap	498

PART VI
System Management 501

CHAPTER 33
Managing Time Zones 503

Time Zones	503
Setting the Time Zone	503
Adding an NTP Server	505
Deleting an NTP Server	506
Setting the System Clock Manually	506

CHAPTER 34
Managing the Chassis 507

Guidelines for Removing and Decommissioning Chassis	507
Acknowledging a Chassis	508
Decommissioning a Chassis	508
Removing a Chassis	509
Recommissioning a Chassis	509
Renumbering a Chassis	510
Toggling the Locator LED	512
Turning On the Locator LED for a Chassis	512
Turning Off the Locator LED for a Chassis	512

CHAPTER 35
Managing Blade Servers 513

Blade Server Management	513
Guidelines for Removing and Decommissioning Blade Servers	514
Booting a Blade Server	514
Shutting Down a Blade Server	515

Power Cycling a Blade Server	516
Performing a Hard Reset on a Blade Server	516
Avoiding Unexpected Server Power Changes	517
Acknowledging a Blade Server	518
Removing a Blade Server from a Chassis	519
Decommissioning a Blade Server	519
Turning On the Locator LED for a Blade Server	520
Turning Off the Locator LED for a Blade Server	520
Resetting the CMOS for a Blade Server	520
Resetting the CIMC for a Blade Server	521
Recovering the Corrupt BIOS on a Blade Server	522
Issuing an NMI from a Blade Server	522

CHAPTER 36
Managing Rack-Mount Servers 525

Rack-Mount Server Management	525
Guidelines for Removing and Decommissioning Rack-Mount Servers	526
Booting a Rack-Mount Server	526
Shutting Down a Rack-Mount Server	527
Power Cycling a Rack-Mount Server	528
Performing a Hard Reset on a Rack-Mount Server	528
Avoiding Unexpected Server Power Changes	529
Acknowledging a Rack-Mount Server	530
Decommissioning a Rack-Mount Server	531
Renumbering a Rack-Mount Server	531
Removing a Rack-Mount Server	532
Turning On the Locator LED for a Rack-Mount Server	533
Turning Off the Locator LED for a Rack-Mount Server	534
Resetting the CMOS for a Rack-Mount Server	534
Resetting the CIMC for a Rack-Mount Server	535
Recovering the Corrupt BIOS on a Rack-Mount Server	535
Showing the Status for a Rack-Mount Server	536
Issuing an NMI from a Rack-Mount Server	536

CHAPTER 37
Managing the I/O Modules 539

I/O Module Management in Cisco UCS Manager GUI	539
--	-----

Resetting the I/O Module 539

CHAPTER 38

Backing Up and Restoring the Configuration 541

Backup and Export Configuration 541

Backup Types 541

Considerations and Recommendations for Backup Operations 542

Import Configuration 543

Import Methods 543

System Restore 543

Required User Role for Backup and Import Operations 544

Backup Operations 544

 Creating a Backup Operation 544

 Running a Backup Operation 545

 Modifying a Backup Operation 545

 Deleting a Backup Operation 547

Import Operations 548

 Creating an Import Operation 548

 Running an Import Operation 549

 Modifying an Import Operation 550

 Deleting an Import Operation 551

Restoring the Configuration for a Fabric Interconnect 552

Erasing the Configuration 553

CHAPTER 39

Recovering a Lost Password 555

Password Recovery for the Admin Account 555

Determining the Leadership Role of a Fabric Interconnect 556

Recovering the Admin Account Password in a Standalone Configuration 556

Recovering the Admin Account Password in a Cluster Configuration 557

PART VII

System Monitoring 561

CHAPTER 40

Monitoring Traffic 563

Traffic Monitoring 563

Guidelines and Recommendations for Traffic Monitoring 564

Creating an Ethernet Traffic Monitoring Session 565

Creating a Fibre Channel Traffic Monitoring Session	566
Adding Traffic Sources to a Monitoring Session	567
Adding an Uplink Source Port to a Monitoring Session	567
Adding a vNIC or vHBA Source to a Monitoring Session	568
Adding a VLAN or VSAN Source to a Monitoring Session	570
Adding a Storage Port Source to a Monitoring Session	571
Activating a Traffic Monitoring Session	572
Deleting a Traffic Monitoring Session	573

CHAPTER 41**Monitoring Hardware 575**

Monitoring Fan Modules	575
Monitoring Management Interfaces	577
Management Interfaces Monitoring Policy	577
Configuring the Management Interfaces Monitoring Policy	578
Server Disk Drive Monitoring	579
Support for Disk Drive Monitoring	580
Prerequisites for Disk Drive Monitoring	580
Viewing the Status of a Disk Drive	580
Interpreting the Status of a Monitored Disk Drive	582

CHAPTER 42**Configuring Statistics-Related Policies 585**

Configuring Statistics Collection Policies	585
Statistics Collection Policy	585
Configuring a Statistics Collection Policy	586
Configuring Statistics Threshold Policies	586
Statistics Threshold Policy	586
Server and Server Component Statistics Threshold Policy Configuration	587
Configuring a Server and Server Component Statistics Threshold Policy	587
Deleting a Server and Server Component Statistics Threshold Policy	588
Configuring a Server and Server Component Statistics Threshold Policy Class	588
Deleting a Server and Server Component Statistics Threshold Policy Class	590
Uplink Ethernet Port Statistics Threshold Policy Configuration	590
Configuring an Uplink Ethernet Port Statistics Threshold Policy	590
Configuring an Uplink Ethernet Port Statistics Threshold Policy Class	591
Deleting an Uplink Ethernet Port Statistics Threshold Policy Class	593

Server Port, Chassis, and Fabric Interconnect Statistics Threshold Policy	
Configuration	593
Configuring a Server Port, Chassis, and Fabric Interconnect Statistics Threshold Policy	593
Configuring a Server Port, Chassis, and Fabric Interconnect Statistics Threshold Policy Class	594
Deleting a Server Port, Chassis, and Fabric Interconnect Statistics Threshold Policy Class	596
Fibre Channel Port Statistics Threshold Policy Configuration	596
Configuring a Fibre Channel Port Statistics Threshold Policy	596
Configuring a Fibre Channel Port Statistics Threshold Policy Class	597
Deleting an Uplink Fibre Channel Port Statistics Threshold Policy Class	599

CHAPTER 43
Configuring Call Home 601

Call Home	601
Call Home Considerations and Guidelines	603
Cisco UCS Faults and Call Home Severity Levels	604
Cisco Smart Call Home	605
Configuring Call Home	606
Disabling Call Home	608
Enabling Call Home	608
Configuring System Inventory Messages	609
Configuring System Inventory Messages	609
Sending a System Inventory Message	609
Configuring Call Home Profiles	610
Call Home Profiles	610
Configuring a Call Home Profile	611
Deleting a Call Home Profile	612
Sending a Test Call Home Alert	613
Configuring Call Home Policies	614
Call Home Policies	614
Configuring a Call Home Policy	614
Disabling a Call Home Policy	615
Enabling a Call Home Policy	615
Deleting a Call Home Policy	616

Example: Configuring Call Home for Smart Call Home	616
Configuring Smart Call Home	616
Configuring the Default Cisco TAC-1 Profile	618
Configuring a System Inventory Message for Smart Call Home	619
Registering Smart Call Home	620

CHAPTER 44**Managing the System Event Log 623**

System Event Log	623
Viewing the System Event Log for a Server	624
Viewing the System Event Log for an Individual Server	624
Viewing the System Event Log for All of the Servers in a Chassis	624
Configuring the SEL Policy	625
Backing Up the System Event Log for a Server	627
Backing Up the System Event Log for an Individual Server	627
Backing Up the System Event Log for All of the Servers in a Chassis	627
Clearing the System Event Log for a Server	628
Clearing the System Event Log for an Individual Server	628
Clearing the System Event Log for All of the Servers in a Chassis	628

CHAPTER 45**Configuring Settings for Faults, Events, and Logs 631**

Configuring Settings for the Fault Collection Policy	631
Fault Collection Policy	631
Configuring the Fault Collection Policy	632
Configuring Settings for the Core File Exporter	632
Core File Exporter	632
Configuring the Core File Exporter	633
Disabling the Core File Exporter	634
Configuring the Syslog	634



Preface

This preface includes the following sections:

- [Audience, page xxxi](#)
- [Conventions, page xxxi](#)
- [Related Cisco UCS Documentation, page xxxiii](#)
- [Documentation Feedback, page xxxiii](#)
- [Obtaining Documentation and Submitting a Service Request, page xxxiii](#)

Audience

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

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

Conventions

This document uses the following conventions:

Convention	Indication
bold font	Commands, keywords, GUI elements, and user-entered text appear in bold font .
<i>italic font</i>	Document titles, new or emphasized terms, and arguments for which you supply values are in <i>italic font</i> .

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

**Note**

Means *reader take note*.

**Tip**

Means *the following information will help you solve a problem*.

**Caution**

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

**Timesaver**

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

**Warning**

Means *reader be warned*. In this situation, you might perform an action that could result in bodily injury.

Related Cisco UCS Documentation

Documentation Roadmaps

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

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

Other Documentation Resources

An ISO file containing all B and C-Series documents is available at the following URL: <http://www.cisco.com/cisco/software/type.html?mdfid=283853163&flowid=25821>. From this page, click **Unified Computing System (UCS) Documentation Roadmap Bundle**.

The ISO file is updated after every major documentation release.

Follow [Cisco UCS Docs on Twitter](#) to receive document update notifications.

Documentation Feedback

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

Obtaining Documentation and Submitting a Service Request

For information on obtaining documentation, submitting a service request, and gathering additional information, see the monthly [What's New in Cisco Product Documentation](#), which also lists all new and revised Cisco technical documentation.

Subscribe to the *What's New in Cisco Product Documentation* as a Really Simple Syndication (RSS) feed and set content to be delivered directly to your desktop using a reader application. The RSS feeds are a free service and Cisco currently supports RSS version 2.0.

Follow [Cisco UCS Docs on Twitter](#) to receive document update notifications.



PART **I**

Introduction

- [New and Changed Information, page 3](#)
- [Overview of Cisco Unified Computing System, page 9](#)
- [Overview of Cisco UCS Manager, page 43](#)
- [Overview of Cisco UCS Manager CLI, page 47](#)



CHAPTER 1

New and Changed Information

This chapter includes the following sections:

- [New and Changed Information for this Release, page 3](#)

New and Changed Information for this Release

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

Table 1: New Features and Significant Behavioral Changes in Cisco UCS, Release 2.0(3)

Feature	Description	Where Documented
Cipher Suite	Adds support for Cipher Suite in HTTPS configuration.	Configuring Communication Services, on page 103
Web Session Refresh	Enables you to configure the web session refresh period and timeout for authentication domains.	Configuring Authentication, on page 125
BIOS Settings	Adds support for new BIOS settings that can be included in BIOS policies and configured from Cisco UCS Manager.	Configuring Server-Related Policies, on page 347
Overview of enabling MPIO	High level information added for how to enable MPIO with iSCSI boot.	Enabling MPIO on Windows, on page 420

Table 2: New Features and Significant Behavioral Changes in Cisco UCS, Release 2.0(2)

Feature	Description	Where Documented
IQN Pools	Adds support for IQN pools in Cisco UCS domains configured for iSCSI boot.	iSCSI Boot, on page 418
Adapter Port Channels	Enables you to group all the physical links from a Cisco UCS Virtual Interface Card (VIC) to an I/O Module into one logical link. (Requires supported hardware.)	Configuring Ports and Port Channels, on page 71
Unified Port Support for 6296 Fabric Interconnect	Enables you to use the Configure Unified Ports wizard to configure ports on a 6296 fabric interconnect.	Unified Ports on the 6200 Series Fabric Interconnect, on page 72
Renumbering for Rack-Mount Servers	Enables you to renumber an integrated rack-mount server.	Managing Rack-Mount Servers, on page 525
Changes to Behavior for Power State Synchronization	Adds information and a caution about power state synchronization, including use of the physical power button or the reset feature on a blade server or an integrated rack-mount server.	Managing Blade Servers, on page 513 Managing Rack-Mount Servers, on page 525
BIOS Settings	Adds support for new BIOS settings that can be included in BIOS policies and configured from Cisco UCS Manager.	Configuring Server-Related Policies, on page 347

Table 3: New Features in Cisco UCS, Release 2.0(1)

Feature	Description	Where Documented
Disk Drive Monitoring Support	Support for disk drive monitoring on certain blade servers and a specific LSI storage controller firmware level.	Monitoring Hardware, on page 575
Fabric Port Channels	Enables you to group several of the physical links from a IOM to a fabric interconnect into one logical link for redundancy and bandwidth sharing. (Requires supported hardware.)	Configuring Ports and Port Channels, on page 71
Firmware Bundle Option	Enables you to select a bundle instead of a version when updating firmware using the Cisco UCS Manager GUI.	Managing Firmware, on page 185

Feature	Description	Where Documented
iSCSI Boot	iSCSI boot enables a server to boot its operating system from an iSCSI target machine located remotely over a network.	iSCSI Boot, on page 418
Licensing	Updated information for new UCS hardware.	Licenses, on page 251
Pre-login Banner	Displays user-defined banner text prior to login when a user logs into Cisco UCS Manager using the GUI or CLI.	Pre-Login Banner, on page 52
Unified Ports	Unified ports are ports on the 6200 series fabric interconnect that can be configured to carry either Ethernet or Fibre Channel traffic.	Unified Ports on the 6200 Series Fabric Interconnect, on page 72
Upstream Disjoint Layer-2 Networks	Enables you to configure Cisco UCS to communicate with upstream disjoint layer-2 networks.	Configuring Upstream Disjoint Layer-2 Networks, on page 301
Virtual Interfaces	The number of vNICs and vHBAs configurable for a service profile is determined by adapter capability and the amount of virtual interface (VIF) namespace available on the adapter.	Managing Virtual Interfaces, on page 259

Feature	Description	Where Documented
Virtual Interface Card Drivers	Cisco UCS Virtual Interface Card (VIC) drivers facilitate communication between supported operating systems and Cisco UCS Virtual Interface Cards (VICs).	<p>This feature is now documented in the following installation guides:</p> <ul style="list-style-type: none"> • <i>Cisco UCS Manager Interface Card Drivers for ESX Installation Guide</i> • <i>Cisco UCS Manager Interface Card Drivers for Linux Installation Guide</i> • <i>Cisco UCS Manager Interface Card Drivers for Windows Installation Guide</i> <p>The VIC driver installation guides can be found here: http://www.cisco.com/en/US/products/ps10281/prod_installation_guides_list.html</p>
VM-FEX Integration for VMware	<p>Cisco Virtual Machine Fabric Extender (VM-FEX) for VMware provides management integration and network communication between Cisco UCS Manager and VMware vCenter.</p> <p>In previous releases, this functionality was known as VN-Link in Hardware.</p>	<p>This feature is now documented in the following configuration guides:</p> <ul style="list-style-type: none"> • <i>Cisco UCS Manager VM-FEX for VMware GUI Configuration Guide</i> • <i>Cisco UCS Manager VM-FEX for VMware CLI Configuration Guide</i> <p>The VM-FEX configuration guides can be found here: http://www.cisco.com/en/US/products/ps10281/products_installation_and_configuration_guides_list.html</p>

Feature	Description	Where Documented
VM-FEX Integration for KVM (Red Hat Linux)	Cisco Virtual Machine Fabric Extender (VM-FEX) for VMware provides external switching for virtual machines running on a KVM Linux-based hypervisor in a Cisco UCS domain.	<p>This feature is documented in the following configuration guides:</p> <ul style="list-style-type: none">• <i>Cisco UCS Manager VM-FEX for KVM GUI Configuration Guide</i>• <i>Cisco UCS Manager VM-FEX for KVM CLI Configuration Guide</i> <p>The VM-FEX configuration guides can be found here: http://www.cisco.com/en/US/products/ps10281/products_installation_and_configuration_guides_list.html</p>



CHAPTER 2

Overview of Cisco Unified Computing System

This chapter includes the following sections:

- [About Cisco Unified Computing System](#) , page 9
- [Unified Fabric](#), page 10
- [Server Architecture and Connectivity](#), page 12
- [Traffic Management](#), page 33
- [Opt-In Features](#), page 38
- [Virtualization in Cisco UCS](#) , page 40

About Cisco Unified Computing System

Cisco Unified Computing System (Cisco UCS) fuses access layer networking and servers. This high-performance, next-generation server system provides a data center with a high degree of workload agility and scalability.

The hardware and software components support Cisco's unified fabric, which runs multiple types of data center traffic over a single converged network adapter.

Architectural Simplification

The simplified architecture of Cisco UCS reduces the number of required devices and centralizes switching resources. By eliminating switching inside a chassis, network access-layer fragmentation is significantly reduced.

Cisco UCS implements Cisco unified fabric within racks and groups of racks, supporting Ethernet and Fibre Channel protocols over 10 Gigabit Cisco Data Center Ethernet and Fibre Channel over Ethernet (FCoE) links.

This radical simplification reduces the number of switches, cables, adapters, and management points by up to two-thirds. All devices in a Cisco UCS domain remain under a single management domain, which remains highly available through the use of redundant components.

High Availability

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

Scalability

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

Flexibility

A Cisco UCS domain allows you to quickly align computing resources in the data center with rapidly changing business requirements. This built-in flexibility is determined by whether you choose to fully implement the stateless computing feature.

Pools of servers and other system resources can be applied as necessary to respond to workload fluctuations, support new applications, scale existing software and business services, and accommodate both scheduled and unscheduled downtime. Server identity can be abstracted into a mobile service profile that can be moved from server to server with minimal downtime and no need for additional network configuration.

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

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

Optimized for Server Virtualization

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

Unified Fabric

With unified fabric, multiple types of data center traffic can run over a single Data Center Ethernet (DCE) network. Instead of having a series of different host bus adapters (HBAs) and network interface cards (NICs) present in a server, unified fabric uses a single converged network adapter. This type of adapter can carry LAN and SAN traffic on the same cable.

Cisco UCS uses Fibre Channel over Ethernet (FCoE) to carry Fibre Channel and Ethernet traffic on the same physical Ethernet connection between the fabric interconnect and the server. This connection terminates at a converged network adapter on the server, and the unified fabric terminates on the uplink ports of the fabric interconnect. On the core network, the LAN and SAN traffic remains separated. Cisco UCS does not require that you implement unified fabric across the data center.

The converged network adapter presents an Ethernet interface and Fibre Channel interface to the operating system. At the server, the operating system is not aware of the FCoE encapsulation because it sees a standard Fibre Channel HBA.

At the fabric interconnect, the server-facing Ethernet port receives the Ethernet and Fibre Channel traffic. The fabric interconnect (using Ethertype to differentiate the frames) separates the two traffic types. Ethernet frames and Fibre Channel frames are switched to their respective uplink interfaces.

Fibre Channel over Ethernet

Cisco UCS leverages Fibre Channel over Ethernet (FCoE) standard protocol to deliver Fibre Channel. The upper Fibre Channel layers are unchanged, so the Fibre Channel operational model is maintained. FCoE network management and configuration is similar to a native Fibre Channel network.

FCoE encapsulates Fibre Channel traffic over a physical Ethernet link. FCoE is encapsulated over Ethernet with the use of a dedicated Ethertype, 0x8906, so that FCoE traffic and standard Ethernet traffic can be carried on the same link. FCoE has been standardized by the ANSI T11 Standards Committee.

Fibre Channel traffic requires a lossless transport layer. Instead of the buffer-to-buffer credit system used by native Fibre Channel, FCoE depends upon the Ethernet link to implement lossless service.

Ethernet links on the fabric interconnect provide two mechanisms to ensure lossless transport for FCoE traffic:

- Link-level flow control
- Priority flow control

Link-Level Flow Control

IEEE 802.3x link-level flow control allows a congested receiver to signal the endpoint to pause data transmission for a short time. This link-level flow control pauses all traffic on the link.

The transmit and receive directions are separately configurable. By default, link-level flow control is disabled for both directions.

On each Ethernet interface, the fabric interconnect can enable either priority flow control or link-level flow control (but not both).

Priority Flow Control

The priority flow control (PFC) feature applies pause functionality to specific classes of traffic on the Ethernet link. For example, PFC can provide lossless service for the FCoE traffic, and best-effort service for the standard Ethernet traffic. PFC can provide different levels of service to specific classes of Ethernet traffic (using IEEE 802.1p traffic classes).

PFC decides whether to apply pause based on the IEEE 802.1p CoS value. When the fabric interconnect enables PFC, it configures the connected adapter to apply the pause functionality to packets with specific CoS values.

By default, the fabric interconnect negotiates to enable the PFC capability. If the negotiation succeeds, PFC is enabled and link-level flow control remains disabled (regardless of its configuration settings). If the PFC negotiation fails, you can either force PFC to be enabled on the interface or you can enable IEEE 802.x link-level flow control.

Server Architecture and Connectivity

Overview of Service Profiles

Service profiles are the central concept of Cisco UCS. Each service profile serves a specific purpose: ensuring that the associated server hardware has the configuration required to support the applications it will host.

The service profile maintains configuration information about the server hardware, interfaces, fabric connectivity, and server and network identity. This information is stored in a format that you can manage through Cisco UCS Manager. All service profiles are centrally managed and stored in a database on the fabric interconnect.

Every server must be associated with a service profile.



Important

At any given time, each server can be associated with only one service profile. Similarly, each service profile can be associated with only one server at a time.

After you associate a service profile with a server, the server is ready to have an operating system and applications installed, and you can use the service profile to review the configuration of the server. If the server associated with a service profile fails, the service profile does not automatically fail over to another server.

When a service profile is disassociated from a server, the identity and connectivity information for the server is reset to factory defaults.

Network Connectivity through Service Profiles

Each service profile specifies the LAN and SAN network connections for the server through the Cisco UCS infrastructure and out to the external network. You do not need to manually configure the network connections for Cisco UCS servers and other components. All network configuration is performed through the service profile.

When you associate a service profile with a server, the Cisco UCS internal fabric is configured with the information in the service profile. If the profile was previously associated with a different server, the network infrastructure reconfigures to support identical network connectivity to the new server.

Configuration through Service Profiles

A service profile can take advantage of resource pools and policies to handle server and connectivity configuration.

Hardware Components Configured by Service Profiles

When a service profile is associated with a server, the following components are configured according to the data in the profile:

- Server, including BIOS and CIMC
- Adapters
- Fabric interconnects

You do not need to configure these hardware components directly.

Server Identity Management through Service Profiles

You can use the network and device identities burned into the server hardware at manufacture or you can use identities that you specify in the associated service profile either directly or through identity pools, such as MAC, WWN, and UUID.

The following are examples of configuration information that you can include in a service profile:

- Profile name and description
- Unique server identity (UUID)
- LAN connectivity attributes, such as the MAC address
- SAN connectivity attributes, such as the WWN

Operational Aspects configured by Service Profiles

You can configure some of the operational functions for a server in a service profile, such as the following:

- Firmware packages and versions
- Operating system boot order and configuration
- IPMI and KVM access

vNIC Configuration by Service Profiles

A vNIC is a virtualized network interface that is configured on a physical network adapter and appears to be a physical NIC to the operating system of the server. The type of adapter in the system determines how many vNICs you can create. For example, a converged network adapter has two NICs, which means you can create a maximum of two vNICs for each adapter.

A vNIC communicates over Ethernet and handles LAN traffic. At a minimum, each vNIC must be configured with a name and with fabric and network connectivity.

vHBA Configuration by Service Profiles

A vHBA is a virtualized host bus adapter that is configured on a physical network adapter and appears to be a physical HBA to the operating system of the server. The type of adapter in the system determines how many vHBAs you can create. For example, a converged network adapter has two HBAs, which means you can create a maximum of two vHBAs for each of those adapters. In contrast, a network interface card does not have any HBAs, which means you cannot create any vHBAs for those adapters.

A vHBA communicates over FCoE and handles SAN traffic. At a minimum, each vHBA must be configured with a name and fabric connectivity.

Service Profiles that Override Server Identity

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

You can disassociate this service profile from one server and then associate it with another server. This re-association can be done either manually or through an automated server pool policy. The burned-in settings,

such as UUID and MAC address, on the new server are overwritten with the configuration in the service profile. As a result, the change in server is transparent to your network. You do not need to reconfigure any component or application on your network to begin using the new server.

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

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

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

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



Note

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

Service Profiles that Inherit Server Identity

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

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

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

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



Important

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

Service Profile Templates

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

**Tip**

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

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

Cisco UCS supports the following types of service profile templates:

Initial template

Service profiles created from an initial template inherit all the properties of the template. However, after you create the profile, it is no longer connected to the template. If you need to make changes to one or more profiles created from this template, you must change each profile individually.

Updating template

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

Policies

Policies determine how Cisco UCS components will act in specific circumstances. You can create multiple instances of most policies. For example, you might want different boot policies, so that some servers can PXE boot, some can SAN boot, and others can boot from local storage.

Policies allow separation of functions within the system. A subject matter expert can define policies that are used in a service profile, which is created by someone without that subject matter expertise. For example, a LAN administrator can create adapter policies and quality of service policies for the system. These policies can then be used in a service profile that is created by someone who has limited or no subject matter expertise with LAN administration.

You can create and use two types of policies in Cisco UCS Manager:

- Configuration policies that configure the servers and other components
- Operational policies that control certain management, monitoring, and access control functions

Configuration Policies

Boot Policy

The boot policy determines the following:

- Configuration of the boot device

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

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

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



Important

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

Chassis Discovery Policy

The chassis discovery policy determines how the system reacts when you add a new chassis. Cisco UCS Manager uses the settings in the chassis discovery policy to determine the minimum threshold for the number of links between the chassis and the fabric interconnect and whether to group links from the IOM to the fabric interconnect in a fabric port channel.

Chassis Links

If you have a Cisco UCS domain that has some chassis wired with 1 link, some with 2 links, some with 4 links, and some with 8 links we recommend that you configure the chassis discovery policy for the minimum number links in the domain so that Cisco UCS Manager can discover all chassis.



Tip

For Cisco UCS implementations that mix IOMs with different numbers of links, we recommend using the platform max value. Using platform max insures that Cisco UCS Manager uses the maximum number of IOM uplinks available.

After the initial discovery, you must reacknowledge the chassis that are wired for a greater number of links and Cisco UCS Manager configures the chassis to use all available links.

Cisco UCS Manager cannot discover any chassis that is wired for fewer links than are configured in the chassis discovery policy. For example, if the chassis discovery policy is configured for 4 links, Cisco UCS Manager cannot discover any chassis that is wired for 1 link or 2 links. Reacknowledgement of the chassis does not resolve this issue.

The following table provides an overview of how the chassis discovery policy works in a multi-chassis Cisco UCS domain:

Table 4: Chassis Discovery Policy and Chassis Links

Number of Links Wired for the Chassis	1-Link Chassis Discovery Policy	2-Link Chassis Discovery Policy	4-Link Chassis Discovery Policy	8-Link Chassis Discovery Policy	Platform-Max Discovery Policy
1 link between IOM and fabric interconnects	Chassis is discovered by Cisco UCS Manager and added to the Cisco UCS domain as a chassis wired with 1 link.	Chassis cannot be discovered by Cisco UCS Manager and is not added to the Cisco UCS domain.	Chassis cannot be discovered by Cisco UCS Manager and is not added to the Cisco UCS domain.	Chassis cannot be discovered by Cisco UCS Manager and is not added to the Cisco UCS domain.	Chassis is discovered by Cisco UCS Manager and added to the Cisco UCS domain as a chassis wired with 1 link.
2 links between IOM and fabric interconnects	Chassis is discovered by Cisco UCS Manager and added to the Cisco UCS domain as a chassis wired with 1 link. After initial discovery, reacknowledge the chassis and Cisco UCS Manager recognizes and uses the additional links.	Chassis is discovered by Cisco UCS Manager and added to the Cisco UCS domain as a chassis wired with 2 link.	Chassis cannot be discovered by Cisco UCS Manager and is not added to the Cisco UCS domain.	Chassis cannot be discovered by Cisco UCS Manager and is not added to the Cisco UCS domain.	Chassis cannot be discovered by Cisco UCS Manager and is not added to the Cisco UCS domain.

Number of Links Wired for the Chassis	1-Link Chassis Discovery Policy	2-Link Chassis Discovery Policy	4-Link Chassis Discovery Policy	8-Link Chassis Discovery Policy	Platform-Max Discovery Policy
4 links between IOM and fabric interconnects	Chassis is discovered by Cisco UCS Manager and added to the Cisco UCS domain as a chassis wired with 1 link. After initial discovery, reacknowledge the chassis and Cisco UCS Manager recognizes and uses the additional links.	Chassis is discovered by Cisco UCS Manager and added to the Cisco UCS domain as a chassis wired with 2 links. After initial discovery, reacknowledge the chassis and Cisco UCS Manager recognizes and uses the additional links.	Chassis is discovered by Cisco UCS Manager and added to the Cisco UCS domain as a chassis wired with 4 link.	Chassis cannot be discovered by Cisco UCS Manager and is not added to the Cisco UCS domain.	If the IOM has 4 links, the chassis is discovered by Cisco UCS Manager and added to the Cisco UCS domain as a chassis wired with 4 links. If the IOM has 8 links, the chassis is not fully discovered by Cisco UCS Manager.
8 links between IOM and fabric interconnects	Chassis is discovered by Cisco UCS Manager and added to the Cisco UCS domain as a chassis wired with 1 link. After initial discovery, reacknowledge the chassis and Cisco UCS Manager recognizes and uses the additional links.	Chassis is discovered by Cisco UCS Manager and added to the Cisco UCS domain as a chassis wired with 2 links. After initial discovery, reacknowledge the chassis and Cisco UCS Manager recognizes and uses the additional links.	Chassis is discovered by Cisco UCS Manager and added to the Cisco UCS domain as a chassis wired with 4 links. After initial discovery, reacknowledge the chassis and Cisco UCS Manager recognizes and uses the additional links.	Chassis is discovered by Cisco UCS Manager and added to the Cisco UCS domain as a chassis wired with 8 links.	Chassis is discovered by Cisco UCS Manager and added to the Cisco UCS domain as a chassis wired with 8 links.

Link Grouping

For hardware configurations that support fabric port channels, link grouping determines whether all of the links from the IOM to the fabric interconnect are grouped into a fabric port channel during chassis discovery. If the link grouping preference is set to port channel, all of the links from the IOM to the fabric interconnect

are grouped in a fabric port channel. If set to no group, links from the IOM to the fabric interconnect are not grouped in a fabric port channel.

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



Note The link grouping preference only takes effect if both sides of the links between an IOM or FEX and the fabric interconnect support fabric port channels. If one side of the links does not support fabric port channels, this preference is ignored and the links are not grouped in a port channel.

Dynamic vNIC Connection Policy

The dynamic vNIC connection policy determines how the connectivity between VMs and dynamic vNICs is configured. This policy is required for Cisco UCS domains that include servers with VIC adapters on which you have installed VMs and configured dynamic vNICs.

Each dynamic vNIC connection policy includes an Ethernet adapter policy and designates the number of vNICs that can be configured for any server associated with a service profile that includes the policy.



Note If you migrate a server that is configured with dynamic vNICs, the dynamic interface used by the vNICs fails and Cisco UCS Manager notifies you of that failure.

When the server comes back up, Cisco UCS Manager assigns new dynamic vNICs to the server. If you are monitoring traffic on the dynamic vNIC, you must reconfigure the monitoring source.

Ethernet and Fibre Channel Adapter Policies

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

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

**Note**

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

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

Operating System Specific Adapter Policies

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

**Important**

We recommend that you use the values in these policies for the applicable operating system. Do not modify any of the values in the default policies unless directed to do so by Cisco Technical Support.

However, if you are creating an Ethernet adapter policy for a Windows OS (instead of using the default Windows adapter policy), you must use the following formulas to calculate values that work with Windows:

$$\text{Completion Queues} = \text{Transmit Queues} + \text{Receive Queues}$$

$$\text{Interrupt Count} = (\text{Completion Queues} + 2) \text{ rounded up to nearest power of } 2$$

For example, if Transmit Queues = 1 and Receive Queues = 8 then:

$$\text{Completion Queues} = 1 + 8 = 9$$

$$\text{Interrupt Count} = (9 + 2) \text{ rounded up to the nearest power of } 2 = 16$$

Global Cap Policy

The global cap policy is a global policy that specifies whether policy-driven chassis group power capping or manual blade-level power capping will be applied to all servers in a chassis.

We recommend that you use the default power capping method: policy-driven chassis group power capping.

**Important**

Any change to the manual blade-level power cap configuration will result in the loss of any groups or configuration options set for policy-driven chassis group power capping.

Host Firmware Package

This policy enables you to specify a set of firmware versions that make up the host firmware package (also known as the host firmware pack). The host firmware includes the following firmware for server and adapter endpoints:

- **Adapter**
- **BIOS**
- **Board Controller**
- **FC Adapters**
- **HBA Option ROM**
- **Storage Controller**

**Tip**

You can include more than one type of firmware in the same host firmware package. For example, a host firmware package can include both BIOS firmware and storage controller firmware or adapter firmware for two different models of adapters. However, you can only have one firmware version with the same type, vendor, and model number. The system recognizes which firmware version is required for an endpoint and ignores all other firmware versions.

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

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

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

Prerequisites

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

IPMI Access Profile

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

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

Local Disk Configuration Policy

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

The local disk modes include the following:

- **No Local Storage**—For a diskless server or a SAN only configuration. If you select this option, you cannot associate any service profile which uses this policy with a server that has a local disk.
- **RAID 0 Striped**—Data is striped across all disks in the array, providing fast throughput. There is no data redundancy, and all data is lost if any disk fails.
- **RAID 1 Mirrored**—Data is written to two disks, providing complete data redundancy if one disk fails. The maximum array size is equal to the available space on the smaller of the two drives.
- **Any Configuration**—For a server configuration that carries forward the local disk configuration without any changes.
- **No RAID**—For a server configuration that removes the RAID and leaves the disk MBR and payload unaltered.
- **RAID 5 Striped Parity**—Data is striped across all disks in the array. Part of the capacity of each disk stores parity information that can be used to reconstruct data if a disk fails. RAID 5 provides good data throughput for applications with high read request rates.
- **RAID 6 Striped Dual Parity**—Data is striped across all disks in the array and two parity disks are used to provide protection against the failure of up to two physical disks. In each row of data blocks, two sets of parity data are stored.
- **RAID10 Mirrored and Striped**— RAID 10 uses mirrored pairs of disks to provide complete data redundancy and high throughput rates.

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

Management Firmware Package

This policy enables you to specify a set of firmware versions that make up the management firmware package (also known as a management firmware pack). The management firmware package includes the Cisco Integrated Management Controller (CIMC) on the server. You do not need to use this package if you upgrade the CIMC directly.

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

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

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

Management Interfaces Monitoring Policy

This policy defines how the mgmt0 Ethernet interface on the fabric interconnect should be monitored. If Cisco UCS detects a management interface failure, a failure report is generated. If the configured number of failure reports is reached, the system assumes that the management interface is unavailable and generates a fault. By default, the management interfaces monitoring policy is disabled.

If the affected management interface belongs to a fabric interconnect which is the managing instance, Cisco UCS confirms that the subordinate fabric interconnect's status is up, that there are no current failure reports logged against it, and then modifies the managing instance for the end-points.

If the affected fabric interconnect is currently the primary inside of a high availability setup, a failover of the management plane is triggered. The data plane is not affected by this failover.

You can set the following properties related to monitoring the management interface:

- Type of mechanism used to monitor the management interface.
- Interval at which the management interface's status is monitored.
- Maximum number of monitoring attempts that can fail before the system assumes that the management is unavailable and generates a fault message.



Important

In the event of a management interface failure on a fabric interconnect, the managing instance may not change if one of the following occurs:

- A path to the end-point through the subordinate fabric interconnect does not exist.
- The management interface for the subordinate fabric interconnect has failed.
- The path to the end-point through the subordinate fabric interconnect has failed.

Network Control Policy

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

- Whether the Cisco Discovery Protocol (CDP) is enabled or disabled
- How the VIF behaves if no uplink port is available in end-host mode
- The action that Cisco UCS Manager takes on the remote Ethernet interface, vEthernet interface, or vFibreChannel interface when the associated border port fails
- Whether the server can use different MAC addresses when sending packets to the fabric interconnect
- Whether MAC registration occurs on a per-VNIC basis or for all VLANs.

Action on Uplink Fail

By default, the **Action on Uplink Fail** property in the network control policy is configured with a value of link-down. For adapters such as the Cisco UCS M81KR Virtual Interface Card, this default behavior directs Cisco UCS Manager to bring the vEthernet or vFibreChannel interface down if the associated border port fails. For Cisco UCS systems using a non-VM-FEX capable converged network adapter that supports both Ethernet and FCoE traffic, such as Cisco UCS CNA M72KR-Q and the Cisco UCS CNA M72KR-E, this

default behavior directs Cisco UCS Manager to bring the remote Ethernet interface down if the associated border port fails. In this scenario, any vFibreChannel interfaces that are bound to the remote Ethernet interface are brought down as well.



Note Cisco UCS Manager, release 1.4(2) and earlier did not enforce the **Action on Uplink Fail** property for those types of non-VM-FEX capable converged network adapters mentioned above. If the **Action on Uplink Fail** property was set to link-down, Cisco UCS Manager would ignore this setting and instead issue a warning. In the current version of Cisco UCS Manager this setting is enforced. Therefore, if your implementation includes one of those converged network adapters and the adapter is expected to handle both Ethernet and FCoE traffic, we recommend that you configure the **Action on Uplink Fail** property with a value of warning.

Please note that this configuration may result in an Ethernet teaming driver not being able to detect a link failure when the border port goes down.

MAC Registration Mode

In Cisco UCS Manager, releases 1.4 and earlier, MAC addresses were installed on all of the VLANs belonging to an interface. Starting in release 2.0, MAC addresses are installed only on the native VLAN by default. In most implementations this maximizes the VLAN port count.



Note If a trunking driver is being run on the host and the interface is in promiscuous mode, we recommend that you set the Mac Registration Mode to All VLANs.

Power Control Policy

Cisco UCS uses the priority set in the power control policy, along with the blade type and configuration, to calculate the initial power allocation for each blade within a chassis. During normal operation, the active blades within a chassis can borrow power from idle blades within the same chassis. If all blades are active and reach the power cap, service profiles with higher priority power control policies take precedence over service profiles with lower priority power control policies.

Priority is ranked on a scale of 1-10, where 1 indicates the highest priority and 10 indicates lowest priority. The default priority is 5.

For mission-critical application a special priority called no-cap is also available. Setting the priority to no-cap prevents Cisco UCS from leveraging unused power from that particular blade server. The server is allocated the maximum amount of power that that blade can reach.



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

Power Policy

The power policy is a global policy that specifies the redundancy for power supplies in all chassis in the Cisco UCS domain. This policy is also known as the PSU policy.

For more information about power supply redundancy, see *Cisco UCS 5108 Server Chassis Hardware Installation Guide*.

Quality of Service Policy

A quality of service (QoS) policy assigns a system class to the outgoing traffic for a vNIC or vHBA. This system class determines the quality of service for that traffic. For certain adapters you can also specify additional controls on the outgoing traffic, such as burst and rate.

You must include a QoS policy in a vNIC policy or vHBA policy and then include that policy in a service profile to configure the vNIC or vHBA.

Rack Server Discovery Policy

The rack server discovery policy determines how the system reacts when you add a new rack-mount server. Cisco UCS Manager uses the settings in the rack server discovery policy to determine whether any data on the hard disks are scrubbed and whether server discovery occurs immediately or needs to wait for explicit user acknowledgement.

Cisco UCS Manager cannot discover any rack-mount server that has not been correctly cabled and connected to the fabric interconnects. For information about how to integrate a supported Cisco UCS rack-mount server with Cisco UCS Manager, see the hardware installation guide for that server.

Server Autoconfiguration Policy

Cisco UCS Manager uses this policy to determine how to configure a new server. If you create a server autoconfiguration policy, the following occurs when a new server starts:

- 1 The qualification in the server autoconfiguration policy is executed against the server.
- 2 If the server meets the required qualifications, the server is associated with a service profile created from the service profile template configured in the server autoconfiguration policy. The name of that service profile is based on the name given to the server by Cisco UCS Manager.
- 3 The service profile is assigned to the organization configured in the server autoconfiguration policy.

Server Discovery Policy

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

If you create a server discovery policy, the following occurs when a new server starts:

- 1 The qualification in the server discovery policy is executed against the server.
- 2 If the server meets the required qualifications, Cisco UCS Manager applies the following to the server:
 - Depending upon the option selected for the action, either discovers the new server immediately or waits for a user to acknowledge the new server
 - Applies the scrub policy to the server

Server Inheritance Policy

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

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

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

Server Pool Policy

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

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

Server Pool Policy Qualifications

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

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

You can use the server pool policy qualifications to qualify servers according to the following criteria:

- Adapter type
- Chassis location
- Memory type and configuration
- Power group
- CPU cores, type, and configuration
- Storage configuration and capacity
- Server model

Depending upon the implementation, you may configure several policies with server pool policy qualifications including the following:

- Autoconfiguration policy
- Chassis discovery policy
- Server discovery policy

- Server inheritance policy
- Server pool policy

vHBA Template

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

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

VM Lifecycle Policy

The VM lifecycle policy determines how long Cisco UCS Manager retains offline VMs and offline dynamic vNICs in its database. If a VM or dynamic vNIC remains offline after that period, Cisco UCS Manager deletes the object from its database.

All virtual machines (VMs) on Cisco UCS servers are managed by vCenter. Cisco UCS Manager cannot determine whether an inactive VM is temporarily shutdown, has been deleted, or is in some other state that renders it inaccessible. Therefore, Cisco UCS Manager considers all inactive VMs to be in an offline state.

Cisco UCS Manager considers a dynamic vNIC to be offline when the associated VM is shutdown, or the link between the fabric interconnect and the I/O module fails. On rare occasions, an internal error can also cause Cisco UCS Manager to consider a dynamic vNIC to be offline.

The default VM and dynamic vNIC retention period is 15 minutes. You can set that for any period of time between 1 minute and 7200 minutes (or 5 days).



Note

The VMs that Cisco UCS Manager displays are for information and monitoring only. You cannot manage VMs through Cisco UCS Manager. Therefore, when you delete a VM from the Cisco UCS Manager database, you do not delete the VM from the server or from vCenter.

vNIC Template

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

Beginning in Cisco UCS, Release 2.0(2), Cisco UCS Manager does not automatically create a VM-FEX port profile with the correct settings when you create a vNIC template. If you want to create a VM-FEX port profile, you must configure the target of the vNIC template as a VM.

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



Note

If your server has two Emulex or QLogic NICs (Cisco UCS CNA M71KR-E or Cisco UCS CNA M71KR-Q), you must configure vNIC policies for both adapters in your service profile to get a user-defined MAC address for both NICs. If you do not configure policies for both NICs, Windows still detects both of them in the PCI bus. Then because the second eth is not part of your service profile, Windows assigns it a hardware MAC address. If you then move the service profile to a different server, Windows sees additional NICs because one NIC did not have a user-defined MAC address.

vNIC/vHBA Placement Policies

vNIC/vHBA placement policies are used to determine what types of vNICs or vHBAs can be assigned to the physical adapters on a server. Each vNIC/vHBA placement policy contains four virtual network interface connections (vCons) that are virtual representations of the physical adapters. When a vNIC/vHBA placement policy is assigned to a service profile, and the service profile is associated with a server, the vCons in the vNIC/vHBA placement policy are assigned to the physical adapters.

If you do not include a vNIC/vHBA placement policy in the service profile or you use the default configuration for a server with two adapters, Cisco UCS Manager defaults to the **All** configuration and equally distributes the vNICs and vHBAs between the adapters.

You can use this policy to assign vNICs or vHBAs to either of the two vCons. Cisco UCS Manager uses the vCon assignment to determine how to assign the vNICs and vHBAs to the physical adapter during service profile association.

- **All**—All configured vNICs and vHBAs can be assigned to the vCon, whether they are explicitly assigned to it, unassigned, or dynamic.
- **Assigned Only**—vNICs and vHBAs must be explicitly assigned to the vCon. You can assign them explicitly through the service profile or the properties of the vNIC or vHBA.
- **Exclude Dynamic**—Dynamic vNICs and vHBAs cannot be assigned to the vCon. The vCon can be used for all static vNICs and vHBAs, whether they are unassigned or explicitly assigned to it.
- **Exclude Unassigned**—Unassigned vNICs and vHBAs cannot be assigned to the vCon. The vCon can be used for dynamic vNICs and vHBAs and for static vNICs and vHBAs that are explicitly assigned to it.

Operational Policies

Fault Collection Policy

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

A fault in Cisco UCS has the following lifecycle:

- 1 A condition occurs in the system and Cisco UCS Manager raises a fault. This is the active state.
- 2 When the fault is alleviated, it enters a flapping or soaking interval that is designed to prevent flapping. Flapping occurs when a fault is raised and cleared several times in rapid succession. During the flapping interval, the fault retains its severity for the length of time specified in the fault collection policy.
- 3 If the condition reoccurs during the flapping interval, the fault returns to the active state. If the condition does not reoccur during the flapping interval, the fault is cleared.
- 4 The cleared fault enters the retention interval. This interval ensures that the fault reaches the attention of an administrator even if the condition that caused the fault has been alleviated and the fault has not been deleted prematurely. The retention interval retains the cleared fault for the length of time specified in the fault collection policy.
- 5 If the condition reoccurs during the retention interval, the fault returns to the active state. If the condition does not reoccur, the fault is deleted.

Flow Control Policy

Flow control policies determine whether the uplink Ethernet ports in a Cisco UCS domain send and receive IEEE 802.3x pause frames when the receive buffer for a port fills. These pause frames request that the transmitting port stop sending data for a few milliseconds until the buffer clears.

For flow control to work between a LAN port and an uplink Ethernet port, you must enable the corresponding receive and send flow control parameters for both ports. For Cisco UCS, the flow control policies configure these parameters.

When you enable the send function, the uplink Ethernet port sends a pause request to the network port if the incoming packet rate becomes too high. The pause remains in effect for a few milliseconds before traffic is reset to normal levels. If you enable the receive function, the uplink Ethernet port honors all pause requests from the network port. All traffic is halted on that uplink port until the network port cancels the pause request.

Because you assign the flow control policy to the port, changes to the policy have an immediate effect on how the port reacts to a pause frame or a full receive buffer.

Maintenance Policy

A maintenance policy determines how Cisco UCS Manager reacts when a change that requires a server reboot is made to a service profile associated with a server or to an updating service profile bound to one or more service profiles.

The maintenance policy specifies how Cisco UCS Manager deploys the service profile changes. The deployment can occur in one of the following ways:

- Immediately
- When acknowledged by a user with admin privileges
- Automatically at the time specified in the schedule

If the maintenance policy is configured to deploy the change during a scheduled maintenance window, the policy must include a valid schedule. The schedule deploys the changes in the first available maintenance window.

Scrub Policy

This policy determines what happens to local data and to the BIOS settings on a server during the discovery process and when the server is disassociated from a service profile. Depending upon how you configure a scrub policy, the following can occur at those times:

Disk Scrub

One of the following occurs to the data on any local drives on disassociation:

- If enabled, destroys all data on any local drives
- If disabled, preserves all data on any local drives, including local storage configuration

BIOS Settings Scrub

One of the following occurs to the BIOS settings when a service profile containing the scrub policy is disassociated from a server:

- If enabled, erases all BIOS settings for the server and resets them to the BIOS defaults for that server type and vendor
- If disabled, preserves the existing BIOS settings on the server

Serial over LAN Policy

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

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

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

Statistics Collection Policy

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

For NIC statistics, Cisco UCS Manager displays the average, minimum, and maximum of the change since the last collection of statistics. If the values are 0, there has been no change since the last collection.

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

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



Note

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

Statistics Threshold Policy

A statistics threshold policy monitors statistics about certain aspects of the system and generates an event if the threshold is crossed. You can set both minimum and maximum thresholds. For example, you can configure

the policy to raise an alarm if the CPU temperature exceeds a certain value, or if a server is overutilized or underutilized.

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

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

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


Note

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

Pools

Pools are collections of identities, or physical or logical resources, that are available in the system. All pools increase the flexibility of service profiles and allow you to centrally manage your system resources.

You can use pools to segment unconfigured servers or available ranges of server identity information into groupings that make sense for the data center. For example, if you create a pool of unconfigured servers with similar characteristics and include that pool in a service profile, you can use a policy to associate that service profile with an available, unconfigured server.

If you pool identifying information, such as MAC addresses, you can pre-assign ranges for servers that will host specific applications. For example, all database servers could be configured within the same range of MAC addresses, UUIDs, and WWNs.

Server Pools

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

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

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

MAC Pools

A MAC pool is a collection of network identities, or MAC addresses, that are unique in their layer 2 environment and are available to be assigned to vNICs on a server. If you use MAC pools in service profiles,

you do not have to manually configure the MAC addresses to be used by the server associated with the service profile.

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

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

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

UUID Suffix Pools

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

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

WWN Pools

A 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 vHBA



Important

A WWN pool can include only WWNNs or WWPNS in the ranges from 20:00:00:00:00:00:00:00 to 20:FF:FF:FF:FF:FF:FF:FF or from 50:00:00:00:00:00:00:00 to 5F:FF:FF:FF:FF:FF:FF:FF. All other WWN ranges are reserved. To ensure the uniqueness of the Cisco UCS WWNNs and WWPNS in the SAN fabric, we recommend that you use the following WWN prefix for all blocks in a pool:
20:00:00:25:B5:XX:XX:XX

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

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

WWNN Pools

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

WWPN Pools

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

Management IP Pool

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

You can configure service profiles and service profile templates to use IP addresses from the management IP pool. You cannot configure servers to use the management IP pool.

All IP addresses in the management IP pool must be in the same subnet as the IP address of the fabric interconnect.

**Note**

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

Traffic Management

Oversubscription

Oversubscription occurs when multiple network devices are connected to the same fabric interconnect port. This practice optimizes fabric interconnect use, since ports rarely run at maximum speed for any length of time. As a result, when configured correctly, oversubscription allows you to take advantage of unused bandwidth. However, incorrectly configured oversubscription can result in contention for bandwidth and a lower quality of service to all services that use the oversubscribed port.

For example, oversubscription can occur if four servers share a single uplink port, and all four servers attempt to send data at a cumulative rate higher than available bandwidth of uplink port.

Oversubscription Considerations

The following elements can impact how you configure oversubscription in a Cisco UCS domain:

Ratio of Server-Facing Ports to Uplink Ports

You need to know what how many server-facing ports and uplink ports are in the system, because that ratio can impact performance. For example, if your system has twenty ports that can communicate down to the servers and only two ports that can communicate up to the network, your uplink ports will be oversubscribed. In this situation, the amount of traffic created by the servers can also affect performance.

Number of Uplink Ports from Fabric Interconnect to Network

You can choose to add more uplink ports between the Cisco UCS fabric interconnect and the upper layers of the LAN to increase bandwidth. In Cisco UCS, you must have at least one uplink port per fabric interconnect to ensure that all servers and NICs to have access to the LAN. The number of LAN uplinks should be determined by the aggregate bandwidth needed by all Cisco UCS servers.

For the 6100 series fabric interconnects, Fibre Channel uplink ports are available on the expansion slots only. You must add more expansion slots to increase number of available Fibre Channel uplinks. Ethernet uplink ports can exist on the fixed slot and on expansion slots.

For the 6200 series fabric interconnects running Cisco UCS Manager, version 2.0 and higher, Ethernet uplink ports and Fibre Channel uplink ports are both configurable on the base module, as well as on the expansion module.

For example, if you have two Cisco UCS 5100 series chassis that are fully populated with half width Cisco UCS B200-M1 servers, you have 16 servers. In a cluster configuration, with one LAN uplink per fabric interconnect, these 16 servers share 20GbE of LAN bandwidth. If more capacity is needed, more uplinks from the fabric interconnect should be added. We recommend that you have symmetric configuration of the uplink in cluster configurations. In the same example, if 4 uplinks are used in each fabric interconnect, the 16 servers are sharing 80 GB of bandwidth, so each has approximately 5 GB of capacity. When multiple uplinks are used on a Cisco UCS fabric interconnect the network design team should consider using a port channel to make best use of the capacity.

Number of Uplink Ports from I/O Module to Fabric Interconnect

You can choose to add more bandwidth between I/O module and fabric interconnect by using more uplink ports and increasing the number of cables. In Cisco UCS, you can have one, two, or four cables connecting a I/O module to a Cisco UCS 6100 series fabric interconnect. You can have up to eight cables if you're connecting a 2208 I/O module and a 6248 fabric interconnect. The number of cables determines the number of active uplink ports and the oversubscription ratio.

Number of Active Links from Server to Fabric Interconnect

The amount of non-oversubscribed bandwidth available to each server depends on the number of I/O modules used and the number of cables used to connect those I/O modules to the fabric interconnects. Having a second I/O module in place provides additional bandwidth and redundancy to the servers. This level of flexibility in design ensures that you can provide anywhere from 80 Gbps (two I/O modules with four links each) to 10 Gbps (one I/O module with one link) to the chassis.

With 80 Gbps to the chassis, each half-width server in the Cisco UCS domain can get up to 10 Gbps in a non-oversubscribed configuration, with an ability to use up to 20 Gbps with 2:1 oversubscription.

Guidelines for Estimating Oversubscription

When you estimate the optimal oversubscription ratio for a fabric interconnect port, consider the following guidelines:

Cost/Performance Slider

The prioritization of cost and performance is different for each data center and has a direct impact on the configuration of oversubscription. When you plan hardware usage for oversubscription, you need to know where the data center is located on this slider. For example, oversubscription can be minimized if the data center is more concerned with performance than cost. However, cost is a significant factor in most data centers, and oversubscription requires careful planning.

Bandwidth Usage

The estimated bandwidth that you expect each server to actually use is important when you determine the assignment of each server to a fabric interconnect port and, as a result, the oversubscription ratio of the ports. For oversubscription, you must consider how many GBs of traffic the server will consume on average, the ratio of configured bandwidth to used bandwidth, and the times when high bandwidth use will occur.

Network Type

The network type is only relevant to traffic on uplink ports, because FCoE does not exist outside Cisco UCS. The rest of the data center network only differentiates between LAN and SAN traffic. Therefore, you do not need to take the network type into consideration when you estimate oversubscription of a fabric interconnect port.

Pinning

Pinning in Cisco UCS is only relevant to uplink ports. You can pin Ethernet or FCoE traffic from a given server to a specific uplink Ethernet port or uplink FC port.

When you pin the NIC and HBA of both physical and virtual servers to uplink ports, you give the fabric interconnect greater control over the unified fabric. This control ensures more optimal utilization of uplink port bandwidth.

Cisco UCS uses pin groups to manage which NICs, vNICs, HBAs, and vHBAs are pinned to an uplink port. To configure pinning for a server, you can either assign a pin group directly, or include a pin group in a vNIC policy, and then add that vNIC policy to the service profile assigned to that server. All traffic from the vNIC or vHBA on the server travels through the I/O module to the same uplink port.

Pinning Server Traffic to Server Ports

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

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



Note

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

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

Chassis with One I/O Module (Not Configured for Fabric Port Channels)



Note

If the adapter in a server supports and is configured for adapter port channels, those port channels are pinned to the same link as described in the following table. If the I/O module in the chassis supports and is configured for fabric port channels, the server slots are pinned to a fabric port channel rather than to an individual link.

Links on Chassis	Link 1 / Fabric Port Channel	Link 2	Link 3	Link 4	Link 5	Link 6	Link 7	Link 8
1 link	All server slots	None	None	None	None	None	None	None

Links on Chassis	Link 1 / Fabric Port Channel	Link 2	Link 3	Link 4	Link 5	Link 6	Link 7	Link 8
2 links	Server slots 1, 3, 5, and 7	Server slots 2, 4, 6, and 8	None	None	None	None	None	None
4 links	Server slots 1 and 5	Server slots 2 and 6	Server slots 3 and 7	Server slots 4 and 8	None	None	None	None
8 links	Server slot 1	Server slot 2	Server slot 3	Server slot 4	Server slot 5	Server slot 6	Server slot 7	Server slot 8
Fabric Port Channel	All server slots	N/A	N/A	N/A	N/A	N/A	N/A	N/A

Chassis with Two I/O Modules

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

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

Guidelines for Pinning

When you determine the optimal configuration for pin groups and pinning for an uplink port, consider the estimated bandwidth usage for the servers. If you know that some servers in the system will use a lot of bandwidth, ensure that you pin these servers to different uplink ports.

Quality of Service

Cisco UCS provides the following methods to implement quality of service:

- System classes that specify the global configuration for certain types of traffic across the entire system
- QoS policies that assign system classes for individual vNICs
- Flow control policies that determine how uplink Ethernet ports handle pause frames

System Classes

Cisco UCS uses Data Center Ethernet (DCE) to handle all traffic inside a Cisco UCS 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 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. This provides a level of traffic management, even in an oversubscribed system. For example, you can configure the Fibre Channel Priority system class to determine the percentage of DCE bandwidth allocated to FCoE traffic.

The following table describes the system classes that you can configure:

Table 5: System Classes

System Class	Description
Platinum Gold Silver Bronze	A configurable set of system classes that you can include in the QoS policy for a service profile. Each system class manages one lane of traffic. All properties of these system classes are available for you to assign custom settings and policies.
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.

Quality of Service Policy

A quality of service (QoS) policy assigns a system class to the outgoing traffic for a vNIC or vHBA. This system class determines the quality of service for that traffic. For certain adapters you can also specify additional controls on the outgoing traffic, such as burst and rate.

You must include a QoS policy in a vNIC policy or vHBA policy and then include that policy in a service profile to configure the vNIC or vHBA.

Flow Control Policy

Flow control policies determine whether the uplink Ethernet ports in a Cisco UCS domain send and receive IEEE 802.3x pause frames when the receive buffer for a port fills. These pause frames request that the transmitting port stop sending data for a few milliseconds until the buffer clears.

For flow control to work between a LAN port and an uplink Ethernet port, you must enable the corresponding receive and send flow control parameters for both ports. For Cisco UCS, the flow control policies configure these parameters.

When you enable the send function, the uplink Ethernet port sends a pause request to the network port if the incoming packet rate becomes too high. The pause remains in effect for a few milliseconds before traffic is reset to normal levels. If you enable the receive function, the uplink Ethernet port honors all pause requests from the network port. All traffic is halted on that uplink port until the network port cancels the pause request.

Because you assign the flow control policy to the port, changes to the policy have an immediate effect on how the port reacts to a pause frame or a full receive buffer.

Opt-In Features

Each Cisco UCS domain is licensed for all functionality. Depending upon how the system is configured, you can decide to opt in to some features or opt out of them for easier integration into existing environment. If a process change happens, you can change your system configuration and include one or both of the opt-in features.

The opt-in features are as follows:

- Stateless computing, which takes advantage of mobile service profiles with pools and policies where each component, such as a server or an adapter, is stateless.
- Multi-tenancy, which uses organizations and role-based access control to divide the system into smaller logical segments.

Stateless Computing

Stateless computing allows you to use a service profile to apply the personality of one server to a different server in the same Cisco UCS domain. The personality of the server includes the elements that identify that server and make it unique in the Cisco UCS domain. If you change any of these elements, the server could lose its ability to access, use, or even achieve booted status.

The elements that make up a server's personality include the following:

- Firmware versions
- UUID (used for server identification)

- MAC address (used for LAN connectivity)
- World Wide Names (used for SAN connectivity)
- Boot settings

Stateless computing creates a dynamic server environment with highly flexible servers. Every physical server in a Cisco UCS domain remains anonymous until you associate a service profile with it, then the server gets the identity configured in the service profile. If you no longer need a business service on that server, you can shut it down, disassociate the service profile, and then associate another service profile to create a different identity for the same physical server. The "new" server can then host another business service.

To take full advantage of the flexibility of statelessness, the optional local disks on the servers should only be used for swap or temp space and not to store operating system or application data.

You can choose to fully implement stateless computing for all physical servers in a Cisco UCS domain, to not have any stateless servers, or to have a mix of the two types.

If You Opt In to Stateless Computing

Each physical server in the Cisco UCS domain is defined through a service profile. Any server can be used to host one set of applications, then reassigned to another set of applications or business services, if required by the needs of the data center.

You create service profiles that point to policies and pools of resources that are defined in the Cisco UCS domain. The server pools, WWN pools, and MAC pools ensure that all unassigned resources are available on an as-needed basis. For example, if a physical server fails, you can immediately assign the service profile to another server. Because the service profile provides the new server with the same identity as the original server, including WWN and MAC address, the rest of the data center infrastructure sees it as the same server and you do not need to make any configuration changes in the LAN or SAN.

If You Opt Out of Stateless Computing

Each server in the Cisco UCS domain is treated as a traditional rack mount server.

You create service profiles that inherit the identify information burned into the hardware and use these profiles to configure LAN or SAN connectivity for the server. However, if the server hardware fails, you cannot reassign the service profile to a new server.

Multi-Tenancy

Multi-tenancy allows you to divide up the large physical infrastructure of an Cisco UCS domain into logical entities known as organizations. As a result, you can achieve a logical isolation between organizations without providing a dedicated physical infrastructure for each organization.

You can assign unique resources to each tenant through the related organization, in the multi-tenant environment. These resources can include different policies, pools, and quality of service definitions. You can also implement locales to assign or restrict user privileges and roles by organization, if you do not want all users to have access to all organizations.

If you set up a multi-tenant environment, all organizations are hierarchical. The top-level organization is always root. The policies and pools that you create in root are system-wide and are available to all organizations in the system. However, any policies and pools created in other organizations are only available to organizations that are above it in the same hierarchy. For example, if a system has organizations named Finance and HR that are not in the same hierarchy, Finance cannot use any policies in the HR organization, and HR cannot

access any policies in the Finance organization. However, both Finance and HR can use policies and pools in the root organization.

If you create organizations in a multi-tenant environment, you can also set up one or more of the following for each organization or for a sub-organization in the same hierarchy:

- Resource pools
- Policies
- Service profiles
- Service profile templates

If You Opt In to Multi-Tenancy

Each Cisco UCS domain is divided into several distinct organizations. The types of organizations you create in a multi-tenancy implementation depends upon the business needs of the company. Examples include organizations that represent the following:

- Enterprise groups or divisions within a company, such as marketing, finance, engineering, or human resources
- Different customers or name service domains, for service providers

You can create locales to ensure that users have access only to those organizations that they are authorized to administer.

If You Opt Out of Multi-Tenancy

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

Virtualization in Cisco UCS

Overview of Virtualization

Virtualization allows the creation of multiple virtual machines (VMs) to run in isolation, side by side on the same physical machine.

Each virtual machine has its own set of virtual hardware (RAM, CPU, NIC) upon which an operating system and fully configured applications are loaded. The operating system sees a consistent, normalized set of hardware regardless of the actual physical hardware components.

In a virtual machine, both hardware and software are encapsulated in a single file for rapid copying, provisioning, and moving between physical servers. You can move a virtual machine, within seconds, from one physical server to another for zero-downtime maintenance and continuous workload consolidation.

The virtual hardware makes it possible for many servers, each running in an independent virtual machine, to run on a single physical server. The advantages of virtualization include better use of computing resources, greater server density, and seamless server migration.

Overview of Cisco Virtual Machine Fabric Extender

A virtualized server implementation consists of one or more VMs running as 'guests' on a single physical server. The guest VMs are hosted and managed by a software layer called the hypervisor or virtual machine manager (VMM). The hypervisor typically presents a virtual network interface to each VM and performs Layer 2 switching of traffic from a VM to other local VMs or to a physical interface to the external network.

Working with a Cisco virtual interface card (VIC) adapter, Cisco Virtual Machine Fabric Extender (VM-FEX) bypasses software-based switching of VM traffic by the hypervisor in favor of external hardware-based switching in the fabric interconnect. This method results in a reduced load on the server CPU, faster switching, and the ability to apply a rich set of network management features to local and remote traffic.

VM-FEX extends the (prestandard) IEEE 802.1Qbh port extender architecture to the VMs, providing each VM interface with a virtual Peripheral Component Interconnect Express (PCIe) device and a virtual port on a switch. This solution allows precise rate limiting and quality of service (QoS) guarantees on the VM interface.

Virtualization with Network Interface Cards and Converged Network Adapters

Network interface card (NIC) and converged network adapters support virtualized environments with the standard VMware integration with ESX installed on the server and all virtual machine management performed through the VC.

Portability of Virtual Machines

If you implement service profiles you retain the ability to easily move a server identity from one server to another. After you image the new server, the ESX treats that server as if it were the original.

Communication between Virtual Machines on the Same Server

These adapters implement the standard communications between virtual machines on the same server. If an ESX host includes multiple virtual machines, all communications must go through the virtual switch on the server.

If the system uses the native VMware drivers, the virtual switch is out of the network administrator's domain and is not subject to any network policies. As a result, for example, QoS policies on the network are not applied to any data packets traveling from VM1 to VM2 through the virtual switch.

If the system includes another virtual switch, such as the Nexus 1000, that virtual switch is subject to the network policies configured on that switch by the network administrator.

Virtualization with a Virtual Interface Card Adapter

A Cisco VIC adapter, such as the Cisco UCS M81KR Virtual Interface Card, is a converged network adapter (CNA) designed for both single-OS and VM-based deployments. The VIC adapter supports static or dynamic virtualized interfaces, including up to 128 virtual network interface cards (vNICs).

VIC adapters support VM-FEX to provide hardware-based switching of traffic to and from virtual machine interfaces.



CHAPTER 3

Overview of Cisco UCS Manager

This chapter includes the following sections:

- [About Cisco UCS Manager](#) , page 43
- [Tasks You Can Perform in Cisco UCS Manager](#) , page 44
- [Tasks You Cannot Perform in Cisco UCS Manager](#) , page 46
- [Cisco UCS Manager in a High Availability Environment](#), page 46

About Cisco UCS Manager

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

Multiple Management Interfaces

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

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

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

However, you cannot do the following:

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

Centralized Management

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

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

Support for Virtual and Physical Servers

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

Role-Based Administration and Multi-Tenancy Support

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

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

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

Tasks You Can Perform in Cisco UCS Manager

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

Cisco UCS Hardware Management

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

- Chassis
- Servers
- Fabric interconnects
- Fans

- Ports
- Interface cards
- I/O modules

Cisco UCS Resource Management

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

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

Server Administration

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

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

Network Administration

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

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

Storage Administration

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

- Configure ports, port channels, and SAN PIN groups
- Create VSANs
- Configure the quality of service classes and definitions

- Create the pools and policies related to the network configuration, such as WWN pools and Fibre Channel adapter profiles

Tasks You Cannot Perform in Cisco UCS Manager

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

No Cross-System Management

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

No Operating System or Application Provisioning or Management

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

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

Cisco UCS Manager in a High Availability Environment

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

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



CHAPTER 4

Overview of Cisco UCS Manager CLI

This chapter includes the following sections:

- [Managed Objects, page 47](#)
- [Command Modes, page 47](#)
- [Object Commands, page 49](#)
- [Complete a Command, page 50](#)
- [Command History, page 50](#)
- [Committing, Discarding, and Viewing Pending Commands, page 50](#)
- [Online Help for the CLI, page 51](#)
- [CLI Session Limits, page 51](#)
- [Web Session Limits, page 51](#)
- [Pre-Login Banner, page 52](#)

Managed Objects

Cisco UCS uses a managed object model, where managed objects are abstract representations of physical or logical entities that can be managed. For example, servers, chassis, I/O cards, and processors are physical entities represented as managed objects, and resource pools, user roles, service profiles, and policies are logical entities represented as managed objects.

Managed objects may have one or more associated properties that can be configured.

Command Modes

The CLI is organized into a hierarchy of command modes, with the EXEC mode being the highest-level mode of the hierarchy. Higher-level modes branch into lower-level modes. You use **create**, **enter**, and **scope** commands to move from higher-level modes to modes in the next lower level, and you use the **exit** command to move up one level in the mode hierarchy.

**Note**

Most command modes are associated with managed objects, so you must create an object before you can access the mode associated with that object. You use **create** and **enter** commands to create managed objects for the modes being accessed. The **scope** commands do not create managed objects and can only access modes for which managed objects already exist.

Each mode contains a set of commands that can be entered in that mode. Most of the commands available in each mode pertain to the associated managed object. Depending on your assigned role and locale, you may have access to only a subset of the commands available in a mode; commands to which you do not have access are hidden.

The CLI prompt for each mode shows the full path down the mode hierarchy to the current mode. This helps you to determine where you are in the command mode hierarchy, and it can be an invaluable tool when you need to navigate through the hierarchy.

The following table lists the main command modes, the commands used to access each mode, and the CLI prompt associated with each mode.

Table 6: Main Command Modes and Prompts

Mode Name	Commands Used to Access	Mode Prompt
EXEC	top command from any mode	#
adapter	scope adapter command from EXEC mode	/adapter #
chassis	scope chassis command from EXEC mode	/chassis #
Ethernet server	scope eth-server command from EXEC mode	/eth-server #
Ethernet uplink	scope eth-uplink command from EXEC mode	/eth-uplink #
fabric-interconnect	scope fabric-interconnect command from EXEC mode	/fabric-interconnect #
Fibre Channel uplink	scope fc-uplink command from EXEC mode	/fc-uplink #
firmware	scope firmware command from EXEC mode	/firmware #
Host Ethernet interface	scope host-eth-if command from EXEC mode	/host-eth-if #
Host Fibre Channel interface	scope host-fc-if command from EXEC mode	/host-fc-if #

Mode Name	Commands Used to Access	Mode Prompt
monitoring	scope monitoring command from EXEC mode	/monitoring #
organization	scope org command from EXEC mode	/org #
security	scope security command from EXEC mode	/security #
server	scope server command from EXEC mode	/server #
service-profile	scope service-profile command from EXEC mode	/service-profile #
system	scope system command from EXEC mode	/system #
virtual HBA	scope vhba command from EXEC mode	/vhba #
virtual NIC	scope vnic command from EXEC mode	/vnic #

Object Commands

Four general commands are available for object management:

- **create** *object*
- **delete** *object*
- **enter** *object*
- **scope** *object*

You can use the **scope** command with any managed object, whether a permanent object or a user-instantiated object. The other commands allow you to create and manage user-instantiated objects. For every **create object** command, a corresponding **delete object** and **enter object** command exists.

In the management of user-instantiated objects, the behavior of these commands depends on whether the object exists, as described in the following tables:

Table 7: Command behavior if the object does not exist

Command	Behavior
create object	The object is created and its configuration mode, if applicable, is entered.

Command	Behavior
<code>delete object</code>	An error message is generated.
<code>enter object</code>	The object is created and its configuration mode, if applicable, is entered.
<code>scope object</code>	An error message is generated.

Table 8: Command behavior if the object exists

Command	Behavior
<code>create object</code>	An error message is generated.
<code>delete object</code>	The object is deleted.
<code>enter object</code>	The configuration mode, if applicable, of the object is entered.
<code>scope object</code>	The configuration mode of the object is entered.

Complete a Command

You can use the Tab key in any mode to complete a command. Partially typing a command name and pressing Tab causes the command to be displayed in full or to the point where another keyword must be chosen or an argument value must be entered.

Command History

The CLI stores all commands used in the current session. You can step through the previously used commands by using the Up Arrow or Down Arrow keys. The Up Arrow key steps to the previous command in the history, and the Down Arrow key steps to the next command in the history. If you get to the end of the history, pressing the Down Arrow key does nothing.

All commands in the history can be entered again by simply stepping through the history to recall the desired command and pressing Enter. The command is entered as if you had manually typed it. You can also recall a command and change it before you enter it.

Committing, Discarding, and Viewing Pending Commands

When you enter a configuration command in the CLI, the command is not applied until you enter the **commit-buffer** command. Until committed, a configuration command is pending and can be discarded by entering a **discard-buffer** command.

You can accumulate pending changes in multiple command modes and apply them together with a single **commit-buffer** command. You can view the pending commands by entering the **show configuration pending** command in any command mode.



Note Committing multiple commands together is not an atomic operation. If any command fails, the successful commands are applied despite the failure. Failed commands are reported in an error message.

While any commands are pending, an asterisk (*) appears before the command prompt. The asterisk disappears when you enter the **commit-buffer** command, as shown in this example:

```
switch-1# scope chassis 1
switch-1 /chassis # enable locator-led
switch-1 /chassis* # show configuration pending
  scope chassis 1
+   enable locator-led
  exit
switch-1 /chassis* # commit-buffer
switch-1 /chassis #
```

Online Help for the CLI

At any time, you can type the ? character to display the options available at the current state of the command syntax. If you have not typed anything at the prompt, typing ? lists all available commands for the mode you are in. If you have partially typed a command, typing ? lists all available keywords and arguments available at your current position in the command syntax.

CLI Session Limits

Cisco UCS Manager limits the number of CLI sessions that can be active at one time to 32 total sessions. This value is not configurable.

Web Session Limits

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

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

Setting the Web Session Limit for Cisco UCS Manager from the CLI

Procedure

	Command or Action	Purpose
Step 1	UCS-A# scope system	Enters system mode.
Step 2	UCS-A /system # scope services	Enters system services mode.

	Command or Action	Purpose
Step 3	UCS-A /system/services # scope web-session-limits	Enters system services web session limits mode.
Step 4	UCS-A /system/services/web-session-limits # set total num-of-logins-total	The maximum number of concurrent HTTP and HTTPS sessions allowed for all users within the system. Enter an integer between 1 and 256.
Step 5	UCS-A /system/services/web-session-limits # commit-buffer	Commits the transaction to the system configuration.

The following example sets the maximum number of HTTP and HTTPS sessions allowed by the system to 200 and commits the transaction:

```
UCS-A# scope system
UCS-A /system # scope services
UCS-A /system/services # scope web-session-limits
UCS-A /system/services/web-session-limits* # set total 200
UCS-A /system/services/web-session-limits* # commit-buffer
UCS-A /system/services/web-session-limits #
```

Pre-Login Banner

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

Creating the Pre-Login Banner

Procedure

	Command or Action	Purpose
Step 1	UCS-A# scope security	Enters security mode.
Step 2	UCS-A /security # scope banner	Enters banner security mode.
Step 3	UCS-A /security/banner # create pre-login-banner	Creates a pre login banner.
Step 4	UCS-A /security/banner/pre-login-banner # set message	Specifies the message that Cisco UCS Manager displays to the user before it displays the login prompt for the Cisco UCS Manager GUI or CLI. You can enter any standard ASCII character in this field.

	Command or Action	Purpose
		Launches a dialog for entering the pre-login banner message text.
Step 5	At the prompt, type a pre-login banner message and press Enter .	On the line following your input, type ENDOFBUF to finish. Press Ctrl and C to cancel out of the set message dialog.
Step 6	UCS-A /security/banner/pre-login-banner # commit-buffer	Commits the transaction to the system configuration.

The following example creates the pre-login banner:

```
UCS-A# scope security
UCS-A /security # scope banner
UCS-A /security/banner # create pre-login-banner
UCS-A /security/banner/pre-login-banner* # set message
Enter lines one at a time. Enter ENDOFBUF to finish. Press ^C to abort.
Enter prelogin banner:
>Welcome to UCS System 1
>ENDOFBUF
UCS-A /security/banner/pre-login-banner* # commit-buffer
UCS-A /security/banner/pre-login-banner #
```

Modifying the Pre-Login Banner

Procedure

	Command or Action	Purpose
Step 1	UCS-A# scope security	Enters security mode.
Step 2	UCS-A /security # scope banner	Enters banner security mode.
Step 3	UCS-A /security/banner # scope pre-login-banner	Enters pre-login-banner banner security mode.
Step 4	UCS-A /security/banner/pre-login-banner # set message	Specifies the message that Cisco UCS Manager displays to the user before it displays the login prompt for the Cisco UCS Manager GUI or CLI. You can enter any standard ASCII character in this field. Launches a dialog for entering the pre-login banner message text.
Step 5	At the prompt, modify the pre-login banner message and press Enter .	On the line following your input, type ENDOFBUF to finish. Press Ctrl and C to cancel out of the set message dialog.

	Command or Action	Purpose
Step 6	UCS-A /security/banner/pre-login-banner # commit-buffer	Commits the transaction to the system configuration.

The following example modifies the pre-login banner:

```
UCS-A# scope security
UCS-A /security # scope banner
UCS-A /security/banner # create pre-login-banner
UCS-A /security/banner/pre-login-banner* # set message
Enter lines one at a time. Enter ENDOFBUF to finish. Press ^C to abort.
Enter prelogin banner:
Welcome to UCS System 1
ENDOFBUF
UCS-A /security/banner/pre-login-banner* # commit-buffer
UCS-A /security/banner/pre-login-banner #
```

Deleting the Pre-Login Banner

Procedure

	Command or Action	Purpose
Step 1	UCS-A# scope security	Enters security mode.
Step 2	UCS-A /security # scope banner	Enters banner security mode.
Step 3	UCS-A /security/banner # delete pre-login-banner	Deletes the pre-login banner from the system.
Step 4	UCS-A /security/banner # commit-buffer	Commits the transaction to the system configuration.

The following example deletes the pre-login banner:

```
UCS-A# scope security
UCS-A /security # scope banner
UCS-A /security/banner # delete pre-login-banner
UCS-A /security/banner* # commit-buffer
UCS-A /security/banner #
```




PART **II**

System Configuration

- [Configuring the Fabric Interconnects, page 57](#)
- [Configuring Ports and Port Channels, page 71](#)
- [Configuring Communication Services, page 103](#)
- [Configuring Authentication, page 125](#)
- [Configuring Organizations, page 153](#)
- [Configuring Role-Based Access Control, page 159](#)
- [Managing Firmware, page 185](#)
- [Configuring DNS Servers, page 241](#)
- [Configuring System-Related Policies, page 243](#)
- [Managing Licenses, page 251](#)
- [Managing Virtual Interfaces, page 259](#)



CHAPTER 5

Configuring the Fabric Interconnects

This chapter includes the following sections:

- [Initial System Setup, page 57](#)
- [Performing an Initial System Setup for a Standalone Configuration, page 59](#)
- [Initial System Setup for a Cluster Configuration, page 61](#)
- [Enabling a Standalone Fabric Interconnect for Cluster Configuration, page 64](#)
- [Changing the System Name, page 65](#)
- [Changing the Management Subnet of a Cluster, page 65](#)
- [Ethernet Switching Mode, page 66](#)
- [Configuring Ethernet Switching Mode, page 67](#)
- [Fibre Channel Switching Mode, page 68](#)
- [Configuring Fibre Channel Switching Mode, page 68](#)

Initial System Setup

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

- Installation method (GUI or CLI)
- Setup mode (restore from full system backup or initial setup)
- System configuration type (standalone or cluster configuration)
- System name
- Admin password
- Management port IP address and subnet mask
- Default gateway IP address
- DNS Server IP address

- Default domain name

Setup Mode

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

System Configuration Type

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

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

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



Note

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

To use the cluster configuration, the two fabric interconnects must be directly connected together using Ethernet cables between the L1 (L1-to-L1) and L2 (L2-to-L2) high availability ports, with no other fabric interconnects in between. This allows the two fabric interconnects to continuously monitor the status of each other and quickly know when one has failed.

Both fabric interconnects in a cluster configuration must go through the initial setup process. The first fabric interconnect to be set up must be enabled for a cluster configuration. Then, when the second fabric interconnect is set up, it detects the first fabric interconnect as a peer fabric interconnect in the cluster.

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

Management Port IP Address

In a standalone configuration, you must specify only one IP address and the subnet mask for the single management port on the fabric interconnect.

In a cluster configuration, you must specify the following three IP addresses in the same subnet:

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

Performing an Initial System Setup for a Standalone Configuration

Before You Begin

- 1 Verify the following physical connections on the fabric interconnect:

- The console port is physically connected to a computer terminal or console server
- The management Ethernet port (mgmt0) is connected to an external hub, switch, or router

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

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

- 9600 baud
- 8 data bits
- No parity
- 1 stop bit

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

- System name.
- Password for the admin account. Choose a strong password that meets the guidelines for Cisco UCS Manager passwords. This password cannot be blank.
- Management port IP address and subnet mask.
- Default gateway IP address.
- DNS server IP address (optional).
- Domain name for the system (optional).

Procedure

- 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 **setup** to continue as an initial system setup.
- Step 5** Enter **y** to confirm that you want to continue the initial setup.
- Step 6** Enter the password for the admin account.
- Step 7** To confirm, re-enter the password for the admin account.
- Step 8** Enter **no** to continue the initial setup for a standalone configuration.
- Step 9** Enter the system name.
- Step 10** Enter the IP address for the management port on the fabric interconnect.
- Step 11** Enter the subnet mask for the management port on the fabric interconnect.
- Step 12** Enter the IP address for the default gateway.
- Step 13** Enter **yes** if you want to specify the IP address for the DNS server, or **no** if you do not.
- Step 14** (Optional) Enter the IP address for the DNS server.
- Step 15** Enter **yes** if you want to specify the default domain name, or **no** if you do not.
- Step 16** (Optional) Enter the default domain name.
- Step 17** Review the setup summary and enter **yes** to save and apply the settings, or enter **no** to go through the Setup wizard again to change some of the settings.
If you choose to go through the Setup wizard again, it provides the values you previously entered, and the values appear in brackets. To accept previously entered values, press Enter.

The following example sets up a standalone configuration using the console:

```

Enter the installation method (console/gui)? console
Enter the setup mode (restore from backup or initial setup) [restore/setup]? setup
You have chosen to setup a new switch. Continue? (y/n): y
Enter the password for "admin": adminpassword%958
Confirm the password for "admin": adminpassword%958
Do you want to create a new cluster on this switch (select 'no' for standalone setup or if
you want this switch to be added to an existing cluster)? (yes/no) [n]: no
Enter the system name: foo
Mgmt0 IPv4 address: 192.168.10.10
Mgmt0 IPv4 netmask: 255.255.255.0
IPv4 address of the default gateway: 192.168.10.1
Configure the DNS Server IPv4 address? (yes/no) [n]: yes
DNS IPv4 address: 20.10.20.10
Configure the default domain name? (yes/no) [n]: yes
Default domain name: domainname.com
Following configurations will be applied:
Switch Fabric=A
System Name=foo
Management IP Address=192.168.10.10
Management IP Netmask=255.255.255.0
Default Gateway=192.168.10.1
DNS Server=20.10.20.10
Domain Name=domainname.com
Apply and save the configuration (select 'no' if you want to re-enter)? (yes/no): yes

```

Initial System Setup for a Cluster Configuration

Performing an Initial System Setup for the First Fabric Interconnect

Before You Begin

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

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

- 2 Verify that the console port parameters on the computer terminal (or console server) attached to the console port are as follows:
 - 9600 baud
 - 8 data bits
 - No parity
 - 1 stop bit
- 3 Collect the following information that you will need to supply during the initial setup:
 - System name.
 - Password for the admin account. Choose a strong password that meets the guidelines for Cisco UCS Manager passwords. This password cannot be blank.
 - Three static IP addresses: two for the management port on both fabric interconnects (one per fabric interconnect) and one for the cluster IP address used by Cisco UCS Manager.
 - Subnet mask for the three static IP addresses.
 - Default gateway IP address.
 - DNS server IP address (optional).
 - Domain name for the system (optional).

Procedure

- 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 **setup** to continue as an initial system setup.
- Step 5** Enter **y** to confirm that you want to continue the initial setup.
- Step 6** Enter the password for the admin account.
- Step 7** To confirm, re-enter the password for the admin account.
- Step 8** Enter **yes** to continue the initial setup for a cluster configuration.
- Step 9** Enter the fabric interconnect fabric (either **A** or **B**).
- Step 10** Enter the system name.
- Step 11** Enter the IP address for the management port on the fabric interconnect.
- Step 12** Enter the subnet mask for the management port on the fabric interconnect.
- Step 13** Enter the IP address for the default gateway.
- Step 14** Enter the virtual IP address.
- Step 15** Enter **yes** if you want to specify the IP address for the DNS server, or **no** if you do not.
- Step 16** (Optional) Enter the IP address for the DNS server.
- Step 17** Enter **yes** if you want to specify the default domain name, or **no** if you do not.
- Step 18** (Optional) Enter the default domain name.
- Step 19** Review the setup summary and enter **yes** to save and apply the settings, or enter **no** to go through the Setup wizard again to change some of the settings.

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

The following example sets up the first fabric interconnect for a cluster configuration using the console:

```

Enter the installation method (console/gui)? console
Enter the setup mode (restore from backup or initial setup) [restore/setup]? setup
You have chosen to setup a new switch. Continue? (y/n): y
Enter the password for "admin": adminpassword%958
Confirm the password for "admin": adminpassword%958
Do you want to create a new cluster on this switch (select 'no' for standalone setup or if
you want this switch to be added to an existing cluster)? (yes/no) [n]: yes
Enter the switch fabric (A/B): A
Enter the system name: foo
Mgmt0 IPv4 address: 192.168.10.10
Mgmt0 IPv4 netmask: 255.255.255.0
IPv4 address of the default gateway: 192.168.10.1
Virtual IPv4 address : 192.168.10.12
Configure the DNS Server IPv4 address? (yes/no) [n]: yes
DNS IPv4 address: 20.10.20.10
Configure the default domain name? (yes/no) [n]: yes
Default domain name: domainname.com
Following configurations will be applied:
Switch Fabric=A
System Name=foo
Management IP Address=192.168.10.10
Management IP Netmask=255.255.255.0
Default Gateway=192.168.10.1
Cluster Enabled=yes
Virtual Ip Address=192.168.10.12
DNS Server=20.10.20.10
Domain Name=domainname.com
Apply and save the configuration (select 'no' if you want to re-enter)? (yes/no): yes

```


Performing an Initial System Setup for the Second Fabric Interconnect

Before You Begin

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

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

- 2 Verify that the console port parameters on the computer terminal (or console server) attached to the console port are as follows:
 - 9600 baud
 - 8 data bits
 - No parity
 - 1 stop bit
- 3 Collect the following information that you will need to supply during the initial setup:
 - Password for the admin account of the peer fabric interconnect. Choose a strong password that meets the guidelines for Cisco UCS Manager passwords. This password cannot be blank.
 - Management port IP address in the same subnet as the peer fabric interconnect.

Procedure

-
- 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.

Note The fabric interconnect should detect the peer fabric interconnect in the cluster. If it does not, check the physical connections between the L1 and L2 ports, and verify that the peer fabric interconnect has been enabled for a cluster configuration.
 - Step 4** Enter **y** to add the subordinate fabric interconnect 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 the subordinate fabric interconnect.
 - Step 7** Review the setup summary and enter **yes** to save and apply the settings, or enter **no** to go through the Setup wizard again to change some of the settings.

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

The following example sets up the second fabric interconnect for a cluster configuration using the console:

```
Enter the installation method (console/gui)? console
Installer has detected the presence of a peer switch. This switch will be added to the
cluster. Continue?[y/n] y
Enter the admin password of the peer switch: adminpassword%958
Mgmt0 IPv4 address: 192.168.10.11
Management Ip Address=192.168.10.11
Apply and save the configuration (select 'no' if you want to re-enter)? (yes/no): yes
```

Enabling a Standalone Fabric Interconnect for Cluster Configuration

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

Procedure

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

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

```
UCS-A# connect local-mgmt
UCS-A(local-mgmt)# enable cluster 192.168.1.101
This command will enable cluster mode on this setup. You cannot change it
back to stand-alone. Are you sure you want to continue? (yes/no): yes
UCS-A(local-mgmt)#
```

What to Do Next

Add the second fabric interconnect to the cluster.

Changing the System Name

Procedure

	Command or Action	Purpose
Step 1	UCS-A # scope system	Enters system mode.
Step 2	UCS-A /system # set name name	Sets the system name.
Step 3	UCS-A /system # commit-buffer	Commits the transaction to the system configuration.

The name is updated on both fabric interconnects within about 30 seconds after the transaction is committed.

The following example changes the system name and commits the transaction:

```
UCS-A# scope system
UCS-A /system* # set name SanJose5
UCS-A /system* # commit-buffer
UCS-A /system #
```

Changing the Management Subnet of a Cluster

When changing the management subnet in a cluster configuration, you must change the following three IP addresses simultaneously and you must configure all three in the same subnet:

- Management port IP address for fabric interconnect A
- Management port IP address for fabric interconnect B
- Cluster IP (virtual IP) address

Procedure

	Command or Action	Purpose
Step 1	UCS-A# scope fabric-interconnect a	Enters fabric interconnect mode for fabric A.
Step 2	UCS-A /fabric-interconnect # set out-of-band ip ip-address netmask netmask gw gateway-ip-address	Sets the IP address, netmask, and gateway IP address of the fabric interconnect.
Step 3	UCS-A /fabric-interconnect # scope fabric-interconnect b	Enters fabric interconnect mode for fabric B.
Step 4	UCS-A /fabric-interconnect # set out-of-band ip ip-address netmask netmask gw gateway-ip-address	Sets the IP address, netmask, and gateway IP address of the fabric interconnect.

	Command or Action	Purpose
Step 5	UCS-A /fabric-interconnect # scope system	Enters system mode.
Step 6	UCS-A /system # set virtual-ip vip-address	Sets the virtual IP address for the cluster.
Step 7	UCS-A /system # commit-buffer	Commits the transaction to the system configuration.

When you commit the transaction, you are disconnected from the management session. Reconnect at the new management IP address.

This example changes both fabric-interconnect IP addresses, changes the virtual IP address, and commits the transaction, disconnecting the session:

```
UCS-A# scope fabric-interconnect a
UCS-A /fabric-interconnect # set out-of-band ip 192.0.2.111 netmask 255.255.255.0 gw 192.0.2.1
UCS-A /fabric-interconnect* # scope fabric-interconnect b
UCS-A /fabric-interconnect* # set out-of-band ip 192.0.2.112 netmask 255.255.255.0 gw
192.0.2.1
UCS-A /fabric-interconnect* # scope system
UCS-A /system* # set virtual-ip 192.0.2.113
UCS-A /system* # commit-buffer
```

Ethernet Switching Mode

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

End-Host Mode

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

- Layer 2 switching for L2 aggregation
- Virtual Switching System (VSS) aggregation layer



Note

When end-host mode is enabled, if a vNIC is hard pinned to an uplink port and this uplink port goes down, the system cannot re-pin the vNIC, and the vNIC remains down.

Switch Mode

Switch mode is the traditional Ethernet switching mode. The fabric interconnect runs STP to avoid loops, and broadcast and multicast packets are handled in the traditional way. Switch mode is not the default Ethernet

switching mode, and should be used only if the fabric interconnect is directly connected to a router, or if either of the following are used upstream:

- Layer 3 aggregation
- VLAN in a box



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

Configuring Ethernet Switching Mode



Important When you change the Ethernet switching mode, Cisco UCS Manager logs you out and restarts the fabric interconnect. For a cluster configuration, Cisco UCS Manager restarts both fabric interconnects sequentially. The second fabric interconnect can take several minutes to complete the change in Ethernet switching mode and become system ready. The configuration is retained.

Procedure

	Command or Action	Purpose
Step 1	UCS-A# scope eth-uplink	Enters Ethernet uplink mode.
Step 2	UCS-A /eth-uplink # set mode {end-host switch}	Sets the fabric interconnect to the specified switching mode.
Step 3	UCS-A /eth-uplink # commit-buffer	Commits the transaction to the system configuration. Cisco UCS Manager restarts the fabric interconnect, logs you out, and disconnects Cisco UCS Manager CLI.

The following example sets the fabric interconnect to end-host mode and commits the transaction:

```
UCS-A# scope eth-uplink
UCS-A /eth-uplink # set mode end-host
Warning: When committed, this change will cause the switch to reboot
UCS-A /eth-uplink* # commit-buffer
UCS-A /eth-uplink #
```

Fibre Channel Switching Mode

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

End-Host Mode

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

End-host mode is synonymous with NPV mode. This is the default Fibre Channel Switching mode.



Note

When end-host mode is enabled, if a vHBA is hard pinned to a uplink Fibre Channel port and this uplink port goes down, the system cannot re-pin the vHBA, and the vHBA remains down.

Switch Mode

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



Note

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

Switch mode is not the default Fibre Channel switching mode.

Configuring Fibre Channel Switching Mode



Important

When you change the Fibre Channel switching mode, Cisco UCS Manager logs you out and restarts the fabric interconnect. For a cluster configuration, Cisco UCS Manager restarts both fabric interconnects sequentially. The second fabric interconnect can take several minutes to complete the change in Fibre Channel switching mode and become system ready.

Procedure

	Command or Action	Purpose
Step 1	UCS-A# scope fc-uplink	Enters Fibre Channel uplink mode.

	Command or Action	Purpose
Step 2	UCS-A /fc-uplink # set mode { end-host switch }	Sets the fabric interconnect to the specified switching mode.
Step 3	UCS-A /fc-uplink # commit-buffer	Commits the transaction to the system configuration. Cisco UCS Manager restarts the fabric interconnect, logs you out, and disconnects Cisco UCS Manager CLI.

The following example sets the fabric interconnect to end-host mode and commits the transaction:

```
UCS-A # scope fc-uplink
UCS-A /fc-uplink # set mode end-host
UCS-A /fc-uplink* # commit-buffer
UCS-A /fc-uplink #
```




CHAPTER 6

Configuring Ports and Port Channels

This chapter includes the following sections:

- [Server and Uplink Ports on the 6100 Series Fabric Interconnect, page 71](#)
- [Unified Ports on the 6200 Series Fabric Interconnect, page 72](#)
- [Server Ports, page 79](#)
- [Uplink Ethernet Ports, page 80](#)
- [Appliance Ports, page 81](#)
- [FCoE and Fibre Channel Storage Ports, page 84](#)
- [Uplink Ethernet Port Channels, page 87](#)
- [Appliance Port Channels, page 90](#)
- [Fibre Channel Port Channels, page 94](#)
- [Adapter Port Channels, page 98](#)
- [Fabric Port Channels, page 99](#)

Server and Uplink Ports on the 6100 Series Fabric Interconnect

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

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



Note

Ports on the 6100 series fabric interconnect are not unified. For more information on Unified Ports, see [Unified Ports on the 6200 Series Fabric Interconnect](#).

Each fabric interconnect can include the following port types:

Server Ports

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

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

Uplink Ethernet Ports

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

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

- Uplink
- FCoE
- Appliance

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

Uplink Fibre Channel Ports

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

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

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

Unified Ports on the 6200 Series Fabric Interconnect

Unified ports are ports on the 6200 series fabric interconnect that can be configured to carry either Ethernet or Fibre Channel traffic. These ports are not reserved. They cannot be used by a Cisco UCS domain until you configure them.

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

Port Modes

The port mode determines whether a unified port on the fabric interconnect is configured to carry Ethernet or Fibre Channel traffic. The port mode is not automatically discovered by the fabric interconnect; it is configured in Cisco UCS Manager.

Changing the port mode results in the existing port configuration being deleted and replaced by a new logical port. Any objects associated with that port configuration, such as VLANs and VSANS, are removed. There is no restriction on the number of times the port mode can be changed for a unified port.

Port Types

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

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

By default, unified ports changed to Ethernet port mode are set to uplink Ethernet port type. unified ports changed to Fibre Channel port mode are set to the Fibre Channel uplink port type. Fibre Channel ports cannot be unconfigured.

Changing the port type does not require a reboot.

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

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



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

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

- Fibre Channel uplink ports
- Fibre Channel port channel members
- Fibre Channel storage ports
- SPAN destination ports
- SPAN source ports



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

Beacon LEDs for Unified Ports

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

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



Note

For unified ports on the expansion module, the Beacon LED property may be reset to the default value of Off during expansion module reboot.

Guidelines for Configuring Unified Ports

Consider the following guidelines and restrictions when configuring unified ports:

Hardware and Software Requirements

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

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

Port Mode Placement

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

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

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

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

**Note**

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

Special Considerations for UCS Manager CLI Users

Because the Cisco UCS Manager CLI does not validate port mode changes until you commit the buffer to the system configuration, it is easy to violate the grouping restrictions if you attempt to commit the buffer before creating at least two new interfaces. To prevent errors, we recommend that you wait to commit your changes to the system configuration until you have created new interfaces for all of the unified ports changing from one port mode to another.

Committing the buffer before configuring multiple interfaces will result in an error, but you do not need to start over. You can continue to configure unified ports until the configuration satisfies the aforementioned requirements.

Effect of Port Mode Changes on Data Traffic

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

**Tip**

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

Impact of Port Mode Changes on an Expansion Module

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

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

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

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

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

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

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

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

Configuring the Port Mode



Caution

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

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

In the Cisco UCS Manager CLI, there are no new commands to support Unified Ports. Instead, you change the port mode by scoping to the mode for the desired port type and then creating a new interface. When you create a new interface for an already configured slot ID and port ID, UCS Manager deletes the previously configured interface and creates a new one. If a port mode change is required because you configure a port that previously operated in Ethernet port mode to a port type in Fibre Channel port mode, UCS Manager notes the change.

Procedure

	Command or Action	Purpose
Step 1	UCS-A# <i>scope port-type-mode</i>	<p>Enters the specified port type mode for one of the following port types:</p> <p>eth-server For configuring server ports.</p> <p>eth-storage For configuring Ethernet storage ports and Ethernet storage port channels.</p> <p>eth-traffic-mon For configuring Ethernet SPAN ports.</p> <p>eth-uplink For configuring Ethernet uplink ports.</p> <p>fc-storage For configuring Fibre Channel storage ports.</p>

	Command or Action	Purpose
		<p>fc-traffic-mon</p> <p>For configuring Fibre Channel SPAN ports.</p> <p>fc-uplink</p> <p>For configuring Fibre Channel uplink ports and Fibre Channel uplink port channels.</p>
Step 2	UCS-A <i>/port-type-mode</i> # scope fabric {a b}	Enters the specified port type mode for the specified fabric.
Step 3	UCS-A <i>/port-type-mode/fabric</i> # create interface slot-id port-id	<p>Creates an interface for the specified port type.</p> <p>If you are changing the port type from Ethernet port mode to Fibre Channel port mode, or vice-versa, the following warning appears:</p> <pre>Warning: This operation will change the port mode (from Ethernet to FC or vice-versa). When committed, this change will require the module to restart.</pre>
Step 4	Create new interfaces for other ports belonging to the Ethernet or Fibre Channel port block.	There are several restrictions that govern how Ethernet and Fibre Channel ports can be arranged on a fixed or expansion module. Among other restrictions, it is required that you change ports in groups of two. Violating any of the restrictions outlined in the Guidelines for Configuring Unified Ports section will result in an error.
Step 5	UCS-A <i>/port-type-mode/fabric/interface</i> # commit-buffer	Commits the transaction to the system configuration.

Depending upon the module for which you configured the port modes, data traffic for the Cisco UCS domain is interrupted as follows:

- Fixed module—The fabric interconnect reboots. All data traffic through that fabric interconnect is interrupted. In a cluster configuration that provides high availability and includes servers with vNICs that are configured for failover, traffic fails over to the other fabric interconnect and no interruption occurs. Changing the port mode for both sides at once results in both fabric interconnects rebooting simultaneously and a complete loss of traffic until both fabric interconnects are brought back up.

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

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

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

The following example changes ports 9 and 10 on slot 1 from Ethernet uplink ports in Ethernet port mode to uplink Fibre Channel ports in Fibre Channel port mode:

```
UCS-A# scope fc-uplink
UCS-A /fc-uplink # scope fabric a
UCS-A /fc-uplink/fabric # create interface 1 9
```

```
Warning: This operation will change the port mode (from Ethernet to FC or vice-versa).
When committed, this change will require the fixed module to restart.
UCS-A /fc-uplink/fabric/interface* # up
UCS-A /fc-uplink/fabric* #create interface 1 10
Warning: This operation will change the port mode (from Ethernet to FC or vice-versa).
When committed, this change will require the fixed module to restart.
UCS-A /fc-uplink/fabric/interface* #commit-buffer
```

Configuring the Beacon LEDs for Unified Ports

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

Procedure

	Command or Action	Purpose
Step 1	UCS-A# scope fabric-interconnect {a b}	Enters fabric interconnect mode for the specified fabric.
Step 2	UCS-A /fabric # scope card slot-id	Enters card mode for the specified fixed or expansion module.
Step 3	UCS-A /fabric/card # scope beacon-led	Enters beacon LED mode.
Step 4	UCS-A /fabric/card/beacon-led # set admin-state {eth fc off}	Specifies which port mode is represented by illuminated beacon LED lights. eth All of the Unified Ports configured in Ethernet mode illuminate. fc All of the Unified Ports configured in Fibre Channel mode illuminate. off Beacon LED lights for all ports on the module are turned off.
Step 5	UCS-A /fabric/card/beacon-led # commit-buffer	Commits the transaction to the system configuration.

The following example illuminates all of the beacon lights for Unified Ports in Ethernet port mode and commits the transaction:

```
UCS-A# scope fabric-interconnect a
UCS-A /fabric # scope card 1
UCS-A /fabric/card # scope beacon-led
UCS-A /fabric/card/beacon-led # set admin-state eth
UCS-A /fabric/card/beacon-led* # commit-buffer
UCS-A /fabric/card/beacon-led #
```


Server Ports

Configuring a Server Port

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

Procedure

	Command or Action	Purpose
Step 1	UCS-A# scope eth-server	Enters Ethernet server mode.
Step 2	UCS-A /eth-server # scope fabric {a b}	Enters Ethernet server fabric mode for the specified fabric.
Step 3	UCS-A /eth-server/fabric # create interface slot-num port-num	Creates an interface for the specified Ethernet server port.
Step 4	UCS-A /eth-server/fabric # commit-buffer	Commits the transaction to the system configuration.

The following example creates an interface for Ethernet server port 12 on slot 1 of fabric B and commits the transaction:

```
UCS-A# scope eth-server
UCS-A /eth-server # scope fabric b
UCS-A /eth-server/fabric # create interface 1 12
UCS-A /eth-server/fabric* # commit-buffer
UCS-A /eth-server/fabric #
```

Unconfiguring a Server Port

Procedure

	Command or Action	Purpose
Step 1	UCS-A# scope eth-server	Enters Ethernet server mode.
Step 2	UCS-A /eth-server # scope fabric {a b}	Enters Ethernet server fabric mode for the specified fabric.
Step 3	UCS-A /eth-server/fabric # delete interface slot-num port-num	Deletes the interface for the specified Ethernet server port.
Step 4	UCS-A /eth-server/fabric # commit-buffer	Commits the transaction to the system configuration.

The following example unconfigures Ethernet server port 12 on slot 1 of fabric B and commits the transaction:

```
UCS-A# scope eth-server
UCS-A /eth-server # scope fabric b
UCS-A /eth-server/fabric # delete interface 1 12
UCS-A /eth-server/fabric* # commit-buffer
UCS-A /eth-server/fabric #
```

Uplink Ethernet Ports

Configuring an Uplink Ethernet Port

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

Procedure

	Command or Action	Purpose
Step 1	UCS-A# scope eth-uplink	Enters Ethernet uplink mode.
Step 2	UCS-A /eth-uplink # scope fabric {a b}	Enters Ethernet uplink fabric mode for the specified fabric.
Step 3	UCS-A /eth-uplink/fabric # create interface slot-num port-num	Creates an interface for the specified Ethernet uplink port.
Step 4	UCS-A /eth-uplink/fabric # set speed {10gbps 1gbps}	(Optional) Sets the speed for the specified Ethernet uplink port. Note For the 6100 series fabric interconnects, the admin speed is only configurable for the first eight ports on a 20-port fabric interconnect and the first 16 ports on a 40-port fabric interconnect.
Step 5	UCS-A /eth-uplink/fabric # commit-buffer	Commits the transaction to the system configuration.

The following example creates an interface for Ethernet uplink port 3 on slot 2 of fabric B, sets the speed to 10 gbps, and commits the transaction:

```
UCS-A# scope eth-uplink
UCS-A /eth-uplink # scope fabric b
UCS-A /eth-uplink/fabric # create interface 2 3
UCS-A /eth-uplink/fabric # set speed 10gbps
UCS-A /eth-uplink/fabric* # commit-buffer
UCS-A /eth-uplink/fabric #
```

Unconfiguring an Uplink Ethernet Port

Procedure

	Command or Action	Purpose
Step 1	UCS-A# scope eth-uplink	Enters Ethernet uplink mode.
Step 2	UCS-A /eth-uplink # scope fabric {a b}	Enters Ethernet uplink fabric mode for the specified fabric.
Step 3	UCS-A /eth-uplink/fabric # delete interface slot-num port-num	Deletes the interface for the specified Ethernet uplink port.
Step 4	UCS-A /eth-uplink/fabric # commit-buffer	Commits the transaction to the system configuration.

The following example unconfigures Ethernet uplink port 3 on slot 2 of fabric B and commits the transaction:

```
UCS-A# scope eth-uplink
UCS-A /eth-uplink # scope fabric b
UCS-A /eth-uplink/fabric # delete interface 2 3
UCS-A /eth-uplink/fabric* # commit-buffer
UCS-A /eth-uplink/fabric #
```

Appliance Ports

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

Configuring an Appliance Port

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

Procedure

	Command or Action	Purpose
Step 1	UCS-A# scope eth-storage	Enters Ethernet storage mode.
Step 2	UCS-A /eth-storage # scope fabric {a b}	Enters Ethernet storage mode for the specified fabric.
Step 3	UCS-A /eth-storage/fabric # create interface slot-num port-num	Creates an interface for the specified appliance port.
Step 4	UCS-A /eth-storage/fabric/interface # set portmode {access trunk}	(Optional) Specifies whether the port mode is access or trunk. By default, the mode is set to trunk.

	Command or Action	Purpose
Step 5	UCS-A /eth-storage/fabric/interface # set pingroupname <i>pin-group name</i>	(Optional) Specifies the appliance pin target to the specified fabric and port, or fabric and port channel.
Step 6	UCS-A /eth-storage/fabric/interface # set prio <i>sys-class-name</i>	(Optional) Specifies the QoS class for the appliance port. By default, the priority is set to best-effort. The <i>sys-class-name</i> argument can be one of the following class keywords: <ul style="list-style-type: none"> • Fc—Use this priority for QoS policies that control vHBA traffic only. • Platinum—Use this priority for QoS policies that control vNIC traffic only. • Gold—Use this priority for QoS policies that control vNIC traffic only. • Silver—Use this priority for QoS policies that control vNIC traffic only. • Bronze—Use this priority for QoS policies that control vNIC traffic only. • Best Effort—Do not use this priority. It is reserved for the Basic Ethernet traffic lane. If you assign this priority to a QoS policy and configure another system class as CoS 0, Cisco UCS Manager does not default to this system class. It defaults to the priority with CoS 0 for that traffic.
Step 7	UCS-A /eth-storage/fabric/interface # set adminspeed { 10gbps 1gbps }	(Optional) Specifies the admin speed for the interface. By default, the admin speed is set to 10gbps.
Step 8	UCS-A /eth-storage/fabric/interface # commit buffer	Commits the transaction to the system configuration.

The following example creates an interface for an appliance port 2 on slot 3 of fabric B, sets the port mode to access, pins the appliance port to a pin group called pingroup1, sets the QoS class to fc, sets the admin speed to 10 gbps, and commits the transaction:

```
UCS-A# scope eth-storage
UCS-A /eth-storage # scope fabric b
UCS-A /eth-storage/fabric # create interface 3 2
UCS-A /eth-storage/fabric* # set portmode access
UCS-A /eth-storage/fabric* # set pingroupname pingroup1
UCS-A /eth-storage/fabric* # set prio fc
UCS-A /eth-storage/fabric* # set adminspeed 10gbps
UCS-A /eth-storage/fabric* # commit-buffer
UCS-A /eth-storage/fabric #
```

What to Do Next

Assign a VLAN or target MAC address for the appliance port.

Assigning a Target MAC Address to an Appliance Port or Appliance Port Channel

The following procedure assigns a target MAC address to an appliance port. To assign a target MAC address to an appliance port channel, scope to the port channel instead of the interface.

Procedure

	Command or Action	Purpose
Step 1	UCS-A# scope eth-storage	Enters Ethernet storage mode.
Step 2	UCS-A /eth-storage# scope fabric {a b}	Enters Ethernet storage mode for the specified fabric.
Step 3	UCS-A /eth-storage/fabric# scope interface <i>slot-id port-id</i>	Enters Ethernet interface mode for the specified interface. Note To assign a target MAC address to an appliance port channel, use the scope port-channel command instead of scope interface .
Step 4	UCS-A /eth-storage/fabric/interface# create eth-target <i>eth-target name</i>	Specifies the name for the specified MAC address target.
Step 5	UCS-A /eth-storage/fabric/interface/eth-target# set mac-address <i>mac-address</i>	Specifies the MAC address in nn:nn:nn:nn:nn:nn format.

The following example assigns a target MAC address for an appliance device on port 3, slot 2 of fabric B and commits the transaction:

```
UCS-A# scope eth-storage
UCS-A /eth-storage* # scope fabric b
UCS-A /eth-storage/fabric* # scope interface 2 3
UCS-A /eth-storage/fabric/interface* # create eth-target macname
UCS-A /eth-storage/fabric/interface* # set mac-address 01:23:45:67:89:ab
UCS-A /eth-storage/fabric/interface* # commit-buffer
UCS-A /eth-storage/fabric #
```

The following example assigns a target MAC address for appliance devices on port channel 13 of fabric B and commits the transaction:

```
UCS-A# scope eth-storage
UCS-A /eth-storage* # scope fabric b
UCS-A /eth-storage/fabric* # scope port-channel 13
UCS-A /eth-storage/fabric/port-channel* # create eth-target macname
UCS-A /eth-storage/fabric/port-channel* # set mac-address 01:23:45:67:89:ab
UCS-A /eth-storage/fabric/port-channel* # commit-buffer
UCS-A /eth-storage/fabric #
```

Unconfiguring an Appliance Port

Procedure

	Command or Action	Purpose
Step 1	UCS-A # scope eth-storage	Enters Ethernet storage mode.
Step 2	UCS-A /eth-storage # scope fabric {a b}	Enters Ethernet storage mode for the specified fabric.
Step 3	UCS-A /eth-storage/fabric # delete eth-interface slot-num port-num	Deletes the interface for the specified appliance port.
Step 4	UCS-A /eth-storage/fabric # commit-buffer	Commits the transaction to the system configuration.

The following example unconfigures appliance port 3 on slot 2 of fabric B and commits the transaction:

```
UCS-A# scope eth-storage
UCS-A /eth-storage # scope fabric b
UCS-A /eth-storage/fabric # delete eth-interface 2 3
UCS-A /eth-storage/fabric* # commit-buffer
UCS-A /eth-storage/fabric #
```

FCoE and Fibre Channel Storage Ports

Configuring a Fibre Channel Storage or FCoE Port

Procedure

	Command or Action	Purpose
Step 1	UCS-A# scope fc-storage	Enters Fibre Channel storage mode.
Step 2	UCS-A /fc-storage # scope fabric {a b}	Enters Fibre Channel storage mode for the specified fabric.
Step 3	UCS-A /fc-storage/fabric # create interface {fc fcoe} slot-num port-num	Creates an interface for the specified Fibre Channel storage port.
Step 4	UCS-A /fc-storage/fabric # commit-buffer	Commits the transaction.

The following example creates an interface for Fibre Channel storage port 10 on slot 2 of fabric A and commits the transaction:

```
UCS-A# scope fc-storage
UCS-A /fc-storage # scope fabric a
```

```
UCS-A /fc-storage/fabric* # create interface fc 2 10
UCS-A /fc-storage/fabric # commit-buffer
```

What to Do Next

Assign a VSAN.

Unconfiguring a Fibre Channel Storage or FCoE Port

Procedure

	Command or Action	Purpose
Step 1	UCS-A# scope fc-storage	Enters Fibre Channel storage mode.
Step 2	UCS-A /fc-storage # scope fabric {a b}	Enters Fibre Channel storage mode for the specified fabric.
Step 3	UCS-A /fc-storage/fabric # delete interface {fc fcoe} slot-num port-num	Deletes the interface for the specified Fibre Channel or FCoE storage port.
Step 4	UCS-A /fc-storage/fabric # commit-buffer	Commits the transaction.

The following example unconfigures Fibre Channel storage port 10 on slot 2 of fabric A and commits the transaction:

```
UCS-A# scope fc-storage
UCS-A /fc-storage # scope fabric a
UCS-A /fc-storage/fabric* # delete interface fc 2 10
UCS-A /fc-storage/fabric # commit-buffer
```

Restoring a Fibre Channel Storage Port Back to an Uplink Fibre Channel Port

Procedure

	Command or Action	Purpose
Step 1	UCS-A# scope fc-uplink	Enters Fibre Channel uplink mode.
Step 2	UCS-A /fc-uplink # scope fabric {a b}	Enters Fibre Channel uplink mode for the specified fabric.
Step 3	UCS-A /fc-uplink/fabric # create interface slot-num port-num	Creates an interface for the specified Fibre Channel uplink port.
Step 4	UCS-A /fc-uplink/fabric # commit-buffer	Commits the transaction.

The following example creates an interface for Fibre Channel uplink port 10 on slot 2 of fabric A and commits the transaction:

```
UCS-A# scope fc-uplink
UCS-A /fc-uplink # scope fabric a
UCS-A /fc-uplink/fabric* # create interface 2 10
UCS-A /fc-uplink/fabric # commit-buffer
```

Default Zoning

Zoning allows you to set up access control between hosts and storage devices. When a zone is configured or the configuration is updated, this information is propagated to all the other switches in the fabric.

In Cisco UCS, the zoning configuration is inherited from an upstream switch. You cannot configure zoning or view information about your zoning configuration through Cisco UCS Manager. The only configurable zoning option in Cisco UCS Manager is whether the default zone in a VSAN (nodes not assigned to any zone) permits or denies access among its members.

When default zoning is **enabled**, all traffic is permitted among members of the default zone.

When default zoning is **disabled**, all traffic is denied among members of the default zone.

Default zoning is applied on a per-VSAN basis. You cannot enable default zoning at the fabric level.



Note

Default zoned configurations are not recommended for production deployments, which must always use direct connect Fibre Channel topologies with upstream MDS or Nexus 5000 switches.

Enabling Default Zoning

Procedure

	Command or Action	Purpose
Step 1	UCS-A# scope fc-storage	Enters Fibre Channel storage mode.
Step 2	UCS-A /fc-storage # scope vsan vsan-name vsan-id fcoe-id	Enters Fibre Channel storage VSAN mode.
Step 3	UCS-A /fc-storage/vsan # set default-zoning enabled	Enables default zoning. When default zoning is enabled , all traffic is permitted among members of the default zone.
Step 4	UCS-A /fc-storage/vsan # commit-buffer	Commits the transaction.

The following example enables default zoning for a VSAN called accounting with VSAN ID 2112 and FCoE VSAN ID 4021 and commits the transaction:

```
UCS-A# scope fc-storage
UCS-A /fc-storage # scope vsan accounting 2112 4021
UCS-A /fc-storage/vsan # set default-zoning enabled
UCS-A /fc-storage/vsan* # commit-buffer
UCS-A /fc-storage/vsan #
```


Disabling Default Zoning

Procedure

	Command or Action	Purpose
Step 1	UCS-A# scope fc-storage	Enters Fibre Channel storage mode.
Step 2	UCS-A /fc-storage # scope vsan vsan-name vsan-id fcoe-id	Enters Fibre Channel storage VSAN mode.
Step 3	UCS-A /fc-storage/vsan # set default-zoning disabled	Disables default zoning. When default zoning is disabled , all traffic is denied among members of the default zone.
Step 4	UCS-A /fc-storage/vsan # commit-buffer	Commits the transaction.

The following example disables default zoning for a VSAN called accounting with VSAN ID 2112 and FCoE VSAN ID 4021 and commits the transaction:

```
UCS-A# scope fc-storage
UCS-A /fc-storage # scope vsan accounting 2112 4021
UCS-A /fc-storage # set default-zoning disabled
UCS-A /fc-storage* # commit-buffer
UCS-A /fc-storage #
```

Uplink Ethernet Port Channels

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



Note

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

Configuring an Uplink Ethernet Port Channel

Procedure

	Command or Action	Purpose
Step 1	UCS-A# scope eth-uplink	Enters Ethernet uplink mode.

	Command or Action	Purpose
Step 2	UCS-A /eth-uplink # scope fabric {a b }	Enters Ethernet uplink fabric mode for the specified fabric.
Step 3	UCS-A /eth-uplink/fabric # create port-channel <i>port-num</i>	Creates a port channel on the specified Ethernet uplink port, and enters Ethernet uplink fabric port channel mode.
Step 4	UCS-A /eth-uplink/fabric/port-channel # { enable disable }	(Optional) Enables or disables the administrative state of the port channel. The port channel is disabled by default.
Step 5	UCS-A /eth-uplink/fabric/port-channel # set name <i>port-chan-name</i>	(Optional) Specifies the name for the port channel.
Step 6	UCS-A /eth-uplink/fabric/port-channel # set flow-control-policy <i>policy-name</i>	(Optional) Assigns the specified flow control policy to the port channel.
Step 7	UCS-A /eth-uplink/fabric/port-channel # commit-buffer	Commits the transaction to the system configuration.

The following example creates a port channel on port 13 of fabric A, sets the name to portchan13a, enables the administrative state, assigns the flow control policy named flow-con-pol432 to the port channel, and commits the transaction:

```
UCS-A# scope eth-uplink
UCS-A /eth-uplink # scope fabric a
UCS-A /eth-uplink/fabric # create port-channel 13
UCS-A /eth-uplink/fabric/port-channel* # enable
UCS-A /eth-uplink/fabric/port-channel* # set name portchan13a
UCS-A /eth-uplink/fabric/port-channel* # set flow-control-policy flow-con-pol432
UCS-A /eth-uplink/fabric/port-channel* # commit-buffer
UCS-A /eth-uplink/fabric/port-channel #
```

Unconfiguring an Uplink Ethernet Port Channel

Procedure

	Command or Action	Purpose
Step 1	UCS-A# scope eth-uplink	Enters Ethernet uplink mode.
Step 2	UCS-A /eth-uplink # scope fabric {a b }	Enters Ethernet uplink fabric mode for the specified fabric.
Step 3	UCS-A /eth-uplink/fabric # delete port-channel <i>port-num</i>	Deletes the port channel on the specified Ethernet uplink port.

	Command or Action	Purpose
Step 4	UCS-A /eth-uplink/fabric # commit-buffer	Commits the transaction to the system configuration.

The following example unconfigures the port channel on port 13 of fabric A and commits the transaction:

```
UCS-A# scope eth-uplink
UCS-A /eth-uplink # scope fabric a
UCS-A /eth-uplink/fabric # delete port-channel 13
UCS-A /eth-uplink/fabric* # commit-buffer
UCS-A /eth-uplink/fabric #
```

Adding a Member Port to an Uplink Ethernet Port Channel

Procedure

	Command or Action	Purpose
Step 1	UCS-A# scope eth-uplink	Enters Ethernet uplink mode.
Step 2	UCS-A /eth-uplink # scope fabric {a b }	Enters Ethernet uplink fabric mode for the specified fabric.
Step 3	UCS-A /eth-uplink/fabric # scope port-channel port-num	Enters Ethernet uplink fabric port channel mode for the specified port channel.
Step 4	UCS-A /eth-uplink/fabric/port-channel # create member-port slot-num port-num	Creates the specified member port from the port channel and enters Ethernet uplink fabric port channel member port mode.
Step 5	UCS-A /eth-uplink/fabric/port-channel # commit-buffer	Commits the transaction to the system configuration.

The following example adds the member port on slot 1, port 7 to the port channel on port 13 of fabric A and commits the transaction.

```
UCS-A# scope eth-uplink
UCS-A /eth-uplink # scope fabric a
UCS-A /eth-uplink/fabric # scope port-channel 13
UCS-A /eth-uplink/fabric/port-channel # create member-port 1 7
UCS-A /eth-uplink/fabric/port-channel* # commit-buffer
UCS-A /eth-uplink/fabric/port-channel #
```

Deleting a Member Port from an Uplink Ethernet Port Channel

Procedure

	Command or Action	Purpose
Step 1	UCS-A# scope eth-uplink	Enters Ethernet uplink mode.
Step 2	UCS-A /eth-uplink # scope fabric {a b }	Enters Ethernet uplink fabric mode for the specified fabric.
Step 3	UCS-A /eth-uplink/fabric # scope port-channel port-num	Enters Ethernet uplink fabric port channel mode for the specified port channel.
Step 4	UCS-A /eth-uplink/fabric/port-channel # delete member-port slot-num port-num	Deletes the specified member port from the port channel.
Step 5	UCS-A /eth-uplink/fabric/port-channel # commit-buffer	Commits the transaction to the system configuration.

The following example deletes a member port from the port channel on port 13 of fabric A and commits the transaction:

```
UCS-A# scope eth-uplink
UCS-A /eth-uplink # scope fabric a
UCS-A /eth-uplink/fabric # scope port-channel 13
UCS-A /eth-uplink/fabric/port-channel # delete member-port 1 7
UCS-A /eth-uplink/fabric/port-channel* # commit-buffer
UCS-A /eth-uplink/fabric/port-channel #
```

Appliance Port Channels

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

Configuring an Appliance Port Channel

Procedure

	Command or Action	Purpose
Step 1	UCS-A# scope eth-storage	Enters Ethernet storage mode.
Step 2	UCS-A /eth-storage # scope fabric {a b }	Enters Ethernet storage fabric mode for the specified fabric.

	Command or Action	Purpose
Step 3	UCS-A /eth-storage/fabric # create port-channel <i>port-num</i>	Creates a port channel on the specified Ethernet storage port, and enters Ethernet storage fabric port channel mode.
Step 4	UCS-A /eth-storage/fabric/port-channel # { enable disable }	(Optional) Enables or disables the administrative state of the port channel. The port channel is disabled by default.
Step 5	UCS-A /eth-storage/fabric/port-channel # set name <i>port-chan-name</i>	(Optional) Specifies the name for the port channel.
Step 6	UCS-A /eth-storage/fabric/port-channel # set pingroupname <i>pin-group name</i>	(Optional) Specifies the appliance pin target to the specified fabric and port, or fabric and port channel.
Step 7	UCS-A /eth-storage/fabric/port-channel # set portmode { access trunk }	(Optional) Specifies whether the port mode is access or trunk. By default, the mode is set to trunk.
Step 8	UCS-A /eth-storage/fabric/port-channel # set prio <i>sys-class-name</i>	(Optional) Specifies the QoS class for the appliance port. By default, the priority is set to best-effort. The <i>sys-class-name</i> argument can be one of the following class keywords: <ul style="list-style-type: none"> • Fc—Use this priority for QoS policies that control vHBA traffic only. • Platinum—Use this priority for QoS policies that control vNIC traffic only. • Gold—Use this priority for QoS policies that control vNIC traffic only. • Silver—Use this priority for QoS policies that control vNIC traffic only. • Bronze—Use this priority for QoS policies that control vNIC traffic only. • Best Effort—Do not use this priority. It is reserved for the Basic Ethernet traffic lane. If you assign this priority to a QoS policy and configure another system class as CoS 0, Cisco UCS Manager does not default to this system class. It defaults to the priority with CoS 0 for that traffic.
Step 9	UCS-A /eth-storage/fabric/port-channel # set speed { 1gbps 2gbps 4gbps 8gbps auto }	(Optional) Specifies the speed for the port channel.

	Command or Action	Purpose
Step 10	UCS-A /eth-storage/fabric/port-channel # commit-buffer	Commits the transaction to the system configuration.

The following example creates a port channel on port 13 of fabric A and commits the transaction:

```
UCS-A# scope eth-storage
UCS-A /eth-storage # scope fabric a
UCS-A /eth-storage/fabric # create port-channel 13
UCS-A /eth-storage/fabric/port-channel* # enable
UCS-A /eth-storage/fabric/port-channel* # set name portchan13a
UCS-A /eth-storage/fabric/port-channel* # set pingroupname pingroup1
UCS-A /eth-storage/fabric/port-channel* # set portmode access
UCS-A /eth-storage/fabric/port-channel* # set prio fc
UCS-A /eth-storage/fabric/port-channel* # set speed 2gbps
UCS-A /eth-storage/fabric/port-channel* # commit-buffer
UCS-A /eth-storage/fabric/port-channel #
```

Unconfiguring an Appliance Port Channel

Procedure

	Command or Action	Purpose
Step 1	UCS-A# scope eth-storage	Enters Ethernet storage mode.
Step 2	UCS-A /eth-storage # scope fabric {a b }	Enters Ethernet storage fabric mode for the specified fabric.
Step 3	UCS-A /eth-storage/fabric # delete port-channel port-num	Deletes the port channel from the specified Ethernet storage port.
Step 4	UCS-A /eth-storage/fabric # commit-buffer	Commits the transaction to the system configuration.

The following example unconfigures the port channel on port 13 of fabric A and commits the transaction:

```
UCS-A# scope eth-storage
UCS-A /eth-storage # scope fabric a
UCS-A /eth-storage/fabric # delete port-channel 13
UCS-A /eth-storage/fabric* # commit-buffer
UCS-A /eth-storage/fabric #
```

Enabling or Disabling an Appliance Port Channel

Procedure

	Command or Action	Purpose
Step 1	UCS-A# scope eth-storage	Enters Ethernet storage mode.
Step 2	UCS-A /eth-storage # scope fabric {a b }	Enters Ethernet storage mode for the specified fabric.
Step 3	UCS-A /eth-storage/fabric # scope port-channel port-chan-name	Enters Ethernet storage port channel mode.
Step 4	UCS-A /eth-storage/fabric/port-channel # {enable disable }	Enables or disables the administrative state of the port channel. The port channel is disabled by default.
Step 5	UCS-A /eth-storage/fabric/port-channel # commit-buffer	Commits the transaction to the system configuration.

The following example enables port channel 13 on fabric A and commits the transaction:

```
UCS-A# scope eth-storage
UCS-A /eth-storage # scope fabric a
UCS-A /eth-storage/fabric # scope port-channel 13
UCS-A /eth-storage/fabric/port-channel* # enable
UCS-A /eth-storage/fabric/port-channel* # commit-buffer
UCS-A /eth-storage/fabric/port-channel #
```

Adding a Member Port to an Appliance Port Channel

Procedure

	Command or Action	Purpose
Step 1	UCS-A# scope eth-storage	Enters Ethernet storage mode.
Step 2	UCS-A /eth-storage # scope fabric {a b }	Enters Ethernet storage fabric mode for the specified fabric.
Step 3	UCS-A /eth-storage/fabric # scope port-channel port-num	Enters Ethernet storage fabric port channel mode for the specified port channel.
Step 4	UCS-A /eth-storage/fabric/port-channel # create member-port slot-num port-num	Creates the specified member port from the port channel and enters Ethernet storage fabric port channel member port mode.
Step 5	UCS-A /eth-storage/fabric/port-channel # commit-buffer	Commits the transaction to the system configuration.

The following example adds the member port on slot 1, port 7 to the port channel on port 13 of fabric A and commits the transaction.

```
UCS-A# scope eth-storage
UCS-A /eth-storage # scope fabric a
UCS-A /eth-storage/fabric # scope port-channel 13
UCS-A /eth-storage/fabric/port-channel # create member-port 1 7
UCS-A /eth-storage/fabric/port-channel* # commit-buffer
UCS-A /eth-storage/fabric/port-channel #
```

Deleting a Member Port from an Appliance Port Channel

Procedure

	Command or Action	Purpose
Step 1	UCS-A# scope eth-storage	Enters Ethernet storage mode.
Step 2	UCS-A /eth-storage # scope fabric {a b }	Enters Ethernet storage fabric mode for the specified fabric.
Step 3	UCS-A /eth-storage/fabric # scope port-channel port-num	Enters Ethernet storage fabric port channel mode for the specified port channel.
Step 4	UCS-A /eth-storage/fabric/port-channel # delete member-port slot-num port-num	Deletes the specified member port from the port channel.
Step 5	UCS-A /eth-storage/fabric/port-channel # commit-buffer	Commits the transaction to the system configuration.

The following example deletes a member port from the port channel on port 13 of fabric A and commits the transaction:

```
UCS-A# scope eth-storage
UCS-A /eth-storage # scope fabric a
UCS-A /eth-storage/fabric # scope port-channel 13
UCS-A /eth-storage/fabric/port-channel # delete member-port 1 7
UCS-A /eth-storage/fabric/port-channel* # commit-buffer
UCS-A /eth-storage/fabric/port-channel #
```

Fibre Channel Port Channels

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

You can create up to four Fibre Channel port channels in each Cisco UCS domain. Each Fibre Channel port channel can include a maximum of 16 uplink Fibre Channel ports.

Configuring a Fibre Channel Port Channel



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

Procedure

	Command or Action	Purpose
Step 1	UCS-A# scope fc-uplink	Enters Fibre Channel uplink mode.
Step 2	UCS-A /fc-uplink # scope fabric {a b }	Enters Fibre Channel uplink fabric mode for the specified fabric.
Step 3	UCS-A /fc-uplink/fabric # create port-channel port-num	Creates a port channel on the specified Fibre Channel uplink port, and enters Fibre Channel uplink fabric port channel mode.
Step 4	UCS-A /fc-uplink/fabric/port-channel # {enable disable}	(Optional) Enables or disables the administrative state of the port channel. The port channel is disabled by default.
Step 5	UCS-A /fc-uplink/fabric/port-channel # set name port-chan-name	(Optional) Specifies the name for the port channel.
Step 6	UCS-A /fc-uplink/fabric/port-channel # set speed {1gbps 2gbps 4gbps 8gbps auto}	(Optional) Specifies the speed for the port channel.
Step 7	UCS-A /fc-uplink/fabric/port-channel # commit-buffer	Commits the transaction to the system configuration.

The following example creates port channel 13 on fabric A, sets the name to portchan13a, enables the administrative state, sets the speed to 2 Gbps, and commits the transaction:

```
UCS-A# scope fc-uplink
UCS-A /fc-uplink # scope fabric a
UCS-A /fc-uplink/fabric # create port-channel 13
UCS-A /fc-uplink/fabric/port-channel* # enable
UCS-A /fc-uplink/fabric/port-channel* # set name portchan13a
UCS-A /fc-uplink/fabric/port-channel* # set speed 2gbps
UCS-A /fc-uplink/fabric/port-channel* # commit-buffer
UCS-A /fc-uplink/fabric/port-channel #
```

Unconfiguring a Fibre Channel Port Channel

Procedure

	Command or Action	Purpose
Step 1	UCS-A# scope fc-uplink	Enters Fibre Channel uplink mode.
Step 2	UCS-A /fc-uplink # scope fabric {a b }	Enters Fibre Channel uplink fabric mode for the specified fabric.
Step 3	UCS-A /fc-uplink/fabric # delete port-channel port-num	Deletes the port channel on the specified Fibre Channel uplink port.
Step 4	UCS-A /fc-uplink/fabric # commit-buffer	Commits the transaction to the system configuration.

The following example unconfigures port channel 13 on fabric A and commits the transaction:

```
UCS-A# scope fc-uplink
UCS-A /fc-uplink # scope fabric a
UCS-A /fc-uplink/fabric # delete port-channel 13
UCS-A /fc-uplink/fabric* # commit-buffer
UCS-A /fc-uplink/fabric #
```

Enabling or Disabling a Fibre Channel Port Channel

Procedure

	Command or Action	Purpose
Step 1	UCS-A# scope fc-uplink	Enters Fibre Channel uplink mode.
Step 2	UCS-A /fc-uplink # scope fabric {a b }	Enters Fibre Channel uplink mode for the specified fabric.
Step 3	UCS-A /fc-uplink/fabric # scope port-channel port-chan-name	Enters Fibre Channel uplink port channel mode.
Step 4	UCS-A /fc-uplink/fabric/port-channel # {enable disable }	Enables or disables the administrative state of the port channel. The port channel is disabled by default.

The following example enables port channel 13 on fabric A and commits the transaction:

```
UCS-A# scope fc-uplink
UCS-A /fc-uplink # scope fabric a
UCS-A /fc-uplink/fabric # scope port-channel 13
UCS-A /fc-uplink/fabric/port-channel* # enable
UCS-A /fc-uplink/fabric/port-channel* # commit-buffer
UCS-A /fc-uplink/fabric/port-channel #
```

Adding a Member Port to a Fibre Channel Port Channel

Procedure

	Command or Action	Purpose
Step 1	UCS-A# scope fc-uplink	Enters Fibre Channel uplink mode.
Step 2	UCS-A /fc-uplink # scope fabric {a b}	Enters Fibre Channel uplink fabric mode for the specified fabric.
Step 3	UCS-A /fc-uplink/fabric # scope port-channel port-num	Enters Fibre Channel uplink fabric port channel mode for the specified port channel.
Step 4	UCS-A /fc-uplink/fabric/port-channel # create member-port slot-num port-num	Creates the specified member port from the port channel and enters Fibre Channel uplink fabric port channel member port mode.
Step 5	UCS-A /fc-uplink/fabric/port-channel # commit-buffer	Commits the transaction to the system configuration.

The following example adds the member port on slot 1, port 7 to port channel 13 on fabric A and commits the transaction.

```
UCS-A# scope fc-uplink
UCS-A /fc-uplink # scope fabric a
UCS-A /fc-uplink/fabric # scope port-channel 13
UCS-A /fc-uplink/fabric # create member-port 1 7
UCS-A /fc-uplink/fabric/port-channel* # commit-buffer
UCS-A /fc-uplink/fabric/port-channel #
```

Deleting a Member Port from a Fibre Channel Port Channel

Procedure

	Command or Action	Purpose
Step 1	UCS-A# scope fc-uplink	Enters Fibre Channel uplink mode.
Step 2	UCS-A /fc-uplink # scope fabric {a b}	Enters Fibre Channel uplink fabric mode for the specified fabric.
Step 3	UCS-A /fc-uplink/fabric # scope port-channel port-num	Enters Fibre Channel uplink fabric port channel mode for the specified port channel.
Step 4	UCS-A /fc-uplink/fabric/port-channel # delete member-port slot-num port-num	Deletes the specified member port from the port channel.
Step 5	UCS-A /fc-uplink/fabric/port-channel # commit-buffer	Commits the transaction to the system configuration.

The following example deletes a member port from port channel 13 on fabric A and commits the transaction:

```
UCS-A# scope fc-uplink
UCS-A /fc-uplink # scope fabric a
UCS-A /fc-uplink/fabric # scope port-channel 13
UCS-A /fc-uplink/fabric # delete member-port 1 7
UCS-A /fc-uplink/fabric/port-channel* # commit-buffer
UCS-A /fc-uplink/fabric/port-channel #
```

Adapter Port Channels

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

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

Viewing Adapter Port Channels

Procedure

	Command or Action	Purpose
Step 1	UCS-A# scope eth-server	Enters Ethernet server mode.
Step 2	UCS-A /eth-server # scope fabric {a b}	Enters Ethernet server fabric mode for the specified fabric.
Step 3	UCS-A /eth-server/fabric # show host-port-channel [detail expand]	Displays fabric port channels on the specified fabric interconnect.

This following example shows how to display information on host port channels within a fabric in the Ethernet server mode:

```
UCS-A # scope eth-server
UCS-A /eth-server # scope fabric a
UCS-A /eth-server/fabric # show host-port-channel
```

Host Port channel:

```

Port Channel ID  Chassis ID  Admin State  Oper State  State Reason
-----
1125             2             Enabled      Up
UCS-A /eth-server/fabric #
```

Fabric Port Channels

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

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

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


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

Cabling Considerations for Fabric Port Channels

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

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

Figure 1: Port Groups for Fabric Port Channels



Caution

Adding or removing links from a fabric port channel is disruptive and may affect the available amount of VIF namespace.

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

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

Configuring a Fabric Port Channel

Procedure

-
- Step 1** To include all links from the IOM to the fabric interconnect in a fabric port channel during chassis discovery, set the link grouping preference in the chassis discovery policy to port channel.
[Configuring the Chassis Discovery Policy](#)
- Step 2** To include links from individual chassis in a fabric port channel during chassis discovery, set the link grouping preference in the chassis connectivity policy to port channel.
[Configuring a Chassis Connectivity Policy](#)
- Step 3** After chassis discovery, enable or disable additional fabric port channel member ports.
[Enabling or Disabling a Fabric Port Channel Member Port](#)
-

What to Do Next

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

Viewing Fabric Port Channels

Procedure

	Command or Action	Purpose
Step 1	UCS-A# scope eth-server	Enters Ethernet server mode.
Step 2	UCS-A /eth-server # scope fabric {a b}	Enters Ethernet server fabric mode for the specified fabric.
Step 3	UCS-A /eth-server/fabric # show fabric-port-channel [detail expand]	Displays fabric port channels on the specified fabric interconnect.

The following example displays information about configured fabric port channels on fabric interconnect A:

```
UCS-A# scope eth-server
UCS-A /eth-server # scope fabric a
UCS-A /eth-server/fabric # show fabric-port-channel
Fabric Port Channel:
  Port Channel Id Chassis Id Admin State Oper State      State Reason
```

```

-----
      1025 1      Enabled  Failed  No operational members
      1026 2      Enabled  Up
UCS-A /eth-server/fabric #

```

Enabling or Disabling a Fabric Port Channel Member Port

Procedure

	Command or Action	Purpose
Step 1	UCS-A# scope eth-server	Enters Ethernet server mode.
Step 2	UCS-A /eth-server # scope fabric {a b}	Enters Ethernet server fabric mode for the specified fabric.
Step 3	UCS-A /eth-server/fabric # scope fabric-port-channel port-chan-id	Enters Ethernet server fabric, fabric port channel mode for the specified fabric.
Step 4	UCS-A /eth-server/fabric/fabric-port-channel # scope member-port slot-id port-id	Enters Ethernet server fabric, fabric port channel mode for the specified member port.
Step 5	UCS-A /eth-server/fabric/fabric-port-channel # {enable disable}	Enables or disables the specified member port.
Step 6	UCS-A /eth-server/fabric/fabric-port-channel # commit-buffer	Commits the transaction to the system configuration.

The following example disables fabric channel member port 1 31 on fabric port channel 1025 and commits the transaction:

```

UCS-A# scope eth-server
UCS-A /eth-server # scope fabric a
UCS-A /eth-server/fabric # scope fabric-port-channel 1025
UCS-A /eth-server/fabric/fabric-port-channel # scope member-port 1 31
UCS-A /eth-server/fabric/fabric-port-channel/member-port # disable
UCS-A /eth-server/fabric/fabric-port-channel/member-port* # commit-buffer
UCS-A /eth-server/fabric/fabric-port-channel/member-port #

```




CHAPTER 7

Configuring Communication Services

This chapter includes the following sections:

- [Communication Services, page 103](#)
- [Configuring CIM XML, page 104](#)
- [Configuring HTTP, page 105](#)
- [Unconfiguring HTTP, page 106](#)
- [Configuring HTTPS, page 106](#)
- [Enabling HTTP Redirection, page 116](#)
- [Configuring SNMP, page 116](#)
- [Enabling Telnet, page 123](#)
- [Disabling Communication Services, page 124](#)

Communication Services

You can use the following communication services to interface third-party applications with Cisco UCS:

Communication Service	Description
CIM XML	This service is disabled by default and is only available in read-only mode. The default port is 5988. This common information model is one of the standards defined by the Distributed Management Task Force.

Communication Service	Description
HTTP	<p>This service is enabled on port 80 by default.</p> <p>You must enable either HTTP or HTTPS to run Cisco UCS Manager GUI. If you select HTTP, all data is exchanged in clear text mode.</p> <p>For security purposes, we recommend that you enable HTTPS and disable HTTP.</p> <p>By default, Cisco UCS redirects any attempt to communicate via HTTP to the HTTPS equivalent. We recommend that you do not change this behavior.</p> <p>Note If you are upgrading to Cisco UCS, version 1.4(1), this does not happen by default. If you want to redirect any attempt to communicate via HTTP to an HTTPS equivalent, you should enable Redirect HTTP to HTTPS in Cisco UCS Manager.</p>
HTTPS	<p>This service is enabled on port 443 by default.</p> <p>With HTTPS, all data is exchanged in encrypted mode through a secure server.</p> <p>For security purposes, we recommend that you only use HTTPS and either disable or redirect HTTP communications.</p>
SMASH CLP	<p>This service is enabled for read-only access and supports a limited subset of the protocols, such as the show command. You cannot disable it.</p> <p>This shell service is one of the standards defined by the Distributed Management Task Force.</p>
SNMP	<p>This service is disabled by default. If enabled, the default port is 161. You must configure the community and at least one SNMP trap.</p> <p>Enable this service only if your system includes integration with an SNMP server.</p>
SSH	<p>This service is enabled on port 22. You cannot disable it, nor can you change the default port.</p> <p>This service provides access to the Cisco UCS Manager CLI.</p>
Telnet	<p>This service is disabled by default.</p> <p>This service provides access to the Cisco UCS Manager CLI.</p>

Configuring CIM XML

Procedure

	Command or Action	Purpose
Step 1	UCS-A# scope system	Enters system mode.
Step 2	UCS-A /system # scope services	Enters system services mode.

	Command or Action	Purpose
Step 3	UCS-A /system/services # enable cimxml	Enables the CIM XML service.
Step 4	UCS-A /system/services # set cimxml port <i>port-num</i>	Specifies the port to be used for the CIM XML connection.
Step 5	UCS-A /system/services # commit-buffer	Commits the transaction to the system configuration.

The following example enables CIM XML, sets the port number to 5988, and commits the transaction:

```
UCS-A# scope system
UCS-A /system # scope services
UCS-A /system/services # enable cimxml
UCS-A /system/services* # set cimxml port 5988
UCS-A /system/services* # commit-buffer
UCS-A /system/services #
```

Configuring HTTP

Procedure

	Command or Action	Purpose
Step 1	UCS-A# scope system	Enters system mode.
Step 2	UCS-A /system # scope services	Enters system services mode.
Step 3	UCS-A /system/services # enable http	Enables the HTTP service.
Step 4	UCS-A /system/services # set http port <i>port-num</i>	Specifies the port to be used for the HTTP connection.
Step 5	UCS-A /system/services # commit-buffer	Commits the transaction to the system configuration.

The following example enables HTTP, sets the port number to 80, and commits the transaction:

```
UCS-A# scope system
UCS-A /system # scope services
UCS-A /system/services # enable http
UCS-A /system/services* # set http port 80
Warning: When committed, this closes all the web sessions.
UCS-A /system/services* # commit-buffer
UCS-A /system/services #
```

Unconfiguring HTTP

Procedure

	Command or Action	Purpose
Step 1	UCS-A# scope system	Enters system mode.
Step 2	UCS-A /system # scope services	Enters system services mode.
Step 3	UCS-A /system/services # disable http	Disables the HTTP service.
Step 4	UCS-A /system/services # commit-buffer	Commits the transaction to the system configuration.

The following example disables HTTP and commits the transaction:

```
UCS-A# scope system
UCS-A /system # scope services
UCS-A /system/services # disable http
UCS-A /system/services* # commit-buffer
UCS-A /system/services #
```

Configuring HTTPS

Certificates, Key Rings, and Trusted Points

HTTPS uses components of the Public Key Infrastructure (PKI) to establish secure communications between two devices, such as a client's browser and Cisco UCS Manager.

Encryption Keys and Key Rings

Each PKI device holds a pair of asymmetric Rivest-Shamir-Adleman (RSA) encryption keys, one kept private and one made public, stored in an internal key ring. A message encrypted with either key can be decrypted with the other key. To send an encrypted message, the sender encrypts the message with the receiver's public key, and the receiver decrypts the message using its own private key. A sender can also prove its ownership of a public key by encrypting (also called 'signing') a known message with its own private key. If a receiver can successfully decrypt the message using the public key in question, the sender's possession of the corresponding private key is proven. Encryption keys can vary in length, with typical lengths from 512 bits to 2048 bits. In general, a longer key is more secure than a shorter key. Cisco UCS Manager provides a default key ring with an initial 1024-bit key pair, and allows you to create additional key rings.

The default key ring certificate must be manually regenerated if the cluster name changes or the certificate expires.

This operation is only available in the UCS Manager CLI.

Certificates

To prepare for secure communications, two devices first exchange their digital certificates. A certificate is a file containing a device's public key along with signed information about the device's identity. To merely

support encrypted communications, a device can generate its own key pair and its own self-signed certificate. When a remote user connects to a device that presents a self-signed certificate, the user has no easy method to verify the identity of the device, and the user's browser will initially display an authentication warning. By default, Cisco UCS Manager contains a built-in self-signed certificate containing the public key from the default key ring.

Trusted Points

To provide stronger authentication for Cisco UCS Manager, you can obtain and install a third-party certificate from a trusted source, or trusted point, that affirms the identity of your device. The third-party certificate is signed by the issuing trusted point, which can be a root certificate authority (CA) or an intermediate CA or trust anchor that is part of a trust chain that leads to a root CA. To obtain a new certificate, you must generate a certificate request through Cisco UCS Manager and submit the request to a trusted point.



Important The certificate must be in Base64 encoded X.509 (CER) format.

Creating a Key Ring

Cisco UCS Manager supports a maximum of 8 key rings, including the default key ring.

Procedure

	Command or Action	Purpose
Step 1	UCS-A# scope security	Enters security mode.
Step 2	UCS-A /security # create keyring <i>keyring-name</i>	Creates and names the key ring.
Step 3	UCS-A /security/keyring # set modulus { mod1024 mod1536 mod2048 mod512 }	Sets the SSL key length in bits.
Step 4	UCS-A /security/keyring # commit-buffer	Commits the transaction.

The following example creates a keyring with a key size of 1024 bits:

```
UCS-A# scope security
UCS-A /security # create keyring kr220
UCS-A /security/keyring* # set modulus mod1024
UCS-A /security/keyring* # commit-buffer
UCS-A /security/keyring #
```

What to Do Next

Create a certificate request for this key ring.

Regenerating the Default Key Ring

The default key ring certificate must be manually regenerated if the cluster name changes or the certificate expires.

Procedure

	Command or Action	Purpose
Step 1	UCS-A# scope security	Enters security mode.
Step 2	UCS-A /security # scope keyring default	Enters key ring security mode for the default key ring.
Step 3	UCS-A /security/keyring # set regenerate yes	Regenerates the default key ring.
Step 4	UCS-A /security/keyring # commit-buffer	Commits the transaction.

The following example regenerates the default key ring:

```
UCS-A# scope security
UCS-A /security # scope keyring default
UCS-A /security/keyring* # set regenerate yes
UCS-A /security/keyring* # commit-buffer
UCS-A /security/keyring #
```

Creating a Certificate Request for a Key Ring

Creating a Certificate Request for a Key Ring with Basic Options

Procedure

	Command or Action	Purpose
Step 1	UCS-A# scope security	Enters security mode.
Step 2	UCS-A /security # scope keyring <i>keyring-name</i>	Enters configuration mode for the key ring.
Step 3	UCS-A /security/keyring # create certreq { ip ip-address subject-name name }	Creates a certificate request using the IP address or name of the fabric interconnect. You are prompted to enter a password for the certificate request.
Step 4	UCS-A /security/keyring/certreq* # commit-buffer	Commits the transaction.
Step 5	UCS-A /security/keyring # show certreq	Displays the certificate request, which you can copy and send to a trust anchor or certificate authority.

The following example creates and displays a certificate request for a key ring with basic options:

```
UCS-A# scope security
UCS-A /security # scope keyring kr220
UCS-A /security/keyring # create certreq ip 192.168.200.123 subject-name sjc04
Certificate request password:
```

```

Confirm certificate request password:
UCS-A /security/keyring* # commit-buffer
UCS-A /security/keyring # show certreq
Certificate request subject name: sjc04
Certificate request ip address: 192.168.200.123
Certificate request e-mail name:
Certificate request country name:
State, province or county (full name):
Locality (eg, city):
Organization name (eg, company):
Organization Unit name (eg, section):
Request:
-----BEGIN CERTIFICATE REQUEST-----
MIIBfTCB5wIBADARMQ8wDQYDVQQDEwZzYW1jMDQwZ8wDQYJKoZIhvcNAQEBBQAD
gY0AMIGJAoGBALpKnlt8qMZ04UGqILKFXQQc2c8b/vW2rnRF80PhKbhghLA1YZ1F
JqcYEG5Y11+vgohLBTd45s0GC8m4RTLJWHO4SwccAUXQ5Zngf45YtX1Wsy1wUWV4
0re/zgTk/WCd56RfOBvWR2Dtztu2pGA14sd761zLxt29K7R8mzj6CAUVAgMBAAGg
LTArBgkqhkiG9w0BCQ4xHjAcMBoGA1UdEQEB/wQQMA6CBnNhbmWmWnIcECsEiXjAN
BgkqhkiG9w0BAQQFAAOBgQCcsxN0qUHYGFoQw56RwQueLTNPnrndqUwuZHUO03Teg
nhsyu4satpyiPqVV9viKZ+spvc6x5PWIcTWgHhH8BimOb/0OKuG8kwfIGGsED1Av
TTYvUP+BZ90FiPbRIA718S+V8ndXr1HejiQGx1DNqoN+odCXPC5kjoXD01ZTL09H
BA==
-----END CERTIFICATE REQUEST-----
UCS-A /security/keyring #

```

Creating a Certificate Request for a Key Ring with Advanced Options

Procedure

	Command or Action	Purpose
Step 1	UCS-A# scope security	Enters security mode.
Step 2	UCS-A /security # scope keyring <i>keyring-name</i>	Enters configuration mode for the key ring.
Step 3	UCS-A /security/keyring # create certreq	Creates a certificate request.
Step 4	UCS-A /security/keyring/certreq* # set country <i>country name</i>	Specifies the country code of the country in which the company resides.
Step 5	UCS-A /security/keyring/certreq* # set dns <i>DNS Name</i>	Specifies the Domain Name Server (DNS) address associated with the request.
Step 6	UCS-A /security/keyring/certreq* # set e-mail <i>E-mail name</i>	Specifies the email address associated with the certificate request.
Step 7	UCS-A /security/keyring/certreq* # set ip <i>certificate request ip address</i>	Specifies the IP address of the Fabric Interconnect.
Step 8	UCS-A /security/keyring/certreq* # set locality <i>locality name (eg, city)</i>	Specifies the city or town in which the company requesting the certificate is headquartered.
Step 9	UCS-A /security/keyring/certreq* # set org-name <i>organization name</i>	Specifies the organization requesting the certificate.
Step 10	UCS-A /security/keyring/certreq* # set org-unit-name <i>organizational unit name</i>	Specifies the organizational unit.

	Command or Action	Purpose
Step 11	UCS-A /security/keyring/certreq* # set password <i>certificate request password</i>	Specifies an optional password for the certificate request.
Step 12	UCS-A /security/keyring/certreq* # set state <i>state, province or county</i>	Specifies the state or province in which the company requesting the certificate is headquartered.
Step 13	UCS-A /security/keyring/certreq* # set subject-name <i>certificate request name</i>	Specifies the fully qualified domain name of the Fabric Interconnect.
Step 14	UCS-A /security/keyring/certreq* # commit-buffer	Commits the transaction.
Step 15	UCS-A /security/keyring # show certreq	Displays the certificate request, which you can copy and send to a trust anchor or certificate authority.

The following example creates and displays a certificate request for a key ring with advanced options:

```
UCS-A# scope security
UCS-A /security # scope keyring kr220
UCS-A /security/keyring # create certreq
UCS-A /security/keyring/certreq* # set ip 192.168.200.123
UCS-A /security/keyring/certreq* # set subject-name sjc04
UCS-A /security/keyring/certreq* # set country US
UCS-A /security/keyring/certreq* # set dns bg1-samc-15A
UCS-A /security/keyring/certreq* # set email test@cisco.com
UCS-A /security/keyring/certreq* # set locality new york city
UCS-A /security/keyring/certreq* # set org-name "Cisco Systems"
UCS-A /security/keyring/certreq* # set org-unit-name Testing
UCS-A /security/keyring/certreq* # set state new york
UCS-A /security/keyring/certreq* # commit-buffer
UCS-A /security/keyring/certreq # show certreq
Certificate request subject name: sjc04
Certificate request ip address: 192.168.200.123
Certificate request e-mail name: test@cisco.com
Certificate request country name: US
State, province or county (full name): New York
Locality name (eg, city): new york city
Organization name (eg, company): Cisco
Organization Unit name (eg, section): Testing
Request:
-----BEGIN CERTIFICATE REQUEST-----
MIIBfTCB5wIBADARMQ8wDQYDVQQDEwZzYW1jMDQwgZ8wDQYJKoZIhvcNAQEBBQAD
gY0AMIGJAoGBALpKn1t8qMZO4UGqILKFXQQc2c8b/vW2rnRF80PhKbhghLALYZ1F
JqcYEG5Y11+vgohLBTd45s0GC8m4RTLJWHO4SwccAUXQ5Zngf45YtX1Wsy1wUWV4
Ore/zgTk/WCd56RFOBvWR2Dtztu2pGA14sd761zLxt29K7R8mzj6CAUVAgMBAAGg
LTArBqkqhkiG9w0BCQ4xHjAcMBoGA1UdEQEB/wQMA6CBnNhbWMwNlECSEiXjAN
BgkqhkiG9w0BAQQFAAQBgcqCsxN0QUHYGFoQw56RwQueLTNPnrndqUwuZHU003Teg
nhsyu4satpyiPqVV9viKZ+spvc6x5PWlCTWgHhH8BimOb/00KuG8kwfIGGsED1Av
TTYvUP+BZ9OFiPbRIA718S+V8ndXrlHejiQGx1DNqon+odCXPc5kjoXD01ZTL09H
BA==
-----END CERTIFICATE REQUEST-----

UCS-A /security/keyring/certreq #
```


What to Do Next

- Copy the text of the certificate request, including the BEGIN and END lines, and save it in a file. Send the file with the certificate request to a trust anchor or certificate authority to obtain a certificate for the key ring.
- Create a trusted point and set the certificate chain for the certificate of trust received from the trust anchor.

Creating a Trusted Point

Procedure

	Command or Action	Purpose
Step 1	UCS-A# scope security	Enters security mode.
Step 2	UCS-A /security # create trustpoint <i>name</i>	Creates and names a trusted point.
Step 3	UCS-A /security/trustpoint # set certchain [<i>certchain</i>]	Specifies certificate information for this trusted point. If you do not specify certificate information in the command, you are prompted to enter a certificate or a list of trustpoints defining a certification path to the root certificate authority (CA). On the next line following your input, type ENDOFBUF to finish. Important The certificate must be in Base64 encoded X.509 (CER) format.
Step 4	UCS-A /security/trustpoint # commit-buffer	Commits the transaction.

The following example creates a trusted point and provides a certificate for the trusted point:

```
UCS-A# scope security
UCS-A /security # create trustpoint tPoint10
UCS-A /security/trustpoint* # set certchain
Enter lines one at a time. Enter ENDOFBUF to finish. Press ^C to abort.
Trustpoint Certificate Chain:
> -----BEGIN CERTIFICATE-----
> MIIDMDCCApmgAwIBAgIBADANBgkqhkiG9w0BAQQFADB0MQswCQYDVQQGEwJVUzEL
> BxMMU2FuIEpvc2UsIENBMRUwEwYDVQQKEwxFeGFtcGx1IEluYy4xEzARBgnVBAsT
> C1R1c3QgR3JvdXAxGTAXBgNVBAMTEHRlc3QuZXhhbXBsZS5jb20xHzAdBgkqhkiG
> 9w0BCQEWVzZXJAZXhhbXBsZS5jb20wgZ8wDQYJKoZIhvcNAQEBBQADgY0AMIGJ
> AoGBAMZw4nTepNIDhVzb0j7Z2Je4xAG56zmSHRMQeOGHemdh66u2/XAoLx7YCcYU
> ZgAMivyCsKgb/6CjQtsofvtrmC/eAehuK3/SINv7wd6Vv2pBt6ZpXgd4VBKOND1
> GmbkPayV1Qjbg4MD2dx2+H8EH3lMtdZrgKvPxPTE+bf5wZVNAgMBAAGgJTAjBgkq
> hkiG9w0BCQcxFhMUQSBjaGFsbGVuZ2UgcGFzc3dvcmQwDQYJKoZIhvcNAQEFBQAD
> gYEAG61CaJoJavMhzC190306Mg51zq1zXcz75+VFj2I6rH9asckC1d3mkOVx5gJU
> Ptt5CVQpNgNldvbDPSsXretysOhqHmp9+CLv8FDuy1CDYfuaLtv1WvfhevskV0j6
> jtcEMyZ+f7+3yh42lido3n04MIGeBgnVHSMegZYygZOAFLlNjtcEMyZ+f7+3yh42
> lido3n04oXikdjbOMQswCQYDVQQGEwJVUzELMAkGA1UECBMCQ0ExFDASBgNVBACT
> C1NhbndhRhIENsYXJhbnRwCQYDVQQKEwJ0dW92YSBTeXN0ZW1zIEluYy4xZDASBgNV
> BAsTC0Vuz21uZWVyaW5nMQ8wDQYDVQQDEwZ0ZXN0Q0GCAQAwDAYDVR0TBAUwAwEB
> /zANBgkqhkiG9w0BAQQFAAOBgQAhWaRwXNR6B4g6Lsnr+fptHv+VvhB5fKqGQqXc
> wR4pYiO4z42/j9Ijjenh75tCKMhW51az8copP1EBmOcyuhf5C6vasrenn1ddkkYt4
```

```
> PR0vxGc40whui0zBolesmsmjBbedUCwQgdFDWhDIZJwK5+N3x/kfa2EHU6id1avt
> 4YL5Jg==
> -----END CERTIFICATE-----
> ENDOFBUF
UCS-A /security/trustpoint* # commit-buffer
UCS-A /security/trustpoint #
```

What to Do Next

Obtain a key ring certificate from the trust anchor or certificate authority and import it into the key ring.

Importing a Certificate into a Key Ring

Before You Begin

- Configure a trusted point that contains the certificate chain for the key ring certificate.
- Obtain a key ring certificate from a trust anchor or certificate authority.

Procedure

	Command or Action	Purpose
Step 1	UCS-A# scope security	Enters security mode.
Step 2	UCS-A /security # scope keyring <i>keyring-name</i>	Enters configuration mode for the key ring that will receive the certificate.
Step 3	UCS-A /security/keyring # set trustpoint <i>name</i>	Specifies the trusted point for the trust anchor or certificate authority from which the key ring certificate was obtained.
Step 4	UCS-A /security/keyring # set cert	<p>Launches a dialog for entering and uploading the key ring certificate.</p> <p>At the prompt, paste the certificate text that you received from the trust anchor or certificate authority. On the next line following the certificate, type ENDOFBUF to complete the certificate input.</p> <p>Important The certificate must be in Base64 encoded X.509 (CER) format.</p>
Step 5	UCS-A /security/keyring # commit-buffer	Commits the transaction.

The following example specifies the trust point and imports a certificate into a key ring:

```
UCS-A# scope security
UCS-A /security # scope keyring kr220
UCS-A /security/keyring # set trustpoint tPoint10
UCS-A /security/keyring* # set cert
Enter lines one at a time. Enter ENDOFBUF to finish. Press ^C to abort.
Keyring certificate:
> -----BEGIN CERTIFICATE-----
> MIIB/zCCAWgCAQAwZkxkCzAJBgNVBAYTA1VMTQswCOYDVOQIEwJDQTEVMBMGAlUE
> BxMMU2FuIEpvc2UsIENBMRUwEwyDVQKEwxFeGFtcGx1IEluYy4xEzARBGNVBAsT
> ClRlc3QgR3JvdXAxGTAXBgNVBAMTEHRlc3QuZXhhbXBsZS5jb20xHzAdBgkqhkiG
> 9w0BCQEWZHVzZXJAZXhhbXBsZS5jb20wgZ8wDQYJKoZIhvcNAQEBBQADgY0AMIGJ
```

```

> AoGBAMZw4nTepNIDhVzb0j7Z2Je4xAG56zmSHRMQeOGHemdh66u2/XAoLx7YCcYU
> ZgAMivYCsKgb/6CjQtsofvtrmC/eAehuK3/SINv7wd6Vv2pBt6ZpXgd4VBNKOND1
> GmbkPayVlQjbG4MD2dx2+H8EH3LmtdZrgKvPxPTE+bf5wZVNAgMBAAGgJTAjBgkq
> hkiG9w0BCQcxFhMUQSBjaGFsbGVuZ2UgcGFzc3dvcmQwDQYJKoZIhvcNAQEFBQAD
> gYEAG61CaJoJaVMhzC190306Mg51zq1zXcz75+VFj2I6rH9asckClD3mkOVx5gJU
> Ptt5CVQpNgNldvbDPSsXretysOhqHmp9+CLv8FDuy1CDYfuaLtv1WvfhevskV0j6
> mK3Ku+YiORnv6DhxrOoqau8r/hyI/L43l7IPN1HhOi3oha4=
> -----END CERTIFICATE-----
> ENDOFBUF
UCS-A /security/keyring* # commit-buffer
UCS-A /security/keyring #

```

What to Do Next

Configure your HTTPS service with the key ring.

Configuring HTTPS



Caution

After you complete the HTTPS configuration, including changing the port and key ring to be used by HTTPS, all current HTTP and HTTPS sessions are closed without warning as soon as you save or commit the transaction.

Procedure

	Command or Action	Purpose
Step 1	UCS-A# scope system	Enters system mode.
Step 2	UCS-A /system # scope services	Enters system services mode.
Step 3	UCS-A /system/services # enable https	Enables the HTTPS service.
Step 4	UCS-A /system/services # set https port <i>port-num</i>	(Optional) Specifies the port to be used for the HTTPS connection.
Step 5	UCS-A /system/services # set https keyring <i>keyring-name</i>	(Optional) Specifies the name of the key ring you created for HTTPS.
Step 6	UCS-A /system/services # set https cipher-suite-mode <i>cipher-suite-mode</i>	(Optional) The level of Cipher Suite security used by the Cisco UCS domain. <i>cipher-suite-mode</i> can be one of the following keywords: <ul style="list-style-type: none"> • high-strength • medium-strength • low-strength • custom—Allows you to specify a user-defined Cipher Suite specification string.

	Command or Action	Purpose
Step 7	UCS-A /system/services # set https cipher-suite cipher-suite-spec-string	(Optional) Specifies a custom level of Cipher Suite security for this Cisco UCS domain if cipher-suite-mode is set to custom . <i>cipher-suite-spec-string</i> can contain up to 256 characters and must conform to the OpenSSL Cipher Suite specifications. You cannot use any spaces or special characters except ! (exclamation point), + (plus sign), - (hyphen), and : (colon). For details, see http://httpd.apache.org/docs/2.0/mod/mod_ssl.html#sslcipher-suite.. For example, the medium strength specification string Cisco UCS Manager uses as the default is: ALL:!ADH:!EXPORT56:!LOW:RC4+RSA:+HIGH:+MEDIUM:+EXP:+eNULL Note This option is ignored if cipher-suite-mode is set to anything other than custom .
Step 8	UCS-A /system/services # commit-buffer	Commits the transaction to the system configuration.

The following example enables HTTPS, sets the port number to 443, sets the key ring name to kring7984, sets the Cipher Suite security level to high, and commits the transaction:

```
UCS-A# scope system
UCS-A /system # scope services
UCS-A /system/services # enable https
UCS-A /system/services* # set https port 443
Warning: When committed, this closes all the web sessions.
UCS-A /system/services* # set https keyring kring7984
UCS-A /system/services* # set https cipher-suite-mode high
UCS-A /system/services* # commit-buffer
UCS-A /system/services #
```

Deleting a Key Ring

Procedure

	Command or Action	Purpose
Step 1	UCS-A# scope security	Enters security mode.
Step 2	UCS-A /security # delete keyring name	Deletes the named key ring.
Step 3	UCS-A /security # commit-buffer	Commits the transaction.

The following example deletes a key ring:

```
UCS-A# scope security
UCS-A /security # delete keyring key10
UCS-A /security* # commit-buffer
UCS-A /security #
```

Deleting a Trusted Point

Before You Begin

Ensure that the trusted point is not used by a key ring.

Procedure

	Command or Action	Purpose
Step 1	UCS-A# scope security	Enters security mode.
Step 2	UCS-A /security # delete trustpoint name	Deletes the named trusted point.
Step 3	UCS-A /security # commit-buffer	Commits the transaction.

The following example deletes a trusted point:

```
UCS-A# scope security
UCS-A /security # delete trustpoint tPoint10
UCS-A /security* # commit-buffer
UCS-A /security #
```

Unconfiguring HTTPS

Before You Begin

Disable HTTP to HTTPS redirection.

Procedure

	Command or Action	Purpose
Step 1	UCS-A# scope system	Enters system mode.
Step 2	UCS-A /system # scope services	Enters system services mode.
Step 3	UCS-A /system/services # disable https	Disables the HTTPS service.
Step 4	UCS-A /system/services # commit-buffer	Commits the transaction to the system configuration.

The following example disables HTTPS and commits the transaction:

```
UCS-A# scope system
UCS-A /system # scope services
UCS-A /system/services # disable https
UCS-A /system/services* # commit-buffer
UCS-A /system/services #
```

Enabling HTTP Redirection

Before You Begin

Enable both HTTP and HTTPS.

Procedure

	Command or Action	Purpose
Step 1	UCS-A# scope system	Enters system mode.
Step 2	UCS-A /system # scope services	Enters system services mode.
Step 3	UCS-A /system/services # enable http-redirect	Enables the HTTP redirect service. If enabled, all attempts to communicate via HTTP are redirected to the equivalent HTTPS address. This option effectively disables HTTP access to this Cisco UCS domain.
Step 4	UCS-A /system/services # commit-buffer	Commits the transaction to the system configuration.

The following example enables HTTP to HTTPS redirection and commits the transaction:

```
UCS-A# scope system
UCS-A /system # scope services
UCS-A /system/services # enable http-redirect
Warning: When committed, this closes all the web sessions.
UCS-A /system/services* # commit-buffer
UCS-A /system/services #
```

Configuring SNMP

Information about SNMP

The Simple Network Management Protocol (SNMP) is an application-layer protocol that provides a message format for communication between SNMP managers and agents. SNMP provides a standardized framework and a common language used for the monitoring and management of devices in a network.

SNMP Functional Overview

The SNMP framework consists of three parts:

- An SNMP manager—The system used to control and monitor the activities of network devices using SNMP.
- An SNMP agent—The software component within Cisco UCS, the managed device, that maintains the data for Cisco UCS and reports the data, as needed, to the SNMP manager. Cisco UCS includes the

agent and a collection of MIBs. To enable the SNMP agent and create the relationship between the manager and agent, enable and configure SNMP in Cisco UCS Manager.

- A managed information base (MIB)—The collection of managed objects on the SNMP agent. Cisco UCS release 1.4(1) and higher support a larger number of MIBs than earlier releases.

Cisco UCS supports SNMPv1, SNMPv2c and SNMPv3. Both SNMPv1 and SNMPv2c use a community-based form of security. SNMP is defined in the following:

- RFC 3410 (<http://tools.ietf.org/html/rfc3410>)
- RFC 3411 (<http://tools.ietf.org/html/rfc3411>)
- RFC 3412 (<http://tools.ietf.org/html/rfc3412>)
- RFC 3413 (<http://tools.ietf.org/html/rfc3413>)
- RFC 3414 (<http://tools.ietf.org/html/rfc3414>)
- RFC 3415 (<http://tools.ietf.org/html/rfc3415>)
- RFC 3416 (<http://tools.ietf.org/html/rfc3416>)
- RFC 3417 (<http://tools.ietf.org/html/rfc3417>)
- RFC 3418 (<http://tools.ietf.org/html/rfc3418>)
- RFC 3584 (<http://tools.ietf.org/html/rfc3584>)

SNMP Notifications

A key feature of SNMP is the ability to generate notifications from an SNMP agent. These notifications do not require that requests be sent from the SNMP manager. Notifications can indicate improper user authentication, restarts, the closing of a connection, loss of connection to a neighbor router, or other significant events.

Cisco UCS Manager generates SNMP notifications as either traps or informs. Traps are less reliable than informs because the SNMP manager does not send any acknowledgment when it receives a trap, and Cisco UCS Manager cannot determine if the trap was received. An SNMP manager that receives an inform request acknowledges the message with an SNMP response protocol data unit (PDU). If the Cisco UCS Manager does not receive the PDU, it can send the inform request again.

SNMP Security Levels and Privileges

SNMPv1, SNMPv2c, and SNMPv3 each represent a different security model. The security model combines with the selected security level to determine the security mechanism applied when the SNMP message is processed.

The security level determines the privileges required to view the message associated with an SNMP trap. The privilege level determines whether the message needs to be protected from disclosure or authenticated. The supported security level depends upon which security model is implemented. SNMP security levels support one or more of the following privileges:

- noAuthNoPriv—No authentication or encryption
- authNoPriv—Authentication but no encryption
- authPriv—Authentication and encryption

SNMPv3 provides for both security models and security levels. A security model is an authentication strategy that is set up for a user and the role in which the user resides. A security level is the permitted level of security within a security model. A combination of a security model and a security level determines which security mechanism is employed when handling an SNMP packet.

Supported Combinations of SNMP Security Models and Levels

The following table identifies what the combinations of security models and levels mean.

Table 9: SNMP Security Models and Levels

Model	Level	Authentication	Encryption	What Happens
v1	noAuthNoPriv	Community string	No	Uses a community string match for authentication.
v2c	noAuthNoPriv	Community string	No	Uses a community string match for authentication.
v3	noAuthNoPriv	Username	No	Uses a username match for authentication.
v3	authNoPriv	HMAC-MD5 or HMAC-SHA	No	Provides authentication based on the Hash-Based Message Authentication Code (HMAC) Message Digest 5 (MD5) algorithm or the HMAC Secure Hash Algorithm (SHA).
v3	authPriv	HMAC-MD5 or HMAC-SHA	DES	Provides authentication based on the HMAC-MD5 or HMAC-SHA algorithms. Provides Data Encryption Standard (DES) 56-bit encryption in addition to authentication based on the Cipher Block Chaining (CBC) DES (DES-56) standard.

SNMPv3 Security Features

SNMPv3 provides secure access to devices by a combination of authenticating and encrypting frames over the network. SNMPv3 authorizes management operations only by configured users and encrypts SNMP messages. The SNMPv3 User-Based Security Model (USM) refers to SNMP message-level security and offers the following services:

- Message integrity—Ensures that messages have not been altered or destroyed in an unauthorized manner and that data sequences have not been altered to an extent greater than can occur non-maliciously.
- Message origin authentication—Ensures that the claimed identity of the user on whose behalf received data was originated is confirmed.
- Message confidentiality and encryption—Ensures that information is not made available or disclosed to unauthorized individuals, entities, or processes.

SNMP Support in Cisco UCS

Cisco UCS provides the following support for SNMP:

Support for MIBs

Cisco UCS supports read-only access to MIBs.

For information about the specific MIBs available for Cisco UCS and where you can obtain them, see the [MIB Quick Reference for Cisco UCS](#).

Authentication Protocols for SNMPv3 Users

Cisco UCS supports the following authentication protocols for SNMPv3 users:

- HMAC-MD5-96 (MD5)
- HMAC-SHA-96 (SHA)

AES Privacy Protocol for SNMPv3 Users

Cisco UCS uses Advanced Encryption Standard (AES) as one of the privacy protocols for SNMPv3 message encryption and conforms with RFC 3826.

The privacy password, or `priv` option, offers a choice of DES or 128-bit AES encryption for SNMP security encryption. If you enable AES-128 configuration and include a privacy password for an SNMPv3 user, Cisco UCS Manager uses the privacy password to generate a 128-bit AES key. The AES privacy password can have a minimum of eight characters. If the passphrases are specified in clear text, you can specify a maximum of 64 characters.

Enabling SNMP and Configuring SNMP Properties

SNMP messages from a Cisco UCS domain display the fabric interconnect name rather than the system name.

Procedure

	Command or Action	Purpose
Step 1	UCS-A# scope monitoring	Enters monitoring mode.
Step 2	UCS-A /monitoring # enable snmp	Enables SNMP.
Step 3	UCS-A /monitoring # set snmp community <i>community-name</i>	Specifies SNMP community. The community name can be any alphanumeric string up to 32 characters.
Step 4	UCS-A /monitoring # set snmp syscontact <i>system-contact-name</i>	Specifies the system contact person responsible for the SNMP. The system contact name can be any alphanumeric string up to 255 characters, such as an email address or name and telephone number.
Step 5	UCS-A /monitoring # set snmp syslocation <i>system-location-name</i>	Specifies the location of the host on which the SNMP agent (server) runs. The system location name can be any alphanumeric string up to 512 characters.
Step 6	UCS-A /monitoring # commit-buffer	Commits the transaction to the system configuration.

The following example enables SNMP, configures an SNMP community named `SnmCommSystem2`, configures a system contact named `contactperson`, configures a contact location named `systemlocation`, and commits the transaction:

```
UCS-A# scope monitoring
UCS-A /monitoring # enable snmp
UCS-A /monitoring* # set snmp community SnmCommSystem2
UCS-A /monitoring* # set snmp syscontact contactperson1
UCS-A /monitoring* # set snmp syslocation systemlocation
UCS-A /monitoring* # commit-buffer
UCS-A /monitoring #
```

What to Do Next

Create SNMP traps and users.

Creating an SNMP Trap**Procedure**

	Command or Action	Purpose
Step 1	UCS-A# scope monitoring	Enters monitoring mode.
Step 2	UCS-A /monitoring # enable snmp	Enables SNMP.
Step 3	UCS-A /monitoring # create snmp-trap <i>{hostname ip-addr}</i>	Creates an SNMP trap host with the specified hostname or IP address.

	Command or Action	Purpose
Step 4	UCS-A /monitoring/snmp-trap # set community <i>community-name</i>	Specifies the SNMP community name to be used for the SNMP trap.
Step 5	UCS-A /monitoring/snmp-trap # set port <i>port-num</i>	Specifies the port to be used for the SNMP trap.
Step 6	UCS-A /monitoring/snmp-trap # set version {v1 v2c v3}	Specifies the SNMP version and model used for the trap.
Step 7	UCS-A /monitoring/snmp-trap # set notification type {traps informs}	(Optional) If you select v2c or v3 for the version, the type of trap to send.
Step 8	UCS-A /monitoring/snmp-trap # set v3 privilege {auth noauth priv}	(Optional) If you select v3 for the version, the privilege associated with the trap. This can be: <ul style="list-style-type: none"> • auth—Authentication but no encryption • noauth—No authentication or encryption • priv—Authentication and encryption
Step 9	UCS-A /monitoring/snmp-trap # commit-buffer	Commits the transaction to the system configuration.

The following example enables SNMP, creates an SNMP trap, specifies that the trap will use the SnmpCommSystem2 community on port 2, sets the version to v3, sets the notification type to traps, sets the v3 privilege to priv, and commits the transaction:

```
UCS-A# scope monitoring
UCS-A /monitoring # enable snmp
UCS-A /monitoring* # create snmp-trap 192.168.100.112
UCS-A /monitoring/snmp-trap* # set community SnmpCommSystem2
UCS-A /monitoring/snmp-trap* # set port 2
UCS-A /monitoring/snmp-trap* # set version v3
UCS-A /monitoring/snmp-trap* # set notificationtype traps
UCS-A /monitoring/snmp-trap* # set v3 privilege priv
UCS-A /monitoring/snmp-trap* # commit-buffer
UCS-A /monitoring/snmp-trap #
```

Deleting an SNMP Trap

Procedure

	Command or Action	Purpose
Step 1	UCS-A# scope monitoring	Enters monitoring mode.

	Command or Action	Purpose
Step 2	UCS-A /monitoring # delete snmp-trap { <i>hostname</i> <i>ip-addr</i> }	Deletes the specified SNMP trap host with the specified hostname or IP address.
Step 3	UCS-A /monitoring # commit-buffer	Commits the transaction to the system configuration.

The following example deletes the SNMP trap at IP address 192.168.100.112 and commits the transaction:

```
UCS-A# scope monitoring
UCS-A /monitoring # delete snmp-trap 192.168.100.112
UCS-A /monitoring* # commit-buffer
UCS-A /monitoring #
```

Creating an SNMPv3 User

Procedure

	Command or Action	Purpose
Step 1	UCS-A# scope monitoring	Enters monitoring mode.
Step 2	UCS-A /monitoring # enable snmp	Enables SNMP.
Step 3	UCS-A /monitoring # create snmp-user <i>user-name</i>	Creates the specified SNMPv3 user. An SNMP username cannot be the same as a local username. Choose an SNMP username that does not match a local username.
Step 4	UCS-A /monitoring/snmp-user # set aes-128 { <i>no</i> <i>yes</i> }	Enables or disables the use of AES-128 encryption.
Step 5	UCS-A /monitoring/snmp-user # set auth { <i>md5</i> <i>sha</i> }	Specifies the use of MD5 or DHA authentication.
Step 6	UCS-A /monitoring/snmp-user # set password	Specifies the user password. After you enter the set password command, you are prompted to enter and confirm the password.
Step 7	UCS-A /monitoring/snmp-user # set priv-password	Specifies the user privacy password. After you enter the set priv-password command, you are prompted to enter and confirm the privacy password.
Step 8	UCS-A /monitoring/snmp-user # commit-buffer	Commits the transaction to the system configuration.

The following example enables SNMP, creates an SNMPv3 user named snmp-user14, disables AES-128 encryption, specifies the use of MD5 authentication, sets the password and privacy password, and commits the transaction:

```
UCS-A# scope monitoring
UCS-A /monitoring # enable snmp
UCS-A /monitoring* # create snmp-user snmp-user14
UCS-A /monitoring/snmp-user* # set aes-128 no
UCS-A /monitoring/snmp-user* # set auth md5
UCS-A /monitoring/snmp-user* # set password
Enter a password:
Confirm the password:
UCS-A /monitoring/snmp-user* # set priv-password
Enter a password:
Confirm the password:
UCS-A /monitoring/snmp-user* # commit-buffer
UCS-A /monitoring/snmp-user #
```

Deleting an SNMPv3 User

Procedure

	Command or Action	Purpose
Step 1	UCS-A# scope monitoring	Enters monitoring mode.
Step 2	UCS-A /monitoring # delete snmp-user <i>user-name</i>	Deletes the specified SNMPv3 user.
Step 3	UCS-A /monitoring # commit-buffer	Commits the transaction to the system configuration.

The following example deletes the SNMPv3 user named snmp-user14 and commits the transaction:

```
UCS-A# scope monitoring
UCS-A /monitoring # delete snmp-user snmp-user14
UCS-A /monitoring* # commit-buffer
UCS-A /monitoring #
```

Enabling Telnet

Procedure

	Command or Action	Purpose
Step 1	UCS-A# scope system	Enters system mode.
Step 2	UCS-A /system # scope services	Enters system services mode.
Step 3	UCS-A /services # enable telnet-server	Enables the Telnet service.
Step 4	UCS-A /services # commit-buffer	Commits the transaction to the system configuration.

The following example enables Telnet and commits the transaction:

```
UCS-A# scope system
UCS-A /system # scope services
UCS-A /services # enable telnet-server
UCS-A /services* # commit-buffer
UCS-A /services #
```

Disabling Communication Services

Procedure

	Command or Action	Purpose
Step 1	UCS-A# scope system	Enters system mode.
Step 2	UCS-A /system # scope services	Enters system services mode.
Step 3	UCS-A /system/services # disable <i>service-name</i>	Disables the specified service, where the <i>service-name</i> argument is one of the following keywords: <ul style="list-style-type: none"> • <i>cimxml</i> —Disables CIM XML service • http —Disables HTTP service • https —Disables HTTPS service • telnet-server —Disables Telnet service
Step 4	UCS-A /system/services # commit-buffer	Commits the transaction to the system configuration.

The following example disables CIM XML and commits the transaction:

```
UCS-A# scope system
UCS-A# scope services
UCS-A /system/services # disable cimxml
UCS-A /system/services* # commit-buffer
UCS-A /system/services #
```



CHAPTER 8

Configuring Authentication

This chapter includes the following sections:

- [Authentication Services, page 125](#)
- [Guidelines and Recommendations for Remote Authentication Providers, page 125](#)
- [User Attributes in Remote Authentication Providers, page 126](#)
- [LDAP Group Rule, page 128](#)
- [Configuring LDAP Providers, page 128](#)
- [Configuring RADIUS Providers, page 135](#)
- [Configuring TACACS+ Providers, page 138](#)
- [Configuring Multiple Authentication Systems, page 141](#)
- [Selecting a Primary Authentication Service, page 148](#)

Authentication Services

Cisco UCS supports two methods to authenticate user logins:

- Through user accounts local to Cisco UCS Manager
- Remotely through one of the following protocols:
 - LDAP
 - RADIUS
 - TACACS+

Guidelines and Recommendations for Remote Authentication Providers

If a system is configured for one of the supported remote authentication services, you must create a provider for that service to ensure that Cisco UCS Manager can communicate with it. In addition, you need to be aware of the following guidelines that impact user authorization:

User Accounts in Remote Authentication Services

User accounts can exist locally in Cisco UCS Manager or in the remote authentication server.

The temporary sessions for users who log in through remote authentication services can be viewed through Cisco UCS Manager GUI or Cisco UCS Manager CLI.

User Roles in Remote Authentication Services

If you create user accounts in the remote authentication server, you must ensure that the accounts include the roles those users require for working in Cisco UCS Manager and that the names of those roles match the names used in Cisco UCS Manager. Depending on the role policy, a user may not be allowed to log in or will be granted only read-only privileges.

User Attributes in Remote Authentication Providers

For RADIUS and TACACS+ configurations, you must configure a user attribute for Cisco UCS in each remote authentication provider through which users log in to Cisco UCS Manager. This user attribute holds the roles and locales assigned to each user.



Note

This step is not required for LDAP configurations that use LDAP Group Mapping to assign roles and locales.

When a user logs in, Cisco UCS Manager does the following:

- 1 Queries the remote authentication service.
- 2 Validates the user.
- 3 If the user is validated, checks for the roles and locales assigned to that user.

The following table contains a comparison of the user attribute requirements for the remote authentication providers supported by Cisco UCS.

Table 10: Comparison of User Attributes by Remote Authentication Provider

Authentication Provider	Custom Attribute	Schema Extension	Attribute ID Requirements
LDAP	Not required if group mapping is used Optional if group mapping is not used	Optional. You can choose to do either of the following: <ul style="list-style-type: none"> • Do not extend the LDAP schema and configure an existing, unused attribute that meets the requirements. • Extend the LDAP schema and create a custom attribute with a unique name, such as CiscoAVPair. 	The Cisco LDAP implementation requires a unicode type attribute. If you choose to create the CiscoAVPair custom attribute, use the following attribute ID: 1.3.6.1.4.1.9.287247.1 A sample OID is provided in the following section.

Authentication Provider	Custom Attribute	Schema Extension	Attribute ID Requirements
RADIUS	Optional	<p>Optional. You can choose to do either of the following:</p> <ul style="list-style-type: none"> Do not extend the RADIUS schema and use an existing, unused attribute that meets the requirements. Extend the RADIUS schema and create a custom attribute with a unique name, such as cisco-avpair. 	<p>The vendor ID for the Cisco RADIUS implementation is 009 and the vendor ID for the attribute is 001.</p> <p>The following syntax example shows how to specify multiples user roles and locales if you choose to create the cisco-avpair attribute: shell:roles="admin,aaa" shell:locales="L1,abc". Use a comma "," as the delimiter to separate multiple values.</p>
TACACS+	Required	<p>Required. You must extend the schema and create a custom attribute with the name cisco-av-pair.</p>	<p>The cisco-av-pair name is the string that provides the attribute ID for the TACACS+ provider.</p> <p>The following syntax example shows how to specify multiples user roles and locales when you create the cisco-av-pair attribute: cisco-av-pair=shell:roles="admin aaa" shell:locales*"L1 abc". Using an asterisk (*) in the cisco-av-pair attribute syntax flags the locale as optional, preventing authentication failures for other Cisco devices that use the same authorization profile. Use a space as the delimiter to separate multiple values.</p>

Sample OID for LDAP User Attribute

The following is a sample OID for a custom CiscoAVPair attribute:

```

CN=CiscoAVPair,CN=Schema,
CN=Configuration,CN=X
objectClass: top
objectClass: attributeSchema
cn: CiscoAVPair
distinguishedName: CN=CiscoAVPair,CN=Schema,CN=Configuration,CN=X
instanceType: 0x4
uSNCreated: 26318654
attributeID: 1.3.6.1.4.1.9.287247.1
attributeSyntax: 2.5.5.12
isSingleValued: TRUE
showInAdvancedViewOnly: TRUE
adminDisplayName: CiscoAVPair
adminDescription: UCS User Authorization Field
oMSyntax: 64

```

```

LDAPDisplayName: CiscoAVPair
name: CiscoAVPair
objectCategory: CN=Attribute-Schema,CN=Schema,CN=Configuration,CN=X

```

LDAP Group Rule

The LDAP group rule is used to determine whether Cisco UCS should use LDAP groups when assigning user roles and locales to a remote user.

Configuring LDAP Providers

Configuring Properties for LDAP Providers

The properties that you configure in this task are the default settings for all provider connections of this type defined in Cisco UCS Manager. If an individual provider includes a setting for any of these properties, Cisco UCS uses that setting and ignores the default setting.

If you are using Active Directory as your LDAP server, create a user account in the Active Directory server to bind with Cisco UCS. This account should be given a non-expiring password.

Procedure

	Command or Action	Purpose
Step 1	UCS-A# scope security	Enters security mode.
Step 2	UCS-A /security # scope ldap	Enters security LDAP mode.
Step 3	UCS-A /security/ldap # set attribute attribute	Restricts database searches to records that contain the specified attribute.
Step 4	UCS-A /security/ldap # set basedn distinguished-name	Restricts database searches to records that contain the specified distinguished name.
Step 5	UCS-A /security/ldap # set filter filter	Restricts database searches to records that contain the specified filter.
Step 6	UCS-A /security/ldap # set timeout seconds	(Optional) Sets the time interval the system waits for a response from the LDAP server before noting the server as down.
Step 7	UCS-A /security/ldap # commit-buffer	Commits the transaction to the system configuration.

The following example sets the LDAP attribute to CiscoAvPair, the base distinguished name to "DC=cisco-ucsm-aaa3,DC=qalab,DC=com", the filter to sAMAccountName=\$userid, and the timeout interval to 5 seconds, and commits the transaction:

```

UCS-A# scope security
UCS-A /security # scope ldap

```

```

UCS-A /security/ldap # set attribute CiscoAvPair
UCS-A /security/ldap* # set basedn "DC=cisco-ucsm-aaa3,DC=qalab,DC=com"
UCS-A /security/ldap* # set filter sAMAccountName=$userid
UCS-A /security/ldap* # set timeout 5
UCS-A /security/ldap* # commit-buffer
UCS-A /security/ldap #

```

What to Do Next

Create an LDAP provider.

Creating an LDAP Provider

Cisco UCS Manager supports a maximum of 16 LDAP providers.

Before You Begin

If you are using Active Directory as your LDAP server, create a user account in the Active Directory server to bind with Cisco UCS. This account should be given a non-expiring password.

- In the LDAP server, perform one of the following configurations:
 - Configure LDAP groups. LDAP groups contain user role and locale information.
 - Configure users with the attribute that holds the user role and locale information for Cisco UCS Manager. You can choose whether to extend the LDAP schema for this attribute. If you do not want to extend the schema, use an existing LDAP attribute to hold the Cisco UCS user roles and locales. If you prefer to extend the schema, create a custom attribute, such as the CiscoAVPair attribute.

The Cisco LDAP implementation requires a unicode type attribute.

If you choose to create the CiscoAVPair custom attribute, use the following attribute ID:
1.3.6.1.4.1.9.287247.1
 - For a cluster configuration, add the management port IP addresses for both fabric interconnects. This configuration ensures that remote users can continue to log in if the first fabric interconnect fails and the system fails over to the second fabric interconnect. All login requests are sourced from these IP addresses, not the virtual IP address used by Cisco UCS Manager.
- If you want to use secure communications, create a trusted point containing the certificate of the root certificate authority (CA) of the LDAP server in Cisco UCS Manager.

Procedure

	Command or Action	Purpose
Step 1	UCS-A# scope security	Enters security mode.
Step 2	UCS-A /security # scope ldap	Enters security LDAP mode.
Step 3	UCS-A /security/ldap # create server <i>server-name</i>	Creates an LDAP server instance and enters security LDAP server mode. If SSL is enabled, the <i>server-name</i> , typically an IP address or FQDN, must exactly match a Common Name (CN) in the LDAP server's security certificate. Unless an IP address is specified, a DNS server must be configured in Cisco UCS Manager.

	Command or Action	Purpose
Step 4	UCS-A /security/ldap/server # set attribute <i>attr-name</i>	(Optional) An LDAP attribute that stores the values for the user roles and locales. This property is always a name-value pair. The system queries the user record for the value that matches this attribute name. If you do not want to extend your LDAP schema, you can configure an existing, unused LDAP attribute with the Cisco UCS roles and locales. Alternatively, you can create an attribute named CiscoAVPair in the remote authentication service with the following attribute ID: 1.3.6.1.4.1.9.287247.1 This value is required unless a default attribute has been set on the LDAP General tab.
Step 5	UCS-A /security/ldap/server # set basedn <i>basedn-name</i>	(Optional) The specific distinguished name in the LDAP hierarchy where the server should begin a search when a remote user logs in and the system attempts to get the user's DN based on their username. The maximum supported string length is 127 characters. This value is required unless a default base DN has been set on the LDAP General tab.
Step 6	UCS-A /security/ldap/server # set binddn <i>binddn-name</i>	(Optional) The distinguished name (DN) for an LDAP database account that has read and search permissions for all objects under the base DN. The maximum supported string length is 127 ASCII characters.
Step 7	UCS-A /security/ldap/server # set filter <i>filter-value</i>	(Optional) The LDAP search is restricted to those usernames that match the defined filter. This value is required unless a default filter has been set on the LDAP General tab.
Step 8	UCS-A /security/ldap/server # set password	The password for the LDAP database account specified in the Bind DN field. You can enter any standard ASCII characters except for space, § (section sign), ? (question mark), or = (equal sign). To set the password, press Enter after typing the set password command and enter the key value at the prompt.
Step 9	UCS-A /security/ldap/server # set order <i>order-num</i>	(Optional) The order in which Cisco UCS uses this provider to authenticate users.
Step 10	UCS-A /security/ldap/server # set port <i>port-num</i>	(Optional) The port through which Cisco UCS communicates with the LDAP database. The standard port number is 389.

	Command or Action	Purpose
Step 11	UCS-A /security/ldap/server # set ssl {yes no}	Enables or disables the use of encryption when communicating with the LDAP server. The options are as follows: <ul style="list-style-type: none"> • yes —Encryption is required. If encryption cannot be negotiated, the connection fails. • no —Encryption is disabled. Authentication information is sent as clear text. <p>LDAP uses STARTTLS. This allows encrypted communication using port 389.</p>
Step 12	UCS-A /security/ldap/server # set timeout timeout-num	The length of time in seconds the system should spend trying to contact the LDAP database before it times out. <p>Enter an integer from 1 to 60 seconds, or enter 0 (zero) to use the global timeout value specified on the LDAP General tab. The default is 30 seconds.</p>
Step 13	UCS-A /security/ldap/server # commit-buffer	Commits the transaction to the system configuration.

The following example creates an LDAP server instance named 10.193.169.246, configures the binddn, password, order, port, and SSL settings, and commits the transaction:

```
UCS-A# scope security
UCS-A /security # scope ldap
UCS-A /security/ldap* # create server 10.193.169.246
UCS-A /security/ldap/server* # set binddn
"cn=Administrator,cn=Users,DC=cisco-ucsm-aaa3,DC=qalab,DC=com"
UCS-A /security/ldap/server* # set password
Enter the password:
Confirm the password:
UCS-A /security/ldap/server* # set order 2
UCS-A /security/ldap/server* # set port 389
UCS-A /security/ldap/server* # set ssl yes
UCS-A /security/ldap/server* # set timeout 30
UCS-A /security/ldap/server* # commit-buffer
UCS-A /security/ldap/server #
```

What to Do Next

For implementations involving a single LDAP database, select LDAP as the authentication service.

For implementations involving multiple LDAP databases, configure an LDAP provider group.

Changing the LDAP Group Rule for an LDAP Provider

Procedure

	Command or Action	Purpose
Step 1	UCS-A# scope security	Enters security mode.
Step 2	UCS-A /security # scope ldap	Enters security LDAP mode.
Step 3	UCS-A /security/ldap # scope server ldap-provider	Enters security LDAP provider mode.
Step 4	UCS-A /security/ldap/server # scope ldap-group-rule	Enters LDAP group rule mode.
Step 5	UCS-A /security/ldap/server/ldap-group-rule # set authorization {enable disable}	<p>Specifies whether Cisco UCS searches LDAP groups when assigning user roles and locales to a remote user.</p> <ul style="list-style-type: none"> • disable—Cisco UCS does not access any LDAP groups. • enable—Cisco UCS searches the LDAP provider groups mapped in this Cisco UCS domain. If the remote user is found, Cisco UCS assigns the user roles and locales defined for that LDAP group in the associated LDAP group map. <p>Note Role and locale assignment is cumulative. If a user is included in multiple groups, or has a role or locale specified in the LDAP attribute, Cisco UCS assigns that user all the roles and locales mapped to any of those groups or attributes.</p>
Step 6	UCS-A /security/ldap/server/ldap-group-rule # set member-of-attribute attr-name	<p>The attribute Cisco UCS uses to determine group membership in the LDAP database.</p> <p>The supported string length is 63 characters. The default string is memberOf.</p>
Step 7	UCS-A /security/ldap/server/ldap-group-rule # set traversal {non-recursive recursive}	<p>Specifies whether Cisco UCS takes the settings for a group member's parent group, if necessary. This can be:</p> <ul style="list-style-type: none"> • non-recursive—Cisco UCS only searches those groups that the user belongs to. • recursive—Cisco UCS searches all the ancestor groups belonging to the user.
Step 8	UCS-A /security/ldap/server/ldap-group-rule # commit-buffer	Commits the transaction to the system configuration.

The following example sets the LDAP group rule to enable authorization, sets the member of attribute to memberOf, sets the traversal to non-recursive, and commits the transaction:

```
UCS-A# scope security
UCS-A /security # scope ldap
UCS-A /security/ldap # scope server ldaprovider
UCS-A /security/ldap/server # scope ldap-group-rule
UCS-A /security/ldap/server/ldap-group-rule # set authorization enable
UCS-A /security/ldap/server/ldap-group-rule* # set member-of-attribute memberOf
UCS-A /security/ldap/server/ldap-group-rule* # set traversal non-recursive
UCS-A /security/ldap/server/ldap-group-rule* # commit-buffer
UCS-A /security/ldap/server/ldap-group-rule #
```

Deleting an LDAP Provider

Procedure

	Command or Action	Purpose
Step 1	UCS-A# scope security	Enters security mode
Step 2	UCS-A /security # scope ldap	Enters security LDAP mode
Step 3	UCS-A /security/ldap # delete server <i>serv-name</i>	Deletes the specified server.
Step 4	UCS-A /security/ldap # commit-buffer	Commits the transaction to the system configuration.

The following example deletes the LDAP server called ldap1 and commits the transaction:

```
UCS-A# scope security
UCS-A /security # scope ldap
UCS-A /security/ldap # delete server ldap1
UCS-A /security/ldap* # commit-buffer
UCS-A /security/ldap #
```

LDAP Group Mapping

For organizations that already use LDAP groups to restrict access to LDAP databases, group membership information can be used by UCSM to assign a role or locale to an LDAP user during login. This eliminates the need to define role or locale information in the LDAP user object when Cisco UCS Manager is deployed.

When a user logs in to Cisco UCS Manager, information about the user's role and locale are pulled from the LDAP group map. If the role and locale criteria match the information in the policy, access is granted.

Role and locale definitions are configured locally in Cisco UCS Manager and do not update automatically based on changes to an LDAP directory. When deleting or renaming LDAP groups in an LDAP directory, it is important that you update Cisco UCS Manager with the change.

An LDAP group map can be configured to include any of the following combinations of roles and locales:

- Roles only
- Locales only
- Both roles and locales

For example, consider an LDAP group representing a group of server administrators at a specific location. The LDAP group map might be configured to include user roles like server-profile and server-equipment. To restrict access to server administrators at a specific location, the locale could be set to a particular site name.

**Note**

Cisco UCS Manager includes many out-of-the-box user roles but does not include any locales. Mapping an LDAP provider group to a locale requires that you create a custom locale.

Creating an LDAP Group Map

Before You Begin

- Create an LDAP group in the LDAP server.
- Configure the distinguished name for the LDAP group in the LDAP server.
- Create locales in Cisco UCS Manager (optional).
- Create custom roles in Cisco UCS Manager (optional).

Procedure

	Command or Action	Purpose
Step 1	UCS-A# scope security	Enters security mode.
Step 2	UCS-A /security # scope ldap	Enters security LDAP mode.
Step 3	UCS-A /security/ldap # create ldap-group group-dn	Creates an LDAP group map for the specified DN.
Step 4	UCS-A /security/ldap/ldap-group # create locale locale-name	Maps the LDAP group to the specified locale.
Step 5	UCS-A /security/ldap/ldap-group # create role role-name	Maps the LDAP group to the specified role.
Step 6	UCS-A /security/ldap/ldap-group # commit-buffer	Commits the transaction to the system configuration.

The following example maps the LDAP group mapped to a DN, sets the locale to pacific, sets the role to admin, and commits the transaction:

```
UCS-A# scope security
UCS-A /security # scope ldap
UCS-A /security/ldap # create ldap-group cn=security,cn=users,dc=lab,dc=com
UCS-A /security/ldap/ldap-group* # create locale pacific
UCS-A /security/ldap/ldap-group* # create role admin
UCS-A /security/ldap/ldap-group* # commit-buffer
UCS-A /security/ldap/ldap-group #
```

What to Do Next

Set the LDAP group rule.

Deleting an LDAP Group Map

Procedure

	Command or Action	Purpose
Step 1	UCS-A# scope security	Enters security mode.
Step 2	UCS-A /security # scope ldap	Enters security LDAP mode.
Step 3	UCS-A /security/ldap # delete ldap-group <i>group-dn</i>	Deletes the LDAP group map for the specified DN.
Step 4	UCS-A /security/ldap # commit-buffer	Commits the transaction to the system configuration.

The following example deletes an LDAP group map and commits the transaction:

```
UCS-A# scope security
UCS-A /security # scope ldap
UCS-A /security/ldap # delete ldap-group cn=security,cn=users,dc=lab,dc=com
UCS-A /security/ldap* # commit-buffer
UCS-A /security/ldap #
```

Configuring RADIUS Providers

Configuring Properties for RADIUS Providers

The properties that you configure in this task are the default settings for all provider connections of this type defined in Cisco UCS Manager. If an individual provider includes a setting for any of these properties, Cisco UCS uses that setting and ignores the default setting.

Procedure

	Command or Action	Purpose
Step 1	UCS-A# scope security	Enters security mode.
Step 2	UCS-A /security # scope radius	Enters security RADIUS mode.
Step 3	UCS-A /security/radius # set retries <i>retry-num</i>	(Optional) Sets the number of times to retry communicating with the RADIUS server before noting the server as down.
Step 4	UCS-A /security/radius # set timeout <i>seconds</i>	(Optional) Sets the time interval that the system waits for a response from the RADIUS server before noting the server as down.

	Command or Action	Purpose
Step 5	UCS-A /security/radius # commit-buffer	Commits the transaction to the system configuration.

The following example sets the RADIUS retries to 4, sets the timeout interval to 30 seconds, and commits the transaction:

```
UCS-A# scope security
UCS-A /security # scope radius
UCS-A /security/radius # set retries 4
UCS-A /security/radius* # set timeout 30
UCS-A /security/radius* # commit-buffer
UCS-A /security/radius #
```

What to Do Next

Create a RADIUS provider.

Creating a RADIUS Provider

Cisco UCS Manager supports a maximum of 16 RADIUS providers.

Before You Begin

Perform the following configuration in the RADIUS server:

- Configure users with the attribute that holds the user role and locale information for Cisco UCS Manager. You can choose whether to extend the RADIUS schema for this attribute. If you do not want to extend the schema, use an existing RADIUS attribute to hold the Cisco UCS user roles and locales. If you prefer to extend the schema, create a custom attribute, such as the `cisco-avpair` attribute.

The vendor ID for the Cisco RADIUS implementation is 009 and the vendor ID for the attribute is 001.

The following syntax example shows how to specify multiples user roles and locales if you choose to create the `cisco-avpair` attribute: `shell:roles="admin,aaa" shell:locales="L1,abc"`. Use a comma "," as the delimiter to separate multiple values.

- For a cluster configuration, add the management port IP addresses for both fabric interconnects. This configuration ensures that remote users can continue to log in if the first fabric interconnect fails and the system fails over to the second fabric interconnect. All login requests are sourced from these IP addresses, not the virtual IP address used by Cisco UCS Manager.

Procedure

	Command or Action	Purpose
Step 1	UCS-A# scope security	Enters security mode.
Step 2	UCS-A /security # scope radius	Enters security RADIUS mode.
Step 3	UCS-A /security/radius # create server <i>server-name</i>	Creates a RADIUS server instance and enters security RADIUS server mode

	Command or Action	Purpose
Step 4	UCS-A /security/radius/server # set authport <i>authport-num</i>	(Optional) Specifies the port used to communicate with the RADIUS server.
Step 5	UCS-A /security/radius/server # set key	Sets the RADIUS server key. To set the key value, press Enter after typing the set key command and enter the key value at the prompt.
Step 6	UCS-A /security/radius/server # set order <i>order-num</i>	(Optional) Specifies when in the order this server will be tried.
Step 7	UCS-A /security/radius # set retries <i>retry-num</i>	(Optional) Sets the number of times to retry communicating with the RADIUS server before noting the server as down.
Step 8	UCS-A /security/radius # set timeout <i>seconds</i>	(Optional) Sets the time interval that the system waits for a response from the RADIUS server before noting the server as down.
Step 9	UCS-A /security/radius/server # commit-buffer	Commits the transaction to the system configuration.

The following example creates a server instance named `radiuserv7`, sets the authentication port to 5858, sets the key to `radiuskey321`, sets the order to 2, sets the retries to 4, sets the timeout to 30, and commits the transaction:

```
UCS-A# scope security
UCS-A /security # scope radius
UCS-A /security/radius # create server radiusserv7
UCS-A /security/radius/server* # set authport 5858
UCS-A /security/radius/server* # set key
Enter the key: radiuskey321
Confirm the key: radiuskey321
UCS-A /security/radius/server* # set order 2
UCS-A /security/radius/server* # set retries 4
UCS-A /security/radius/server* # set timeout 30
UCS-A /security/radius/server* # commit-buffer
UCS-A /security/radius/server #
```

What to Do Next

For implementations involving a single RADIUS database, select RADIUS as the primary authentication service.

For implementations involving multiple RADIUS databases, configure a RADIUS provider group.

Deleting a RADIUS Provider

Procedure

	Command or Action	Purpose
Step 1	UCS-A# scope security	Enters security mode.
Step 2	UCS-A /security # scope RADIUS	Enters security RADIUS mode.
Step 3	UCS-A /security/radius # delete server <i>serv-name</i>	Deletes the specified server.
Step 4	UCS-A /security/radius # commit-buffer	Commits the transaction to the system configuration.

The following example deletes the RADIUS server called radius1 and commits the transaction:

```
UCS-A# scope security
UCS-A /security # scope radius
UCS-A /security/radius # delete server radius1
UCS-A /security/radius* # commit-buffer
UCS-A /security/radius #
```

Configuring TACACS+ Providers

Configuring Properties for TACACS+ Providers

The properties that you configure in this task are the default settings for all provider connections of this type defined in Cisco UCS Manager. If an individual provider includes a setting for any of these properties, Cisco UCS uses that setting and ignores the default setting.

Procedure

	Command or Action	Purpose
Step 1	UCS-A# scope security	Enters security mode.
Step 2	UCS-A /security # scope tacacs	Enters security TACACS+ mode.
Step 3	UCS-A /security/tacacs # set timeout <i>seconds</i>	(Optional) Sets the time interval that the system waits for a response from the TACACS+ server before noting the server as down.
Step 4	UCS-A /security/tacacs # commit-buffer	Commits the transaction to the system configuration.

The following example sets the TACACS+ timeout interval to 45 seconds and commits the transaction:

```
UCS-A# scope security
UCS-A /security # scope tacacs
UCS-A /security/tacacs # set timeout 45
UCS-A /security/tacacs* # commit-buffer
UCS-A /security/tacacs #
```

What to Do Next

Create a TACACS+ provider.

Creating a TACACS+ Provider

Cisco UCS Manager supports a maximum of 16 TACACS+ providers.

Before You Begin

Perform the following configuration in the TACACS+ server:

- Create the cisco-av-pair attribute. You cannot use an existing TACACS+ attribute.

The cisco-av-pair name is the string that provides the attribute ID for the TACACS+ provider.

The following syntax example shows how to specify multiples user roles and locales when you create the cisco-av-pair attribute: `cisco-av-pair=shell:roles="admin aaa" shell:locales*"L1 abc"`.

Using an asterisk (*) in the cisco-av-pair attribute syntax flags the locale as optional, preventing authentication failures for other Cisco devices that use the same authorization profile. Use a space as the delimiter to separate multiple values.

- For a cluster configuration, add the management port IP addresses for both fabric interconnects. This configuration ensures that remote users can continue to log in if the first fabric interconnect fails and the system fails over to the second fabric interconnect. All login requests are sourced from these IP addresses, not the virtual IP address used by Cisco UCS Manager.

Procedure

	Command or Action	Purpose
Step 1	UCS-A# scope security	Enters security mode.
Step 2	UCS-A /security # scope tacacs	Enters security TACACS+ mode.
Step 3	UCS-A /security/tacacs # create server <i>server-name</i>	Creates an TACACS+ server instance and enters security TACACS+ server mode
Step 4	UCS-A /security/tacacs/server # set key	(Optional) Sets the TACACS+ server key. To set the key value, press Enter after typing the set key command and enter the key value at the prompt.
Step 5	UCS-A /security/tacacs/server # set order <i>order-num</i>	(Optional) Specifies when in the order this server will be tried.
Step 6	UCS-A /security/tacacs/server # set port <i>port-num</i>	Specifies the port used to communicate with the TACACS+ server.

	Command or Action	Purpose
Step 7	UCS-A /security/tacacs/server # commit-buffer	Commits the transaction to the system configuration.

The following example creates a server instance named tacacsserv680, sets the key to tacacskey321, sets the order to 4, sets the authentication port to 5859, and commits the transaction:

```
UCS-A# scope security
UCS-A /security # scope tacacs
UCS-A /security/tacacs # create server tacacsserv680
UCS-A /security/tacacs/server* # set key
Enter the key: tacacskey321
Confirm the key: tacacskey321
UCS-A /security/tacacs/server* # set order 4
UCS-A /security/tacacs/server* # set port 5859
UCS-A /security/tacacs/server* # commit-buffer
UCS-A /security/tacacs/server #
```

What to Do Next

For implementations involving a single TACACS+ database, select TACACS+ as the primary authentication service.

For implementations involving multiple TACACS+ databases, configure a TACACS+ provider group.

Deleting a TACACS+ Provider

Procedure

	Command or Action	Purpose
Step 1	UCS-A# scope security	Enters security mode.
Step 2	UCS-A /security # scope tacacs	Enters security TACACS mode.
Step 3	UCS-A /security/tacacs # delete server <i>serv-name</i>	Deletes the specified server.
Step 4	UCS-A /security/tacacs # commit-buffer	Commits the transaction to the system configuration.

The following example deletes the TACACS server called tacacs1 and commits the transaction:

```
UCS-A# scope security
UCS-A /security # scope tacacs
UCS-A /security/tacacs # delete server TACACS1
UCS-A /security/tacacs* # commit-buffer
UCS-A /security/tacacs #
```

Configuring Multiple Authentication Systems

Multiple Authentication Systems

You can configure Cisco UCS to use multiple authentication systems by configuring the following features:

- Provider groups
- Authentication domains

Once provider groups and authentication domains have been configured in Cisco UCS Manager, the following syntax can be used to log in to the system using Cisco UCS Manager CLI: **ucs: auth-domain \user-name**

When multiple authentication domains and native authentication are configured with a remote authentication service, use one of the following syntax examples to log in with SSH or Putty:

From a Linux terminal:

- **ssh ucs-auth-domain\username@UCSM-ip-address**

```
ssh ucs-example\jsmith@192.0.20.11
```
- **ssh -l ucs-auth-domain\username {UCSM-ip-address | UCSM-host-name}**

```
ssh -l ucs-example\jsmith 192.0.20.11
```
- **ssh {UCSM-ip-address | UCSM-host-name} -l ucs-auth-domain\username**

```
ssh 192.0.20.11 -l ucs-example\jsmith
```

From a Putty client:

- Login as: **ucs-auth-domain\username**

```
Login as: ucs-example\jsmith
```

Provider Groups

A provider group is a set of providers that will be used by Cisco UCS during the authentication process. Cisco UCS Manager allows you to create a maximum of 16 provider groups, with a maximum of eight providers allowed per group.

During authentication, all the providers within a provider group are tried in order. If all of the configured servers are unavailable or unreachable, Cisco UCS Manager automatically falls back to the local authentication method using the local username and password.

Creating an LDAP Provider Group

Creating an LDAP provider group allows you to authenticate using multiple LDAP databases.



Note

Authenticating with a single LDAP database does not require you to set up an LDAP provider group.

Before You Begin

Create one or more LDAP providers.

Procedure

	Command or Action	Purpose
Step 1	UCS-A# scope security	Enters security mode.
Step 2	UCS-A /security # scope ldap	Enters security LDAP mode.
Step 3	UCS-A /security/ldap # create auth-server-group <i>auth-server-group-name</i>	Creates an LDAP provider group and enters authentication server group security LDAP mode.
Step 4	UCS-A /security/ldap/auth-server-group # create server-ref <i>ldap-provider-name</i>	Adds the specified LDAP provider to the LDAP provider group and enters server reference authentication server group security LDAP mode.
Step 5	UCS-A /security/ldap/auth-server-group/server-ref # set order <i>order-num</i>	Specifies the order in which Cisco UCS uses this provider to authenticate users. Valid values include no-value and 0-16, with the lowest value indicating the highest priority. Setting the order to no-value is equivalent to giving that server reference the highest priority.
Step 6	UCS-A /security/ldap/auth-server-group/server-ref # commit-buffer	Commits the transaction to the system configuration.

The following example creates an LDAP provider group called ldapgroup, adds two previously configured providers called ldap1 and ldap2 to the provider group, sets the order, and commits the transaction:

```
UCS-A# scope security
UCS-A /security # scope ldap
UCS-A /security/ldap # create auth-server-group ldapgroup
UCS-A /security/ldap/auth-server-group* # create server-ref ldap1
UCS-A /security/ldap/auth-server-group/server-ref* # set order 1
UCS-A /security/ldap/auth-server-group/server-ref* # up
UCS-A /security/ldap/auth-server-group* # create server-ref ldap2
UCS-A /security/ldap/auth-server-group/server-ref* # set order 2
UCS-A /security/ldap/auth-server-group/server-ref* # commit-buffer
UCS-A /security/ldap/auth-server-group/server-ref #
```

What to Do Next

Configure an authentication domain or select a default authentication service.

Deleting an LDAP Provider Group**Before You Begin**

Remove the provider group from an authentication configuration.

Procedure

	Command or Action	Purpose
Step 1	UCS-A# scope security	Enters security mode.
Step 2	UCS-A /security # scope ldap	Enters security LDAP mode.
Step 3	UCS-A /security/ldap # delete auth-server-group auth-server-group-name	Deletes the LDAP provider group.
Step 4	UCS-A /security/ldap # commit-buffer	Commits the transaction to the system configuration.

The following example deletes an LDAP provider group called ldapgroup and commits the transaction:

```
UCS-A# scope security
UCS-A /security # scope ldap
UCS-A /security/ldap # delete auth-server-group ldapgroup
UCS-A /security/ldap* # commit-buffer
UCS-A /security/ldap #
```

Creating a RADIUS Provider Group

Creating a RADIUS provider group allows you to authenticate using multiple RADIUS databases.

**Note**

Authenticating with a single RADIUS database does not require you to set up a RADIUS provider group.

Before You Begin

Create one or more RADIUS providers.

Procedure

	Command or Action	Purpose
Step 1	UCS-A# scope security	Enters security mode.
Step 2	UCS-A /security # scope radius	Enters security RADIUS mode.
Step 3	UCS-A /security/radius # create auth-server-group auth-server-group-name	Creates a RADIUS provider group and enters authentication server group security RADIUS mode.
Step 4	UCS-A /security/RADIUS/auth-server-group # create server-ref radius-provider-name	Adds the specified RADIUS provider to the RADIUS provider group and enters server reference authentication server group security RADIUS mode.

	Command or Action	Purpose
Step 5	UCS-A /security/radius/auth-server-group/server-ref # set order <i>order-num</i>	Specifies the order in which Cisco UCS uses this provider to authenticate users. Valid values include no-value and 0-16, with the lowest value indicating the highest priority. Setting the order to no-value is equivalent to giving that server reference the highest priority.
Step 6	UCS-A /security/radius/auth-server-group/server-ref # commit-buffer	Commits the transaction to the system configuration.

The following example creates a RADIUS provider group called radiusgroup, adds two previously configured providers called radius1 and radius2 to the provider group, sets the order, and commits the transaction:

```
UCS-A# scope security
UCS-A /security # scope radius
UCS-A /security/radius # create auth-server-group radiusgroup
UCS-A /security/radius/auth-server-group* # create server-ref radius1
UCS-A /security/radius/auth-server-group/server-ref* # set order 1
UCS-A /security/radius/auth-server-group/server-ref* # up
UCS-A /security/radius/auth-server-group* # create server-ref radius2
UCS-A /security/radius/auth-server-group/server-ref* # set order 2
UCS-A /security/radius/auth-server-group/server-ref* # commit-buffer
UCS-A /security/radius/auth-server-group/server-ref #
```

What to Do Next

Configure an authentication domain or select a default authentication service.

Deleting a RADIUS Provider Group

Remove the provider group from an authentication configuration.

Procedure

	Command or Action	Purpose
Step 1	UCS-A# scope security	Enters security mode.
Step 2	UCS-A /security # scope radius	Enters security RADIUS mode.
Step 3	UCS-A /security/radius # delete auth-server-group <i>auth-server-group-name</i>	Deletes the RADIUS provider group.
Step 4	UCS-A /security/radius # commit-buffer	Commits the transaction to the system configuration.

The following example deletes a RADIUS provider group called radiusgroup and commits the transaction:

```
UCS-A# scope security
UCS-A /security # scope radius
```

```
UCS-A /security/radius # delete auth-server-group radiusgroup
UCS-A /security/radius* # commit-buffer
UCS-A /security/radius #
```

Creating a TACACS Provider Group

Creating a TACACS+ provider group allows you to authenticate using multiple TACACS+ databases.



Note

Authenticating with a single TACACS+ database does not require you to set up a TACACS+ provider group.

Before You Begin

Create a TACACS provider.

Procedure

	Command or Action	Purpose
Step 1	UCS-A# scope security	Enters security mode.
Step 2	UCS-A /security # scope tacacs	Enters security TACACS mode.
Step 3	UCS-A /security/tacacs # create auth-server-group <i>auth-server-group-name</i>	Creates a TACACS provider group and enters authentication server group security TACACS mode.
Step 4	UCS-A /security/tacacs/auth-server-group # create server-ref <i>tacacs-provider-name</i>	Adds the specified TACACS provider to the TACACS provider group and enters server reference authentication server group security TACACS mode.
Step 5	UCS-A /security/tacacs/auth-server-group/server-ref # set order <i>order-num</i>	Specifies the order in which Cisco UCS uses this provider to authenticate users. Valid values include no-value and 0-16, with the lowest value indicating the highest priority. Setting the order to no-value is equivalent to giving that server reference the highest priority.
Step 6	UCS-A /security/tacacs/auth-server-group/server-ref # commit-buffer	Commits the transaction to the system configuration.

The following example creates a TACACS provider group called tacacsgroup, adds two previously configured providers called tacacs1 and tacacs2 to the provider group, sets the order, and commits the transaction:

```
UCS-A# scope security
UCS-A /security # scope tacacs
UCS-A /security/tacacs # create auth-server-group tacacsgroup
UCS-A /security/tacacs/auth-server-group* # create server-ref tacacs1
UCS-A /security/tacacs/auth-server-group/server-ref* # set order 1
UCS-A /security/tacacs/auth-server-group/server-ref* # up
```

```
UCS-A /security/tacacs/auth-server-group* # create server-ref tacacs2
UCS-A /security/tacacs/auth-server-group/server-ref* # set order 2
UCS-A /security/tacacs/auth-server-group/server-ref* # commit-buffer
UCS-A /security/tacacs/auth-server-group/server-ref #
```

What to Do Next

Configure an authentication domain or select a default authentication service.

Deleting a TACACS Provider Group

Remove the provider group from an authentication configuration.

Procedure

	Command or Action	Purpose
Step 1	UCS-A# scope security	Enters security mode.
Step 2	UCS-A /security # scope tacacs	Enters security TACACS mode.
Step 3	UCS-A /security/tacacs # delete auth-server-group <i>auth-server-group-name</i>	Deletes the TACACS provider group.
Step 4	UCS-A /security/tacacs # commit-buffer	Commits the transaction to the system configuration.

The following example deletes a TACACS provider group called tacacsgroup and commits the transaction:

```
UCS-A# scope security
UCS-A /security # scope tacacs
UCS-A /security/tacacs # delete auth-server-group tacacsgroup
UCS-A /security/tacacs* # commit-buffer
UCS-A /security/tacacs #
```

Authentication Domains

Authentication domains are used by Cisco UCS Manager to leverage multiple authentication systems. Each authentication domain is specified and configured during login. If no authentication domain is specified, the default authentication service configuration is used.

You can create up to eight authentication domains. Each authentication domain is associated with a provider group and realm in Cisco UCS Manager. If no provider group is specified, all servers within the realm are used.

Creating an Authentication Domain

Procedure

	Command or Action	Purpose
Step 1	UCS-A# scope security	Enters security mode.
Step 2	UCS-A /security # create auth-domain <i>domain-name</i>	Creates an authentication domain and enters authentication domain mode. Note For systems using RADIUS as their preferred authentication protocol, the authentication domain name is considered part of the user name and counts toward the 32 character limit for locally created user names. Because Cisco UCS inserts 5 characters for formatting, authentication will fail if the combined total of the domain name plus the user name is more than 27 characters.
Step 3	UCS-A /security/auth-domain # set refresh-period <i>seconds</i>	(Optional) When a web client connects to Cisco UCS Manager, the client needs to send refresh requests to Cisco UCS Manager to keep the web session active. This option specifies the maximum amount of time allowed between refresh requests for a user in this domain. If this time limit is exceeded, Cisco UCS Manager considers the web session to be inactive, but it does not terminate the session. Specify an integer between 60 and 172800. The default is 600 seconds.
Step 4	UCS-A /security/auth-domain # set session-timeout <i>seconds</i>	(Optional) The maximum amount of time that can elapse after the last refresh request before Cisco UCS Manager considers a web session to have ended. If this time limit is exceeded, Cisco UCS Manager automatically terminates the web session. Specify an integer between 60 and 172800. The default is 7200 seconds.
Step 5	UCS-A /security/auth-domain # create default-auth	(Optional) Creates a default authentication for the specified authentication domain.
Step 6	UCS-A /security/auth-domain/default-auth # set auth-server-group <i>auth-serv-group-name</i>	(Optional) Specifies the provider group for the specified authentication domain.
Step 7	UCS-A /security/auth-domain/default-auth	Specifies the realm for the specified authentication domain.

	Command or Action	Purpose
	# set realm {ldap local radius tacacs}	
Step 8	UCS-A /security/auth-domain/default-auth # commit-buffer	Commits the transaction to the system configuration.

The following example creates an authentication domain called domain1 with a web refresh period of 3600 seconds (1 hour) and a session timeout period of 14400 seconds (4 hours). It then configures domain1 to use the providers in ldapgroup1, sets the realm type to ldap, and commits the transaction.

```
UCS-A# scope security
UCS-A /security # create auth-domain domain1
UCS-A /security/auth-domain* # set refresh-period 3600
UCS-A /security/auth-domain* # set session-timeout 14400
UCS-A /security/auth-domain* # create default-auth
UCS-A /security/auth-domain/auth-domain* # set auth-server-group ldapgroup1
UCS-A /security/auth-domain/auth-domain* # set realm ldap
UCS-A /security/auth-domain/auth-domain* # commit-buffer
UCS-A /security/auth-domain/auth-domain #
```

Selecting a Primary Authentication Service

Selecting the Console Authentication Service

Before You Begin

If the system uses a remote authentication service, create a provider for that authentication service. If the system uses only local authentication through Cisco UCS, you do not need to create a provider first.

Procedure

	Command or Action	Purpose
Step 1	UCS-A# scope security	Enters security mode.
Step 2	UCS-A /security # scope console-auth	Enters console authorization security mode.
Step 3	UCS-A /security/console-auth # set realm <i>auth-type</i>	Specifies the console authentication, where the <i>auth-type</i> argument is one of the following keywords: <ul style="list-style-type: none"> • ldap —Specifies LDAP authentication • local —Specifies local authentication • none —Allows local users to log on without specifying a password • radius —Specifies RADIUS authentication • tacacs —Specifies TACACS+ authentication

	Command or Action	Purpose
Step 4	UCS-A /security/console-auth # set auth-server-group <i>auth-serv-group-name</i>	(Optional) The associated provider group, if any.
Step 5	UCS-A /security/console-auth # commit-buffer	Commits the transaction to the system configuration.

The following example sets the authentication to LDAP, sets the console authentication provider group to provider1, and commits the transaction:

```
UCS-A# scope security
UCS-A /security # scope console-auth
UCS-A /security/console-auth # set realm local
UCS-A /security/console-auth # set auth-server-group provider1
UCS-A /security/console-auth* # commit-buffer
UCS-A /security/console-auth #
```

Selecting the Default Authentication Service

Procedure

	Command or Action	Purpose
Step 1	UCS-A# scope security	Enters security mode.
Step 2	UCS-A /security # scope default-auth	Enters default authorization security mode.
Step 3	UCS-A /security/default-auth # set realm <i>auth-type</i>	Specifies the default authentication, where <i>auth-type</i> is one of the following keywords: <ul style="list-style-type: none"> • ldap—Specifies LDAP authentication • local—Specifies local authentication • none—Allows local users to log on without specifying a password • radius—Specifies RADIUS authentication • tacacs—Specifies TACACS+ authentication
Step 4	UCS-A /security/default-auth # set auth-server-group <i>auth-serv-group-name</i>	(Optional) The associated provider group, if any.
Step 5	UCS-A /security/default-auth # set refresh-period <i>seconds</i>	(Optional) When a web client connects to Cisco UCS Manager, the client needs to send refresh requests to Cisco UCS Manager to keep the web session active. This option specifies the maximum amount of time allowed between refresh requests for a user in this domain.

	Command or Action	Purpose
		If this time limit is exceeded, Cisco UCS Manager considers the web session to be inactive, but it does not terminate the session. Specify an integer between 60 and 172800. The default is 600 seconds.
Step 6	UCS-A /security/default-auth # set session-timeout <i>seconds</i>	(Optional) The maximum amount of time that can elapse after the last refresh request before Cisco UCS Manager considers a web session to have ended. If this time limit is exceeded, Cisco UCS Manager automatically terminates the web session. Specify an integer between 60 and 172800. The default is 7200 seconds.
Step 7	UCS-A /security/default-auth # commit-buffer	Commits the transaction to the system configuration.

The following example sets the default authentication to LDAP, the default authentication provider group to provider1, the refresh period to 7200 seconds (2 hours), and the session timeout period to 28800 seconds (8 hours). It then commits the transaction.

```

UCS-A# scope security
UCS-A /security # scope default-auth
UCS-A /security/default-auth # set realm ldap
UCS-A /security/default-auth* # set auth-server-group provider1
UCS-A /security/default-auth* # set refresh-period 7200
UCS-A /security/default-auth* # set session-timeout 28800
UCS-A /security/default-auth* # commit-buffer
UCS-A /security/default-auth #
    
```

Role Policy for Remote Users

By default, if user roles are not configured in Cisco UCS Manager read-only access is granted to all users logging in to Cisco UCS Manager from a remote server using the LDAP, RADIUS, or TACACS protocols. For security reasons, it might be desirable to restrict access to those users matching an established user role in Cisco UCS Manager.

You can configure the role policy for remote users in the following ways:

assign-default-role

Does not restrict user access to Cisco UCS Manager based on user roles. Read-only access is granted to all users unless other user roles have been defined in Cisco UCS Manager.

This is the default behavior.

no-login

Restricts user access to Cisco UCS Manager based on user roles. If user roles have not been assigned for the remote authentication system, access is denied.

Configuring the Role Policy for Remote Users

Procedure

	Command or Action	Purpose
Step 1	UCS-A# scope security	Enters security mode.
Step 2	UCS-A /security # set remote-user default-role {assign-default-role no-login}	Specifies whether user access to Cisco UCS Manager is restricted based on user roles.
Step 3	UCS-A /security # commit-buffer	Commits the transaction to the system configuration.

The following example sets the role policy for remote users and commits the transaction:

```
UCS-A# scope security
UCS-A /security # set remote-user default-role assign-default-role
UCS-A /security* # commit-buffer
UCS-A /security #
```




CHAPTER 9

Configuring Organizations

This chapter includes the following sections:

- [Organizations in a Multi-Tenancy Environment, page 153](#)
- [Hierarchical Name Resolution in a Multi-Tenancy Environment, page 154](#)
- [Configuring an Organization Under the Root Organization, page 156](#)
- [Configuring an Organization Under an Organization that is not Root, page 156](#)
- [Deleting an Organization, page 157](#)

Organizations in a Multi-Tenancy Environment

Multi-tenancy allows you to divide up the large physical infrastructure of an Cisco UCS domain into logical entities known as organizations. As a result, you can achieve a logical isolation between organizations without providing a dedicated physical infrastructure for each organization.

You can assign unique resources to each tenant through the related organization, in the multi-tenant environment. These resources can include different policies, pools, and quality of service definitions. You can also implement locales to assign or restrict user privileges and roles by organization, if you do not want all users to have access to all organizations.

If you set up a multi-tenant environment, all organizations are hierarchical. The top-level organization is always root. The policies and pools that you create in root are system-wide and are available to all organizations in the system. However, any policies and pools created in other organizations are only available to organizations that are above it in the same hierarchy. For example, if a system has organizations named Finance and HR that are not in the same hierarchy, Finance cannot use any policies in the HR organization, and HR cannot access any policies in the Finance organization. However, both Finance and HR can use policies and pools in the root organization.

If you create organizations in a multi-tenant environment, you can also set up one or more of the following for each organization or for a sub-organization in the same hierarchy:

- Resource pools
- Policies
- Service profiles

- Service profile templates

The root organization is always the top level organization.

Hierarchical Name Resolution in a Multi-Tenancy Environment

In a multi-tenant environment, Cisco UCS uses the hierarchy of an organization to resolve the names of policies and resource pools. When Cisco UCS Manager searches for details of a policy or a resource assigned to a pool, the following occurs:

- 1 Cisco UCS Manager checks for policies and pools with the specified name within the organization assigned to the service profile or policy.
- 2 If a policy is found or an available resource is inside a pool, Cisco UCS Manager uses that policy or resource. If the pool does not have any available resources at the local level, Cisco UCS Manager moves up in the hierarchy to the parent organization and searches for a pool with the same name. Cisco UCS Manager repeats this step until the search reaches the root organization.
- 3 If the search reaches the root organization and has not found an available resource or policy, Cisco UCS Manager returns to the local organization and begins to search for a default policy or available resource in the default pool.
- 4 If an applicable default policy or available resource in a default pool is found, Cisco UCS Manager uses that policy or resource. If the pool does not have any available resources, Cisco UCS Manager moves up in the hierarchy to the parent organization and searches for a default pool. Cisco UCS Manager repeats this step until the search reaches the root organization.
- 5 If Cisco UCS Manager cannot find an applicable policy or available resource in the hierarchy, it returns an allocation error.

Example: Server Pool Name Resolution in a Single-Level Hierarchy

In this example, all organizations are at the same level below the root organization. For example, a service provider creates separate organizations for each customer. In this configuration, organizations only have access to the policies and resource pools assigned to that organization and to the root organization.

In this example, a service profile in the XYZcustomer organization is configured to use servers from the XYZcustomer server pool. When resource pools and policies are assigned to the service profile, the following occurs:

- 1 Cisco UCS Manager checks for an available server in the XYZcustomer server pool.
- 2 If the XYZcustomer server pool has an available server, Cisco UCS Manager associates that server with the service profile and discontinues the search. If the pool does not have an available server, Cisco UCS Manager checks the root organization for a server pool with the same name.
- 3 If the root organization includes an XYZcustomer server pool and that pool has an available server, Cisco UCS Manager associates that server with the service profile and discontinues the search. If the pool does not have an available server, Cisco UCS Manager returns to the XYZcustomer organization to check the default server pool.
- 4 If the default pool in the XYZcustomer organization has an available server, Cisco UCS Manager associates that server with the service profile and discontinues the search. If the default pool does not have an available server, Cisco UCS Manager checks the default server pool in the root organization.

- 5 If the default server pool in the root organization has an available server, Cisco UCS Manager associates that server with the service profile and discontinues the search. If the default pool does not have an available server, Cisco UCS Manager returns an allocation error.

Example: Server Pool Name Resolution in a Multi-Level Hierarchy

In this example, each organization includes at least one suborganization. For example, a company could create organizations for each major division in the company and for subdivisions of those divisions. In this configuration, each organization has access to its local policies and resource pools and to the resource pools in the parent hierarchy.

In this example, the Finance organization includes two sub-organizations, AccountsPayable and AccountsReceivable. A service profile in the AccountsPayable organization is configured to use servers from the AP server pool. When resource pools and policies are assigned to the service profile, the following occurs:

- 1 Cisco UCS Manager checks for an available server in the AP server pool defined in the service profile.
- 2 If the AP server pool has an available server, Cisco UCS Manager associates that server with the service profile and discontinues the search. If the pool does not have an available server, Cisco UCS Manager moves one level up the hierarchy and checks the Finance organization for a pool with the same name.
- 3 If the Finance organization includes a pool with the same name and that pool has an available server, Cisco UCS Manager associates that server with the service profile and discontinues the search. If the pool does not have an available server, Cisco UCS Manager moves one level up in the hierarchy and checks the root organization for a pool with the same name.
- 4 If the root organization includes a pool with the same name and that pool has an available server, Cisco UCS Manager associates that server with the service profile and discontinues the search. If the pool does not have an available server, Cisco UCS Manager returns to the AccountsPayable organization to check the default server pool.
- 5 If the default pool in the AccountsPayable organization has an available server, Cisco UCS Manager associates that server with the service profile and discontinues the search. If the default pool does not have an available server, Cisco UCS Manager moves one level up in the hierarchy and checks the default server pool in the Finance organization.
- 6 If the default pool in the Finance organization has an available server, Cisco UCS Manager associates that server with the service profile and discontinues the search. If the default pool does not have an available server, Cisco UCS Manager moves one level up in the hierarchy and checks the default server pool in the root organization.
- 7 If the default server pool in the root organization has an available server, Cisco UCS Manager associates that server with the service profile and discontinues the search. If the default pool does not have an available server, Cisco UCS Manager returns an allocation error.

Configuring an Organization Under the Root Organization

Procedure

	Command or Action	Purpose
Step 1	UCS-A# scope org /	Enters the root organization mode.
Step 2	UCS-A /org # create org <i>org-name</i>	Creates the specified organization under the root organization and enters organization mode for the specified organization. Note When you move from one organization mode to another, the command prompt does not change.
Step 3	UCS-A /org # commit-buffer	Commits the transaction to the system configuration.

The following example creates an organization named Finance under the root organization and commits the transaction:

```
UCS-A# scope org /
UCS-A /org # create org Finance
UCS-A /org* # commit-buffer
UCS-A /org #
```

Configuring an Organization Under an Organization that is not Root

Procedure

	Command or Action	Purpose
Step 1	UCS-A# scope org /	Enters the root organization mode.
Step 2	UCS-A /org # scope org <i>org-name</i>	Enters organization mode for the specified organization. Note When you move from one organization mode to another, the command prompt does not change.
Step 3	UCS-A /org # create org <i>org-name</i>	Creates the specified organization under the previously configured non-root organization and enters organization mode for the specified organization.
Step 4	UCS-A /org # commit-buffer	Commits the transaction to the system configuration.

The following example creates an organization named Finance under the NorthAmerica organization and commits the transaction:

```
UCS-A# scope org /
UCS-A /org # scope org NorthAmerica
UCS-A /org # create org Finance
```

```
UCS-A /org* # commit-buffer
UCS-A /org #
```

Deleting an Organization

Procedure

	Command or Action	Purpose
Step 1	UCS-A# scope org /	Enters the root organization mode.
Step 2	UCS-A /org # delete org <i>org-name</i>	Deletes the specified organization.
Step 3	UCS-A /org # commit-buffer	Commits the transaction to the system configuration.

The following example deletes the organization under the root organization named Finance and commits the transaction:

```
UCS-A# scope org /
UCS-A /org # delete org Finance
UCS-A /org* # commit-buffer
UCS-A /org #
```




CHAPTER 10

Configuring Role-Based Access Control

This chapter includes the following sections:

- [Role-Based Access Control, page 159](#)
- [User Accounts for Cisco UCS Manager , page 159](#)
- [User Roles, page 162](#)
- [User Locales, page 166](#)
- [Configuring User Roles, page 167](#)
- [Configuring Locales, page 169](#)
- [Configuring Locally Authenticated User Accounts, page 171](#)
- [Password Profile for Locally Authenticated Users, page 179](#)
- [Monitoring User Sessions, page 182](#)

Role-Based Access Control

Role-Based Access Control (RBAC) is a method of restricting or authorizing system access for users based on user roles and locales. A role defines the privileges of a user in the system and the locale defines the organizations (domains) that a user is allowed access. Because users are not directly assigned privileges, management of individual user privileges is simply a matter of assigning the appropriate roles and locales.

A user is granted write access to desired system resources only if the assigned role grants the access privileges and the assigned locale allows access. For example, a user with the Server Administrator role in the Engineering organization could update server configurations in the Engineering organization but could not update server configurations in the Finance organization unless the locales assigned to the user include the Finance organization.

User Accounts for Cisco UCS Manager

User accounts are used to access the system. Up to 48 user accounts can be configured in each Cisco UCS domain. Each user account must have a unique username and password.

A user account can be set with a SSH public key. The public key can be set in either of the two formats: OpenSSH and SECSH.

Admin Account

Each Cisco UCS domain has an admin account. The admin account is a default user account and cannot be modified or deleted. This account is the system administrator or superuser account and has full privileges. There is no default password assigned to the admin account; you must choose the password during the initial system setup.

The admin account is always active and does not expire. You cannot configure the admin account as inactive.

Locally Authenticated User Accounts

A locally authenticated user account is authenticated directly through the fabric interconnect and can be enabled or disabled by anyone with admin or aaa privileges. Once a local user account is disabled, the user cannot log in. Configuration details for disabled local user accounts are not deleted by the database. If you re-enable a disabled local user account, the account becomes active again with the existing configuration, including username and password.

Remotely Authenticated User Accounts

A remotely authenticated user account is any user account that is authenticated through LDAP, RADIUS, or TACACS+.

If a user maintains a local user account and a remote user account simultaneously, the roles defined in the local user account override those maintained in the remote user account.

Expiration of User Accounts

User accounts can be configured to expire at a predefined time. When the expiration time is reached, the user account is disabled.

By default, user accounts do not expire.



Note

After you configure a user account with an expiration date, you cannot reconfigure the account to not expire. You can, however, configure the account with the latest expiration date available.

Guidelines for Cisco UCS Manager Usernames

The username is also used as the login ID for Cisco UCS Manager. When you assign usernames to Cisco UCS Manager user accounts, consider the following guidelines and restrictions:

- The login ID can contain between 1 and 32 characters, including the following:
 - Any alphabetic character
 - Any digit
 - _ (underscore)
 - - (dash)
 - . (dot)

- The unique username for each user account cannot be all-numeric. You cannot create a local user with an all-numeric username.
- The unique username must start with an alphabetic character. It cannot start with a number or a special character, such as an underscore.

After you create a user account, you cannot change the username. You must delete the user account and create a new one.

Reserved Words: Locally Authenticated User Accounts

The following words cannot be used when creating a local user account in Cisco UCS Manager.

- root
- bin
- daemon
- adm
- ip
- sync
- shutdown
- halt
- news
- uucp
- operator
- games
- gopher
- nobody
- nscd
- mailnull
- mail
- rpcuser
- rpc
- mtsuser
- ftpuser
- ftp
- man
- sys
- sandme

- debug

Guidelines for Cisco UCS Manager Passwords

A password is required for each locally authenticated user account. A user with admin or aaa privileges can configure Cisco UCS Manager to perform a password strength check on user passwords. If the password strength check is enabled, each user must have a strong password.

Cisco recommends that each user have a strong password. If you enable the password strength check for locally authenticated users, Cisco UCS Manager rejects any password that does not meet the following requirements:

- Must contain a minimum of 8 characters and a maximum of 64 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 3 times consecutively, such as aaabbb.
- Must not be identical to the username or the reverse of the username.
- Must pass a password dictionary check. For example, the password must not be based on a standard dictionary word.
- Must not contain the following symbols: \$ (dollar sign), ? (question mark), and = (equals sign).
- Should not be blank for local user and admin accounts.

Web Session Limits for User Accounts

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

By default, the number of concurrent web sessions allowed by Cisco UCS Manager is set to 32; although this value can be configured up to the system maximum of 256.

User Roles

User roles contain one or more privileges that define the operations allowed for the user who is assigned the role. A user can be assigned one or more roles. A user assigned multiple roles has the combined privileges of all assigned roles. For example, if Role1 has storage related privileges, and Role2 has server related privileges, users who are assigned to both Role1 and Role2 have storage and server related privileges.

A Cisco UCS domain can contain up to 48 user roles, including the default user roles.

All roles include read access to all configuration settings in the Cisco UCS domain. The difference between the read-only role and other roles is that a user who is only assigned the read-only role cannot modify the system state. A user assigned another role can modify the system state in that user's assigned area or areas.

Roles can be created, modified to add new or remove existing privileges, or deleted. When a role is modified, the new privileges are applied to all users assigned to that role. Privilege assignment is not restricted to the privileges defined for the default roles. That is, you can use a custom set of privileges to create a unique role. For example, the default Server Administrator and Storage Administrator roles have different set of privileges, but a new Server and Storage Administrator role can be created that combines the privileges of both roles.

If a role is deleted after it has been assigned to users, it is also deleted from those user accounts.

User profiles on AAA servers (RADIUS or TACACS+) should be modified to add the roles corresponding to the privileges granted to that user. The attribute is used to store the role information. The AAA servers return this attribute with the request and parse it to get the roles. LDAP servers return the roles in the user profile attributes.

**Note**

If a local user account and a remote user account have the same username, any roles assigned to the remote user are overridden by those assigned to the local user.

Default User Roles

The system contains the following default user roles:

AAA Administrator

Read-and-write access to users, roles, and AAA configuration. Read access to the rest of the system.

Administrator

Complete read-and-write access to the entire system. The default admin account is assigned this role by default and it cannot be changed.

Facility Manager

Read-and-write access to power management operations through the power-mgmt privilege. Read access to the rest of the system.

Network Administrator

Read-and-write access to fabric interconnect infrastructure and network security operations. Read access to the rest of the system.

Operations

Read-and-write access to systems logs, including the syslog servers, and faults. Read access to the rest of the system.

Read-Only

Read-only access to system configuration with no privileges to modify the system state.

Server Equipment Administrator

Read-and-write access to physical server related operations. Read access to the rest of the system.

Server Profile Administrator

Read-and-write access to logical server related operations. Read access to the rest of the system.

Server Security Administrator

Read-and-write access to server security related operations. Read access to the rest of the system.

Storage Administrator

Read-and-write access to storage operations. Read access to the rest of the system.

Reserved Words: User Roles

The following words cannot be used when creating custom roles in Cisco UCS Manager.

- network-admin
- network-operator
- vdc-admin
- vdc-operator
- server-admin

Privileges

Privileges give users assigned to user roles access to specific system resources and permission to perform specific tasks. The following table lists each privilege and the user role given that privilege by default.

Table 11: User Privileges

Privilege	Description	Default Role Assignment
aaa	System security and AAA	AAA Administrator
admin	System administration	Administrator
ext-lan-config	External LAN configuration	Network Administrator
ext-lan-policy	External LAN policy	Network Administrator
ext-lan-qos	External LAN QoS	Network Administrator
ext-lan-security	External LAN security	Network Administrator
ext-san-config	External SAN configuration	Storage Administrator
ext-san-policy	External SAN policy	Storage Administrator
ext-san-qos	External SAN QoS	Storage Administrator

Privilege	Description	Default Role Assignment
ext-san-security	External SAN security	Storage Administrator
fault	Alarms and alarm policies	Operations
operations	Logs and Smart Call Home	Operations
pod-config	Pod configuration	Network Administrator
pod-policy	Pod policy	Network Administrator
pod-qos	Pod QoS	Network Administrator
pod-security	Pod security	Network Administrator
power-mgmt	Read-and-write access to power management operations	Facility Manager
read-only	Read-only access Read-only cannot be selected as a privilege; it is assigned to every user role.	Read-Only
server-equipment	Server hardware management	Server Equipment Administrator
server-maintenance	Server maintenance	Server Equipment Administrator
server-policy	Server policy	Server Equipment Administrator
server-security	Server security	Server Security Administrator
service-profile-config	Service profile configuration	Server Profile Administrator
service-profile-config-policy	Service profile configuration policy	Server Profile Administrator
service-profile-ext-access	Service profile end point access	Server Profile Administrator
service-profile-network	Service profile network	Network Administrator
service-profile-network-policy	Service profile network policy	Network Administrator
service-profile-qos	Service profile QoS	Network Administrator
service-profile-qos-policy	Service profile QoS policy	Network Administrator
service-profile-security	Service profile security	Server Security Administrator
service-profile-security-policy	Service profile security policy	Server Security Administrator

Privilege	Description	Default Role Assignment
service-profile-server	Service profile server management	Server Profile Administrator
service-profile-server-oper	Service profile consumer	Server Profile Administrator
service-profile-server-policy	Service profile pool policy	Server Security Administrator
service-profile-storage	Service profile storage	Storage Administrator
service-profile-storage-policy	Service profile storage policy	Storage Administrator

User Locales

A user can be assigned one or more locales. Each locale defines one or more organizations (domains) the user is allowed access, and access would be limited to the organizations specified in the locale. One exception to this rule is a locale without any organizations, which gives unrestricted access to system resources in all organizations.

A Cisco UCS domain can contain up to 48 user locales.

Users with AAA privileges (AAA Administrator role) can assign organizations to the locale of other users. The assignment of organizations is restricted to only those in the locale of the user assigning the organizations. For example, if a locale contains only the Engineering organization then a user assigned that locale can only assign the Engineering organization to other users.



Note

You cannot assign a locale to users with one or more of the following privileges:

- aaa
- admin
- operations

You can hierarchically manage organizations. A user that is assigned at a top level organization has automatic access to all organizations under it. For example, an Engineering organization can contain a Software Engineering organization and a Hardware Engineering organization. A locale containing only the Software Engineering organization has access to system resources only within that organization; however, a locale that contains the Engineering organization has access to the resources for both the Software Engineering and Hardware Engineering organizations.

Configuring User Roles

Creating a User Role

Procedure

	Command or Action	Purpose
Step 1	UCS-A# scope security	Enters security mode.
Step 2	UCS-A /security # create role name	Creates the user role and enters security role mode.
Step 3	UCS-A /security/role # add privilege privilege-name	<p>Adds one or more privileges to the role.</p> <p>Note You can specify more than one <i>privilege-name</i> on the same command line to add multiple privileges to the role, or you can add privileges to the same role using multiple add commands.</p>
Step 4	UCS-A /security/role # commit-buffer	Commits the transaction to the system configuration.

The following example creates the service-profile-security-admin role, adds the service profile security and service profile security policy privileges to the role, and commits the transaction:

```
UCS-A# scope security
UCS-A /security # create role ls-security-admin
UCS-A /security/role* # add privilege service-profile-security service-profile-security-policy
UCS-A /security/role* # commit-buffer
UCS-A /security/role #
```

Adding Privileges to a User Role

Procedure

	Command or Action	Purpose
Step 1	UCS-A# scope security	Enters security mode.
Step 2	UCS-A /security # scope role name	Enters security role mode for the specified role.
Step 3	UCS-A /security/role # add privilege privilege-name	<p>Adds one or more privileges to the existing privileges of the user role.</p> <p>Note You can specify more than one <i>privilege-name</i> on the same command line to add multiple privileges to the role, or you can add privileges to the same role using multiple add privilege commands.</p>

	Command or Action	Purpose
Step 4	UCS-A /security/role # commit-buffer	Commits the transaction to the system configuration.

The following example adds the server security and server policy privileges to the service-profile-security-admin role and commits the transaction:

```
UCS-A# scope security
UCS-A /security # scope role service-profile-security-admin
UCS-A /security/role # add privilege server-security server-policy
UCS-A /security/role* # commit-buffer
UCS-A /security/role #
```

Removing Privileges from a User Role

Procedure

	Command or Action	Purpose
Step 1	UCS-A# scope security	Enters security mode.
Step 2	UCS-A /security # scope role name	Enters security role mode for the specified role.
Step 3	UCS-A /security/role # remove privilege privilege-name	Removes one or more privileges from the existing user role privileges. Note You can specify more than one <i>privilege-name</i> on the same command line to remove multiple privileges from the role, or you can remove privileges from the same role using multiple remove privilege commands.
Step 4	UCS-A /security/role # commit-buffer	Commits the transaction to the system configuration.

The following example removes the server security and server policy privileges from the service-profile-security-admin role and commits the transaction:

```
UCS-A# scope security
UCS-A /security # scope role service-profile-security-admin
UCS-A /security/role # remove privilege server-security server-policy
UCS-A /security/role* # commit-buffer
UCS-A /security/role #
```

Deleting a User Role

Procedure

	Command or Action	Purpose
Step 1	UCS-A# scope security	Enters security mode.
Step 2	UCS-A /security # delete role name	Deletes the user role.
Step 3	UCS-A /security # commit-buffer	Commits the transaction to the system configuration.

The following example deletes the service-profile-security-admin role and commits the transaction:

```
UCS-A# scope security
UCS-A /security # delete role service-profile-security-admin
UCS-A /security* # commit-buffer
UCS-A /security #
```

Configuring Locales

Creating a Locale

Procedure

	Command or Action	Purpose
Step 1	UCS-A# scope security	Enters security mode.
Step 2	UCS-A /security # create locale locale-name	Creates a locale and enters security locale mode.
Step 3	UCS-A /security/locale # create org-ref org-ref-name orgdn orgdn-name	References (binds) an organization to the locale. The <i>org-ref-name</i> argument is the name used to identify the organization reference, and the <i>orgdn-name</i> argument is the distinguished name of the organization being referenced.
Step 4	UCS-A /security/locale # commit-buffer	Commits the transaction to the system configuration.

The following example creates the western locale, references the finance organization to the locale, names the reference finance-ref, and commits the transaction:

```
UCS-A# scope security
UCS-A /security # create locale western
UCS-A /security/locale* # create org-ref finance-ref orgdn finance
UCS-A /security/locale* # commit-buffer
UCS-A /security/locale #
```

Assigning an Organization to a Locale

Procedure

	Command or Action	Purpose
Step 1	UCS-A# scope security	Enters security mode.
Step 2	UCS-A# scope locale <i>locale-name</i>	Enters security locale mode.
Step 3	UCS-A /security/locale # create org-ref <i>org-ref-name</i> orgdn <i>orgdn-name</i>	References (binds) an organization to the locale. The <i>org-ref-name</i> argument is the name used to identify the organization reference, and the <i>orgdn-name</i> argument is the distinguished name of the organization being referenced.
Step 4	UCS-A /security/locale # commit-buffer	Commits the transaction to the system configuration.

The following example enters the western locale, adds (references) the marketing organization to the locale, names the reference marketing-ref, and commits the transaction:

```
UCS-A# scope security
UCS-A /security # scope locale western
UCS-A /security/locale* # create org-ref marketing-ref orgdn marketing
UCS-A /security/locale* # commit-buffer
UCS-A /security/locale #
```

Deleting an Organization from a Locale

Procedure

	Command or Action	Purpose
Step 1	UCS-A# scope security	Enters security mode.
Step 2	UCS-A /security # scope locale <i>locale-name</i>	Enters security locale mode.
Step 3	UCS-A /security/locale # delete org-ref <i>org-ref-name</i>	Deletes the organization from the locale.
Step 4	UCS-A /security/locale # commit-buffer	Commits the transaction to the system configuration.

The following example deletes the finance organization from the western locale and commits the transaction:

```
UCS-A# scope security
UCS-A /security # scope locale western
UCS-A /security/locale # delete org-ref finance-ref
UCS-A /security/locale* # commit-buffer
UCS-A /security/locale #
```

Deleting a Locale

Procedure

	Command or Action	Purpose
Step 1	UCS-A# scope security	Enters security mode.
Step 2	UCS-A /security # delete locale locale-name	Deletes the locale.
Step 3	UCS-A /security # commit-buffer	Commits the transaction to the system configuration.

The following example deletes the western locale and commits the transaction:

```
UCS-A# scope security
UCS-A /security # delete locale western
UCS-A /security* # commit-buffer
UCS-A /security #
```

Configuring Locally Authenticated User Accounts

Creating a User Account

At a minimum, we recommend that you create the following users:

- Server administrator account
- Network administrator account
- Storage administrator

Before You Begin

Perform the following tasks, if the system includes any of the following:

- Remote authentication services, ensure the users exist in the remote authentication server with the appropriate roles and privileges.
- Multi-tenancy with organizations, create one or more locales. If you do not have any locales, all users are created in root and are assigned roles and privileges in all organizations.
- SSH authentication, obtain the SSH key.

Procedure

	Command or Action	Purpose
Step 1	UCS-A# scope security	Enters security mode.

	Command or Action	Purpose
Step 2	UCS-A /security # create local-user <i>local-user-name</i>	Creates a user account for the specified local user and enters security local user mode.
Step 3	UCS-A /security/local-user # set account-status { active inactive }	Specifies whether the local user account is enabled or disabled. If the account status for a local user account is set to inactive, the user is prevented from logging into the system using their existing credentials.
Step 4	UCS-A /security/local-user # set password <i>password</i>	Sets the password for the user account
Step 5	UCS-A /security/local-user # set firstname <i>first-name</i>	(Optional) Specifies the first name of the user.
Step 6	UCS-A /security/local-user # set lastname <i>last-name</i>	(Optional) Specifies the last name of the user.
Step 7	UCS-A /security/local-user # set expiration <i>month day-of-month year</i>	(Optional) Specifies the date that the user account expires. The <i>month</i> argument is the first three letters of the month name. Note After you configure a user account with an expiration date, you cannot reconfigure the account to not expire. You can, however, configure the account with the latest expiration date available.
Step 8	UCS-A /security/local-user # set email <i>email-addr</i>	(Optional) Specifies the user e-mail address.
Step 9	UCS-A /security/local-user # set phone <i>phone-num</i>	(Optional) Specifies the user phone number.
Step 10	UCS-A /security/local-user # set sshkey <i>ssh-key</i>	(Optional) Specifies the SSH key used for passwordless access.
Step 11	UCS-A security/local-user # commit-buffer	Commits the transaction.

The following example creates the user account named kikipopo, enables the user account, sets the password to foo12345, and commits the transaction:

```
UCS-A# scope security
UCS-A /security # create local-user kikipopo
UCS-A /security/local-user* # set account-status active
UCS-A /security/local-user* # set password
Enter a password:
Confirm the password:
UCS-A /security/local-user* # commit-buffer
UCS-A /security/local-user #
```

The following example creates the user account named lincey, enables the user account, sets an OpenSSH key for passwordless access, and commits the transaction.

```
UCS-A# scope security
UCS-A /security # create local-user lincey
UCS-A /security/local-user* # set account-status active
UCS-A /security/local-user* # set sshkey "ssh-rsa
AAAAB3NzaC1yc2EAAAABIwAAAIEAuo9VQ2CmWBI9/S1f30k1CWjnV3lgdXMzO0WU15iPw851kdQqap+NFuNmHcb4K
iaQB8X/PDdmtlxQQcawclj+k8f4VcOelBx1sGk5luq51s1ob1VOIEwKEL/h51rdbN1I8y3SS9I/gGiBZ9ARlop9LDpD
m8HPh2LOgyH7Ei1MI8="
UCS-A /security/local-user* # commit-buffer
UCS-A /security/local-user #
```

The following example creates the user account named jforlenz, enables the user account, sets a Secure SSH key for passwordless access, and commits the transaction.

```
UCS-A# scope security
UCS-A /security # create local-user jforlenz
UCS-A /security/local-user* # set account-status active
UCS-A /security/local-user* # set sshkey
Enter lines one at a time. Enter ENDOFBUF to finish. Press ^C to abort.
User's SSH key:
> ---- BEGIN SSH2 PUBLIC KEY ----
>AAAAB3NzaC1yc2EAAAABIwAAAIEAuo9VQ2CmWBI9/S1f30k1CWjnV3lgdXMzO0WU15iPw8
>51kdQqap+NFuNmHcb4KiaQB8X/PDdmtlxQQcawclj+k8f4VcOelBx1sGk5luq51s1ob1VO
>IEwKEL/h51rdbN1I8y3SS9I/gGiBZ9ARlop9LDpDm8HPh2LOgyH7Ei1MI8=
> ---- END SSH2 PUBLIC KEY ----
> ENDOFBUF
UCS-A /security/local-user* # commit-buffer
UCS-A /security/local-user #
```

Enabling the Password Strength Check for Locally Authenticated Users

You must be a user with admin or aaa privileges to enable the password strength check. If the password strength check is enabled, Cisco UCS Manager does not permit a user to choose a password that does not meet the guidelines for a strong password.

Procedure

	Command or Action	Purpose
Step 1	UCS-A# scope security	Enters security mode.
Step 2	UCS-A /security # enforce-strong-password {yes no}	Specifies whether the password strength check is enabled or disabled.

The following example enables the password strength check:

```
UCS-A# scope security
UCS-A /security # set enforce-strong-password yes
UCS-A /security #
```

Setting Web Session Limits for User Accounts

Procedure

	Command or Action	Purpose
Step 1	UCS-A# scope system	Enters system mode.
Step 2	UCS-A /system # scope services	Enters system services mode.
Step 3	UCS-A /system/services # scope web-session-limits	Enters system services web session limits mode.
Step 4	UCS-A /system/services/web-session-limits # set peruser num-of-logins-per-user	Sets the maximum number of concurrent HTTP and HTTPS sessions allowed for each user. Enter an integer between 1 and 256. By default, this value is set to 32.
Step 5	UCS-A /system/services/web-session-limits # commit-buffer	Commits the transaction to the system configuration.

The following example sets the maximum number of HTTP and HTTPS sessions allowed by each user account to 60 and commits the transaction:

```
UCS-A# scope system
UCS-A /system # scope services
UCS-A /system/services # scope web-session-limits
UCS-A /system/services/web-session-limits* # set peruser 60
UCS-A /system/services/web-session-limits* # commit-buffer
UCS-A /system/services/web-session-limits #
```

Assigning a Role to a User Account

Changes in user roles and privileges do not take effect until the next time the user logs in. If a user is logged in when you assign a new role to or remove an existing role from a user account, the active session continues with the previous roles and privileges.

Procedure

	Command or Action	Purpose
Step 1	UCS-A# scope security	Enters security mode.
Step 2	UCS-A /security # scope local-user local-user-name	Enters security local user mode for the specified local user account.
Step 3	UCS-A /security/local-user # create role role-name	Assigns the specified role to the user account . Note The create role command can be entered multiple times to assign more than one role to a user account.

	Command or Action	Purpose
Step 4	UCS-A security/local-user # commit-buffer	Commits the transaction.

The following example assigns the operations role to the kikipopo local user account and commits the transaction:

```
UCS-A# scope security
UCS-A /security # scope local-user kikipopo
UCS-A /security/local-user # create role operations
UCS-A /security/local-user* # commit-buffer
UCS-A /security/local-user #
```

Assigning a Locale to a User Account



Note Do not assign locales to users with an admin or aaa role.

Procedure

	Command or Action	Purpose
Step 1	UCS-A# scope security	Enters security mode.
Step 2	UCS-A /security # scope local-user <i>local-user-name</i>	Enters security local user mode for the specified local user account.
Step 3	UCS-A /security/local-user # create locale <i>locale-name</i>	Assigns the specified locale to the user account. Note The create locale command can be entered multiple times to assign more than one locale to a user account.
Step 4	UCS-A security/local-user # commit-buffer	Commits the transaction.

The following example assigns the western locale to the kikipopo local user account and commits the transaction:

```
UCS-A# scope security
UCS-A /security # scope local-user kikipopo
UCS-A /security/local-user # create locale western
UCS-A /security/local-user* # commit-buffer
UCS-A /security/local-user #
```

Removing a Role from a User Account

Changes in user roles and privileges do not take effect until the next time the user logs in. If a user is logged in when you assign a new role to or remove an existing role from a user account, the active session continues with the previous roles and privileges.

Procedure

	Command or Action	Purpose
Step 1	UCS-A# scope security	Enters security mode.
Step 2	UCS-A /security # scope local-user <i>local-user-name</i>	Enters security local user mode for the specified local user account.
Step 3	UCS-A /security/local-user # delete role <i>role-name</i>	Removes the specified role from the user account . Note The delete role command can be entered multiple times to remove more than one role from a user account.
Step 4	UCS-A security/local-user # commit-buffer	Commits the transaction.

The following example removes the operations role from the kikipopo local user account and commits the transaction:

```
UCS-A# scope security
UCS-A /security # scope local-user kikipopo
UCS-A /security/local-user # delete role operations
UCS-A /security/local-user* # commit-buffer
UCS-A /security/local-user #
```

Removing a Locale from a User Account

Procedure

	Command or Action	Purpose
Step 1	UCS-A# scope security	Enters security mode.
Step 2	UCS-A /security # scope local-user <i>local-user-name</i>	Enters security local user mode for the specified local user account.
Step 3	UCS-A /security/local-user # delete locale <i>locale-name</i>	Removes the specified locale from the user account. Note The delete locale command can be entered multiple times to remove more than one locale from a user account.

	Command or Action	Purpose
Step 4	UCS-A security/local-user # commit-buffer	Commits the transaction.

The following example removes the western locale from the kikipopo local user account and commits the transaction:

```
UCS-A# scope security
UCS-A /security # scope local-user kikipopo
UCS-A /security/local-user # delete locale western
UCS-A /security/local-user* # commit-buffer
UCS-A /security/local-user #
```

Enabling or Disabling a User Account

You must be a user with admin or aaa privileges to enable or disable a local user account.

Before You Begin

Create a local user account.

Procedure

	Command or Action	Purpose
Step 1	UCS-A# scope security	Enters security mode.
Step 2	UCS-A /security # scope local-user	Enters local-user security mode.
Step 3	UCS-A /security/local-user # set account-status {active inactive}	Specifies whether the local user account is enabled or disabled. The admin user account is always set to active. It cannot be modified. Note If you set the account status to inactive, the configuration is not deleted from the database.

The following example enables a local user account called accounting:

```
UCS-A# scope security
UCS-A /security # scope local-user accounting
UCS-A /security/local-user # set account-status active
```

Clearing the Password History for a Locally Authenticated User

Procedure

	Command or Action	Purpose
Step 1	UCS-A# scope security	Enters security mode.
Step 2	UCS-A /security # scope local-user <i>user-name</i>	Enters local user security mode for the specified user account.
Step 3	UCS-A /security/local-user # set clear password-history yes	Clears the password history for the specified user account.
Step 4	UCS-A /security/local-user # commit-buffer	Commits the transaction to the system configuration.

The following example configures the password history count and commits the transaction:

```
UCS-A # scope security
UCS-A /security # scope local-user admin
UCS-A /security/local-user # set clear password-history yes
UCS-A /security/local-user* # commit-buffer
UCS-A /security/local-user #
```

Deleting a User Account

Procedure

	Command or Action	Purpose
Step 1	UCS-A# scope security	Enters security mode.
Step 2	UCS-A /security # delete local-user <i>local-user-name</i>	Deletes the local-user account.
Step 3	UCS-A /security # commit-buffer	Commits the transaction to the system configuration.

The following example deletes the foo user account and commits the transaction:

```
UCS-A# scope security
UCS-A /security # delete local-user foo
UCS-A /security* # commit-buffer
UCS-A /security #
```

Password Profile for Locally Authenticated Users

The password profile contains the password history and password change interval properties for all locally authenticated users of Cisco UCS Manager. You cannot specify a different password profile for each locally authenticated user.


Note

You must have admin or aaa privileges to change the password profile properties. Except for password history, these properties do not apply to users with admin or aaa privileges.

Password History Count

The password history count allows you to prevent locally authenticated users from reusing the same password over and over again. When this property is configured, Cisco UCS Manager stores passwords that were previously used by locally authenticated users up to a maximum of 15 passwords. The passwords are stored in reverse chronological order with the most recent password first to ensure that the only the oldest password can be reused when the history count threshold is reached.

A user must create and use the number of passwords configured in the password history count before being able to reuse one. For example, if you set the password history count to 8, a locally authenticated user cannot reuse the first password until after the ninth password has expired.

By default, the password history is set to 0. This value disables the history count and allows users to reuse previously passwords at any time.

If necessary, you can clear the password history count for a locally authenticated user and enable reuse of previous passwords.

Password Change Interval

The password change interval enables you to restrict the number of password changes a locally authenticated user can make within a given number of hours. The following table describes the two configuration options for the password change interval.

Interval Configuration	Description	Example
No password change allowed	This option does not passwords for locally authenticated users to be changed within a specified number of hours after a password change. You can specify a no change interval between 1 and 745 hours. By default, the no change interval is 24 hours.	For example, to prevent passwords from being changed within 48 hours after a locally authenticated user changes his or her password, set the following: <ul style="list-style-type: none"> • Change during interval to disable • No change interval to 48

Interval Configuration	Description	Example
Password changes allowed within change interval	<p>This option specifies the maximum number of times that passwords for locally authenticated users can be changed within a pre-defined interval.</p> <p>You can specify a change interval between 1 and 745 hours and a maximum number of password changes between 0 and 10. By default, a locally authenticated user is permitted a maximum of 2 password changes within a 48 hour interval.</p>	<p>For example, to allow to be changed a maximum of once within 24 hours after a locally authenticated user changes his or her password, set the following:</p> <ul style="list-style-type: none"> • Change during interval to enable • Change count to 1 • Change interval to 24

Configuring the Maximum Number of Password Changes for a Change Interval

You must have admin or aaa privileges to change the password profile properties. Except for password history, these properties do not apply to users with admin or aaa privileges.

Procedure

	Command or Action	Purpose
Step 1	UCS-A# scope security	Enters security mode.
Step 2	UCS-A /security # scope password-profile	Enters password profile security mode.
Step 3	UCS-A /security/password-profile # set change-during-interval enable	Restricts the number of password changes a locally authenticated user can make within a given number of hours.
Step 4	UCS-A /security/password-profile # set change-count <i>pass-change-num</i>	<p>Specifies the maximum number of times a locally authenticated user can change his or her password during the Change Interval.</p> <p>This value can be anywhere from 0 to 10.</p>
Step 5	UCS-A /security/password-profile # set change-interval <i>num-of-hours</i>	<p>Specifies the maximum number of hours over which the number of password changes specified in the Change Count field are enforced.</p> <p>This value can be anywhere from 1 to 745 hours.</p> <p>For example, if this field is set to 48 and the Change Count field is set to 2, a locally authenticated user can make no more than 2 password changes within a 48 hour period.</p>
Step 6	UCS-A /security/password-profile # commit-buffer	Commits the transaction to the system configuration.

The following example enables the change during interval option, sets the change count to 5, sets the change interval to 72 hours, and commits the transaction:

```
UCS-A # scope security
UCS-A /security # scope password-profile
UCS-A /security/password-profile # set change-during-interval enable
UCS-A /security/password-profile* # set change-count 5
UCS-A /security/password-profile* # set change-interval 72
UCS-A /security/password-profile* # commit-buffer
UCS-A /security/password-profile #
```

Configuring a No Change Interval for Passwords

You must have admin or aaa privileges to change the password profile properties. Except for password history, these properties do not apply to users with admin or aaa privileges.

Procedure

	Command or Action	Purpose
Step 1	UCS-A# scope security	Enters security mode.
Step 2	UCS-A /security # scope password-profile	Enters password profile security mode.
Step 3	UCS-A /security/password-profile # set change-during-interval disable	Disables the change during interval feature.
Step 4	UCS-A /security/password-profile # set no-change-interval min-num-hours	Specifies the minimum number of hours that a locally authenticated user must wait before changing a newly created password. This value can be anywhere from 1 to 745 hours. This interval is ignored if the Change During Interval property is not set to Disable .
Step 5	UCS-A /security/password-profile # commit-buffer	Commits the transaction to the system configuration.

The following example disables the change during interval option, sets the no change interval to 72 hours, and commits the transaction:

```
UCS-A # scope security
UCS-A /security # scope password-profile
UCS-A /security/password-profile # set change-during-interval disable
UCS-A /security/password-profile* # set no-change-interval 72
UCS-A /security/password-profile* # commit-buffer
UCS-A /security/password-profile #
```

Configuring the Password History Count

You must have admin or aaa privileges to change the password profile properties.

Procedure

	Command or Action	Purpose
Step 1	UCS-A# scope security	Enters security mode.
Step 2	UCS-A /security # scope password-profile	Enters password profile security mode.
Step 3	UCS-A /security/password-profile # set history-count num-of-passwords	Specifies the number of unique passwords that a locally authenticated user must create before that user can reuse a previously used password This value can be anywhere from 0 to 15. By default, the History Count field is set to 0, which disables the history count and allows users to reuse previously used passwords at any time.
Step 4	UCS-A /security/password-profile # commit-buffer	Commits the transaction to the system configuration.

The following example configures the password history count and commits the transaction:

```
UCS-A # scope security
UCS-A /security # scope password-profile
UCS-A /security/password-profile # set history-count 5
UCS-A /security/password-profile* # commit-buffer
UCS-A /security/password-profile #
```

Monitoring User Sessions

Procedure

	Command or Action	Purpose
Step 1	UCS-A# scope security	Enters security mode.
Step 2	UCS-A /security # show user-session {local remote} [detail]	Displays session information for all users logged in to the system. An asterisk (*) next to the session ID denotes the current login session.

The following example lists all local users logged in to the system. The asterisk indicates which session is the current login session.

```
UCS-A# scope security
UCS-A /security # show user-session local
Session Id      User           Host           Login Time
-----
pts_25_1_31264*  steve         192.168.100.111 2009-05-09T14:06:59
ttyS0_1_3532    jeff          console         2009-05-02T15:11:08
web_25277_A     faye          192.168.100.112 2009-05-15T22:11:25
```


The following example displays detailed information on all local users logged in to the system:

```
UCS-A# scope security
UCS-A /security # show user-session local detail
Session Id pts_25_1_31264:
  Fabric Id: A
  Term: pts/25
  User: steve
  Host: 64.101.53.93
  Pid: 31264
  Login Time: 2009-05-09T14:06:59

Session Id ttyS0_1_3532:
  Fabric Id: A
  Term: ttyS0
  User: jeff
  Host: console
  Pid: 3532
  Login Time: 2009-05-02T15:11:08

Session Id web_25277_A:
  Fabric Id: A
  Term: web_25277
  User: faye
  Host: 192.168.100.112
  Pid: 3518
  Login Time: 2009-05-15T22:11:25
```




CHAPTER 11

Managing Firmware

This chapter includes the following sections:

- [Overview of Firmware, page 185](#)
- [Firmware Image Management, page 186](#)
- [Firmware Versions, page 188](#)
- [Firmware Upgrades, page 189](#)
- [Firmware Downgrades, page 201](#)
- [Completing the Prerequisites for Upgrading the Firmware, page 201](#)
- [Downloading and Managing Firmware Packages, page 206](#)
- [Directly Upgrading Firmware at Endpoints, page 214](#)
- [Updating Firmware through Service Profiles, page 225](#)
- [Managing the Capability Catalog, page 232](#)
- [Updating Management Extensions, page 238](#)

Overview of Firmware

Cisco UCS uses firmware obtained from and certified by Cisco to support the endpoints in a Cisco UCS domain. Each endpoint is a component in the Cisco UCS domain that requires firmware to function. The upgrade order for the endpoints in a Cisco UCS domain depends upon the upgrade path, but includes the following:

- Cisco UCS Manager
- I/O modules
- Fabric interconnects
- Endpoints physically located on adapters, including NIC and HBA firmware, and Option ROM (where applicable) that can be upgraded through firmware packages included in a service profile

- Endpoints physically located on servers, such as the BIOS, storage controller (RAID controller), and Cisco Integrated Management Controller (CIMC) that can be upgraded through firmware packages included in a service profile

See the required order of steps for your upgrade path to determine the appropriate order in which to upgrade the endpoints in your Cisco UCS domain.

**Note**

Beginning with Cisco UCS, Release 1.4(1), Cisco is releasing firmware upgrades in multiple bundles, rather than one large firmware package. For more information see [Firmware Image Management, on page 186](#).

Cisco maintains a set of best practices for managing firmware images and updates in this document and in the following technical note: [Unified Computing System Firmware Management Best Practices](#).

This document uses the following definitions for managing firmware:

Upgrade

Changes the firmware running on an endpoint to another image, such as a release or patch. Upgrade includes both update and activation.

Update

Copies the firmware image to the backup partition on an endpoint.

Activate

Sets the firmware in the backup partition as the active firmware version on the endpoint. Activation can require or cause the reboot of an endpoint.

For Management Extensions and Capability Catalog upgrades, update and activate occur simultaneously. You only need to update or activate those upgrades. You do not need to perform both steps.

Firmware Image Management

Cisco delivers all firmware updates to Cisco UCS components in bundles of images. Cisco UCS firmware updates are available to be downloaded in the following bundles:

Cisco UCS Infrastructure Software Bundle

This bundle includes the following firmware images that are required to update the following components:

- Cisco UCS Manager software
- Kernel and system firmware for the fabric interconnects
- I/O module firmware

Cisco UCS B-Series Blade Server Software Bundle

This bundle includes the following firmware images that are required to update the firmware for the blade servers in a Cisco UCS domain. In addition to the bundles created for a release, these bundles can also be released between infrastructure bundles to enable Cisco UCS Manager to support a blade server that is not included in the most recent infrastructure bundle.

- CIMC firmware
- BIOS firmware
- Adapter firmware
- Board controller firmware
- Third-party firmware images required by the new server

Cisco UCS C-Series Rack-Mount Server Software Bundle

This bundle includes the following firmware images that are required to update components on rack-mount servers that have been integrated with and are managed by Cisco UCS Manager:

- CIMC firmware
- BIOS firmware
- Adapter firmware
- Storage controller firmware



Note

You cannot use this bundle for standalone C-series servers. The firmware management system in those servers cannot interpret the header required by Cisco UCS Manager. For information on how to upgrade standalone C-series servers, see the C-series configuration guides.

Cisco also provides release notes, which you can obtain on the same website from which you obtained the bundles.

Firmware Image Headers

Every firmware image has a header, which includes the following:

- Checksum
- Version information
- Compatibility information that the system can use to verify the compatibility of component images and any dependencies

Firmware Image Catalog

Cisco UCS Manager provides you with two views of the catalog of firmware images and their contents that have been downloaded to the fabric interconnect:

Packages

This view provides you with a read-only representation of the firmware bundles that have been downloaded onto the fabric interconnect. This view is sorted by image, not by the contents of the image. For packages, you can use this view to see which component images are in each downloaded firmware bundle.

Images

The images view lists the component images available on the system. You cannot use this view to see complete firmware bundles or to group the images by bundle. The information available about each component image includes the name of the component, the image size, the image version, and the vendor and model of the component.

You can use this view to identify the firmware updates available for each component. You can also use this view to delete obsolete and unneeded images. Cisco UCS Manager deletes a package after all images in the package have been deleted.



Tip

Cisco UCS Manager stores the images in bootflash on the fabric interconnect. In a cluster system, space usage in bootflash on both fabric interconnects is the same, because all images are synchronized between them. If Cisco UCS Manager reports that the bootflash is out of space, delete obsolete images to free up space.

Firmware Versions

The firmware version terminology used depends upon the type of endpoint, as follows:

Firmware Versions in CIMC, I/O Modules, and Adapters

Each CIMC, I/O module, and adapter has two slots for firmware in flash. Each slot holds a version of firmware. One slot is active and the other is the backup slot. A component boots from whichever slot is designated as active.

The following firmware version terminology is used in Cisco UCS Manager:

Running Version

The running version is the firmware that is active and in use by the endpoint.

Startup Version

The startup version is the firmware that will be used when the endpoint next boots up. Cisco UCS Manager uses the activate operation to change the startup version.

Backup Version

The backup version is the firmware in the other slot and is not in use by the endpoint. This version can be firmware that you have updated to the endpoint but have not yet activated, or it can be an older firmware version that was replaced by a recently activated version. Cisco UCS Manager uses the update operation to replace the image in the backup slot.

If the endpoint cannot boot from the startup version, it boots from the backup version.

Firmware Versions in the Fabric Interconnect and Cisco UCS Manager

You can only activate the fabric interconnect firmware and Cisco UCS Manager on the fabric interconnect. The fabric interconnect and Cisco UCS Manager firmware do not have backup versions, because all the images are stored on the fabric interconnect. As a result, the number of bootable fabric interconnect images is not limited to two, like the server CIMC and adapters. Instead, the number of bootable fabric interconnect images is limited by the available space in the memory of the fabric interconnect and the number of images stored there.

The fabric interconnect and Cisco UCS Manager firmware have running and startup versions of the kernel and system firmware. The kernel and system firmware must run the same versions of firmware.

Firmware Upgrades

Cisco UCS firmware is upgraded through a combination of the following methods:

- Direct upgrade at the endpoints. For a cluster configuration with two fabric interconnects, a direct upgrade can be minimally disruptive to data traffic. However, it requires that the Cisco UCS domain does not include firmware policies for those endpoints that you upgrade directly. You cannot avoid disruption to traffic in a Cisco UCS domain with only one fabric interconnect.



Note Direct upgrade is not available for all endpoints, including the server BIOS, storage controller, HBA firmware, and HBA option ROM. You must upgrade those endpoints through the host firmware package included in the service profile associated with the server.

- Upgrades to server endpoints through service profiles that include a host firmware package, a management firmware package, or both. This method can be disruptive to data traffic and should be performed during a maintenance window.



Note The Cisco UCS Manager CLI does not allow you to upgrade hardware that is not supported in the release to which you are upgrading, Cisco UCS Manager CLI displays an error message if you attempt to upgrade hardware to an unsupported release.

Cautions, Guidelines, and Best Practices for Firmware Upgrades

Before you upgrade the firmware for any endpoint in a Cisco UCS domain, consider the following cautions, guidelines, and best practices:



Note The Cisco UCS Manager CLI does not allow you to upgrade hardware that is not supported in the release to which you are upgrading, Cisco UCS Manager CLI displays an error message if you attempt to upgrade hardware to an unsupported release.

Configuration Changes and Settings that Can Impact Upgrades

Depending upon the configuration of your Cisco UCS domain, the following changes may require you to make configuration changes after you upgrade. To avoid faults and other issues, we recommend that you make any required changes before you upgrade.

Overlapping FCoE VLAN IDs and Ethernet VLAN IDs Are No Longer Allowed with Cisco UCS Release 2.0



Caution

In Cisco UCS 1.4 and earlier releases, Ethernet VLANs and FCoE VLANs could have overlapping VLAN IDs. However, starting with Cisco UCS release 2.0, overlapping VLAN IDs are not allowed. If Cisco UCS Manager detects overlapping VLAN IDs during an upgrade, it raises a critical fault. If you do not reconfigure your VLAN IDs, Cisco UCS Manager raises a critical fault and drops Ethernet traffic on the overlapped VLANs. Therefore, we recommend that you ensure there are no overlapping Ethernet and FCoE VLAN IDs before you upgrade to Cisco UCS release 2.0.

If you did not explicitly configure the FCoE VLAN ID for a VSAN in Cisco UCS 1.4 and earlier releases, Cisco UCS Manager assigned VLAN 1 as the default FCoE VLAN for the default VSAN (with default VSAN ID 1). In those releases, VLAN 1 was also used as the default VLAN for Ethernet traffic. Therefore, if you accepted the default VLAN ID for the FCoE VLAN and one or more Ethernet VLANs, you must reconfigure the VLAN IDs for either the FCoE VLAN(s) on the VSAN(s) or the Ethernet VLAN(s).

For a new installation of Cisco UCS release 2.0, the default VLAN IDs are as follows:

- The default Ethernet VLAN ID is 1.
- The default FCoE VLAN ID is 4048.

After an upgrade from Cisco UCS release 1.4, where VLAN ID 4048 was used for FCoE storage port native VLAN, to release 2.0, the default VLAN IDs are as follows:

- The default Ethernet VLAN ID is 1.
- The current default FCoE VLAN ID is preserved. Cisco UCS Manager raises a critical fault on the conflicting Ethernet VLAN, if any. You must change one of the VLAN IDs to a VLAN ID that is not used or reserved.



Note

If a Cisco UCS domain uses one of the default VLAN IDs, which results in overlapping VLANs, you can change one or more of the default VLAN IDs to any VLAN ID that is not used or reserved. In release 2.0, VLANs with IDs from 3968 to 4047 are reserved.

VSANs with IDs in the Reserved Range are not Operational

A VSAN with an ID in the reserved range is not operational after an upgrade. Make sure that none of the VSANs configured in Cisco UCS Manager are in the reserved range, as follows:

- If you plan to use FC switch mode in a Cisco UCS domain, do not configure VSANs with an ID in the range from 3040 to 4078.
- If you plan to use FC end-host mode in a Cisco UCS domain, do not configure VSANs with an ID in the range from 3840 to 4079.

If a VSAN has an ID in the reserved range, change that VSAN ID to any VSAN ID that is not used or reserved.

All Connectivity May Be Lost During Upgrades if vNIC Failover and NIC Teaming Are Both Enabled

All connectivity may be lost during firmware upgrades if you have configured both **Enable Failover** on one or more vNICs and you have also configured NIC teaming/bonding at the host operating system level. Please design for availability by using one or the other method, but never both.

To determine whether you have enabled failover for one or more vNICs in a Cisco UCS domain, verify the configuration of the vNICs within each service profile associated with a server. For more information, see the [Cisco UCS Manager configuration guide](#) for the release that you are running.

Impact of Upgrade from a Release Prior to Release 1.3(1i)

An upgrade from an earlier Cisco UCS firmware release to release 1.3(1i) or higher has the following impact on the Protect Configuration property of the local disk configuration policy the first time servers are associated with service profiles after the upgrade:

Unassociated Servers

After you upgrade the Cisco UCS domain, the initial server association proceeds without configuration errors whether or not the local disk configuration policy matches the server hardware. Even if you enable the Protect Configuration property, Cisco UCS does not protect the user data on the server if there are configuration mismatches between the local disk configuration policy on the previous service profile and the policy in the new service profile.



Note If you enable the Protect Configuration property and the local disk configuration policy encounters mismatches between the previous service profile and the new service profile, all subsequent service profile associations with the server are blocked.

Associated Servers

Any servers that are already associated with service profiles do not reboot after the upgrade. Cisco UCS Manager does not report any configuration errors if there is a mismatch between the local disk configuration policy and the server hardware.

When a service profile is disassociated from a server and a new service profile associated, the setting for the Protect Configuration property in the new service profile takes precedence and overwrites the setting in the previous service profile.

Hardware-Related Guidelines and Best Practices for Firmware Upgrades

The hardware in a Cisco UCS domain can impact how you upgrade. Before you upgrade any endpoint, consider the following guidelines and best practices:

No Server or Chassis Maintenance



Caution

Do not remove the hardware that contains the endpoint or perform any maintenance on it until the update process has completed. If the hardware is removed or otherwise unavailable due to maintenance, the firmware update fails. This failure may corrupt the backup partition. You cannot update the firmware on an endpoint with a corrupted backup partition.

Avoid Replacing RAID-Configured Hard Disks Prior to Upgrade

Under the following circumstances, Cisco UCS Manager may scrub all data on a hard disk as part of the RAID synchronization process during an upgrade of the server firmware:

- The hard disks in the server are configured for RAID.
- One or more of the RAID-configured hard disks in the server are removed.
- The hard disk or disks are replaced with hard disks that are configured with a pre-existing RAID and the local disk configuration policy included in the service profile on the server is not used to configure those hard disks.
- The server firmware is upgraded, causing the server to reboot and Cisco UCS Manager to begin the RAID synchronization process.

If the original hard disks contained vital data that needs to be preserved, avoid inserting new hard disks that are already configured for RAID.

Always Upgrade Cisco UCS Gen-2 Adapters through a Host Firmware Package

You cannot upgrade Cisco UCS Gen-2 adapters directly at the endpoints. You must upgrade the firmware on those adapters through a host firmware package.

Cannot Upgrade Cisco UCS 82598KR-CI 10-Gigabit Ethernet Adapter

The firmware on the Cisco UCS 82598KR-CI 10-Gigabit Ethernet Adapter (N20-AI0002), Intel-based adapter card, is burned into the hardware at manufacture. You cannot upgrade the firmware on this adapter.

Number of Fabric Interconnects

For a cluster configuration with two fabric interconnects, you can take advantage of the failover between the fabric interconnects and perform a direct firmware upgrade of the endpoints without disrupting data traffic. However, you cannot avoid disrupting data traffic for those endpoints which must be upgraded through a host or management firmware package.

For a standalone configuration with a single fabric interconnect, you can minimize the disruption to data traffic when you perform a direct firmware upgrade of the endpoints. However, you must reboot the fabric interconnect to complete the upgrade and, therefore, cannot avoid disrupting traffic.

Firmware- and Software-Related Best Practices for Upgrades

Before you upgrade any endpoint, consider the following guidelines and best practices:

Determine Appropriate Type of Firmware Upgrade for Each Endpoint

Some endpoints, such as adapters and the server CIMC, can be upgraded through either a direct firmware upgrade or a firmware package included in a service profile. The configuration of a Cisco UCS domain determines how you upgrade these endpoints. If the service profiles associated with the servers include a host firmware package, upgrade the adapters for those servers through the firmware package. In the same way, if the service profiles associated with the servers include a management firmware package, upgrade the CIMC for those servers through the firmware package.

Upgrades of a CIMC through a management firmware package or an adapter through a firmware package in the service profile associated with the server take precedence over direct firmware upgrades. You cannot directly upgrade an endpoint if the service profile associated with the server includes a firmware package. To perform a direct upgrade, you must remove the firmware package from the service profile.

Do Not Activate All Endpoints Simultaneously in Cisco UCS Manager GUI

If you use Cisco UCS Manager GUI to update the firmware, do not select **ALL** from the **Filter** drop-down list in the **Activate Firmware** dialog box to activate all endpoints simultaneously. Many firmware releases and patches have dependencies that require the endpoints to be activated in a specific order for the firmware update to succeed. This order can change depending upon the contents of the release or patch. Activating all endpoints does not guarantee that the updates occur in the required order and can disrupt communications between the endpoints and the fabric interconnects and Cisco UCS Manager. For information about the dependencies in a specific release or patch, see the release notes provided with that release or patch.

Impact of Activation for Adapters and I/O Modules

During a direct upgrade, you should configure **Set Startup Version Only** for an adapter. With this setting, the activated firmware moves into the pending-next-boot state, and the server is not immediately rebooted. The activated firmware does not become the running version of firmware on the adapter until the server is rebooted. You cannot configure **Set Startup Version Only** for an adapter in the host firmware package.

If a server is not associated with a service profile, the activated firmware remains in the pending-next-boot state. Cisco UCS Manager does not reboot the endpoints or activate the firmware until the server is associated with a service profile. If necessary, you can manually reboot or reset an unassociated server to activate the firmware.

When you configure **Set Startup Version Only** for an I/O module, the I/O module is rebooted when the fabric interconnect in its data path is rebooted. If you do not configure **Set Startup Version Only** for an I/O module, the I/O module reboots and disrupts traffic. In addition, if Cisco UCS Manager detects a protocol and firmware version mismatch between the fabric interconnect and the I/O module, Cisco UCS Manager automatically updates the I/O module with the firmware version that matches the firmware in the fabric interconnect and then activates the firmware and reboots the I/O module again.

Select Ignore Compatibility Check When Upgrading

During a direct upgrade to a newer release, we recommend that you choose **Ignore Compatibility Check**. Newer releases may have incompatible code with older releases. This option ensures that the upgrade can proceed and avoids compatibility issues.

Disable Call Home before Upgrading to Avoid Unnecessary Alerts (Optional)

When you upgrade a Cisco UCS domain, Cisco UCS Manager restarts the components to complete the upgrade process. This restart causes events that are identical to service disruptions and component failures that trigger Call Home alerts to be sent. If you do not disable Call Home before you begin the upgrade, you can ignore the alerts generated by the upgrade-related component restarts.

Required Order of Components for Firmware Activation

If you upgrade firmware by individual components in a Cisco UCS domain, activate the updates in the required order for quicker activation and to avoid potential issues with conflicting firmware versions.

Summary of Steps for Upgrading from Cisco UCS, Release 1.0(2) and Later

- 1 Download the following firmware images:
 - Cisco UCS Infrastructure Software Bundle—Required for all Cisco UCS domains.
 - Cisco UCS B-Series Blade Server Software Bundle—Required for all Cisco UCS domains that include blade servers.
 - Cisco UCS C-Series Rack-Mount Server Software Bundle—Only required for Cisco UCS domains that include integrated rack-mount servers. This bundle contains firmware to enable Cisco UCS Manager to manage those servers and is not applicable to standalone C-Series rack-mount servers.
- 2 (Optional) Disable Call Home—If the Cisco UCS domain includes Call Home or Smart Call Home, disable Call Home to ensure you do not receive unnecessary alerts when Cisco UCS Manager restarts components.
- 3 Update adapters, /BMC, and IOMs—If you prefer, you can upgrade the adapters in a host firmware package as part of the last upgrade step.
- 4 Activate adapters—Choose **Ignore Compatibility Check** and **Set Startup Version Only** when performing this step.
- 5 Activate /BMC—Choose **Ignore Compatibility Check** when performing this step.
- 6 Activate Cisco UCS Manager—Choose **Ignore Compatibility Check** when performing this step.
- 7 Activate I/O modules—Choose **Ignore Compatibility Check** and **Set Startup Version Only** when performing this step.
- 8 Activate subordinate fabric interconnect—Choose **Ignore Compatibility Check** when performing this step.
- 9 Activate primary fabric interconnect—Choose **Ignore Compatibility Check** when performing this step.
- 10 Update host firmware package(s) for servers—Must be the last firmware upgraded. We recommend that you upgrade the board controller firmware during this step to avoid an additional reboot of servers with that firmware. You must upgrade the following firmware in a host firmware package:
 - BIOS
 - Storage controller
 - Certain adapters
- 11 (Optional) Enable Call Home—If you disabled Call Home before the upgrading the firmware, enable Call Home.

Summary of Upgrading from Cisco UCS, Release 1.0(1)

- 1 Download the following firmware images:
 - Cisco UCS Infrastructure Software Bundle—Required for all Cisco UCS domains.

- Cisco UCS B-Series Blade Server Software Bundle—Required for all Cisco UCS domains that include blade servers.
 - Cisco UCS C-Series Rack-Mount Server Software Bundle—Only required for Cisco UCS domains that include integrated rack-mount servers. This bundle contains firmware to enable Cisco UCS Manager to manage those servers and is not applicable to standalone C-Series rack-mount servers.
- 2 (Optional) Disable Call Home—If the Cisco UCS domain includes Call Home or Smart Call Home, disable Call Home to ensure you do not receive unnecessary alerts when Cisco UCS Manager restarts components.
 - 3 Update adapters, BMC, and IOMs—If you prefer, you can upgrade the adapters in a host firmware package as part of the last upgrade step.
 - 4 Activate adapters—Choose **Ignore Compatibility Check** and **Set Startup Version Only** when performing this step.
 - 5 BMC—Choose **Ignore Compatibility Check** when performing this step.
 - 6 Activate I/O modules—Choose **Ignore Compatibility Check** and **Set Startup Version Only** when performing this step.
 - 7 Activate subordinate fabric interconnect—Choose **Ignore Compatibility Check** when performing this step.
 - 8 Activate primary fabric interconnect—Choose **Ignore Compatibility Check** when performing this step.
 - 9 Activate Cisco UCS Manager—Choose **Ignore Compatibility Check** when performing this step.
 - 10 Update host firmware package(s) for servers—Must be the last firmware upgraded. We recommend that you upgrade the board controller firmware during this step to avoid an additional reboot of servers with that firmware. You must upgrade the following firmware in a host firmware package:
 - BIOS
 - Storage controller
 - Certain adapters
 - 11 (Optional) Enable Call Home—If you disabled Call Home before the upgrading the firmware, enable Call Home.

Required Order for Adding Support for Previously Unsupported Servers

From Cisco UCS, Release 1.4(1) and later, the method for adding support for previously unsupported type of servers, such as a new blade server or a rack-mount server, to an existing Cisco UCS domain requires the following additional steps after you upgrade your existing firmware to the new release.

Adding Support for a Previously Unsupported Cisco UCS Blade Server

After you upgrade the firmware for the existing components, you can add support for a previously unsupported server that was released between infrastructure bundle releases. When you add the first server of a previously unsupported type of blade server, you must perform the steps to enable Cisco UCS Manager to support that type of server in the following order:

- 1 Insert the blade server into the chassis as described in the server installation guide. Cisco UCS Manager cannot discover the server as it is unsupported, and the finite state machine (FSM) for the discovery fails with an unsupported server error.
- 2 Obtain the B-Series server bundle for the new blade server from Cisco.com and download it to the fabric interconnect.
- 3 Activate the Capability Catalog image from the server bundle.
- 4 Activate the Management Extension from the server bundle.
- 5 Wait for Cisco UCS Manager to retry discovery of the new server. If server discovery does not begin within a few minutes, acknowledge the server.

**Note**

You only need to perform these steps for the first server of a previously unsupported type of blade server. Cisco UCS Manager discovers all subsequent servers of that type automatically.

Integrating Cisco UCS Rack-Mount Servers

After you upgrade the firmware for the existing components, you can integrate one or more Cisco UCS rack-mount servers. When you integrate rack-mount servers, you must perform the steps in the following order:

- 1 If you have not already done so, configure the rack server discovery policy in Cisco UCS Manager.
- 2 Follow the instructions in the appropriate [rack-mount server installation guide](#) for installing and integrating a rack-mount server in a system managed by Cisco UCS Manager.
- 3 Wait for Cisco UCS Manager to discover the new server. If server discovery does not begin within a few minutes, acknowledge the server.

Direct Firmware Upgrade at Endpoints

If you follow the correct procedure and apply the upgrades in the correct order, a direct firmware upgrade and the activation of the new firmware version on the endpoints is minimally disruptive to traffic in a Cisco UCS domain.

You can directly upgrade the firmware on the following endpoints:

- Adapters
- CIMCs
- I/O modules
- Board controllers
- Cisco UCS Manager
- Fabric interconnects

The adapter and board controller firmware can also be upgraded through the host firmware package in the service profile. If you use a host firmware package to upgrade this firmware, you can reduce the number of times a server needs to be rebooted during the firmware upgrade process.

**Note**

Upgrades of a CIMC through a management firmware package or an adapter through a firmware package in the service profile associated with the server take precedence over direct firmware upgrades. You cannot directly upgrade an endpoint if the service profile associated with the server includes a firmware package. To perform a direct upgrade, you must remove the firmware package from the service profile.

Stages of a Direct Firmware Upgrade

Cisco UCS Manager separates the direct upgrade process into two stages to ensure that you can push the firmware to an endpoint while the system is running without affecting uptime on the server or other endpoints.

Update

During this stage, the system copies the selected firmware version from the primary fabric interconnect to the backup partition in the endpoint and verifies that the firmware image is not corrupt. The update process always overwrites the firmware in the backup slot.

The update stage applies only to the following endpoints:

- Adapters
- CIMCs
- I/O modules

**Caution**

Do not remove the hardware that contains the endpoint or perform any maintenance on it until the update process has completed. If the hardware is removed or otherwise unavailable due to maintenance, the firmware update fails. This failure may corrupt the backup partition. You cannot update the firmware on an endpoint with a corrupted backup partition.

Activate

During this stage, the system sets the specified image version (normally the backup version) as the startup version and, if you do not specify **Set Startup Version Only**, immediately reboots the endpoint. When the endpoint is rebooted, the backup partition becomes the active partition, and the active partition becomes the backup partition. The firmware in the new active partition becomes the startup version and the running version.

The following endpoints only require activation because the specified firmware image already exists on the endpoint:

- Cisco UCS Manager
- Fabric interconnects
- Board controllers on those servers that support them

When the firmware is activated, the endpoint is rebooted and the new firmware becomes the active kernel version and system version. If the endpoint cannot boot from the startup firmware, it defaults to the backup version and raises a fault.

**Caution**

When you configure **Set Startup Version Only** for an I/O module, the I/O module is rebooted when the fabric interconnect in its data path is rebooted. If you do not configure **Set Startup Version Only** for an I/O module, the I/O module reboots and disrupts traffic. In addition, if Cisco UCS Manager detects a protocol and firmware version mismatch between the fabric interconnect and the I/O module, Cisco UCS Manager automatically updates the I/O module with the firmware version that matches the firmware in the fabric interconnect and then activates the firmware and reboots the I/O module again.

Outage Impacts of Direct Firmware Upgrades

When you perform a direct firmware upgrade on an endpoint, you can disrupt traffic or cause an outage in one or more of the endpoints in the Cisco UCS domain.

Outage Impact of a Fabric Interconnect Firmware Upgrade

When you upgrade the firmware for a fabric interconnect, you cause the following outage impacts and disruptions:

- The fabric interconnect reboots.
- The corresponding I/O modules reboot.

Outage Impact of a Cisco UCS Manager Firmware Upgrade

A firmware upgrade to Cisco UCS Manager causes the following disruptions:

- Cisco UCS Manager GUI—All users logged in to Cisco UCS Manager GUI are logged out and their sessions ended.
Any unsaved work in progress is lost.
- Cisco UCS Manager CLI—All users logged in through telnet are logged out and their sessions ended.

Outage Impact of an I/O Module Firmware Upgrade

When you upgrade the firmware for an I/O module, you cause the following outage impacts and disruptions:

- For a standalone configuration with a single fabric interconnect, data traffic is disrupted when the I/O module reboots. For a cluster configuration with two fabric interconnects, data traffic fails over to the other I/O module and the fabric interconnect in its data path.
- If you activate the new firmware as the startup version only, the I/O module reboots when the corresponding fabric interconnect is rebooted.
- If you activate the new firmware as the running and startup version, the I/O module reboots immediately.
- An I/O module can take up to ten minutes to become available after a firmware upgrade.

Outage Impact of a CIMC Firmware Upgrade

When you upgrade the firmware for a CIMC in a server, you impact only the CIMC and internal processes. You do not interrupt server traffic. This firmware upgrade causes the following outage impacts and disruptions to the CIMC:

- Any activities being performed on the server through the KVM console and vMedia are interrupted.
- Any monitoring or IPMI polling is interrupted.

Outage Impact of an Adapter Firmware Upgrade

If you activate the firmware for an adapter and do not configure the **Set Startup Version Only** option, you cause the following outage impacts and disruptions:

- The server reboots.
- Server traffic is disrupted.

Firmware Upgrades through Service Profiles

You can use service profiles to upgrade the server and adapter firmware, including the BIOS on the server, by defining the following policies and including them in the service profile associated with a server:

- Host Firmware Package policy
- Management Firmware Package policy



Note

You cannot upgrade the firmware on an I/O module, fabric interconnect, or Cisco UCS Manager through service profiles. You must upgrade the firmware on those endpoints directly.

Host Firmware Package

This policy enables you to specify a set of firmware versions that make up the host firmware package (also known as the host firmware pack). The host firmware includes the following firmware for server and adapter endpoints:

- **Adapter**
- **BIOS**
- **Board Controller**
- **FC Adapters**
- **HBA Option ROM**
- **Storage Controller**



Tip

You can include more than one type of firmware in the same host firmware package. For example, a host firmware package can include both BIOS firmware and storage controller firmware or adapter firmware for two different models of adapters. However, you can only have one firmware version with the same type, vendor, and model number. The system recognizes which firmware version is required for an endpoint and ignores all other firmware versions.

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

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

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

Prerequisites

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

Management Firmware Package

This policy enables you to specify a set of firmware versions that make up the management firmware package (also known as a management firmware pack). The management firmware package includes the Cisco Integrated Management Controller (CIMC) on the server. You do not need to use this package if you upgrade the CIMC directly.

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

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

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

Stages of a Firmware Upgrade through Service Profiles

You can use the host and management firmware package policies in service profiles to upgrade server and adapter firmware.



Caution

If you modify a host firmware package by adding an endpoint or changing firmware versions for an existing endpoint, Cisco UCS Manager upgrades the endpoints and reboots all servers associated with that firmware package as soon as the changes are saved, disrupting data traffic to and from the servers.

New Service Profile

For a new service profile, this upgrade takes place over the following stages:

Firmware Package Policy Creation

During this stage, you create the host and/or management firmware packages and include them in the appropriate firmware policies.

Service Profile Association

During this stage, you include the firmware packages in a service profile, and then associate the service profile with a server. The system pushes the selected firmware versions to the endpoints. For a host firmware package, the server is rebooted to ensure that the endpoints are running the versions specified in the firmware package.

Existing Service Profile

If the service profile is already associated with a server, Cisco UCS Manager upgrades the firmware as soon as you save the changes to the host firmware packages. For a host firmware package, Cisco UCS Manager reboots the server as soon as the change is saved.

Firmware Downgrades

You downgrade firmware in a Cisco UCS domain in the same way that you upgrade firmware. The package or version that you select when you update the firmware determines whether you are performing an upgrade or a downgrade.



Note

The Cisco UCS Manager CLI does not allow you to downgrade hardware that is not supported in the release to which you are downgrading. Cisco UCS Manager CLI displays an error message if you attempt to downgrade hardware to an unsupported release.

Completing the Prerequisites for Upgrading the Firmware

Prerequisites for Upgrading and Downgrading Firmware

All endpoints in a Cisco UCS domain must be fully functional and all processes must be complete before you begin a firmware upgrade or downgrade on those endpoints. You cannot upgrade or downgrade an endpoint that is not in a functional state. For example, the firmware on a server that has not been discovered cannot be upgraded or downgraded. An incomplete process, such as an FSM that has failed after the maximum number of retries, can cause the upgrade or downgrade on an endpoint to fail. If an FSM is in progress, Cisco UCS Manager queues up the update and activation and runs them when the FSM has completed successfully.

Before you upgrade or downgrade firmware in a Cisco UCS domain, complete the following prerequisites:

- Review the Release Notes.
- Back up the configuration into an All Configuration backup file.
- For a cluster configuration, verify that the high availability status of the fabric interconnects shows that both are up and running.
- For a standalone configuration, verify that the Overall Status of the fabric interconnect is Operable.
- Verify that the data path is up and running. For more information, see [Verifying that the Data Path is Ready](#).

- Verify that all servers, I/O modules, and adapters are fully functional. An inoperable server cannot be upgraded.
- Verify that the Cisco UCS domain does not include any critical or major faults. If such faults exist, you must resolve them before you upgrade the system. A critical or major fault may cause the upgrade to fail.
- Verify that all servers have been discovered. They do not need to be powered on or associated with a service profile.
- If you want to integrate a rack-mount server into the Cisco UCS domain, follow the instructions in the appropriate [rack-mount server installation guide](#) for installing and integrating a rack-mount server in a system managed by Cisco UCS Manager.

Creating an All Configuration Backup File

This procedure assumes that you do not have an existing backup operation for an All Configuration backup file.

Before You Begin

Obtain the backup server IP address and authentication credentials.

Procedure

	Command or Action	Purpose
Step 1	UCS-A# scope system	Enters system mode.
Step 2	UCS-A /system # create backup URL all-configuration enabled	Creates an enabled All Configuration backup operation that runs as soon as you enter the commit-buffer command. The all-configuration option backs up the server, fabric, and system related configuration. Specify the URL for the backup file using one of the following syntax: <ul style="list-style-type: none"> • ftp:// username@hostname / path • scp:// username@hostname / path • sftp:// username@hostname / path • tftp:// hostname : port-num / path
Step 3	UCS-A /system # commit-buffer	Commits the transaction.

The following example uses SCP to create an All Configuration backup file on the host named host35 and commits the transaction:

```
UCS-A# scope system
UCS-A /system* # create backup scp://user@host35/backups/all-config.bak all-configuration
enabled
Password:
UCS-A /system* # commit-buffer
UCS-A /system #
```

Verifying the Operability of a Fabric Interconnect

If your Cisco UCS domain is running in a high availability cluster configuration, you must verify the operability of both fabric interconnects.

Procedure

	Command or Action	Purpose
Step 1	UCS-A# scope fabric-interconnect {a b}	Enters fabric interconnect mode for the specified fabric interconnect.
Step 2	UCS-A /fabric-interconnect # show	Displays information about the fabric interconnect. Verify that the operability of the fabric interconnects is in the Operable state. If the operability is not in the Operable state, run a show tech-support command and contact Cisco Technical Support. Do not proceed with the firmware upgrade. For more information about the show tech-support command, see the <i>Cisco UCS Manager B-Series Troubleshooting Guide</i> .

The following example displays that the operability for both fabric interconnects is in the Operable state:

```
UCS-A# scope fabric-interconnect a
UCS-A /fabric-interconnect # show
Fabric Interconnect:
  ID OOB IP Addr      OOB Gateway      OOB Netmask      Operability
  ---
  A  192.168.100.10    192.168.100.20   255.255.255.0    Operable

UCS-A /fabric-interconnect # exit
UCS-A# scope fabric-interconnect b
UCS-A /fabric-interconnect # show
Fabric Interconnect:
  ID OOB IP Addr      OOB Gateway      OOB Netmask      Operability
  ---
  B  192.168.100.11    192.168.100.20   255.255.255.0    Operable
```

Verifying the High Availability Status and Roles of a Cluster Configuration

The high availability status is the same for both fabric interconnects in a cluster configuration.

Procedure

	Command or Action	Purpose
Step 1	UCS-A# show cluster state	Displays the operational state and leadership role for both fabric interconnects in a high availability cluster. Verify that both fabric interconnects (A and B) are in the Up state and HA is in the Ready state. If the fabric interconnects are not in the Up state or HA is not in the Ready state, run a show tech-support command and contact Cisco Technical Support. Do not proceed with the firmware upgrade. For more

	Command or Action	Purpose
		information about the show tech-support command, see the <i>Cisco UCS Troubleshooting Guide</i> . Also note which fabric interconnect has the primary role and which has the subordinate role; you will need to know this information to upgrade the firmware on the fabric interconnects.

The following example displays that both fabric interconnects are in the Up state, HA is in the Ready state, fabric interconnect A has the primary role, and fabric interconnect B has the subordinate role:

```
UCS-A# show cluster state
Cluster Id: 0x4432f72a371511de-0xb97c000de1b1ada4

A: UP, PRIMARY
B: UP, SUBORDINATE

HA READY
```

Verifying the Status of an I/O Module

If your Cisco UCS is running in a high availability cluster configuration, you must verify the status for both I/O modules in all chassis.

Procedure

	Command or Action	Purpose
Step 1	UCS-A# scope chassis <i>chassis-id</i>	Enters chassis mode for the specified chassis.
Step 2	UCS-A /chassis # scope iom <i>iom-id</i>	Enters chassis I/O module mode for the selected I/O module.
Step 3	UCS-A # show	Shows the status of the specified I/O module on the specified chassis. Verify that the overall status of the I/O module is in the Operable state. If the overall status is not in the Operable state, run a show tech-support command and contact Cisco Technical Support. Do not proceed with the firmware upgrade. For more information about the show tech-support command, see the <i>Cisco UCS Troubleshooting Guide</i> .

The following example displays that the overall status for both I/O modules on chassis 1 is in the Operable state:

```
UCS-A# scope chassis 1
UCS-A /chassis # scope iom 1
UCS-A /chassis/iom # show
IOM:
  ID           Side  Fabric ID Overall Status
  -----
-----
```

```

1 Left A Operable
UCS-A /chassis/iom # exit
UCS-A /chassis # scope iom 2
UCS-A /chassis/iom # show
IOM:
-----
ID Side Fabric ID Overall Status
-----
2 Right B Operable
    
```

Verifying the Status of a Server

Procedure

	Command or Action	Purpose
Step 1	UCS-A# scope server <i>chassis-id</i> / <i>server-id</i>	Enters chassis server mode for the specified server in the specified chassis.
Step 2	UCS-A /chassis/server # show status detail	Shows the status detail of the server. Verify that the overall status of the server is Ok, Unavailable, or any value that does not indicate a failure. If the overall status is in a state that indicates a failure, such as Discovery Failed, the endpoints on that server cannot be upgraded.

The following example displays that the overall status for server 7 on chassis 1 is in the Ok state:

```

UCS-A# scope server 1/7
UCS-A /chassis/server # show status detail
Server 1/7:
  Slot Status: Equipped
  Conn Path: A,B
  Conn Status: A,B
  Managing Instance: B
  Availability: Unavailable
  Admin State: In Service
  Overall Status: Ok
  Oper Qualifier: N/A
  Discovery: Complete
  Current Task:
    
```

Verifying the Status of Adapters on Servers in a Chassis

Procedure

	Command or Action	Purpose
Step 1	UCS-A# scope server <i>chassis-id</i> / <i>server-id</i>	Enters chassis server mode for the specified server in the specified chassis
Step 2	UCS-A /chassis/server # show adapter status	Displays the status of the adapter. Verify that the overall status of the adapter is in the Operable state. If the overall status of the adapter is in any state other than Operable, you cannot upgrade it. However, you can

	Command or Action	Purpose
		proceed with the upgrade for the other adapters in the Cisco UCS domain.

The following example displays that the overall status for the adapter in server 7 on chassis 1 is in the Operable state:

```
UCS-A# scope server 1/7
UCS-A /chassis/server # show adapter status
Server 1/1:
  Overall Status
  -----
  Operable
```

Downloading and Managing Firmware Packages

Obtaining Software Bundles from Cisco

Before You Begin

Determine which of the following software bundles you need to update the Cisco UCS domain:

- Cisco UCS Infrastructure Software Bundle—Required for all Cisco UCS domains.
- Cisco UCS B-Series Blade Server Software Bundle—Required for all Cisco UCS domains that include blade servers.
- Cisco UCS C-Series Rack-Mount Server Software Bundle—Only required for Cisco UCS domains that include integrated rack-mount servers. This bundle contains firmware to enable Cisco UCS Manager to manage those servers and is not applicable to standalone C-Series rack-mount servers.

Procedure

- Step 1** In a web browser, navigate to Cisco.com.
- Step 2** Under **Support**, click **All Downloads**.
- Step 3** In the center pane, click **Unified Computing and Servers**.
- Step 4** If prompted, enter your Cisco.com username and password to log in.
- Step 5** In the right pane, click the link for the software bundles you require, as follows:

Bundle	Navigation Path
Cisco UCS Infrastructure Software Bundle	Click Cisco UCS Infrastructure and UCS Manager Software > Unified Computing System (UCS) Infrastructure Software Bundle .
Cisco UCS B-Series Blade Server Software Bundle	Click Cisco UCS B-Series Blade Server Software > Unified Computing System (UCS) Server Software Bundle .

Bundle	Navigation Path
Cisco UCS C-Series Rack-Mount Server Software Bundle	Click Cisco UCS C-Series Rack-Mount UCS-Managed Server Software > Unified Computing System (UCS) Server Software Bundle .

Tip The Unified Computing System (UCS) Documentation Roadmap Bundle, which is accessible through these paths, is a downloadable ISO image of all Cisco UCS documentation.

Step 6 On the first page from which you download a software bundle, click the **Release Notes** link to download the latest version of the Release Notes.

Step 7 For each software bundle that you want to download, do the following:

- a) Click the link for the release you want to download the latest release 2.0 software bundle. The release number is followed by a number and a letter in parentheses. The number identifies the maintenance release level, and the letter differentiates between patches of that maintenance release. For more information about what is in each maintenance release and patch, see the latest version of the Release Notes.
- b) Click one of the following buttons and follow the instructions provided:
 - **Download Now**—Allows you to download the software bundle immediately.
 - **Add to Cart**—Adds the software bundle to your cart to be downloaded at a later time.
- c) Follow the prompts to complete your download of the software bundle(s).

Step 8 Read the Release Notes before upgrading your Cisco UCS domain.

What to Do Next

Download the software bundles to the fabric interconnect.

Downloading Firmware Images to the Fabric Interconnect from a Remote Location



Note In a cluster setup, the image file for the firmware bundle is downloaded to both fabric interconnects, regardless of which fabric interconnect is used to initiate the download. Cisco UCS Manager maintains all firmware packages and images in both fabric interconnects in sync. If one fabric interconnect is down, the download still finishes successfully. The images are synced to the other fabric interconnect when it comes back online.

Before You Begin

Obtain the required firmware bundles from Cisco.

Procedure

	Command or Action	Purpose
Step 1	UCS-A# scope firmware	Enters firmware mode.
Step 2	UCS-A /firmware # download image URL	Downloads the firmware bundle for Cisco UCS. Using the download path provided by Cisco, specify the URL with one of the following syntax: <ul style="list-style-type: none"> • ftp:// server-ip-addr / path • scp:// username@server-ip-addr / path • sftp:// username@server-ip-addr / path • tftp:// server-ip-addr : port-num / path <p>Note TFTP has a file size limitation of 32 MB. Because firmware bundles can be much larger than that, we recommend that you do not select TFTP for firmware downloads.</p> <p>If you use a hostname rather than an IP address, configure a DNS server in Cisco UCS Manager.</p>
Step 3	Enter the password for the remote server.	The password for the remote server username. This field does not apply if the protocol is tftp.
Step 4	UCS-A /firmware # show download-task	Displays the status for your download task. When your image is completely downloaded, the task state changes from Downloading to Downloaded. The CLI does not automatically refresh, so you may have to enter the show download-task command multiple times until the task state displays Downloaded.
Step 5	Repeat this task until all of the firmware bundles have been downloaded to the fabric interconnect.	

The following example uses SCP to download the ucs-k9-bundle.1.0.0.988.gbin firmware package.

```
UCS-A# scope firmware
UCS-A /firmware # download image scp://user1@192.168.10.10/images/ucs-k9-bundle.1.0.0.988.gbin
Password: yourpassword
UCS-A /firmware # show download-task
UCS-A /firmware #
```

What to Do Next

After the image file for the firmware bundles have downloaded completely, update the firmware on the endpoints.

Displaying the Firmware Package Download Status

After a firmware download operation has been started, you can check the download status to see if the package is still downloading or if it has completely downloaded.

Procedure

	Command or Action	Purpose
Step 1	UCS-A# scope firmware	Enters firmware mode.
Step 2	UCS-A /firmware # show download-task	Displays the status for your download task. When your image is completely downloaded, the task state changes from Downloading to Downloaded. The CLI does not automatically refresh, so you may have to enter the show download-task command multiple times until the task state displays Downloaded.

The following example displays the download status for the ucs-k9-bundle.1.0.0.988.gbin firmware package. The **show download-task** command is entered multiple times until the download state indicates that the firmware package has been downloaded:

```
UCS-A# scope firmware
UCS-A /firmware # show download-task

Download task:
File Name Protocol Server      Userid      State
-----
ucs-k9-bundle.1.0.0.988.gbin
      Scp      10.193.32.11  user1      Downloading

UCS-A /firmware # show download-task

Download task:
File Name Protocol Server      Userid      State
-----
ucs-k9-bundle.1.0.0.988.gbin
      Scp      10.193.32.11  user1      Downloading

UCS-A /firmware # show download-task

Download task:
File Name Protocol Server      Userid      State
-----
ucs-k9-bundle.1.0.0.988.gbin
      Scp      10.193.32.11  user1      Downloaded
```

Canceling an Image Download

You can cancel the download task for an image only while it is in progress. After the image has downloaded, deleting the download task does not delete the image that was downloaded. You cannot cancel the FSM related to the image download task.

Procedure

	Command or Action	Purpose
Step 1	UCS-A# scope firmware	Enters firmware mode.
Step 2	UCS-A /firmware # delete download-task task-name	Deletes the specified download task.
Step 3	UCS-A /firmware # commit-buffer	Commits the transaction to the system configuration.

The following example cancels an image download:

```
UCS-A# scope firmware
UCS-A /firmware # delete download-task taskname
UCS-A /firmware* # commit-buffer
UCS-A /firmware* #
```

Displaying All Available Software Images on the Fabric Interconnect

This procedure is optional and displays the available software images on the fabric interconnect for all endpoints. You can also use the **show image** command in each endpoint mode to display the available software images for that endpoint.

Procedure

	Command or Action	Purpose
Step 1	UCS-A# scope firmware	Enters firmware mode.
Step 2	UCS-A /firmware # show image	Displays all software images downloaded onto the fabric interconnect. Note You must provide the software version number when directly updating an endpoint. If you intend to directly update firmware at an endpoint, note its version number in the right column.

The following example displays all available software images on the fabric interconnect:

```
UCS-A# scope firmware
UCS-A /firmware # show image
Name                                                    Type                Version
-----
ucs-2100.1.0.0.988.gbin                                Iom                  1.0 (0.988)
ucs-6100-k9-kickstart.4.0.1a.N2.1.0.988.gbin          Switch Kernel
4.0 (1a)N2 (1.0.988)
ucs-6100-k9-system.4.0.1a.N2.1.0.988.gbin             Switch Software
4.0 (1a)N2 (1.0.988)
ucs-b200-m1-bios.S5500.86B.01.00.0030-978a.021920.gbin Server Bios
S5500.86B.01.00.0030-978a.021920
ucs-b200-m1-k9-bmc.1.0.0.988.gbin                     Bmc                  1.0 (0.988)
ucs-b200-m1-sasctlr.2009.02.09.gbin                   Storage Controller  2009.02.09
ucs-m71kr-e-cna.1.0.0.988.gbin                       Adapter              1.0 (0.988)
ucs-m71kr-e-hba.zf280a4.gbin                         Host Hba              zf280a4
```

```

ucs-m71kr-e-optionrom.ZN502N5.gbin      Host Hba Optionrom  ZN502N5
ucs-m71kr-q-cna.1.0.0.988.gbin          Adapter             1.0(0.988)
ucs-m71kr-q-optionrom.1.69.gbin         Host Hba Optionrom  1.69
ucs-m81kr-vic.1.0.0.988.gbin           Adapter             1.0(0.988)
ucs-manager-k9.1.0.0.988.gbin          System              1.0(0.988)
    
```

Displaying All Available Packages on the Fabric Interconnect

This procedure is optional and displays the available software packages on the fabric interconnect for all endpoints.. You can also use the **show package** command in each endpoint mode to display the available software images for that endpoint.

Procedure

	Command or Action	Purpose
Step 1	UCS-A# scope firmware	Enters firmware mode.
Step 2	UCS-A /firmware # show package	Displays all software packages downloaded onto the fabric interconnect. Note You must provide the software version number when directly updating an endpoint. If you intend to directly update firmware at an endpoint, note its version number in the right column.

The following example displays all available software packages on the fabric interconnect:

```

UCS-A# scope firmware
UCS-A /firmware # show package
Name                                     Version
-----
ucs-k9-bundle.1.3.0.221.bin
ucs-k9-bundle.1.4.0.292.gbin
ucs-k9-bundle.1.4.0.357.gbin
ucs-k9-bundle.1.4.0.378.gbin             1.4(0.378)
ucs-k9-bundle.1.4.0.390.gbin             1.4(0.390)
Pubs-A /firmware #
    
```

Determining the Contents of a Firmware Package

Procedure

	Command or Action	Purpose
Step 1	UCS-A# scope firmware	Enters firmware mode.
Step 2	UCS-A /firmware # show package package-name expand	Displays the contents of the specified firmware package.

The following example displays the contents of a firmware package:

```

UCS-A# scope firmware
UCS-A /firmware # show package ucs-k9-bundle.1.4.0.390.gbin expand
    
```

```

Package ucs-k9-bundle.1.4.0.390.gbin:
  Images:
    ucs-2100.1.4.0.390.gbin
    ucs-6100-k9-kickstart.4.2.1.N1.1.3.390.gbin
    ucs-6100-k9-system.4.2.1.N1.1.3.390.gbin
    ucs-b200-m1-bios.S5500.1.4.0.6.090220101221.gbin
    ucs-b200-m1-k9-cimc.1.4.0.390.gbin
    ucs-b200-m1-sasctrlr.01.28.03.00_06.28.00.00_03.12.00.00.gbin
    ucs-b200-m2-bios.S5500.1.4.0.6.090220101221.gbin
    ucs-b230-m1-bios.B230M1.1.4.0.35.090220101135.gbin
    ucs-b230-m1-k9-cimc.1.4.0.390.gbin
    ucs-b230-m1-mrsasctrlr.20.7.1-0020_4.18.00_NA.gbin
    ucs-b230-m1-pld.B2301008.gbin
    ucs-b250-m1-bios.S5500.1.4.0.6.090220101735.gbin
    ucs-b250-m1-k9-cimc.1.4.0.390.gbin
    ucs-b250-m2-bios.S5500.1.4.0.6.090220101735.gbin
    ucs-b440-m1-bios.B440M1.1.4.0.3.090120101140.gbin
    ucs-b440-m1-k9-cimc.1.4.0.390.gbin
    ucs-b440-m1-mrsasctrlr.12.4.0-0028_3.13.00_NA.gbin
    ucs-b440-m1-pld.B440100C-B4402006.gbin
    ucs-c-pci-n2xx-acpci01.1.4.0.390.gbin
    ucs-c200-bios.C200.1.2.1.3.082520100537.gbin
    ucs-c200-k9-cimc.1.4.0.390.gbin
    ucs-c250-bios.C250.1.2.1.3.082520102328.gbin
    ucs-c250-k9-cimc.1.4.0.390.gbin
    ucs-m51kr-b.5.2.7.12.1.gbin
    ucs-m61kr-i.2.1.60.1.1.gbin
    ucs-m71kr-e-cna.1.4.0.390.gbin
    ucs-m71kr-e-hba.2.80A4.gbin
    ucs-m71kr-e-optionrom.5.03A8.gbin
    ucs-m71kr-q-cna.1.4.0.390.gbin
    ucs-m71kr-q-optionrom.2.02.gbin
    ucs-m72kr-e.2.702.200.1702.gbin
    ucs-m72kr-q.01.02.08.gbin
    ucs-m81kr-vic.1.4.0.390.gbin
    ucs-manager-k9.1.4.0.390.gbin
UCS-A /firmware #
    
```

Checking the Available Space on a Fabric Interconnect

If an image download fails, check whether the bootflash on the fabric interconnect or fabric interconnects in the Cisco UCS has sufficient available space.

Procedure

	Command or Action	Purpose
Step 1	UCS-A# scope fabric-interconnect {a b}	Enters fabric interconnect mode for the specified fabric.
Step 2	UCS-A /fabric-interconnect # show storage [detail expand]	Displays the available space for the specified fabric. Note When you download a firmware image bundle, a fabric interconnect needs at least twice as much available space as the size of the firmware image bundle. If the bootflash does not have sufficient space, delete the obsolete firmware, core files, and other unneeded objects from the fabric interconnect.

The following example displays the available space for a fabric interconnect:

```
UCS-A# scope fabric-interconnect
UCS-A /fabric-interconnect # show storage
Storage on local flash drive of fabric interconnect:
  Partition      Size (MBytes)  Used Percentage
-----
  bootflash     8658           50
  opt            1917           2
  workspace     277            4
UCS-A /fabric-interconnect #
```

Deleting Firmware Packages from a Fabric Interconnect

Use this procedure if you want to delete an entire package. If you prefer, you can also delete only a single image from a package.

Procedure

	Command or Action	Purpose
Step 1	UCS-A# scope firmware	Enters firmware mode.
Step 2	UCS-A /firmware # delete package <i>package-name</i>	Deletes the specified firmware package.
Step 3	UCS-A /firmware # commit-buffer	Commits the transaction to the system configuration.

Cisco UCS Manager deletes the selected package or packages and all images contained within each package.

The following example deletes a firmware package and commits the transaction:

```
UCS-A# scope firmware
UCS-A /firmware # delete image ucs-k9-bundle.1.4.0.433m.gbin
UCS-A /firmware* # commit-buffer
UCS-A /firmware #
```

Deleting Firmware Images from a Fabric Interconnect

Use this procedure if you want to delete only a single image from a package.

Procedure

	Command or Action	Purpose
Step 1	UCS-A# scope firmware	Enters firmware mode.
Step 2	UCS-A /firmware # delete image <i>image-name</i>	Deletes the specified firmware image.
Step 3	UCS-A /firmware # commit-buffer	Commits the transaction to the system configuration.

The following example deletes a firmware image and commits the transaction:

```
UCS-A# scope firmware
UCS-A /firmware # delete image ucs-2100.1.4.0.433k.gbin
UCS-A /firmware* # commit-buffer
UCS-A /firmware #
```

Directly Upgrading Firmware at Endpoints

Updating and Activating the Firmware on an Adapter



Caution

Do not remove the hardware that contains the endpoint or perform any maintenance on it until the update process has completed. If the hardware is removed or otherwise unavailable due to maintenance, the firmware update fails. This failure may corrupt the backup partition. You cannot update the firmware on an endpoint with a corrupted backup partition.

Procedure

	Command or Action	Purpose
Step 1	UCS-A# scope adapter <i>chassis-id / blade-id / adapter-id</i>	Enters chassis server adapter mode for the specified adapter.
Step 2	UCS-A /chassis/server/adapter # show image	Displays the available software images for the adapter.
Step 3	UCS-A /chassis/server/adapter # update firmware <i>version-num</i>	Updates the selected firmware version on the adapter.
Step 4	UCS-A /chassis/server/adapter # commit-buffer	(Optional) Commits the transaction. Use this step only if you intend to use the show firmware command in Step 5 to verify that the firmware update completed successfully before activating the firmware in Step 6. You can skip this step and commit the update-firmware and activate-firmware commands in the same transaction; however, if the firmware update does not complete successfully, the firmware activation does not start. Cisco UCS Manager copies the selected firmware image to the backup memory partition and verifies that image is not corrupt. The image remains as the backup version until you explicitly activate it.
Step 5	UCS-A /chassis/server/adapter # show firmware	(Optional) Displays the status of the firmware update. Use this step only if you want to verify that the firmware update completed successfully. The firmware update is complete when the update status is Ready. The CLI does not automatically refresh, so you may have to enter the show firmware command multiple times

	Command or Action	Purpose
		until the task state changes from Updating to Ready. Continue to Step 6 when the update status is Ready.
Step 6	UCS-A /chassis/server/adapter # activate firmware <i>version-num</i> [ignorecompcheck [set-startup-only] set-startup-only]	Activates the selected firmware version on the adapter. Use the set-startup-only keyword if you want to move the activated firmware into the pending-next-boot state and not immediately reboot the server. The activated firmware does not become the running version of firmware on the adapter until the server is rebooted. You cannot use the set-startup-only keyword for an adapter in the host firmware package. Use the ignorecompcheck keyword if you want to ignore the compatibility check and activate the firmware regardless of any possible incompatibilities or currently executing tasks.
Step 7	UCS-A /chassis/server/adapter # commit-buffer	Commits the transaction. If a server is not associated with a service profile, the activated firmware remains in the pending-next-boot state. Cisco UCS Manager does not reboot the endpoints or activate the firmware until the server is associated with a service profile. If necessary, you can manually reboot or reset an unassociated server to activate the firmware.
Step 8	UCS-A /chassis/server/adapter # show firmware	(Optional) Displays the status of the firmware activation. Use this step only if you want to verify that the firmware activation completed successfully. The CLI does not automatically refresh, so you may have to enter the show firmware command multiple times until the task state changes from Activating to Ready.

The following example updates and activates the adapter firmware to version 1.2(1) in the same transaction, without verifying that the firmware update and firmware activation completed successfully:

```
UCS-A# scope adapter 1/1/1
UCS-A# /chassis/server/adapter # show image
Name                                     Type                               Version                            State
-----
ucs-m81kr-vic.1.2.1.gbin                 Adapter                             1.2(1)                             Active

UCS-A# /chassis/server/adapter # update firmware 1.2(1)
UCS-A# /chassis/server/adapter* # activate firmware 1.2(1) ignorecompcheck set-startup-only
UCS-A# /chassis/server/adapter* # commit-buffer
UCS-A# /chassis/server/adapter #
```

The following example updates the adapter firmware to version 1.2(1), verifies that the firmware update completed successfully before starting the firmware activation, activates the adapter firmware, and verifies that the firmware activation completed successfully:

```
UCS-A# scope adapter 1/1/1
UCS-A# /chassis/server/adapter # show image
Name                                     Type                               Version                            State
-----
```

```

ucs-m81kr-vic.1.2.1.gbin                               Adapter          1.2(1)          Active

UCS-A# /chassis/server/adapter # update firmware 1.2(1)
UCS-A# /chassis/server/adapter* # commit-buffer
UCS-A# /chassis/server/adapter # show firmware
Adapter 1:
  Running-Vers: 1.1(1)
  Update-Status: Updating
  Activate-Status: Ready

UCS-A# /chassis/server/adapter # show firmware
Adapter 1:
  Running-Vers: 1.1(1)
  Update-Status: Ready
  Activate-Status: Ready

UCS-A# /chassis/server/adapter # activate firmware 1.2(1) ignorecompcheck
UCS-A# /chassis/server/adapter* # commit-buffer
UCS-A# /chassis/server/adapter # show firmware
Adapter 1:
  Running-Vers: 1.1(1)
  Update-Status: Ready
  Activate-Status: Activating

UCS-A# /chassis/server/adapter # show firmware
Adapter 1:
  Running-Vers: 1.2(1)
  Update-Status: Ready
  Activate-Status: Ready
    
```

Updating and Activating the BIOS Firmware on a Server



Caution

Do not remove the hardware that contains the endpoint or perform any maintenance on it until the update process has completed. If the hardware is removed or otherwise unavailable due to maintenance, the firmware update fails. This failure may corrupt the backup partition. You cannot update the firmware on an endpoint with a corrupted backup partition.

Procedure

	Command or Action	Purpose
Step 1	UCS-A# scope server <i>chassis-id / blade-id</i>	Enters chassis server mode for the specified server.
Step 2	UCS-A /chassis/server # scope bios	Enters chassis server BIOS mode.
Step 3	UCS-A /chassis/server/bios # show image	Displays the available BIOS firmware images.
Step 4	UCS-A /chassis/server/bios # update firmware <i>version-num</i>	Updates the selected BIOS firmware for the server.
Step 5	UCS-A /chassis/server/bios # commit-buffer	(Optional) Commits the transaction. Use this step only if you intend to use the show firmware command in Step 6 to verify that the firmware update completed successfully before activating the firmware in Step 7. You can

	Command or Action	Purpose
		<p>skip this step and commit the update-firmware and activate-firmware commands in the same transaction; however, if the firmware update does not complete successfully, the firmware activation does not start.</p> <p>Cisco UCS Manager copies the selected firmware image to the backup memory partition and verifies that image is not corrupt. The image remains as the backup version until you explicitly activate it.</p>
Step 6	UCS-A /chassis/server/bios # show firmware	<p>(Optional) Displays the status of the firmware update.</p> <p>Use this step only if you want to verify that the firmware update completed successfully. The firmware update is complete when the update status is Ready. The CLI does not automatically refresh, so you may have to enter the show firmware command multiple times until the task state changes from Updating to Ready. Continue to Step 7 when the update status is Ready.</p>
Step 7	UCS-A /chassis/server/bios # activate firmware <i>version-num</i> [ignorecompcheck]	<p>Activates the selected server BIOS firmware version.</p> <p>Use the ignorecompcheck keyword if you want to ignore the compatibility check and activate the firmware regardless of any possible incompatibilities or currently executing tasks.</p>
Step 8	UCS-A /chassis/server/bios # commit-buffer	Commits the transaction.
Step 9	UCS-A /chassis/bios # show firmware	<p>(Optional) Displays the status of the firmware activation.</p> <p>Use this step only if you want to verify that the firmware activation completed successfully. The CLI does not automatically refresh, so you may have to enter the show firmware command multiple times until the task state changes from Activating to Ready.</p>

The following example updates and activates the BIOS firmware in the same transaction, without verifying that the firmware update and activation completed successfully:

```
UCS-A# scope server 1/1
UCS-A# /chassis/server # scope bios
UCS-A# /chassis/server/bios # show image
Name                                     Type          Version
-----
ucs-b230-m1-bios.B230.2.0.1.1.49.gbin  Server Bios  B230.2.0.1.1.49
ucs-b230-m1-bios.B230.2.0.2.0.00.gbin  Server Bios  B230.2.0.2.0.00

UCS-A# /chassis/server/bios # update firmware B230.2.0.2.0.00
UCS-A# /chassis/server/bios* # activate firmware B230.2.0.2.0.00 ignorecompcheck
UCS-A# /chassis/server/bios* # commit-buffer
UCS-A# /chassis/server/bios #
```

Updating and Activating the CIMC Firmware on a Server

The activation of firmware for a CIMC does not disrupt data traffic. However, it will interrupt all KVM sessions and disconnect any vMedia attached to the server.



Caution

Do not remove the hardware that contains the endpoint or perform any maintenance on it until the update process has completed. If the hardware is removed or otherwise unavailable due to maintenance, the firmware update fails. This failure may corrupt the backup partition. You cannot update the firmware on an endpoint with a corrupted backup partition.

Procedure

	Command or Action	Purpose
Step 1	UCS-A# scope server <i>chassis-id / blade-id</i>	Enters chassis server mode for the specified server.
Step 2	UCS-A /chassis/server # scope cimc	Enters chassis server CIMC mode.
Step 3	UCS-A /chassis/server/cimc # show image	Displays the available software images for the adapter.
Step 4	UCS-A /chassis/server/cimc # update firmware <i>version-num</i>	Updates the selected firmware version on the CIMC in the server.
Step 5	UCS-A /chassis/server/cimc # commit-buffer	(Optional) Commits the transaction. Use this step only if you intend to use the show firmware command in Step 6 to verify that the firmware update completed successfully before activating the firmware in Step 7. You can skip this step and commit the update-firmware and activate-firmware commands in the same transaction; however, if the firmware update does not complete successfully, the firmware activation does not start. Cisco UCS Manager copies the selected firmware image to the backup memory partition and verifies that image is not corrupt. The image remains as the backup version until you explicitly activate it.
Step 6	UCS-A /chassis/server/cimc # show firmware	(Optional) Displays the status of the firmware update. Use this step only if you want to verify that the firmware update completed successfully. The firmware update is complete when the update status is Ready. The CLI does not automatically refresh, so you may have to enter the show firmware command multiple times until the task state changes from Updating to Ready. Continue to Step 7 when the update status is Ready.

	Command or Action	Purpose
Step 7	UCS-A /chassis/server/cimc # activate firmware <i>version-num</i> [ignorecompcheck]	Activates the selected firmware version on the CIMC in the server. Use the ignorecompcheck keyword if you want to ignore the compatibility check and activate the firmware regardless of any possible incompatibilities or currently executing tasks.
Step 8	UCS-A /chassis/server/cimc # commit-buffer	Commits the transaction.
Step 9	UCS-A /chassis/server/cimc # show firmware	(Optional) Displays the status of the firmware activation. Use this step only if you want to verify that the firmware activation completed successfully. The CLI does not automatically refresh, so you may have to enter the show firmware command multiple times until the task state changes from Activating to Ready.

The following example updates and activates the CIMC firmware to version 1.2(1) in the same transaction, without verifying that the firmware update and firmware activation completed successfully:

```
UCS-A# scope server 1/1
UCS-A# /chassis/server # scope cimc
UCS-A# /chassis/server/cimc # show image
Name                                     Type                               Version                               State
-----
ucs-b200-m1-k9-cimc.1.2.1.gbin          Bmc                                1.2 (1)
Active

UCS-A# /chassis/server/cimc # update firmware 1.2(1)
UCS-A# /chassis/server/cimc* # activate firmware 1.2(1) ignorecompcheck set-startup-only
UCS-A# /chassis/server/cimc* # commit-buffer
UCS-A# /chassis/server/cimc #
```

The following example updates the CIMC firmware to version 1.2(1), verifies that the firmware update completed successfully before starting the firmware activation, activates the CIMC firmware, and verifies that the firmware activation completed successfully:

```
UCS-A# scope server 1/1
UCS-A# /chassis/server # scope cimc
UCS-A# /chassis/server/cimc # show image
Name                                     Type                               Version                               State
-----
ucs-b200-m1-k9-cimc.1.2.1.gbin          Bmc                                1.2 (1)
Active

UCS-A# /chassis/server/cimc # update firmware 1.2(1)
UCS-A# /chassis/server/cimc* # commit-buffer
UCS-A# /chassis/server/cimc # show firmware
Running-Vers  Update-Status  Activate-Status
-----
1.1 (1)       Updating         Ready

UCS-A# /chassis/server/cimc # show firmware
Running-Vers  Update-Status  Activate-Status
```

```

-----
1.1(1)          Ready          Ready

UCS-A# /chassis/server/cimc # activate firmware 1.2(1) ignorecompcheck
UCS-A# /chassis/server/cimc* # commit-buffer
UCS-A# /chassis/server/cimc # show firmware
Running-Vers   Update-Status   Activate-Status
-----
1.1(1)          Ready          Activating

UCS-A# /chassis/server/cimc # show firmware
Running-Vers   Update-Status   Activate-Status
-----
1.2(1)          Ready          Ready

```

Updating and Activating the Firmware on an IOM

If your system is running in a high availability cluster configuration, you must update and activate both I/O modules.



Caution

Do not remove the hardware that contains the endpoint or perform any maintenance on it until the update process has completed. If the hardware is removed or otherwise unavailable due to maintenance, the firmware update fails. This failure may corrupt the backup partition. You cannot update the firmware on an endpoint with a corrupted backup partition.

Procedure

	Command or Action	Purpose
Step 1	UCS-A# scope chassis <i>chassis-id</i>	Enters chassis mode for the specified chassis.
Step 2	UCS-A /chassis # scope iom <i>iom-id</i>	Enters chassis I/O module mode for the selected I/O module.
Step 3	UCS-A /chassis/iom # show image	Displays the available software images for the I/O module.
Step 4	UCS-A /chassis/iom # update firmware <i>version-num</i>	Updates the selected firmware version on the I/O module.
Step 5	UCS-A /chassis/iom # commit-buffer	(Optional) Commits the transaction. Use this step only if you intend to use the show firmware command in Step 6 to verify that the firmware update completed successfully before activating the firmware in Step 7. You can skip this step and commit the update-firmware and activate-firmware commands in the same transaction; however, if the firmware update does not complete successfully, the firmware activation does not start. Cisco UCS Manager copies the selected firmware image to the backup memory partition and verifies that image is not corrupt. The image remains as the backup version until you explicitly activate it.

	Command or Action	Purpose
Step 6	UCS-A /chassis/iom # show firmware	(Optional) Displays the status of the firmware update. Use this step only if you want to verify that the firmware update completed successfully. The firmware update is complete when the update status is Ready. The CLI does not automatically refresh, so you may have to enter the show firmware command multiple times until the task state changes from Updating to Ready. Continue to Step 7 when the update status is Ready.
Step 7	UCS-A /chassis/iom # activate firmware <i>version-num</i> [ignorecompcheck [set-startup-only] set-startup-only]	Activates the selected firmware version on the I/O module. Use the set-startup-only keyword if you want to reboot the I/O module only when the fabric interconnect in its data path reboots. If you do not use the set-startup-only keyword, the I/O module reboots and disrupts traffic. In addition, if Cisco UCS Manager detects a protocol and firmware version mismatch between it and the I/O module, it updates the I/O module with the firmware version that matches its own and then activates the firmware and reboots the I/O module again. Use the ignorecompcheck keyword if you want to ignore the compatibility check and activate the firmware regardless of any possible incompatibilities or currently executing tasks.
Step 8	UCS-A /chassis/iom # commit-buffer	Commits the transaction.
Step 9	UCS-A /chassis/iom # show firmware	(Optional) Displays the status of the firmware activation. Use this step only if you want to verify that the firmware activation completed successfully. The CLI does not automatically refresh, so you may have to enter the show firmware command multiple times until the task state changes from Activating to Ready.

The following example updates and activates the I/O module firmware to version 1.2(1) in the same transaction, without verifying that the firmware update and firmware activation completed successfully:

```
UCS-A# scope chassis 1
UCS-A# /chassis # scope iom 1
UCS-A# /chassis/iom # show image
Name                                     Type                Version             State
-----
ucs-2100.1.2.1.gbin                     Iom                 1.2(1)             Active

UCS-A# /chassis/iom # update firmware 1.2(1)
UCS-A# /chassis/iom* # activate firmware 1.2(1) ignorecompcheck set-startup-only
UCS-A# /chassis/iom* # commit-buffer
UCS-A# /chassis/iom #
```

The following example updates the I/O module firmware to version 1.2(1), verifies that the firmware update completed successfully before starting the firmware activation, activates the I/O module firmware, and verifies that the firmware activation completed successfully:

```
UCS-A# scope chassis 1
UCS-A# /chassis # scope iom 1
UCS-A# /chassis/iom # show image
-----
Name                               Type                               Version                               State
-----
ucs-2100.1.2.1.gbin                Iom                                1.2(1)                               Active

UCS-A# /chassis/iom # update firmware 1.2(1)
UCS-A# /chassis/iom* # commit-buffer
UCS-A# /chassis/iom # show firmware
IOM      Fabric ID Running-Vers    Update-Status    Activate-Status
-----
1 A      1.1(1)                        Updating         Ready

UCS-A# /chassis/iom # show firmware
IOM      Fabric ID Running-Vers    Update-Status    Activate-Status
-----
1 A      1.1(1)                        Ready            Ready

UCS-A# /chassis/iom # activate firmware 1.2(1) ignorecompcheck
UCS-A# /chassis/iom* # commit-buffer
UCS-A# /chassis/iom # show firmware
IOM      Fabric ID Running-Vers    Update-Status    Activate-Status
-----
1 A      1.1(1)                        Ready            Activating

UCS-A# /chassis/iom # show firmware
IOM      Fabric ID Running-Vers    Update-Status    Activate-Status
-----
1 A      1.2(1)                        Ready            Ready
```

Activating the Board Controller Firmware on a Server

Only certain servers, such as the Cisco UCS B440 High Performance blade server and the Cisco UCS B230 blade server, have board controller firmware. The board controller firmware controls many of the server functions, including eUSBs, LEDs, and I/O connectors.



Note

This activation procedure causes the server to reboot. Depending upon whether or not the service profile associated with the server includes a maintenance policy, the reboot can occur immediately. To reduce the number of times a server needs to be rebooted during the upgrade process, we recommend that you upgrade the board controller firmware through the host firmware package in the service profile as the last step of upgrading a Cisco UCS domain, along with the server BIOS.

Procedure

	Command or Action	Purpose
Step 1	UCS-A# scope server <i>chassis-id / server-id</i>	Enters chassis server mode for the specified server.
Step 2	UCS-A /chassis/server # scope boardcontroller	Enters board controller mode for the server.

	Command or Action	Purpose
Step 3	UCS-A /chassis/server/boardcontroller # show image	(Optional) Displays the available software images for the board controller.
Step 4	UCS-A /chassis/server/boardcontroller # show firmware	(Optional) Displays the current running software image for the board controller.
Step 5	UCS-A /chassis/server/boardcontroller # activate firmware <i>version-num</i> [ignorecompcheck]	Activates the selected firmware version on the board controller in the server. Use the ignorecompcheck keyword if you want to ignore the compatibility check and activate the firmware regardless of any possible incompatibilities or currently executing tasks.
Step 6	UCS-A /chassis/server/boardcontroller # commit-buffer	Commits the transaction to the system configuration.

Cisco UCS Manager disconnects all active sessions, logs out all users, and activates the software. When the upgrade is complete, you are prompted to log back in. If you are prompted to re-login immediately after being disconnected, the login will fail. You must wait until the activation of Cisco UCS Manager is completed, which takes a few minutes.

The following example activates the board controller firmware:

```
UCS-A# scope server 1/1
UCS-A# /chassis/server # scope boardcontroller
UCS-A# /chassis/server/boardcontroller # show image
Name                                     Type                               Version                             State
-----
ucs-b440-m1-pld.B440100C-B4402006.bin  Board Controller                   B440100C-B4402006                   Active

UCS-A# /chassis/server/boardcontroller # show firmware
BoardController:
  Running-Vers: B440100C-B4402006
  Activate-Status: Ready

UCS-A# /chassis/server/boardcontroller # activate firmware B440100C-B4402006
ignorecompcheck
UCS-A# /chassis/server/boardcontroller* # commit-buffer
```

Activating the Cisco UCS Manager Software

Procedure

	Command or Action	Purpose
Step 1	UCS-A# scope system	Enters system mode.

	Command or Action	Purpose
Step 2	UCS-A /system # show image	Displays the available software images for Cisco UCS Manager (system).
Step 3	UCS-A /system # activate firmware <i>version-num</i> [ignorecompcheck]	Activates the selected firmware version on the system. Use the ignorecompcheck keyword if you want to ignore the compatibility check and activate the firmware regardless of any possible incompatibilities or currently executing tasks. Note Activating Cisco UCS Manager does not require rebooting the fabric interconnect; however, management services will briefly go down and all VSH shells will be terminated as part of the activation.
Step 4	UCS-A /system # commit-buffer	Commits the transaction. Cisco UCS Manager makes the selected version the startup version and schedules the activation to occur when the fabric interconnects are upgraded.

The following example upgrades Cisco UCS Manager to version 1.2(1) and commits the transaction:

```
UCS-A# scope system
UCS-A# /system # show image
Name                                     Type           Version        State
-----
ucs-manager-k9.1.2.1.gbin                System         1.2(1)        Active

UCS-A# /system # activate firmware 1.2(1)
UCS-A# /system* # commit-buffer
UCS-A# /system #
```

Activating the Firmware on a Fabric Interconnect

When updating the firmware on two fabric interconnects in a high availability cluster configuration, you must activate the subordinate fabric interconnect before activating the primary fabric interconnect. For more information about determining the role for each fabric interconnect, see [Verifying the High Availability Status and Roles of a Cluster Configuration](#), on page 203.

For a standalone configuration with a single fabric interconnect, you can minimize the disruption to data traffic when you perform a direct firmware upgrade of the endpoints. However, you must reboot the fabric interconnect to complete the upgrade and, therefore, cannot avoid disrupting traffic.



Tip

If you ever need to recover the password to the admin account that was created when you configured the fabric interconnects for the Cisco UCS domain, you must know the running kernel version and the running system version. If you do not plan to create additional accounts, we recommend that you save the path to these firmware versions in a text file so that you can access them if required.

Procedure

	Command or Action	Purpose
Step 1	UCS-A# scope fabric-interconnect {a b}	Enters fabric interconnect mode for the specified fabric interconnect.
Step 2	UCS-A /fabric-interconnect # show image	Displays the available software images for the fabric interconnect.
Step 3	UCS-A /fabric-interconnect # activate firmware {kernel-version <i>kernel-ver-num</i> system-version <i>system-ver-num</i> }	Activates the selected firmware version on the fabric interconnect.
Step 4	UCS-A /fabric-interconnect # commit-buffer	Commits the transaction. Cisco UCS Manager updates and activates the firmware, and then reboots the fabric interconnect and any I/O module in the data path to that fabric interconnect, disrupting data traffic to and from that fabric interconnect.

The following example upgrades the fabric interconnect to version 4.0(1a)N2(1.2.1) and commits the transaction:

```
UCS-A# scope fabric-interconnect a
UCS-A /fabric-interconnect # show image
Name                                                    Type                Version              State
-----
ucs-6100-k9-kickstart.4.0.1a.N2.1.2.1.gbin           Fabric Interconnect  4.0(1a)N2(1.2.1)    Active
ucs-6100-k9-system.4.0.1a.N2.1.2.1.gbin              Fabric Interconnect  4.0(1a)N2(1.2.1)    Active

UCS-A /fabric-interconnect # activate firmware kernel-version 4.0(1a)N2(1.2.1) system-version
4.0(1a)N2(1.2.1)
UCS-A /fabric-interconnect* # commit-buffer
UCS-A /fabric-interconnect #
```

Updating Firmware through Service Profiles

Host Firmware Package

This policy enables you to specify a set of firmware versions that make up the host firmware package (also known as the host firmware pack). The host firmware includes the following firmware for server and adapter endpoints:

- **Adapter**
- **BIOS**
- **Board Controller**
- **FC Adapters**

- **HBA Option ROM**
- **Storage Controller**

**Tip**

You can include more than one type of firmware in the same host firmware package. For example, a host firmware package can include both BIOS firmware and storage controller firmware or adapter firmware for two different models of adapters. However, you can only have one firmware version with the same type, vendor, and model number. The system recognizes which firmware version is required for an endpoint and ignores all other firmware versions.

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

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

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

Prerequisites

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

Management Firmware Package

This policy enables you to specify a set of firmware versions that make up the management firmware package (also known as a management firmware pack). The management firmware package includes the Cisco Integrated Management Controller (CIMC) on the server. You do not need to use this package if you upgrade the CIMC directly.

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

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

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

Effect of Updates to Host Firmware Packages and Management Firmware Packages

To update firmware through a host firmware package or a management firmware package, you need to update the firmware in the package. What happens after you save the changes to a host or management firmware package depends upon how the Cisco UCS domain is configured.

The following table describes the most common options for upgrading servers with a host or management firmware package.

Service Profile	Maintenance Policy	Upgrade Actions
<p>Host or management firmware package is not included in a service profile or an updating service profile template.</p> <p>OR</p> <p>You want to upgrade the firmware without making any changes to the existing service profile or updating service profile template.</p>	<p>No maintenance policy</p>	<p>After you update the firmware package, do one of the following:</p> <ul style="list-style-type: none"> • To reboot and upgrade some or all servers simultaneously, follow the procedure in the Cisco UCS Manager configuration guides for the appropriate release to add the firmware package to one or more service profiles that are associated with servers or to an updating service profile template. • To reboot and upgrade one server at a time, do the following for each server: <ol style="list-style-type: none"> 1 Create a new service profile and include the firmware package in that service profile. 2 Dissociate the server from its service profile. 3 Associate the server with the new service profile. 4 After the server has been rebooted and the firmware upgraded, disassociate the server from the new service profile and associate it with its original service profile. <p>Caution If the original service profile includes a scrub policy, this procedure may result in data loss when the disk or the BIOS is scrubbed upon association with the new service profile.</p>
<p>Host or management firmware package is included in one or more service profiles, and the service profiles are associated with one or more servers.</p> <p>OR</p> <p>Host or management firmware package is included in an updating service profile template, and the service profiles created from that template are associated with one or more servers.</p>	<p>No maintenance policy</p> <p>OR</p> <p>A maintenance policy configured for immediate updates.</p>	<p>The following occurs when you update the firmware package:</p> <ol style="list-style-type: none"> 1 The changes to the firmware package take effect as soon as you save them. 2 Cisco UCS Manager verifies the model numbers and vendor against all servers associated with service profiles that include this policy. If the model numbers and vendor match a firmware version in the policy, Cisco UCS Manager reboots the servers and updates the firmware. <p>All servers associated with service profiles that include the firmware package are rebooted at the same time.</p>

Service Profile	Maintenance Policy	Upgrade Actions
<p>Host or management firmware package is included in one or more service profiles, and the service profiles are associated with one or more servers.</p> <p>OR</p> <p>Host or management firmware package is included in an updating service profile template, and the service profiles created from that template are associated with one or more servers.</p>	<p>Configured for user acknowledgment</p>	<p>The following occurs when you update the firmware package:</p> <ol style="list-style-type: none"> 1 Cisco UCS Manager asks you to confirm your change and advises that a user-acknowledged reboot of the servers is required. 2 Click the flashing Pending Activities button to select the servers you want to reboot and apply the new firmware. 3 Cisco UCS Manager verifies the model numbers and vendor against all servers associated with service profiles that include this policy. If the model numbers and vendor match a firmware version in the policy, Cisco UCS Manager reboots the server and updates the firmware. <p>A manual reboot of the servers does not cause Cisco UCS Manager to apply the firmware package, nor does it cancel the pending activities. You must acknowledge or cancel the pending activity through the Pending Activities button.</p>
<p>Host or management firmware package is included in one or more service profiles, and the service profiles are associated with one or more servers.</p> <p>OR</p> <p>Host or management firmware package is included in an updating service profile template, and the service profiles created from that template are associated with one or more servers.</p>	<p>Configured for changes to take effect during a specific maintenance window.</p>	<p>The following occurs when you update the firmware package:</p> <ol style="list-style-type: none"> 1 Cisco UCS Manager asks you to confirm your change and advises that a user-acknowledged reboot of the servers is required. 2 Click the flashing Pending Activities button to select the servers you want to reboot and apply the new firmware. 3 Cisco UCS Manager verifies the model numbers and vendor against all servers associated with service profiles that include this policy. If the model numbers and vendor match a firmware version in the policy, Cisco UCS Manager reboots the server and updates the firmware. <p>A manual reboot of the servers does not cause Cisco UCS Manager to apply the firmware package, nor does it cancel the scheduled maintenance activities.</p>

Creating or Updating a Host Firmware Package

If the policy is included in one or more service profiles associated with a server and those service profiles do not include maintenance policies, Cisco UCS Manager updates and activates the firmware in the server and adapter with the new versions and reboots the server as soon as you save the host firmware package policy.



Tip You can include more than one type of firmware in the same host firmware package. For example, a host firmware package can include both BIOS firmware and storage controller firmware or adapter firmware for two different models of adapters. However, you can only have one firmware version with the same type, vendor, and model number. The system recognizes which firmware version is required for an endpoint and ignores all other firmware versions.

Before You Begin

Ensure that the appropriate firmware has been downloaded to the fabric interconnect.

Procedure

	Command or Action	Purpose
Step 1	UCS-A# scope org <i>org-name</i>	Enters organization mode for the specified organization. To enter the root organization mode, type / as the <i>org-name</i> .
Step 2	UCS-A org/ # create fw-host-pack <i>pack-name</i>	Creates a host firmware package with the specified package name and enters organization firmware host package mode.
Step 3	UCS-A /org/fw-host-pack # set descr <i>description</i>	(Optional) Provides a description for the host firmware package. Note If your description includes spaces, special characters, or punctuation, you must begin and end your description with quotation marks. The quotation marks will not appear in the description field of any show command output.
Step 4	UCS-A org/fw-host-pack # create pack-image <i>hw-vendor-name</i> <i>hw-model</i> { adapter host-hba host-hba-combined host-hba-optionrom host-nic server-bios storage-controller unspecified } <i>version-num</i>	Creates a package image for the host firmware package and enters organization firmware host package image mode. The <i>hw-vendor-name</i> and <i>hw-model</i> values are labels that help you easily identify the package image when you enter the show image detail command. The <i>version-num</i> value specifies the version number of the firmware being used for the package image. The model and model number (PID) must match the servers that are associated with this firmware package. If you select the wrong model or model number, Cisco UCS Manager cannot install the firmware update.
Step 5	UCS-A org/fw-host-pack/pack-image # set version <i>version-num</i>	(Optional) Specifies the package image version number. Changing this number triggers firmware updates on all components using the firmware through a service profile. Use this step only when updating a host firmware package, not when creating a package.

	Command or Action	Purpose
		Note The host firmware package can contain multiple package images. Repeat Step 4 , and Step 5 , to create additional package images for other components.
Step 6	UCS-A org/fw-host-pack/pack-image # commit-buffer	Commits the transaction. Cisco UCS Manager verifies the model numbers and vendor against all servers associated with service profiles that include this policy. If the model numbers and vendor match a firmware version in the policy, Cisco UCS Manager updates the firmware according to the settings in the maintenance policies included in the service profiles.

The following example creates the app1 host firmware package, creates a storage controller package image with version 2009.02.09 firmware, and commits the transaction:

```
UCS-A# scope org /
UCS-A /org # create fw-host-pack app1
UCS-A /org/fw-host-pack* # set descr "This is a host firmware package example."
UCS-A /org/fw-host-pack* # create pack-image Cisco UCS storage-controller 2009.02.09
UCS-A /org/fw-host-pack/pack-image* # commit-buffer
UCS-A /org/fw-host-pack/pack-image #
```

What to Do Next

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

Creating or Updating a Management Firmware Package



Caution

If the policy is included in one or more service profiles associated with a server and those service profiles do not include maintenance policies, Cisco UCS Manager updates and activates the management firmware in the server with the new versions and reboots the server as soon as you save the management firmware package policy

Before You Begin

Ensure that the appropriate firmware has been downloaded to the fabric interconnect.

Procedure

	Command or Action	Purpose
Step 1	UCS-A# scope org <i>org-name</i>	Enters organization mode for the specified organization. To enter the root organization mode, type <i>/</i> as the <i>org-name</i> .
Step 2	UCS-A org/ # create fw-mgmt-pack <i>pack-name</i>	Creates a management firmware package with the specified package name and enters organization firmware management package mode.

	Command or Action	Purpose
Step 3	UCS-A /org/fw-mgmt-pack # set descr <i>description</i>	(Optional) Provides a description for the management firmware package. Note If your description includes spaces, special characters, or punctuation, you must begin and end your description with quotation marks. The quotation marks will not appear in the description field of any show command output.
Step 4	UCS-A org/fw-mgmt-pack # create pack-image <i>hw-vendor-name hw-model</i> bmc <i>version-num</i>	Creates a package image for the management firmware package and enters organization firmware management package image mode. The <i>hw-vendor-name</i> and <i>hw-model</i> values are labels that help you easily identify the package image when you enter the show image detail command. The <i>version-num</i> value specifies the version number of the firmware being used for the package image. The model and model number (PID) must match the servers that are associated with this firmware package. If you select the wrong model or model number, Cisco UCS Manager cannot install the firmware update.
Step 5	UCS-A org/fw-mgmt-pack/pack-image # set version <i>version-num</i>	(Optional) Specifies the package image version number. Changing this number triggers firmware updates on all components using the firmware through a service profile. Use this step only when updating a firmware package, not when creating a package.
Step 6	UCS-A org/fw-mgmt-pack/pack-image # commit-buffer	Commits the transaction. Cisco UCS Manager verifies the model numbers and vendor against all servers associated with service profiles that include this policy. If the model numbers and vendor match a firmware version in the policy, Cisco UCS Manager updates the firmware according to the settings in the maintenance policies included in the service profiles.

The following example creates the cimc1 host firmware package, creates a CIMC package image with version 1.0(0.988) firmware, and commits the transaction:

```
UCS-A# scope org /
UCS-A /org # create fw-mgmt-pack cimc1
UCS-A /org/fw-mgmt-pack* # set descr "This is a management firmware package example."
UCS-A /org/fw-mgmt-pack* # create pack-image Cisco UCS cimc 1.0(0.988)
UCS-A /org/fw-mgmt-pack/pack-image* # commit-buffer
UCS-A /org/fw-mgmt-pack/pack-image #
```

What to Do Next

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

Managing the Capability Catalog

Capability Catalog

The Capability Catalog is a set of tunable parameters, strings, and rules. Cisco UCS Manager uses the catalog to update the display and configurability of components such as newly qualified DIMMs and disk drives for servers.

The catalog is divided by hardware components, such as the chassis, CPU, local disk, and I/O module. You can use the catalog to view the list of providers available for that component. There is one provider per hardware component. Each provider is identified by the vendor, model (PID), and revision. For each provider, you can also view details of the equipment manufacturer and the form factor.

For information about which hardware components are dependent upon a particular catalog release, see the component support tables in the [Service Notes for the B- Series servers](#). For information about which components are introduced in a specific release, see the Cisco UCS [Release Notes](#).

Contents of the Capability Catalog

The contents of the Capability Catalog include the following:

Implementation-Specific Tunable Parameters

- Power and thermal constraints
- Slot ranges and numbering
- Adapter capacities

Hardware-Specific Rules

- Firmware compatibility for components such as the BIOS, CIMC, RAID controller, and adapters
- Diagnostics
- Hardware-specific reboot

User Display Strings

- Part numbers, such as the CPN, PID/VID
- Component descriptions
- Physical layout/dimensions
- OEM information

Updates to the Capability Catalog

Capability Catalog updates are included in each Cisco UCS Manager update. Unless otherwise instructed by Cisco Technical Support, you only need to activate the Capability Catalog update after you've downloaded, updated, and activated an Cisco UCS Infrastructure Software Bundle.

As soon as you activate a Capability Catalog update, Cisco UCS Manager immediately updates to the new baseline catalog. You do not have to perform any further tasks. Updates to the Capability Catalog do not require you to reboot any component in the Cisco UCS domain or to reinstall Cisco UCS Manager.

Each Cisco UCS Manager release contains a baseline catalog. In rare circumstances, Cisco releases an update to the Capability Catalog and makes it available on the same site where you download firmware images. The catalog update is compatible with Cisco UCS, Release 1.3(1) and later.



Note

The Capability Catalog version is determined by the version of Cisco UCS Manager that you are using. Cisco UCS Manager 2.0 releases work with any 2.0 release of the Capability Catalog, but not with any 1.0 releases of the Capability Catalog. For information about Capability Catalog releases supported by specific Cisco UCS Manager releases, see the *Release Notes for Cisco UCS Manager* accessible through the *Cisco UCS B-Series Servers Documentation Roadmap* available at the following URL: <http://www.cisco.com/go/unifiedcomputing/b-series-doc>.

Activating a Capability Catalog Update

Procedure

	Command or Action	Purpose
Step 1	UCS-A# scope system	Enters system mode.
Step 2	UCS-A /system # scope capability	Enters system capability mode.
Step 3	UCS-A /system/capability # activate firmware firmware-version	Activates the specified Capability Catalog version.
Step 4	UCS-A /system/capability # commit-buffer	Commits the transaction to the system configuration.

The following example activates a Capability Catalog update and commits the transaction:

```
UCS-A# scope system
UCS-A /system # scope capability
UCS-A /system/capability # activate firmware 1.0(3)
UCS-A /system/capability* # commit-buffer
UCS-A /system/capability #
```

Verifying that the Capability Catalog is Current

Procedure

	Command or Action	Purpose
Step 1	UCS-A# scope system	Enters system mode.
Step 2	UCS-A /system # scope capability	Enters system capability mode.
Step 3	UCS-A /system/capability # show version	Displays the current Capability Catalog version.
Step 4	On Cisco.com , determine the most recent release of the Capability Catalog available.	For more information about the location of Capability Catalog updates, see Obtaining Capability Catalog Updates from Cisco, on page 237 .
Step 5	If a more recent version of the Capability Catalog is available on Cisco.com, update the Capability Catalog with that version.	

The following example displays the current Capability Catalog version:

```
UCS-A# scope system
UCS-A /system # scope capability
UCS-A /system/capability # show version
Catalog:
  Running-Vers: 1.0 (8.35)
  Activate-Status: Ready
UCS-A /system/capability #
```

Restarting a Capability Catalog Update

You can restart a failed Capability Catalog file update, modifying the update parameters if necessary.

Procedure

	Command or Action	Purpose
Step 1	UCS-A# scope system	Enters system command mode.
Step 2	UCS-A /system # scope capability	Enters capability command mode.
Step 3	UCS-A /system/capability # show cat-updater [filename]	(Optional) Displays the update history for Capability Catalog file update operations.
Step 4	UCS-A /system/capability # scope cat-updater filename	Enters the command mode for the Capability Catalog file update operation.
Step 5	UCS-A /system/capability/cat-updater # set userid username	(Optional) Specifies the username for the remote server.

	Command or Action	Purpose
Step 6	UCS-A /system/capability/cat-updater # set password <i>password</i>	(Optional) Specifies the password for the remote server username. If no password is configured, you are prompted for a password when you start the update.
Step 7	UCS-A /system/capability/cat-updater # set protocol {ftp scp sftp tftp}	(Optional) Specifies the file transfer protocol for the remote server. Note TFTP has a file size limitation of 32 MB. Because catalog images can be much larger than that, we recommend that you do not use TFTP for catalog image downloads.
Step 8	UCS-A /system/capability/cat-updater # set server {hostname ip-address}	(Optional) Specifies the hostname or IP address of the remote server.
Step 9	UCS-A /system/capability/cat-updater # set path <i>pathname/filename</i>	(Optional) Specifies the path and file name of the Capability Catalog file on the remote server.
Step 10	UCS-A /system/capability/cat-updater # restart	Restarts the Capability Catalog file update operation.

The following example changes the server IP address and restarts the Capability Catalog file update operation:

```
UCS-A# scope system
UCS-A /system # scope capability
UCS-A /system/capability # show cat-updater

Catalog Updater:
  File Name Protocol Server          Userid          Status
  -----
  ucs-catalog.1.0.0.4.bin
    Scp          192.0.2.111    user1          Failed

UCS-A /system/capability # scope cat-updater ucs-catalog.1.0.0.4.bin
UCS-A /system/capability/cat-updater # set server 192.0.2.112
UCS-A /system/capability/cat-updater # restart
UCS-A /system/capability/cat-updater #
```

Viewing a Capability Catalog Provider

Procedure

	Command or Action	Purpose
Step 1	UCS-A# scope system	Enters system command mode.
Step 2	UCS-A /system # scope capability	Enters capability command mode.
Step 3	UCS-A /system/capability # show {chassis cpu disk fan fru iom	Displays vendor, model, and revision information for all components in the specified component category.

	Command or Action	Purpose
	<code>memory psu server} [vendor model revision] [detail expand]</code>	To view manufacturing and form factor details for a specific component, specify the <i>vendor</i> , <i>model</i> , and <i>revision</i> with the expand keyword. If any of these fields contains spaces, you must enclose the field with quotation marks.



Note

If the server contains one or more SATA devices, such as a hard disk drive or solid state drive, the **show disk** command displays ATA in the Vendor field. Use the **expand** keyword to display additional vendor information.

The following example lists the installed fans and displays detailed information from the Capability Catalog about a specific fan:

```
UCS-A# scope system
UCS-A /system # scope capability
UCS-A /system/capability # show fan
```

```
Fan Module:
Vendor                Model                Revision
-----
Cisco Systems, Inc.   N10-FAN1             0
Cisco Systems, Inc.   N10-FAN2             0
Cisco Systems, Inc.   N20-FAN5             0
```

```
UCS-A /system/capability # show fan "Cisco Systems, Inc." N10-FAN1 0 expand
```

```
Fan Module:
Vendor: Cisco Systems, Inc.
Model: N10-FAN1
Revision: 0

Equipment Manufacturing:
Name: Fan Module for UCS 6140 Fabric Interconnect
PID: N10-FAN1
VID: NA
Caption: Fan Module for UCS 6140 Fabric Interconnect
Part Number: N10-FAN1
SKU: N10-FAN1
CLEI:
Equipment Type:

Form Factor:
Depth (C): 6.700000
Height (C): 1.600000
Width (C): 4.900000
Weight (C): 1.500000
```

```
UCS-A /system/capability #
```

Downloading Individual Capability Catalog Updates

Obtaining Capability Catalog Updates from Cisco

Procedure

-
- Step 1** In a web browser, navigate to Cisco.com.
- Step 2** Under **Support**, click **All Downloads**.
- Step 3** In the center pane, click **Unified Computing and Servers**.
- Step 4** If prompted, enter your Cisco.com username and password to log in.
- Step 5** In the right pane, click **Cisco UCS Infrastructure and UCS Manager Software > Unified Computing System (UCS) Manager Capability Catalog**.
- Step 6** Click the link for the latest release of the Capability Catalog.
- Step 7** Click one of the following buttons and follow the instructions provided:
- **Download Now**—Allows you to download the catalog update immediately
 - **Add to Cart**—Adds the catalog update to your cart to be downloaded at a later time
- Step 8** Follow the prompts to complete your download of the catalog update.
-

What to Do Next

Update the Capability Catalog.

Updating the Capability Catalog from a Remote Location

You cannot perform a partial update to the Capability Catalog. When you update the Capability Catalog, all components included in the catalog image are updated.

Procedure

	Command or Action	Purpose
Step 1	UCS-A# scope system	Enters system command mode.
Step 2	UCS-A /system # scope capability	Enters capability command mode.
Step 3	UCS-A /system/capability # update catalog URL	Imports and applies the specified Capability Catalog file. Specify the URL for the operation using one of the following syntax: <ul style="list-style-type: none"> • ftp:// username@hostname / path • scp:// username@hostname / path • sftp:// username@hostname / path • tftp:// hostname : port-num / path

	Command or Action	Purpose
		When a username is specified, you are prompted for a password.
Step 4	UCS-A /system/capability # show version	(Optional) Displays the catalog update version.
Step 5	UCS-A /system/capability # show cat-updater [filename]	(Optional) Displays the update history for a Capability Catalog file, if specified, or for all Capability Catalog file update operations.

Cisco UCS Manager downloads the image and updates the Capability Catalog. You do not need to reboot any hardware components.

The following example uses SCP to import a Capability Catalog file:

```
UCS-A# scope system
UCS-A /system # scope capability
UCS-A /system/capability # update catalog
scp://user1@192.0.2.111/catalogs/ucs-catalog.1.0.0.4.bin
Password:
UCS-A /system/capability # show version
Catalog:
    Update Version: 1.0(0.4)

UCS-A /system/capability # show cat-updater

Catalog Updater:
  File Name Protocol Server          Userid          Status
  -----
  ucs-catalog.1.0.0.4.bin
           Scp          192.0.2.111    user1          Success

UCS-A /system/capability #
```

Updating Management Extensions

Management Extensions

Management Extension updates are included in each Cisco UCS Manager update. Unless otherwise instructed by Cisco Technical Support, you only need to activate the Management Extension update after you've downloaded, updated, and activated an Cisco UCS Infrastructure Software Bundle.

Management Extensions enable you to add support for previously unsupported servers and other hardware to Cisco UCS Manager. For example, you may need to activate a Management Extension if you want to add a new, previously unsupported server to an existing Cisco UCS domain.

The Management Extension image contains the images, information, and firmware required by Cisco UCS Manager to be able to manage the new hardware.

Cisco UCS Manager may need to access a Management Extension when you activate. Therefore, the Management Extension is locked during the activation and update process.

Activating a Management Extension

The Management Extension is included in the server bundle that you have already downloaded. You do not need to download the Management Extension separately.

To verify the Management Extension version, issue the **show version** command.

Procedure

	Command or Action	Purpose
Step 1	UCS-A# scope system	Enters system mode.
Step 2	UCS-A /system # scope management-extension	Enters system Management Extension mode.
Step 3	UCS-A /system/management-extension # activate firmware <i>firmware-version</i> [force-activation]	Activates the specified Management Extension. Use the force-activation keyword to activate the firmware regardless of any possible incompatibilities or currently executing tasks.
Step 4	UCS-A /system/management-extension # commit-buffer	Commits the transaction to the system configuration.

The following example activates the Management Extension and commits the transaction:

```
UCS-A# scope system
UCS-A /system # scope management-extension
UCS-A /system/management-extension # activate firmware 1.0(4)
CS-A /system/management-extension* # commit-buffer
```




CHAPTER 12

Configuring DNS Servers

This chapter includes the following sections:

- [DNS Servers in Cisco UCS](#), page 241
- [Configuring a DNS Server](#), page 241
- [Deleting a DNS Server](#), page 242

DNS Servers in Cisco UCS

You need to specify an external DNS server for each Cisco UCS domain to use if the system requires name resolution of hostnames. For example, you cannot use a name such as `www.cisco.com` when you are configuring a setting on a fabric interconnect if you do not configure a DNS server. You would need to use the IP address of the server.

Configuring a DNS Server

Procedure

	Command or Action	Purpose
Step 1	UCS-A# scope system	Enters system mode.
Step 2	UCS-A /system # scope services	Enters system services mode.
Step 3	UCS-A /system/services # create dns <i>ip-addr</i>	Configures the system to use the DNS server with the specified IP address.
Step 4	UCS-A /system/services # commit-buffer	Commits the transaction to the system configuration.

The following example configures a DNS server with the IP address 192.168.200.105 and commits the transaction:

```
UCS-A# scope system
UCS-A /system # scope services
UCS-A /system/services # create dns 192.168.200.105
UCS-A /system/services* # commit-buffer
UCS-A /system/services #
```

Deleting a DNS Server

Procedure

	Command or Action	Purpose
Step 1	UCS-A# scope system	Enters system mode.
Step 2	UCS-A /system # scope services	Enters system services mode.
Step 3	UCS-A /system/services # delete dns <i>ip-addr</i>	Deletes the NTP server with the specified IP address.
Step 4	UCS-A /system/services # commit-buffer	Commits the transaction to the system configuration.

The following example deletes the DNS server with the IP address 192.168.200.105 and commits the transaction:

```
UCS-A# scope system
UCS-A /system # scope services
UCS-A /system/services # delete dns 192.168.200.105
UCS-A /system/services* # commit-buffer
UCS-A /system/services #
```



CHAPTER 13

Configuring System-Related Policies

This chapter includes the following sections:

- [Configuring the Chassis Discovery Policy, page 243](#)
- [Configuring the Chassis Connectivity Policy, page 247](#)
- [Configuring the Rack Server Discovery Policy, page 248](#)
- [Configuring the Aging Time for the MAC Address Table, page 250](#)

Configuring the Chassis Discovery Policy

Chassis Discovery Policy

The chassis discovery policy determines how the system reacts when you add a new chassis. Cisco UCS Manager uses the settings in the chassis discovery policy to determine the minimum threshold for the number of links between the chassis and the fabric interconnect and whether to group links from the IOM to the fabric interconnect in a fabric port channel.

Chassis Links

If you have a Cisco UCS domain that has some chassis wired with 1 link, some with 2 links, some with 4 links, and some with 8 links we recommend that you configure the chassis discovery policy for the minimum number links in the domain so that Cisco UCS Manager can discover all chassis.



Tip

For Cisco UCS implementations that mix IOMs with different numbers of links, we recommend using the platform max value. Using platform max insures that Cisco UCS Manager uses the maximum number of IOM uplinks available.

After the initial discovery, you must reacknowledge the chassis that are wired for a greater number of links and Cisco UCS Manager configures the chassis to use all available links.

Cisco UCS Manager cannot discover any chassis that is wired for fewer links than are configured in the chassis discovery policy. For example, if the chassis discovery policy is configured for 4 links, Cisco UCS Manager

cannot discover any chassis that is wired for 1 link or 2 links. Reacknowledgement of the chassis does not resolve this issue.

The following table provides an overview of how the chassis discovery policy works in a multi-chassis Cisco UCS domain:

Table 12: Chassis Discovery Policy and Chassis Links

Number of Links Wired for the Chassis	1-Link Chassis Discovery Policy	2-Link Chassis Discovery Policy	4-Link Chassis Discovery Policy	8-Link Chassis Discovery Policy	Platform-Max Discovery Policy
1 link between IOM and fabric interconnects	Chassis is discovered by Cisco UCS Manager and added to the Cisco UCS domain as a chassis wired with 1 link.	Chassis cannot be discovered by Cisco UCS Manager and is not added to the Cisco UCS domain.	Chassis cannot be discovered by Cisco UCS Manager and is not added to the Cisco UCS domain.	Chassis cannot be discovered by Cisco UCS Manager and is not added to the Cisco UCS domain.	Chassis is discovered by Cisco UCS Manager and added to the Cisco UCS domain as a chassis wired with 1 link.
2 links between IOM and fabric interconnects	Chassis is discovered by Cisco UCS Manager and added to the Cisco UCS domain as a chassis wired with 1 link. After initial discovery, reacknowledge the chassis and Cisco UCS Manager recognizes and uses the additional links.	Chassis is discovered by Cisco UCS Manager and added to the Cisco UCS domain as a chassis wired with 2 link.	Chassis cannot be discovered by Cisco UCS Manager and is not added to the Cisco UCS domain.	Chassis cannot be discovered by Cisco UCS Manager and is not added to the Cisco UCS domain.	Chassis cannot be discovered by Cisco UCS Manager and is not added to the Cisco UCS domain.

Number of Links Wired for the Chassis	1-Link Chassis Discovery Policy	2-Link Chassis Discovery Policy	4-Link Chassis Discovery Policy	8-Link Chassis Discovery Policy	Platform-Max Discovery Policy
4 links between IOM and fabric interconnects	Chassis is discovered by Cisco UCS Manager and added to the Cisco UCS domain as a chassis wired with 1 link. After initial discovery, reacknowledge the chassis and Cisco UCS Manager recognizes and uses the additional links.	Chassis is discovered by Cisco UCS Manager and added to the Cisco UCS domain as a chassis wired with 2 links. After initial discovery, reacknowledge the chassis and Cisco UCS Manager recognizes and uses the additional links.	Chassis is discovered by Cisco UCS Manager and added to the Cisco UCS domain as a chassis wired with 4 link.	Chassis cannot be discovered by Cisco UCS Manager and is not added to the Cisco UCS domain.	If the IOM has 4 links, the chassis is discovered by Cisco UCS Manager and added to the Cisco UCS domain as a chassis wired with 4 links. If the IOM has 8 links, the chassis is not fully discovered by Cisco UCS Manager.
8 links between IOM and fabric interconnects	Chassis is discovered by Cisco UCS Manager and added to the Cisco UCS domain as a chassis wired with 1 link. After initial discovery, reacknowledge the chassis and Cisco UCS Manager recognizes and uses the additional links.	Chassis is discovered by Cisco UCS Manager and added to the Cisco UCS domain as a chassis wired with 2 links. After initial discovery, reacknowledge the chassis and Cisco UCS Manager recognizes and uses the additional links.	Chassis is discovered by Cisco UCS Manager and added to the Cisco UCS domain as a chassis wired with 4 links. After initial discovery, reacknowledge the chassis and Cisco UCS Manager recognizes and uses the additional links.	Chassis is discovered by Cisco UCS Manager and added to the Cisco UCS domain as a chassis wired with 8 links.	Chassis is discovered by Cisco UCS Manager and added to the Cisco UCS domain as a chassis wired with 8 links.

Link Grouping

For hardware configurations that support fabric port channels, link grouping determines whether all of the links from the IOM to the fabric interconnect are grouped into a fabric port channel during chassis discovery. If the link grouping preference is set to port channel, all of the links from the IOM to the fabric interconnect

are grouped in a fabric port channel. If set to no group, links from the IOM to the fabric interconnect are not grouped in a fabric port channel.

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



Note The link grouping preference only takes effect if both sides of the links between an IOM or FEX and the fabric interconnect support fabric port channels. If one side of the links does not support fabric port channels, this preference is ignored and the links are not grouped in a port channel.

Configuring the Chassis Discovery Policy

Procedure

	Command or Action	Purpose
Step 1	UCS-A# scope org /	Enters the root organization mode. Note The chassis discovery policy can be accessed only from the root organization.
Step 2	UCS-A /org # scope chassis-disc-policy	Enters organization chassis discovery policy mode.
Step 3	UCS-A /org/chassis-disc-policy # set action {1-link 2-link 4-link 8-link platform-max}	Specifies the minimum threshold for the number of links between the chassis and the fabric interconnect.
Step 4	UCS-A /org/chassis-disc-policy # set descr <i>description</i>	(Optional) Provides a description for the chassis discovery policy. Note If your description includes spaces, special characters, or punctuation, you must begin and end your description with quotation marks. The quotation marks will not appear in the description field of any show command output.
Step 5	UCS-A /org/chassis-disc-policy # set link-aggregation-pref {none port-channel}	Specifies whether the links from the IOMs to the fabric interconnects are grouped in a port channel. Note The link grouping preference only takes effect if both sides of the links between an IOM or FEX and the fabric interconnect support fabric port channels. If one side of the links does not support fabric port channels, this preference is ignored and the links are not grouped in a port channel.
Step 6	UCS-A /org/chassis-disc-policy # set qualifier <i>qualifier</i>	(Optional) Uses the specified server pool policy qualifications to associate this policy with a server pool.
Step 7	UCS-A /org/chassis-disc-policy # commit-buffer	Commits the transaction to the system configuration.

The following example scopes to the default chassis discovery policy, sets it to discover chassis with four links to a fabric interconnect, provides a description for the policy, specifies the server pool policy qualifications that will be used to qualify the chassis, and commits the transaction:

```
UCS-A# scope org /
UCS-A /org # scope chassis-disc-policy
UCS-A /org/chassis-disc-policy* # set action 4-link
UCS-A /org/chassis-disc-policy* # set descr "This is an example chassis discovery policy."
UCS-A /org/chassis-disc-policy* # set qualifier ExampleQual
UCS-A /org/chassis-disc-policy* # commit-buffer
UCS-A /org/chassis-disc-policy #
```

The following example scopes to the default chassis discovery policy, sets it to discover chassis with eight links to a fabric interconnect, provides a description for the policy, sets the link grouping preference to port channel, specifies the server pool policy qualifications that will be used to qualify the chassis, and commits the transaction:

```
UCS-A# scope org /
UCS-A /org # scope chassis-disc-policy
UCS-A /org/chassis-disc-policy* # set action 8-link
UCS-A /org/chassis-disc-policy* # set descr "This is an example chassis discovery policy."
UCS-A /org/chassis-disc-policy* # set link-aggregation-pref port-channel
UCS-A /org/chassis-disc-policy* # set qualifier ExampleQual
UCS-A /org/chassis-disc-policy* # commit-buffer
UCS-A /org/chassis-disc-policy #
```

What to Do Next

To customize fabric port channel connectivity for a specific chassis, configure the chassis connectivity policy.

Configuring the Chassis Connectivity Policy

Chassis Connectivity Policy

The chassis connectivity policy determines the whether a specific chassis is included in a fabric port channel after chassis discovery. This policy is helpful for users who want to configure one or more chassis differently from what is specified in the global chassis discovery policy. The chassis connectivity policy also allows for different connectivity modes per fabric interconnect, further expanding the level of control offered with regards to chassis connectivity.

By default, the chassis connectivity policy is set to global. This means that connectivity control is configured when the chassis is newly discovered, using the settings configured in the chassis discovery policy. Once the chassis is discovered, the chassis connectivity policy controls whether the connectivity control is set to none or port channel.



Note

The chassis connectivity policy is created by Cisco UCS Manager only when the hardware configuration supports fabric port channels. At this time, only the 6200 series fabric interconnects and the 2200 series IOMs support this feature. For all other hardware combinations, Cisco UCS Manager does not create a chassis connectivity policy.

Configuring a Chassis Connectivity Policy

Changing the connectivity mode for a chassis could result in decreased VIF namespace.



Caution

Changing the connectivity mode for a chassis results in chassis reacknowledgement. Traffic may be disrupted during this time.

Procedure

	Command or Action	Purpose
Step 1	UCS-A# scope org org-name	Enters organization mode for the specified organization. To enter the root organization mode, type / as the <i>org-name</i> .
Step 2	UCS-A /org # scope chassis-conn-policy chassis-num [a b]	Enters chassis connection policy organization mode for the specified chassis and fabric.
Step 3	UCS-A /org/chassis-conn-policy # set link-aggregation-pref {global none port-channel}	Specifies whether the links from the IOMs to the fabric interconnects are grouped in a port channel. <ul style="list-style-type: none"> • None—No links are grouped in a port channel • Port Channel—All links from an IOM to a fabric interconnect are grouped in a port channel. • Global—The chassis inherits this configuration from the chassis discovery policy. This is the default value.
Step 4	UCS-A /org/chassis-conn-policy # commit-buffer	Commits the transaction to the system configuration.

The following example shows how to change the fabric port channel connectivity for two chassis. Chassis 6, fabric A is changed to port channel and chassis 12, fabric B is changed to discrete links:

```
UCS-A# scope org /
UCS-A /org # scope chassis-conn-policy 6 a
UCS-A /org/chassis-conn-policy # set link-aggregation-pref port-channel
UCS-A /org/chassis-conn-policy* # up
UCS-A /org* # scope chassis-conn-policy 12 b
UCS-A /org/chassis-conn-policy* # set link-aggregation-pref none
UCS-A /org/chassis-conn-policy* # commit-buffer
UCS-A /org/chassis-conn-policy #
```

Configuring the Rack Server Discovery Policy

Rack Server Discovery Policy

The rack server discovery policy determines how the system reacts when you add a new rack-mount server. Cisco UCS Manager uses the settings in the rack server discovery policy to determine whether any data on

the hard disks are scrubbed and whether server discovery occurs immediately or needs to wait for explicit user acknowledgement.

Cisco UCS Manager cannot discover any rack-mount server that has not been correctly cabled and connected to the fabric interconnects. For information about how to integrate a supported Cisco UCS rack-mount server with Cisco UCS Manager, see the hardware installation guide for that server.

Configuring the Rack Server Discovery Policy

Procedure

	Command or Action	Purpose
Step 1	UCS-A# scope org /	Enters the root organization mode. Note The rack server discovery policy can be accessed only from the root organization.
Step 2	UCS-A /org # scope rackserver-disc-policy	Enters organization rack server discovery policy mode.
Step 3	UCS-A /org/rackserver-disc-policy # set action {immediate user-acknowledged}	Specifies the way the system reacts when you add a new rack server.
Step 4	UCS-A /org/rackserver-disc-policy # set descr <i>description</i>	(Optional) Provides a description for the rack server discovery policy. Note If your description includes spaces, special characters, or punctuation, you must begin and end your description with quotation marks. The quotation marks will not appear in the description field of any show command output.
Step 5	UCS-A /org/rackserver-disc-policy # set scrub-policy <i>scrub-pol-name</i>	Specifies the scrub policy that should run on a newly discovered rack server.
Step 6	UCS-A /org/rackserver-disc-policy # commit-buffer	Commits the transaction to the system configuration.

The following example scopes to the default rack server discovery policy, sets it to immediately discover new rack servers, provides a description for the policy, specifies a scrub policy called scrubpoll, and commits the transaction:

```
UCS-A# scope org /
UCS-A /org # scope rackserver-disc-policy
UCS-A /org/rackserver-disc-policy* # set action immediate
UCS-A /org/rackserver-disc-policy* # set descr "This is an example rackserver discovery
policy."
UCS-A /org/rackserver-disc-policy* # set scrub-policy scrubpoll
UCS-A /org/rackserver-disc-policy* # commit-buffer
UCS-A /org/rackserver-disc-policy #
```

Configuring the Aging Time for the MAC Address Table

Aging Time for the MAC Address Table

To efficiently switch packets between ports, the fabric interconnect maintains a MAC address table. It dynamically builds the MAC address table by using the MAC source address from the packets received and the associated port on which the packets were learned. The fabric interconnect uses an aging mechanism, defined by a configurable aging timer, to determine how long an entry remains in the MAC address table. If an address remains inactive for a specified number of seconds, it is removed from the MAC address table.

You can configure the amount of time (age) that a MAC address entry (MAC address and associated port) remains in the MAC address table.

Configuring the Aging Time for the MAC Address Table

Procedure

	Command or Action	Purpose
Step 1	UCS-A# scope eth-uplink	Enters Ethernet uplink mode.
Step 2	UCS-A /eth-uplink # set mac-aging { <i>dd hh mm ss</i> mode-default never }	Specifies the aging time for the MAC address table. Use the mode-default keyword to set the aging time to a default value dependent on the configured Ethernet switching mode. Use the never keyword to never remove MAC addresses from the table regardless of how long they have been idle.
Step 3	UCS-A /eth-uplink # commit-buffer	Commits the transaction to the system configuration.

The following example sets the aging time for the MAC address table to one day and 12 hours and commits the transaction:

```
UCS-A# scope eth-uplink
UCS-A /eth-uplink # set mac-aging 01 12 00 00
UCS-A /eth-uplink* # commit-buffer
UCS-A /eth-uplink #
```



CHAPTER 14

Managing Licenses

This chapter includes the following sections:

- [Licenses, page 251](#)
- [Obtaining the Host ID for a Fabric Interconnect, page 252](#)
- [Obtaining a License, page 253](#)
- [Installing a License, page 254](#)
- [Viewing the Licenses Installed on a Fabric Interconnect, page 254](#)
- [Viewing License Usage for a Fabric Interconnect, page 255](#)
- [Uninstalling a License, page 257](#)

Licenses

Each Cisco UCS fabric interconnect comes with several port licenses that are factory installed and shipped with the hardware. Fabric interconnects can be purchased fully licensed or partially licensed. Additional licenses can also be purchased after delivery.

At a minimum, each fabric interconnect ships with the following counted licenses pre-installed:

- Cisco UCS 6120XP fabric interconnect—pre-installed licenses for the first eight Ethernet ports enabled in Cisco UCS Manager and any Fibre Channel ports on expansion modules
- Cisco UCS 6140XP fabric interconnect—pre-installed licenses for the first sixteen Ethernet ports enabled in Cisco UCS Manager and any Fibre Channel ports on expansion modules
- Cisco UCS 6248 fabric interconnect—pre-installed licenses for the first twelve unified ports enabled in Cisco UCS Manager. Expansion modules come with eight licenses that can be used on the expansion module or the base module.
- Cisco UCS 6296 fabric interconnect—pre-installed licenses for the first eighteen unified ports enabled in Cisco UCS Manager. Expansion modules come with eight licenses that can be used on the expansion module or the base module.

**Note**

The eight default licenses that come with a 6200 series fabric interconnect expansion module can be used to enable ports on the base module, but will travel with the expansion module if it is removed. Upon removal of an expansion module, any default expansion module licenses being used by the base module are removed from the ports on the base module, resulting in unlicensed ports.

Port licenses are not bound to physical ports. When you disable a licensed port, that license is then retained for use with the next enabled port. If you want to use additional fixed ports, you must purchase and install licenses for those ports.

**Important**

Licenses are not portable across product generations. Licenses purchased for 6100 series fabric interconnects cannot be used to enable ports on 6200 series fabric interconnects or vice-versa.

Grace Period

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

**Note**

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

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

High Availability Configurations

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

Obtaining the Host ID for a Fabric Interconnect

The host ID is also known as the serial number.

Procedure

	Command or Action	Purpose
Step 1	UCS-A# scope license	Enters license mode.
Step 2	UCS-A /license # show server-host-id	Obtains the host ID or serial number for the fabric interconnect. Tip Use the entire host ID that displays after the equal (=) sign.

	Command or Action	Purpose
--	-------------------	---------

The following example obtains the host ID for a fabric interconnect:

```
UCS-A# scope license
UCS-A /license # show server-host-id
Server host id:
  Scope Host Id
  -----
  A      VDH=SSII2121212
  B      VDH=SSII3131313
UCS-A /license #
```

What to Do Next

Obtain the required licenses from Cisco.

Obtaining a License



Note

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

Before You Begin

Obtain the following:

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

Procedure

- Step 1** Obtain the product authorization key (PAK) from the claim certificate or other proof of purchase document.
- Step 2** Locate the website URL in the claim certificate or proof of purchase document.
- Step 3** Access the website URL for the fabric interconnect and enter the serial number and the PAK. Cisco sends you the license file by email. The license file is digitally signed to authorize use on only the requested fabric interconnect. The requested features are also enabled once Cisco UCS Manager accesses the license file.

What to Do Next

Install the license on the fabric interconnect.

Installing a License


Note

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

Before You Begin

Obtain the required licenses from Cisco.

Procedure

	Command or Action	Purpose
Step 1	UCS-A# scope license	Enters license mode.
Step 2	UCS-A /license # download license <i>from-filesystem</i>	Downloads the license from its source location. For the <i>from-filesystem:</i> argument, use one of the following syntaxes: <ul style="list-style-type: none"> • ftp:// <i>server-ip-addr</i> • scp:// <i>username@server-ip-addr</i> • sftp:// <i>username@server-ip-addr</i> • tftp:// <i>server-ip-addr : port-num</i>
Step 3	UCS-A /license # install file <i>license_filename</i>	Installs the license.

The following example uses FTP to download and install a license:

```
UCS-A # scope license
UCS-A /license # download license ftp://192.168.10.10/license/port9.lic
UCS-A /license # install file port9.lic
UCS-A /license #
```

Viewing the Licenses Installed on a Fabric Interconnect

Procedure

	Command or Action	Purpose
Step 1	UCS-A# scope license	Enters license mode.

	Command or Action	Purpose
Step 2	UCS-A /license # show file [<i>license_filename</i> detail]	Displays the licenses installed on the fabric interconnect with the level of detail specified in the command.

The following example displays the full details for the licenses installed on a fabric interconnect:

```
UCS-A# scope license
UCS-A /license # show file detail

License file: UCSFEAT20100928112305377.lic
  Id: 1212121212121212
  Version: 1.0
  Scope: A
  State: Installed
  Features
  Feature Name: ETH_PORT_ACTIVATION_PKG
  Vendor: cisco
  Version: 1.0
  Quantity: 24
  Lines
    Line Id: 1
    Type: Increment
    Expiry Date: Never
    Pak:
    Quantity: 24
    Signature: B10101010101

License file: UCSFEAT20100928112332175.lic
  Id: 1313131313131313
  Version: 1.0
  Scope: B
  State: Installed
  Features
  Feature Name: ETH_PORT_ACTIVATION_PKG
  Vendor: cisco
  Version: 1.0
  Quantity: 24
  Lines
    Line Id: 1
    Type: Increment
    Expiry Date: Never
    Pak:
    Quantity: 24
    Signature: F302020202020

UCS-A /license #
```

Viewing License Usage for a Fabric Interconnect

Procedure

	Command or Action	Purpose
Step 1	UCS-A# scope license	Enters license mode.

	Command or Action	Purpose
Step 2	UCS-A /license # show usage	<p>Displays the license usage table for all license files installed on the fabric interconnect.</p> <p>This following are included:</p> <ul style="list-style-type: none"> • Feat Name The name of the feature to which the license applies. • Scope The fabric associated with the license. • Default The default number of licenses provided for this Cisco UCS domain. • Total Quant The total number of licenses available. This value is the sum of the number of default licenses plus the number of purchased licenses. • Used Quant The number of licenses currently being used by the system. If this value exceeds the total number of licenses available, then some ports will stop functioning after their associated grace period expires. • State The operational state of the license. • Peer License Count Comparison The number of licenses on the peer fabric interconnect compared to this fabric interconnect. This can be one of the following: <ul style="list-style-type: none"> • exceeds—the peer fabric interconnect has more licenses installed than this fabric interconnect • lacks—the peer fabric interconnect has fewer licenses installed than this fabric interconnect • matching—the same number of licenses are installed on both fabric interconnects

	Command or Action	Purpose
		<ul style="list-style-type: none"> • Grace Used The number of grace period days that this license has used. After the grace period ends, Cisco UCS sends alert messages until a new license is purchased.

The following example displays full details of the licenses installed on a fabric interconnect:

```
UCS-A# scope license
UCS-A /license # show usage
Feat Name                               Scope Default Total Quant Used Quant State      Peer Count
Comparison    Grace Used
-----
ETH_PORT_ACTIVATION_PKG                 A             16           40           11 License Ok Matching
0
ETH_PORT_ACTIVATION_PKG                 B             16           40           11 License Ok Matching
0
UCS-A /license #
```

Uninstalling a License



Note

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

Before You Begin

- Back up the Cisco UCS Manager configuration.
- Disable the feature or port associated with the license you want to uninstall.

Procedure

	Command or Action	Purpose
Step 1	UCS-A# scope license	Enters license mode.
Step 2	UCS-A /license # clear file license-filename	Uninstalls the specified license.

Cisco UCS Manager deactivates the license, removes the license from the list of licenses, and deletes the license from the fabric interconnect. In a cluster setup, you must uninstall the license from the other fabric interconnect.

The following example shows the uninstallation of port9.lic:

```
UCS-A # scope license
UCS-A /license # clear file port9.lic
Clearing license port9.lic:
SERVER this_host ANY
VENDOR cisco
INCREMENT ETH_PORT_ACTIVATION_PKG cisco 1.0 permanent 1 \
  VENDOR_STRING=<LIC_SOURCE>UCS_SWIFT</LIC_SOURCE><SKU>N10-L001=</SKU> \
  HOSTID=VDH=FLC12360025 \
  NOTICE="<LicFileID>20090519200954833</LicFileID><LicLineID>1</LicLineID> \
  <PAK></PAK>" SIGN=C01FAE4E87FA

Clearing license .....done
UCS-A /license #
```



CHAPTER 15

Managing Virtual Interfaces

This chapter includes the following sections:

- [Virtual Interfaces, page 259](#)
- [Virtual Interface Subscription Management and Error Handling, page 259](#)

Virtual Interfaces

In a blade server environment, the number of vNICs and vHBAs configurable for a service profile is determined by adapter capability and the amount of virtual interface (VIF) namespace available on the adapter. In Cisco UCS, portions of VIF namespace are allotted in chunks called VIFs. Depending on your hardware, the maximum number of VIFs are allocated on a predefined, per-port basis.

The maximum number of VIFs varies based on hardware capability and port connectivity. For each configured vNIC or vHBA, one or two VIFs are allocated. Stand-alone vNICs and vHBAs use one VIF and failover vNICs and vHBAs use two.

The following variables affect the number of VIFs available to a blade server, and therefore, how many vNICs and vHBAs you can configure for a service profile.

- Maximum number of VIFs supported on your fabric interconnect
- How the fabric interconnects are cabled
- If your fabric interconnect and IOM are configured in fabric port channel mode

For more information about the maximum number of VIFs supported by your hardware configuration, see *Cisco UCS 6100 and 6200 Series Configuration Limits for Cisco UCS Manager* for your software release.

Virtual Interface Subscription Management and Error Handling

For fabric interconnects grouped in a port-channel, changes to the way you connect the fabric interconnect to the I/O module could result in a drastic change to the number of VIFs available to a blade server. To help you track the effect of these changes, Cisco UCS Manager maintains the following metrics:

- Maximum number of VIFs supported by hardware
- Connectivity type

If you change your configuration in a way that decreases the number of VIFs available to a blade, UCS Manager will display a warning and ask you if you want to proceed. This includes several scenarios, including times where adding or moving a connection decreases the number of VIFs.



PART

Network Configuration

- [Configuring VLANs, page 263](#)
- [Configuring LAN Pin Groups, page 275](#)
- [Configuring MAC Pools, page 277](#)
- [Configuring Quality of Service, page 281](#)
- [Configuring Network-Related Policies, page 291](#)
- [Configuring Upstream Disjoint Layer-2 Networks, page 301](#)



CHAPTER 16

Configuring VLANs

This chapter includes the following sections:

- [Named VLANs, page 263](#)
- [Private VLANs, page 264](#)
- [VLAN Port Limitations, page 265](#)
- [Configuring Named VLANs, page 266](#)
- [Configuring Private VLANs, page 271](#)
- [Viewing the VLAN Port Count, page 274](#)

Named VLANs

A named VLAN creates a connection to a specific external LAN. The VLAN isolates traffic to that external LAN, including broadcast traffic.

The name that you assign to a VLAN ID adds a layer of abstraction that allows you to globally update all servers associated with service profiles that use the named VLAN. You do not need to reconfigure the servers individually to maintain communication with the external LAN.

You can create more than one named VLAN with the same VLAN ID. For example, if servers that host business services for HR and Finance need to access the same external LAN, you can create VLANs named HR and Finance with the same VLAN ID. Then, if the network is reconfigured and Finance is assigned to a different LAN, you only have to change the VLAN ID for the named VLAN for Finance.

In a cluster configuration, you can configure a named VLAN to be accessible only to one fabric interconnect or to both fabric interconnects.

Guidelines for VLAN IDs



Important

You cannot create VLANs with IDs from 3968 to 4047. This range of VLAN IDs is reserved.

VLANs in the LAN cloud and FCoE VLANs in the SAN cloud must have different IDs. Using the same ID for a VLAN and an FCoE VLAN in a VSAN results in a critical fault and traffic disruption for all vNICs and uplink ports using that VLAN. Ethernet traffic is dropped on any VLAN which has an ID that overlaps with an FCoE VLAN ID.

VLAN 4048 is user-configurable. However, Cisco UCS Manager uses VLAN 4048 for the following default values. If you want to assign 4048 to a VLAN, you must reconfigure these values:

- After an upgrade to Cisco UCS, release 2.0: The FCoE storage port native VLAN uses VLAN 4048 by default. If the default FCoE VSAN was set to use VLAN 1 before the upgrade, you must change it to a VLAN ID that is not used or reserved. For example, consider changing the default to 4049 if that VLAN ID is not in use.
- After a fresh install of Cisco UCS, release 2.0: The FCoE VLAN for the default VSAN uses VLAN 4048 by default. The FCoE storage port native VLAN uses VLAN 4049.

The VLAN name is case sensitive.

Private VLANs

A private VLAN (PVLAN) partitions the Ethernet broadcast domain of a VLAN into subdomains and allows you to isolate some ports. Each subdomain in a PVLAN includes a primary VLAN and one or more secondary VLANs. All secondary VLANs in a PVLAN must share the same primary VLAN. The secondary VLAN ID differentiates one subdomain from another.

Isolated VLANs

All secondary VLANs in a Cisco UCS domain must be isolated VLANs. Cisco UCS does not support community VLANs.

Ports on Isolated VLANs

Communications on an isolated VLAN can only use the associated port in the primary VLAN. These ports are isolated ports and are not configurable in Cisco UCS Manager. If the primary VLAN includes multiple secondary VLANs, those isolated VLANs cannot communicate directly with each other.

An isolated port is a host port that belongs to an isolated secondary VLAN. This port has complete isolation from other ports within the same private VLAN domain. PVLANS block all traffic to isolated ports except traffic from promiscuous ports. Traffic received from an isolated port is forwarded only to promiscuous ports. You can have more than one isolated port in a specified isolated VLAN. Each port is completely isolated from all other ports in the isolated VLAN.

Guidelines for Uplink Ports

When you create PVLANS, be aware of the following guidelines:

- The uplink Ethernet port channel cannot be in promiscuous mode.
- Each primary VLAN can have only one isolated VLAN.

- VIFs on VNTAG adapters can have only one isolated VLAN.

Guidelines for VLAN IDs



Important

You cannot create VLANs with IDs from 3968 to 4047. This range of VLAN IDs is reserved.

VLANs in the LAN cloud and FCoE VLANs in the SAN cloud must have different IDs. Using the same ID for a VLAN and an FCoE VLAN in a VSAN results in a critical fault and traffic disruption for all vNICs and uplink ports using that VLAN. Ethernet traffic is dropped on any VLAN which has an ID that overlaps with an FCoE VLAN ID.

VLAN 4048 is user-configurable. However, Cisco UCS Manager uses VLAN 4048 for the following default values. If you want to assign 4048 to a VLAN, you must reconfigure these values:

- After an upgrade to Cisco UCS, release 2.0: The FCoE storage port native VLAN uses VLAN 4048 by default. If the default FCoE VSAN was set to use VLAN 1 before the upgrade, you must change it to a VLAN ID that is not used or reserved. For example, consider changing the default to 4049 if that VLAN ID is not in use.
- After a fresh install of Cisco UCS, release 2.0: The FCoE VLAN for the default VSAN uses VLAN 4048 by default. The FCoE storage port native VLAN uses VLAN 4049.

The VLAN name is case sensitive.

VLAN Port Limitations

Cisco UCS Manager limits the number of VLAN port instances that can be configured under border and server domains on a fabric interconnect to 6000.

Types of Ports Included in the VLAN Port Count

The following types of ports are counted in the VLAN port calculation:

- Border uplink Ethernet ports
- Border uplink Ether-channel member ports
- FCoE ports in a SAN cloud
- Ethernet ports in a NAS cloud
- Static and dynamic vNICs created through service profiles
- VM vNICs created as part of a port profile in a hypervisor in hypervisor domain

Based on the number of VLANs configured for these ports, Cisco UCS Manager keeps track of the cumulative count of VLAN port instances and enforces the VLAN port limit during validation. Cisco UCS Manager reserves some pre-defined VLAN port resources for control traffic. These include management VLANs configured under HIF and NIF ports.

VLAN Port Limit Enforcement

Cisco UCS Manager validates VLAN port availability during the following operations.

- Configuring and unconfiguring border ports and border port channels
- Adding or removing VLANs from a cloud
- Configuring or unconfiguring SAN or NAS ports
- Associating or disassociating service profiles that contain configuration changes
- Configuring or unconfiguring VLANs under vNICs or vHBAs
- Upon receiving creation or deleting notifications from a VMWare vNIC, from an ESX hypervisor



Note This is outside the control of Cisco UCS Manager

- Fabric interconnect reboot
- Cisco UCS Manager upgrade or downgrade

Cisco UCS Manager strictly enforces the VLAN port limit on service profile operations. If Cisco UCS Manager detects that you have exceeded the VLAN port limit service profile configuration will fail during deployment.

Exceeding the VLAN port count in a border domain is less disruptive. When the VLAN port count is exceeded in a border domain Cisco UCS Manager changes the allocation status to Exceeded. In order to change the status back to Available, you should complete one of the following actions:

- Unconfigure one or more border ports
- Remove VLANs from the LAN cloud
- Unconfigure one or more vNICs or vHBAs

Configuring Named VLANs

Creating a Named VLAN Accessible to Both Fabric Interconnects (Uplink Ethernet Mode)



Important

You cannot create VLANs with IDs from 3968 to 4047. This range of VLAN IDs is reserved.

VLANs in the LAN cloud and FCoE VLANs in the SAN cloud must have different IDs. Using the same ID for a VLAN and an FCoE VLAN in a VSAN results in a critical fault and traffic disruption for all vNICs and uplink ports using that VLAN. Ethernet traffic is dropped on any VLAN which has an ID that overlaps with an FCoE VLAN ID.

Procedure

	Command or Action	Purpose
Step 1	UCS-A# scope eth-uplink	Enters Ethernet uplink mode.
Step 2	UCS-A /eth-uplink # create vlan <i>vlan-name vlan-id</i>	Creates a named VLAN, specifies the VLAN name and VLAN ID, and enters Ethernet uplink VLAN mode.

	Command or Action	Purpose
		The VLAN name is case sensitive.
Step 3	UCS-A /eth-uplink/fabric/vlan # set sharing {isolated none primary}	<p>Sets the sharing for the specified VLAN.</p> <p>This can be one of the following:</p> <ul style="list-style-type: none"> • isolated —This is a secondary VLAN associated with a primary VLAN. This VLAN is private. • none —This VLAN does not have any secondary or private VLANs. • primary —This VLAN can have one or more secondary VLANs.
Step 4	UCS-A /eth-uplink/vlan # commit-buffer	Commits the transaction to the system configuration.

The following example creates a named VLAN for both fabric interconnects, names the VLAN accounting, assigns the VLAN ID 2112, sets the sharing to none, and commits the transaction:

```
UCS-A# scope eth-uplink
UCS-A /eth-uplink # create vlan accounting 2112
UCS-A /eth-uplink/vlan* # set sharing none
UCS-A /eth-uplink/vlan* # commit-buffer
UCS-A /eth-uplink/vlan #
```

Creating a Named VLAN Accessible to Both Fabric Interconnects (Ethernet Storage Mode)



Important

You cannot create VLANs with IDs from 3968 to 4047. This range of VLAN IDs is reserved.

VLANs in the LAN cloud and FCoE VLANs in the SAN cloud must have different IDs. Using the same ID for a VLAN and an FCoE VLAN in a VSAN results in a critical fault and traffic disruption for all vNICs and uplink ports using that VLAN. Ethernet traffic is dropped on any VLAN which has an ID that overlaps with an FCoE VLAN ID.

Procedure

	Command or Action	Purpose
Step 1	UCS-A# scope eth-storage	Enters Ethernet storage mode.
Step 2	UCS-A /eth-storage # create vlan vlan-name vlan-id	<p>Creates a named VLAN, specifies the VLAN name and VLAN ID, and enters Ethernet storage VLAN mode.</p> <p>The VLAN name is case sensitive.</p>

	Command or Action	Purpose
Step 3	UCS-A /eth-storage/vlan # create member-port {a b} <i>slot-id port-id</i>	Creates a member port for the specified VLAN on the specified fabric.
Step 4	UCS-A /eth-storage/vlan/member-port # commit-buffer	Commits the transaction to the system configuration.

The following example creates a named VLAN for both fabric interconnects, names the VLAN accounting, assigns the VLAN ID 2112, creates a member port on slot 2, port 20, and commits the transaction:

```
UCS-A# scope eth-storage
UCS-A /eth-storage # create vlan accounting 2112
UCS-A /eth-storage/vlan* # create member-port a 2 20
UCS-A /eth-storage/vlan/member-port* # commit-buffer
UCS-A /eth-storage/vlan/member-port #
```

Creating a Named VLAN Accessible to One Fabric Interconnect (Uplink Ethernet Mode)



Important

You cannot create VLANs with IDs from 3968 to 4047. This range of VLAN IDs is reserved.

VLANs in the LAN cloud and FCoE VLANs in the SAN cloud must have different IDs. Using the same ID for a VLAN and an FCoE VLAN in a VSAN results in a critical fault and traffic disruption for all vNICs and uplink ports using that VLAN. Ethernet traffic is dropped on any VLAN which has an ID that overlaps with an FCoE VLAN ID.

Procedure

	Command or Action	Purpose
Step 1	UCS-A# scope eth-uplink	Enters Ethernet uplink mode.
Step 2	UCS-A /eth-uplink # scope fabric {a b}	Enters Ethernet uplink fabric interconnect mode for the specified fabric interconnect (A or B).
Step 3	UCS-A /eth-uplink/fabric # create vlan <i>vlan-name vlan-id</i>	Creates a named VLAN, specifies the VLAN name and VLAN ID, and enters Ethernet uplink fabric interconnect VLAN mode. The VLAN name is case sensitive.
Step 4	UCS-A /eth-uplink/fabric/vlan # set sharing {isolated none primary}	Sets the sharing for the specified VLAN. This can be one of the following: <ul style="list-style-type: none"> • isolated —This is a secondary VLAN associated with a primary VLAN. This VLAN is private. • none —This VLAN does not have any secondary or private VLANs.

	Command or Action	Purpose
		<ul style="list-style-type: none"> • primary —This VLAN can have one or more secondary VLANs.
Step 5	UCS-A /eth-uplink/fabric/vlan # commit-buffer	Commits the transaction to the system configuration.

The following example creates a named VLAN for fabric interconnect A, names the VLAN finance, assigns the VLAN ID 3955, sets the sharing to none, and commits the transaction:

```
UCS-A# scope eth-uplink
UCS-A /eth-uplink # scope fabric a
UCS-A /eth-uplink/fabric # create vlan finance 3955
UCS-A /eth-uplink/fabric/vlan* # set sharing none
UCS-A /eth-uplink/fabric/vlan* # commit-buffer
UCS-A /eth-uplink/fabric/vlan #
```

Creating a Named VLAN Accessible to One Fabric Interconnect (Ethernet Storage Mode)



Important

You cannot create VLANs with IDs from 3968 to 4047. This range of VLAN IDs is reserved.

VLANs in the LAN cloud and FCoE VLANs in the SAN cloud must have different IDs. Using the same ID for a VLAN and an FCoE VLAN in a VSAN results in a critical fault and traffic disruption for all vNICs and uplink ports using that VLAN. Ethernet traffic is dropped on any VLAN which has an ID that overlaps with an FCoE VLAN ID.

Procedure

	Command or Action	Purpose
Step 1	UCS-A# scope eth-storage	Enters Ethernet storage mode.
Step 2	UCS-A /eth-storage # scope fabric {a b}	Enters Ethernet storage fabric interconnect mode for the specified fabric interconnect.
Step 3	UCS-A /eth-storage/fabric # create vlan <i>vlan-name vlan-id</i>	Creates a named VLAN, specifies the VLAN name and VLAN ID, and enters Ethernet storage fabric interconnect VLAN mode. The VLAN name is case sensitive.
Step 4	UCS-A /eth-storage/vlan # create member-port {a b} slot-id port-id	Creates a member port for the specified VLAN on the specified fabric.
Step 5	UCS-A /eth-storage/fabric/vlan/member-port # commit-buffer	Commits the transaction to the system configuration.

The following example creates a named VLAN for fabric interconnect A, names the VLAN finance, assigns the VLAN ID 3955, creates a member port on slot 2, port 20, and commits the transaction:

```
UCS-A# scope eth-storage
UCS-A /eth-storage # scope fabric a
UCS-A /eth-storage/fabric # create vlan finance 3955
UCS-A /eth-storage/fabric/vlan* # create member-port a 2 20
UCS-A /eth-storage/fabric/vlan/member-port* # commit-buffer
UCS-A /eth-storage/fabric/vlan/member-port #
```

Deleting a Named VLAN

If Cisco UCS Manager includes a named VLAN with the same VLAN ID as the one you delete, the VLAN is not removed from the fabric interconnect configuration until all named VLANs with that ID are deleted.

If you are deleting a private primary VLAN, make sure to reassign the secondary VLANs to another working primary VLAN.

Before You Begin

Before you delete a VLAN from a fabric interconnect, ensure that the VLAN has been removed from all vNICs and vNIC templates.



Note

If you delete a VLAN that is assigned to a vNIC or vNIC template, the vNIC could allow that VLAN to flap.

Procedure

	Command or Action	Purpose
Step 1	UCS-A# scope eth-uplink	Enters Ethernet uplink mode.
Step 2	UCS-A /eth-uplink # scope fabric {a b}	(Optional) Enters Ethernet uplink fabric mode. Use this command when you want to delete a named VLAN only from the specified fabric (a or b).
Step 3	UCS-A /eth-uplink # delete vlan <i>vlan-name</i>	Deletes the specified named VLAN.
Step 4	UCS-A /eth-uplink # commit-buffer	Commits the transaction to the system configuration.

The following example deletes a named VLAN accessible to both fabric interconnects and commits the transaction:

```
UCS-A# scope eth-uplink
UCS-A /eth-uplink # delete vlan accounting
UCS-A /eth-uplink* # commit-buffer
UCS-A /eth-uplink #
```


The following example deletes a named VLAN accessible to one fabric interconnect and commits the transaction:

```
UCS-A# scope eth-uplink
UCS-A /eth-uplink # scope fabric a
UCS-A /eth-uplink/fabric # delete vlan finance
UCS-A /eth-uplink/fabric* # commit-buffer
UCS-A /eth-uplink/fabric #
```

Configuring Private VLANs

Creating a Primary VLAN for a Private VLAN (Accessible to Both Fabric Interconnects)



Important

You cannot create VLANs with IDs from 3968 to 4047. This range of VLAN IDs is reserved.

VLANs in the LAN cloud and FCoE VLANs in the SAN cloud must have different IDs. Using the same ID for a VLAN and an FCoE VLAN in a VSAN results in a critical fault and traffic disruption for all vNICs and uplink ports using that VLAN. Ethernet traffic is dropped on any VLAN which has an ID that overlaps with an FCoE VLAN ID.

Procedure

	Command or Action	Purpose
Step 1	UCS-A# scope eth-uplink	Enters Ethernet uplink mode.
Step 2	UCS-A /eth-uplink # create vlan <i>vlan-name vlan-id</i>	Creates a named VLAN, specifies the VLAN name and VLAN ID, and enters Ethernet uplink VLAN mode. The VLAN name is case sensitive.
Step 3	UCS-A /eth-uplink/vlan # set sharing primary	Sets the VLAN as the primary VLAN.
Step 4	UCS-A /eth-uplink/vlan # commit-buffer	Commits the transaction to the system configuration.

The following example creates a named VLAN for both fabric interconnects, names the VLAN accounting, assigns the VLAN ID 2112, makes this VLAN the primary VLAN, and commits the transaction:

```
UCS-A# scope eth-uplink
UCS-A /eth-uplink # create vlan accounting 2112
UCS-A /eth-uplink/vlan* # set sharing primary
UCS-A /eth-uplink/vlan* # commit-buffer
UCS-A /eth-uplink/vlan #
```

Creating a Primary VLAN for a Private VLAN (Accessible to One Fabric Interconnect)



Important

You cannot create VLANs with IDs from 3968 to 4047. This range of VLAN IDs is reserved.

VLANs in the LAN cloud and FCoE VLANs in the SAN cloud must have different IDs. Using the same ID for a VLAN and an FCoE VLAN in a VSAN results in a critical fault and traffic disruption for all vNICs and uplink ports using that VLAN. Ethernet traffic is dropped on any VLAN which has an ID that overlaps with an FCoE VLAN ID.

Procedure

	Command or Action	Purpose
Step 1	UCS-A# scope eth-uplink	Enters Ethernet uplink mode.
Step 2	UCS-A /eth-uplink # scope fabric {a b}	Enters Ethernet uplink fabric interconnect mode for the specified fabric interconnect.
Step 3	UCS-A /eth-uplink/fabric # create vlan <i>vlan-name vlan-id</i>	Creates a named VLAN, specifies the VLAN name and VLAN ID, and enters Ethernet uplink fabric interconnect VLAN mode. The VLAN name is case sensitive.
Step 4	UCS-A /eth-uplink/fabric/vlan # set sharing primary	Sets the VLAN as the primary VLAN.
Step 5	UCS-A /eth-uplink/fabric/vlan # commit-buffer	Commits the transaction to the system configuration.

The following example creates a named VLAN for fabric interconnect A, names the VLAN finance, assigns the VLAN ID 3955, makes this VLAN the primary VLAN, and commits the transaction:

```
UCS-A# scope eth-uplink
UCS-A /eth-uplink # scope fabric a
UCS-A /eth-uplink/fabric # create vlan finance 3955
UCS-A /eth-uplink/fabric/vlan* # set sharing primary
UCS-A /eth-uplink/fabric/vlan* # commit-buffer
UCS-A /eth-uplink/fabric/vlan #
```

Creating a Secondary VLAN for a Private VLAN (Accessible to Both Fabric Interconnects)



Important

You cannot create VLANs with IDs from 3968 to 4047. This range of VLAN IDs is reserved.

VLANs in the LAN cloud and FCoE VLANs in the SAN cloud must have different IDs. Using the same ID for a VLAN and an FCoE VLAN in a VSAN results in a critical fault and traffic disruption for all vNICs and uplink ports using that VLAN. Ethernet traffic is dropped on any VLAN which has an ID that overlaps with an FCoE VLAN ID.

Procedure

	Command or Action	Purpose
Step 1	UCS-A# scope eth-uplink	Enters Ethernet uplink mode.
Step 2	UCS-A /eth-uplink # create vlan <i>vlan-name vlan-id</i>	Creates a named VLAN, specifies the VLAN name and VLAN ID, and enters Ethernet uplink VLAN mode. The VLAN name is case sensitive.
Step 3	UCS-A /eth-uplink/vlan # set sharing isolated	Sets the VLAN as the secondary VLAN.
Step 4	UCS-A /eth-uplink/vlan # set pubnwnname <i>primary-vlan-name</i>	Specifies the primary VLAN to be associated with this secondary VLAN.
Step 5	UCS-A /eth-uplink/vlan # commit-buffer	Commits the transaction to the system configuration.

The following example creates a named VLAN for both fabric interconnects, names the VLAN accounting, assigns the VLAN ID 2112, makes this VLAN the secondary VLAN, associates the secondary VLAN with the primary VLAN, and commits the transaction:

```
UCS-A# scope eth-uplink
UCS-A /eth-uplink # create vlan accounting 2112
UCS-A /eth-uplink/vlan* # set sharing isolated
UCS-A /eth-uplink/vlan* # set pubnwnname pvlan1000
UCS-A /eth-uplink/vlan* # commit-buffer
UCS-A /eth-uplink/vlan #
```

Creating a Secondary VLAN for a Private VLAN (Accessible to One Fabric Interconnect)



Important

You cannot create VLANs with IDs from 3968 to 4047. This range of VLAN IDs is reserved.

VLANs in the LAN cloud and FCoE VLANs in the SAN cloud must have different IDs. Using the same ID for a VLAN and an FCoE VLAN in a VSAN results in a critical fault and traffic disruption for all vNICs and uplink ports using that VLAN. Ethernet traffic is dropped on any VLAN which has an ID that overlaps with an FCoE VLAN ID.

Procedure

	Command or Action	Purpose
Step 1	UCS-A# scope eth-uplink	Enters Ethernet uplink mode.
Step 2	UCS-A /eth-uplink # scope fabric {a b}	Enters Ethernet uplink fabric interconnect mode for the specified fabric interconnect (A or B).

	Command or Action	Purpose
Step 3	UCS-A /eth-uplink/fabric # create vlan <i>vlan-name vlan-id</i>	Creates a named VLAN, specifies the VLAN name and VLAN ID, and enters Ethernet uplink fabric interconnect VLAN mode. The VLAN name is case sensitive.
Step 4	UCS-A /eth-uplink/vlan # set sharing isolated	Sets the VLAN as the secondary VLAN.
Step 5	UCS-A /eth-uplink/vlan # set pubnwnname <i>primary-vlan-name</i>	Specifies the primary VLAN to be associated with this secondary VLAN.
Step 6	UCS-A /eth-uplink/fabric/vlan/member-port # commit-buffer	Commits the transaction to the system configuration.

The following example creates a named VLAN for fabric interconnect A, names the VLAN finance, assigns the VLAN ID 3955, makes this VLAN the secondary VLAN, associates the secondary VLAN with the primary VLAN, and commits the transaction:

```
UCS-A# scope eth-uplink
UCS-A /eth-uplink # scope fabric a
UCS-A /eth-uplink/fabric # create vlan finance 3955
UCS-A /eth-uplink/fabric/vlan* # set sharing isolated
UCS-A /eth-uplink/fabric/vlan* # set pubnwnname pvlan1000
UCS-A /eth-uplink/fabric/vlan* # commit-buffer
UCS-A /eth-uplink/fabric/vlan #
```

Viewing the VLAN Port Count

Procedure

	Command or Action	Purpose
Step 1	UCS-A# scope fabric-interconnect {a b}	Enters fabric interconnect mode for the specified fabric interconnect.
Step 2	UCS-A /fabric-interconnect # show vlan-port-count	Displays the VLAN port count.

The following example displays the VLAN port count for fabric interconnect A:

```
UCS-A# scope fabric-interconnect a
UCS-A /fabric-interconnect # show vlan-port-count

VLAN-Port Count:
VLAN-Port Limit      Access VLAN-Port Count      Border VLAN-Port Count      Alloc Status
-----
6000                  3                            0                            Available
```



CHAPTER 17

Configuring LAN Pin Groups

This chapter includes the following sections:

- [LAN Pin Groups, page 275](#)
- [Configuring a LAN Pin Group, page 275](#)

LAN Pin Groups

Cisco UCS uses LAN pin groups to pin Ethernet traffic from a vNIC on a server to an uplink Ethernet port or port channel on the fabric interconnect. You can use this pinning to manage the distribution of traffic from the servers.

To configure pinning for a server, you must include the LAN pin group in a vNIC policy. The vNIC policy is then included in the service profile assigned to that server. All traffic from the vNIC travels through the I/O module to the specified uplink Ethernet port.



Note

If you do not assign a pin group to a server interface through a vNIC policy, Cisco UCS Manager chooses an uplink Ethernet port or port channel for traffic from that server interface dynamically. This choice is not permanent. A different uplink Ethernet port or port channel may be used for traffic from that server interface after an interface flap or a server reboot.

Configuring a LAN Pin Group

In a system with two fabric interconnects, you can associate the pin group with only one fabric interconnect or with both fabric interconnects.

Before You Begin

Configure the ports and port channels with which you want to configure the pin group. You can only include ports and port channels configured as uplink ports in a LAN pin group.

Procedure

	Command or Action	Purpose
Step 1	UCS-A# scope eth-uplink	Enters Ethernet uplink mode.
Step 2	UCS-A /eth-uplink # create pin-group <i>pin-group-name</i>	Creates an Ethernet (LAN) pin group with the specified name, and enters Ethernet uplink pin group mode.
Step 3	UCS-A /eth-uplink/pin-group # set descr <i>description</i>	(Optional) Provides a description for the pin group. Note If your description includes spaces, special characters, or punctuation, you must begin and end your description with quotation marks. The quotation marks will not appear in the description field of any show command output.
Step 4	UCS-A /eth-uplink/pin-group # set target { a b dual } { port slot-num / <i>port-num</i> port-channel <i>port-num</i> }	(Optional) Sets the Ethernet pin target to the specified fabric and port, or fabric and port channel.
Step 5	UCS-A /eth-uplink/pin-group # commit-buffer	Commits the transaction to the system configuration.

The following example creates a LAN pin group named pingroup54 on fabric A, provides a description for the pin group, sets the pin group target to port channel 28, and commits the transaction:

```
UCS-A# scope eth-uplink
UCS-A /eth-uplink # create pin-group pingroup54
UCS-A /eth-uplink/pin-group* # set descr "This is my pin group #54"
UCS-A /eth-uplink/pin-group* # set target a port-channel 28
UCS-A /eth-uplink/pin-group* # commit-buffer
UCS-A /eth-uplink/pin-group #
```

What to Do Next

Include the pin group in a vNIC template.



CHAPTER 18

Configuring MAC Pools

This chapter includes the following sections:

- [MAC Pools, page 277](#)
- [Creating a MAC Pool, page 277](#)
- [Deleting a MAC Pool, page 278](#)

MAC Pools

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

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

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

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

Creating a MAC Pool

Procedure

	Command or Action	Purpose
Step 1	UCS-A# scope org <i>org-name</i>	Enters organization mode for the specified organization. To enter the root organization mode, type / as the <i>org-name</i> .
Step 2	UCS-A /org # create mac-pool <i>mac-pool-name</i>	Creates a MAC pool with the specified name, and enters organization MAC pool mode.

	Command or Action	Purpose
Step 3	UCS-A /org/mac-pool # set descr <i>description</i>	(Optional) Provides a description for the MAC pool. Note If your description includes spaces, special characters, or punctuation, you must begin and end your description with quotation marks. The quotation marks will not appear in the description field of any show command output.
Step 4	UCS-A /org/mac-pool # create block <i>first-mac-addr last-mac-addr</i>	Creates a block (range) of MAC addresses, and enters organization MAC pool block mode. You must specify the first and last MAC addresses in the address range using the form <i>nn:nn:nn:nn:nn:nn</i> , with the addresses separated by a space. Note A MAC pool can contain more than one MAC address block. To create multiple MAC address blocks, you must enter multiple create block commands from organization MAC pool mode.
Step 5	UCS-A /org/mac-pool # commit-buffer	Commits the transaction to the system configuration.

The following example creates a MAC pool named pool37, provides a description for the pool, defines a MAC address block by specifying the first and last MAC addresses in the block, and commits the transaction:

```
UCS-A# scope org /
UCS-A /org # create mac-pool pool37
UCS-A /org/mac-pool* # set descr "This is my MAC pool"
UCS-A /org/mac-pool* # create block 00:A0:D7:42:00:01 00:A0:D7:42:01:00
UCS-A /org/mac-pool/block* # commit-buffer
UCS-A /org/mac-pool/block #
```

What to Do Next

Include the MAC pool in a vNIC template.

Deleting a MAC Pool

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

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

Procedure

	Command or Action	Purpose
Step 1	UCS-A# scope org <i>org-name</i>	Enters organization mode for the specified organization. To enter the root organization mode, type / as the <i>org-name</i> .
Step 2	UCS-A /org # delete mac-pool <i>pool-name</i>	Deletes the specified MAC pool.
Step 3	UCS-A /org # commit-buffer	Commits the transaction to the system configuration.

The following example deletes the MAC pool named pool4 and commits the transaction:

```
UCS-A# scope org /
UCS-A /org # delete mac-pool pool4
UCS-A /org* # commit-buffer
UCS-A /org #
```




CHAPTER 19

Configuring Quality of Service

This chapter includes the following sections:

- [Quality of Service, page 281](#)
- [Configuring System Classes, page 281](#)
- [Configuring Quality of Service Policies, page 284](#)
- [Configuring Flow Control Policies, page 287](#)

Quality of Service

Cisco UCS provides the following methods to implement quality of service:

- System classes that specify the global configuration for certain types of traffic across the entire system
- QoS policies that assign system classes for individual vNICs
- Flow control policies that determine how uplink Ethernet ports handle pause frames

Configuring System Classes

System Classes

Cisco UCS uses Data Center Ethernet (DCE) to handle all traffic inside a Cisco UCS domain. This industry standard enhancement to Ethernet divides the bandwidth of the Ethernet pipe into eight virtual lanes. Two virtual lanes are reserved for internal system and management traffic. You can configure quality of service 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. This provides a level of traffic management, even in an oversubscribed system. For example, you can configure the Fibre Channel Priority system class to determine the percentage of DCE bandwidth allocated to FCoE traffic.

The following table describes the system classes that you can configure:

Table 13: System Classes

System Class	Description
Platinum Gold Silver Bronze	A configurable set of system classes that you can include in the QoS policy for a service profile. Each system class manages one lane of traffic. All properties of these system classes are available for you to assign custom settings and policies.
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.

Configuring a System Class

The type of adapter in a server may limit the maximum MTU supported. For example, network MTU above the maximums may cause the packet to be dropped for the following adapters:

- The Cisco UCS M71KR CNA adapter, which supports a maximum MTU of 9216.
- The Cisco UCS 82598KR-CI adapter, which supports a maximum MTU of 14000.

Procedure

	Command or Action	Purpose
Step 1	UCS-A# scope eth-server	Enters Ethernet server mode.
Step 2	UCS-A /eth-server # scope qos	Enters Ethernet server QoS mode.
Step 3	UCS-A /eth-server/qos # scope eth-classified {bronze gold platinum silver}	Enters Ethernet server QoS Ethernet classified mode for the specified system class.
Step 4	UCS-A /eth-server/qos/eth-classified # enable	Enables the specified system class.
Step 5	UCS-A /eth-server/qos/eth-classified # set cos cos-value	Specifies the class of service for the specified system class. Valid class of service values are 0 to 6; higher values indicate more important traffic.

	Command or Action	Purpose
Step 6	UCS-A /eth-server/qos/eth-classified # set drop {drop no-drop}	Specifies whether the channel can drop packets or not. Note Only one system class can use the no-drop option.
Step 7	UCS-A /eth-server/qos/eth-classified # set mtu {mtu-value fc normal}	The maximum transmission unit, or packet size, that this vNIC accepts. Note If the vNIC has an associated QoS policy, the MTU specified here must be equal to or less than the MTU specified in the associated QoS system class. If this MTU value exceeds the MTU value in the QoS system class, packets may be dropped during data transmission.
Step 8	UCS-A /eth-server/qos/eth-classified # set multicast-optimize {no yes}	Specifies whether the class is optimized to for sending multicast packets.
Step 9	UCS-A /eth-server/qos/eth-classified # set weight {weight-value best-effort none}	Specifies the relative weight for the specified system class. Valid weight values are 0 to 10.
Step 10	UCS-A /eth-server/qos/eth-classified # commit-buffer	Commits the transaction to the system configuration.

The following example enables the platinum system class, allows the channel to drop packets, sets the class of service to 6, sets the MTU to normal, optimizes the class for sending multicast packets, sets the relative weight to 5, and commits the transaction:

```
UCS-A# scope eth-server
UCS-A /eth-server # scope qos
UCS-A /eth-server/qos # scope eth-classified platinum
UCS-A /eth-server/qos/eth-classified # enable
UCS-A /eth-server/qos/eth-classified* # set drop drop
UCS-A /eth-server/qos/eth-classified* # set cos 6
UCS-A /eth-server/qos/eth-classified* # set mtu normal
UCS-A /eth-server/qos/eth-classified* # set multicast-optimize yes
UCS-A /eth-server/qos/eth-classified* # set weight 5
UCS-A /eth-server/qos/eth-classified* # commit-buffer
UCS-A /eth-server/qos/eth-classified #
```

Disabling a System Class

If you disable a system class that is used in a QoS policy, Cisco UCS Manager uses the system class configured with CoS 0 for traffic on servers that are configured with the QoS policy. If no system class is configured as CoS 0, the Best Effort system class is used. You cannot disable the Best Effort or Fibre Channel system classes.

Procedure

	Command or Action	Purpose
Step 1	UCS-A# scope eth-server	Enters Ethernet server mode.
Step 2	UCS-A /eth-server # scope qos	Enters Ethernet server QoS mode.
Step 3	UCS-A /eth-server/qos # scope eth-classified { bronze gold platinum silver }	Enters Ethernet server QoS Ethernet classified mode for the specified system class.
Step 4	UCS-A /eth-server/qos/eth-classified # disable	Disables the specified system class.
Step 5	UCS-A /eth-server/qos/eth-classified # commit-buffer	Commits the transaction to the system configuration.

The following example disables the platinum system class and commits the transaction:

```
UCS-A# scope eth-server
UCS-A /eth-server # scope qos
UCS-A /eth-server/qos # scope eth-classified platinum
UCS-A /eth-server/qos/eth-classified # disable
UCS-A /eth-server/qos/eth-classified* # commit-buffer
UCS-A /eth-server/qos/eth-classified #
```

Configuring Quality of Service Policies

Quality of Service Policy

A quality of service (QoS) policy assigns a system class to the outgoing traffic for a vNIC or vHBA. This system class determines the quality of service for that traffic. For certain adapters you can also specify additional controls on the outgoing traffic, such as burst and rate.

You must include a QoS policy in a vNIC policy or vHBA policy and then include that policy in a service profile to configure the vNIC or vHBA.

Configuring a QoS Policy

Procedure

	Command or Action	Purpose
Step 1	Switch-A# scope org <i>org-name</i>	Enters org mode for the specified organization. To enter the default org mode, type / as the <i>org-name</i> .
Step 2	Switch-A /org # create qos-policy <i>policy-name</i>	Creates the specified QoS policy, and enters org QoS policy mode.

	Command or Action	Purpose
Step 3	Switch-A /org/qos-policy # create egress-policy	Creates the egress policy (for both vNICs and vHBAs) to be used by the QoS policy, and enters org QoS policy egress policy mode.
Step 4	Switch-A /org/qos-policy/egress-policy # set host-cos-control {full none}	(Optional) Specifies whether the host or Cisco UCS Manager controls the class of service (CoS). Use the full keyword to have the host control the CoS. If the packet has a valid CoS value, the host uses that value. Otherwise, it uses the CoS value associated with the specified class priority. Use the none keyword to have Cisco UCS Manager use the CoS value associated with the specified priority.
Step 5	Switch-A /org/qos-policy/egress-policy # set prio sys-class-name	Specifies the system class to be used for the egress policy. The <i>sys-class-name</i> argument can be one of the following class keywords: <ul style="list-style-type: none"> • Fc—Use this priority for QoS policies that control vHBA traffic only. • Platinum—Use this priority for QoS policies that control vNIC traffic only. • Gold—Use this priority for QoS policies that control vNIC traffic only. • Silver—Use this priority for QoS policies that control vNIC traffic only. • Bronze—Use this priority for QoS policies that control vNIC traffic only. • Best Effort—Do not use this priority. It is reserved for the Basic Ethernet traffic lane. If you assign this priority to a QoS policy and configure another system class as CoS 0, Cisco UCS Manager does not default to this system class. It defaults to the priority with CoS 0 for that traffic.
Step 6	Switch-A /org/qos-policy/egress-policy # set rate {line-rate kbps} burst bytes	Specifies the rate limit for egress traffic by defining the average traffic rate and burst size. The line-rate keyword sets the rate limit to the physical line rate. Rate limiting is not applicable to all adapters. For example, this setting is not supported on the Cisco UCS M82-8P Virtual Interface Card.
Step 7	Switch-A /org/qos-policy/egress-policy # committ-buffer	Commits the transaction to the system configuration.

The following example creates a QoS policy for vNIC traffic, assigns the platinum system class and sets the rate limit (traffic rate and burst size) for the egress policy, and commits the transaction:

```
Switch-A# scope org /
Switch-A /org # create qos-policy VnicPolicy34
Switch-A /org/qos-policy* # create egress-policy
Switch-A /org/qos-policy/egress-policy* # set prio platinum
Switch-A /org/qos-policy/egress-policy* # set rate 5000000 burst 65000
Switch-A /org/qos-policy/egress-policy* # commit-buffer
Switch-A /org/qos-policy/egress-policy #
```

The following example creates a QoS policy for vHBA traffic, assigns the fc (Fibre Channel) system class and sets the rate limit (traffic rate and burst size) for the egress policy, and commits the transaction:

```
Switch-A# scope org /
Switch-A /org # create qos-policy VhbaPolicy12
Switch-A /org/qos-policy* # create egress-policy
Switch-A /org/qos-policy/egress-policy* # set prio fc
Switch-A /org/qos-policy/egress-policy* # set rate 5000000 burst 65000
Switch-A /org/qos-policy/egress-policy* # commit-buffer
Switch-A /org/qos-policy/egress-policy #
```

What to Do Next

Include the QoS policy in a vNIC or vHBA template.

Deleting a QoS Policy

If you delete a QoS policy that is in use or you disable a system class that is used in a QoS policy, any vNIC or vHBA that uses that QoS policy is assigned to the Best Effort system class or to the system class with a CoS of 0. In a system that implements multi-tenancy, Cisco UCS Manager first attempts to find a matching QoS policy in the organization hierarchy.

Procedure

	Command or Action	Purpose
Step 1	UCS-A# scope org <i>org-name</i>	Enters organization mode for the specified organization. To enter the root organization mode, type <i>/</i> as the <i>org-name</i> .
Step 2	UCS-A /org # delete qos-policy <i>policy-name</i>	Deletes the specified QoS policy.
Step 3	UCS-A /org # commit-buffer	Commits the transaction to the system configuration.

The following deletes the QoS policy named QosPolicy34 and commits the transaction:

```
UCS-A# scope org /
UCS-A /org # delete qos-policy QosPolicy34
UCS-A /org* # commit-buffer
UCS-A /org #
```


Configuring Flow Control Policies

Flow Control Policy

Flow control policies determine whether the uplink Ethernet ports in a Cisco UCS domain send and receive IEEE 802.3x pause frames when the receive buffer for a port fills. These pause frames request that the transmitting port stop sending data for a few milliseconds until the buffer clears.

For flow control to work between a LAN port and an uplink Ethernet port, you must enable the corresponding receive and send flow control parameters for both ports. For Cisco UCS, the flow control policies configure these parameters.

When you enable the send function, the uplink Ethernet port sends a pause request to the network port if the incoming packet rate becomes too high. The pause remains in effect for a few milliseconds before traffic is reset to normal levels. If you enable the receive function, the uplink Ethernet port honors all pause requests from the network port. All traffic is halted on that uplink port until the network port cancels the pause request.

Because you assign the flow control policy to the port, changes to the policy have an immediate effect on how the port reacts to a pause frame or a full receive buffer.

Configuring a Flow Control Policy

Before You Begin

Configure the network port with the corresponding setting for the flow control that you need. For example, if you enable the send setting for flow-control pause frames in the policy, make sure that the receive parameter in the network port is set to on or desired. If you want the Cisco UCS port to receive flow-control frames, make sure that the network port has a send parameter set to on or desired. If you do not want to use flow control, you can set the send and receive parameters on the network port to off.

Procedure

	Command or Action	Purpose
Step 1	UCS-A# scope eth-uplink	Enters Ethernet uplink mode.
Step 2	UCS-A /eth-uplink # scope flow-control	Enters Ethernet uplink flow control mode.
Step 3	UCS-A /eth-uplink/flow-control # create policy <i>policy-name</i>	Creates the specified flow control policy.
Step 4	UCS-A /eth-uplink/flow-control/policy # set prio <i>prio-option</i>	Specifies one of the following flow control priority options: <ul style="list-style-type: none"> • auto —The Cisco UCS system and the network negotiate whether PPP will be used on this fabric interconnect. • on —PPP is enabled on this fabric interconnect.

	Command or Action	Purpose
Step 5	UCS-A /eth-uplink/flow-control/policy # set receive <i>receive-option</i>	Specifies one of the following flow control receive options: <ul style="list-style-type: none"> • off —Pause requests from the network are ignored and traffic flow continues as normal. • on —Pause requests are honored and all traffic is halted on that uplink port until the network cancels the pause request.
Step 6	UCS-A /eth-uplink/flow-control/policy # set send <i>send-option</i>	Specifies one of the following flow control send options: <ul style="list-style-type: none"> • off —Traffic on the port flows normally regardless of the packet load. • on —The Cisco UCS system sends a pause request to the network if the incoming packet rate becomes too high. The pause remains in effect for a few milliseconds before traffic is reset to normal levels.
Step 7	UCS-A /org/qos-policy/vnic-egress-policy # commit-buffer	Commits the transaction to the system configuration.

The following configures a flow control policy and commits the transaction:

```
UCS-A# scope eth-uplink
UCS-A /eth-uplink # scope flow-control
UCS-A /eth-uplink/flow-control # create policy FlowControlPolicy23
UCS-A /eth-uplink/flow-control/policy* # set prio auto
UCS-A /eth-uplink/flow-control/policy* # set receive on
UCS-A /eth-uplink/flow-control/policy* # set send on
UCS-A /eth-uplink/flow-control/policy* # commit-buffer
UCS-A /eth-uplink/flow-control/policy #
```

What to Do Next

Associate the flow control policy with an uplink Ethernet port or port channel.

Deleting a Flow Control Policy

Procedure

	Command or Action	Purpose
Step 1	UCS-A# scope eth-uplink	Enters Ethernet uplink mode.
Step 2	UCS-A /eth-uplink # scope flow-control	Enters Ethernet uplink flow control mode.
Step 3	UCS-A /eth-uplink/flow-control # delete policy <i>policy-name</i>	Deletes the specified flow control policy.

	Command or Action	Purpose
Step 4	UCS-A /eth-uplink/flow-control # commit-buffer	Commits the transaction to the system configuration.

The following example deletes the flow control policy named FlowControlPolicy23 and commits the transaction:

```
UCS-A# scope eth-uplink
UCS-A /eth-uplink # scope flow-control
UCS-A /eth-uplink/flow-control # delete policy FlowControlPolicy23
UCS-A /eth-uplink/flow-control* # commit-buffer
UCS-A /eth-uplink/flow-control #
```




CHAPTER 20

Configuring Network-Related Policies

This chapter includes the following sections:

- [Configuring vNIC Templates, page 291](#)
- [Configuring Ethernet Adapter Policies, page 294](#)
- [Configuring Network Control Policies, page 297](#)

Configuring vNIC Templates

vNIC Template

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

Beginning in Cisco UCS, Release 2.0(2), Cisco UCS Manager does not automatically create a VM-FEX port profile with the correct settings when you create a vNIC template. If you want to create a VM-FEX port profile, you must configure the target of the vNIC template as a VM.

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



Note

If your server has two Emulex or QLogic NICs (Cisco UCS CNA M71KR-E or Cisco UCS CNA M71KR-Q), you must configure vNIC policies for both adapters in your service profile to get a user-defined MAC address for both NICs. If you do not configure policies for both NICs, Windows still detects both of them in the PCI bus. Then because the second eth is not part of your service profile, Windows assigns it a hardware MAC address. If you then move the service profile to a different server, Windows sees additional NICs because one NIC did not have a user-defined MAC address.

Configuring a vNIC Template

Procedure

	Command or Action	Purpose
Step 1	UCS-A# scope org <i>org-name</i>	Enters organization mode for the specified organization. To enter the root organization mode, type <i>/</i> as the <i>org-name</i> .
Step 2	UCS-A /org # create vnic-templ <i>vnic-templ-name</i> [eth-if <i>vlan-name</i>] [fabric { a b }] [target [adapter vm]]	<p>Creates a vNIC template and enters organization vNIC template mode.</p> <p>The target you choose determines whether or not Cisco UCS Manager automatically creates a VM-FEX port profile with the appropriate settings for the vNIC template. This can be one of the following:</p> <ul style="list-style-type: none"> • Adapter—The vNICs apply to all adapters. No VM-FEX port profile is created if you choose this option. • VM—The vNICs apply to all virtual machines. A VM-FEX port profile is created if you choose this option.
Step 3	UCS-A /org/vnic-templ # set descr <i>description</i>	(Optional) Provides a description for the vNIC template.
Step 4	UCS-A /org/vnic-templ # set fabric { a a-b b b-a }	<p>(Optional)</p> <p>Specifies the fabric to use for the vNIC. If you did not specify the fabric when creating the vNIC template in Step 2, you have the option to specify it with this command.</p> <p>If you want this vNIC to be able to access the second fabric interconnect if the default one is unavailable, choose a-b (A is the primary) or b-a (B is the primary) .</p> <p>Note Do not enable fabric failover for the vNIC under the following circumstances:</p> <ul style="list-style-type: none"> • If the Cisco UCS domain is running in Ethernet Switch Mode. vNIC fabric failover is not supported in that mode. If all Ethernet uplinks on one fabric interconnect fail, the vNICs do not fail over to the other. • if you plan to associate this vNIC with a server that has an adapter which does not support fabric failover, such as the Cisco UCS 82598KR-CI 10-Gigabit Ethernet Adapter. If you do so, Cisco UCS Manager generates a configuration fault when you associate the service profile with the server.

	Command or Action	Purpose
Step 5	UCS-A /org/vnic-templ # set mac-pool <i>mac-pool-name</i>	The MAC address pool that vNICs created from this vNIC template should use.
Step 6	UCS-A /org/vnic-templ # set mtu <i>mtu-value</i>	The maximum transmission unit, or packet size, that vNICs created from this vNIC template should use. Enter an integer between 1500 and 9216. Note If the vNIC template has an associated QoS policy, the MTU specified here must be equal to or less than the MTU specified in the associated QoS system class. If this MTU value exceeds the MTU value in the QoS system class, packets may be dropped during data transmission.
Step 7	UCS-A /org/vnic-templ # set nw-control-policy <i>policy-name</i>	The network control policy that vNICs created from this vNIC template should use.
Step 8	UCS-A /org/vnic-templ # set pin-group <i>group-name</i>	The LAN pin group that vNICs created from this vNIC template should use.
Step 9	UCS-A /org/vnic-templ # set qos-policy <i>policy-name</i>	The quality of service policy that vNICs created from this vNIC template should use.
Step 10	UCS-A /org/vnic-templ # set stats-policy <i>policy-name</i>	The statistics collection policy that vNICs created from this vNIC template should use.
Step 11	UCS-A /org/vnic-templ # set type { initial-template updating-template }	Specifies the vNIC template update type. If you do not want vNIC instances created from this template to be automatically updated when the template is updated, use the initial-template keyword; otherwise, use the updating-template keyword to ensure that all vNIC instances are updated when the vNIC template is updated.
Step 12	UCS-A /org/vnic-templ # commit-buffer	Commits the transaction to the system configuration.

The following example configures a vNIC template and commits the transaction:

```
UCS-A# scope org /
UCS-A /org* # create vnic template VnicTempFoo
UCS-A /org/vnic-templ* # set descr "This is a vNIC template example."
UCS-A /org/vnic-templ* # set fabric a
UCS-A /org/vnic-templ* # set mac-pool pool137
UCS-A /org/vnic-templ* # set mtu 8900
UCS-A /org/vnic-templ* # set nw-control-policy ncp5
UCS-A /org/vnic-templ* # set pin-group PinGroup54
UCS-A /org/vnic-templ* # set qos-policy QosPol5
UCS-A /org/vnic-templ* # set stats-policy ServStatsPolicy
UCS-A /org/vnic-templ* # set type updating-template
UCS-A /org/vnic-templ* # commit-buffer
UCS-A /org/vnic-templ #
```

Deleting a vNIC Template

Procedure

	Command or Action	Purpose
Step 1	UCS-A# scope org <i>org-name</i>	Enters organization mode for the specified organization. To enter the root organization mode, type / as the <i>org-name</i> .
Step 2	UCS-A /org # delete vnic-templ <i>vnic-templ-name</i>	Deletes the specified vNIC template.
Step 3	UCS-A /org # commit-buffer	Commits the transaction to the system configuration.

The following example deletes the vNIC template named VnicTempFoo and commits the transaction:

```
UCS-A# scope org /
UCS-A /org # delete vnic template VnicTempFoo
UCS-A /org* # commit-buffer
UCS-A /org #
```

Configuring Ethernet Adapter Policies

Ethernet and Fibre Channel Adapter Policies

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

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

**Note**

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

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

Operating System Specific Adapter Policies

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

**Important**

We recommend that you use the values in these policies for the applicable operating system. Do not modify any of the values in the default policies unless directed to do so by Cisco Technical Support.

However, if you are creating an Ethernet adapter policy for a Windows OS (instead of using the default Windows adapter policy), you must use the following formulas to calculate values that work with Windows:

$$\text{Completion Queues} = \text{Transmit Queues} + \text{Receive Queues}$$

$$\text{Interrupt Count} = (\text{Completion Queues} + 2) \text{ rounded up to nearest power of } 2$$

For example, if Transmit Queues = 1 and Receive Queues = 8 then:

$$\text{Completion Queues} = 1 + 8 = 9$$

$$\text{Interrupt Count} = (9 + 2) \text{ rounded up to the nearest power of } 2 = 16$$

Configuring an Ethernet Adapter Policy

Procedure

	Command or Action	Purpose
Step 1	UCS-A# scope org <i>org-name</i>	Enters organization mode for the specified organization. To enter the root organization mode, type / as the <i>org-name</i> .

	Command or Action	Purpose
Step 2	UCS-A /org# create eth-policy <i>policy-name</i>	Creates the specified Ethernet adapter policy and enters organization Ethernet policy mode.
Step 3	UCS-A /org/eth-policy # set comp-queue count <i>count</i>	(Optional) Configures the Ethernet completion queue.
Step 4	UCS-A /org/eth-policy # set descr <i>description</i>	(Optional) Provides a description for the policy. Note If your description includes spaces, special characters, or punctuation, you must begin and end your description with quotation marks. The quotation marks will not appear in the description field of any show command output.
Step 5	UCS-A /org/eth-policy # set failover timeout <i>timeout-sec</i>	(Optional) Configures the Ethernet failover.
Step 6	UCS-A /org/eth-policy # set interrupt { coalescing-time <i>sec</i> coalescing-type { idle min } count <i>count</i> mode { intx msi msi-x }}	(Optional) Configures the Ethernet interrupt.
Step 7	UCS-A /org/eth-policy # set offload { large-receive tcp-rx-checksum tcp-segment tcp-tx-checksum } { disabled enabled }	(Optional) Configures the Ethernet offload.
Step 8	UCS-A /org/eth-policy # set recv-queue { count <i>count</i> ring-size <i>size-num</i> }	(Optional) Configures the Ethernet receive queue.
Step 9	UCS-A /org/eth-policy # set rss receivesidescaling { disabled enabled }	(Optional) Configures the RSS.
Step 10	UCS-A /org/eth-policy # set trans-queue { count <i>count</i> ring-size <i>size-num</i> }	(Optional) Configures the Ethernet transmit queue.
Step 11	UCS-A /org/eth-policy # commit-buffer	Commits the transaction to the system configuration.

The following example configures an Ethernet adapter policy, and commits the transaction:

```
UCS-A# scope org /
UCS-A /org* # create eth-policy EthPolicy19
UCS-A /org/eth-policy* # set comp-queue count 16
UCS-A /org/eth-policy* # set descr "This is an Ethernet adapter policy example."
UCS-A /org/eth-policy* # set failover timeout 300
UCS-A /org/eth-policy* # set interrupt count 64
UCS-A /org/eth-policy* # set offload large-receive disabled
UCS-A /org/eth-policy* # set recv-queue count 32
UCS-A /org/eth-policy* # set rss receivesidescaling enabled
UCS-A /org/eth-policy* # set trans-queue
UCS-A /org/eth-policy* # commit-buffer
UCS-A /org/eth-policy #
```

Deleting an Ethernet Adapter Policy

Procedure

	Command or Action	Purpose
Step 1	UCS-A# scope org <i>org-name</i>	Enters organization mode for the specified organization. To enter the root organization mode, type <i>/</i> as the <i>org-name</i> .
Step 2	UCS-A /org # delete eth-policy <i>policy-name</i>	Deletes the specified Ethernet adapter policy.
Step 3	UCS-A /org # commit-buffer	Commits the transaction to the system configuration.

The following example deletes the Ethernet adapter policy named EthPolicy19 and commits the transaction:

```
UCS-A# scope org /
UCS-A /org # delete eth-policy EthPolicy19
UCS-A /org* # commit-buffer
UCS-A /org #
```

Configuring Network Control Policies

Network Control Policy

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

- Whether the Cisco Discovery Protocol (CDP) is enabled or disabled
- How the VIF behaves if no uplink port is available in end-host mode
- The action that Cisco UCS Manager takes on the remote Ethernet interface, vEthernet interface, or vFibreChannel interface when the associated border port fails
- Whether the server can use different MAC addresses when sending packets to the fabric interconnect
- Whether MAC registration occurs on a per-VNIC basis or for all VLANs.

Action on Uplink Fail

By default, the **Action on Uplink Fail** property in the network control policy is configured with a value of link-down. For adapters such as the Cisco UCS M81KR Virtual Interface Card, this default behavior directs Cisco UCS Manager to bring the vEthernet or vFibreChannel interface down if the associated border port fails. For Cisco UCS systems using a non-VM-FEX capable converged network adapter that supports both Ethernet and FCoE traffic, such as Cisco UCS CNA M72KR-Q and the Cisco UCS CNA M72KR-E, this default behavior directs Cisco UCS Manager to bring the remote Ethernet interface down if the associated border port fails. In this scenario, any vFibreChannel interfaces that are bound to the remote Ethernet interface are brought down as well.



Note

Cisco UCS Manager, release 1.4(2) and earlier did not enforce the **Action on Uplink Fail** property for those types of non-VM-FEX capable converged network adapters mentioned above. If the **Action on Uplink Fail** property was set to link-down, Cisco UCS Manager would ignore this setting and instead issue a warning. In the current version of Cisco UCS Manager this setting is enforced. Therefore, if your implementation includes one of those converged network adapters and the adapter is expected to handle both Ethernet and FCoE traffic, we recommend that you configure the **Action on Uplink Fail** property with a value of warning.

Please note that this configuration may result in an Ethernet teaming driver not being able to detect a link failure when the border port goes down.

MAC Registration Mode

In Cisco UCS Manager, releases 1.4 and earlier, MAC addresses were installed on all of the VLANs belonging to an interface. Starting in release 2.0, MAC addresses are installed only on the native VLAN by default. In most implementations this maximizes the VLAN port count.



Note

If a trunking driver is being run on the host and the interface is in promiscuous mode, we recommend that you set the Mac Registration Mode to All VLANs.

Configuring a Network Control Policy

MAC address-based port security for Emulex converged Network Adapters (N20-AE0102) is not supported. When MAC address-based port security is enabled, the fabric interconnect restricts traffic to packets that contain the MAC address that it first learns. This is either the source MAC address used in the FCoE Initialization Protocol packet, or the MAC address in an ethernet packet, whichever is sent first by the adaptor. This configuration can result in either FCoE or Ethernet packets being dropped.

Procedure

	Command or Action	Purpose
Step 1	UCS-A# scope org /	Enters the root organization mode.
Step 2	UCS-A /org # create nw-ctrl-policy <i>policy-name</i>	Creates the specified network control policy, and enters organization network control policy mode.
Step 3	UCS-A /org/nw-ctrl-policy # {disable enable} cdp	Disables or enables Cisco Discovery Protocol (CDP).
Step 4	UCS-A /org/nw-ctrl-policy # set uplink-fail-action {link-down warning}	Specifies the action to be taken when no uplink port is available in end-host mode. Use the link-down keyword to change the operational state of a vNIC to down when uplink connectivity is lost on the fabric interconnect, and facilitate fabric failover for vNICs. Use the warning keyword to maintain server-to-server connectivity even when no uplink port is

	Command or Action	Purpose
		available, and disable fabric failover when uplink connectivity is lost on the fabric interconnect. The default uplink failure action is link-down.
Step 5	UCS-A /org/nw-ctrl-policy # set mac-registration-mode {all-host-vlans only-native-vlan	Whether adapter-registered MAC addresses are added only to the native VLAN associated with the interface or added to all VLANs associated with the interface. This can be one of the following: <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 <i>not</i> running in Promiscuous mode.
Step 6	UCS-A /org/nw-ctrl-policy # create mac-security	Enters organization network control policy MAC security mode
Step 7	UCS-A /org/nw-ctrl-policy/mac-security # set forged-transmit {allow deny}	Allows or denies the forging of MAC addresses when sending traffic. MAC security is disabled when forged MAC addresses are allowed, and MAC security is enabled when forged MAC addresses are denied. By default, forged MAC addresses are allowed (MAC security is disabled).
Step 8	UCS-A /org/nw-ctrl-policy/mac-security # commit-buffer	Commits the transaction to the system configuration.

The following example creates a network control policy named ncp5, enables CDP, sets the uplink fail action to link-down, denies forged MAC addresses (enables MAC security), and commits the transaction:

```
UCS-A# scope org /
UCS-A /org # create nw-ctrl-policy ncp5
UCS-A /org/nw-ctrl-policy* # enable cdp
UCS-A /org/nw-ctrl-policy* # set uplink-fail-action link-down
UCS-A /org/nw-ctrl-policy* # create mac-security
UCS-A /org/nw-ctrl-policy/mac-security* # set forged-transmit deny
UCS-A /org/nw-ctrl-policy/mac-security* # commit-buffer
UCS-A /org/nw-ctrl-policy/mac-security #
```

Deleting a Network Control Policy

Procedure

	Command or Action	Purpose
Step 1	UCS-A# scope org /	Enters the root organization mode.

	Command or Action	Purpose
Step 2	UCS-A /org # delete nwctrl-policy <i>policy-name</i>	Deletes the specified network control policy.
Step 3	UCS-A /org # commit-buffer	Commits the transaction to the system configuration.

The following example deletes the network control policy named ncp5 and commits the transaction:

```
UCS-A# scope org /  
UCS-A /org # delete nwctrl-policy ncp5  
UCS-A /org* # commit-buffer  
UCS-A /org #
```



CHAPTER 21

Configuring Upstream Disjoint Layer-2 Networks

This chapter includes the following sections:

- [Upstream Disjoint Layer-2 Networks, page 301](#)
- [Guidelines for Configuring Upstream Disjoint L2 Networks, page 302](#)
- [Pinning Considerations for Upstream Disjoint L2 Networks, page 303](#)
- [Configuring Cisco UCS for Upstream Disjoint L2 Networks, page 304](#)
- [Assigning Ports and Port Channels to VLANs, page 305](#)
- [Removing Ports and Port Channels from VLANs, page 306](#)
- [Viewing Ports and Port Channels Assigned to VLANs, page 307](#)

Upstream Disjoint Layer-2 Networks

Upstream disjoint layer-2 networks (disjoint L2 networks) are required if you have two or more Ethernet “clouds” that never connect, but must be accessed by servers or virtual machines located in the same Cisco UCS domain. For example, you could configure disjoint L2 networks if you require one of the following:

- Servers or virtual machines to access a public network and a backup network
- In a multi-tenant system, servers or virtual machines for more than one customer are located in the same Cisco UCS domain and need to access the L2 networks for both customers.



Note

By default, data traffic in Cisco UCS works on a principle of mutual inclusion. All traffic for all VLANs and upstream networks travels along all uplink ports and port channels. If you have upgraded from a release that does not support upstream disjoint layer-2 networks, you must assign the appropriate uplink interfaces to your VLANs, or traffic for those VLANs continues to flow along all uplink ports and port channels.

The configuration for disjoint L2 networks works on a principle of selective exclusion. Traffic for a VLAN that is designated as part of a disjoint network can only travel along an uplink Ethernet port or port channel that is specifically assigned to that VLAN, and is selectively excluded from all other uplink ports and port channels. However, traffic for VLANs that are not specifically assigned to an uplink Ethernet port or port

channel can still travel on all uplink ports or port channels, including those that carry traffic for the disjoint L2 networks.

In Cisco UCS, the VLAN represents the upstream disjoint L2 network. When you design your network topology for disjoint L2 networks, you must assign uplink interfaces to VLANs not the reverse.

For information about the maximum number of supported upstream disjoint L2 networks, see [Cisco UCS 6100 and 6200 Series Configuration Limits for Cisco UCS Manager, Release 2.0](#).

Guidelines for Configuring Upstream Disjoint L2 Networks

When you plan your configuration for upstream disjoint L2 networks, consider the following:

Ethernet Switching Mode Must Be End-Host Mode

Cisco UCS only supports disjoint L2 networks when the Ethernet switching mode of the fabric interconnects is configured for end-host mode. You cannot connect to disjoint L2 networks if the Ethernet switching mode of the fabric interconnects is switch mode.

Symmetrical Configuration Is Recommended for High Availability

If a Cisco UCS domain is configured for high availability with two fabric interconnects, we recommend that both fabric interconnects are configured with the same set of VLANs.

VLAN Validity Criteria Are the Same for Uplink Ethernet Ports and Port Channels

The VLAN used for the disjoint L2 networks must be configured and assigned to an uplink Ethernet port or uplink Ethernet port channel. If the port or port channel does not include the VLAN, Cisco UCS Manager considers the VLAN invalid and does the following:

- Displays a configuration warning in the **Status Details** area for the server.
- Ignores the configuration for the port or port channel and drops all traffic for that VLAN.



Note

The validity criteria are the same for uplink Ethernet ports and uplink Ethernet port channels. Cisco UCS Manager does not differentiate between the two.

Overlapping VLANs Are Not Supported

Cisco UCS does not support overlapping VLANs in disjoint L2 networks. You must ensure that each VLAN only connects to one upstream disjoint L2 domain.

Each vNIC Can Only Communicate with One Disjoint L2 Network

A vNIC can only communicate with one disjoint L2 network. If a server needs to communicate with multiple disjoint L2 networks, you must configure a vNIC for each of those networks.

To communicate with more than two disjoint L2 networks, a server must have a Cisco VIC adapter that supports more than two vNICs.

Appliance Port Must Be Configured with the Same VLAN as Uplink Ethernet Port or Port Channel

For an appliance port to communicate with a disjoint L2 network, you must ensure that at least one uplink Ethernet port or port channel is in the same network and is therefore assigned to the same VLANs that are used by the appliance port. If Cisco UCS Manager cannot identify an uplink Ethernet port or port channel that includes all VLANs that carry traffic for an appliance port, the appliance port experiences a pinning failure and goes down.

For example, a Cisco UCS domain includes a global VLAN named `vlan500` with an ID of 500. `vlan500` is created as a global VLAN on the uplink Ethernet port. However, Cisco UCS Manager does not propagate this VLAN to appliance ports. To configure an appliance port with `vlan500`, you must create another VLAN named `vlan500` with an ID of 500 for the appliance port. You can create this duplicate VLAN in the **Appliances** node on the **LAN** tab of the Cisco UCS Manager GUI or the **eth-storage** scope in the Cisco UCS Manager CLI. If you are prompted to check for VLAN Overlap, accept the overlap and Cisco UCS Manager creates the duplicate VLAN for the appliance port.

Default VLAN 1 Cannot Be Configured Explicitly on an Uplink Ethernet Port or Port Channel

Cisco UCS Manager implicitly assigns default VLAN 1 to all uplink ports and port channels. Even if you do not configure any other VLANs, Cisco UCS uses default VLAN 1 to handle data traffic for all uplink ports and port channels.



Note

After you configure VLANs in a Cisco UCS domain, default VLAN 1 remains implicitly on all uplink ports and port channels. You cannot explicitly assign default VLAN 1 to an uplink port or port channel, nor can you remove it from an uplink port or port channel.

If you attempt to assign default VLAN 1 to a specific port or port channel, Cisco UCS Manager raises an Update Failed fault.

Therefore, if you configure a Cisco UCS domain for disjoint L2 networks, do not configure any vNICs with default VLAN 1 unless you want all data traffic for that server to be carried on all uplink Ethernet ports and port channels and sent to all upstream networks.

Pinning Considerations for Upstream Disjoint L2 Networks

Communication with an upstream disjoint L2 network requires that you ensure that the pinning is properly configured. Whether you implement soft pinning or hard pinning, a VLAN membership mismatch causes traffic for one or more VLANs to be dropped.

Soft Pinning

Soft pinning is the default behavior in Cisco UCS. If you plan to implement soft pinning, you do not need to create LAN pin groups to specify a pin target for a vNIC. Instead, Cisco UCS Manager pins the vNIC to an uplink Ethernet port or port channel according to VLAN membership criteria.

With soft pinning, Cisco UCS Manager validates data traffic from a vNIC against the VLAN membership of all uplink Ethernet ports and port channels. If you have configured disjoint L2 networks, Cisco UCS Manager must be able to find an uplink Ethernet port or port channel that is assigned to all VLANs on the vNIC. If no uplink Ethernet port or port channel is configured with all VLANs on the vNIC, Cisco UCS Manager does the following:

- Brings the link down.

- Drops the traffic for all of the VLANs on the vNIC.
- Raises the following faults:
 - Link Down
 - VIF Down

Cisco UCS Manager does not raise a fault or warning about the VLAN configuration.

For example, a vNIC on a server is configured with VLANs 101, 102, and 103. Interface 1/3 is assigned only to VLAN 102. Interfaces 1/1 and 1/2 are not explicitly assigned to a VLAN, which makes them available for traffic on VLANs 101 and 103. As a result of this configuration, the Cisco UCS domain does not include a border port interface that can carry traffic for all three VLANs for which the vNIC is configured. As a result, Cisco UCS Manager brings down the vNIC, drops traffic for all three VLANs on the vNIC, and raises the Link Down and VIF Down faults.

Hard Pinning

Hard pinning occurs when you use LAN pin groups to specify the pinning target for the traffic intended for the disjoint L2 networks. In turn, the uplink Ethernet port or port channel that is the pinning target must be configured to communicate with the appropriate disjoint L2 network.

With hard pinning, Cisco UCS Manager validates data traffic from a vNIC against the VLAN membership of all uplink Ethernet ports and port channels, and validates the LAN pin group configuration to ensure it includes the VLAN and the uplink Ethernet port or port channel. If the validation fails at any point, Cisco UCS Manager does the following:

- Raises a Pinning VLAN Mismatch fault with a severity of Warning.
- Drops traffic for the VLAN.
- Does not bring the link down, so that traffic for other VLANs can continue to flow along it.

For example, if you want to configure hard pinning for an upstream disjoint L2 network that uses VLAN 177, do the following:

- Create a LAN pin group with the uplink Ethernet port or port channel that carries the traffic for the disjoint L2 network.
- Configure at least one vNIC in the service profile with VLAN 177 and the LAN pin group.
- Assign VLAN 177 to an uplink Ethernet port or port channel included in the LAN pin group

If the configuration fails at any of these three points, then Cisco UCS Manager warns for a VLAN mismatch for VLAN 177 and drops the traffic for that VLAN only.

Configuring Cisco UCS for Upstream Disjoint L2 Networks

When you configure a Cisco UCS domain to connect with upstream disjoint L2 networks, you need to ensure that you complete all of the following steps.

Before You Begin

Before you begin this configuration, ensure that the ports on the fabric interconnects are properly cabled to support your disjoint L2 networks configuration.

Procedure

	Command or Action	Purpose
Step 1	Configure Ethernet switching mode for both fabric interconnects in Ethernet End-Host Mode.	The Ethernet switching mode must be in End-Host Mode for Cisco UCS to be able to communicate with upstream disjoint L2 networks. See Configuring Ethernet Switching Mode .
Step 2	Configure the ports and port channels that you require to carry traffic for the disjoint L2 networks.	See Configuring Ports and Port Channels , on page 71.
Step 3	Configure the LAN pin groups required to pin the traffic for the appropriate uplink Ethernet ports or port channels.	(Optional) See Configuring LAN Pin Groups , on page 275.
Step 4	Create one or more VLANs.	These can be named VLANs or private VLANs. For a cluster configuration, we recommend that you create the VLANs in Uplink Ethernet Mode and accessible to both fabric interconnects. See Configuring VLANs , on page 263.
Step 5	Assign the desired ports or port channels to the VLANs for the disjoint L2 networks.	When this step is completed, traffic for those VLANs can only be sent through the trunks for the assigned ports and/or port channels. Assigning Ports and Port Channels to VLANs , on page 305
Step 6	Ensure that the service profiles for all servers that need to communicate with the disjoint L2 networks include the correct LAN connectivity configuration to ensure the vNICs send the traffic to the appropriate VLAN.	You can complete this configuration through one or more vNIC templates or when you configure the networking options for the service profile. See Configuring Service Profiles , on page 465.

Assigning Ports and Port Channels to VLANs

Procedure

	Command or Action	Purpose
Step 1	UCS-A# scope eth-uplink	Enters Ethernet uplink mode.
Step 2	UCS-A /eth-uplink # scope vlan vlan-name	Enters Ethernet uplink VLAN mode for the specified VLAN.

	Command or Action	Purpose
Step 3	UCS-A /eth-uplink/vlan # create member-port <i>fabric-interconnect slot-id port-id</i>	Assigns the specified VLAN to the specified uplink Ethernet port.
Step 4	UCS-A /eth-uplink/vlan # create member-port-channel <i>fabric-interconnect member-port-chan-id</i>	Assigns the specified VLAN to the specified uplink Ethernet port channel.
Step 5	UCS-A /eth-uplink/vlan # commit-buffer	Commits the transaction to the system configuration. After a port or port channel is assigned to one or more VLANs, it is removed from all other VLANs.

The following example assigns uplink Ethernet ports to a named VLAN called VLAN100 on fabric interconnect A and commits the transaction:

```
UCS-A# scope eth-uplink
UCS-A /eth-uplink # scope vlan VLAN100
UCS-A /eth-uplink/vlan # create member-port a 2
UCS-A /eth-uplink/vlan # create member-port a 4
UCS-A /eth-uplink/vlan* # commit-buffer
UCS-A /eth-uplink/vlan #
```

Removing Ports and Port Channels from VLANs

Procedure

	Command or Action	Purpose
Step 1	UCS-A# scope eth-uplink	Enters Ethernet uplink mode.
Step 2	UCS-A /eth-uplink # scope vlan <i>vlan-name</i>	Enters Ethernet uplink VLAN mode for the specified VLAN.
Step 3	UCS-A /eth-uplink/vlan # delete member-port <i>fabric-interconnect slot-id port-id</i>	Deletes the specified Uplink Ethernet member port assignment from the VLAN.
Step 4	UCS-A /eth-uplink/vlan # delete member-port-channel <i>fabric-interconnect member-port-chan-id</i>	Deletes the specified Uplink Ethernet port channel assignment from the VLAN.
Step 5	UCS-A /eth-uplink/vlan # commit-buffer	Commits the transaction to the system configuration.

	Command or Action	Purpose
		Important If you remove all port or port channel interfaces from a VLAN, the VLAN returns to the default behavior and data traffic on that VLAN flows on all uplink ports and port channels. Depending upon the configuration in the Cisco UCS domain, this default behavior can cause Cisco UCS Manager to drop traffic for that VLAN. To avoid this occurrence, we recommend that you either assign at least one interface to the VLAN or delete the VLAN.

The following example deletes the association between uplink Ethernet port 2 on fabric interconnect A and the named VLAN called MyVLAN and commits the transaction:

```
UCS-A# scope eth-uplink
UCS-A /eth-uplink # scope vlan MyVLAN
UCS-A /eth-uplink/vlan # delete member-port a 2
UCS-A /eth-uplink/vlan* # commit-buffer
UCS-A /eth-uplink/vlan #
```

Viewing Ports and Port Channels Assigned to VLANs

Procedure

	Command or Action	Purpose
Step 1	UCS-A# scope eth-uplink	Enters Ethernet uplink mode.
Step 2	UCS-A /eth-uplink # scope vlan <i>vlan-name</i>	Enters Ethernet uplink VLAN mode for the specified VLAN.
Step 3	UCS-A /eth-uplink/vlan # show member-port [detail expand]	Shows member ports assigned to the specified VLAN.
Step 4	UCS-A /eth-uplink/vlan # show member-port-channel [detail expand]	Shows member port channels assigned to the specified VLAN.
Step 5	UCS-A /eth-uplink/vlan # commit-buffer	Commits the transaction to the system configuration.

The following example displays the full details for uplink Ethernet ports assigned to a named VLAN called MyVLAN:

```
UCS-A# scope eth-uplink
UCS-A /eth-uplink # scope vlan MyVLAN
UCS-A /eth-uplink/vlan # show member-port detail
Member Port:
    Fabric ID: A
```

```
Slot ID: 1  
Port ID: 2  
Mark Native Vlan: No  
UCS-A /eth-uplink/vlan #
```



PART **IV**

Storage Configuration

- [Configuring Named VSANs, page 311](#)
- [Configuring SAN Pin Groups, page 321](#)
- [Configuring WWN Pools, page 323](#)
- [Configuring Storage-Related Policies, page 327](#)



CHAPTER 22

Configuring Named VSANs

This chapter includes the following sections:

- [Named VSANs, page 311](#)
- [Fibre Channel Uplink Trunking for Named VSANs, page 312](#)
- [Guidelines and Recommendations for VSANs, page 312](#)
- [Creating a Named VSAN Accessible to Both Fabric Interconnects \(Fibre Channel Uplink Mode\), page 314](#)
- [Creating a Named VSAN Accessible to Both Fabric Interconnects \(Fibre Channel Storage Mode\), page 315](#)
- [Creating a Named VSAN Accessible to One Fabric Interconnect \(Fibre Channel Uplink Mode\), page 316](#)
- [Creating a Named VSAN Accessible to One Fabric Interconnect \(Fibre Channel Storage Mode\), page 317](#)
- [Deleting a Named VSAN, page 318](#)
- [Changing the VLAN ID for the FCoE Native VLAN for a Named VSAN, page 318](#)
- [Changing the VLAN ID for the FCoE Native VLAN for a Storage VSAN, page 319](#)
- [Enabling or Disabling Fibre Channel Uplink Trunking, page 319](#)

Named VSANs

A named VSAN creates a connection to a specific external SAN. The VSAN isolates traffic to that external SAN, including broadcast traffic. The traffic on one named VSAN knows that the traffic on another named VSAN exists, but cannot read or access that traffic.

Like a named VLAN, the name that you assign to a VSAN ID adds a layer of abstraction that allows you to globally update all servers associated with service profiles that use the named VSAN. You do not need to reconfigure the servers individually to maintain communication with the external SAN. You can create more than one named VSAN with the same VSAN ID.

Named VSANs in Cluster Configurations

In a cluster configuration, a named VSAN can be configured to be accessible only to the Fibre Channel uplink ports on one fabric interconnect or to the Fibre Channel uplink ports on both fabric interconnects.

Named VSANs and the FCoE VLAN ID

You must configure each named VSAN with an FCoE VLAN ID. This property determines which VLAN is used for transporting the VSAN and its Fibre Channel packets.

For FIP capable, converged network adapters, such as the Cisco UCS CNA M72KR-Q and the Cisco UCS CNA M72KR-E, the named VSAN must be configured with a named VLAN that is not the native VLAN for the FCoE VLAN ID. This configuration ensures that FCoE traffic can pass through these adapters.

In the following sample configuration, a service profile with a vNIC and vHBA mapped to fabric A is associated with a server that has FIP capable, converged network adapters:

- The vNIC is configured to use VLAN 10.
- VLAN 10 is also designated as the native VLAN for the vNIC.
- The vHBA is configured to use VSAN 2.
- Therefore, VSAN 2 cannot be configured with VLAN 10 as the FCoE VLAN ID. VSAN 2 can be mapped to any other VLAN configured on fabric A.

Fibre Channel Uplink Trunking for Named VSANs

You can configure Fibre Channel uplink trunking for the named VSANs on each fabric interconnect. If you enable trunking on a fabric interconnect, all named VSANs in a Cisco UCS domain are allowed on all Fibre Channel uplink ports on that fabric interconnect.

Guidelines and Recommendations for VSANs

The following guidelines and recommendations apply to all named VSANs, including storage VSANs.

VSAN 4079 is a Reserved VSAN ID

Do not configure a VSAN as 4079. This VSAN is reserved and cannot be used in either FC switch mode or FC end-host mode.

If you create a named VSAN with ID 4079, Cisco UCS Manager marks that VSAN with an error and raises a fault.

Reserved VSAN Range for Named VSANs in FC Switch Mode

If you plan to use FC switch mode in a Cisco UCS domain, do not configure VSANs with an ID in the range from 3040 to 4078.

VSANs in that range are not operational if the fabric interconnects are configured to operate in FC switch mode. Cisco UCS Manager marks that VSAN with an error and raises a fault.

Reserved VSAN Range for Named VSANs in FC End-Host Mode

If you plan to use FC end-host mode in a Cisco UCS domain, do not configure VSANs with an ID in the range from 3840 to 4079.

VSANs in that range are not operational if the following conditions exist in a Cisco UCS domain:

- The fabric interconnects are configured to operate in FC end-host mode.
- The Cisco UCS domain is configured with Fibre Channel trunking or SAN port channels.

If these configurations exist, Cisco UCS Manager does the following:

- 1 Renders all VSANs with an ID in the range from 3840 to 4079 non-operational.
- 2 Raises a fault against the non-operational VSANs.
- 3 Transfers all non-operational VSANs to the default VSAN.
- 4 Transfers all vHBAs associated with the non-operational VSANs to the default VSAN.

If you disable Fibre Channel trunking and delete any existing SAN port channels, Cisco UCS Manager returns all VSANs in the range from 3840 to 4078 to an operational state and restores any associated vHBAs back to those VSANs.

Range Restrictions for Named VSAN IDs in FC Switch Mode

If you plan to use FC switch mode in a Cisco UCS domain, do not configure VSANs in the range from 3040 to 4078.

When a fabric interconnect operating in FC switch mode is connected to MDS as the upstream switch, VSANs configured in Cisco UCS Manager in the range from 3040 to 4078 and assigned as port VSANs cannot be created in MDS. This configuration results in a possible port VSAN mismatch.

Guidelines for FCoE VLAN IDs



Note

FCoE VLANs in the SAN cloud and VLANs in the LAN cloud must have different IDs. Using the same ID for an FCoE VLAN in a VSAN and a VLAN results in a critical fault and traffic disruption for all vNICs and uplink ports using that FCoE VLAN. Ethernet traffic is dropped on any VLAN which has an ID that overlaps with an FCoE VLAN ID.

VLAN 4048 is user-configurable. However, Cisco UCS Manager uses VLAN 4048 for the following default values. If you want to assign 4048 to a VLAN, you must reconfigure these values:

The default FCoE VLAN varies according to the type of VSAN and whether Cisco UCS is a fresh installation or an upgrade, as follows:

- After an upgrade to Cisco UCS, release 2.0: The FCoE storage port native VLAN uses VLAN 4048 by default. If the default FCoE VSAN was set to use VLAN 1 before the upgrade, you must change it to a VLAN ID that is not used or reserved. For example, consider changing the default to 4049 if that VLAN ID is not in use.
- After a fresh install of Cisco UCS, release 2.0: The FCoE VLAN for the default VSAN uses VLAN 4048 by default. The FCoE storage port native VLAN uses VLAN 4049.

Creating a Named VSAN Accessible to Both Fabric Interconnects (Fibre Channel Uplink Mode)



Note FCoE VLANs in the SAN cloud and VLANs in the LAN cloud must have different IDs. Using the same ID for an FCoE VLAN in a VSAN and a VLAN results in a critical fault and traffic disruption for all vNICs and uplink ports using that FCoE VLAN. Ethernet traffic is dropped on any VLAN which has an ID that overlaps with an FCoE VLAN ID.

Procedure

	Command or Action	Purpose
Step 1	UCS-A# scope fc-uplink	Enters Fibre Channel uplink mode.
Step 2	UCS-A /fc-uplink # create vsan vsan-name vsan-id fcoe-id	Creates the specified named VSAN, specifies the VSAN name, VSAN ID and FCoE VLAN ID, and enters Fibre Channel uplink VSAN mode. <ul style="list-style-type: none"> • After an upgrade to Cisco UCS, release 2.0: The FCoE storage port native VLAN uses VLAN 4048 by default. If the default FCoE VSAN was set to use VLAN 1 before the upgrade, you must change it to a VLAN ID that is not used or reserved. For example, consider changing the default to 4049 if that VLAN ID is not in use. • After a fresh install of Cisco UCS, release 2.0: The FCoE VLAN for the default VSAN uses VLAN 4048 by default. The FCoE storage port native VLAN uses VLAN 4049.
Step 3	UCS-A /fc-uplink/vsan # commit-buffer	Commits the transaction to the system configuration.

The following example creates a named VSAN for both fabric interconnects, names the VSAN accounting, assigns the VSAN ID 2112, assigns the FCoE VLAN ID 4021, and commits the transaction:

```
UCS-A# scope fc-uplink
UCS-A /fc-uplink* # create vsan accounting 2112 4021
UCS-A /fc-uplink/vsan* # commit-buffer
UCS-A /fc-uplink/vsan #
```

Creating a Named VSAN Accessible to Both Fabric Interconnects (Fibre Channel Storage Mode)



Note FCoE VLANs in the SAN cloud and VLANs in the LAN cloud must have different IDs. Using the same ID for an FCoE VLAN in a VSAN and a VLAN results in a critical fault and traffic disruption for all vNICs and uplink ports using that FCoE VLAN. Ethernet traffic is dropped on any VLAN which has an ID that overlaps with an FCoE VLAN ID.

Procedure

	Command or Action	Purpose
Step 1	UCS-A# scope fc-storage	Enters Fibre Channel storage mode.
Step 2	UCS-A /fc-storage # create vsan vsan-name vsan-id fcoe-id	Creates the specified named VSAN, specifies the VSAN name, VSAN ID, and FCoE VLAN ID, and enters Fibre Channel storage VSAN mode. <ul style="list-style-type: none"> • After an upgrade to Cisco UCS, release 2.0: The FCoE storage port native VLAN uses VLAN 4048 by default. If the default FCoE VSAN was set to use VLAN 1 before the upgrade, you must change it to a VLAN ID that is not used or reserved. For example, consider changing the default to 4049 if that VLAN ID is not in use. • After a fresh install of Cisco UCS, release 2.0: The FCoE VLAN for the default VSAN uses VLAN 4048 by default. The FCoE storage port native VLAN uses VLAN 4049.
Step 3	UCS-A /fc-storage/vsan # create member-port {a b} slot-id port-id	Creates a member port; specifies the slot ID and port ID.
Step 4	UCS-A /fc-storage/vsan # commit-buffer	Commits the transaction to the system configuration.

The following example creates a named VSAN, names the VSAN finance, assigns the VSAN ID 3955, assigns the FCoE VLAN ID 4021, creates a member port and assigns the it to member port A, slot 1 port 40, and commits the transaction:

```
UCS-A# scope fc-storage
UCS-A /fc-storage/ # create VSAN finance 3955 4021
UCS-A /fc-storage/vsan # create member-port fcoe a 1 40
UCS-A /fc-storage/vsan/member-port* # commit-buffer
UCS-A /fc-storage/vsan/member-port #
```

Creating a Named VSAN Accessible to One Fabric Interconnect (Fibre Channel Uplink Mode)



Note FCoE VLANs in the SAN cloud and VLANs in the LAN cloud must have different IDs. Using the same ID for an FCoE VLAN in a VSAN and a VLAN results in a critical fault and traffic disruption for all vNICs and uplink ports using that FCoE VLAN. Ethernet traffic is dropped on any VLAN which has an ID that overlaps with an FCoE VLAN ID.

Procedure

	Command or Action	Purpose
Step 1	UCS-A# scope fc-uplink	Enters Fibre Channel uplink mode.
Step 2	UCS-A /fc-uplink # scope fabric {a b}	Enters Fibre Channel uplink fabric interconnect mode for the specified fabric interconnect (A or B).
Step 3	UCS-A /fc-uplink/fabric # create vsan vsan-name vsan-id fcoe-id	Creates the specified named VSAN, specifies the VSAN name, VSAN ID, and FCoE VLAN ID, and enters Fibre Channel uplink VSAN mode. <ul style="list-style-type: none"> • After an upgrade to Cisco UCS, release 2.0: The FCoE storage port native VLAN uses VLAN 4048 by default. If the default FCoE VSAN was set to use VLAN 1 before the upgrade, you must change it to a VLAN ID that is not used or reserved. For example, consider changing the default to 4049 if that VLAN ID is not in use. • After a fresh install of Cisco UCS, release 2.0: The FCoE VLAN for the default VSAN uses VLAN 4048 by default. The FCoE storage port native VLAN uses VLAN 4049.
Step 4	UCS-A /fc-uplink/fabric/vsan # commit-buffer	Commits the transaction to the system configuration.

The following example creates a named VSAN for fabric interconnect A, names the VSAN finance, assigns the VSAN ID 3955, assigns the FCoE VLAN ID 2221, and commits the transaction:

```
UCS-A# scope fc-uplink
UCS-A /fc-uplink # scope fabric a
UCS-A /fc-uplink/fabric # create vsan finance 3955 2221
UCS-A /fc-uplink/fabric/vsan* # commit-buffer
UCS-A /fc-uplink/fabric/vsan #
```

Creating a Named VSAN Accessible to One Fabric Interconnect (Fibre Channel Storage Mode)



Note FCoE VLANs in the SAN cloud and VLANs in the LAN cloud must have different IDs. Using the same ID for an FCoE VLAN in a VSAN and a VLAN results in a critical fault and traffic disruption for all vNICs and uplink ports using that FCoE VLAN. Ethernet traffic is dropped on any VLAN which has an ID that overlaps with an FCoE VLAN ID.

Procedure

	Command or Action	Purpose
Step 1	UCS-A# scope fc-storage	Enters Fibre Channel storage mode.
Step 2	UCS-A /fc-storage # scope fabric {a b}	Enters Fibre Channel storage mode for the specified fabric interconnect.
Step 3	UCS-A /fc-storage/fabric # create vsan vsan-name vsan-id fcoe-id	Creates the specified named VSAN, specifies the VSAN name, VSAN ID, and FCoE VLAN ID, and enters Fibre Channel storage VSAN mode. <ul style="list-style-type: none"> • After an upgrade to Cisco UCS, release 2.0: The FCoE storage port native VLAN uses VLAN 4048 by default. If the default FCoE VSAN was set to use VLAN 1 before the upgrade, you must change it to a VLAN ID that is not used or reserved. For example, consider changing the default to 4049 if that VLAN ID is not in use. • After a fresh install of Cisco UCS, release 2.0: The FCoE VLAN for the default VSAN uses VLAN 4048 by default. The FCoE storage port native VLAN uses VLAN 4049.
Step 4	UCS-A /fc-storage/fabric/vsan # create member-port {a b} slot-id port-id	Creates a member port on the specified VSAN.
Step 5	UCS-A /fc-storage/fabric/vsan # commit-buffer	Commits the transaction to the system configuration.

The following example creates a named VSAN on fabric A, names the VSAN finance, assigns the VSAN ID 3955, assigns the FCoE VLAN ID 2221, creates a member port and assigns the it to member port A, slot 1 port 40, and commits the transaction:

```
UCS-A# scope fc-storage
UCS-A /fc-storage/ # scope fabric a
UCS-A /fc-storage/fabric # create VSAN finance 3955 2221
UCS-A /fc-storage/fabric/vsan # create member-port a 1 40
```

```
UCS-A /fc-storage/fabric/vsan/member-port* # commit-buffer
UCS-A /fc-storage/fabric/vsan/member-port #
```

Deleting a Named VSAN

If Cisco UCS Manager includes a named VSAN with the same VSAN ID as the one you delete, the VSAN is not removed from the fabric interconnect configuration until all named VSANs with that ID are deleted.

Procedure

	Command or Action	Purpose
Step 1	UCS-A# scope fc-uplink	Enters Fibre Channel uplink mode.
Step 2	UCS-A /fc-uplink # delete vsan vsan-name	Deletes the specified named VSAN.
Step 3	UCS-A /fc-uplink # commit-buffer	Commits the transaction to the system configuration.

The following example deletes a named VSAN and commits the transaction:

```
UCS-A# scope fc-uplink
UCS-A /fc-uplink # delete vsan finance
UCS-A /fc-uplink* # commit-buffer
UCS-A /fc-uplink #
```

Changing the VLAN ID for the FCoE Native VLAN for a Named VSAN



Note

FCoE VLANs in the SAN cloud and VLANs in the LAN cloud must have different IDs. Using the same ID for an FCoE VLAN in a VSAN and a VLAN results in a critical fault and traffic disruption for all vNICs and uplink ports using that FCoE VLAN. Ethernet traffic is dropped on any VLAN which has an ID that overlaps with an FCoE VLAN ID.

Procedure

	Command or Action	Purpose
Step 1	UCS-A# scope fc-uplink	Enters Fibre Channel uplink mode.
Step 2	UCS-A /fc-uplink # scope vsan vsan-name	Enters VSAN mode for the specified named VSAN.
Step 3	UCS-A /fc-uplink/vsan # set fcoe-vlan fcoe-vlan-id	Sets the unique identifier assigned to the VLAN used for Fibre Channel connections.
Step 4	UCS-A /fc-uplink/vsan # commit-buffer	Commits the transaction to the system configuration.

The following example changes the VLAN ID for the FCoE Native VLAN on a named VSAN called finance to 4000 and commits the transaction:

```
UCS-A# scope fc-uplink
UCS-A /fc-uplink # scope vsan finance
UCS-A /fc-uplink/vsan # set fcoe-vlan 4000
UCS-A /fc-uplink/vsan* # commit-buffer
UCS-A /fc-uplink/vsan #
```

Changing the VLAN ID for the FCoE Native VLAN for a Storage VSAN



Note FCoE VLANs in the SAN cloud and VLANs in the LAN cloud must have different IDs. Using the same ID for an FCoE VLAN in a VSAN and a VLAN results in a critical fault and traffic disruption for all vNICs and uplink ports using that FCoE VLAN. Ethernet traffic is dropped on any VLAN which has an ID that overlaps with an FCoE VLAN ID.

Procedure

	Command or Action	Purpose
Step 1	UCS-A# scope fc-storage	Enters Fibre Channel storage mode.
Step 2	UCS-A /fc-storage # set fcoe-storage-native-vlan fcoe-id	Sets the unique identifier assigned to the VLAN used for Fibre Channel connections.
Step 3	UCS-A /fc-storage # commit-buffer	Commits the transaction to the system configuration.

The following example changes the VLAN ID for the FCoE Native VLAN on a storage VSAN called finance to 4000 and commits the transaction:

```
UCS-A# scope fc-storage
UCS-A /fc-storage # set fcoe-storage-native-vlan 4000
UCS-A /fc-storage* # commit-buffer
UCS-A /fc-storage #
```

Enabling or Disabling Fibre Channel Uplink Trunking



Note If the fabric interconnects are configured for Fibre Channel end-host mode, enabling Fibre Channel uplink trunking renders all VSANs with an ID in the range from 3840 to 4079 non-operational.

Procedure

	Command or Action	Purpose
Step 1	UCS-A# scope fc-uplink	Enters Fibre Channel uplink mode.
Step 2	UCS-A /fc-uplink # scope fabric {a b }	Enters Fibre Channel uplink mode for the specified fabric.
Step 3	UCS-A /fc-uplink/fabric # set uplink-trunking {enabled disabled }	Enables or disables uplink trunking.
Step 4	UCS-A /fc-uplink/fabric # commit-buffer	Commits the transaction to the system configuration.

The following example enables Fibre Channel uplink trunking for fabric A and commits the transaction:

```
UCS-A# scope fc-uplink
UCS-A /fc-uplink # scope fabric a
UCS-A /fc-uplink/fabric # set uplink-trunking enabled
UCS-A /fc-uplink/fabric* # commit-buffer
UCS-A /fc-uplink/fabric #
```



CHAPTER 23

Configuring SAN Pin Groups

This chapter includes the following sections:

- [SAN Pin Groups, page 321](#)
- [Configuring a SAN Pin Group, page 321](#)

SAN Pin Groups

Cisco UCS uses SAN pin groups to pin Fibre Channel traffic from a vHBA on a server to an uplink Fibre Channel port on the fabric interconnect. You can use this pinning to manage the distribution of traffic from the servers.



Note

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

To configure pinning for a server, you must include the SAN pin group in a vHBA policy. The vHBA policy is then included in the service profile assigned to that server. All traffic from the vHBA will travel through the I/O module to the specified uplink Fibre Channel port.

You can assign the same pin group to multiple vHBA policies. As a result, you do not need to manually pin the traffic for each vHBA.



Important

Changing the target interface for an existing SAN pin group disrupts traffic for all vHBAs which use that pin group. The fabric interconnect performs a log in and log out for the Fibre Channel protocols to re-pin the traffic.

Configuring a SAN Pin Group

In a system with two fabric interconnects, you can associate the pin group with only one fabric interconnect or with both fabric interconnects.

Procedure

	Command or Action	Purpose
Step 1	UCS-A# scope fc-uplink	Enters Fibre Channel uplink mode.
Step 2	UCS-A /fc-uplink # create pin-group <i>pin-group-name</i>	Creates a Fibre Channel (SAN) pin group with the specified name, and enters Fibre Channel uplink pin group mode.
Step 3	UCS-A /fc-uplink/pin-group # set descr <i>description</i>	(Optional) Provides a description for the pin group. Note If your description includes spaces, special characters, or punctuation, you must begin and end your description with quotation marks. The quotation marks will not appear in the description field of any show command output.
Step 4	UCS-A /fc-uplink/pin-group # set target {a b dual} port <i>slot-num / port-num</i>	(Optional) Sets the Fibre Channel pin target to the specified fabric and port.
Step 5	UCS-A /fc-uplink/pin-group # commit-buffer	Commits the transaction to the system configuration.

The following example creates a SAN pin group named fcpingroup12, provides a description for the pin group, sets the pin group target to slot 2, port 1, and commits the transaction:

```
UCS-A# scope fc-uplink
UCS-A /fc-uplink # create pin-group fcpingroup12
UCS-A /fc-uplink/pin-group* # set descr "This is my pin group #12"
UCS-A /fc-uplink/pin-group* # set target a port 2/1
UCS-A /fc-uplink/pin-group* # commit-buffer
UCS-A /fc-uplink/pin-group #
```

What to Do Next

Include the pin group in a vHBA template.



CHAPTER 24

Configuring WWN Pools

This chapter includes the following sections:

- [WWN Pools, page 323](#)
- [Creating a WWN Pool, page 324](#)
- [Deleting a WWN Pool, page 325](#)

WWN Pools

A 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 vHBA



Important

A WWN pool can include only WWNNs or WWPNS in the ranges from 20:00:00:00:00:00:00:00 to 20:FF:FF:FF:FF:FF:FF:FF or from 50:00:00:00:00:00:00:00 to 5F:FF:FF:FF:FF:FF:FF:FF. All other WWN ranges are reserved. To ensure the uniqueness of the Cisco UCS WWNNs and WWPNS in the SAN fabric, we recommend that you use the following WWN prefix for all blocks in a pool:
20:00:00:25:B5:XX:XX:XX

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

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

WWNN Pools

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

WWPN Pools

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

Creating a WWN Pool



Important

A WWN pool can include only WWNNs or WWPNs in the ranges from 20:00:00:00:00:00:00 to 20:FF:FF:FF:FF:FF:FF or from 50:00:00:00:00:00:00 to 5F:FF:FF:FF:FF:FF:FF. All other WWN ranges are reserved. To ensure the uniqueness of the Cisco UCS WWNNs and WWPNs in the SAN fabric, we recommend that you use the following WWN prefix for all blocks in a pool:
20:00:00:25:B5:XX:XX:XX

Procedure

	Command or Action	Purpose
Step 1	UCS-A# scope org <i>org-name</i>	Enters organization mode for the specified organization. To enter the root organization mode, type / as the <i>org-name</i>
Step 2	UCS-A /org # create wwn-pool <i>wwn-pool-name</i> { node-wwn-assignment port-wwn-assignment }	Creates a WWN pool with the specified name and purpose, and enters organization WWN pool mode. The purpose of the WWN pool can be one of the following: <ul style="list-style-type: none"> • To assign world wide node names (WWNNs) and world wide port names (WWPNs) • To assign only WWNNs • To assign only WWPNs
Step 3	UCS-A /org/wwn-pool # set descr <i>description</i>	(Optional) Provides a description for the WWN pool. Note If your description includes spaces, special characters, or punctuation, you must begin and end your description with quotation marks. The quotation marks will not appear in the description field of any show command output.
Step 4	UCS-A /org/wwn-pool # create block <i>first-wwn last-wwn</i>	(Optional) Creates a block (range) of WWNs, and enters organization WWN pool block mode. You must specify the first and last WWN in the block using the form <i>nn:nn:nn:nn:nn:nn:nn:nn</i> , with the WWNs separated by a space. Note A WWN pool can contain more than one WWN block. To create multiple WWN blocks, you must enter multiple create block commands from organization WWN pool mode.

	Command or Action	Purpose
Step 5	UCS-A /org/wwn-pool # create initiator <i>wwn wwn</i>	(Optional) Creates a single initiator, and enters organization WWN pool initiator mode. You must specify the initiator using the form <i>nn:nn:nn:nn:nn:nn:nn:nn</i> . Note A WWN pool can contain more than one initiator. To create multiple initiators, you must enter multiple create initiator commands from organization WWN pool mode.
Step 6	UCS-A /org/wwn-pool/block# commit-buffer	Commits the transaction to the system configuration.

The following example creates a WWN pool named sanpool, provides a description for the pool, specifies a block of WWNs and an initiator to be used for the pool, and commits the transaction:

```
UCS-A# scope org /
UCS-A /org # create wwn-pool sanpool node-wwn-assignment
UCS-A /org/wwn-pool* # set descr "This is my WWNN pool"
UCS-A /org/wwn-pool* # create block 23:00:00:05:AD:1E:00:01 23:00:00:05:AD:1E:01:00
UCS-A /org/wwn-pool/block* # exit
UCS-A /org/wwn-pool* # create initiator 23:00:00:05:AD:1E:02:00
UCS-A /org/wwn-pool/initiator* # commit-buffer
UCS-A /org/wwn-pool/initiator #
```

What to Do Next

Include the WWPN pool in a vHBA template.

Deleting a WWN Pool

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

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

Procedure

	Command or Action	Purpose
Step 1	UCS-A# scope org <i>org-name</i>	Enters organization mode for the specified organization. To enter the root organization mode, type / as the <i>org-name</i> .
Step 2	UCS-A /org # delete wwn-pool <i>pool-name</i>	Deletes the specified WWN pool.

	Command or Action	Purpose
Step 3	UCS-A /org # commit-buffer	Commits the transaction to the system configuration.

The following example deletes the WWN pool named pool4 and commits the transaction:

```
UCS-A# scope org /  
UCS-A /org # delete wwn-pool pool4  
UCS-A /org* # commit-buffer  
UCS-A /org #
```




CHAPTER 25

Configuring Storage-Related Policies

This chapter includes the following sections:

- [Configuring vHBA Templates, page 327](#)
- [Configuring Fibre Channel Adapter Policies, page 329](#)

Configuring vHBA Templates

vHBA Template

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

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

Configuring a vHBA Template

Procedure

	Command or Action	Purpose
Step 1	UCS-A# scope org <i>org-name</i>	Enters organization mode for the specified organization. To enter the root organization mode, type / as the <i>org-name</i> .
Step 2	UCS-A /org # create vhma-templ <i>vhba-templ-name</i> [fabric { a b }] [fc-if <i>vsan-name</i>]	Creates a vHBA template and enters organization vHBA template mode.
Step 3	UCS-A /org/vhma-templ # set descr <i>description</i>	(Optional) Provides a description for the vHBA template.
Step 4	UCS-A /org/vhma-templ # set fabric { a b }	(Optional) Specifies the fabric to use for the vHBA. If you did not specify the fabric when creating the vHBA template in

	Command or Action	Purpose
		Step 2, then you have the option to specify it with this command.
Step 5	UCS-A /org/vhba-templ # set fc-if <i>vsan-name</i>	(Optional) Specifies the Fibre Channel interface (named VSAN) to use for the vHBA template. If you did not specify the Fibre Channel interface when creating the vHBA template in Step 2, you have the option to specify it with this command.
Step 6	UCS-A /org/vhba-templ # set max-field-size <i>size-num</i>	Specifies the maximum size of the Fibre Channel frame payload (in bytes) that the vHBA supports.
Step 7	UCS-A /org/vhba-templ # set pin-group <i>group-name</i>	Specifies the pin group to use for the vHBA template.
Step 8	UCS-A /org/vhba-templ # set qos-policy <i>mac-pool-name</i>	Specifies the QoS policy to use for the vHBA template.
Step 9	UCS-A /org/vhba-templ # set stats-policy <i>policy-name</i>	Specifies the server and server component statistics threshold policy to use for the vHBA template.
Step 10	UCS-A /org/vhba-templ # set type { initial-template updating-template }	Specifies the vHBA template update type. If you do not want vHBA instances created from this template to be automatically updated when the template is updated, use the initial-template keyword; otherwise, use the updating-template keyword to ensure that all vHBA instances are updated when the vHBA template is updated.
Step 11	UCS-A /org/vhba-templ # set wwpn-pool <i>pool-name</i>	Specifies the WWPN pool to use for the vHBA template.
Step 12	UCS-A /org/vhba-templ # commit-buffer	Commits the transaction to the system configuration.

The following example configures a vHBA template and commits the transaction:

```
UCS-A# scope org /
UCS-A /org* # create vhma template VhmaTempFoo
UCS-A /org/vhma-templ* # set descr "This is a vHBA template example."
UCS-A /org/vhma-templ* # set fabric a
UCS-A /org/vhma-templ* # set fc-if accounting
UCS-A /org/vhma-templ* # set max-field-size 2112
UCS-A /org/vhma-templ* # set pin-group FcPinGroup12
UCS-A /org/vhma-templ* # set qos-policy policy34foo
UCS-A /org/vhma-templ* # set stats-policy ServStatsPolicy
UCS-A /org/vhma-templ* # set type updating-template
UCS-A /org/vhma-templ* # set wwpn-pool SanPool7
UCS-A /org/vhma-templ* # commit-buffer
UCS-A /org/vhma-templ #
```

Deleting a vHBA Template

Procedure

	Command or Action	Purpose
Step 1	UCS-A# scope org <i>org-name</i>	Enters organization mode for the specified organization. To enter the root organization mode, type <i>/</i> as the <i>org-name</i> .
Step 2	UCS-A /org # delete vhma-templ <i>vhba-templ-name</i>	Deletes the specified vHBA template.
Step 3	UCS-A /org # commit-buffer	Commits the transaction to the system configuration.

The following example deletes the vHBA template named VhbaTempFoo and commits the transaction:

```
UCS-A# scope org /
UCS-A /org # delete vhma template VhbaTempFoo
UCS-A /org* # commit-buffer
UCS-A /org #
```

Configuring Fibre Channel Adapter Policies

Ethernet and Fibre Channel Adapter Policies

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

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



Note

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

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

Operating System Specific Adapter Policies

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



Important

We recommend that you use the values in these policies for the applicable operating system. Do not modify any of the values in the default policies unless directed to do so by Cisco Technical Support.

However, if you are creating an Ethernet adapter policy for a Windows OS (instead of using the default Windows adapter policy), you must use the following formulas to calculate values that work with Windows:

$$\text{Completion Queues} = \text{Transmit Queues} + \text{Receive Queues}$$

$$\text{Interrupt Count} = (\text{Completion Queues} + 2) \text{ rounded up to nearest power of } 2$$

For example, if Transmit Queues = 1 and Receive Queues = 8 then:

$$\text{Completion Queues} = 1 + 8 = 9$$

$$\text{Interrupt Count} = (9 + 2) \text{ rounded up to the nearest power of } 2 = 16$$

Configuring a Fibre Channel Adapter Policy

Procedure

	Command or Action	Purpose
Step 1	UCS-A# scope org <i>org-name</i>	Enters organization mode for the specified organization. To enter the root organization mode, type / as the <i>org-name</i> .

	Command or Action	Purpose
Step 2	UCS-A /org # create fc-policy <i>policy-name</i>	Creates the specified Fibre Channel adapter policy and enters organization Fibre Channel policy mode.
Step 3	UCS-A /org/fc-policy # set descr <i>description</i>	(Optional) Provides a description for the policy. Note If your description includes spaces, special characters, or punctuation, you must begin and end your description with quotation marks. The quotation marks will not appear in the description field of any show command output.
Step 4	UCS-A /org/fc-policy # set error-recovery { fc-error-recovery { disabled enabled } link-down-timeout <i>timeout-msec</i> port-down-io-retry-count <i>retry-count</i> port-down-timeout <i>timeout-msec</i> }	(Optional) Configures the Fibre Channel error recovery.
Step 5	UCS-A /org/fc-policy # set interrupt mode { intx msi msi-x }	(Optional) Configures the driver interrupt mode.
Step 6	UCS-A /org/fc-policy # set port { io-throttle-count <i>throttle-count</i> max-luns <i>max-num</i> }	(Optional) Configures the Fibre Channel port.
Step 7	UCS-A /org/fc-policy # set port-f-logi { retries <i>retry-count</i> timeout <i>timeout-msec</i> }	(Optional) Configures the Fibre Channel port fabric login (FLOGI).
Step 8	UCS-A /org/fc-policy # set port-p-logi { retries <i>retry-count</i> timeout <i>timeout-msec</i> }	(Optional) Configures the Fibre Channel port-to-port login (PLOGI).
Step 9	UCS-A /org/fc-policy # set recv-queue { count <i>count</i> ring-size <i>size-num</i> }	(Optional) Configures the Fibre Channel receive queue.
Step 10	UCS-A /org/fc-policy # set scsi-io { count <i>count</i> ring-size <i>size-num</i> }	(Optional) Configures the Fibre Channel SCSI I/O.
Step 11	UCS-A /org/fc-policy # set trans-queue ring-size <i>size-num</i> }	(Optional) Configures the Fibre Channel transmit queue.
Step 12	UCS-A /org/fc-policy # commit-buffer	Commits the transaction to the system configuration.

The following example configures a Fibre Channel adapter policy and commits the transaction:

```
UCS-A# scope org /
UCS-A /org* # create fc-policy FcPolicy42
UCS-A /org/fc-policy* # set descr "This is a Fibre Channel adapter policy example."
UCS-A /org/fc-policy* # set error-recovery error-detect-timeout 2500
```

```

UCS-A /org/fc-policy* # set port max-luns 4
UCS-A /org/fc-policy* # set port-f-logi retries 250
UCS-A /org/fc-policy* # set port-p-logi timeout 5000
UCS-A /org/fc-policy* # set recv-queue count 1
UCS-A /org/fc-policy* # set scsi-io ring-size 256
UCS-A /org/fc-policy* # set trans-queue ring-size 256
UCS-A /org/fc-policy* # commit-buffer
UCS-A /org/fc-policy #

```

Deleting a Fibre Channel Adapter Policy

Procedure

	Command or Action	Purpose
Step 1	UCS-A# scope org <i>org-name</i>	Enters organization mode for the specified organization. To enter the root organization mode, type / as the <i>org-name</i> .
Step 2	UCS-A /org # delete fc-policy <i>policy-name</i>	Deletes the specified Fibre Channel adapter policy.
Step 3	UCS-A /org # commit-buffer	Commits the transaction to the system configuration.

The following example deletes the Fibre Channel adapter policy named FcPolicy42 and commits the transaction:

```

UCS-A# scope org /
UCS-A /org # delete fc-policy FcPolicy42
UCS-A /org* # commit-buffer
UCS-A /org #

```



PART **V**

Server Configuration

- [Configuring Server-Related Pools, page 335](#)
- [Setting the Management IP Address, page 339](#)
- [Configuring Server-Related Policies, page 347](#)
- [Configuring Server Boot, page 413](#)
- [Deferring Deployment of Service Profile Updates, page 453](#)
- [Configuring Service Profiles, page 465](#)
- [Managing Power in Cisco UCS, page 491](#)



CHAPTER 26

Configuring Server-Related Pools

This chapter includes the following sections:

- [Server Pool Configuration, page 335](#)
- [UUID Suffix Pool Configuration, page 337](#)

Server Pool Configuration

Server Pools

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

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

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

Creating a Server Pool

Procedure

	Command or Action	Purpose
Step 1	UCS-A# scope org <i>org-name</i>	Enters organization mode for the specified organization. To enter the root organization mode, type / as the <i>org-name</i> .

	Command or Action	Purpose
Step 2	UCS-A /org # create server-pool <i>server-pool-name</i>	Creates a server pool with the specified name, and enters organization server pool mode.
Step 3	UCS-A /org/server-pool # create server <i>chassis-num/slot-num</i>	Creates a server for the server pool. Note A server pool can contain more than one server. To create multiple servers for the pool, you must enter multiple create server commands from organization server pool mode.
Step 4	UCS-A /org/server-pool # commit-buffer	Commits the transaction to the system configuration.

The following example creates a server pool named ServPool2, creates two servers for the server pool, and commits the transaction:

```
UCS-A# scope org /
UCS-A /org # create server-pool ServPool2
UCS-A /org/server-pool* # create server 1/1
UCS-A /org/server-pool* # create server 1/4
UCS-A /org/server-pool* # commit-buffer
UCS-A /org/server-pool #
```

Deleting a Server Pool

Procedure

	Command or Action	Purpose
Step 1	UCS-A# scope org <i>org-name</i>	Enters organization mode for the specified organization. To enter the root organization mode, type / as the <i>org-name</i> .
Step 2	UCS-A /org # delete server-pool <i>server-pool-name</i>	Deletes the specified server pool.
Step 3	UCS-A /org # commit-buffer	Commits the transaction to the system configuration.

The following example deletes the server pool named ServPool2 and commits the transaction:

```
UCS-A# scope org /
UCS-A /org # delete server-pool ServPool2
UCS-A /org* # commit-buffer
UCS-A /org #
```

UUID Suffix Pool Configuration

UUID Suffix Pools

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

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

Creating a UUID Suffix Pool

Procedure

	Command or Action	Purpose
Step 1	UCS-A# scope org <i>org-name</i>	Enters organization mode for the specified organization. To enter the root organization mode, type / as the <i>org-name</i> .
Step 2	UCS-A /org # create uuid-suffix-pool <i>pool-name</i>	Creates a UUID suffix pool with the specified pool name and enters organization UUID suffix pool mode.
Step 3	UCS-A /org/uuid-suffix-pool # set descr <i>description</i>	(Optional) Provides a description for the UUID suffix pool. Note If your description includes spaces, special characters, or punctuation, you must begin and end your description with quotation marks. The quotation marks will not appear in the description field of any show command output.
Step 4	UCS-A /org/uuid-suffix-pool # create block <i>first-uuid last-uuid</i>	Creates a block (range) of UUID suffixes, and enters organization UUID suffix pool block mode. You must specify the first and last UUID suffixes in the block using the form <i>nnnn-nnnnnnnnnnnnn</i> , with the UUID suffixes separated by a space. Note A UUID suffix pool can contain more than one UUID suffix block. To create multiple blocks, you must enter multiple create block commands from organization UUID suffix pool mode.
Step 5	UCS-A /org/uuid-suffix-pool/block # commit-buffer	Commits the transaction to the system configuration.

The following example creates a UUID suffix pool named pool4, provides a description for the pool, specifies a block of UUID suffixes to be used for the pool, and commits the transaction:

```
UCS-A# scope org /
UCS-A /org # create uuid-suffix-pool pool4
UCS-A /org/uuid-suffix-pool* # set descr "This is UUID suffix pool 4"
UCS-A /org/uuid-suffix-pool* # create block 1000-000000000001 1000-000000000010
UCS-A /org/uuid-suffix-pool/block* # commit-buffer
UCS-A /org/uuid-suffix-pool/block #
```

What to Do Next

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

Deleting a UUID Suffix Pool

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

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

Procedure

	Command or Action	Purpose
Step 1	UCS-A# scope org <i>org-name</i>	Enters organization mode for the specified organization. To enter the root organization mode, type / as the <i>org-name</i> .
Step 2	UCS-A /org # delete uuid-suffix-pool <i>pool-name</i>	Deletes the specified UUID suffix pool.
Step 3	UCS-A /org # commit-buffer	Commits the transaction to the system configuration.

The following example deletes the UUID suffix pool named pool4 and commits the transaction:

```
UCS-A# scope org /
UCS-A /org # delete uuid-suffix-pool pool4
UCS-A /org* # commit-buffer
UCS-A /org #
```



CHAPTER 27

Setting the Management IP Address

This chapter includes the following sections:

- [Management IP Address, page 339](#)
- [Configuring the Management IP Address on a Blade Server, page 340](#)
- [Configuring the Management IP Address on a Rack Server, page 341](#)
- [Setting the Management IP Address on a Service Profile or Service Profile Template, page 343](#)
- [Configuring the Management IP Pool, page 344](#)

Management IP Address

Each server in a Cisco UCS domain must have a management IP address assigned to its Cisco Integrated Management Controller (CIMC) or to the service profile associated with the server. Cisco UCS Manager uses this IP address for external access that terminates in the CIMC. This external access can be through one of the following:

- KVM console
- Serial over LAN
- An IPMI tool

The management IP address used to access the CIMC on a server can be one of the following:

- A static IPv4 address assigned directly to the server.
- A static IPv4 address assigned to a service profile. You cannot configure a service profile template with a static IP address.
- An IP address drawn from the management IP address pool and assigned to a service profile or service profile template.

You can assign a management IP address to each CIMC on the server and to the service profile associated with the server. If you do so, you must use different IP addresses for each of them.

**Note**

You cannot assign a static IP address to a server or service profile if that IP address has already been assigned to a server or service profile in the Cisco UCS domain. If you attempt to do so, Cisco UCS Manager warns you that the IP address is already in use and rejects the configuration.

A management IP address that is assigned to a service profile moves with the service profile. If a KVM or SoL session is active when you migrate the service profile to another server, Cisco UCS Manager terminates that session and does not restart it after the migration is completed. You configure this IP address when you create or modify a service profile.

Configuring the Management IP Address on a Blade Server

Configuring a Blade Server to Use a Static IP Address

Procedure

	Command or Action	Purpose
Step 1	UCS-A# scope server <i>chassis-id / blade-id</i>	Enters chassis server mode for the specified server.
Step 2	UCS-A /chassis/server # scope cimc	Enters chassis server CIMC mode.
Step 3	UCS-A /chassis/server/cimc # create ext-static-ip	Creates a static management IP address for the specified server.
Step 4	UCS-A /chassis/server/cimc/ext-static-ip # set addr <i>ip-addr</i>	Specifies the static IPv4 address to be assigned to the server.
Step 5	UCS-A /chassis/server/cimc/ext-static-ip # set default-gw <i>ip-addr</i>	Specifies the default gateway that the IP address should use.
Step 6	UCS-A /chassis/server/cimc/ext-static-ip # set subnet <i>ip-addr</i>	Specifies the subnet mask for the IP address.
Step 7	UCS-A /chassis/server/cimc/ext-static-ip # commit-buffer	Commits the transaction to the system configuration.

The following example configures a static management IP address for chassis 1 server 1, sets the static IPv4 address, sets the default gateway, sets the subnet mask, and commits the transaction:

```
UCS-A# scope server 1/1
UCS-A /chassis/server # scope cimc
UCS-A /chassis/server/cimc # create ext-static-ip
UCS-A /chassis/server/cimc/ext-static-ip* # set addr 192.168.10.10
UCS-A /chassis/server/cimc/ext-static-ip* # set default-gw 192.168.10.1
UCS-A /chassis/server/cimc/ext-static-ip* # set subnet 255.255.255.0
UCS-A /chassis/server/cimc/ext-static-ip* # commit-buffer
UCS-A /chassis/server/cimc/ext-static-ip #
```

Configuring a Blade Server to Use the Management IP Pool

Deleting the static management IP address returns the specified server to the management IP pool.

Procedure

	Command or Action	Purpose
Step 1	UCS-A# scope server <i>chassis-id / blade-id</i>	Enters chassis server mode for the specified server.
Step 2	UCS-A /chassis/server # scope cimc	Enters chassis server CIMC mode.
Step 3	UCS-A /chassis/server/cimc # delete ext-static-ip	Deletes the external static IP address and returns the blade server to the management IP pool.
Step 4	UCS-A /chassis/server/cimc/ # commit-buffer	Commits the transaction to the system configuration.

The following example deletes the static management IP address for chassis 1 server 1 and commits the transaction:

```
UCS-A# scope server 1/1
UCS-A /chassis/server # scope cimc
UCS-A /chassis/server/cimc # delete ext-static-ip
UCS-A /chassis/server/cimc* # commit-buffer
UCS-A /chassis/server/cimc/ #
```

Configuring the Management IP Address on a Rack Server

Configuring a Rack Server to Use a Static IP Address

Procedure

	Command or Action	Purpose
Step 1	UCS-A# scope server <i>blade-id</i>	Enters server mode for the specified server.
Step 2	UCS-A /server # scope cimc	Enters server CIMC mode.
Step 3	UCS-A /server/cimc # create ext-static-ip	Creates a static management IP address for the specified server.
Step 4	UCS-A /server/cimc/ext-static-ip # set addr <i>ip-addr</i>	Specifies the static IPv4 address to be assigned to the server.
Step 5	UCS-A /server/cimc/ext-static-ip # set default-gw <i>ip-addr</i>	Specifies the default gateway that the IP address should use.

	Command or Action	Purpose
Step 6	UCS-A /server/cimc/ext-static-ip # set subnet ip-addr	Specifies the subnet mask for the IP address.
Step 7	UCS-A /server/cimc/ext-static-ip # commit-buffer	Commits the transaction to the system configuration.

The following example configures a static management IP address for rack server 1, sets the static IPv4 address, sets the default gateway, sets the subnet mask, and commits the transaction:

```
UCS-A# scope server 1
UCS-A /server # scope cimc
UCS-A /server/cimc # create ext-static-ip
UCS-A /server/cimc/ext-static-ip* # set addr 192.168.10.10
UCS-A /server/cimc/ext-static-ip* # set default-gw 192.168.10.1
UCS-A /server/cimc/ext-static-ip* # set subnet 255.255.255.0
UCS-A /server/cimc/ext-static-ip* # commit-buffer
UCS-A /server/cimc/ext-static-ip #
```

Configuring a Rack Server to Use the Management IP Pool

Deleting the static management IP address returns the specified server to the management IP pool.

Procedure

	Command or Action	Purpose
Step 1	UCS-A# scope server blade-id	Enters server mode for the specified server.
Step 2	UCS-A /server # scope cimc	Enters server CIMC mode.
Step 3	UCS-A /server/cimc # delete ext-static-ip	Deletes the external static IP address and returns the rack server to the management IP pool.
Step 4	UCS-A /server/cimc/ # commit-buffer	Commits the transaction to the system configuration.

The following example deletes the static management IP address for rack server 1 and commits the transaction:

```
UCS-A# scope server 1
UCS-A /server # scope cimc
UCS-A /server/cimc # delete ext-static-ip
UCS-A /server/cimc* # commit-buffer
UCS-A /server/cimc/ #
```


Setting the Management IP Address on a Service Profile or Service Profile Template

Procedure

	Command or Action	Purpose
Step 1	UCS-A# scope org <i>org-name</i>	Enters organization mode for the specified organization. To enter the root organization mode, type / as the org-name.
Step 2	UCS-A /org # scope service-profile <i>profile-name</i>	Enters organization service profile mode for the specified service.
Step 3	UCS-A /org/service-profile # set ext-mgmt-ip-state { none pooled static }	Specifies how the management IP address will be assigned to the service profile. You can set the management IP address policy using the following options: <ul style="list-style-type: none"> • None--The service profile is not assigned an IP address. • Pooled--The service profile is assigned an IP address from the management IP pool. • Static--The service profile is assigned the configured static IP address. <p>Note Setting the ext-management-ip-state to static for a service profile template is not supported and will result in an error.</p>
Step 4	UCS-A /org/service-profile # commit-buffer	Commits the transaction to the system configuration.

The following example sets the management IP address policy for a service profile called accounting to static and then commits the transaction:

```
UCS-A# scope org /
UCS-A /org # scope service-profile accounting
UCS-A /org/service-profile # set ext-mgmt-ip-state static
UCS-A /org/service-profile* # commit-buffer
UCS-A /org/service-profile #
```

What to Do Next

If you have set the management IP address to static, configure a server to use a static IP address.

Configuring the Management IP Pool

Management IP Pool

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

You can configure service profiles and service profile templates to use IP addresses from the management IP pool. You cannot configure servers to use the management IP pool.

All IP addresses in the management IP pool must be in the same subnet as the IP address of the fabric interconnect.



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

Configuring an IP Address Block for the Management IP Pool

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

Procedure

	Command or Action	Purpose
Step 1	UCS-A# scope org /	Enters root organization mode.
Step 2	UCS-A /org # scope ip-pool ext-mgmt	Enters organization IP pool mode. Note You cannot create (or delete) a management IP pool. You can only enter (scope to) the existing default pool.
Step 3	UCS-A /org/ip-pool # set descr description	(Optional) Provides a description for the management IP pool. This description applies to all address blocks in the management IP pool. Note If your description includes spaces, special characters, or punctuation, you must begin and end your description with quotation marks. The quotation marks will not appear in the description field of any show command output.
Step 4	UCS-A /org/ip-pool # create block first-ip-addr last-ip-addr gateway-ip-addr subnet-mask	Creates a block (range) of IP addresses, and enters organization IP pool block mode. You must specify the first and last IP addresses in the address range, the gateway IP address, and subnet mask.

	Command or Action	Purpose
		Note A IP pool can contain more than one IP address block. To create multiple IP address blocks, you must enter multiple create block commands from organization IP pool mode.
Step 5	UCS-A /org/ip-pool/block # commit-buffer	Commits the transaction to the system configuration.

The following example configures an IP address block for the management IP pool and commits the transaction:

```
UCS-A# scope org /
UCS-A /org # scope ip-pool ext-mgmt
UCS-A /org/ip-pool* # set descr "This is a management IP pool example."
UCS-A /org/ip-pool* # create block 192.168.100.1 192.168.100.200 192.168.100.10 255.255.255.0
UCS-A /org/ip-pool/block* # commit-buffer
UCS-A /org/ip-pool/block #
```

What to Do Next

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

Deleting an IP Address Block from the Management IP Pool

Procedure

	Command or Action	Purpose
Step 1	UCS-A# scope org <i>org-name</i>	Enters organization mode for the specified organization. To enter the root organization mode, type <i>/</i> as the <i>org-name</i> .
Step 2	UCS-A /org # scope ip-pool ext-mgmt	Enters the management IP pool.
Step 3	UCS-A /org/ip-pool # delete block <i>first-ip-addr last-ip-addr</i>	Deletes the specified block (range) of IP addresses.
Step 4	UCS-A /org/ip-pool # commit-buffer	Commits the transaction to the system configuration.

The following example configures an IP address block for the management IP pool and commits the transaction:

```
UCS-A# scope org /
UCS-A /org # scope ip-pool ext-mgmt
UCS-A /org/ip-pool # delete block 192.168.100.1 192.168.100.200
UCS-A /org/ip-pool* # commit-buffer
UCS-A /org/ip-pool #
```




CHAPTER 28

Configuring Server-Related Policies

This chapter includes the following sections:

- [Configuring BIOS Settings, page 347](#)
- [Configuring IPMI Access Profiles, page 372](#)
- [Configuring Local Disk Configuration Policies, page 376](#)
- [Configuring Scrub Policies, page 381](#)
- [Configuring Serial over LAN Policies, page 383](#)
- [Configuring Server Autoconfiguration Policies, page 385](#)
- [Configuring Server Discovery Policies, page 387](#)
- [Configuring Server Inheritance Policies, page 389](#)
- [Configuring Server Pool Policies, page 391](#)
- [Configuring Server Pool Policy Qualifications, page 393](#)
- [Configuring vNIC/vHBA Placement Policies, page 405](#)

Configuring BIOS Settings

Server BIOS Settings

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

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

Depending upon the needs of the data center, you can configure BIOS policies for some service profiles and use the BIOS defaults in other service profiles in the same Cisco UCS domain, or you can use only one of them. You can also use Cisco UCS Manager to view the actual BIOS settings on a server and determine whether they are meeting current needs.

**Note**

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

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

Main BIOS Settings

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

Name	Description
Reboot on BIOS Settings Change set reboot-on-update	When the server is rebooted after you change one or more BIOS settings. yes —If you enable this setting, the server is rebooted according to the maintenance policy in the server's service profile. For example, if the maintenance policy requires user acknowledgment, the server is not rebooted and the BIOS changes are not applied until a user acknowledges the pending activity. no —If you do not enable this setting, the BIOS changes are not applied until the next time the server is rebooted, whether as a result of another server configuration change or a manual reboot.
Quiet Boot set quiet-boot-config quiet-boot	What the BIOS displays during Power On Self-Test (POST). This can be one of the following: <ul style="list-style-type: none"> • disabled—The BIOS displays all messages and Option ROM information during boot. • enabled—The BIOS displays the logo screen, but does not display any messages or Option ROM information during boot. • platform-default—The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.

Name	Description
Post Error Pause set post-error-pause-config post-error-pause	What happens when the server encounters a critical error during POST. This can be one of the following: <ul style="list-style-type: none"> • disabled—The BIOS continues to attempt to boot the server. • enabled—The BIOS pauses the attempt to boot the server and opens the Error Manager when a critical error occurs during POST. • platform-default—The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.
Resume Ac On Power Loss set resume-ac-on-power-loss-config resume-action	How the server behaves when power is restored after an unexpected power loss. This can be one of the following: <ul style="list-style-type: none"> • stay-off—The server remains off until manually powered on. • last-state—The server is powered on and the system attempts to restore its last state. • reset—The server is powered on and automatically reset. • platform-default—The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.
Front Panel Lockout set front-panel-lockout-config front-panel-lockout	Whether the power and reset buttons on the front panel are ignored by the server. This can be one of the following: <ul style="list-style-type: none"> • disabled—The power and reset buttons on the front panel are active and can be used to affect the server. • enabled—The power and reset buttons are locked out. The server can only be reset or powered on or off from the CIMC GUI. • platform-default—The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.

Name	Description
<p>ACPI10 Support set acpi10-support-config acpi10-support</p>	<p>Whether the BIOS publishes the ACPI 1.0 version of FADT in the Root System Description table. This version may be required for compatibility with OS versions that only support ACPI 1.0. This can be one of the following:</p> <ul style="list-style-type: none"> • disabled—ACPI 1.0 version is not published. • enabled—ACPI 1.0 version is published. • platform-default—The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.

Processor BIOS Settings

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

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

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

Name	Description
<p>Core Multi Processing set core-multi-processing-config multi-processing</p>	<p>Sets the state of logical processor cores in a package. If you disable this setting, Hyper Threading is also disabled. This can be one of the following:</p> <ul style="list-style-type: none"> • all—Enables multi processing on all logical processor cores. • 1 through 10—Specifies the number of logical processor cores that can run on the server. To disable multi processing and have only one logical processor core running on the server, select 1. • platform-default—The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor. <p>We recommend that you contact your operating system vendor to make sure the operating system supports this feature.</p>
<p>Execute Disabled Bit set execute-disable bit</p>	<p>Classifies memory areas on the server to specify where where application code can execute. As a result of this classification, the processor disables code execution if a malicious worm attempts to insert code in the buffer. This setting helps to prevent damage, worm propagation, and certain classes of malicious buffer overflow attacks. This can be one of the following:</p> <ul style="list-style-type: none"> • disabled—The processor does not classify memory areas. • enabled—The processor classifies memory areas. • platform-default—The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor. <p>We recommend that you contact your operating system vendor to make sure the operating system supports this feature.</p>
<p>Virtualization Technology (VT) set intel-vt-config vt</p>	<p>Whether the processor uses Intel Virtualization Technology, which allows a platform to run multiple operating systems and applications in independent partitions. This can be one of the following:</p> <ul style="list-style-type: none"> • disabled—The processor does not permit virtualization. • enabled—The processor allows multiple operating systems in independent partitions. • platform-default—The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor. <p>Note If you change this option, you must power cycle the server before the setting takes effect.</p>

Name	Description
Direct Cache Access set direct-cache-access-config access	<p>Allows processors to increase I/O performance by placing data from I/O devices directly into the processor cache. This setting helps to reduce cache misses. This can be one of the following:</p> <ul style="list-style-type: none"> • disabled—Data from I/O devices is not placed directly into the processor cache. • enabled—Data from I/O devices is placed directly into the processor cache. • platform-default—The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.
Processor C State set processor-c-state-config c-state	<p>Whether the system can enter a power savings mode during idle periods. This can be one of the following:</p> <ul style="list-style-type: none"> • disabled—The system remains in high performance state even when idle. • enabled—The system can reduce power to system components such as the DIMMs and CPUs. • platform-default—The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor. <p>We recommend that you contact your operating system vendor to make sure the operating system supports this feature.</p>
Processor C1E set processor-c1e-config c1e	<p>Allows the processor to transition to its minimum frequency upon entering C1. This setting does not take effect until after you have rebooted the server. This can be one of the following:</p> <ul style="list-style-type: none"> • disabled—The CPU continues to run at its maximum frequency in C1 state. • enabled—The CPU transitions to its minimum frequency. This option saves the maximum amount of power in C1 state. • platform-default—The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.

Name	Description
<p>Processor C3 Report set processor-c3-report-config processor-c3-report</p>	<p>Whether the processor sends the C3 report to the operating system. This can be one of the following:</p> <ul style="list-style-type: none"> • disabled—The processor does not send the C3 report. • acpi-c2—The processor sends the C3 report using the ACPI C2 format. • acpi-c3—The processor sends the C3 report using the ACPI C3 format. • platform-default—The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor. <p>On the B440 server, the BIOS Setup menu uses enabled and disabled for these options. If you specify acpi-c2 or acpi-c2, the server sets the BIOS value for that option to enabled.</p>
<p>Processor C6 Report set processor-c6-report-config processor-c6-report</p>	<p>Whether the processor sends the C6 report to the operating system. This can be one of the following:</p> <ul style="list-style-type: none"> • disabled—The processor does not send the C6 report. • enabled—The processor sends the C6 report. • platform-default—The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.
<p>Processor C7 Report set processor-c7-report-config processor-c7-report</p>	<p>Whether the processor sends the C7 report to the operating system. This can be one of the following:</p> <ul style="list-style-type: none"> • disabled—The processor does not send the C7 report. • enabled—The processor sends the C7 report. • platform-default—The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.

Name	Description
<p>CPU Performance set cpu-performance-config cpu-performance</p>	<p>Sets the CPU performance profile for the server. This can be one of the following:</p> <ul style="list-style-type: none"> • enterprise—All prefetchers and data reuse are disabled. • high-throughput—All prefetchers are enabled, and data reuse is disabled. • hpc—All prefetchers and data reuse are enabled. This setting is also known as high performance computing. • platform-default—The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.
<p>Max Variable MTRR Setting set max-variable-mtrr-setting-config processor-mtrr</p>	<p>Allows you to select the number of MTRR variables. This can be one of the following:</p> <ul style="list-style-type: none"> • auto-max—The BIOS uses the default value for the processor. • 8—The BIOS uses the number specified for the variable MTRR. • platform-default—The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.

Name	Description
Package C State Limit set package-c-state-limit-config package-c-state-limit	<p>The amount of power available to the server components when they are idle. This can be one of the following:</p> <ul style="list-style-type: none"> • c0—The server provides all server components with full power at all times. This option maintains the highest level of performance and requires the greatest amount of power. • c1—When the CPU is idle, the system slightly reduces the power consumption. This option requires less power than C0 and allows the server to return quickly to high performance mode. • c3—When the CPU is idle, the system reduces the power consumption further than with the C1 option. This requires less power than C1 or C0, but it takes the server slightly longer to return to high performance mode. • c6—When the CPU is idle, the system reduces the power consumption further than with the C3 option. This option saves more power than C0, C1, or C3, but there may be performance issues until the server returns to full power. • no-limit—The server may enter any available C state. • platform-default—The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.

Intel Directed I/O BIOS Settings

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

Name	Description
VT for Directed IO set intel-vt-directed-io-config vtd	<p>Whether the processor uses Intel Virtualization Technology for Directed I/O (VT-d). This can be one of the following:</p> <ul style="list-style-type: none"> • disabled—The processor does not use virtualization technology. • enabled—The processor uses virtualization technology. • platform-default—The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor. <p>Note This option must be enabled if you want to change any of the other Intel Directed I/O BIOS settings.</p>

Name	Description
Interrupt Remap set intel-vt-directed-io-config interrupt-remapping	Whether the processor supports Intel VT-d Interrupt Remapping. This can be one of the following: <ul style="list-style-type: none"> • disabled—The processor does not support remapping. • enabled—The processor uses VT-d Interrupt Remapping as required. • platform-default—The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.
Coherency Support set intel-vt-directed-io-config coherency-support	Whether the processor supports Intel VT-d Coherency. This can be one of the following: <ul style="list-style-type: none"> • disabled—The processor does not support coherency. • enabled—The processor uses VT-d Coherency as required. • platform-default—The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.
ATS Support set intel-vt-directed-io-config ats-support	Whether the processor supports Intel VT-d Address Translation Services (ATS). This can be one of the following: <ul style="list-style-type: none"> • disabled—The processor does not support ATS. • enabled—The processor uses VT-d ATS as required. • platform-default—The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.
Pass Through DMA Support set intel-vt-directed-io-config passthrough-dma	Whether the processor supports Intel VT-d Pass-through DMA. This can be one of the following: <ul style="list-style-type: none"> • disabled—The processor does not support pass-through DMA. • enabled—The processor uses VT-d Pass-through DMA as required. • platform-default—The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.

RAS Memory BIOS Settings

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

Name	Description
<p>Memory RAS Config set memory-ras-config ras-config</p>	<p>How the memory reliability, availability, and serviceability (RAS) is configured for the server. This can be one of the following:</p> <ul style="list-style-type: none"> • maximum performance—System performance is optimized. • mirroring—System reliability is optimized by using half the system memory as backup. • lockstep—If the DIMM pairs in the server have an identical type, size, and organization and are populated across the SMI channels, you can enable lockstep mode to minimize memory access latency and provide better performance. Lockstep is enabled by default for B440 servers. • sparing—System reliability is enhanced with a degree of memory redundancy while making more memory available to the operating system than mirroring. • platform-default—The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.
<p>NUMA set numa-config numa-optimization</p>	<p>Whether the BIOS supports NUMA. This can be one of the following:</p> <ul style="list-style-type: none"> • disabled—The BIOS does not support NUMA. • enabled—The BIOS includes the ACPI tables that are required for NUMA-aware operating systems. If you enable this option, the system must disable Inter-Socket Memory interleaving on some platforms. • platform-default—The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.

Name	Description
Mirroring Mode set memory-mirroring-mode mirroring-mode	<p>Memory mirroring enhances system reliability by keeping two identical data images in memory.</p> <p>This option is only available if you choose the mirroring option for Memory RAS Config. It can be one of the following:</p> <ul style="list-style-type: none"> • inter-socket—Memory is mirrored between two Integrated Memory Controllers (IMCs) across CPU sockets. • intra-socket—One IMC is mirrored with another IMC in the same socket. • platform-default—The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.
Sparing Mode set memory-sparing-mode sparing-mode	<p>Sparing optimizes reliability by holding memory in reserve so that it can be used in case other DIMMs fail. This option provides some memory redundancy, but does not provide as much redundancy as mirroring. The available sparing modes depend on the current memory population.</p> <p>This option is only available if you choose sparing option for Memory RAS Config. It can be one of the following:</p> <ul style="list-style-type: none"> • dimmm-sparing—One DIMM is held in reserve. If a DIMM fails, the contents of a failing DIMM are transferred to the spare DIMM. • rank-sparing—A spare rank of DIMMs is held in reserve. If a rank of DIMMs fails, the contents of the failing rank are transferred to the spare rank. • platform-default—The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.
LV DDR Mode set lv-dimm-support-config lv-ddr-mode	<p>Whether the system prioritizes low voltage or high frequency memory operations. This can be one of the following:</p> <ul style="list-style-type: none"> • power-saving-mode—The system prioritizes low voltage memory operations over high frequency memory operations. This mode may lower memory frequency in order to keep the voltage low. • performance-mode—The system prioritizes high frequency operations over low voltage operations. • platform-default—The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.

Serial Port BIOS Settings

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

Name	Description
Serial Port A set serial-port-a-config serial-port-a	Whether serial port A is enabled or disabled. This can be one of the following: <ul style="list-style-type: none"> • disabled—The serial port is disabled. • enabled—The serial port is enabled. • platform-default—The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.

USB BIOS Settings

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

Name	Description
Make Device Non Bootable set usb-boot-config make-device-non-bootable	Whether the server can boot from a USB device. This can be one of the following: <ul style="list-style-type: none"> • disabled—The server can boot from a USB device. • enabled—The server cannot boot from a USB device. • platform-default—The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.

Name	Description
<p>USB System Idle Power Optimizing Setting</p> <p>set usb-system-idle-power-optimizing-setting-config usb-idle-power-optimizing</p>	<p>Whether the USB System Idle Power Optimizing setting is used to reduce USB EHCI idle power consumption. Depending upon the value you choose, this setting can have an impact on performance. This can be one of the following:</p> <ul style="list-style-type: none"> • high-performance—The USB System Idle Power Optimizing setting is disabled, because optimal performance is preferred over power savings. Selecting this option can significantly improve performance. We recommend you select this option unless your site has server power restrictions. • lower-idle-power—The USB System Idle Power Optimizing setting is enabled, because power savings are preferred over optimal performance. • platform-default—The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.
<p>USB Front Panel Access Lock</p> <p>set usb-front-panel-access-lock-config usb-front-panel-lock</p>	<p>USB front panel lock is configured to enable or disable the front panel access to USB ports. This can be one of the following:</p> <ul style="list-style-type: none"> • disabled • enabled • platform-default—The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.

PCI Configuration BIOS Settings

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

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

Boot Options BIOS Settings

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

Name	Description
Boot Option Retry set boot-option-retry-config retry	Whether the BIOS retries NON-EFI based boot options without waiting for user input. This can be one of the following: <ul style="list-style-type: none"> • disabled—Waits for user input before retrying NON-EFI based boot options. • enabled—Continually retries NON-EFI based boot options without waiting for user input. • platform-default—The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.

Name	Description
Intel Entry SAS RAID set intel-entry-sas-raid-config sas-raid	Whether the Intel SAS Entry RAID Module is enabled. This can be one of the following: <ul style="list-style-type: none"> • disabled—The Intel SAS Entry RAID Module is disabled. • enabled—The Intel SAS Entry RAID Module is enabled. • platform-default—The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.
Intel Entry SAS RAID Module set intel-entry-sas-raid-config sas-raid-module	How the Intel SAS Entry RAID Module is configured. This can be one of the following: <ul style="list-style-type: none"> • it-ir-raid—Configures the RAID module to use Intel IT/IR RAID. • intel-esrtii—Configures the RAID module to use Intel Embedded Server RAID Technology II. • platform-default—The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.
Onboard SCU Storage Support set onboard-sas-storage-config onboard-sas-ctrl	Whether the onboard software RAID controller is available to the server. This can be one of the following: <ul style="list-style-type: none"> • disabled—The software RAID controller is not available. • enabled—The software RAID controller is available. • platform-default—The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.

Server Management BIOS Settings

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

General Settings

Name	Description
Assert Nmi on Serr set assert-nmi-on-serr-config assertion	<p>Whether the BIOS generates a non-maskable interrupt (NMI) and logs an error when a system error (SERR) occurs. This can be one of the following:</p> <ul style="list-style-type: none"> • disabled—The BIOS does not generate an NMI or log an error when a SERR occurs. • enabled—The BIOS generates an NMI and logs an error when a SERR occurs. You must enable this setting if you want to enable Assert Nmi on Perr. • platform-default—The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.
Assert Nmi on Perr set assert-nmi-on-perr-config assertion	<p>Whether the BIOS generates a non-maskable interrupt (NMI) and logs an error when a processor bus parity error (PERR) occurs. This can be one of the following:</p> <ul style="list-style-type: none"> • disabled—The BIOS does not generate an NMI or log an error when a PERR occurs. • enabled—The BIOS generates an NMI and logs an error when a PERR occurs. You must enable Assert Nmi on Serr to use this setting. • platform-default—The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.
OS Boot Watchdog Timer set os-boot-watchdog-timer-config os-boot-watchdog-timer	<p>Whether the BIOS programs the watchdog timer with a predefined timeout value. If the operating system does not complete booting before the timer expires, the CIMC resets the system and an error is logged. This can be one of the following:</p> <ul style="list-style-type: none"> • disabled—The watchdog timer is not used to track how long the server takes to boot. • enabled—The watchdog timer tracks how long the server takes to boot. If the server does not boot within the predefined length of time, the CIMC resets the system and logs an error. • platform-default—The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor. <p>This feature requires either operating system support or Intel Management software.</p>

Name	Description
<p>OS Boot Watchdog Timer Timeout Policy</p> <p><code>set os-boot-watchdog-timer-policy-config</code> <code>os-boot-watchdog-timer-policy</code></p>	<p>What action the system takes if the watchdog timer expires. This can be one of the following:</p> <ul style="list-style-type: none"> • power-off—The server is powered off if the watchdog timer expires during OS boot. • reset—The server is reset if the watchdog timer expires during OS boot. • platform-default—The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor. <p>This option is only available if you enable the OS Boot Watchdog Timer.</p>
<p>OS Boot Watchdog Timer Timeout</p> <p><code>set</code> <code>os-boot-watchdog-timer-timeout-config</code> <code>os-boot-watchdog-timer-timeout</code></p>	<p>What timeout value the BIOS uses to configure the watchdog timer. This can be one of the following:</p> <ul style="list-style-type: none"> • 5-minutes—The watchdog timer expires 5 minutes after the OS begins to boot. • 10-minutes—The watchdog timer expires 10 minutes after the OS begins to boot. • 15-minutes—The watchdog timer expires 15 minutes after the OS begins to boot. • 20-minutes—The watchdog timer expires 20 minutes after the OS begins to boot. • platform-default—The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor. <p>This option is only available if you enable the OS Boot Watchdog Timer.</p>

Console Redirection Settings

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

Name	Description
BAUD Rate set console-redirect-config baud-rate	<p>What BAUD rate is used for the serial port transmission speed. If you disable Console Redirection, this option is not available. This can be one of the following:</p> <ul style="list-style-type: none"> • 9600—A 9600 BAUD rate is used. • 19200—A 19200 BAUD rate is used. • 38400—A 38400 BAUD rate is used. • 57600—A 57600 BAUD rate is used. • 115200—A 115200 BAUD rate is used. • platform-default—The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor. <p>Note This setting must match the setting on the remote terminal application.</p>
Terminal Type set console-redirect-config terminal-type	<p>What type of character formatting is used for console redirection. This can be one of the following:</p> <ul style="list-style-type: none"> • pc-ansi—The PC-ANSI terminal font is used. • vt100—A supported vt100 video terminal and its character set are used. • vt100-plus—A supported vt100-plus video terminal and its character set are used. • vt-utf8—A video terminal with the UTF-8 character set is used. • platform-default—The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor. <p>Note This setting must match the setting on the remote terminal application.</p>
Legacy OS Redirect set console-redirect-config legacy-os-redirect	<p>Whether redirection from a legacy operating system, such as DOS, is enabled on the serial port. This can be one of the following:</p> <ul style="list-style-type: none"> • disabled—The serial port enabled for console redirection is hidden from the legacy operating system. • enabled—The serial port enabled for console redirection is visible to the legacy operating system. • platform-default—The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.

BIOS Policy

The BIOS policy is a policy that automates the configuration of BIOS settings for a server or group of servers. You can create global BIOS policies available to all servers in the root organization, or you can create BIOS policies in sub-organizations that are only available to that hierarchy.

To use a BIOS policy, do the following:

- 1 Create the BIOS policy in Cisco UCS Manager.
- 2 Assign the BIOS policy to one or more service profiles.
- 3 Associate the service profile with a server.

During service profile association, Cisco UCS Manager modifies the BIOS settings on the server to match the configuration in the BIOS policy. If you do not create and assign a BIOS policy to a service profile, the server uses the default BIOS settings for that server platform.

Default BIOS Settings

Cisco UCS Manager includes a set of default BIOS settings for each type of server supported by Cisco UCS. The default BIOS settings are available only in the root organization and are global. Only one set of default BIOS settings can exist for each server platform supported by Cisco UCS. You can modify the default BIOS settings, but you cannot create an additional set of default BIOS settings.

Each set of default BIOS settings are designed for a particular type of supported server and are applied to all servers of that specific type which do not have a BIOS policy included in their service profiles.

Unless a Cisco UCS implementation has specific needs that are not met by the server-specific settings, we recommend that you use the default BIOS settings that are designed for each type of server in the Cisco UCS domain.

Cisco UCS Manager applies these server platform-specific BIOS settings as follows:

- The service profile associated with a server does not include a BIOS policy.
- The BIOS policy is configured with the platform-default option for a specific setting.

You can modify the default BIOS settings provided by Cisco UCS Manager. However, any changes to the default BIOS settings apply to all servers of that particular type or platform. If you want to modify the BIOS settings for only certain servers, we recommend that you use a BIOS policy.

Creating a BIOS Policy



Note

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

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

Procedure

	Command or Action	Purpose
Step 1	UCS-A# scope org <i>org-name</i>	Enters org mode for the specified organization. To enter the default org mode, type <i>/</i> as the <i>org-name</i> .
Step 2	UCS-A /org # create bios-policy <i>policy-name</i>	Creates a BIOS policy with the specified policy name, and enters org BIOS policy mode.
Step 3	Configure the BIOS settings.	For the CLI commands, descriptions and information about the options for each BIOS setting, see the following topics: <ul style="list-style-type: none"> • Main page: Main BIOS Settings, on page 348 • Processor page: Processor BIOS Settings, on page 350 • Intel Directed IO page: Intel Directed I/O BIOS Settings, on page 356 • RAS Memory page: RAS Memory BIOS Settings, on page 358 • Serial Port page: Serial Port BIOS Settings, on page 360 • USB page: USB BIOS Settings, on page 360 • PCI Configuration page: PCI Configuration BIOS Settings, on page 361 • Boot Options page: Boot Options BIOS Settings, on page 362 • Server Management page: Server Management BIOS Settings, on page 363
Step 4	UCS-A /org/bios-policy # commit-buffer	Commits the transaction to the system configuration.

The following example creates a BIOS policy under the root organization and commits the transaction:

```
UCS-A# scope org /
UCS-A /org # create bios-policy biosPolicy3
UCS-A /org/bios-policy* # set numa-config numa-optimization enabled
```

```
UCS-A /org/bios-policy* # commit-buffer
UCS-A /org/bios-policy #
```

Modifying BIOS Defaults

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

Unless a Cisco UCS implementation has specific needs that are not met by the server-specific settings, we recommend that you use the default BIOS settings that are designed for each type of server in the Cisco UCS domain.

Procedure

	Command or Action	Purpose
Step 1	UCS-A# scope system	Enters system mode.
Step 2	UCS-A /system # scope server-defaults	Enters server defaults mode.
Step 3	UCS-A /system/server-defaults # show platform	(Optional) Displays platform descriptions for all servers.
Step 4	UCS-A /system/server-defaults # scope platform platform-description	Enters server defaults mode for the server specified. For the <i>platform-description</i> argument, enter the server description displayed by the show platform command using the following format: " <i>vendor</i> " <i>model revision</i> . Tip You must enter the vendor exactly as shown in the show platform command, including all punctuation marks.
Step 5	UCS-A /system/server-defaults/platform # scope bios-settings	Enters server defaults BIOS settings mode for the server.
Step 6	Reconfigure the BIOS settings.	For the CLI commands, descriptions and information about the options for each BIOS setting, see the following topics: <ul style="list-style-type: none"> • Main page: Main BIOS Settings, on page 348 • Processor page: Processor BIOS Settings, on page 350 • Intel Directed IO page: Intel Directed I/O BIOS Settings, on page 356 • RAS Memory page: RAS Memory BIOS Settings, on page 358 • Serial Port page: Serial Port BIOS Settings, on page 360 • USB page: USB BIOS Settings, on page 360

	Command or Action	Purpose
		<ul style="list-style-type: none"> • PCI Configuration page: PCI Configuration BIOS Settings, on page 361 • Boot Options page: Boot Options BIOS Settings, on page 362 • Server Management page: Server Management BIOS Settings, on page 363
Step 7	UCS-A /system/server-defaults/platform/bios-settings # commit-buffer	Commits the transaction to the system configuration.

The following example shows how to change the NUMA default BIOS setting for a platform and commit the transaction:

```
UCS-A# scope system
UCS-A /system # scope server-defaults
UCS-A /system/server-defaults # show platform

Platform:
  Product Name  Vendor      Model      Revision
  -----
  Cisco B200-M1 Cisco Systems, Inc.
                    N20-B6620-1
                    0

UCS-A /system/server-defaults # scope platform "Cisco Systems, Inc." N20-B6620-1 0
UCS-A /system/server-defaults/platform # scope bios-settings
UCS-A /system/server-defaults/platform/bios-settings # set numa-config numa-optimization
disabled
UCS-A /system/server-defaults/platform/bios-settings* # commit-buffer
UCS-A /system/server-defaults/platform/bios-settings #
```

Viewing the Actual BIOS Settings for a Server

Follow this procedure to see the actual BIOS settings on a server.

Procedure

	Command or Action	Purpose
Step 1	UCS-A# scope server <i>chassis-id / server-id</i>	Enters chassis server mode for the specified server.
Step 2	UCS-A /chassis/server # scope bios	Enters BIOS mode for the specified server.
Step 3	UCS-A /chassis/server/bios # scope bios-settings	Enters BIOS settings mode for the specified server.

	Command or Action	Purpose
Step 4	UCS-A /chassis/server/bios/bios-settings # show setting	Displays the BIOS setting. Enter show ? to display a list of allowed values for <i>setting</i> .

The following example displays a BIOS setting for blade 3 in chassis 1:

```
UCS-A# scope server 1/3
UCS-A /chassis/server # scope bios
UCS-A /chassis/server/bios # scope bios-settings
UCS-A /chassis/server/bios/bios-settings # show intel-vt-config

Intel Vt Config:
  Vt
  --
  Enabled

UCS-A /chassis/server/bios/bios-settings #
```

Configuring IPMI Access Profiles

IPMI Access Profile

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

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

Configuring an IPMI Access Profile

Before You Begin

Obtain the following:

- Username with appropriate permissions that can be authenticated by the operating system of the server
- Password for the username
- Permissions associated with the username

Procedure

	Command or Action	Purpose
Step 1	UCS-A# scope org org-name	Enters organization mode for the specified organization. To enter the root organization mode, type <i>/</i> as the <i>org-name</i> .

	Command or Action	Purpose
Step 2	UCS-A /org # create ipmi-access-profile <i>profile-name</i>	Creates the specified IPMI access profile and enters organization IPMI access profile mode.
Step 3	UCS-A /org/ipmi-access-profile # create ipmi-user <i>ipmi-user-name</i>	Creates the specified endpoint user and enters organization IPMI access profile endpoint user mode. Note More than one endpoint user can be created within an IPMI access profile, with each endpoint user having its own password and privileges.
Step 4	UCS-A /org/ipmi-access-profile/ipmi-user # set password	Sets the password for the endpoint user. After entering the set password command, you are prompted to enter and confirm the password. For security purposes, the password that you type does not appear in the CLI.
Step 5	UCS-A /org/ipmi-access-profile/ipmi-user # set privilege {admin readonly}	Specifies whether the endpoint user has administrative or read-only privileges.
Step 6	UCS-A /org/ipmi-access-profile/ipmi-user # commit-buffer	Commits the transaction to the system configuration.

The following example creates an IPMI access profile named ReadOnly, creates an endpoint user named bob, sets the password and the privileges for bob, and commits the transaction:

```
UCS-A# scope org /
UCS-A /org # create ipmi-access-profile ReadOnly
UCS-A /org/ipmi-access-profile* # create ipmi-user bob
UCS-A /org/ipmi-access-profile/ipmi-user* # set password
Enter a password:
Confirm the password:
UCS-A /org/ipmi-access-profile/ipmi-user* # set privilege readonly
UCS-A /org/ipmi-access-profile/ipmi-user* # commit-buffer
UCS-A /org/ipmi-access-profile/ipmi-user #
```

What to Do Next

Include the IPMI profile in a service profile and/or template.

Deleting an IPMI Access Profile

Procedure

	Command or Action	Purpose
Step 1	UCS-A# scope org <i>org-name</i>	Enters organization mode for the specified organization. To enter the root organization mode, type / as the <i>org-name</i> .
Step 2	UCS-A /org # delete ipmi-access-profile <i>profile-name</i>	Deletes the specified IPMI access profile.
Step 3	UCS-A /org # commit-buffer	Commits the transaction to the system configuration.

The following example deletes the IPMI access profile named ReadOnly and commits the transaction:

```
UCS-A# scope org /
UCS-A /org # delete ipmi-access-profile ReadOnly
UCS-A /org* # commit-buffer
UCS-A /org #
```

Adding an Endpoint User to an IPMI Access Profile

Procedure

	Command or Action	Purpose
Step 1	UCS-A# scope org <i>org-name</i>	Enters organization mode for the specified organization. To enter the root organization mode, type / as the <i>org-name</i> .
Step 2	UCS-A /org # scope ipmi-access-profile <i>profile-name</i>	Enters organization IPMI access profile mode for the specified IPMI access profile.
Step 3	UCS-A /org/ipmi-access-profile # create ipmi-user <i>ipmi-user-name</i>	Creates the specified endpoint user and enters organization IPMI access profile endpoint user mode. Note More than one endpoint user can be created within an IPMI access profile, with each endpoint user having its own password and privileges.
Step 4	UCS-A /org/ipmi-access-profile/ipmi-user # set password	Sets the password for the endpoint user. After entering the set password command, you are prompted to enter and confirm the password. For security purposes, the password that you type does not appear in the CLI.

	Command or Action	Purpose
Step 5	UCS-A /org/ipmi-access-profile/ipmi-user # set privilege {admin readonly}	Specifies whether the endpoint user has administrative or read-only privileges.
Step 6	UCS-A /org/ipmi-access-profile/ipmi-user # commit-buffer	Commits the transaction to the system configuration.

The following example adds an endpoint user named alice to the IPMI access profile named ReadOnly and commits the transaction:

```
UCS-A# scope org /
UCS-A /org* # scope ipmi-access-profile ReadOnly
UCS-A /org/ipmi-access-profile* # create ipmi-user alice
UCS-A /org/ipmi-access-profile/ipmi-user* # set password
Enter a password:
Confirm the password:
UCS-A /org/ipmi-access-profile/ipmi-user* # set privilege readonly
UCS-A /org/ipmi-access-profile/ipmi-user* # commit-buffer
UCS-A /org/ipmi-access-profile/ipmi-user #
```

Deleting an Endpoint User from an IPMI Access Profile

Procedure

	Command or Action	Purpose
Step 1	UCS-A# scope org <i>org-name</i>	Enters organization mode for the specified organization. To enter the root organization mode, type / as the <i>org-name</i> .
Step 2	UCS-A /org # scope ipmi-access-profile <i>profile-name</i>	Enters organization IPMI access profile mode for the specified IPMI access profile.
Step 3	UCS-A /org/ipmi-access-profile # delete ipmi-user <i>epuser-name</i>	Deletes the specified endpoint user from the IPMI access profile.
Step 4	UCS-A /org/ipmi-access-profile # commit-buffer	Commits the transaction to the system configuration.

The following example deletes the endpoint user named alice from the IPMI access profile named ReadOnly and commits the transaction:

```
UCS-A# scope org /
UCS-A /org # scope ipmi-access-profile ReadOnly
UCS-A /org/ipmi-access-profile # delete ipmi-user alice
UCS-A /org/ipmi-access-profile* # commit-buffer
UCS-A /org/ipmi-access-profile #
```

Configuring Local Disk Configuration Policies

Local Disk Configuration Policy

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

The local disk modes include the following:

- **No Local Storage**—For a diskless server or a SAN only configuration. If you select this option, you cannot associate any service profile which uses this policy with a server that has a local disk.
- **RAID 0 Striped**—Data is striped across all disks in the array, providing fast throughput. There is no data redundancy, and all data is lost if any disk fails.
- **RAID 1 Mirrored**—Data is written to two disks, providing complete data redundancy if one disk fails. The maximum array size is equal to the available space on the smaller of the two drives.
- **Any Configuration**—For a server configuration that carries forward the local disk configuration without any changes.
- **No RAID**—For a server configuration that removes the RAID and leaves the disk MBR and payload unaltered.
- **RAID 5 Striped Parity**—Data is striped across all disks in the array. Part of the capacity of each disk stores parity information that can be used to reconstruct data if a disk fails. RAID 5 provides good data throughput for applications with high read request rates.
- **RAID 6 Striped Dual Parity**—Data is striped across all disks in the array and two parity disks are used to provide protection against the failure of up to two physical disks. In each row of data blocks, two sets of parity data are stored.
- **RAID10 Mirrored and Striped**— RAID 10 uses mirrored pairs of disks to provide complete data redundancy and high throughput rates.

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

Guidelines for all Local Disk Configuration Policies

Before you create a local disk configuration policy, consider the following guidelines:

No Mixed HDDs and SSDs

Do not include HDDs and SSDs in a single server or RAID configuration.

Do Not Assign a Service Profile with the Default Local Disk Configuration Policy from a B200 M1 or M2 to a B200 M3

Due to the differences in the RAID/JBOD support provided by the storage controllers of B200 M1 and M2 servers and those of the B200 M3 server, you cannot assign or re-assign a service profile that includes the default local disk configuration policy from a B200M1 or M2 server to a B200 M3 server. The default local disk configuration policy includes those with Any Configuration or JBOD configuration.

Impact of Upgrade from a Release Prior to Release 1.3(1i)

An upgrade from an earlier Cisco UCS firmware release to release 1.3(1i) or higher has the following impact on the Protect Configuration property of the local disk configuration policy the first time servers are associated with service profiles after the upgrade:

Unassociated Servers

After you upgrade the Cisco UCS domain, the initial server association proceeds without configuration errors whether or not the local disk configuration policy matches the server hardware. Even if you enable the Protect Configuration property, Cisco UCS does not protect the user data on the server if there are configuration mismatches between the local disk configuration policy on the previous service profile and the policy in the new service profile.



Note If you enable the Protect Configuration property and the local disk configuration policy encounters mismatches between the previous service profile and the new service profile, all subsequent service profile associations with the server are blocked.

Associated Servers

Any servers that are already associated with service profiles do not reboot after the upgrade. Cisco UCS Manager does not report any configuration errors if there is a mismatch between the local disk configuration policy and the server hardware.

When a service profile is disassociated from a server and a new service profile associated, the setting for the Protect Configuration property in the new service profile takes precedence and overwrites the setting in the previous service profile.

Guidelines for Local Disk Configuration Policies Configured for RAID

No Mixed HDDs and SSDs

Do not include HDDs and SSDs in a single RAID configuration.

Do Not Use the Any Configuration Mode on Servers with MegaRAID Storage Controllers

If a blade server or rack-mount server in a Cisco UCS domain includes a MegaRAID storage controller, do not configure the local disk configuration policy in the service profile for that server with the **Any Configuration** mode. If you use this mode for servers with a MegaRAID storage controller, the installer for the operating system cannot detect any local storage on the server.

If you want to install an operating system on local storage on a server with a MegaRAID storage controller, you must configure the local disk configuration policy with a mode that creates a RAID LUN (RAID volume) on the server.

Server May Not Boot After RAID1 Cluster Migration if Any Configuration Mode Specified in Service Profile

After RAID1 clusters are migrated, you need to associate a service profile with the server. If the local disk configuration policy in the service profile is configured with **Any Configuration** mode rather than **RAID1**, the RAID LUN remains in "inactive" state during and after association. As a result, the server cannot boot.

To avoid this issue, ensure that the service profile you associate with the server contains the identical local disk configuration policy as the original service profile before the migration and does not include the **Any Configuration** mode.

Configure RAID Settings in Local Disk Configuration Policy for Servers with MegaRAID Storage Controllers

If a blade server or integrated rack-mount server has a MegaRAID controller, you must configure RAID settings for the drives in the Local Disk Configuration policy included in the service profile for that server.

If you do not configure your RAID LUNs before installing the OS, disk discovery failures might occur during the installation and you might see error messages such as “No Device Found.”

Do Not Use JBOD Mode on Servers with MegaRAID Storage Controllers

Do not configure or use JBOD mode or JBOD operations on any blade server or integrated rack-mount server with a MegaRAID storage controllers. JBOD mode and operations are not intended for nor are they fully functional on these servers.

Maximum of One RAID Volume and One RAID Controller in Integrated Rack-Mount Servers

A rack-mount server that has been integrated with Cisco UCS Manager can have a maximum of one RAID volume irrespective of how many hard drives are present on the server.

All the local hard drives in an integrated rack-mount server must be connected to only one RAID Controller. Integration with Cisco UCS Manager does not support the connection of local hard drives to multiple RAID Controllers in a single rack-mount server. We therefore recommend that you request a single RAID Controller configuration when you order rack-mount servers to be integrated with Cisco UCS Manager.

In addition, do not use third party tools to create multiple RAID LUNs on rack-mount servers. Cisco UCS Manager does not support that configuration.

Maximum of One RAID Volume and One RAID Controller in Blade Servers

A blade server can have a maximum of one RAID volume irrespective of how many drives are present in the server. All the local hard drives must be connected to only one RAID controller. For example, a B200 M3 server has an LSI controller and an Intel Patsburg controller, but only the LSI controller can be used as a RAID controller.

In addition, do not use third party tools to create multiple RAID LUNs on blade servers. Cisco UCS Manager does not support that configuration.

Number of Disks Selected in Mirrored RAID Should Not Exceed Two

If the number of disks selected in the Mirrored RAID exceed two, RAID 1 is created as a RAID 10 LUN. This issue can occur with the Cisco UCS B440 M1 and B440 M2 servers.

B420 M3 Server Does Not Support All Configuration Modes

The B420 M3 server does not support the following configuration modes in a local disk configuration policy:

- No RAID
- RAID 6 Striped Dual Parity

In addition, the B420 M3 does not support JBOD modes or operations.

Creating a Local Disk Configuration Policy

Procedure

	Command or Action	Purpose
Step 1	UCS-A# scope org <i>org-name</i>	Enters organization mode for the specified organization. To enter the root organization mode, type / as the <i>org-name</i> .
Step 2	UCS-A /org # create local-disk-config-policy <i>policy-name</i>	Creates a local disk configuration policy and enters local disk configuration policy mode.
Step 3	UCS-A /org/local-disk-config-policy # set descr <i>description</i>	(Optional) Provides a description for the local disk configuration policy.
Step 4	UCS-A /org/local-disk-config-policy # set mode { any-configuration no-local-storage no-raid raid-0-striped raid-1-mirrored raid-5-striped-parity raid-6-striped-dual-parity raid-10-mirrored-and-striped }	Specifies the mode for the local disk configuration policy.
Step 5	UCS-A /org/local-disk-config-policy # set protect { yes no }	Specifies whether the server retains the configuration in the local disk configuration policy even if the server is disassociated from the service profile. Caution Protect Configuration becomes non-functional if one or more disks in the server are defective or faulty. When a service profile is disassociated from a server and a new service profile associated, the setting for the Protect Configuration property in the new service profile takes precedence and overwrites the setting in the previous service profile. Note If you disassociate the server from a service profile with this option enabled and then associate it with a new service profile that includes a local disk configuration policy with different properties, the server returns a configuration mismatch error and the association fails.
Step 6	UCS-A /org/local-disk-config-policy # commit-buffer	Commits the transaction to the system configuration.

The following example configures a local disk configuration policy and commits the transaction:

```
UCS-A# scope org /
UCS-A /org # create local-disk-config-policy DiskPolicy7
UCS-A /org/local-disk-config-policy* # set mode raid-1-mirrored
```

```
UCS-A /org/local-disk-config-policy* # set protect yes
UCS-A /org/local-disk-config-policy* # commit-buffer
UCS-A /org/local-disk-config-policy #
```

Viewing a Local Disk Configuration Policy

Procedure

	Command or Action	Purpose
Step 1	UCS-A# scope org <i>org-name</i>	Enters organization mode for the specified organization. To enter the root organization mode, type / as the <i>org-name</i> .
Step 2	UCS-A /org # show local-disk-config-policy <i>policy-name</i>	Displays the local disk policy. If you have not configured a local disk policy, the local disk configuration (created by the create local-disk-config command) displays. Displays the local disk definition (set by the create local-disk-config command). If the serial over LAN definition is not set, and if a policy is set (using the set local-disk-config-policy command), then the policy will be displayed.

The following example shows how to display local disk policy information for a local disk configuration policy called DiskPolicy7:

```
UCS-A# scope org /
UCS-A /org # show local-disk-config-policy DiskPolicy7

Local Disk Config Policy:
Name: DiskPolicy7
Mode: Raid 1 Mirrored
Description:
Protect Configuration: Yes
```

Deleting a Local Disk Configuration Policy

Procedure

	Command or Action	Purpose
Step 1	UCS-A# scope org <i>org-name</i>	Enters organization mode for the specified organization. To enter the root organization mode, type / as the <i>org-name</i> .
Step 2	UCS-A /org # delete local-disk-config-policy <i>policy-name</i>	Deletes the specified local disk configuration policy.
Step 3	UCS-A /org # commit-buffer	Commits the transaction to the system configuration.

The following example deletes the local disk configuration policy named DiskPolicy7 and commits the transaction:

```
UCS-A# scope org /
UCS-A /org # delete local-disk-config-policy DiskPolicy7
UCS-A /org* # commit-buffer
UCS-A /org #
```

Configuring Scrub Policies

Scrub Policy

This policy determines what happens to local data and to the BIOS settings on a server during the discovery process and when the server is disassociated from a service profile. Depending upon how you configure a scrub policy, the following can occur at those times:

Disk Scrub

One of the following occurs to the data on any local drives on disassociation:

- If enabled, destroys all data on any local drives
- If disabled, preserves all data on any local drives, including local storage configuration

BIOS Settings Scrub

One of the following occurs to the BIOS settings when a service profile containing the scrub policy is disassociated from a server:

- If enabled, erases all BIOS settings for the server and resets them to the BIOS defaults for that server type and vendor
- If disabled, preserves the existing BIOS settings on the server

Creating a Scrub Policy

Procedure

	Command or Action	Purpose
Step 1	UCS-A# scope org <i>org-name</i>	Enters organization mode for the specified organization. To enter the root organization mode, type <i>/</i> as the <i>org-name</i> .
Step 2	UCS-A /org # create scrub-policy <i>policy-name</i>	Creates a scrub policy with the specified policy name, and enters organization scrub policy mode.
Step 3	UCS-A /org/scrub-policy # set descr <i>description</i>	(Optional) Provides a description for the scrub policy.

	Command or Action	Purpose
		Note If your description includes spaces, special characters, or punctuation, you must begin and end your description with quotation marks. The quotation marks will not appear in the description field of any show command output.
Step 4	UCS-A /org/scrub-policy # set disk-scrub {no yes}	Disables or enables disk scrubbing on servers using this scrub policy as follows: <ul style="list-style-type: none"> • If enabled, destroys all data on any local drives • If disabled, preserves all data on any local drives, including local storage configuration
Step 5	UCS-A /org/scrub-policy # set bios-settings-scrub {no yes}	Disables or enables BIOS settings scrubbing on servers using this scrub policy as follows: <ul style="list-style-type: none"> • If enabled, erases all BIOS settings for the server and resets them to the BIOS defaults for that server type and vendor • If disabled, preserves the existing BIOS settings on the server
Step 6	UCS-A /org/scrub-policy # commit-buffer	Commits the transaction to the system configuration.

The following example creates a scrub policy named ScrubPolicy2, enables disk scrubbing on servers using the scrub policy, and commits the transaction:

```
UCS-A# scope org /
UCS-A /org # create scrub-policy ScrubPolicy2
UCS-A /org/scrub-policy* # set descr "Scrub disk but not BIOS."
UCS-A /org/scrub-policy* # set disk-scrub yes
UCS-A /org/scrub-policy* # set bios-settings-scrub no
UCS-A /org/scrub-policy* # commit-buffer
UCS-A /org/scrub-policy #
```

Deleting a Scrub Policy

Procedure

	Command or Action	Purpose
Step 1	UCS-A# scope org org-name	Enters organization mode for the specified organization. To enter the root organization mode, type / as the <i>org-name</i> .
Step 2	UCS-A /org # delete scrub-policy policy-name	Deletes the specified scrub policy.

	Command or Action	Purpose
Step 3	UCS-A /org # commit-buffer	Commits the transaction to the system configuration.

The following example deletes the scrub policy named ScrubPolicy2 and commits the transaction:

```
UCS-A# scope org /
UCS-A /org # delete scrub-policy ScrubPolicy2
UCS-A /org* # commit-buffer
UCS-A /org #
```

Configuring Serial over LAN Policies

Serial over LAN Policy

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

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

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

Configuring a Serial over LAN Policy

Procedure

	Command or Action	Purpose
Step 1	UCS-A# scope org <i>org-name</i>	Enters organization mode for the specified organization. To enter the root organization mode, type <i>/</i> as the <i>org-name</i> .
Step 2	UCS-A /org # create sol-policy <i>policy-name</i>	Creates a serial over LAN policy and enters organization serial over LAN policy mode.
Step 3	UCS-A /org/sol-policy # set descr <i>description</i>	(Optional) Provides a description for the policy. Note If your description includes spaces, special characters, or punctuation, you must begin and end your description with quotation marks. The quotation marks will not appear in the description field of any show command output.
Step 4	UCS-A /org/sol-policy # set speed <i>{115200 19200 38400 57600 9600}</i>	Specifies the serial baud rate.

	Command or Action	Purpose
Step 5	UCS-A /org/sol-policy # { disable enable }	Disables or enables the serial over LAN policy. By default, the serial over LAN policy is disabled; you must enable it before it can be applied.
Step 6	UCS-A /org/sol-policy # commit-buffer	Commits the transaction to the system configuration.

The following example creates a serial over LAN policy named Sol9600, provides a description for the policy, sets the speed to 9,600 baud, enables the policy, and commits the transaction:

```
UCS-A# scope org /
UCS-A /org* # create sol-policy Sol9600
UCS-A /org/sol-policy* # set descr "Sets serial over LAN policy to 9600 baud."
UCS-A /org/sol-policy* # set speed 9600
UCS-A /org/sol-policy* # enable
UCS-A /org/sol-policy* # commit-buffer
UCS-A /org/sol-policy #
```

Viewing a Serial over LAN Policy

Procedure

	Command or Action	Purpose
Step 1	UCS-A# scope org org-name	Enters organization mode for the specified organization. To enter the root organization mode, type / as the <i>org-name</i> .
Step 2	UCS-A /org # show sol-policy policy-name	Displays the serial over LAN definition (set by the create sol-config command). If the serial over LAN definition is not set, and if a policy is set (using the set sol-policy command), then the policy will be displayed.

The following example shows how to display serial over LAN information for a serial over LAN policy called Sol9600:

```
UCS-A# scope org /
UCS-A /org # show sol-policy Sol9600

SOL Policy:
Full Name: Sol9600
SOL State: Enable
Speed: 9600
Description:
```

Deleting a Serial over LAN Policy

Procedure

	Command or Action	Purpose
Step 1	UCS-A# scope org <i>org-name</i>	Enters organization mode for the specified organization. To enter the root organization mode, type / as the <i>org-name</i> .
Step 2	UCS-A /org # delete sol-policy <i>policy-name</i>	Deletes the specified serial over LAN policy.
Step 3	UCS-A /org # commit-buffer	Commits the transaction to the system configuration.

The following example deletes the serial over LAN policy named Sol9600 and commits the transaction:

```
UCS-A# scope org /
UCS-A /org* # delete sol-policy Sol9600
UCS-A /org* # commit-buffer
UCS-A /org #
```

Configuring Server Autoconfiguration Policies

Server Autoconfiguration Policy

Cisco UCS Manager uses this policy to determine how to configure a new server. If you create a server autoconfiguration policy, the following occurs when a new server starts:

- 1 The qualification in the server autoconfiguration policy is executed against the server.
- 2 If the server meets the required qualifications, the server is associated with a service profile created from the service profile template configured in the server autoconfiguration policy. The name of that service profile is based on the name given to the server by Cisco UCS Manager.
- 3 The service profile is assigned to the organization configured in the server autoconfiguration policy.

Configuring a Server Autoconfiguration Policy

Procedure

	Command or Action	Purpose
Step 1	UCS-A# scope org <i>org-name</i>	Enters organization mode for the specified organization. To enter the root organization mode, type / as the <i>org-name</i> .

	Command or Action	Purpose
Step 2	UCS-A /org # create server-autoconfig-policy <i>policy-name</i>	Creates a server autoconfiguration policy with the specified policy name, and enters organization server autoconfiguration policy mode.
Step 3	UCS-A /org/server-autoconfig-policy # set descr <i>description</i>	(Optional) Provides a description for the policy. Note If your description includes spaces, special characters, or punctuation, you must begin and end your description with quotation marks. The quotation marks will not appear in the description field of any show command output.
Step 4	UCS-A /org/server-autoconfig-policy # set destination org <i>org-name</i>	(Optional) Specifies the organization for which the server is to be used.
Step 5	UCS-A /org/server-autoconfig-policy # set qualifier <i>server-qual-name</i>	(Optional) Specifies server pool policy qualification to use for qualifying the server.
Step 6	UCS-A /org/server-autoconfig-policy # set template <i>profile-name</i>	(Optional) Specifies a service profile template to use for creating a service profile instance for the server.
Step 7	UCS-A /org/server-autoconfig-policy # commit-buffer	Commits the transaction to the system configuration.

The following example creates a server autoconfiguration policy named AutoConfigFinance, provides a description for the policy, specifies finance as the destination organization, ServPoolQual22 as the server pool policy qualification, and ServTemp2 as the service profile template, and commits the transaction:

```
UCS-A# scope org /
UCS-A /org* # create server-autoconfig-policy AutoConfigFinance
UCS-A /org/server-autoconfig-policy* # set descr "Server Autoconfiguration Policy for Finance"
UCS-A /org/server-autoconfig-policy* # set destination org finance
UCS-A /org/server-autoconfig-policy* # set qualifier ServPoolQual22
UCS-A /org/server-autoconfig-policy* # set template ServTemp2
UCS-A /org/server-autoconfig-policy* # commit-buffer
UCS-A /org/server-autoconfig-policy #
```

Deleting a Server Autoconfiguration Policy

Procedure

	Command or Action	Purpose
Step 1	UCS-A# scope org <i>org-name</i>	Enters organization mode for the specified organization. To enter the root organization mode, type <i>/</i> as the <i>org-name</i> .
Step 2	UCS-A /org # delete server-autoconfig-policy <i>policy-name</i>	Deletes the specified server autoconfiguration policy.
Step 3	UCS-A /org # commit-buffer	Commits the transaction to the system configuration.

The following example deletes the server autoconfiguration policy named AutoConfigFinance and commits the transaction:

```
UCS-A# scope org /
UCS-A /org* # delete server-autoconfig-policy AutoConfigFinance
UCS-A /org* # commit-buffer
UCS-A /org #
```

Configuring Server Discovery Policies

Server Discovery Policy

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

If you create a server discovery policy, the following occurs when a new server starts:

- 1 The qualification in the server discovery policy is executed against the server.
- 2 If the server meets the required qualifications, Cisco UCS Manager applies the following to the server:
 - Depending upon the option selected for the action, either discovers the new server immediately or waits for a user to acknowledge the new server
 - Applies the scrub policy to the server

Configuring a Server Discovery Policy

Before You Begin

If you plan to associate this policy with a server pool, create server pool policy qualifications.

Procedure

	Command or Action	Purpose
Step 1	UCS-A# scope org /	Enters the root organization mode. Note Chassis discovery policies can only be accessed from the root organization.
Step 2	UCS-A /org # create server-disc-policy <i>policy-name</i>	Creates a server discovery policy with the specified policy name, and enters org server discovery policy mode.
Step 3	UCS-A /org/server-disc-policy # set action { diag immediate user-acknowledged }	Specifies when the system will attempt to discover new servers.
Step 4	UCS-A /org/chassis-disc-policy # set descr <i>description</i>	(Optional) Provides a description for the server discovery policy. Note If your description includes spaces, special characters, or punctuation, you must begin and end your description with quotation marks. The quotation marks will not appear in the description field of any show command output.
Step 5	UCS-A /org/server-disc-policy # set qualifier <i>qualifier</i>	(Optional) Uses the specified server pool policy qualifications to associates this policy with a server pool.
Step 6	UCS-A /org/server-disc-policy # set scrub-policy	Specifies the scrub policy to be used by this policy. The scrub policy defines whether the disk drive on a server should be scrubbed clean upon discovery.
Step 7	UCS-A /org/server-disc-policy # commit-buffer	Commits the transaction to the system configuration.

The following example creates a server discovery policy named ServDiscPolExample, sets it to immediately discover new servers, provides a description for the policy, specifies the server pool policy qualifications and scrub policy, and commits the transaction:

```
UCS-A# scope org /
UCS-A /org # create server-disc-policy ServDiscPolExample
UCS-A /org/server-disc-policy* # set action immediate
UCS-A /org/server-disc-policy* # set descr "This is an example server discovery policy."
UCS-A /org/server-disc-policy* # set qualifier ExampleQual
UCS-A /org/server-disc-policy* # set scrub-policy NoScrub
UCS-A /org/server-disc-policy # commit-buffer
```

What to Do Next

Include the server discovery policy in a service profile and/or template.

Deleting a Server Discovery Policy

Procedure

	Command or Action	Purpose
Step 1	UCS-A# scope org <i>org-name</i>	Enters organization mode for the specified organization. To enter the root organization mode, type / as the <i>org-name</i> .
Step 2	UCS-A/org # Delete server-disc-policy <i>policy-name</i>	Deletes the specified server discovery policy.
Step 3	UCS-A /org/server-disc-policy # commit-buffer	Commits the transaction to the system configuration.

The following example deletes the server discovery policy named ServDiscPolExample and commits the transaction:

```
UCS-A# scope org /
UCS-A /org # delete server-disc-policy ServDiscPolExample
UCS-A /org* # commit-buffer
UCS-A /org #
```

Configuring Server Inheritance Policies

Server Inheritance Policy

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

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

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

Configuring a Server Inheritance Policy

A blade server or rack-mount server with a VIC adapter, such as the Cisco UCS M81KR Virtual Interface Card, does not have server identity values burned into the server hardware at manufacture. As a result, the identity of the adapter must be derived from default pools. If the default pools do not include sufficient entries for one to be assigned to the server, service profile association fails with a configuration error.

Procedure

	Command or Action	Purpose
Step 1	UCS-A# scope org <i>org-name</i>	Enters organization mode for the specified organization. To enter the root organization mode, type / as the <i>org-name</i> .
Step 2	UCS-A /org # create server-inherit-policy <i>policy-name</i>	Creates a server inheritance policy with the specified policy name, and enters organization server inheritance policy mode.
Step 3	UCS-A /org/server-inherit-policy # set descr <i>description</i>	(Optional) Provides a description for the policy. Note If your description includes spaces, special characters, or punctuation, you must begin and end your description with quotation marks. The quotation marks will not appear in the description field of any show command output.
Step 4	UCS-A /org/server-inherit-policy # set destination org <i>org-name</i>	(Optional) Specifies the organization for which the server is to be used.
Step 5	UCS-A /org/server-inherit-policy # set qualifier <i>server-qual-name</i>	(Optional) Specifies server pool policy qualification to use for qualifying the server.
Step 6	UCS-A /org/server-inherit-policy # commit-buffer	Commits the transaction to the system configuration.

The following example creates a server inheritance policy named InheritEngineering, provides a description for the policy, specifies engineering as the destination organization and ServPoolQual22 as the server pool policy qualification, and commits the transaction:

```
UCS-A# scope org /
UCS-A /org* # create server-inherit-policy InheritEngineering
UCS-A /org/server-inherit-policy* # set descr "Server Inheritance Policy for Engineering"
UCS-A /org/server-inherit-policy* # set destination org engineering
UCS-A /org/server-inherit-policy* # set qualifier ServPoolQual22
UCS-A /org/server-inherit-policy* # commit-buffer
UCS-A /org/server-inherit-policy #
```


Deleting a Server Inheritance Policy

Procedure

	Command or Action	Purpose
Step 1	UCS-A# scope org <i>org-name</i>	Enters organization mode for the specified organization. To enter the root organization mode, type / as the <i>org-name</i> .
Step 2	UCS-A /org # delete server-inherit-policy <i>policy-name</i>	Deletes the specified server inheritance policy.
Step 3	UCS-A /org # commit-buffer	Commits the transaction to the system configuration.

The following example deletes the server inheritance policy named InheritEngineering and commits the transaction:

```
UCS-A# scope org /
UCS-A /org* # delete server-inherit-policy InheritEngineering
UCS-A /org* # commit-buffer
UCS-A /org #
```

Configuring Server Pool Policies

Server Pool Policy

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

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

Configuring a Server Pool Policy

Procedure

	Command or Action	Purpose
Step 1	UCS-A# scope org <i>org-name</i>	Enters organization mode for the specified organization. To enter the root organization mode, type / as the <i>org-name</i> .
Step 2	UCS-A /org # create pooling-policy <i>policy-name</i>	Creates a server pool policy with the specified name, and enters organization pooling policy mode.
Step 3	UCS-A /org/pooling-policy # set descr <i>description</i>	(Optional) Provides a description for the server pool policy.

	Command or Action	Purpose
		Note If your description includes spaces, special characters, or punctuation, you must begin and end your description with quotation marks. The quotation marks will not appear in the description field of any show command output.
Step 4	UCS-A /org/pooling-policy # set pool <i>pool-distinguished-name</i>	Specifies the server pool to use with the server pool policy. You must specify the full distinguished name for the pool.
Step 5	UCS-A /org/pooling-policy # set qualifier <i>qualifier-name</i>	Specifies the server pool qualifier to use with the server pool policy.
Step 6	UCS-A /org/pooling-policy # commit-buffer	Commits the transaction to the system configuration.

The following example creates a server pool policy named ServerPoolPolicy4 and commits the transaction:

```
UCS-A# scope org /
UCS-A /org # create pooling-policy ServerPoolPolicy4
UCS-A /org/pooling-policy* # set pool org-root/compute-pool-pool3
UCS-A /org/pooling-policy* # set qualifier ServPoolQual8
UCS-A /org/pooling-policy* # commit-buffer
UCS-A /org/pooling-policy #
```

Deleting a Server Pool Policy

Procedure

	Command or Action	Purpose
Step 1	UCS-A# scope org <i>org-name</i>	Enters organization mode for the specified organization. To enter the root organization mode, type / as the <i>org-name</i> .
Step 2	UCS-A /org # delete pooling-policy <i>policy-name</i>	Deletes the specified server pool policy.
Step 3	UCS-A /org # commit-buffer	Commits the transaction to the system configuration.

The following example deletes the server pool policy named ServerPoolPolicy4 and commits the transaction:

```
UCS-A# scope org /
UCS-A /org # delete pooling-policy ServerPoolPolicy4
UCS-A /org/pooling-policy* # commit-buffer
UCS-A /org/pooling-policy #
```

Configuring Server Pool Policy Qualifications

Server Pool Policy Qualifications

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

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

You can use the server pool policy qualifications to qualify servers according to the following criteria:

- Adapter type
- Chassis location
- Memory type and configuration
- Power group
- CPU cores, type, and configuration
- Storage configuration and capacity
- Server model

Depending upon the implementation, you may configure several policies with server pool policy qualifications including the following:

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

Creating a Server Pool Policy Qualification

Procedure

	Command or Action	Purpose
Step 1	UCS-A# scope org <i>org-name</i>	Enters organization mode for the specified organization. To enter the root organization mode, type / as the <i>org-name</i> .

	Command or Action	Purpose
Step 2	UCS-A /org # create server-qual <i>server-qual-name</i>	Creates a server pool qualification with the specified name, and enters organization server qualification mode.
Step 3	UCS-A /org/server-qual # commit-buffer	Commits the transaction to the system configuration.

The following example creates a server pool qualification named ServPoolQual22 and commits the transaction:

```
UCS-A# scope org /
UCS-A /org* # create server-qual ServPoolQual22
UCS-A /org/server-qual* # commit-buffer
UCS-A /org/server-qual #
```

What to Do Next

Configure one or more of the following server component qualifications:

- Adapter qualification
- Chassis qualification
- Memory qualification
- Power group qualification
- Processor qualification
- Storage qualification

Deleting a Server Pool Policy Qualification

Procedure

	Command or Action	Purpose
Step 1	UCS-A# scope org <i>org-name</i>	Enters organization mode for the specified organization. To enter the root organization mode, type / as the <i>org-name</i> .
Step 2	UCS-A /org # delete server-qual <i>server-qual-name</i>	Deletes the specified server pool qualification.
Step 3	UCS-A /org/server-qual # commit-buffer	Commits the transaction to the system configuration.

The following example deletes the server pool qualification named ServPoolQual22 and commits the transaction:

```
UCS-A# scope org /
UCS-A /org* # delete server-qual ServPoolQual22
UCS-A /org* # commit-buffer
UCS-A /org #
```

Creating an Adapter Qualification

Before You Begin

Create a server pool policy qualification.

Procedure

	Command or Action	Purpose
Step 1	UCS-A# scope org <i>org-name</i>	Enters organization mode for the specified organization. To enter the root organization mode, type / as the <i>org-name</i> .
Step 2	UCS-A /org # scope server-qual <i>server-qual-name</i>	Enters organization server qualification mode for the specified server pool policy qualification.
Step 3	UCS-A /org/server-qual # create adapter	Creates an adapter qualification and enters organization server qualification adapter mode.
Step 4	UCS-A /org/server-qual/adapter # create cap-qual <i>adapter-type</i>	Creates an adapter capacity qualification for the specified adapter type and enters organization server qualification adapter capacity qualification mode. The <i>adapter-type</i> argument can be any of the following values: <ul style="list-style-type: none"> • fcoe —Fibre Channel over Ethernet • non-virtualized-eth-if —Non-virtualized Ethernet interface • non-virtualized-fc-if —Non-virtualized Fibre Channel interface • path-encap-consolidated —Path encapsulation consolidated • path-encap-virtual —Path encapsulation virtual • protected-eth-if —Protected Ethernet interface • protected-fc-if —Protected Fibre Channel interface • protected-fcoe —Protected Fibre Channel over Ethernet • virtualized-eth-if —Virtualized Ethernet interface • virtualized-fc-if —Virtualized Fibre Channel interface • virtualized-scsi-if —Virtualized SCSI interface
Step 5	UCS-A /org/server-qual/adapter/cap-qual # set maximum { <i>max-cap</i> unspecified }	Specifies the maximum capacity for the selected adapter type.
Step 6	UCS-A /org/server-qual/adapter/cap-qual # commit-buffer	Commits the transaction to the system configuration.

The following example creates and configures an adapter qualification for a non-virtualized Ethernet interface and commits the transaction:

```
UCS-A# scope org /
UCS-A /org # scope server-qual ServPoolQual22
UCS-A /org/server-qual # create adapter
UCS-A /org/server-qual/adapter* # create cap-qual non-virtualized-eth-if
UCS-A /org/server-qual/adapter/cap-qual* # set maximum 2500000000
UCS-A /org/server-qual/adapter/cap-qual* # commit-buffer
UCS-A /org/server-qual/adapter/cap-qual #
```

Deleting an Adapter Qualification

Procedure

	Command or Action	Purpose
Step 1	UCS-A# scope org <i>org-name</i>	Enters organization mode for the specified organization. To enter the root organization mode, type / as the <i>org-name</i> .
Step 2	UCS-A /org # scope server-qual <i>server-qual-name</i>	Enters organization server qualification mode for the specified server pool policy qualification.
Step 3	UCS-A /org/server-qual # delete adapter	Deletes the adapter qualification from the server pool policy qualification.
Step 4	UCS-A /org/server-qual # commit-buffer	Commits the transaction to the system configuration.

The following example deletes the adapter qualification from the server pool policy qualification named ServPoolQual22 and commits the transaction:

```
UCS-A# scope org /
UCS-A /org # scope server-qual ServPoolQual22
UCS-A /org/server-qual # delete adapter
UCS-A /org/server-qual* # commit-buffer
UCS-A /org/server-qual #
```

Configuring a Chassis Qualification

Before You Begin

Create a server pool policy qualification.

Procedure

	Command or Action	Purpose
Step 1	UCS-A# scope org <i>org-name</i>	Enters organization mode for the specified organization. To enter the root organization mode, type <i>/</i> as the <i>org-name</i> .
Step 2	UCS-A /org # scope server-qual <i>server-qual-name</i>	Enters organization server qualification mode for the specified server pool policy qualification.
Step 3	UCS-A /org/server-qual # create chassis <i>min-chassis-num max-chassis-num</i>	Creates a chassis qualification for the specified chassis range and enters organization server qualification chassis mode.
Step 4	UCS-A /org/server-qual/chassis # create slot <i>min-slot-num max-slot-num</i>	Creates a chassis slot qualification for the specified slot range and enters organization server qualification chassis slot mode.
Step 5	UCS-A /org/server-qual/chassis/slot # commit-buffer	Commits the transaction to the system configuration.

The following example configures a chassis qualification for slots 1 to 4 on chassis 1 and 2 and commits the transaction:

```
UCS-A# scope org /
UCS-A /org* # scope server-qual ServPoolQual22
UCS-A /org/server-qual* # create chassis 1 2
UCS-A /org/server-qual/chassis* # create slot 1 4
UCS-A /org/server-qual/chassis/slot* # commit-buffer
UCS-A /org/server-qual/chassis/slot #
```

Deleting a Chassis Qualification**Procedure**

	Command or Action	Purpose
Step 1	UCS-A# scope org <i>org-name</i>	Enters organization mode for the specified organization. To enter the root organization mode, type <i>/</i> as the <i>org-name</i> .
Step 2	UCS-A /org # scope server-qual <i>server-qual-name</i>	Enters organization server qualification mode for the specified server pool policy qualification.
Step 3	UCS-A /org/server-qual # delete chassis <i>min-chassis-num max-chassis-num</i>	Deletes the chassis qualification for the specified chassis range.
Step 4	UCS-A /org/server-qual # commit-buffer	Commits the transaction to the system configuration.

The following example deletes the chassis qualification for chassis 1 and 2 and commits the transaction:

```
UCS-A# scope org /
UCS-A /org # scope server-qual ServPoolQual22
UCS-A /org/server-qual # delete chassis 1 2
UCS-A /org/server-qual* # commit-buffer
UCS-A /org/server-qual #
```

Creating a CPU Qualification

Before You Begin

Create a server pool policy qualification.

Procedure

	Command or Action	Purpose
Step 1	UCS-A# scope org <i>org-name</i>	Enters organization mode for the specified organization. To enter the root organization mode, type / as the <i>org-name</i> .
Step 2	UCS-A /org # scope server-qual <i>server-qual-name</i>	Enters organization server qualification mode for the specified server pool policy qualification.
Step 3	UCS-A /org/server-qual # create cpu	Creates a CPU qualification and enters organization server qualification processor mode.
Step 4	UCS-A /org/server-qual/cpu # set arch { any dual-core-opteron intel-p4-c opteron pentium-4 turion-64 xeon xeon-mp }	Specifies the processor architecture type.
Step 5	UCS-A /org/server-qual/cpu # set maxcores { <i>max-core-num</i> unspecified }	Specifies the maximum number of processor cores.
Step 6	UCS-A /org/server-qual/cpu # set mincores { <i>min-core-num</i> unspecified }	Specifies the minimum number of processor cores.
Step 7	UCS-A /org/server-qual/cpu # set maxprocs { <i>max-proc-num</i> unspecified }	Specifies the maximum number of processors.
Step 8	UCS-A /org/server-qual/cpu # set minprocs { <i>min-proc-num</i> unspecified }	Specifies the minimum number of processors.
Step 9	UCS-A /org/server-qual/cpu # set maxthreads { <i>max-thread-num</i> unspecified }	Specifies the maximum number of threads.
Step 10	UCS-A /org/server-qual/cpu # set minthreads { <i>min-thread-num</i> unspecified }	Specifies the minimum number of threads.
Step 11	UCS-A /org/server-qual/cpu # set stepping { <i>step-num</i> unspecified }	Specifies the processor stepping number.

	Command or Action	Purpose
Step 12	UCS-A /org/server-qual/cpu # set model-regex <i>regex</i>	Specifies a regular expression that the processor name must match.
Step 13	UCS-A /org/server-qual/cpu # commit-buffer	Commits the transaction to the system configuration.

The following example creates and configures a CPU qualification and commits the transaction:

```
UCS-A# scope org /
UCS-A /org # scope server-qual ServPoolQual22
UCS-A /org/server-qual # create processor
UCS-A /org/server-qual/cpu* # set arch xeon
UCS-A /org/server-qual/cpu* # set maxcores 8
UCS-A /org/server-qual/cpu* # set mincores 4
UCS-A /org/server-qual/cpu* # set maxprocs 2
UCS-A /org/server-qual/cpu* # set minprocs 1
UCS-A /org/server-qual/cpu* # set maxthreads 16
UCS-A /org/server-qual/cpu* # set minthreads 8
UCS-A /org/server-qual/cpu* # set stepping 5
UCS-A /org/server-qual/cpu* # commit-buffer
UCS-A /org/server-qual/cpu #
```

Deleting a CPU Qualification

Procedure

	Command or Action	Purpose
Step 1	UCS-A# scope org <i>org-name</i>	Enters organization mode for the specified organization. To enter the root organization mode, type <i>/</i> as the <i>org-name</i> .
Step 2	UCS-A /org # scope server-qual <i>server-qual-name</i>	Enters organization server qualification mode for the specified server pool policy qualification.
Step 3	UCS-A /org/server-qual # delete cpu	Deletes the processor qualification.
Step 4	UCS-A /org/server-qual # commit-buffer	Commits the transaction to the system configuration.

The following example deletes the processor qualification and commits the transaction:

```
UCS-A# scope org /
UCS-A /org # scope server-qual ServPoolQual22
UCS-A /org/server-qual # delete cpu
UCS-A /org/server-qual* # commit-buffer
UCS-A /org/server-qual #
```

Creating a Power Group Qualification

Before You Begin

Create a server pool policy qualification.

Procedure

	Command or Action	Purpose
Step 1	UCS-A# scope org <i>org-name</i>	Enters organization mode for the specified organization. To enter the root organization mode, type / as the <i>org-name</i> .
Step 2	UCS-A /org # scope server-qual <i>server-qual-name</i>	Enters organization server qualification mode for the specified server pool policy qualification.
Step 3	UCS-A /org/server-qual # create power-group <i>power-group-name</i>	Creates a power group qualification for the specified power group name.
Step 4	UCS-A /org/server-qual # commit-buffer	Commits the transaction to the system configuration.

The following example configures a power group qualification for a power group called powergroup1 and commits the transaction:

```
UCS-A# scope org /
UCS-A /org # scope server-qual ServPoolQual122
UCS-A /org/server-qual # create power-group powergroup1
UCS-A /org/server-qual* # commit-buffer
UCS-A /org/server-qual #
```

Deleting a Power Group Qualification

Procedure

	Command or Action	Purpose
Step 1	UCS-A# scope org <i>org-name</i>	Enters organization mode for the specified organization. To enter the root organization mode, type / as the <i>org-name</i> .
Step 2	UCS-A /org # scope server-qual <i>server-qual-name</i>	Enters organization server qualification mode for the specified server pool policy qualification.
Step 3	UCS-A /org/server-qual # delete power-group <i>power-group-name</i>	Deletes the specified power group qualification.
Step 4	UCS-A /org/server-qual # commit-buffer	Commits the transaction to the system configuration.

The following example deletes a power group qualification for a power group called powergroup1 and commits the transaction:

```
UCS-A# scope org /
UCS-A /org # scope server-qual ServPoolQual22
UCS-A /org/server-qual # delete power-group powergroup1
UCS-A /org/server-qual* # commit-buffer
UCS-A /org/server-qual #
```

Creating a Memory Qualification

Before You Begin

Create a server pool policy qualification.

Procedure

	Command or Action	Purpose
Step 1	UCS-A# scope org <i>org-name</i>	Enters organization mode for the specified organization. To enter the root organization mode, type <i>/</i> as the <i>org-name</i> .
Step 2	UCS-A /org # scope server-qual <i>server-qual-name</i>	Enters organization server qualification mode for the specified server pool policy qualification.
Step 3	UCS-A /org/server-qual # create memory	Creates a memory qualification and enters organization server qualification memory mode.
Step 4	UCS-A /org/server-qual/memory # set clock { <i>clock-num</i> unspec }	Specifies the memory clock speed.
Step 5	UCS-A /org/server-qual/memory # set maxcap { <i>max-cap-num</i> unspec }	Specifies the maximum capacity of the memory array.
Step 6	UCS-A /org/server-qual/memory # set mincap { <i>min-cap-num</i> unspec }	Specifies the minimum capacity of the memory array.
Step 7	UCS-A /org/server-qual/memory # set speed { <i>speed-num</i> unspec }	Specifies the memory data rate.
Step 8	UCS-A /org/server-qual/memory # set units { <i>unit-num</i> unspec }	Specifies the number of memory units (DRAM chips mounted to the memory board).
Step 9	UCS-A /org/server-qual/memory # set width { <i>width-num</i> unspec }	Specifies the bit width of the data bus.
Step 10	UCS-A /org/server-qual/memory # commit-buffer	Commits the transaction to the system configuration.

The following example creates and configures a memory qualification and commits the transaction:

```
UCS-A# scope org /
UCS-A /org # scope server-qual ServPoolQual22
UCS-A /org/server-qual # create memory
UCS-A /org/server-qual/memory* # set clock 1067
```

```

UCS-A /org/server-qual/memory* # set maxcap 4096
UCS-A /org/server-qual/memory* # set mincap 2048
UCS-A /org/server-qual/memory* # set speed unspec
UCS-A /org/server-qual/memory* # set units 16
UCS-A /org/server-qual/memory* # set width 64
UCS-A /org/server-qual/memory* # commit-buffer
UCS-A /org/server-qual/memory #

```

Deleting a Memory Qualification

Procedure

	Command or Action	Purpose
Step 1	UCS-A# scope org <i>org-name</i>	Enters organization mode for the specified organization. To enter the root organization mode, type / as the <i>org-name</i> .
Step 2	UCS-A /org # scope server-qual <i>server-qual-name</i>	Enters organization server qualification mode for the specified server pool policy qualification.
Step 3	UCS-A /org/server-qual # delete memory	Deletes the memory qualification.
Step 4	UCS-A /org/server-qual # commit-buffer	Commits the transaction to the system configuration.

The following example deletes the memory qualification and commits the transaction:

```

UCS-A# scope org /
UCS-A /org # scope server-qual ServPoolQual22
UCS-A /org/server-qual # delete memory
UCS-A /org/server-qual* # commit-buffer
UCS-A /org/server-qual #

```

Creating a Physical Qualification

Before You Begin

Create a server pool policy qualification.

Procedure

	Command or Action	Purpose
Step 1	UCS-A# scope org <i>org-name</i>	Enters organization mode for the specified organization. To enter the root organization mode, type / as the <i>org-name</i> .
Step 2	UCS-A /org # scope server-qual <i>server-qual-name</i>	Enters organization server qualification mode for the specified server pool policy qualification.

	Command or Action	Purpose
Step 3	UCS-A /org/server-qual # create physical-qual	Creates a physical qualification and enters organization server qualification physical mode.
Step 4	UCS-A /org/server-qual/physical-qual # set model-regex regex	Specifies a regular expression that the model name must match.
Step 5	UCS-A /org/server-qual/physical-qual # commit-buffer	Commits the transaction to the system configuration.

The following example creates and configures a physical qualification and commits the transaction:

```
UCS-A# scope org /
UCS-A /org # scope server-qual ServPoolQual22
UCS-A /org/server-qual # create physical-qual
UCS-A /org/server-qual/physical-qual* # set model-regex
UCS-A /org/server-qual/physical-qual* # commit-buffer
UCS-A /org/server-qual/physical-qual #
```

Deleting a Physical Qualification

Procedure

	Command or Action	Purpose
Step 1	UCS-A# scope org org-name	Enters organization mode for the specified organization. To enter the root organization mode, type / as the <i>org-name</i> .
Step 2	UCS-A /org # scope server-qual server-qual-name	Enters organization server qualification mode for the specified server pool policy qualification.
Step 3	UCS-A /org/server-qual # delete physical-qual	Deletes the physical qualification.
Step 4	UCS-A /org/server-qual # commit-buffer	Commits the transaction to the system configuration.

The following example deletes a physical qualification and commits the transaction:

```
UCS-A# scope org /
UCS-A /org # scope server-qual ServPoolQual22
UCS-A /org/server-qual # delete physical-qual
UCS-A /org/server-qual* # commit-buffer
UCS-A /org/server-qual #
```

Creating a Storage Qualification

Before You Begin

Create a server pool policy qualification.

Procedure

	Command or Action	Purpose
Step 1	UCS-A# scope org <i>org-name</i>	Enters organization mode for the specified organization. To enter the root organization mode, type / as the <i>org-name</i> .
Step 2	UCS-A /org # scope server-qual <i>server-qual-name</i>	Enters organization server qualification mode for the specified server pool policy qualification.
Step 3	UCS-A /org/server-qual # create storage	Creates a storage qualification and enters organization server qualification storage mode.
Step 4	UCS-A /org/server-qual/storage # set blocksize { <i>block-size-num</i> unspecified }	Specifies the storage block size.
Step 5	UCS-A /org/server-qual/storage # set maxcap { <i>max-cap-num</i> unspecified }	Specifies the maximum capacity of the storage array.
Step 6	UCS-A /org/server-qual/storage # set mincap { <i>min-cap-num</i> unspecified }	Specifies the minimum capacity of the storage array.
Step 7	UCS-A /org/server-qual/storage # set numberofblocks { <i>block-num</i> unspecified }	Specifies the number of blocks.
Step 8	UCS-A /org/server-qual/storage # set perdiskcap { <i>disk-cap-num</i> unspecified }	Specifies the per-disk capacity.
Step 9	UCS-A /org/server-qual/storage # set units { <i>unit-num</i> unspecified }	Specifies the number of storage units.
Step 10	UCS-A /org/server-qual/storage # commit-buffer	Commits the transaction to the system configuration.

The following example creates and configures a storage qualification and commits the transaction:

```
UCS-A# scope org /
UCS-A /org # scope server-qual ServPoolQual22
UCS-A /org/server-qual # create storage
UCS-A /org/server-qual/storage* # set blocksize 512
UCS-A /org/server-qual/storage* # set maxcap 420000
UCS-A /org/server-qual/storage* # set mincap 140000
UCS-A /org/server-qual/storage* # set numberofblocks 287277984
UCS-A /org/server-qual/storage* # set perdiskcap 140000
UCS-A /org/server-qual/storage* # set units 1
UCS-A /org/server-qual/storage* # commit-buffer
UCS-A /org/server-qual/storage #
```

Deleting a Storage Qualification

Procedure

	Command or Action	Purpose
Step 1	UCS-A# scope org <i>org-name</i>	Enters organization mode for the specified organization. To enter the root organization mode, type <i>/</i> as the <i>org-name</i> .
Step 2	UCS-A /org # scope server-qual <i>server-qual-name</i>	Enters organization server qualification mode for the specified server pool policy qualification.
Step 3	UCS-A /org/server-qual # delete storage	Deletes the storage qualification.
Step 4	UCS-A /org/server-qual/ # commit-buffer	Commits the transaction to the system configuration.

The following example deletes the storage qualification and commits the transaction:

```
UCS-A# scope org /
UCS-A /org # scope server-qual ServPoolQual22
UCS-A /org/server-qual # delete storage
UCS-A /org/server-qual* # commit-buffer
UCS-A /org/server-qual #
```

Configuring vNIC/vHBA Placement Policies

vNIC/vHBA Placement Policies

vNIC/vHBA placement policies are used to determine what types of vNICs or vHBAs can be assigned to the physical adapters on a server. Each vNIC/vHBA placement policy contains four virtual network interface connections (vCons) that are virtual representations of the physical adapters. When a vNIC/vHBA placement policy is assigned to a service profile, and the service profile is associated with a server, the vCons in the vNIC/vHBA placement policy are assigned to the physical adapters.

If you do not include a vNIC/vHBA placement policy in the service profile or you use the default configuration for a server with two adapters, Cisco UCS Manager defaults to the **All** configuration and equally distributes the vNICs and vHBAs between the adapters.

You can use this policy to assign vNICs or vHBAs to either of the two vCons. Cisco UCS Manager uses the vCon assignment to determine how to assign the vNICs and vHBAs to the physical adapter during service profile association.

- **All**—All configured vNICs and vHBAs can be assigned to the vCon, whether they are explicitly assigned to it, unassigned, or dynamic.
- **Assigned Only**—vNICs and vHBAs must be explicitly assigned to the vCon. You can assign them explicitly through the service profile or the properties of the vNIC or vHBA.

- **Exclude Dynamic**—Dynamic vNICs and vHBAs cannot be assigned to the vCon. The vCon can be used for all static vNICs and vHBAs, whether they are unassigned or explicitly assigned to it.
- **Exclude Unassigned**—Unassigned vNICs and vHBAs cannot be assigned to the vCon. The vCon can be used for dynamic vNICs and vHBAs and for static vNICs and vHBAs that are explicitly assigned to it.

vCon to Adapter Placement

Cisco UCS Manager maps every vCon in a service profile to a physical adapter on the server. How that mapping occurs and how the vCons are assigned to a specific adapter in a server with two adapters depends upon the type of server. You must consider this placement when you configure the vNIC/vHBA placement policy to assign vNICs and vHBAs to vCons.



Note

vCon to adapter placement is not dependent upon the PCIE slot number of the adapter. The adapter numbers used for the purpose of vCon placement are not the PCIE slot numbers of the adapters, but the ID assigned to them during server discovery.

vCon to Adapter Placement for N20-B6620-2 and N20-B6625-2 Blade Servers

In these blade servers, the adapters are numbered left to right, but vCons are numbered right to left. If the server has a single adapter, all vCons are assigned to that adapter. However, if the server has two adapters, the vCons are assigned to the adapters in reverse order, as follows:

- Adapter1 is assigned vCon2 and vCon4
- Adapter2 is assigned vCon1 and vCon3

vCon to Adapter Placement for All Other Supported Servers

For all other servers supported by Cisco UCS, the vCon assignment depends upon the number of adapters in the server, as follows:

Table 14: vCon to Adapter Placement by Number of Adapters in Server

Number of Adapters	vCon1 Assignment	vCon2 Assignment	vCon3 Assignment	vCon4 Assignment
1	Adapter1	Adapter1	Adapter1	Adapter1
2	Adapter1	Adapter2	Adapter1	Adapter2
3	Adapter1	Adapter2	Adapter3	Adapter2
4	Adapter1	Adapter2	Adapter3	Adapter4

vNIC/vHBA to vCon Assignment

Cisco UCS Manager provides two options for assigning vNICs and vHBAs to vCons through the vNIC/vHBA placement policy: explicit assignment and implicit assignment.

Explicit Assignment of vNICs and vHBAs

With explicit assignment, you specify the vCon and, therefore, the adapter to which a vNIC or vHBA is assigned. Use this assignment option when you need to determine how the vNICs and vHBAs are distributed between the adapters on a server.

To configure a vCon and the associated vNICs and vHBAs for explicit assignment, do the following:

- Set the vCon configuration to any of the available options. You can configure the vCons through a vNIC/vHBA placement policy or in the service profile associated with the server. If a vCon is configured for **All**, you can still explicitly assign a vNIC or vHBA to that vCon.
- Assign the vNICs and vHBAs to a vCon. You can make this assignment through the Virtual Host Interface Placement properties of the vNIC or vHBA or in the service profile associated with the server.

If you attempt to assign a vNIC or vHBA to a vCon that is not configured for that type of vNIC or vHBA, Cisco UCS Manager displays a message box to advise you of the configuration error.

During service profile association, Cisco UCS Manager validates the configured placement of the vNICs and vHBAs against the number and capabilities of the physical adapters in the server before assigning the vNICs and vHBAs according to the configuration in the policy. Load distribution is based upon the explicit assignments to the vCons and adapters configured in this policy.

If the adapters do not support the assignment of one or more vNICs or vHBAs, Cisco UCS Manager raises a fault against the service profile.



Note

vCon to adapter assignment occurs in a round-robin order. This order means that vNICs are placed on the adapters in the following order: vcon-1, vcon-3, vcon-2, vcon-4. As a result, under the following circumstances, the PCIE order of vNICs can be different than the explicit assignment configured in Cisco UCS Manager:

- In a server with two adapters when vNICs are explicitly assigned to all four vCons.
- When a service profile that includes explicit assignment is migrated from a server with a higher number of adapters to one with a lower number of adapters.

Implicit Assignment of vNICs and vHBAs

With implicit assignment, Cisco UCS Manager determines the vCon and, therefore, the adapter to which a vNIC or vHBA is assigned according to the capability of the adapter. Use this assignment option if the adaptor to which a vNIC or vHBA is assigned is not important to your system configuration.

To configure a vCon for implicit assignment, do the following:

- Set the vCon configuration to **All**, **Exclude Dynamic**, or **Exclude Unassigned**. You can configure the vCons through a vNIC/vHBA placement policy or in the service profile associated with the server.

- Do not set the vCon configuration to **Exclude Assigned**. Implicit assignment cannot be performed with this setting.
- Do not assign any vNICs or vHBAs to a vCon.

During service profile association, Cisco UCS Manager verifies the number and capabilities of the physical adapters in the server and assigns the vNICs and vHBAs accordingly. Load distribution is based upon the capabilities of the adapters, and placement of the vNICs and vHBAs is performed according to the actual order determined by the system. For example, if one adapter can accommodate more vNICs than another, that adapter is assigned more vNICs.

If the adapters cannot support the number of vNICs and vHBAs configured for that server, Cisco UCS Manager raises a fault against the service profile.

Implicit Assignment of vNICs in a Mixed Adapter Environment

The implicit assignment of vNICs functions differently for a server that has mixed adapters, as follows:

- A dual slot server that has one VIC adapter and one non-VIC adapter, which have different capabilities. For example, a server that contains a Cisco UCS M81KR Virtual Interface Card and a Cisco UCS CNA M71KR-E adapter.
- A configuration that includes both dynamic vNICs and static vNICs.

When you assign vNICs implicitly for a dual slot server that has one VIC adapter and non-VIC adapter, Cisco UCS Manager typically assigns one vNIC to each adapter. The remaining vNICs are assigned according to the relative capabilities of the adapters. The following are examples of the relative capabilities of some of the supported adapters:

- Cisco UCS M81KR Virtual Interface Card (128 vNICs) and Cisco UCS CNA M71KR-E Adapter (2 vNICs) have a 64:1 capability ratio
- Cisco UCS M81KR Virtual Interface Card and Cisco UCS CNA M72KR-E have a 64:1 capability ratio
- Cisco UCS CNA M72KR-E and Cisco UCS CNA M72KR-E have a 1:1 capability ratio
- Cisco UCS M82-8P Virtual Interface Card and Cisco UCS CNA M71KR-E adapter have a 128:1 capability ratio
- Cisco UCS M82-8P Virtual Interface Card and Cisco UCS M81KR Virtual Interface Card have a 2:1 capability ratio.

For example, a Cisco UCS M81KR Virtual Interface Card can handle up to 128 vNICs, while a Cisco UCS CNA M71KR-E can only handle 2 vNICs. This difference gives those adapters a 64:1 ratio. If a dual slot blade server has one of each and you choose to allow implicit assignment of vNICs by Cisco UCS Manager, the load balancing ratio assigns the majority of the vNICs to the Cisco UCS M81KR Virtual Interface Card, as follows:

Total Number of vNICs	vNICs Assigned to Cisco UCS M81KR Virtual Interface Card	vNICs Assigned to Cisco UCS CNA M71KR-E Adapter
20	19	1
130	128	2

**Note**

Exceptions to this implicit assignment occur if you configure the vNICs for fabric failover and if you configure dynamic vNICs for the server.

For a configuration that includes vNIC fabric failover where one adapter does not support vNIC failover, Cisco UCS Manager implicitly assigns all vNICs which have fabric failover enabled to the adapter that supports them. If the configuration only includes vNICs that are configured for fabric failover, no vNICs are implicitly assigned to the adapter which does not support them. If some vNICs are configured for fabric failover and some are not, Cisco UCS Manager assigns all failover vNICs to the adapter which supports them and a minimum of one non-failover vNIC to the adapter which does not support them, according to the ratio above.

For a configuration that includes dynamic vNICs, the same implicit assignment would occur. Cisco UCS Manager assigns all dynamic vNICs to the adapter that supports them. However, with a combination of dynamic vNICs and static vNICs, at least one static vNIC is assigned to the adapter that does not support dynamic vNICs.

Configuring a vNIC/vHBA Placement Policy

Procedure

	Command or Action	Purpose
Step 1	UCS-A# scope org <i>org-name</i>	Enters organization mode for the specified organization. To enter the root organization mode, type / as the <i>org-name</i> .
Step 2	UCS-A /org # create vcon-policy <i>policy-name</i>	Creates the specified vNIC/vHBA placement profile and enters organization vcon policy mode.
Step 3	UCS-A /org/vcon-policy # set descr <i>description</i>	(Optional) Provides a description for the vNIC/vHBA Placement Profile. Note If your description includes spaces, special characters, or punctuation, you must begin and end your description with quotation marks. The quotation marks will not appear in the description field of any show command output.
Step 4	UCS-A /org/vcon-policy # set vcon {1 2} selection {all assigned-only exclude-dynamic exclude-unassigned}	Specifies the selection preference for the specified vCon.
Step 5	UCS-A /org/vcon-policy # commit-buffer	Commits the transaction.

The following example creates a vNIC/vHBA placement policy named Adapter1All, places all vNICs and vHBAs on adapter 1, and commits the transaction:

```
UCS-A# scope org /
UCS-A /org # create vcon-policy Adapter1
UCS-A /org/vcon-policy* # set descr "This profile places all vNICs and vHBAs on adapter 1."
UCS-A /org/vcon-policy* # set vcon 1 selection all
UCS-A /org/vcon-policy* # commit-buffer
UCS-A /org/vcon-policy* #
UCS-A /org #
```

Deleting a vNIC/vHBA Placement Policy

Procedure

	Command or Action	Purpose
Step 1	UCS-A# scope org <i>org-name</i>	Enters organization mode for the specified organization. To enter the root organization mode, type / as the <i>org-name</i> .
Step 2	UCS-A /org # delete vcon-policy <i>policy-name</i>	Deletes the specified vNIC/vHBA placement profile.
Step 3	UCS-A /org # commit-buffer	Commits the transaction.

The following example deletes the vNIC/vHBA placement profile named Adapter1All and commits the transaction:

```
UCS-A# scope org /
UCS-A /org # delete vcon-policy Adapter1All
UCS-A /org* # commit-buffer
UCS-A /org #
```

Explicitly Assigning a vNIC to a vCon

Before You Begin

Configure the vCons through a vNIC/vHBA placement policy or in the service profile with one of the following values:

- **Assigned Only**
- **Exclude Dynamic**
- **Exclude Unassigned**

If a vCon is configured for **All**, you can still explicitly assign a vNIC or vHBA to that vCon. However, you have less control with this configuration.

Procedure

	Command or Action	Purpose
Step 1	UCS-A# scope org <i>org-name</i>	Enters organization mode for the organization which contains the service profile whose vNICs you want to explicitly assign to a vCon. To enter the root organization mode, type / as the <i>org-name</i> .
Step 2	UCS-A /org # scope service-profile <i>profile-name</i>	Enters organization service profile mode for the specified service.
Step 3	UCS-A /org/service-profile # scope vnic <i>vnic-name</i>	Enters organization service profile mode for the specified vnic.
Step 4	UCS-A /org/service-profile/vnic # set vcon {1 2 3 4 any}	Sets the virtual network interface connection (vCon) placement for the specified vNIC. Entering a value of any allows Cisco UCS Manager to determine the vCon to which the vNIC is assigned.
Step 5	UCS-A /org/service-profile/vnic # set order { <i>order-num</i> unspecified}	Specifies the desired PCI order for the vNIC. Valid values include 0-128 and unspecified.
Step 6	UCS-A /org/service-profile/vnic # commit-buffer	Commits the transaction to the system configuration.

The following example sets the vCon placement for a vNIC called vnic3 to 2, sets the desired order to 10, and commits the transaction:

```
UCS-A# scope org /
UCS-A /org # scope service-profile accounting
UCS-A /org/service-profile # scope vnic vnic3
UCS-A /org/service-profile/vnic # set vcon 2
UCS-A /org/service-profile/vnic* # set order 10
UCS-A /org/service-profile/vnic* # commit-buffer
UCS-A /org/service-profile/vnic #
```

Explicitly Assigning a vHBA to a vCon

Before You Begin

Configure the vCons through a vNIC/vHBA placement policy or in the service profile with one of the following values:

- **Assigned Only**
- **Exclude Dynamic**
- **Exclude Unassigned**

If a vCon is configured for **All**, you can still explicitly assign a vNIC or vHBA to that vCon. However, you have less control with this configuration.

Procedure

	Command or Action	Purpose
Step 1	UCS-A# scope org <i>org-name</i>	Enters organization mode for the organization which contains the service profile whose vHBAs you want to explicitly assign to a vCon. To enter the root organization mode, type / as the <i>org-name</i> .
Step 2	UCS-A /org # scope service-profile <i>profile-name</i>	Enters organization service profile mode for the specified service.
Step 3	UCS-A /org/service-profile # scope vhma <i>vhba-name</i>	Enters organization service profile mode for the specified vHBA.
Step 4	UCS-A /org/service-profile/vhma # set vcon {1 2 3 4 any}	Sets the virtual network interface connection (vCon) placement for the specified vHBA. Entering a value of any allows Cisco UCS Manager to determine the vCon to which the vHBA is assigned.
Step 5	UCS-A /org/service-profile/vhma # set order { <i>order-num</i> unspecified}	Specifies the desired PCI order for the vHBA. Valid desired order number values include 0-128 and unspecified.
Step 6	UCS-A /org/service-profile/vhma # commit-buffer	Commits the transaction to the system configuration.

The following example sets the vCon placement for a vHBA called vhma3 to 2, sets the desired order to 10, and commits the transaction:

```
UCS-A# scope org /
UCS-A /org # scope service-profile accounting
UCS-A /org/service-profile # scope vhma vhma3
UCS-A /org/service-profile/vhma # set vcon 2
UCS-A /org/service-profile/vhma* # set order 10
UCS-A /org/service-profile/vhma* # commit-buffer
UCS-A /org/service-profile/vhma #
```



CHAPTER 29

Configuring Server Boot

This chapter includes the following sections:

- [Boot Policy, page 413](#)
- [Creating a Boot Policy, page 414](#)
- [SAN Boot, page 415](#)
- [iSCSI Boot, page 418](#)
- [LAN Boot, page 448](#)
- [Local Disk Boot, page 449](#)
- [Virtual Media Boot, page 450](#)
- [Deleting a Boot Policy, page 451](#)

Boot Policy

The boot policy determines the following:

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

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

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



Important

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

Creating a Boot Policy

You can also create a local boot policy that is restricted to a service profile or service profile template. However, except for iSCSI boot, we recommend that you create a global boot policy that can be included in multiple service profiles or service profile templates.

Before You Begin

If you are creating a boot policy that boots the server from a SAN LUN and you require reliable SAN boot operations, you must first remove all local disks from servers associated with a service profile that includes the boot policy.

Procedure

	Command or Action	Purpose
Step 1	UCS-A# scope org <i>org-name</i>	Enters organization mode for the specified organization. To enter the root organization mode, type / as the <i>org-name</i> .
Step 2	UCS-A /org # create boot-policy <i>policy-name</i> [purpose { operational utility }]	Creates a boot policy with the specified policy name, and enters organization boot policy mode. When you create the boot policy, specify the operational option. This ensures that the server boots from the operating system installed on the server. The utility options is reserved and should only be used if instructed to do so by a Cisco representative.
Step 3	UCS-A /org/boot-policy # set descr <i>description</i>	(Optional) Provides a description for the boot policy. Note If your description includes spaces, special characters, or punctuation, you must begin and end your description with quotation marks. The quotation marks do not appear in the description field of any show command output.
Step 4	UCS-A /org/boot-policy # set reboot-on-update { no yes }	Specifies whether the servers using this boot policy are automatically rebooted after you make changes to the boot order.
Step 5	UCS-A /org/boot-policy # set enforce-vnic-name { no yes }	If you choose yes , Cisco UCS Manager uses any vNICs or vHBAs defined in the Boot Order . If you choose no , Cisco UCS Manager uses the priority specified in the vNIC or vHBA.
Step 6	UCS-A /org/boot-policy # commit-buffer	Commits the transaction to the system configuration.

The following example creates a boot policy named boot-policy-LAN, provides a description for the boot policy, specifies that servers using this policy will not be automatically rebooted when the boot order is changed, and commits the transaction:

```
UCS-A# scope org /
UCS-A /org* # create boot-policy boot-policy-LAN purpose operational
UCS-A /org/boot-policy* # set descr "Boot policy that boots from the LAN."
UCS-A /org/boot-policy* # set reboot-on-update no
UCS-A /org/boot-policy* # commit-buffer
UCS-A /org/boot-policy #
```

What to Do Next

Configure one or more of the following boot options for the boot policy and set their boot order:

- **LAN Boot** —Boots from a centralized provisioning server. It is frequently used to install operating systems on a server from that server.

If you choose the LAN Boot option, continue to [Configuring a LAN Boot for a Boot Policy](#), on page 448.

- **Storage Boot** — Boots from an operating system image on the SAN. You can specify a primary and a secondary SAN boot. If the primary boot fails, the server attempts to boot from the secondary.

We recommend that you use a SAN boot, because it offers the most service profile mobility within the system. If you boot from the SAN, when you move a service profile from one server to another, the new server boots from exactly the same operating system image. Therefore, the new server appears to be exactly the same server to the network.

If you choose the Storage Boot option, continue to [Configuring a SAN Boot for a Boot Policy](#), on page 416.

- **Virtual Media Boot** —Mimics the insertion of a physical CD into a server. It is typically used to manually install operating systems on a server.

If you choose the Virtual Media boot option, continue to [Configuring a Virtual Media Boot for a Boot Policy](#), on page 450.



Tip

We recommend that the boot order in a boot policy include either a local disk or a SAN LUN, but not both, to avoid the possibility of the server booting from the wrong storage type. If you configure a local disk and a SAN LUN for the boot order storage type and the operating system or logical volume manager (LVM) is configured incorrectly, the server might boot from the local disk rather than the SAN LUN.

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

Include the boot policy in a service profile and/or template.

SAN Boot

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

We recommend that you use a SAN boot, because it offers the most service profile mobility within the system. If you boot from the SAN when you move a service profile from one server to another, the new server boots from the exact same operating system image. Therefore, the new server appears to be the exact same server to the network.

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

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

Configuring a SAN Boot for a Boot Policy



Tip

We recommend that the boot order in a boot policy include either a local disk or a SAN LUN, but not both, to avoid the possibility of the server booting from the wrong storage type. If you configure a local disk and a SAN LUN for the boot order storage type and the operating system or logical volume manager (LVM) is configured incorrectly, the server might boot from the local disk rather than the SAN LUN.

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

This procedure continues directly from [Creating a Boot Policy](#).

Before You Begin

Create a boot policy to contain the SAN boot configuration.



Note

If you are creating a boot policy that boots the server from a SAN LUN and you require reliable SAN boot operations, we recommend that you first remove all local disks from servers associated with a service profile that includes the boot policy.

Procedure

	Command or Action	Purpose
Step 1	UCS-A# scope org <i>org-name</i>	Enters organization mode for the specified organization. To enter the root organization mode, type / as the <i>org-name</i> .
Step 2	UCS-A /org # scope boot-policy <i>policy-name</i>	Enters organization boot policy mode for the specified boot policy.
Step 3	UCS-A /org/boot-policy # create storage	Creates a SAN boot for the boot policy and enters organization boot policy storage mode.

	Command or Action	Purpose
Step 4	UCS-A /org/boot-policy/storage # set order {1 2 3 4}	Sets the boot order for the SAN boot.
Step 5	UCS-A /org/boot-policy/storage # create san-image {primary secondary}	Creates a SAN image location, and if the san-image option is specified, enters organization boot policy storage SAN image mode. The use of the terms primary or secondary boot devices does not imply a boot order. The effective order of boot devices within the same device class is determined by PCIe bus scan order.
Step 6	UCS-A /org/boot-policy/storage/san-image # set vhba vhba-name	Specifies the vHBA to be used for the SAN boot.
Step 7	UCS-A /org/boot-policy/storage/san-image # create path {primary secondary}	Creates a primary or secondary SAN boot path and enters organization boot policy SAN path mode. The use of the terms primary or secondary boot devices does not imply a boot order. The effective order of boot devices within the same device class is determined by PCIe bus scan order.
Step 8	UCS-A /org/boot-policy/storage/san-image/path # set {lun lun-id wwn wwn-num}	Specifies the LUN or WWN to be used for the SAN path to the boot image.
Step 9	UCS-A /org/boot-policy/storage/san-image/path # commit-buffer	Commits the transaction to the system configuration.

The following example enters the boot policy named lab1-boot-policy, creates a SAN boot for the policy, sets the boot order to 1, creates a primary SAN image, uses a vHBA named vHBA2, creates primary path using LUN 967295200, and commits the transaction:

```
UCS-A# scope org /
UCS-A /org* # scope boot-policy lab1-boot-policy
UCS-A /org/boot-policy* # create storage
UCS-A /org/boot-policy/storage* # set order 1
UCS-A /org/boot-policy/storage* # create san-image primary
UCS-A /org/boot-policy/storage* # set vhba vHBA2
UCS-A /org/boot-policy/storage/san-image* # create path primary
UCS-A /org/boot-policy/storage/san-image/path* # set lun 967295200
UCS-A /org/boot-policy/storage/san-image/path* # commit-buffer
UCS-A /org/boot-policy/storage/san-image/path #
```

What to Do Next

Include the boot policy in a service profile and/or template.

iSCSI Boot

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

iSCSI boot is supported on the following Cisco UCS hardware:

- Cisco UCS server blades that have the Cisco UCS M51KR-B Broadcom BCM57711 network adapter and use the default MAC address provided by Broadcom.
- Cisco UCS M81KR Virtual Interface Card
- Cisco UCS M82-8P Virtual Interface Card

There are prerequisites that must be met before you configure iSCSI boot. For a list of these prerequisites, see [iSCSI Boot Guidelines and Prerequisites](#), on page 418.

For a high-level procedure for implementing iSCSI boot, see [Configuring iSCSI Boot](#), on page 420.

iSCSI Boot Process

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



Note

The iBFT only uses the first iSCSI vNIC and only Target 1 for the initiator-to-target initialization. This scenario is true even if a second target (Target 2) exists for the first iSCSI vNIC.

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




Note

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

iSCSI Boot Guidelines and Prerequisites

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

- To set up iSCSI boot from a Windows 2008 server where the second vNIC (failover vNIC) must boot from an iSCSI LUN, consult Microsoft Knowledge Base Article 976042. Microsoft has a known issue where Windows might fail to boot from an iSCSI drive or cause a bugcheck error if the networking hardware is changed. To work around this issue, follow the resolution recommended by Microsoft.
 - The storage array must be licensed for iSCSI boot and the array side LUN masking must be properly configured.
 - Two IP addresses must be determined, one for each iSCSI initiator. If possible, the IP addresses should be on the same subnet as the storage array. The IP addresses are assigned statically or dynamically using the Dynamic Host Configuration Protocol (DHCP).
 - You cannot configure boot parameters in the Global boot policy. Instead, after configuring boot parameters, you need to include the boot policy in the appropriate service profile.
 - The operating system (OS) must be iSCSI Boot Firmware Table (iBFT) compatible.
 - For Cisco UCS M51KR-B Broadcom BCM57711 network adapters:
 - Blades that use iSCSI boot must contain the Cisco UCS M51KR-B Broadcom BCM57711 network adapter. For information on installing or replacing an adapter card, see the *Cisco UCS B250 Extended Memory Blade Server Installation and Service Note*. The service note is accessible from the *Cisco UCS B-Series Servers Documentation Roadmap* at <http://www.cisco.com/go/unifiedcomputing/b-series-doc>.
 - Set the MAC addresses on the iSCSI device.
 - If you are using the DHCP Vendor ID (Option 43), the MAC address of an iSCSI device needs to be configured in `/etc/dhcpd.conf`.
 - HBA mode (also known as TCP offload) and the boot to target setting are supported. However, only Windows OS supports HBA mode during installation.
 - Before installing the OS, disable the boot to target setting in the iSCSI adapter policy, then after installing the OS, reenable the boot to target setting.
-
- 

Note Each time you change an adapter policy setting, the adapter reboots to apply the new setting.
-
- When installing the OS on the iSCSI target, the iSCSI target must be ordered *before* the device where the OS image resides. For example, if you are installing the OS on the iSCSI target from a CD, the boot order should be the iSCSI target and then the CD.
 - After the server has been iSCSI booted, do not modify the Initiator Name, Target name, LUN, iSCSI device IP, or Netmask/gateway using the Broadcom tool.
 - Do not interrupt the POST (power on self-test) process or the Cisco UCS M51KR-B Broadcom BCM57711 network adapter will fail to initialize.
- For Cisco UCS M81KR Virtual Interface Card and Cisco UCS M82-8P Virtual Interface Card:
 - Do not set MAC addresses on the iSCSI device.
 - HBA mode and the boot to target setting are *not* supported.

- When installing the OS on the iSCSI target, the iSCSI target must be ordered *after* the device where the OS image resides. For example, if you are installing the OS on the iSCSI target from a CD, the boot order should be the CD and then the iSCSI target.
- If you are using the DHCP Vendor ID (Option 43), the MAC address of the overlay vNIC needs to be configured in `/etc/dhcpd.conf`.
- After the server has been iSCSI booted, do not modify the IP details of the overlay vNIC.
- The VMware ESX/ESXi operating system does not support storing a core dump file to an iSCSI boot target LUN. Dump files must be written to a local disk.

Enabling MPIO on Windows



Note

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

Before You Begin

The server on which you enable MPIO must have a Cisco VIC driver.

Procedure

- Step 1** In the service profile associated with the server, configure the primary and secondary iSCSI vNICs. For more information, see [Creating an iSCSI vNIC in a Service Profile, on page 429](#).
- Step 2** Using the primary iSCSI vNIC, install the Windows operating system on the iSCSI target LUN.
- Step 3** After Windows installation is completed, use the Microsoft software iSCSI initiator to connect to the secondary iSCSI vNIC.
- Step 4** Enable MPIO on the host
- Step 5** In the service profile associated with the server, add the secondary iSCSI vNIC to the boot policy. For more information, see [Creating an iSCSI Boot Policy, on page 427](#).

Configuring iSCSI Boot

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

Procedure

	Command or Action	Purpose
Step 1	Configure the iSCSI boot adapter policy.	(Optional) For more information, see Creating an iSCSI Adapter Policy , on page 422
Step 2	Configure the authentication profiles to be used by the initiator and target.	(Optional) For more information, see Creating an Authentication Profile , on page 424
Step 3	If you plan to configure the iSCSI initiator to use an IP address from a pool of IP addresses, add a block of IP addresses to the iSCSI initiator pool.	(Optional) For more information, see Adding a Block of IP Addresses to the Initiator Pool , on page 425
Step 4	Create a boot policy that can be used in any service profile. Alternatively, you can create a local boot policy only for the specific service policy. However, we recommend that you create a boot policy that can be shared with multiple service profiles.	For more information about creating a boot policy that can be used in any service profile, see Creating an iSCSI Boot Policy , on page 427.
Step 5	If you created a boot policy that can be used in any service profile, you need to assign it to the service profile. Otherwise, proceed to the next step.	For more information, see Creating a Service Profile Template , on page 467.
Step 6	Configure an Ethernet vNIC in a service profile.	The Ethernet vNIC is used as the overlay vNIC for the iSCSI device. For more information, see Configuring a vNIC for a Service Profile , on page 474.
Step 7	Create an iSCSI vNIC in a service profile.	For more information, see Creating an iSCSI vNIC in a Service Profile , on page 429
Step 8	Set the iSCSI initiator to boot using a static IP Address, an IP address from an IP pool, or DHCP.	See either Creating an iSCSI Initiator that Boots Using a Static IP Address , on page 432, Creating an iSCSI Initiator that Boots Using an IP Address from an IP Pool , on page 434, or Creating an iSCSI Initiator that Boots Using DHCP , on page 436.
Step 9	Create an iSCSI static or auto target.	For more information, see either Creating an iSCSI Static Target , on page 442 or Creating an iSCSI Auto Target , on page 445.
Step 10	Associate the service profile with a server.	For more information, see Associating a Service Profile with a Blade Server or Server Pool , on page 485.

	Command or Action	Purpose
Step 11	Verify the iSCSI boot operation.	For more information, see Verifying iSCSI Boot, on page 447
Step 12	Install the OS on the server.	For more information, see one of the following guides: <ul style="list-style-type: none"> • <i>Cisco UCS B-Series Blade Servers VMware Installation Guide</i> • <i>Cisco UCS B-Series Blade Servers Linux Installation Guide</i> • <i>Cisco UCS B-Series Blade Servers Windows Installation Guide</i>
Step 13	Boot the server.	

Creating an iSCSI Adapter Policy

Procedure

	Command or Action	Purpose
Step 1	UCS-A# scope org <i>org-name</i>	Enters organization mode for the specified organization. To enter the root organization mode, type / as the <i>org-name</i> .
Step 2	UCS-A /org # create iscsi-policy <i>policy-name</i>	Creates the iSCSI adapter policy.
Step 3	UCS-A /org/iscsi-policy # set descr <i>description</i>	(Optional) Provides a description for the iSCSI adapter policy.
Step 4	UCS-A /org/iscsi-policy # set iscsi-protocol-item connection-timeout <i>timeout-secs</i>	The number of seconds to wait until Cisco UCS assumes that the initial login has failed and the iSCSI adapter is unavailable. Enter an integer between 0 and 255. If you enter 0, Cisco UCS uses the value set in the adapter firmware (default: 15 seconds).
Step 5	UCS-A /org/iscsi-policy # set iscsi-protocol-item dhcp-timeout <i>timeout-secs</i>	The number of seconds to wait before the initiator assumes that the DHCP server is unavailable. Enter an integer between 60 and 300 (default: 60 seconds).
Step 6	UCS-A /org/iscsi-policy # set iscsi-protocol-item lun-busy-retry-count <i>num</i>	The number of times to retry the connection in case of a failure during iSCSI LUN discovery.

	Command or Action	Purpose
		Enter an integer between 0 and 60. If you enter 0, Cisco UCS uses the value set in the adapter firmware (default: 15 seconds).
Step 7	UCS-A /org/iscsi-policy # set iscsi-protocol-item tcp-time-stamp {no yes}	Specifies whether to apply a TCP timestamp. With this setting, transmitted packets are given a time stamp of when the packet was sent so that the packet's round-trip time can be calculated, when needed. This setting applies only to Cisco UCS M51KR-B Broadcom BCM57711 adapters.
Step 8	UCS-A /org/iscsi-policy # set iscsi-protocol-item hbamode {no yes}	Specifies whether to enable HBA mode. This option should only be enabled for servers with the Cisco UCS NIC M51KR-B adapter running the Windows operating system.
Step 9	UCS-A /org/iscsi-policy # set iscsi-protocol-item boottotarget {no yes}	Specifies whether to boot from the iSCSI target. This option only applies to servers with the Cisco UCS NIC M51KR-B adapter. It should be disabled until you have installed an operating system on the server.
Step 10	UCS-A /org/iscsi-policy # commit-buffer	Commits the transaction to the system configuration.

The following example shows how to create an iSCSI adapter policy called `iscsiboot`, set the connection timeout, DHCP timeout, and LUN busy retry count, apply a TCP timestamp, and commit the transaction:

```
UCS-A# scope org /
UCS-A /org # create iscsi-policy iscsiboot
UCS-A /org/iscsi-policy* # set iscsi-protocol-item connection-timeout 60
UCS-A /org/iscsi-policy* # set iscsi-protocol-item dhcp-timeout 200
UCS-A /org/iscsi-policy* # set iscsi-protocol-item lun-busy-retry-count 5
UCS-A /org/iscsi-policy* # set iscsi-protocol-item tcp-time-stamp yes
UCS-A /org/iscsi-policy* # set iscsi-protocol-item hbamode yes
UCS-A /org/iscsi-policy* # set iscsi-protocol-item boottotarget yes
UCS-A /org/iscsi-policy* # commit-buffer
UCS-A /org/iscsi-policy #
```

What to Do Next

Include the adapter policy in a service profile and/or template.

Deleting an iSCSI Adapter Policy

Procedure

	Command or Action	Purpose
Step 1	UCS-A# scope org <i>org-name</i>	Enters organization mode for the specified organization. To enter the root organization mode, type / as the <i>org-name</i> .
Step 2	UCS-A /org # delete iscsi-policy <i>policy-name</i>	Deletes the iSCSI adapter policy.
Step 3	UCS-A /org # commit-buffer	Commits the transaction to the system configuration.

The following example shows how to delete an iSCSI adapter policy named iscsi-adapter-pol and commit the transaction:

```
UCS-A# scope org /
UCS-A /org # delete iscsi-policy iscsi-adapter-pol
UCS-A /org* # commit-buffer
UCS-A /org #
```

Creating an Authentication Profile

If you use authentication for iSCSI boot, you need to create an authentication profile for both the initiator and target.

Procedure

	Command or Action	Purpose
Step 1	UCS-A# scope org <i>org-name</i>	Enters organization mode for the specified organization. To enter the root organization mode, type / as the <i>org-name</i> .
Step 2	UCS-A /org # create auth-profile <i>profile-name</i>	Creates an authentication profile with the specified name. The name can be up to 16 alphanumeric characters.
Step 3	UCS-A /org/auth-profile* # set user-id <i>id-name</i>	Creates a log in for authentication.
Step 4	UCS-A /org/auth-profile* # set password	Creates a password for authentication.
Step 5	UCS-A /org/auth-profile* # commit-buffer	Commits the transaction to the system configuration.
Step 6	UCS-A /org/auth-profile* # exit	Exits the current mode.
Step 7	Repeat steps 2 through 6 to create an authentication profile for the target.	

The following example shows how to create an authentication profile for an initiator and target and commit the transaction:

```
UCS-A# scope org
UCS-A /org # create auth-profile InitAuth
UCS-A /org/auth-profile* # set user-id init
UCS-A /org/auth-profile* # set password
Enter a password:
Confirm the password:
UCS-A /org/auth-profile* # commit-buffer
UCS-A /org/auth-profile # exit
UCS-A /org # create auth-profile TargetAuth
UCS-A /org/auth-profile* # set user-id target
UCS-A /org/auth-profile* # set password
Enter a password:
Confirm the password:
UCS-A /org/auth-profile* # commit-buffer
UCS-A /org/auth-profile # exit
```

What to Do Next

Create an Ethernet vNIC to be used as the overlay vNIC for the iSCSI device, and then create an iSCSI vNIC.

Deleting an Authentication Profile

Procedure

	Command or Action	Purpose
Step 1	UCS-A# scope org <i>org-name</i>	Enters organization mode for the specified organization. To enter the root organization mode, type <i>/</i> as the <i>org-name</i> .
Step 2	UCS-A /org # delete auth-profile <i>auth-profile-name</i>	Deletes the specified authentication profile.
Step 3	UCS-A /org # commit-buffer	Commits the transaction to the system configuration.

The following example shows how to delete an authentication profile called iscsi-auth and commit the transaction:

```
UCS-A# scope org
UCS-A /org # delete auth-profile iscsi-auth
UCS-A /org* # commit-buffer
UCS-A /org #
```

Adding a Block of IP Addresses to the Initiator Pool

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

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

Procedure

	Command or Action	Purpose
Step 1	UCS-A# scope org <i>org-name</i>	Enters organization mode for the specified organization. To enter the root organization mode, type / as the <i>org-name</i> .
Step 2	UCS-A /org# scope ip-pool iscsi-initiator-pool	Enters the mode to specify an iSCSI initiator pool.
Step 3	UCS-A /org/ip-pool# create block <i>from_ip_address to_ip_address</i> <i>default_gateway subnet_mask</i>	Creates a block of IP addresses for the iSCSI initiator.
Step 4	UCS-A/org/ip-pool/block# show detail expand	(Optional) Shows the block of IP addresses that you have created.
Step 5	UCS-A /org/ip-pool/block # commit-buffer	Commits the transaction to the system configuration.

The following example shows how to create an IP initiator pool for the iSCSI vNIC and commit the transaction:

```
UCS-A # scope org /
UCS-A /org # scope ip-pool iscsi-initiator-pool
UCS-A /org/ip-pool # create block 40.40.40.10 40.40.40.50 40.40.40.1 255.0.0.0
UCS-A /org/ip-pool/block # show detail expand
Block of IP Addresses:
  From: 40.40.40.10
  To: 40.40.40.50
  Default Gateway: 40.40.40.1
  Subnet Mask: 255.0.0.0
UCS-A /org/ip-pool/block # commit buffer
```

What to Do Next

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

Deleting a Block of IP Addresses from the Initiator Pool**Procedure**

	Command or Action	Purpose
Step 1	UCS-A# scope org <i>org-name</i>	Enters organization mode for the specified organization. To enter the root organization mode, type / as the <i>org-name</i> .
Step 2	UCS-A /org# scope ip-pool iscsi-initiator-pool	Enters the mode to specify an iSCSI initiator pool.

	Command or Action	Purpose
Step 3	UCS-A /org/ip-pool# delete block <i>from_ip_address to_ip_address</i>	Deletes the specified block of IP addresses from the initiator pool.
Step 4	UCS-A/org/ip-pool/block# show detail expand	(Optional) Shows that the block of IP addresses has been deleted.
Step 5	UCS-A /org/ip-pool# commit buffer	Commits the transaction to the system configuration.

The following example shows how to delete a block of IP addresses from the initiator pool and commit the transaction:

```
UCS-A # scope org /
UCS-A /org # scope ip-pool iscsi-initiator-pool
UCS-A /org/ip-pool # delete block 40.40.40.10 40.40.40.50 40.40.40.1 255.0.0.0
UCS-A /org/ip-pool # show detail expand

IP Pool:
  Name: iscsi-initiator-pool
  Size: 0
  Assigned: 0
  Descr:
UCS-A /org/ip-pool # commit buffer
```

Creating an iSCSI Boot Policy

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

Procedure

	Command or Action	Purpose
Step 1	UCS-A# scope org <i>org-name</i>	Enters organization mode for the specified organization. To enter the root organization mode, type <i>/</i> as the <i>org-name</i> .
Step 2	UCS-A /org# create boot-policy <i>policy-name</i> [purpose { operational utility }]	Creates a boot policy with the specified policy name, and enters organization boot policy mode. This name can be between 1 and 16 alphanumeric characters. You cannot use spaces or any special characters other than - (hyphen), _ (underscore), : (colon), and . (period), and you cannot change this name after the object has been saved. When you create the boot policy, specify the operational option. This ensures that the server boots from the operating system installed on the server. The utility options is reserved and should only be used if instructed to do so by a Cisco representative.

	Command or Action	Purpose
Step 3	UCS-A /org/boot-policy # set descr <i>description</i>	(Optional) Provides a description for the boot policy. Note If your description includes spaces, special characters, or punctuation, you must begin and end your description with quotation marks. The quotation marks do not appear in the description field of any show command output.
Step 4	UCS-A /org/boot-policy # set enforce-vnic-name {no yes}	(Optional) If you choose yes , Cisco UCS Manager reports whether the device name specified in the boot policy matches what is specified in the service profile. If you choose no , Cisco UCS Manager uses any vNIC, vHBA, or iSCSI device from the service profile and does not report whether the device name specified in the boot policy matches what is specified in the service profile.
Step 5	UCS-A /org/boot-policy # set reboot-on-update {no yes}	Specifies whether the servers using this boot policy are automatically rebooted after you make changes to the boot order. In the Cisco UCS Manager GUI, if the Reboot on Boot Order Change check box is checked for a boot policy, and if CD-ROM or Floppy is the last device in the boot order, deleting or adding the device does not directly affect the boot order and the server does not reboot.
Step 6	UCS-A /org/boot-policy # create iscsi	Adds an iSCSI boot to the boot policy.
Step 7	UCS-A /org/boot-policy/iscsi # create path {primary secondary}	Specifies the primary and secondary paths that Cisco UCS Manager uses to reach the iSCSI target. With iSCSI boot, you set up two paths. Cisco UCS Manager uses the primary path first, and if that fails, then it uses the secondary path.
Step 8	UCS-A /org/boot-policy/iscsi/path # create iscsivnicname <i>iscsi-vnic-name</i>	Creates an iSCSI vNIC.
Step 9	UCS-A /org/boot-policy/iscsi/path # exit	Exits iSCSI path mode.
Step 10	UCS-A /org/boot-policy/iscsi/path # set order <i>order-num</i>	Specifies the order for the iSCSI boot in the boot order.
Step 11	Repeat steps 8-10 to create secondary iSCSI vNICs.	(Optional)
Step 12	UCS-A /org/boot-policy/iscsi # commit-buffer	Commits the transaction to the system configuration.

The following example shows how to create an iSCSI boot policy named `iscsi-boot-policy-LAN`, provide a description for the boot policy, specify that servers using this policy are not automatically rebooted when the boot order is changed, set the boot order for iSCSI boot to 2, create an iSCSI boot and associate it with a vNIC called `iscsienic1`, and commit the transaction:

```
UCS-A# scope org /
UCS-A /org* # create boot-policy iscsi-boot-policy-LAN purpose operational
UCS-A /org/boot-policy* # set descr "Boot policy that boots from iSCSI."
UCS-A /org/boot-policy* # set enforce-vnic-name yes
UCS-A /org/boot-policy* # set reboot-on-update no
UCS-A /org/boot-policy* # create iscsi
UCS-A /org/boot-policy/iscsi* # create path primary
UCS-A /org/boot-policy/iscsi/path* # set iscsivnicname iscsienic1
UCS-A /org/boot-policy/iscsi/path* # exit
UCS-A /org/boot-policy/iscsi* # set order 2
UCS-A /org/boot-policy/iscsi* # commit-buffer
UCS-A /org/boot-policy #
```

What to Do Next

Include the boot policy in a service profile and/or template.

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

Deleting iSCSI Devices from a Boot Policy

Procedure

	Command or Action	Purpose
Step 1	UCS-A# <code>scope org org-name</code>	Enters organization mode for the specified organization. To enter the root organization mode, type <code>/</code> as the <code>org-name</code> .
Step 2	UCS-A /org # <code>scope boot-policy boot-pol-name</code>	Enters boot policy organization mode for the specified boot policy.
Step 3	UCS-A /org/boot-policy # <code>delete iscsi</code>	Deletes the iSCSI boot from the boot policy.
Step 4	UCS-A /org/boot-policy # <code>commit-buffer</code>	Commits the transaction to the system configuration.

The following example shows how to delete an iSCSI boot from the boot policy named `boot-policy-iscsi` and commit the transaction:

```
UCS-A# scope org /
UCS-A /org # scope boot-policy boot-policy-iscsi
UCS-A /org/boot-policy # delete iscsi
UCS-A /org/boot-policy* # commit-buffer
UCS-A /org/boot-policy #
```

Creating an iSCSI vNIC in a Service Profile

You can create an iSCSI vNIC in a service profile.

Before You Begin

You must have an Ethernet vNIC in a service policy to be used as the overlay vNIC for the iSCSI device.

Procedure

	Command or Action	Purpose
Step 1	UCS-A # scope org <i>org-name</i>	Enters organization mode for the specified organization. To enter the root organization mode, type / as the <i>org-name</i> .
Step 2	UCS-A /org # scope service-profile <i>profile-name</i>	Enters service profile organization mode for the service profile.
Step 3	UCS-A /org/service-profile # create vnic-iscsi <i>iscsi-vnic-name</i> .	Specifies the iSCSI vNIC name.
Step 4	UCS-A /org/service-profile/vnic-iscsi* # set iscsi-adaptor-policy <i>iscsi-adaptor-name</i>	(Optional) Specifies the iSCSI adaptor policy that you have created for this iSCSI vNIC.
Step 5	UCS-A /org/service-profile/vnic-iscsi* # set auth-name <i>authentication-profile-name</i>	(Optional) Sets the authentication profile to be used by the iSCSI vNIC. The authentication profile must already exist for it to be set. For more information, see Creating an Authentication Profile , on page 424.
Step 6	UCS-A /org/service-profile/vnic-iscsi* # set identity { dynamic-mac { <i>dynamic-mac-address</i> derived } mac-pool <i>mac-pool-name</i> }	Specifies the MAC address for the iSCSI vNIC. Note The MAC address is only set for Cisco UCS NIC M51KR-B adapters.
Step 7	UCS-A /org/service-profile/vnic-iscsi* # set iscsi-identity { initiator-name <i>initiator-name</i> initiator-pool-name <i>iqn-pool-name</i> }	Specifies the name of the iSCSI initiator or the name of an IQN pool from which the iSCSI initiator name will be provided. The iSCSI initiator name can be up to 223 characters.
Step 8	UCS-A /org/service-profile/vnic-iscsi* # set overlay-vnic-name <i>overlay-vnic-name</i>	Specifies the Ethernet vNIC that is used by the iSCSI device as the overlay vNIC. For more information, see Configuring a vNIC for a Service Profile , on page 474.
Step 9	UCS-A /org/service-profile/vnic-iscsi* # create eth-if	Creates an Ethernet interface for a VLAN assigned to the iSCSI vNIC.
Step 10	UCS-A /org/service-profile/vnic-iscsi/eth-if* # set vlnname <i>vlan-name</i> .	Specifies the VLAN name. The default VLAN is default. For the Cisco UCS M81KR Virtual Interface Card and the Cisco UCS M82-8P Virtual Interface Card, the VLAN that you specify must be the same as the native VLAN on the overlay vNIC. For the Cisco UCS M51KR-B Broadcom BCM57711 adapter, the VLAN that you specify can be any VLAN assigned to the overlay vNIC.

	Command or Action	Purpose
Step 11	UCS-A /org/service-profile/vnic-iscsi # commit-buffer	Commits the transaction to the system configuration.

The following example shows how to create an iSCSI vNIC called scsivnic1, add it to an existing service profile called accounting, and commit the transaction:

```
UCS-A# scope org /
UCS-A /org # scope service-profile accounting
UCS-A /org/service-profile # create vnic-iscsi iSCSI1
UCS-A /org/service-profile/vnic-iscsi* # set iscsi-adaptor-policy iscsiboot
UCS-A /org/service-profile/vnic-iscsi* # set auth-name initauth
UCS-A /org/service-profile/vnic-iscsi* # set identity dynamic-mac derived
UCS-A /org/service-profile/vnic-iscsi* # set iscsi-identity initiator-name iSCSI1
UCS-A /org/service-profile/vnic-iscsi* # set overlay-vnic-name eth1
UCS-A /org/service-profile/vnic-iscsi* # create eth-if
UCS-A /org/service-profile/vnic-iscsi/eth-if* # set vllanname default
UCS-A /org/service-profile/vnic-iscsi/eth-if* # commit buffer
```

What to Do Next

Configure an iSCSI initiator to boot using a static IP address, an IP address from a configured IP pool, or DHCP.

Deleting an iSCSI vNIC from a Service Profile

Procedure

	Command or Action	Purpose
Step 1	UCS-A# scope org org-name	Enters organization mode for the specified organization. To enter the root organization mode, type / as the <i>org-name</i> .
Step 2	UCS-A /org # scope service-profile profile-name	Enters service profile organization mode for the service profile.
Step 3	UCS-A /org/service-profile # delete vnic-iscsi iscsi-vnic-name	Deletes the specified iSCSI vNIC from the specified service profile.
Step 4	UCS-A /org/service-profile # commit-buffer	Commits the transaction to the system configuration.

The following example shows how to delete an iSCSI vNIC called scsivnic1 and commit the transaction:

```
UCS-A# scope org /
UCS-A /org # scope service-profile accounting
UCS-A /org/service-profile # delete vnic-iscsi scsivnic1
UCS-A /org/service-profile* # commit-buffer
UCS-A /org/service-profile #
```

Creating an iSCSI Initiator that Boots Using a Static IP Address

In a service profile, you can create an iSCSI initiator and configure it to boot using a static IP address.

Before You Begin

You have completed the following:

- Created iSCSI overlay vNICs in a service profile.
- Created an iSCSI vNIC in a service profile.

Procedure

	Command or Action	Purpose
Step 1	UCS-A # scope org <i>org-name</i>	Enters organization mode for the specified organization. To enter the root organization mode, type / as the <i>org-name</i> .
Step 2	UCS-A /org # scope service-profile <i>profile-name</i>	Enters service profile organization mode for the service profile.
Step 3	UCS-A /org/service-profile # scope vnic-iscsi <i>iscsi-vnic-name</i>	Enters the configuration mode for the specified iSCSI vNIC.
Step 4	UCS-A /org/service-profile/vnic-iscsi* # scope eth-if	Enters the configuration mode for the iSCSI Ethernet interface.
Step 5	UCS-A /org/service-profile/vnic-iscsi/eth-if* # create ip-if	Creates an IP interface.
Step 6	UCS-A /org/service-profile/vnic-iscsi/eth-if/ip-if* # enter static-ip-params	Specifies that you are entering static IP boot parameters.
Step 7	UCS-A /org/service-profile/vnic-iscsi/eth-if/ip-if/static-ip-params* # set addr <i>ip-address</i>	Specifies the static IP address.
Step 8	UCS-A /org/service-profile/vnic-iscsi/eth-if/ip-if/static-ip-params* # set default-gw <i>ip-address</i>	Specifies the default gateway IP address.
Step 9	UCS-A /org/service-profile/vnic-iscsi/eth-if/ip-if/static-ip-params* # set primary-dns <i>ip-address</i>	Specifies the primary DNS IP address.
Step 10	UCS-A /org/service-profile/vnic-iscsi/eth-if/ip-if/static-ip-params* # set secondary-dns <i>ip-address</i>	Specifies the secondary DNS IP address.
Step 11	UCS-A /org/service-profile/vnic-iscsi/eth-if/ip-if/static-ip-params* # set subnet <i>subnet-ip-address</i>	Specifies the subnet mask.

	Command or Action	Purpose
Step 12	UCS-A /org/service-profile/vnic-iscsi/eth-if/ip-if/static-ip-params* # commit buffer	Commits the transaction to the system configuration.

The following example shows how to configure the initiator to boot using a static IP address and commit the transaction:

```
UCS-A # scope org
UCS-A /org # scope service-profile accounting
UCS-A /org/service-profile # scope vnic-iscsi iSCSI1
UCS-A /org/service-profile/vnic-iscsi* # scope eth-if
UCS-A /org/service-profile/vnic-iscsi/eth-if* # create ip-if
UCS-A /org/service-profile/vnic-iscsi/eth-if/ip-if* # enter static-ip-params
UCS-A /org/service-profile/vnic-iscsi/eth-if/ip-if/static-ip-params* # set addr 10.104.105.193

UCS-A /org/service-profile/vnic-iscsi/eth-if/ip-if/static-ip-params* # set default-gw
10.104.105.1
UCS-A /org/service-profile/vnic-iscsi/eth-if/ip-if/static-ip-params* # set primary-dns
11.11.11.100
UCS-A /org/service-profile/vnic-iscsi/eth-if/ip-if/static-ip-params* # set secondary-dns
11.11.11.100
UCS-A /org/service-profile/vnic-iscsi/eth-if/ip-if/static-ip-params* # set subnet
255.255.255.0
UCS-A /org/service-profile/vnic-iscsi/eth-if/ip-if/static-ip-params* # commit-buffer
```

What to Do Next

Create an iSCSI target.

Deleting the Static IP Address Boot Parameters from an iSCSI Initiator

In a service profile, you can delete the static IP address boot parameters from an iSCSI initiator.

Procedure

	Command or Action	Purpose
Step 1	UCS-A # scope org <i>org-name</i>	Enters organization mode for the specified organization. To enter the root organization mode, type / as the <i>org-name</i> .
Step 2	UCS-A /org # scope service-profile <i>profile-name</i>	Enters service profile organization mode for the service profile.
Step 3	UCS-A /org/service-profile # scope vnic-iscsi <i>iscsi-vnic-name</i>	Enters the configuration mode for the specified iSCSI vNIC.
Step 4	UCS-A /org/service-profile/vnic-iscsi* # scope eth-if	Enters the configuration mode for the iSCSI Ethernet interface.

	Command or Action	Purpose
Step 5	UCS-A /org/service-profile/vnic-iscsi/eth-if* # enter ip-if	Enters the configuration mode for an IP interface.
Step 6	UCS-A /org/service-profile/vnic-iscsi/eth-if/ip-if* # delete static-ip-params	Deletes the static IP boot parameters from an initiator.
Step 7	UCS-A /org/service-profile/vnic-iscsi/eth-if/ip-if/static-ip-params* # commit buffer	Commits the transaction to the system configuration.

The following example shows how to delete the static IP address boot parameters from the initiator and commit the transaction:

```
UCS-A # scope org
UCS-A /org # scope service-profile accounting
UCS-A /org/service-profile # scope vnic-iscsi iSCSI1
UCS-A /org/service-profile/vnic-iscsi* # scope eth-if
UCS-A /org/service-profile/vnic-iscsi/eth-if* # enter ip-if
UCS-A /org/service-profile/vnic-iscsi/eth-if/ip-if* # delete static-ip-params
UCS-A /org/service-profile/vnic-iscsi/eth-if/ip-if/static-ip-params* # commit-buffer
```

Creating an iSCSI Initiator that Boots Using an IP Address from an IP Pool

In a service profile, you can create an iSCSI initiator and configure it to boot using an IP address from an IP pool that you have created.

Before You Begin

You have completed the following:

- Created an overlay vNIC in a service profile
- Created an iSCSI vNIC in a service profile.

Procedure

	Command or Action	Purpose
Step 1	UCS-A # scope org <i>org-name</i>	Enters organization mode for the specified organization. To enter the root organization mode, type / as the <i>org-name</i> .
Step 2	UCS-A /org # scope service-profile <i>profile-name</i>	Enters service profile organization mode for the service profile.
Step 3	UCS-A /org/service-profile # scope vnic-iscsi <i>iscsi-vnic-name</i>	Enters the configuration mode for the specified iSCSI vNIC.

	Command or Action	Purpose
Step 4	UCS-A /org/service-profile/vnic-iscsi* # scope eth-if	Enters the configuration mode for the iSCSI Ethernet interface.
Step 5	UCS-A /org/service-profile/vnic-iscsi/eth-if* # create ip-if	Creates an IP interface.
Step 6	UCS-A /org/service-profile/vnic-iscsi/eth-if/ip-if* # enter pooled-ip-params	Specifies that the iSCSI initiator boot using one of the IP addresses from the previously created iSCSI initiator IP pool.
Step 7	UCS-A /org/service-profile/vnic-iscsi/eth-if/ip-if/pooled-ip-params* # commit buffer	Commits the transaction to the system configuration.

The following example shows how to create an iSCSI initiator and configure it to boot using an IP address from an IP pool:

```
UCS-A # scope org
UCS-A /org # scope service-profile accounting
UCS-A /org/service-profile # scope vnic-iscsi iSCSI1
UCS-A /org/service-profile/vnic-iscsi* # scope eth-if
UCS-A /org/service-profile/vnic-iscsi/eth-if* # create ip-if
UCS-A /org/service-profile/vnic-iscsi/eth-if/ip-if* # enter pooled-ip-params
UCS-A /org/service-profile/vnic-iscsi/eth-if/ip-if/pooled-ip-params* # commit buffer
```

What to Do Next

Create an iSCSI target.

Deleting the IP Pool Boot Parameter from an iSCSI Initiator

In a service profile, you can create an iSCSI initiator and configure it to boot using an IP address from an IP pool that you have created.

Procedure

	Command or Action	Purpose
Step 1	UCS-A # scope org <i>org-name</i>	Enters organization mode for the specified organization. To enter the root organization mode, type / as the <i>org-name</i> .
Step 2	UCS-A /org # scope service-profile <i>profile-name</i>	Enters service profile organization mode for the service profile.
Step 3	UCS-A /org/service-profile # scope vnic-iscsi <i>iscsi-vnic-name</i>	Enters the configuration mode for the specified iSCSI vNIC.

	Command or Action	Purpose
Step 4	UCS-A /org/service-profile/vnic-iscsi* # scope eth-if	Enters the configuration mode for the iSCSI Ethernet interface.
Step 5	UCS-A /org/service-profile/vnic-iscsi/eth-if* # enter ip-if	Enters the configuration mode for an IP interface.
Step 6	UCS-A /org/service-profile/vnic-iscsi/eth-if/ip-if* # delete pooled-ip-params	Specifies that the iSCSI initiator does not use an IP address from an IP pool to boot.
Step 7	UCS-A /org/service-profile/vnic-iscsi/eth-if/ip-if/pooled-ip-params* # commit buffer	Commits the transaction to the system configuration.

The following example shows how to delete the boot using an IP address from an IP pool parameter and commit the transaction:

```
UCS-A # scope org
UCS-A /org # scope service-profile accounting
UCS-A /org/service-profile # scope vnic-iscsi iSCSI1
UCS-A /org/service-profile/vnic-iscsi* # scope eth-if
UCS-A /org/service-profile/vnic-iscsi/eth-if* # enter ip-if
UCS-A /org/service-profile/vnic-iscsi/eth-if/ip-if* # delete pooled-ip-params
UCS-A /org/service-profile/vnic-iscsi/eth-if/ip-if/pooled-ip-params* # commit buffer
```

Creating an iSCSI Initiator that Boots Using DHCP

In a service profile, you can create an iSCSI initiator and configure it to boot using DHCP.

Before You Begin

You have completed the following:

- Created iSCSI overlay vNICs in a service profile.
- Created an iSCSI vNIC in a service profile.

Procedure

	Command or Action	Purpose
Step 1	UCS-A # scope org <i>org-name</i>	Enters organization mode for the specified organization. To enter the root organization mode, type / as the <i>org-name</i> .
Step 2	UCS-A /org # scope service-profile <i>profile-name</i>	Enters service profile organization mode for the service profile.

	Command or Action	Purpose
Step 3	UCS-A /org/service-profile # scope vnic-iscsi <i>iscsi-vnic-name</i>	Enters the configuration mode for the specified iSCSI vNIC.
Step 4	UCS-A /org/service-profile/vnic-iscsi* # scope eth-if	Enters the configuration mode for the iSCSI Ethernet interface.
Step 5	UCS-A /org/service-profile/vnic-iscsi/eth-if* # create ip-if	Creates an IP interface.
Step 6	UCS-A /org/service-profile/vnic-iscsi/eth-if/ip-if* # create dhcp-ip-params	Specifies that you are setting the initiator to boot using DHCP.
Step 7	UCS-A /org/service-profile/vnic-iscsi/eth-if/ip-if/dhcp-ip-params* # commit buffer	Commits the transaction to the system configuration.

The following example shows how to configure the initiator to boot using DHCP and commit the transaction:

```
UCS-A # scope org
UCS-A /org # scope service-profile accounting
UCS-A /org/service-profile # scope vnic-iscsi iSCSI1
UCS-A /org/service-profile/vnic-iscsi* # scope eth-if
UCS-A /org/service-profile/vnic-iscsi/eth-if* # create ip-if
UCS-A /org/service-profile/vnic-iscsi/eth-if/ip-if* # create dhcp-ip-params
UCS-A /org/service-profile/vnic-iscsi/eth-if/ip-if/dhcp-ip-params* # commit-buffer
```

What to Do Next

Create an iSCSI target.

Deleting the DHCP Boot Parameter from an iSCSI Initiator

In a service profile, you can remove the DHCP boot parameter from an iSCSI initiator.

Procedure

	Command or Action	Purpose
Step 1	UCS-A # scope org <i>org-name</i>	Enters organization mode for the specified organization. To enter the root organization mode, type / as the <i>org-name</i> .
Step 2	UCS-A /org # scope service-profile <i>profile-name</i>	Enters service profile organization mode for the service profile.
Step 3	UCS-A /org/service-profile # scope vnic-iscsi <i>iscsi-vnic-name</i>	Enters the configuration mode for the specified iSCSI vNIC.
Step 4	UCS-A /org/service-profile/vnic-iscsi* # scope eth-if	Enters the configuration mode for the iSCSI Ethernet interface.

	Command or Action	Purpose
Step 5	UCS-A /org/service-profile/vnic-iscsi/eth-if* # enter ip-if	Enters the configuration mode for an IP interface.
Step 6	UCS-A /org/service-profile/vnic-iscsi/eth-if/ip-if* # delete dhcp-ip-params	Specifies that the initiator does not use DHCP to boot.
Step 7	UCS-A /org/service-profile/vnic-iscsi/eth-if/ip-if/dhcp-ip-params* # commit buffer	Commits the transaction to the system configuration.

The following example shows how to delete the boot using DHCP parameter and commit the transaction:

```
UCS-A # scope org
UCS-A /org # scope service-profile accounting
UCS-A /org/service-profile # scope vnic-iscsi iSCSI1
UCS-A /org/service-profile/vnic-iscsi* # scope eth-if
UCS-A /org/service-profile/vnic-iscsi/eth-if* # enter ip-if
UCS-A /org/service-profile/vnic-iscsi/eth-if/ip-if* # delete dhcp-ip-params
UCS-A /org/service-profile/vnic-iscsi/eth-if/ip-if/dhcp-ip-params* # commit-buffer
```

IQN Pools

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

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

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

Creating an IQN Pool



Note

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

Procedure

	Command or Action	Purpose
Step 1	UCS-A# scope org org-name	Enters organization mode for the specified organization. To enter the root organization mode, type <i>/</i> as the <i>org-name</i> .
Step 2	UCS-A /org # create iqn-pool pool-name	Creates an IQN pool with the specified pool name and enters organization IQN pool mode. The pool name can be up to 32 characters.

	Command or Action	Purpose
Step 3	UCS-A /org/iqn-pool # set iqn-prefix <i>prefix</i>	Specifies the prefix for the IQN block members. Unless limited by the adapter card, the prefix can contain up to 150 characters.
Step 4	UCS-A /org/iqn-pool # set descr <i>description</i>	(Optional) Provides a description for the IQN pool. Enter up to 256 characters. Note If your description includes spaces, special characters, or punctuation, you must begin and end your description with quotation marks. The quotation marks will not appear in the description field of any show command output.
Step 5	UCS-A /org/iqn-pool # create block <i>suffix from to</i>	Creates a block (range) of IQNs, and enters organization IQN pool block mode. You must specify the base suffix, the starting suffix number, and the ending suffix number. The resulting IQN pool members are of the form <i>prefix:suffix:number</i> . The suffix can be up to 64 characters. Note An IQN pool can contain more than one IQN block. To create multiple blocks, enter multiple create block commands from organization IQN pool mode.
Step 6	UCS-A /org/iqn-pool/block # commit-buffer	Commits the transaction to the system configuration.

The following example creates an IQN pool named pool4, provides a description for the pool, specifies a prefix and a block of suffixes to be used for the pool, and commits the transaction:

```
UCS-A# scope org /
UCS-A /org # create iqn-pool pool4
UCS-A /org/iqn-pool* # set iqn-prefix iqn.alpha.com
UCS-A /org/iqn-pool* # set descr "This is IQN pool 4"
UCS-A /org/iqn-pool* # create block beta 3 5
UCS-A /org/iqn-pool/block* # commit-buffer
UCS-A /org/iqn-pool/block # exit
UCS-A /org/iqn-pool # show detail
IQN Pool:
  Name: pool4
  IQN Prefix: iqn.alpha.com
  Size: 3
  Assigned: 0
  Description: This is IQN pool 4

UCS-A /org/iqn-pool #
```

This example creates a range of IQNs from **iqn.alpha.com:beta:3** to **iqn.alpha.com:beta:5**.

What to Do Next

- Add one or more blocks of IQNs to the IQN suffix pool.
- Include the IQN suffix pool in a service profile and/or template.

Adding a Block to an IQN Pool

Procedure

	Command or Action	Purpose
Step 1	UCS-A# scope org <i>org-name</i>	Enters organization mode for the specified organization. To enter the root organization mode, type / as the <i>org-name</i> .
Step 2	UCS-A /org # scope iqn-pool <i>pool-name</i>	Enters organization IQN pool mode for the specified pool.
Step 3	UCS-A /org/iqn-pool # create block <i>suffix from to</i>	Creates a block (range) of IQN suffixes, and enters organization IQN pool block mode. You must specify the base suffix, the starting suffix number, and the ending suffix number. The resulting IQN pool members are of the form <i>prefix:suffix:number</i> . Note An IQN pool can contain more than one IQN block. To create multiple blocks, enter multiple create block commands from organization IQN pool mode.
Step 4	UCS-A /org/iqn-pool/block # commit-buffer	Commits the transaction to the system configuration.
Step 5	UCS-A /org/iqn-pool/block # exit	(Optional) Returns to organization IQN pool mode.
Step 6	UCS-A /org/iqn-pool # show block	(Optional) Displays the blocks of suffixes.

This example creates a block of IQN suffixes in an IQN pool named pool4 and commits the transaction:

```
UCS-A# scope org /
UCS-A /org # scope iqn-pool pool4
UCS-A /org/iqn-pool # create block beta 3 5
UCS-A /org/iqn-pool/block* # commit-buffer
UCS-A /org/iqn-pool/block # exit
UCS-A /org/iqn-pool # show block
Block of IQN Names:
  Suffix      From  To
  -----
  beta                3    5

UCS-A /org/iqn-pool #
```

This example creates a range of IQN suffixes from **beta:3** to **beta:5**.

Deleting a Block from an IQN Pool

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

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

Procedure

	Command or Action	Purpose
Step 1	UCS-A# scope org <i>org-name</i>	Enters organization mode for the specified organization. To enter the root organization mode, type / as the <i>org-name</i> .
Step 2	UCS-A /org # scope iqn-pool <i>pool-name</i>	Enters organization IQN pool mode for the specified pool.
Step 3	UCS-A /org/iqn-pool # delete block <i>suffix from to</i>	Deletes a block (range) of IQNs. You must specify the base suffix and the first and last numbers in the block to be deleted.

This example deletes a block of suffixes from an IQN pool named pool4 and commits the transaction:

```
UCS-A# scope org /
UCS-A /org # scope iqn-pool pool4
UCS-A /org/iqn-pool # delete block beta 0 12
UCS-A /org/iqn-pool* # commit-buffer
UCS-A /org/iqn-pool #
```

Deleting an IQN Pool

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

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

Procedure

	Command or Action	Purpose
Step 1	UCS-A# scope org <i>org-name</i>	Enters organization mode for the specified organization. To enter the root organization mode, type / as the <i>org-name</i> .
Step 2	UCS-A /org # delete iqn-pool <i>pool-name</i>	Deletes the specified IQN pool.

	Command or Action	Purpose
Step 3	UCS-A /org # commit-buffer	Commits the transaction to the system configuration.

The following example deletes the IQN pool named pool4 and commits the transaction:

```
UCS-A# scope org /
UCS-A /org # delete iqn-pool pool4
UCS-A /org* # commit-buffer
UCS-A /org #
```

Viewing IQN Pool Usage

Procedure

	Command or Action	Purpose
Step 1	UCS-A# scope org <i>org-name</i>	Enters organization mode for the specified organization. To enter the root organization mode, type / as the <i>org-name</i> .
Step 2	UCS-A /org # scope iqn-pool <i>pool-name</i>	Enters organization IQN pool mode for the specified pool.
Step 3	UCS-A /org/iqn-pool # show pooled	Displays the assignments of the IQN block members.

This example displays the assignments of suffixes in the IQN pool named pool4:

```
UCS-A# scope org /
UCS-A /org # scope iqn-pool pool4
UCS-A /org/iqn-pool # show pooled
Pooled:
  Name           Assigned Assigned To Dn
  -----
  beta:3         No
  beta:4         No
  beta:5         No

UCS-A /org/iqn-pool #
```

Creating an iSCSI Static Target

You can create a static target.

Before You Begin

You have already created an iSCSI vNIC.

Procedure

	Command or Action	Purpose
Step 1	UCS-A# scope org <i>org-name</i>	Enters organization mode for the specified organization. To enter the root organization mode, type / as the <i>org-name</i> .
Step 2	UCS-A /org # scope service-profile <i>profile-name</i>	Enters service profile organization mode for the service profile to which you want to add an iSCSI target.
Step 3	UCS-A /org/service-profile # scope vnic-iscsi <i>iscsi-vnic-name</i>	Enters the iSCSI vNIC mode for the specified vNIC name.
Step 4	UCS-A /org/service-profile/vnic-iscsi # scope eth-if	Enters the Ethernet interface mode.
Step 5	UCS-A /org/service-profile/vnic-iscsi/eth-if # create static-target-if {1 2}	Creates a static target for the iSCSI vNIC and assigns a priority level to it. Valid priority levels are 1 or 2.
Step 6	UCS-A /org/service-profile/vnic-iscsi/eth-if/static-target-if # set name <i>name</i>	A regular expression that defines the iSCSI Qualified Name (IQN) or Extended Unique Identifier (EUI) name of the iSCSI target. You can enter any alphanumeric characters as well as the following special characters: <ul style="list-style-type: none"> • . (period) • : (colon) • - (dash) <p>Important This name must be properly formatted according to standard IQN or EUI guidelines. The following examples show properly formatted iSCSI target names:</p> <ul style="list-style-type: none"> • iqn.2001-04.com.example • iqn.2001-04.com.example:storage:diskarrays-sn-a86 • iqn.2001-04.com.example:storage.tape1.sys1.x • iqn.2001-04.com.example:storage.disk2.sys1.x • eui.02004567A425678D
Step 7	UCS-A /org/service-profile/vnic-iscsi/eth-if/static-target-if # set port <i>port-num</i>	The port associated with the iSCSI target. Enter an integer between 1 and 65535. The default is 3260.
Step 8	UCS-A /org/service-profile/vnic-iscsi/eth-if/static-target-if # set auth-name <i>auth-profile</i>	(Optional) If you need the target to authenticate itself and have an authentication profile, you need to specify the name of the authentication profile. The name of the associated authentication profile.

	Command or Action	Purpose
Step 9	UCS-A /org/service-profile/vnic-iscsi/eth-if/static-target-if # set ipaddress <i>ipv4-address</i>	The IPv4 address assigned to the iSCSI target.
Step 10	UCS-A /org/service-profile/vnic-iscsi/eth-if/static-target-if # create lun	Creates the LUN that corresponds to the location of the interface.
Step 11	UCS-A /org/service-profile/vnic-iscsi/eth-if/static-target-if/lun* # set id <i>id-number</i>	Specifies the target LUN id. Valid values are from 0 to 65535.
Step 12	UCS-A /org/service-profile/vnic-iscsi/eth-if/static-target-if/lun* # exit	Exits the current configuration mode.
Step 13	UCS-A /org/service-profile/vnic-iscsi/eth-if/static-target-if # exit	Exits the current configuration mode.
Step 14	UCS-A /org/service-profile/vnic-iscsi/eth-if # commit-buffer	Commits the transaction to the system configuration.
Step 15	Repeat steps 5 through 14 to create a second static target.	(Optional)

The following example shows how to create two iSCSI static target interfaces and commit the transaction:

```
UCS-A # scope org test
UCS-A /org # scope service-profile accounting
UCS-A /org/service-profile # scope vnic-iscsi iSCSI1
UCS-A /org/service-profile/vnic-iscsi # scope eth-if
UCS-A /org/service-profile/vnic-iscsi/eth-if # create static-target-if 1
UCS-A /org/service-profile/vnic-iscsi/eth-if/static-target-if* # set name statictarget1
UCS-A /org/service-profile/vnic-iscsi/eth-if/static-target-if* # set port 3260
UCS-A /org/service-profile/vnic-iscsi/eth-if/static-target-if* # set auth-name authprofile1
UCS-A /org/service-profile/vnic-iscsi/eth-if/static-target-if* # set ip-address 192.168.10.10
UCS-A /org/service-profile/vnic-iscsi/eth-if/static-target-if* # create lun
UCS-A /org/service-profile/vnic-iscsi/eth-if/static-target-if/lun* # set id 1
UCS-A /org/service-profile/vnic-iscsi/eth-if/static-target-if/lun* # exit
UCS-A /org/service-profile/vnic-iscsi/eth-if/static-target-if* # exit
UCS-A /org/service-profile/vnic-iscsi/eth-if* # commit-buffer
UCS-A /org/service-profile/vnic-iscsi/eth-if # create static-target-if 2
UCS-A /org/service-profile/vnic-iscsi/eth-if/static-target-if* # set ipaddress 192.168.10.11

UCS-A /org/service-profile/vnic-iscsi/eth-if/static-target-if* # set name statictarget2
UCS-A /org/service-profile/vnic-iscsi/eth-if/static-target-if* # set port 3260
UCS-A /org/service-profile/vnic-iscsi/eth-if/static-target-if* # set auth-name authprofile1
UCS-A /org/service-profile/vnic-iscsi/eth-if/static-target-if* # create lun
UCS-A /org/service-profile/vnic-iscsi/eth-if/static-target-if/lun* # set id 1
UCS-A /org/service-profile/vnic-iscsi/eth-if/static-target-if/lun* # exit
UCS-A /org/service-profile/vnic-iscsi/eth-if/static-target-if* # exit
UCS-A /org/service-profile/vnic-iscsi/eth-if* # commit-buffer
```

What to Do Next

To configure a second iSCSI device, repeat the steps for creating an iSCSI vNIC, initiator, and target.

Deleting an iSCSI Static Target

You can delete an iSCSI static target. However, you must have at least one iSCSI static target remaining after you delete one. Therefore, you must have two iSCSI static targets in order to delete one of them.



Note

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

Procedure

	Command or Action	Purpose
Step 1	UCS-A# scope org <i>org-name</i>	Enters organization mode for the specified organization. To enter the root organization mode, type <i>/</i> as the <i>org-name</i> .
Step 2	UCS-A /org # scope service-profile <i>profile-name</i>	Enters service profile organization mode for the service profile to which you want to add an iSCSI target.
Step 3	UCS-A /org/service-profile # scope vnic-iscsi <i>iscsi-vnic-name</i>	Enters the iSCSI vNIC mode for the specified vNIC name.
Step 4	UCS-A /org/service-profile/vnic-iscsi # scope eth-if	Enters the Ethernet interface mode.
Step 5	UCS-A /org/service-profile/vnic-iscsi/eth-if # delete static-target-if	Deletes the static target for the iSCSI vNIC.
Step 6	UCS-A /org/service-profile/vnic-iscsi/eth-if # commit-buffer	Commits the transaction to the system configuration.

The following example shows how to delete an iSCSI static target and commit the transaction:

```
UCS-A # scope org test
UCS-A /org # scope service-profile sample
UCS-A /org/service-profile # scope vnic-iscsi trial
UCS-A /org/service-profile/vnic-iscsi # scope eth-if
UCS-A /org/service-profile/vnic-iscsi/eth-if # delete static-target-if 1
UCS-A /org/service-profile/vnic-iscsi/eth-if* # commit-buffer
UCS-A /org/service-profile/vnic-iscsi/eth-if #
```

Creating an iSCSI Auto Target

You can create an iSCSI auto target with or without the vendor IDs.

Before You Begin

These prerequisites must be met before creating iSCSI auto target:

- You have already created an iSCSI vNIC in a service profile.

- You have considered the prerequisites for the VIC that you are using. For more information, see [iSCSI Boot Guidelines and Prerequisites](#), on page 418

Procedure

	Command or Action	Purpose
Step 1	UCS-A# scope org <i>org-name</i>	Enters organization mode for the specified organization. To enter the root organization mode, type / as the <i>org-name</i> .
Step 2	UCS-A /org # scope service-profile <i>profile-name</i>	Enters service profile organization mode for the service profile that you want to add an iSCSI target interface to.
Step 3	UCS-A /org/service-profile # scope vnic-iscsi <i>iscsi-vnic-name</i>	Enters iSCSI vNIC service profile organization mode for the specified vNIC name.
Step 4	UCS-A /org/service-profile/vnic-iscsi # scope eth-if	Enters Ethernet interface iSCSI vNIC service profile organization mode.
Step 5	UCS-A /org/service-profile/vnic-iscsi/eth-if # create auto-target-if	Creates an auto target for the iSCSI vNIC. If you plan to use an auto target without the vendor ID, you must configure an initiator name. For more information, see Creating an iSCSI vNIC in a Service Profile , on page 429.
Step 6	UCS-A /org/service-profile/vnic-iscsi/eth-if/auto-target-if* # set dhcp-vendor-id <i>vendor-id</i>	(Optional) Sets a vendor ID for the auto target. The vendor ID can be up to 32 alphanumeric characters.
Step 7	UCS-A /org/service-profile/vnic-iscsi/eth-if/auto-target-if* # exit	Exists the current configuration mode.
Step 8	UCS-A /org/service-profile/vnic-iscsi/eth-if # commit-buffer	Commits the transaction to the system configuration.

The following example shows how to create an iSCSI auto target *without* a vendor ID and commit the transaction:

```
UCS-A # scope org
UCS-A /org # scope service-profile accounting
UCS-A /org/service-profile # scope vnic-iscsi iSCSI1
UCS-A /org/service-profile/vnic-iscsi # scope eth-if
UCS-A /org/service-profile/vnic-iscsi/eth-if # create auto-target-if
UCS-A /org/service-profile/vnic-iscsi/eth-if/auto-target-if* # exit
UCS-A /org/service-profile/vnic-iscsi/eth-if* # commit-buffer
```

The following example shows how to create an iSCSI auto target *with* a vendor ID and commit the transaction:

```
UCS-A # scope org
```



```

UCS-A /org # scope service-profile accounting
UCS-A /org/service-profile # scope vnic-iscsi iSCSI1
UCS-A /org/service-profile/vnic-iscsi # scope eth-if
UCS-A /org/service-profile/vnic-iscsi/eth-if # create auto-target-if
UCS-A /org/service-profile/vnic-iscsi/eth-if/auto-target-if* # set dhcp-vendor-id iSCSI_Vendor
UCS-A /org/service-profile/vnic-iscsi/eth-if/auto-target-if* # exit
UCS-A /org/service-profile/vnic-iscsi/eth-if* # commit-buffer

```

What to Do Next

To configure a second iSCSI device, repeat the steps for creating an iSCSI vNIC, initiator, and target.

Deleting an iSCSI Auto Target

You can delete an auto target only if you have a static target set.

Procedure

	Command or Action	Purpose
Step 1	UCS-A# scope org <i>org-name</i>	Enters organization mode for the specified organization. To enter the root organization mode, type / as the <i>org-name</i> .
Step 2	UCS-A /org # scope service-profile <i>profile-name</i>	Enters the service profile mode for the service profile to which you want to add an iSCSI target.
Step 3	UCS-A /org/service-profile # scope vnic-iscsi <i>iscsi-vnic-name</i>	Enters the iSCSI vNIC mode for the specified vNIC name.
Step 4	UCS-A /org/service-profile/vnic-iscsi # scope eth-if	Enters the Ethernet interface mode.
Step 5	UCS-A /org/service-profile/vnic-iscsi/eth-if # delete auto-target-if	Deletes the auto target.
Step 6	UCS-A /org/service-profile/vnic-iscsi/eth-if # commit-buffer	Commits the transaction to the system configuration.

The following example shows how to delete an iSCSI auto target and commit the transaction:

```

UCS-A # scope org test
UCS-A /org # scope service-profile accounting
UCS-A /org/service-profile # scope vnic-iscsi iSCSI1
UCS-A /org/service-profile/vnic-iscsi # scope eth-if
UCS-A /org/service-profile/vnic-iscsi/eth-if # delete auto-target-if
UCS-A /org/service-profile/vnic-iscsi/eth-if* # commit-buffer

```

Verifying iSCSI Boot

Use the KVM console to view the boot up messages as the adapter is booting. For information on how to access the KVM console, see the *Starting the KVM Console* chapter.

This step can only be performed using the Cisco UCS Manager GUI. For more information, see the *Starting the KVM Console* chapter in the *UCS Manager GUI Configuration Guide*.

- For the Cisco UCS M51KR-B Broadcom BCM57711, the following message appears:
Logging in the 1st iSCSI Target.... Succeeded.
- For the Cisco UCS M81KR Virtual Interface Card, the following message appears:
Option ROM installed successfully.

LAN Boot

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

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

Configuring a LAN Boot for a Boot Policy

Before You Begin

Create a boot policy to contain the LAN boot configuration.

Procedure

	Command or Action	Purpose
Step 1	UCS-A# scope org <i>org-name</i>	Enters organization mode for the specified organization. To enter the root organization mode, type / as the <i>org-name</i> .
Step 2	UCS-A /org # scope boot-policy <i>policy-name</i>	Enters organization boot policy mode for the specified boot policy.
Step 3	UCS-A /org/boot-policy # create lan	Creates a LAN boot for the boot policy and enters organization boot policy LAN mode.
Step 4	UCS-A /org/boot-policy/lan # set order {1 2 3 4}	Specifies the boot order for the LAN boot.
Step 5	UCS-A /org/boot-policy/lan # create path {primary secondary}	Creates a primary or secondary LAN boot path and enters organization boot policy LAN path mode.
Step 6	UCS-A /org/boot-policy/lan/path # set vnic <i>vnic-name</i>	Specifies the vNIC to use for the LAN path to the boot image.
Step 7	UCS-A /org/boot-policy/lan/path # commit-buffer	Commits the transaction to the system configuration.

The following example enters the boot policy named lab2-boot-policy, creates a LAN boot for the policy, sets the boot order to 2, creates primary and secondary paths using the vNICs named vNIC1 and vNIC2, and commits the transaction:

```
UCS-A# scope org /
UCS-A /org* # scope boot-policy lab2-boot-policy
```

```

UCS-A /org/boot-policy* # create lan
UCS-A /org/boot-policy/lan* # set order 2
UCS-A /org/boot-policy/lan* # create path primary
UCS-A /org/boot-policy/lan/path* # set vnic vNIC1
UCS-A /org/boot-policy/lan/path* # exit
UCS-A /org/boot-policy/lan* # create path secondary
UCS-A /org/boot-policy/lan/path* # set vnic vNIC2
UCS-A /org/boot-policy/lan/path* # commit-buffer
UCS-A /org/boot-policy/lan/path #

```

What to Do Next

Include the boot policy in a service profile and/or template.

Local Disk Boot

If a server has a local drive, you can configure a boot policy to boot the server from that drive.



Note

Cisco UCS Manager does not differentiate between the types of local drives. If an operating system has been installed on more than one local drive or on an internal USB drive (eUSB), you cannot specify which of these local drives the server should use as the boot drive.

Configuring a Local Disk Boot for a Boot Policy

You can also create a local boot policy that is restricted to a service profile or service profile template. However, except for iSCSI boot, we recommend that you create a global boot policy that can be included in multiple service profiles or service profile templates.

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

Procedure

	Command or Action	Purpose
Step 1	UCS-A# scope org <i>org-name</i>	Enters organization mode for the specified organization. To enter the root organization mode, type / as the <i>org-name</i> .
Step 2	UCS-A /org # scope boot-policy <i>policy-name</i>	Enters organization boot policy mode for the specified boot policy.
Step 3	UCS-A /org/boot-policy # create storage	Creates a storage boot for the boot policy and enters organization boot policy storage mode.
Step 4	UCS-A /org/boot-policy/storage # create local	Creates a local storage location. The use of the terms primary or secondary boot devices does not imply a boot order. The effective order of boot devices within the same device class is determined by PCIe bus scan order.

	Command or Action	Purpose
Step 5	UCS-A /org/boot-policy/storage # commit-buffer	Commits the transaction to the system configuration.

The following example enters the boot policy named lab1-boot-policy, creates a local boot for the policy, and commits the transaction:

```
UCS-A# scope org /
UCS-A /org* # scope boot-policy lab1-boot-policy
UCS-A /org/boot-policy* # create storage
UCS-A /org/boot-policy/storage* # create local
UCS-A /org/boot-policy/storage* # commit-buffer
UCS-A /org/boot-policy/storage #
```

What to Do Next

Include the boot policy in a service profile and/or template.

Virtual Media Boot

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

Configuring a Virtual Media Boot for a Boot Policy



Note

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

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

Before You Begin

Create a boot policy to contain the virtual media boot configuration.

Procedure

	Command or Action	Purpose
Step 1	UCS-A# scope org <i>org-name</i>	Enters organization mode for the specified organization. To enter the root organization mode, type / as the <i>org-name</i> .

	Command or Action	Purpose
Step 2	UCS-A /org # scope boot-policy <i>policy-name</i>	Enters organization boot policy mode for the specified boot policy.
Step 3	UCS-A /org/boot-policy # create virtual-media {read-only read-write}	Creates a virtual media boot for the boot policy, specifies whether the virtual media is has read-only or read-write privileges, and enters organization boot policy virtual media mode.
Step 4	UCS-A /org/boot-policy/virtual-media # set order {1 2 3 4}	Sets the boot order for the virtual-media boot.
Step 5	UCS-A /org/boot-policy/virtual-media # commit-buffer	Commits the transaction to the system configuration.

The following example enters the boot policy named lab3-boot-policy, creates a virtual media boot with read-only privileges for the policy, sets the boot order to 3, and commits the transaction:

```
UCS-A# scope org /
UCS-A /org* # scope boot-policy lab3-boot-policy
UCS-A /org/boot-policy* # create virtual-media read-only
UCS-A /org/boot-policy/virtual-media* # set order 3
UCS-A /org/boot-policy/virtual-media* # commit-buffer
```

What to Do Next

Include the boot policy in a service profile and/or template.

Deleting a Boot Policy

Procedure

	Command or Action	Purpose
Step 1	UCS-A# scope org <i>org-name</i>	Enters organization mode for the specified organization. To enter the root organization mode, type / as the <i>org-name</i> .
Step 2	UCS-A /org # delete boot-policy <i>policy-name</i>	Deletes the specified boot policy.
Step 3	UCS-A /org # commit-buffer	Commits the transaction to the system configuration.

The following example deletes the boot policy named boot-policy-LAN and commits the transaction:

```
UCS-A# scope org /
UCS-A /org # delete boot-policy boot-policy-LAN
UCS-A /org* # commit-buffer
UCS-A /org #
```




CHAPTER 30

Deferring Deployment of Service Profile Updates

This chapter includes the following sections:

- [Deferred Deployment of Service Profiles, page 453](#)
- [Configuring Schedules, page 456](#)
- [Configuring Maintenance Policies, page 461](#)
- [Managing Pending Activities, page 462](#)

Deferred Deployment of Service Profiles

Some modifications to a service profile or to an updating service profile template can be disruptive and require a reboot of the server. You can, however, configure deferred deployment to control when those disruptive configuration changes are implemented. For example, you can choose to deploy the service profile changes immediately or have them deployed during a specified maintenance window. You can also choose whether or not a service profile deployment requires explicit user acknowledgement.

Deferred deployment is available for all configuration changes that occur through the association of a service profile with a server. These configuration changes can be prompted by a change to a service profile, to a policy that is included in a service profile, or to an updating service profile template. For example, you can defer the upgrade and activation of firmware through host firmware packages and management firmware packages, such as server BIOS, RAID controller, host HBA, and network adapters. However, you cannot defer the direct deployment of firmware images for components that do not use either of the firmware packages, such as Cisco UCS Manager, fabric interconnects, and I/O modules.

Deferred deployment is not available for the following actions which require the reboot of a server:

- Initial association of a service profile with a server
- Final disassociation of a service profile from a server, without associating the service profile with a different server
- Decommissioning a server
- Reacknowledging a server
- Resetting a server

If you want to defer the deployment of service profile changes, you must configure one or more maintenance policies and configure each service profile with a maintenance policy. If you want to define the time period when the deployment should occur, you also need to create at least one schedule with one or more recurring occurrences or one time occurrences, and include that schedule in a maintenance policy.

Deferred Deployment Schedules

A schedule contains a set of occurrences. These occurrences can be one time only or can recur at a specified time and day each week. The options defined in the occurrence, such as the duration of the occurrence or the maximum number of tasks to be run, determine whether a service profile change is deployed. For example, if a change cannot be deployed during a given maintenance window because the maximum duration or number of tasks has been reached, that deployment is carried over to the next maintenance window.

Each schedule checks periodically to see whether the Cisco UCS domain has entered one or more maintenance windows. If it has, the schedule executes the deployments that are eligible according to the constraints specified in the maintenance policy

A schedule contains one or more occurrences, which determine the maintenance windows associated with that schedule. An occurrence can be one of the following:

One Time Occurrence

One time occurrences define a single maintenance window. These windows continue until the maximum duration of the window or the maximum number of tasks that can be run in the window has been reached.

Recurring Occurrence

Recurring occurrences define a series of maintenance windows. These windows continue until the maximum number of tasks or the end of the day specified in the occurrence has been reached.

Maintenance Policy

A maintenance policy determines how Cisco UCS Manager reacts when a change that requires a server reboot is made to a service profile associated with a server or to an updating service profile bound to one or more service profiles.

The maintenance policy specifies how Cisco UCS Manager deploys the service profile changes. The deployment can occur in one of the following ways:

- Immediately
- When acknowledged by a user with admin privileges
- Automatically at the time specified in the schedule

If the maintenance policy is configured to deploy the change during a scheduled maintenance window, the policy must include a valid schedule. The schedule deploys the changes in the first available maintenance window.

Pending Activities

If you configure deferred deployment in a Cisco UCS domain, Cisco UCS Manager enables you to view all pending activities. You can see activities that are waiting for user acknowledgement and those that have been scheduled.

If a Cisco UCS domain has pending activities, Cisco UCS Manager GUI notifies users with admin privileges when they log in.

Cisco UCS Manager displays information about all pending activities, including the following:

- Name of the service profile to be deployed and associated with a server
- Server affected by the deployment
- Disruption caused by the deployment
- Change performed by the deployment



Note

You cannot specify the maintenance window in which a specific pending activity is applied to the server. The maintenance window depends upon how many activities are pending and which maintenance policy is assigned to the service profile. However, any user with admin privileges can manually initiate a pending activity and reboot the server immediately, whether it is waiting for user acknowledgment or for a maintenance window.

Guidelines and Limitations for Deferred Deployment

Cannot Undo All Changes to Service Profiles or Service Profile Templates

If you cancel a pending change, Cisco UCS Manager attempts to roll back the change without rebooting the server. However, for complex changes, Cisco UCS Manager may have to reboot the server a second time to roll back the change. For example, if you delete a vNIC, Cisco UCS Manager reboots the server according to the maintenance policy included in the service profile. You cannot cancel this reboot and change, even if you restore the original vNIC in the service profile. Instead, Cisco UCS Manager schedules a second deployment and reboot of the server.

Association of Service Profile Can Exceed Boundaries of Maintenance Window

After Cisco UCS Manager begins the association of the service profile, the scheduler and maintenance policy do not have any control over the procedure. If the service profile association does not complete within the allotted maintenance window, the process continues until it is completed. For example, this can occur if the association does not complete in time because of retried stages or other issues.

Cannot Specify Order of Pending Activities

Scheduled deployments run in parallel and independently. You cannot specify the order in which the deployments occur. You also cannot make the deployment of one service profile change dependent upon the completion of another.

Cannot Perform Partial Deployment of Pending Activity

Cisco UCS Manager applies all changes made to a service profile in the scheduled maintenance window. You cannot make several changes to a service profile at the same time and then have those changes be spread across several maintenance windows. When Cisco UCS Manager deploys the service profile changes, it updates the service profile to match the most recent configuration in the database.

Configuring Schedules

Creating a Schedule

Procedure

	Command or Action	Purpose
Step 1	UCS-A# scope system	Enters system mode.
Step 2	UCS-A /system # create scheduler <i>sched-name</i>	Creates a scheduler and enters scheduler mode.
Step 3	UCS-A /system/scheduler # commit-buffer	Commits the transaction to the system configuration.

The following example creates a scheduler called maintenancesched and commits the transaction:

```
UCS-A# scope system
UCS-A /system # create scheduler maintenancesched
UCS-A /system/scheduler* # commit-buffer
UCS-A /system/scheduler #
```

What to Do Next

Create a one time occurrence or recurring occurrence for the schedule.

Creating a One Time Occurrence for a Schedule

Procedure

	Command or Action	Purpose
Step 1	UCS-A# scope system	Enters system mode.
Step 2	UCS-A /system # scope schedule <i>sched-name</i>	Enters scheduler system mode.
Step 3	UCS-A /system/scheduler # create occurrence one-time <i>occurrence-name</i>	Creates a one-time occurrence.
Step 4	UCS-A /system/scheduler/one-time # set date <i>month day-of-month year hour minute</i>	Sets the date and time this occurrence should run.

	Command or Action	Purpose
Step 5	UCS-A /system/scheduler/one-time # set concur-tasks { unlimited <i>max-num-concur-tasks</i> }	(Optional) Sets the maximum number of tasks that can run concurrently during this occurrence. If the maximum number of tasks is reached, the scheduler waits for the amount of time set in the minimum interval property before scheduling new tasks.
Step 6	UCS-A /system/scheduler/one-time # set max-duration { none <i>num-of-days num-of-hours num-of-minutes num-of-seconds</i> }	(Optional) Sets the maximum length of time that this schedule occurrence can run. Cisco UCS completes as many scheduled tasks as possible within the specified time.
Step 7	UCS-A /system/scheduler/one-time # set min-interval { none <i>num-of-days num-of-hours num-of-minutes num-of-seconds</i> }	(Optional) Sets the minimum length of time that the system should wait before starting a new task.
Step 8	UCS-A /system/scheduler/one-time # set proc-cap { unlimited <i>max-num-of-tasks</i> }	(Optional) Sets the maximum number of scheduled tasks that can be run during this occurrence.
Step 9	UCS-A /system/scheduler/one-time # commit-buffer	Commits the transaction to the system configuration.

The following example creates a one time occurrence called onetimemaint for a scheduler called maintsched, sets the maximum number of concurrent tasks to 5, sets the start date to April 1, 2011 at 11:00, and commits the transaction:

```
UCS-A# scope system
UCS-A /system # scope scheduler maintsched
UCS-A /system/scheduler # create occurrence one-time onetimemaint
UCS-A /system/scheduler/one-time* # set date apr 1 2011 11 00
UCS-A /system/scheduler/one-time* # set concur-tasks 5
UCS-A /system/scheduler/one-time* # commit-buffer
UCS-A /system/scheduler/one-time #
```

Creating a Recurring Occurrence for a Schedule

Procedure

	Command or Action	Purpose
Step 1	UCS-A# scope system	Enters system mode.
Step 2	UCS-A /system # scope schedule <i>sched-name</i>	Enters scheduler system mode.

	Command or Action	Purpose
Step 3	UCS-A /system/scheduler # create occurrence recurring <i>occurrence-name</i>	Creates a recurring occurrence.
Step 4	UCS-A /system/scheduler/recurring # set day { even-day every-day friday monday never odd-day saturday sunday thursday tuesday wednesday }	(Optional) Specifies the day on which Cisco UCS runs an occurrence of this schedule. By default, this property is set to never.
Step 5	UCS-A /system/scheduler/recurring # set hour <i>hour</i>	(Optional) Specifies the hour at which this occurrence starts. Note Cisco UCS ends all recurring occurrences on the same day in which they start, even if the maximum duration has not been reached. For example, if you specify a start time of 11 p.m. and a maximum duration of 3 hours, Cisco UCS starts the occurrence at 11 p.m. but ends it at 11:59 p.m. after only 59 minutes.
Step 6	UCS-A /system/scheduler/recurring # set minute <i>minute</i>	(Optional) Specifies the minute at which this occurrence starts.
Step 7	UCS-A /system/scheduler/recurring # set concur-tasks { unlimited <i>max-num-concur-tasks</i> }	(Optional) Sets the maximum number of tasks that can run concurrently during this occurrence. If the maximum number of tasks is reached, the scheduler waits for the amount of time set in the minimum interval property before scheduling new tasks.
Step 8	UCS-A /system/scheduler/recurring # set max-duration { none <i>num-of-hours num-of-minutes num-of-seconds</i> }	(Optional) Sets the maximum length of time that this schedule occurrence can run. Cisco UCS completes as many scheduled tasks as possible within the specified time.
Step 9	UCS-A /system/scheduler/recurring # set min-interval { none <i>num-of-days num-of-hours num-of-minutes num-of-seconds</i> }	(Optional) Sets the minimum length of time that the system should wait before starting a new task.
Step 10	UCS-A /system/scheduler/recurring # set proc-cap { unlimited <i>max-num-of-tasks</i> }	(Optional) Sets the maximum number of scheduled tasks that can be run during this occurrence.
Step 11	UCS-A /system/scheduler/recurring # commit-buffer	Commits the transaction to the system configuration.

The following example creates a recurring occurrence called recurringmaint for a scheduler called maintsched, sets the maximum number of concurrent tasks to 5, sets the day this occurrence will run to even days, sets the time it will start to 11:05, and commits the transaction:

```
UCS-A# scope system
UCS-A /system # scope scheduler maintsched
UCS-A /system/scheduler # create occurrence recurring recurringmaint
UCS-A /system/scheduler/recurring* # set day even-day
UCS-A /system/scheduler/recurring* # set hour 11
UCS-A /system/scheduler/recurring* # set minute 5
UCS-A /system/scheduler/recurring* # set concur-tasks 5
UCS-A /system/scheduler/recurring* # commit-buffer
UCS-A /system/scheduler/recurring #
```

Deleting a One Time Occurrence from a Schedule

If this is the only occurrence in a schedule, that schedule is reconfigured with no occurrences. If the schedule is included in a maintenance policy and that policy is assigned to a service profile, any pending activities related to the server associated with the service profile cannot be deployed. You must add a one time occurrence or a recurring occurrence to the schedule to deploy the pending activity.

Procedure

	Command or Action	Purpose
Step 1	UCS-A# scope system	Enters system mode.
Step 2	UCS-A /system # scope scheduler <i>sched-name</i>	Enters scheduler system mode.
Step 3	UCS-A /system/scheduler # delete occurrence <i>one-time occurrence-name</i>	Deletes the specified one-time occurrence.
Step 4	UCS-A /system/scheduler # commit-buffer	Commits the transaction to the system configuration.

The following example deletes a one time occurrence called onetimemaint from scheduler maintsched and commits the transaction:

```
UCS-A# scope system
UCS-A /system # scope scheduler maintsched
UCS-A /system/scheduler # delete occurrence one-time onetimemaint
UCS-A /system/scheduler* # commit-buffer
UCS-A /system/scheduler #
```

Deleting a Recurring Occurrence from a Schedule

If this is the only occurrence in a schedule, that schedule is reconfigured with no occurrences. If the schedule is included in a maintenance policy and that policy is assigned to a service profile, any pending activities related to the server associated with the service profile cannot be deployed. You must add a one time occurrence or a recurring occurrence to the schedule to deploy the pending activity.

Procedure

	Command or Action	Purpose
Step 1	UCS-A# scope system	Enters system mode.
Step 2	UCS-A /system # scope scheduler <i>sched-name</i>	Enters scheduler system mode.
Step 3	UCS-A /system/scheduler # delete occurrence recurring <i>occurrence-name</i>	Deletes the specified recurring occurrence.
Step 4	UCS-A /system/scheduler # commit-buffer	Commits the transaction to the system configuration.

The following example deletes a recurring occurrence called onetimemaint from scheduler maintsched and commits the transaction:

```
UCS-A# scope system
UCS-A /system # scope scheduler maintsched
UCS-A /system/scheduler # delete occurrence recurring onetimemaint
UCS-A /system/scheduler* # commit-buffer
UCS-A /system/scheduler #
```

Deleting a Schedule

If this schedule is included in a maintenance policy, the policy is reconfigured with no schedule. If that policy is assigned to a service profile, any pending activities related to the server associated with the service profile cannot be deployed. You must add a schedule to the maintenance policy to deploy the pending activity.

Procedure

	Command or Action	Purpose
Step 1	UCS-A# scope system	Enters system mode.
Step 2	UCS-A /system # delete scheduler <i>sched-name</i>	Deletes a scheduler and enters scheduler mode.
Step 3	UCS-A /system # commit-buffer	Commits the transaction to the system configuration.

The following example deletes a scheduler called maintenancesched and commits the transaction:

```
UCS-A# scope system
UCS-A /system # delete scheduler maintenancesched
UCS-A /system* # commit-buffer
UCS-A /system #
```

Configuring Maintenance Policies

Creating a Maintenance Policy

Before You Begin

If you plan to configure this maintenance policy for deferred deployment, create a schedule.

Procedure

	Command or Action	Purpose
Step 1	UCS-A# scope org <i>org-name</i>	Enters organization mode for the specified organization. To enter the root organization mode, type <i>/</i> as the <i>org-name</i> .
Step 2	UCS-A /org # create maint-policy <i>policy-name</i>	Creates the specified maintenance policy and enters maintenance policy mode.
Step 3	UCS-A /org/maint-policy # set reboot-policy { immediate timer-automatic user-ack }	When a service profile is associated with a server, the server needs to be rebooted to complete the association. Specifying the reboot-policy command determines when the reboot occurs for all service profiles that include this maintenance policy. Possible values include: <ul style="list-style-type: none"> • immediate--The server reboots as soon as the change is made to the service profile. • timer-automatic --You select the schedule that specifies when maintenance operations can be applied to the server using the set scheduler command. Cisco UCS reboots the server and completes the service profile changes at the scheduled time. • user-ack --The user must explicitly acknowledge the changes by using the apply pending-changes command before changes are applied.
Step 4	UCS-A /org/maint-policy # set scheduler <i>scheduler-name</i>	(Optional) If the reboot-policy property is set to timer-automatic , you must select the schedule that specifies when maintenance operations can be applied to the server. Cisco UCS reboots the server and completes the service profile changes at the scheduled time.
Step 5	UCS-A /org/maint-policy # commit-buffer	Commits the transaction to the system configuration.

The following example creates a maintenance policy called maintenance, sets the system to reboot immediately when a service profile is associated with a server, and commits the transaction:

```
UCS-A# scope org /
UCS-A /org # create maint-policy maintenance
UCS-A /org/maint-policy* # set reboot-policy immediate
```

```
UCS-A /org/maint-policy* # commit-buffer
UCS-A /org/maint-policy #
```

Deleting a Maintenance Policy

Procedure

	Command or Action	Purpose
Step 1	UCS-A# scope org <i>org-name</i>	Enters organization mode for the specified organization. To enter the root organization mode, type / as the <i>org-name</i> .
Step 2	UCS-A /org # delete maint-policy <i>policy-name</i>	Deletes the specified maintenance policy.
Step 3	UCS-A /org # commit-buffer	Commits the transaction to the system configuration.

The following example deletes a maintenance policy called maintenance and commits the transaction:

```
UCS-A# scope org /
UCS-A /org # delete maint-policy maintenance
UCS-A /org/maint-policy* # commit-buffer
UCS-A /org/maint-policy #
```

Managing Pending Activities

Viewing Pending Activities

Procedure

	Command or Action	Purpose
Step 1	UCS-A# scope org <i>org-name</i>	Enters organization mode. To enter the root organization mode, type / as the <i>org-name</i> .
Step 2	UCS-A /org # scope service-profile <i>profile-name</i>	Enters organization service profile mode for the specified service.
Step 3	UCS-A /org/service-profile # show pending-changes [detail expand]	Displays details about pending-changes.

The following example shows how to display pending changes for a service profile called accounting:

```
UCS-A# scope org /
UCS-A /org # scope service-profile accounting
UCS-A /org/service-profile # show pending-changes detail
```

Pending Changes:


```

Scheduler:
Changed by: admin
Aked by:
Mod. date: 2010-09-20T20:36:09.254
State: Untriggered
Admin State: Untriggered
Pend. Changes: 0
Pend. Disr.: 0
UCS-A /org/service-profile #

```

Deploying a Service Profile Change Waiting for User Acknowledgement

Cisco UCS Manager CLI cannot deploy all pending service profile changes (for multiple service profiles) waiting for user acknowledgement. To simultaneously deploy all pending service profile changes for multiple service profiles, use Cisco UCS Manager GUI.



Important You cannot stop Cisco UCS Manager from rebooting the affected server after you acknowledge a pending activity.

Procedure

	Command or Action	Purpose
Step 1	UCS-A# scope org <i>org-name</i>	Enters organization mode. To enter the root organization mode, type / as the <i>org-name</i> .
Step 2	UCS-A /org # scope service-profile <i>profile-name</i>	Enters organization service profile mode for the specified service.
Step 3	UCS-A /org/service-profile # apply pending-changes immediate	Applies the pending changes immediately. Cisco UCS Manager immediately reboots the server affected by the pending activity.

The following example shows how to apply pending changes for a service profile called accounting:

```

UCS-A# scope org /
UCS-A /org # scope service-profile accounting
UCS-A /org/service-profile # apply pending-changes immediate
UCS-A /org/service-profile #

```

Deploying a Scheduled Service Profile Change Immediately

Cisco UCS Manager CLI cannot deploy all scheduled service profile changes (for multiple service profiles) at the same time. To simultaneously deploy all scheduled service profile changes for multiple service profiles, use Cisco UCS Manager GUI.



Important You cannot stop Cisco UCS Manager from rebooting the affected server after you acknowledge a pending activity.

Procedure

	Command or Action	Purpose
Step 1	UCS-A# scope org <i>org-name</i>	Enters organization mode. To enter the root organization mode, type / as the <i>org-name</i> .
Step 2	UCS-A /org # scope service-profile <i>profile-name</i>	Enters organization service profile mode for the specified service.
Step 3	UCS-A /org/service-profile # apply pending-changes immediate	Applies the pending changes immediately. Cisco UCS Manager immediately reboots the server affected by the pending activity.

The following example shows how to apply pending changes for a service profile called accounting:

```
UCS-A# scope org /
UCS-A /org # scope service-profile accounting
UCS-A /org/service-profile # apply pending-changes immediate
UCS-A /org/service-profile #
```



CHAPTER 31

Configuring Service Profiles

This chapter includes the following sections:

- [Service Profiles that Override Server Identity, page 465](#)
- [Service Profiles that Inherit Server Identity, page 466](#)
- [Guidelines and Recommendations for Service Profiles, page 466](#)
- [Service Profile Templates, page 467](#)
- [Creating a Hardware-Based Service Profile, page 471](#)
- [Configuring a vNIC for a Service Profile, page 474](#)
- [Configuring a vHBA for a Service Profile, page 476](#)
- [Configuring a Local Disk for a Service Profile, page 478](#)
- [Configuring Serial over LAN for a Service Profile, page 479](#)
- [Service Profile Boot Definition Configuration, page 480](#)
- [Service Profiles and Service Profile Template Management, page 485](#)

Service Profiles that Override Server Identity

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

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

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

- Virtualized identity information, including pools of MAC addresses, WWN addresses, and UUIDs

- Ethernet and Fibre Channel adapter profile policies
- Firmware package policies
- Operating system boot order policies

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

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

**Note**

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

Service Profiles that Inherit Server Identity

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

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

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

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

**Important**

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

Guidelines and Recommendations for Service Profiles

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

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

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

No Power Capping Support for Rack-Mount Servers

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

Service Profile Templates

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



Tip

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

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

Cisco UCS supports the following types of service profile templates:

Initial template

Service profiles created from an initial template inherit all the properties of the template. However, after you create the profile, it is no longer connected to the template. If you need to make changes to one or more profiles created from this template, you must change each profile individually.

Updating template

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

Creating a Service Profile Template

Procedure

	Command or Action	Purpose
Step 1	UCS-A# scope org <i>org-name</i>	Enters organization mode for the specified organization. To enter the root organization mode, type / as the <i>org-name</i> .
Step 2	UCS-A /org # create service-profile <i>profile-name</i> { initial-template updating-template }	Creates the specified service profile template and enters organization service profile mode. Enter a unique <i>profile-name</i> to identify this service profile template.

	Command or Action	Purpose
		This name can be between 2 and 32 alphanumeric characters. You cannot use spaces or any special characters other than - (hyphen), _ (underscore), : (colon), and . (period), and this name must be unique across all service profiles and service profile templates within the same organization.
Step 3	UCS-A /org/service-profile # set bios-policy <i>policy-name</i>	Associates the specified BIOS policy with the service profile.
Step 4	UCS-A /org/service-profile # set boot-policy <i>policy-name</i>	Associates the specified boot policy with the service profile.
Step 5	UCS-A /org/service-profile # set descr <i>description</i>	(Optional) Provides a description for the service profile. Note If your description includes spaces, special characters, or punctuation, you must begin and end your description with quotation marks. The quotation marks will not appear in the description field of any show command output.
Step 6	UCS-A /org/service-profile # set dynamic-vnic-conn-policy <i>policy-name</i>	Associates the specified dynamic vNIC connection policy with the service profile.
Step 7	UCS-A /org/service-profile # set ext-mgmt-ip-state { none pooled }	Specifies how the management IP address will be assigned to the service profile. You can set the management IP address policy using the following options: <ul style="list-style-type: none"> • None-- The service profile is not assigned an IP address. • Pooled-- The service profile is assigned an IP address from the management IP pool. Note Setting the management IP address to static for a service profile template will result in an error.
Step 8	UCS-A /org/service-profile # set host-fw-policy <i>policy-name</i>	Associates the specified host firmware policy with the service profile.
Step 9	UCS-A /org/service-profile # set identity { dynamic-uuid { <i>uuid</i> derived } dynamic-wwnn { <i>wwnn</i> derived } uuid-pool <i>pool-name</i> wwnn-pool <i>pool-name</i> }	Specifies how the server acquires a UUID or WWNN. You can do one of the following: <ul style="list-style-type: none"> • Create a unique UUID in the form <i>nnnnnnnn-nnnn-nnnn-nnnnnnnnnnnn</i> . • Derive the UUID from the one burned into the hardware at manufacture. • Use a UUID pool.

	Command or Action	Purpose
		<ul style="list-style-type: none"> • Create a unique WWNN in the form <i>hh : hh : hh : hh : hh : hh</i>. • Derive the WWNN from one burned into the hardware at manufacture. • Use a WWNN pool.
Step 10	UCS-A /org/service-profile # set ipmi-access-profile <i>profile-name</i>	Associates the specified IPMI access profile with the service profile.
Step 11	UCS-A /org/service-profile # set local-disk-policy <i>policy-name</i>	Associates the specified local disk policy with the service profile.
Step 12	UCS-A /org/service-profile # set maint-policy <i>policy-name</i>	Associates the specified maintenance policy with the service profile.
Step 13	UCS-A /org/service-profile # set mgmt-fw-policy <i>policy-name</i>	Associates the specified management firmware policy with the service profile.
Step 14	UCS-A /org/service-profile # set power-control-policy <i>policy-name</i>	Associates the specified power control policy with the service profile.
Step 15	UCS-A /org/service-profile # set scrub-policy <i>policy-name</i>	Associates the specified scrub policy with the service profile.
Step 16	UCS-A /org/service-profile # set sol-policy <i>policy-name</i>	Associates the specified serial over LAN policy with the service profile.
Step 17	UCS-A /org/service-profile # set stats-policy <i>policy-name</i>	Associates the specified statistics policy with the service profile.
Step 18	UCS-A /org/service-profile # set user-label <i>label-name</i>	Specifies the user label associated with the service profile.
Step 19	UCS-A /org/service-profile # set vcon {1 2} selection {all assigned-only exclude-dynamic exclude-unassigned}	Specifies the selection preference for the specified vCon.
Step 20	UCS-A /org/service-profile # set vcon-profile <i>policy-name</i>	<p>Associates the specified vNIC/vHBA placement profile with the service profile.</p> <p>Note You can either assign a vNIC/vHBA placement profile to the service profile, or set vCon selection preferences for the service profile, but you do not need to do both.</p>
Step 21	UCS-A /org/service-profile # commit-buffer	Commits the transaction to the system configuration.

The following example shows how to create a service profile template and commit the transaction:

```
UCS-A# scope org /
UCS-A /org* # create service-profile ServTemp2 updating-template
UCS-A /org/service-profile* # set bios-policy biospol1
UCS-A /org/service-profile* # set boot-policy bootpol32
UCS-A /org/service-profile* # set descr "This is a service profile example."
UCS-A /org/service-profile* # set dynamic-vnic-conn-policy mydynvnicconnpolicy
UCS-A /org/service-profile* # set ext-mgmt-ip-state pooled
UCS-A /org/service-profile* # set host-fw-policy ipmi-user987
UCS-A /org/service-profile* # set identity dynamic-uuid derived
UCS-A /org/service-profile* # set ipmi-access-profile ipmiProf16
UCS-A /org/service-profile* # set local-disk-policy localdiskpol33
UCS-A /org/service-profile* # set maint-policy maintpol4
UCS-A /org/service-profile* # set mgmt-fw-policy mgmtfwpol175
UCS-A /org/service-profile* # set power-control-policy powcontrpol13
UCS-A /org/service-profile* # set scrub-policy scrubpol55
UCS-A /org/service-profile* # set sol-policy solpol2
UCS-A /org/service-profile* # set stats-policy statspol4
UCS-A /org/service-profile* # set user-label mylabel
UCS-A /org/service-profile* # vcon-policy myvconnpolicy
UCS-A /org/service-profile* # commit-buffer
UCS-A /org/service-profile #
```

What to Do Next

- (Optional) Configure a boot definition for the service profile. Use this option only if you have not associated a boot policy with the service profile.
- Create a service profile instance from the service profile template.

Creating a Service Profile Instance from a Service Profile Template

Before You Begin

Verify that there is a service profile template from which to create a service profile instance.

Procedure

	Command or Action	Purpose
Step 1	UCS-A# scope org <i>org-name</i>	Enters organization mode for the specified organization. To enter the root organization mode, type <i>/</i> as the <i>org-name</i> .
Step 2	UCS-A /org # create service-profile <i>profile-name instance</i>	Creates the specified service profile instance and enters organization service profile mode. Enter a unique <i>profile-name</i> to identify this service profile template. This name can be between 2 and 32 alphanumeric characters. You cannot use spaces or any special characters other than - (hyphen), _ (underscore), : (colon), and . (period), and this name must be unique across all service profiles and service profile templates within the same organization.
Step 3	UCS-A /org/service-profile # set src-templ-name <i>profile-name</i>	Specifies the source service profile template to apply to the service profile instance. All configuration settings from the

	Command or Action	Purpose
		service profile template will be applied to the service profile instance.
Step 4	UCS-A /org/service-profile # commit-buffer	Commits the transaction to the system configuration.

The following example creates a service profile instance named ServProf34, applies the service profile template named ServTemp2, and commits the transaction:

```
UCS-A# scope org /
UCS-A /org* # create service-profile ServProf34 instance
UCS-A /org/service-profile* # set src-templ-name ServTemp2
UCS-A /org/service-profile* # commit-buffer
UCS-A /org/service-profile #
```

What to Do Next

Associate the service profile to a server, rack server, or server pool.

Creating a Hardware-Based Service Profile

Procedure

	Command or Action	Purpose
Step 1	UCS-A# scope org <i>org-name</i>	Enters organization mode for the specified organization. To enter the root organization mode, type / as the <i>org-name</i> .
Step 2	UCS-A /org # create service-profile <i>profile-name</i> instance	Creates the specified service profile instance and enters organization service profile mode. Enter a unique <i>profile-name</i> to identify this service profile template. This name can be between 2 and 32 alphanumeric characters. You cannot use spaces or any special characters other than - (hyphen), _ (underscore), : (colon), and . (period), and this name must be unique across all service profiles and service profile templates within the same organization.
Step 3	UCS-A /org/service-profile # set bios-policy <i>policy-name</i>	Associates the specified BIOS policy with the service profile.
Step 4	UCS-A /org/service-profile # set boot-policy <i>policy-name</i>	Associates the specified boot policy with the service profile.
Step 5	UCS-A /org/service-profile # set descr <i>description</i>	(Optional) Provides a description for the service profile.

	Command or Action	Purpose
		<p>Note If your description includes spaces, special characters, or punctuation, you must begin and end your description with quotation marks. The quotation marks will not appear in the description field of any show command output.</p>
Step 6	UCS-A /org/service-profile # set dynamic-vnic-conn-policy <i>policy-name</i>	Associates the specified dynamic vNIC connection policy with the service profile.
Step 7	UCS-A /org/service-profile # set ext-mgmt-ip-state { none pooled static }	<p>Specifies how the management IP address will be assigned to the service profile.</p> <p>You can set the management IP address policy using the following options:</p> <ul style="list-style-type: none"> • None-- The service profile is not assigned an IP address. • Pooled-- The service profile is assigned an IP address from the management IP pool. • Static-- The service profile is assigned the configured static IP address.
Step 8	UCS-A /org/service-profile # set host-fw-policy <i>ipmi-user-name</i>	Associates the specified host forwarding policy with the service profile.
Step 9	UCS-A /org/service-profile # set identity { dynamic-uuid { <i>uuid</i> derived } dynamic-wwnn { <i>wwnn</i> derived } uuid-pool <i>pool-name</i> wwnn-pool <i>pool-name</i> }	<p>Specifies how the server acquires a UUID or WWNN. You can do one of the following:</p> <ul style="list-style-type: none"> • Create a unique UUID in the form <i>nnnnnnnnn-nnnnn-nnnnn-nnnnnnnnnnnnn</i>. • Derive the UUID from the one burned into the hardware at manufacture. • Use a UUID pool. • Create a unique WWNN in the form <i>hh : hh : hh : hh : hh : hh</i>. • Derive the WWNN from one burned into the hardware at manufacture. • Use a WWNN pool.
Step 10	UCS-A /org/service-profile # set ipmi-access-profile <i>profile-name</i>	Associates the specified IPMI access profile with the service profile.
Step 11	UCS-A /org/service-profile # set local-disk-policy <i>policy-name</i>	Associates the specified local disk policy with the service profile.

	Command or Action	Purpose
Step 12	UCS-A /org/service-profile # set maint-policy <i>policy-name</i>	Associates the specified maintenance policy with the service profile.
Step 13	UCS-A /org/service-profile # set mgmt-fw-policy <i>policy-name</i>	Associates the specified management forwarding policy with the service profile.
Step 14	UCS-A /org/service-profile # set power-control-policy <i>policy-name</i>	Associates the specified power control policy with the service profile.
Step 15	UCS-A /org/service-profile # set scrub-policy <i>policy-name</i>	Associates the specified scrub policy with the service profile.
Step 16	UCS-A /org/service-profile # set sol-policy <i>policy-name</i>	Associates the specified serial over LAN policy with the service profile.
Step 17	UCS-A /org/service-profile # set stats-policy <i>policy-name</i>	Associates the specified statistics policy with the service profile.
Step 18	UCS-A /org/service-profile # set user-label <i>label-name</i>	Specifies the user label associated with the service profile.
Step 19	UCS-A /org/service-profile # set vcon {1 2} selection {all assigned-only exclude-dynamic exclude-unassigned}	Specifies the selection preference for the specified vCon.
Step 20	UCS-A /org/service-profile # set vcon-policy <i>policy-name</i>	Associates the specified vNIC/vHBA placement policy with the service profile. Note You can either assign a vNIC/vHBA placement profile to the service profile, or set vCon selection preferences for the service profile, but you do not need to do both.
Step 21	UCS-A /org/service-profile # commit-buffer	Commits the transaction to the system configuration.

The following example shows how to create a service profile instance and commit the transaction:

```
UCS-A# scope org /
UCS-A /org* # create service-profile ServInst90 instance
UCS-A /org/service-profile* # set bios-policy biospol1
UCS-A /org/service-profile* # set boot-policy bootpol32
UCS-A /org/service-profile* # set descr "This is a service profile example."
UCS-A /org/service-profile* # set ext-mgmt-ip-state pooled
UCS-A /org/service-profile* # set host-fw-policy ipmi-user987
UCS-A /org/service-profile* # set identity dynamic-uuid derived
UCS-A /org/service-profile* # set ipmi-access-profile ipmiProf16
UCS-A /org/service-profile* # set local-disk-policy localdiskpol133
UCS-A /org/service-profile* # set maint-policy maintpol4
UCS-A /org/service-profile* # set mgmt-fw-policy mgmtfwpol175
UCS-A /org/service-profile* # set power-control-policy powcontrpol113
UCS-A /org/service-profile* # set scrub-policy scrubpol155
UCS-A /org/service-profile* # set sol-policy solpol2
UCS-A /org/service-profile* # set stats-policy statspol4
UCS-A /org/service-profile* # set user-label mylabel
```

```
UCS-A /org/service-profile* # vcon-policy myvconnpolicy
UCS-A /org/service-profile* # commit-buffer
UCS-A /org/service-profile #
```

What to Do Next

- (Optional) Configure a boot definition for the service profile. Use this option only if you have not associated a boot policy with the service profile.
- Associate the service profile with a blade server, server pool, or rack server.

Configuring a vNIC for a Service Profile

Procedure

	Command or Action	Purpose
Step 1	UCS-A# scope org <i>org-name</i>	Enters organization mode for the specified organization. To enter the root organization mode, type / as the <i>org-name</i> .
Step 2	UCS-A /org # scope service-profile <i>profile-name</i>	Enters organization service profile mode for the specified service.
Step 3	UCS-A /org/service-profile # create vnic <i>vnic-name</i> [<i>eth-if eth-if-name</i>] [fabric { a b }]	Creates a vNIC for the specified service profile and enters organization service profile vNIC mode.
Step 4	UCS-A /org/service-profile/vnic # set adapter-policy <i>policy-name</i>	Specifies the adapter policy to use for the vNIC.
Step 5	UCS-A /org/service-profile/vnic # set fabric { a a-b b b-a }	<p>Specifies the fabric to use for the vNIC. If you did not specify the fabric when creating the vNIC template in Step 3, you have the option to specify it with this command.</p> <p>If you want this vNIC to be able to access the second fabric interconnect if the default one is unavailable, choose a-b (A is the primary) or b-a (B is the primary).</p> <p>Note Do not enable fabric failover for the vNIC under the following circumstances:</p> <ul style="list-style-type: none"> • If the Cisco UCS domain is running in Ethernet Switch Mode. vNIC fabric failover is not supported in that mode. If all Ethernet uplinks on one fabric interconnect fail, the vNICs do not fail over to the other. • if you plan to associate this vNIC with a server that has an adapter which does not support fabric failover, such as the Cisco UCS 82598KR-CI 10-Gigabit Ethernet Adapter. If you do so, Cisco UCS Manager generates a configuration fault when you associate the service profile with the server.

	Command or Action	Purpose
Step 6	UCS-A /org/service-profile/vnic # set identity { dynamic-mac { <i>mac-addr</i> derived } mac-pool <i>mac-pool-name</i> }	Specifies the identity (MAC address) for the vNIC. You can set the identity using one of the following options: <ul style="list-style-type: none"> • Create a unique MAC address in the form <i>nn : nn : nn : nn</i>. • Derive the MAC address from one burned into the hardware at manufacture. • Assign a MAC address from a MAC pool.
Step 7	UCS-A /org/service-profile/vnic # set mtu <i>size-num</i>	The maximum transmission unit, or packet size, that this vNIC accepts. Enter an integer between 1500 and 9216. Note If the vNIC has an associated QoS policy, the MTU specified here must be equal to or less than the MTU specified in the associated QoS system class. If this MTU value exceeds the MTU value in the QoS system class, packets may be dropped during data transmission.
Step 8	UCS-A /org/service-profile/vnic # set nw-control-policy <i>policy-name</i>	The network control policy the vNIC should use.
Step 9	UCS-A /org/service-profile/vnic # set order { <i>order-num</i> unspecified }	Specifies the relative order for the vNIC.
Step 10	UCS-A /org/service-profile/vnic # set pin-group <i>group-name</i>	The LAN pin group the vNIC should use.
Step 11	UCS-A /org/service-profile/vnic # set qos-policy <i>policy-name</i>	The quality of service policy the vNIC should use.
Step 12	UCS-A /org/service-profile/vnic # set stats-policy <i>policy-name</i>	The statistics collection policy the vNIC should use.
Step 13	UCS-A /org/service-profile/vnic # set template-name <i>policy-name</i>	Specifies the dynamic vNIC connectivity policy to use for the vNIC.
Step 14	UCS-A /org/service-profile/vnic # set vcon { 1 2 any }	Assigns the vNIC to the specified vCon. Use the any keyword to have Cisco UCS Manager automatically assign the vNIC.
Step 15	UCS-A /org/service-profile/vnic # commit-buffer	Commits the transaction to the system configuration.

The following example configures a vNIC for a service profile and commits the transaction:

```
UCS-A# scope org /
UCS-A /org* # scope service-profile ServInst90
UCS-A /org/service-profile* # create vnic vnic3 fabric a
UCS-A /org/service-profile/vnic* # set adapter-policy AdaptPol2
UCS-A /org/service-profile/vnic* # set fabric a-b
```

```

UCS-A /org/service-profile/vnic* # set identity mac-pool MacPool3
UCS-A /org/service-profile/vnic* # set mtu 8900
UCS-A /org/service-profile/vnic* # set nw-control-policy ncp5
UCS-A /org/service-profile/vnic* # set order 0
UCS-A /org/service-profile/vnic* # set pin-group EthPinGroup12
UCS-A /org/service-profile/vnic* # set qos-policy QosPol5
UCS-A /org/service-profile/vnic* # set stats-policy StatsPol2
UCS-A /org/service-profile/vnic* # set template-name VnicConnPol3
UCS-A /org/service-profile/vnic* # set set vcon any
UCS-A /org/service-profile/vnic* # commit-buffer
UCS-A /org/service-profile/vnic #

```

Configuring a vHBA for a Service Profile

Procedure

	Command or Action	Purpose
Step 1	UCS-A# scope org <i>org-name</i>	Enters organization mode for the specified organization. To enter the root organization mode, type / as the <i>org-name</i> .
Step 2	UCS-A /org # scope service-profile <i>profile-name</i>	Enters organization service profile mode for the specified service.
Step 3	UCS-A /org/service-profile # create vhma <i>vhba-name</i> [fabric {a b}] [fc-if <i>fc-if-name</i>]	Creates a vHBA for the specified service profile and enters organization service profile vHBA mode.
Step 4	UCS-A /org/service-profile/vhba # set adapter-policy <i>policy-name</i>	Specifies the adapter policy to use for the vHBA.
Step 5	UCS-A /org/service-profile/vhba # set admin-vcon {1 2 any}	Assigns the vHBA to one or all virtual network interface connections.
Step 6	UCS-A /org/service-profile/vhba # set identity { dynamic-wwpn { <i>wwpn</i> derived } wwpn-pool <i>wwn-pool-name</i> }	<p>Specifies the storage identity (world wide port name [WWPN]) for the vHBA.</p> <p>You can set the storage identity using one of the following options:</p> <ul style="list-style-type: none"> • Create a unique WWPN in the form <i>hh : hh : hh : hh : hh : hh</i>. You can specify a WWPN in the range from 20:00:00:00:00:00:00 to 20:FF:FF:FF:FF:FF:FF or from 50:00:00:00:00:00:00 to 5F:FF:FF:FF:FF:FF:FF. • Derive the WWPN from one burned into the hardware at manufacture. • Assign a WWPN from a WWN pool.

	Command or Action	Purpose
		Note If you want the WWPN to be compatible with Cisco MDS Fibre Channel switches, use the manual template and WWPN 20:00:00:25:B5:XX:XX:XX .
Step 7	UCS-A /org/service-profile/vhba # set max-field-size <i>size-num</i>	Specifies the maximum size of the Fibre Channel frame payload (in bytes) that the vHBA supports.
Step 8	UCS-A /org/service-profile/vhba # set order { <i>order-num</i> unspecified }	Specifies the PCI scan order for the vHBA.
Step 9	UCS-A /org/service-profile/vhba # set pers-bind { disabled enabled }	Disables or enables persistent binding to Fibre Channel targets.
Step 10	UCS-A /org/service-profile/vhba # set pin-group <i>group-name</i>	Specifies the pin group to use for the vHBA.
Step 11	UCS-A /org/service-profile/vhba # set qos-policy <i>policy-name</i>	Specifies the QoS policy to use for the vHBA.
Step 12	UCS-A /org/service-profile/vhba # set stats-policy <i>policy-name</i>	Specifies the stats policy to use for the vHBA.
Step 13	UCS-A /org/service-profile/vhba # set template-name <i>policy-name</i>	Specifies the vHBA SAN connectivity policy to use for the vHBA.
Step 14	UCS-A /org/service-profile/vhba # commit-buffer	Commits the transaction to the system configuration.

The following example configures a vHBA for a service profile and commits the transaction:

```
UCS-A# scope org /
UCS-A /org* # scope service-profile ServInst90
UCS-A /org/service-profile* # create vhba vhba3 fabric b
UCS-A /org/service-profile/vhba* # set adapter-policy AdaptPol2
UCS-A /org/service-profile/vhba* # set set admin-vcon any
UCS-A /org/service-profile/vhba* # set admin-vcon any
UCS-A /org/service-profile/vhba* # set identity wwpn-pool SanPool7
UCS-A /org/service-profile/vhba* # set max-field-size 2112
UCS-A /org/service-profile/vhba* # set order 0
UCS-A /org/service-profile/vhba* # set pers-bind enabled
UCS-A /org/service-profile/vhba* # set pin-group FcPinGroup12
UCS-A /org/service-profile/vhba* # set qos-policy QosPol5
UCS-A /org/service-profile/vhba* # set stats-policy StatsPol2
UCS-A /org/service-profile/vhba* # set template-name SanConnPol3
UCS-A /org/service-profile/vhba* # commit-buffer
UCS-A /org/service-profile/vhba #
```

Configuring a Local Disk for a Service Profile

Procedure

	Command or Action	Purpose
Step 1	UCS-A# scope org <i>org-name</i>	Enters organization mode for the specified organization. To enter the root organization mode, type / as the <i>org-name</i> .
Step 2	UCS-A /org # scope service-profile <i>profile-name</i>	Enters organization service profile mode for the specified service profile.
Step 3	UCS-A /org/service-profile # create local-disk-config	Creates a local disk configuration for the service profile and enters organization service profile local disk configuration mode.
Step 4	UCS-A /org/service-profile/local-disk-config # set descr <i>description</i>	(Optional) Provides a description for the local disk configuration.
Step 5	UCS-A /org/service-profile/local-disk-config # set mode { any-configuration no-local-storage no-raid raid-0-striped raid-1-mirrored raid-5-striped-parity raid-6-striped-dual-parity raid-10-mirrored-and-striped }	Specifies the mode for the local disk.
Step 6	UCS-A /org/service-profile/local-disk-config # create partition	Creates a partition for the local disk and enters organization service profile local disk configuration partition mode.
Step 7	UCS-A /org/service-profile/local-disk-config/partition # set descr <i>description</i>	(Optional) Provides a description for the partition.
Step 8	UCS-A /org/service-profile/local-disk-config/partition # set size { <i>size-num</i> unspecified }	Specifies the partition size in MBytes.
Step 9	UCS-A /org/service-profile/local-disk-config/partition # set type { ext2 ext3 fat32 none ntfs swap }	Specifies the partition type.
Step 10	UCS-A /org/service-profile/local-disk-config/partition # commit-buffer	Commits the transaction to the system configuration.

The following example configures a local disk for a service profile and commits the transaction:

```
UCS-A# scope org /
UCS-A /org # scope service-profile ServInst90
UCS-A /org/service-profile # scope boot-definition
```



```

UCS-A /org/service-profile # create local-disk-config
UCS-A /org/service-profile/local-disk-config* # set mode raid-1-mirrored
UCS-A /org/service-profile/local-disk-config* # create partition
UCS-A /org/service-profile/local-disk-config/partition* # set size 1000000
UCS-A /org/service-profile/local-disk-config/partition* # set type ntfs
UCS-A /org/service-profile/local-disk-config/partition* # commit-buffer
UCS-A /org/service-profile/local-disk-config/partition #

```

Configuring Serial over LAN for a Service Profile

Procedure

	Command or Action	Purpose
Step 1	UCS-A# scope org <i>org-name</i>	Enters organization mode for the specified organization. To enter the root organization mode, type <i>/</i> as the <i>org-name</i> .
Step 2	UCS-A /org # scope service-profile <i>profile-name</i>	Enters organization service profile mode for the specified service.
Step 3	UCS-A /org/service-profile # create sol-config	Creates a serial over LAN configuration for the service profile and enters organization service profile SoL configuration mode.
Step 4	UCS-A /org/service-profile/sol-config # {disable enable}	Disables or enables the serial over LAN configuration for the service profile.
Step 5	UCS-A /org/service-profile/sol-config # set descr <i>description</i>	(Optional) Provides a description for the serial over LAN configuration.
Step 6	UCS-A /org/service-profile/sol-config # set speed {115200 19200 38400 57600 9600}	Specifies the serial baud rate.
Step 7	UCS-A /org/service-profile/sol-config # commit-buffer	Commits the transaction to the system configuration.

The following example configures serial over LAN for the service profile named ServInst90 and commits the transaction:

```

UCS-A# scope org /
UCS-A /org # scope service-profile ServInst90
UCS-A /org/service-profile # create sol-config
UCS-A /org/service-profile/sol-config* # enable
UCS-A /org/service-profile/sol-config* # set descr "Sets serial over LAN to 9600 baud."
UCS-A /org/service-profile/sol-config* # set speed 9600
UCS-A /org/service-profile/sol-config* # commit-buffer
UCS-A /org/service-profile/sol-config #

```

Service Profile Boot Definition Configuration

Configuring a Boot Definition for a Service Profile

Procedure

	Command or Action	Purpose
Step 1	UCS-A# scope org <i>org-name</i>	Enters organization mode for the specified organization. To enter the root organization mode, type / as the <i>org-name</i> .
Step 2	UCS-A /org # scope service-profile <i>profile-name</i>	Enters organization service profile mode for the the specified service.
Step 3	UCS-A /org/service-profile # create boot-definition	Creates a boot definition for the service profile and enters organization service profile boot definition mode.
Step 4	UCS-A /org/service-profile/boot-definition # set descr <i>description</i>	(Optional) Provides a description for the boot definition.
Step 5	UCS-A /org/service-profile/boot-definition # set reboot-on-update {no yes}	(Optional) Specifies whether to automatically reboot all servers that use this boot definition after changes are made to the boot order. By default, the reboot on update option is disabled.
Step 6	UCS-A /org/service-profile/boot-definition # commit-buffer	Commits the transaction to the system configuration.

The following example configures a boot definition for a service profile and commits the transaction:

```
UCS-A# scope org /
UCS-A /org* # scope service-profile ServInst90
UCS-A /org/service-profile* # create boot-definition
UCS-A /org/service-profile/boot-definition* # set descr "This boot definition reboots on
update."
UCS-A /org/service-profile/boot-definition* # set reboot-on-update yes
UCS-A /org/service-profile/boot-definition* # commit-buffer
UCS-A /org/service-profile/boot-definition #
```

What to Do Next

Configure one or more of the following boot options for the boot definition and set their boot order:

- **LAN Boot** —Boots from a centralized provisioning server. It is frequently used to install operating systems on a server from that server.

If you choose the LAN Boot option, continue to [Configuring a LAN Boot for a Service Profile Boot Definition](#) , on page 481.

- **Storage Boot** — Boots from an operating system image on the SAN. You can specify a primary and a secondary SAN boot. If the primary boot fails, the server attempts to boot from the secondary.

We recommend that you use a SAN boot, because it offers the most service profile mobility within the system. If you boot from the SAN, when you move a service profile from one server to another, the new server boots from exactly the same operating system image. Therefore, the new server appears to be exactly the same server to the network.

If you choose the Storage Boot option, continue to [Configuring a Storage Boot for a Service Profile Boot Definition](#), on page 482.

- **Virtual Media Boot** — Mimics the insertion of a physical CD into a server. It is typically used to manually install operating systems on a server.

If you choose the Virtual Media boot option, continue to [Configuring a Virtual Media Boot for a Service Profile Boot Definition](#), on page 483.

Configuring a LAN Boot for a Service Profile Boot Definition

Before You Begin

Configure a boot definition for a service profile.

Procedure

	Command or Action	Purpose
Step 1	UCS-A# scope org <i>org-name</i>	Enters organization mode for the specified organization. To enter the root organization mode, type <code>/</code> as the <i>org-name</i> .
Step 2	UCS-A /org # scope service-profile <i>profile-name</i>	Enters organization service profile mode for the specified service profile.
Step 3	UCS-A /org/service-profile # scope boot-definition	Enters organization service profile boot definition mode.
Step 4	UCS-A /org/service-profile/boot-definition # create lan	Creates a LAN boot for the service profile boot definition and enters service profile boot definition LAN mode.
Step 5	UCS-A /org/service-profile/boot-definition/lan # set order {1 2 3 4}	Specifies the boot order for the LAN boot.
Step 6	UCS-A /org/service-profile/boot-definition/lan # create path {primary secondary}	Creates a primary or secondary LAN boot path and enters service profile boot definition LAN path mode.
Step 7	UCS-A /org/service-profile/boot-definition/lan/path # set vnic <i>vnic-name</i>	Specifies the vNIC to use for the LAN image path.

	Command or Action	Purpose
Step 8	UCS-A /org/service-profile/boot-definition/lan/path # commit-buffer	Commits the transaction to the system configuration.

The following example enters the service profile named ServInst90, creates a LAN boot for the service profile boot definition, sets the boot order to 2, creates a primary path, and commits the transaction:

```
UCS-A# scope org /
UCS-A /org* # scope service-profile ServInst90
UCS-A /org/service-profile* # scope boot-definition
UCS-A /org/service-profile/boot-definition* # create lan
UCS-A /org/service-profile/boot-definition/lan* # set order 2
UCS-A /org/service-profile/boot-definition/lan* # create path primary
UCS-A /org/service-profile/boot-definition/lan/path* # set vnic vnic3
UCS-A /org/service-profile/boot-definition/lan/path* # commit-buffer
UCS-A /org/service-profile/boot-definition/lan/path #
```

Configuring a Storage Boot for a Service Profile Boot Definition

Before You Begin

Configure a boot definition for a service profile.

Procedure

	Command or Action	Purpose
Step 1	UCS-A# scope org org-name	Enters organization mode for the specified organization. To enter the root organization mode, type / as the <i>org-name</i> .
Step 2	UCS-A /org # scope service-profile profile-name	Enters organization service profile mode for the specified service.
Step 3	UCS-A /org/service-profile # scope boot-definition	Enters organization service profile boot definition mode.
Step 4	UCS-A /org/service-profile/boot-definition # create storage	Creates a storage boot for the service profile boot definition and enters service profile boot definition storage mode.
Step 5	UCS-A /org/service-profile/boot-definition/storage # set order {1 2 3 4}	Specifies the boot order for the storage boot.
Step 6	UCS-A /org/service-profile/boot-definition/storage # create {local san-image {primary secondary}}	Creates a local storage boot or a SAN image boot. If a SAN image boot is created, it enters service profile boot definition storage SAN image mode.

	Command or Action	Purpose
Step 7	UCS-A /org/service-profile/boot-definition/storage/san-image # create path {primary secondary}	Creates a primary or secondary SAN image path and enters service profile boot definition storage SAN image path mode. The use of the terms primary or secondary boot devices does not imply a boot order. The effective order of boot devices within the same device class is determined by PCIe bus scan order.
Step 8	UCS-A /org/service-profile/boot-definition/storage/san-image/path # set lun lun-num	Specifies the LUN used for the SAN image path.
Step 9	UCS-A /org/service-profile/boot-definition/storage/san-image/path # set vhba vhba-name	Specifies the vHBA used for the SAN image path.
Step 10	UCS-A /org/service-profile/boot-definition/storage/san-image/path # set wwn wwn-num	Specifies the WWN used for the SAN image path.
Step 11	UCS-A /org/service-profile/boot-definition/storage/san-image/path # commit-buffer	Commits the transaction to the system configuration.

The following example enters the service profile named ServInst90, creates a storage boot for the service profile boot definition, sets the boot order to 2, creates a primary path, and commits the transaction:

```
UCS-A# scope org /
UCS-A /org* # scope service-profile ServInst90
UCS-A /org/service-profile* # scope boot-definition
UCS-A /org/service-profile/boot-definition* # create storage
UCS-A /org/service-profile/boot-definition/storage* # create san-image primary
UCS-A /org/service-profile/boot-definition/storage* # set order 2
UCS-A /org/service-profile/boot-definition/storage/san-image* # create path primary
UCS-A /org/service-profile/boot-definition/storage/san-image/path* # set lun 27512
UCS-A /org/service-profile/boot-definition/storage/san-image/path* # set vhba vhba3
UCS-A /org/service-profile/boot-definition/storage/san-image/path* # set wwn
20:00:00:00:20:00:00:23
UCS-A /org/service-profile/boot-definition/storage/san-image/path* # commit-buffer
UCS-A /org/service-profile/boot-definition/storage/san-image/path #
```

Configuring a Virtual Media Boot for a Service Profile Boot Definition

Before You Begin

Configure a boot definition for a service profile.

Procedure

	Command or Action	Purpose
Step 1	UCS-A# scope org <i>org-name</i>	Enters organization mode for the specified organization. To enter the root organization mode, type / as the <i>org-name</i> .
Step 2	UCS-A /org # scope service-profile <i>profile-name</i>	Enters organization service profile mode for the specified service.
Step 3	UCS-A /org/service-profile # scope boot-definition	Enters organization service profile boot definition mode.
Step 4	UCS-A /org/service-profile/boot-definition # create virtual-media { read-only read-write }	Creates a read-only or read-write virtual media boot for the service profile boot definition and enters service profile boot definition virtual media mode.
Step 5	UCS-A /org/service-profile/boot-definition/virtual-media # set order { 1 2 3 4 }	Specifies the boot order for the virtual media boot.
Step 6	UCS-A /org/service-profile/boot-definition/virtual-media # commit-buffer	Commits the transaction to the system configuration.

The following example enters the service profile named ServInst90, creates a virtual media boot with read-only privileges for the service profile boot definition, sets the boot order to 3, and commits the transaction:

```
UCS-A# scope org /
UCS-A /org* # scope service-profile ServInst90
UCS-A /org/service-profile* # scope boot-definition
UCS-A /org/service-profile/boot-definition* # create virtual-media read-only
UCS-A /org/service-profile/boot-definition/virtual-media* # set order 1
UCS-A /org/service-profile/boot-definition/virtual-media* # commit-buffer
UCS-A /org/service-profile/boot-definition/virtual-media #
```

Deleting a Boot Definition for a Service Profile**Procedure**

	Command or Action	Purpose
Step 1	UCS-A# scope org <i>org-name</i>	Enters organization mode for the specified organization. To enter the root organization mode, type / as the <i>org-name</i> .
Step 2	UCS-A /org # scope service-profile <i>profile-name</i>	Enters organization service profile mode for the the specified service.

	Command or Action	Purpose
Step 3	UCS-A /org/service-profile # delete boot-definition	Deletes the boot definition for the service profile.
Step 4	UCS-A /org/service-profile # commit-buffer	Commits the transaction to the system configuration.

The following example deletes the boot definition for a service profile and commits the transaction:

```
UCS-A# scope org /
UCS-A /org # scope service-profile ServInst90
UCS-A /org/service-profile # delete boot-definition
UCS-A /org/service-profile* # commit-buffer
UCS-A /org/service-profile #
```

Service Profiles and Service Profile Template Management

Associating a Service Profile with a Blade Server or Server Pool

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

Procedure

	Command or Action	Purpose
Step 1	UCS-A# scope org org-name	Enters organization mode for the specified organization. To enter the root organization mode, type / as the <i>org-name</i> .
Step 2	UCS-A /org # scope service-profile profile-name	Enters organization service profile mode for the specified service profile.
Step 3	UCS-A /org/service-profile # associate {server chassis-id / slot-id server-pool pool-name qualifier} [restrict-migration]	Associates the service profile with a single server, or to the specified server pool with the specified server pool policy qualifications. Adding the optional the restrict-migration command prevents the service profile from being migrated to another server.
Step 4	UCS-A /org/service-profile # commit-buffer	Commits the transaction to the system configuration.

The following example associates the service profile named ServProf34 with the server in slot 4 of chassis 1 and commits the transaction:

```
UCS-A# scope org /
UCS-A /org* # scope service-profile ServProf34
UCS-A /org/service-profile* # associate server 1/4
```

```
UCS-A /org/service-profile* # commit-buffer
UCS-A /org/service-profile #
```

Associating a Service Profile with a Rack Server

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

Procedure

	Command or Action	Purpose
Step 1	UCS-A# scope org <i>org-name</i>	Enters organization mode for the specified organization. To enter the root organization mode, type / as the <i>org-name</i> .
Step 2	UCS-A /org # scope service-profile <i>profile-name</i>	Enters organization service profile mode for the specified service profile.
Step 3	UCS-A /org/service-profile # associate server <i>serv-id</i> [restrict-migration]	Associates the service profile with the specified rack server. Adding the optional the restrict-migration command prevents the service profile from being migrated to another server.
Step 4	UCS-A /org/service-profile # commit-buffer	Commits the transaction to the system configuration.

The following example associates the service profile named ServProf34 with the rack server 1 and commits the transaction:

```
UCS-A# scope org /
UCS-A /org* # scope service-profile ServProf34
UCS-A /org/service-profile* # associate server 1
UCS-A /org/service-profile* # commit-buffer
UCS-A /org/service-profile #
```

Disassociating a Service Profile from a Server or Server Pool

This procedure covers disassociating a service profile from a blade server, rack server, or server pool.

Procedure

	Command or Action	Purpose
Step 1	UCS-A# scope org <i>org-name</i>	Enters organization mode for the specified organization. To enter the root organization mode, type / as the <i>org-name</i> .

	Command or Action	Purpose
Step 2	UCS-A /org # scope service-profile <i>profile-name</i>	Enters organization service profile mode for the specified service profile.
Step 3	UCS-A /org/service-profile # disassociate	Disassociates the service profile from the server or server pool.
Step 4	UCS-A /org/service-profile # commit-buffer	Commits the transaction to the system configuration.

The following example disassociates the service profile named ServProf34 from the server to which it was associated and commits the transaction:

```
UCS-A# scope org /
UCS-A /org* # scope service-profile ServProf34
UCS-A /org/service-profile # disassociate
UCS-A /org/service-profile* # commit-buffer
UCS-A /org/service-profile #
```

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

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

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

Procedure

	Command or Action	Purpose
Step 1	UCS-A# scope org <i>org-name</i>	Enters the command mode for the organization for which you want to reset the UUID. If the system does not include multi-tenancy, type / as the <i>org-name</i> to enter the root organization.
Step 2	UCS-A /org # scope service-profile <i>profile-name</i>	Enters the service profile that requires the UUID for the associated server to be reset to a different UUID suffix pool.
Step 3	UCS-A /org/service-profile # set identity dynamic-uuid derived	Specifies that the service profile will obtain a UUID dynamically from a pool.

	Command or Action	Purpose
Step 4	UCS-A /org/service-profile # commit-buffer	Commits the transaction to the system configuration.

This example resets the UUID of a service profile to a different UUID suffix pool:

```
UCS-A# scope org /
UCS-A /org # scope service-profile ServInst90
UCS-A /org/service-profile # set identity dynamic-uuid derived
UCS-A /org/service-profile* # commit-buffer
UCS-A /org/service-profile #
```

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

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

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

Procedure

	Command or Action	Purpose
Step 1	UCS-A# scope org <i>org-name</i>	Enters the command mode for the organization that contains the service profile for which you want to reset the MAC address. If the system does not include multi-tenancy, type <i>/</i> as the <i>org-name</i> to enter the root organization.
Step 2	UCS-A /org # scope service-profile <i>profile-name</i>	Enters the command mode for the service profile that requires the MAC address of the associated server to be reset to a different MAC address.
Step 3	UCS-A /org/service-profile # scope vnic <i>vnic-name</i>	Enters the command mode for the vNIC for which you want to reset the MAC address.
Step 4	UCS-A /org/service-profile/vnic # set identity dynamic-mac derived	Specifies that the vNIC will obtain a MAC address dynamically from a pool.
Step 5	UCS-A /org/service-profile/vnic # commit-buffer	Commits the transaction to the system configuration.

This example resets the MAC address of a vNIC in a service profile:

```
UCS-A# scope org /
UCS-A /org # scope service-profile ServInst90
UCS-A /org/service-profile # scope vnic dynamic-prot-001
UCS-A /org/service-profile/vnic # set identity dynamic-mac derived
UCS-A /org/service-profile/vnic* # commit-buffer
UCS-A /org/service-profile/vnic #
```

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

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

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

Procedure

	Command or Action	Purpose
Step 1	UCS-A# scope org <i>org-name</i>	Enters the command mode for the organization that contains the service profile for which you want to reset the WWPN. If the system does not include multi-tenancy, type / as the <i>org-name</i> to enter the root organization.
Step 2	UCS-A /org # scope service-profile <i>profile-name</i>	Enters the service profile of the vHBA for which you want to reset the WWPN.
Step 3	UCS-A /org/service-profile # scope vhma <i>vhba-name</i>	Enters the command mode for vHBA for which you want to reset the WWPN.
Step 4	UCS-A /org/service-profile/vhma # set identity dynamic-wwpn derived	Specifies that the vHBA will obtain a WWPN dynamically from a pool.
Step 5	UCS-A /org/service-profile/vhma # commit-buffer	Commits the transaction to the system configuration.

This example resets the WWPN of a vHBA in a service profile:

```
UCS-A# scope org /
UCS-A /org # scope service-profile ServInst90
UCS-A /org/service-profile # scope vhma vhma3
UCS-A /org/service-profile/vhma # set identity dynamic-wwpn derived
UCS-A /org/service-profile/vhma* # commit-buffer
UCS-A /org/service-profile/vhma #
```




CHAPTER 32

Managing Power in Cisco UCS

This chapter includes the following sections:

- [Power Management in Cisco UCS](#) , page 491
- [Rack Server Power Management](#), page 491
- [Power Management Precautions](#), page 491
- [Configuring the Power Policy](#), page 492
- [Configuring the Global Cap Policy](#), page 493
- [Configuring Policy-Driven Chassis Group Power Capping](#), page 493
- [Configuring Manual Blade-Level Power Capping](#), page 497

Power Management in Cisco UCS

You can manage power through Cisco UCS Manager by configuring any of the following features:

- Power supply redundancy for all chassis in a Cisco UCS domain
- Policy-driven chassis-level power capping
- Manual blade-level power capping

Rack Server Power Management

Power capping is not supported for rack servers.

Power Management Precautions

If the CIMC is reset, the power monitoring functions of Cisco UCS become briefly unavailable for as long as it takes for the CIMC to reboot. While this usually only takes 20 seconds, there is a possibility that the peak power cap could be exceeded during that time. To avoid exceeding the configured power cap in a very low power-capped environment, consider staggering the rebooting or activation of CIMCs.

Configuring the Power Policy

Power Policy

The power policy is a global policy that specifies the redundancy for power supplies in all chassis in the Cisco UCS domain. This policy is also known as the PSU policy.

For more information about power supply redundancy, see *Cisco UCS 5108 Server Chassis Hardware Installation Guide*.

Configuring the Power Policy

Procedure

	Command or Action	Purpose
Step 1	UCS-A# scope org <i>org-name</i>	Enters organization mode for the specified organization. To enter the root organization mode, type <i>/</i> as the <i>org-name</i> .
Step 2	UCS-A /org # scope psu-policy	Enters PSU policy mode.
Step 3	UCS-A /org/psu-policy # set redundancy { grid n-plus-1 non-redund }	<p>Specifies one of the following redundancy types:</p> <ul style="list-style-type: none"> • grid —Provides power redundancy when two power sources are used to power the chassis. If one power source fails, the surviving power supplies on the other power circuit continue to provide power to the chassis. • n-plus-1 —Balances the power load for the chassis across the number of power supplies needed to satisfy non-redundancy plus one additional power supply for redundancy. If any additional power supplies are installed, they are recognized and powered off. • non-redund —Balances the power load for the chassis evenly across all installed power supplies. <p>For more information about power redundancy, see the <i>Cisco UCS 5108 Server Chassis Installation Guide</i>.</p>
Step 4	UCS-A /org/psu-policy # commit-buffer	Commits the transaction to the system configuration.

The following example configures the power policy to use grid redundancy and commits the transaction:

```
UCS-A# scope org /
UCS-A /org # scope psu-policy
UCS-A /org/psu-policy # set redundancy grid
UCS-A /org/psu-policy* # commit-buffer
UCS-A /org/psu-policy #
```

Configuring the Global Cap Policy

Global Cap Policy

The global cap policy is a global policy that specifies whether policy-driven chassis group power capping or manual blade-level power capping will be applied to all servers in a chassis.

We recommend that you use the default power capping method: policy-driven chassis group power capping.



Important

Any change to the manual blade-level power cap configuration will result in the loss of any groups or configuration options set for policy-driven chassis group power capping.

Configuring the Global Cap Policy

Procedure

	Command or Action	Purpose
Step 1	UCS-A# scope power-cap-mgmt	Enters power cap management mode.
Step 2	UCS-A /power-cap-mgmt # set cap-policy { manual-blade-level-cap policy-driven-chassis-group-cap }	Sets the global cap policy to the specified power cap management mode. By default, the global cap policy is set to policy driven chassis group cap.
Step 3	UCS-A /power-cap-mgmt # commit-buffer	Commits the transaction to the system configuration.

The following example sets the global cap policy to manual blade power cap and commits the transaction:

```
UCS-A# scope power-cap-mgmt
UCS-A /power-cap-mgmt # set cap-policy manual-blade-level-cap
UCS-A /power-cap-mgmt* # commit-buffer
UCS-A /power-cap-mgmt #
```

Configuring Policy-Driven Chassis Group Power Capping

Policy-Driven Chassis Group Power Capping

When policy-driven power chassis group power capping is selected in the global cap policy, Cisco UCS can maintain the oversubscription of servers without risking costly power failures. This is achieved through a two-tier process. At the chassis level, Cisco UCS divides the amount of power available between members of the power group. At the blade level, the amount of power allotted to a chassis is divided between blades based on priority.

Each time a service profile is associated or disassociated, UCS Manager recalculates the power allotment for each blade server within the chassis. If necessary, power from lower-priority service profiles is redistributed to higher-priority service profiles.

UCS power groups cap power in less than one second in order to safely protect data center circuit breakers. A blade must stay at its cap for 20 seconds before the chassis power distribution is optimized. This is intentionally carried out over a slower timescale to prevent reacting to transient spikes in demand.



Note The system reserves enough power to boot a server in each slot, even if that slot is empty. This reserved power cannot be leveraged by servers requiring more power. Blades that fail to comply with the power cap are penalized or shut down.

Power Groups

A power group is a set of chassis that all draw power from the same power distribution unit (PDU). In Cisco UCS Manager, you can create power groups that include one or more chassis and then set a peak power cap in AC watts for that power grouping.

Instituting power capping at the chassis level requires the following:

- IOM, CIMC, and BIOS version 1.4 or higher
- 2 PSUs

The peak power cap is a static value that represents the maximum power available to all blade servers within a given power group. If you add or remove a blade from a power group, but do not manually modify the peak power value, the power group adjusts the peak power cap to accommodate the basic power-on requirements of all blades within that power group.

A minimum of 3788 AC watts should be set for each chassis. This converts to 3400 watts of DC power, which is the minimum amount of power required to power a fully-populated chassis.

If insufficient power is available, Cisco UCS Manager raises an alert.

Once a chassis is added to a power group, every service profile associated with that chassis also becomes part of that power group. Similarly, if you add a new blade to a chassis, that blade inherently becomes part of the chassis' power group.



Note Creating a power group is not the same as creating a server pool. However, you can populate a server pool with members of the same power group by creating a power qualifier and adding it to server pool policy.

Creating a Power Group

Before You Begin

Make sure the global power allocation policy is set to Policy Driven Chassis Group Cap.

Procedure

	Command or Action	Purpose
Step 1	UCS-A# scope power-cap-mgmt	Enters power cap management mode.
Step 2	UCS-A /power-cap-mgmt # create power-group <i>power-group-name</i>	Creates a power group and enters power group mode.
Step 3	UCS-A /power-cap-mgmt/power-group # set peak { <i>peak-num</i> disabled uninitialized }	Specifies the maximum peak power (in watts) available to the power group.
Step 4	UCS-A /power-cap-mgmt/power-group # create chassis <i>chassis-id</i>	Adds the specified chassis to the power group and enters power group chassis mode.
Step 5	UCS-A /power-cap-mgmt/power-group/chassis # commit-buffer	Commits the transaction to the system configuration.

The following example creates a power group called powergroup1, specifies the maximum peak power for the power group (10000 watts), adds chassis 1 to the group, and commits the transaction:

```
UCS-A# scope power-cap-mgmt
UCS-A /power-cap-mgmt # create power-group powergroup1
UCS-A /power-cap-mgmt/power-group* # set peak 10000
UCS-A /power-cap-mgmt/power-group* # create chassis 1
UCS-A /power-cap-mgmt/power-group/chassis* # commit-buffer
UCS-A /power-cap-mgmt/power-group/chassis #
```

Deleting a Power Group**Procedure**

	Command or Action	Purpose
Step 1	UCS-A# scope power-cap-mgmt	Enters power cap management mode.
Step 2	UCS-A /power-cap-mgmt # delete power-group <i>power-group-name</i>	Deletes the specified power group.
Step 3	UCS-A /power-cap-mgmt/power-group/chassis # commit-buffer	Commits the transaction to the system configuration.

The following example deletes a power group called powergroup1 and commits the transaction:

```
UCS-A# scope power-cap-mgmt
UCS-A /power-cap-mgmt # delete power-group powergroup1
UCS-A /power-cap-mgmt* # commit-buffer
UCS-A /power-cap-mgmt #
```

Power Control Policy

Cisco UCS uses the priority set in the power control policy, along with the blade type and configuration, to calculate the initial power allocation for each blade within a chassis. During normal operation, the active blades within a chassis can borrow power from idle blades within the same chassis. If all blades are active and reach the power cap, service profiles with higher priority power control policies take precedence over service profiles with lower priority power control policies.

Priority is ranked on a scale of 1-10, where 1 indicates the highest priority and 10 indicates lowest priority. The default priority is 5.

For mission-critical application a special priority called no-cap is also available. Setting the priority to no-cap prevents Cisco UCS from leveraging unused power from that particular blade server. The server is allocated the maximum amount of power that that blade can reach.



Note

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

Creating a Power Control Policy

Procedure

	Command or Action	Purpose
Step 1	UCS-A# scope org <i>org-name</i>	Enters organization mode for the specified organization. To enter the root organization mode, type / as the org-name.
Step 2	UCS-A /org# create power-control-policy <i>power-control-pol-name</i>	Creates a power control policy and enters power control policy mode.
Step 3	UCS-A /org/power-control-policy # set priority { <i>priority-num</i> no-cap }	Specifies the priority for the power control policy.
Step 4	UCS-A /org/power-control-policy # commit-buffer	Commits the transaction to the system configuration.

The following example creates a power control policy called powerpolicy15, sets the priority at level 2, and commits the transaction:

```
UCS-A# scope org /
UCS-A /org # create power-control-policy powerpolicy15
UCS-A /org/power-control policy* # set priority 2
UCS-A /org/power-control policy* # commit-buffer
UCS-A /org/power-control policy #
```

What to Do Next

Include the power control policy in a service profile.

Deleting a Power Control Policy

Procedure

	Command or Action	Purpose
Step 1	UCS-A# scope org <i>org-name</i>	Enters organization mode for the specified organization. To enter the root organization mode, type / as the org-name.
Step 2	UCS-A /org # delete power-control-policy <i>power-control-pol-name</i>	Deletes the specified power control policy.
Step 3	UCS-A /org # commit-buffer	Commits the transaction to the system configuration.

The following example deletes a power control policy called powerpolicy15 and commits the transaction:

```
UCS-A# scope org /
UCS-A /org # delete power-control-policy powerpolicy15
UCS-A /org* # commit-buffer
UCS-A /org #
```

Configuring Manual Blade-Level Power Capping

Manual Blade-Level Power Capping

When manual blade-level power capping is configured in the global cap policy, you can set a power cap for each blade server in a Cisco UCS domain.

The following configuration options are available:

Enabled

You can specify the maximum amount of power that the server can consume at one time. This maximum can be any amount between 0 watts and 1100 watts.

Disabled

No power usage limitations are imposed upon the server. The server can use as much power as it requires.

If the server encounters a spike in power usage that meets or exceeds the maximum configured for the server, Cisco UCS Manager does not disconnect or shut down the server. Instead, Cisco UCS Manager reduces the power that is made available to the server. This reduction can slow down the server, including a reduction in CPU speed.

Setting the Blade-Level Power Cap for a Server

Before You Begin

Make sure the global power allocation policy is set to Manual Blade Level Cap.

Procedure

	Command or Action	Purpose
Step 1	UCS-A# scope server <i>chassis-id / server-id</i>	Enters chassis server mode for the specified server.
Step 2	UCS-A /chassis/server # set power-budget committed { disabled <i>watts</i> }	Commits the server to one of the following power usage levels: <ul style="list-style-type: none"> • disabled —Does not impose any power usage limitations on the server. • <i>watts</i> —Allows you to specify the upper level for power usage by the server. If you choose this setting, enter the maximum number of watts that the server can use. The range is 0 to 10000000 watts.
Step 3	UCS-A /chassis/server # commit-buffer	Commits the transaction to the system configuration.
Step 4	UCS-A /chassis/server # show power-budget	(Optional) Displays the power usage level setting.

The following example limits the power usage for a server to 1000 watts and commits the transaction:

```
UCS-A# scope server 2/4
UCS-A /chassis/server # set power-budget committed 1000
UCS-A /chassis/server* # commit-buffer
UCS-A /chassis/server # show power-budget
Power Budget:
  Committed (W): 1100
  Oper Committed (W): Disabled

UCS-A /chassis/server #
```

Viewing the Blade-Level Power Cap

Procedure

	Command or Action	Purpose
Step 1	UCS-A# scope server <i>chassis-id / server-id</i>	Enters chassis server mode for the specified server.

	Command or Action	Purpose
Step 2	UCS-A /chassis/server # show stats	Displays the power usage statistics collected for the server.

The following example shows the server power usage:

```
UCS-A# scope server 2/4
UCS-A /chassis/server # show stats

Mb Power Stats:
  Time Collected: 2010-04-15T21:18:04.992
  Monitored Object: sys/chassis-1/blade-2/board
  Suspect: No
  Consumed Power (W): 118.285194
  Input Voltage (V): 11.948000
  Input Current (A): 9.900000
  Thresholded: Input Voltage Min

UCS-A /chassis/server #
```




PART VI

System Management

- [Managing Time Zones, page 503](#)
- [Managing the Chassis, page 507](#)
- [Managing Blade Servers, page 513](#)
- [Managing Rack-Mount Servers, page 525](#)
- [Managing the I/O Modules, page 539](#)
- [Backing Up and Restoring the Configuration, page 541](#)
- [Recovering a Lost Password, page 555](#)



CHAPTER 33

Managing Time Zones

This chapter includes the following sections:

- [Time Zones, page 503](#)
- [Setting the Time Zone, page 503](#)
- [Adding an NTP Server, page 505](#)
- [Deleting an NTP Server, page 506](#)
- [Setting the System Clock Manually, page 506](#)

Time Zones

Cisco UCS requires a domain-specific time zone setting and an NTP server to ensure the correct time display in Cisco UCS Manager. If you do not configure both of these settings in a Cisco UCS domain, the time does not display correctly.

Setting the Time Zone

Procedure

	Command or Action	Purpose
Step 1	UCS-A# scope system	Enters system mode.
Step 2	UCS-A /system # scope services	Enters system services mode.
Step 3	UCS-A /system/services # set timezone	At this point, you are prompted to enter a number corresponding to your continent, country, and time zone region. Enter the appropriate information at each prompt. When you have finished specifying the location information, you are prompted to confirm that the

	Command or Action	Purpose
		correct time zone information is being set. Enter 1 (yes) to confirm, or 2 (no) to cancel the operation.
Step 4	UCS-A /system/services # commit-buffer	Commits the transaction to the system configuration.
Step 5	UCS-A /system/services # exit	Enters system mode.
Step 6	UCS-A /system/services # exit	Enters EXEC mode.
Step 7	UCS-A /system/services # show timezone	Displays the configured timezone.

The following example configures the timezone to the Pacific time zone region, commits the transaction, and displays the configured timezone:

```
UCS-A# scope system
UCS-A /system # scope services
UCS-A /system/services # set timezone
Please identify a location so that time zone rules can be set correctly.
Please select a continent or ocean.
1) Africa           4) Arctic Ocean      7) Australia        10) Pacific Ocean
2) Americas        5) Asia              8) Europe
3) Antarctica      6) Atlantic Ocean   9) Indian Ocean
#? Arctic ocean
Please enter a number in range.
#? 2
Please select a country.
1) Anguilla         18) Ecuador          35) Paraguay
2) Antigua & Barbuda 19) El Salvador      36) Peru
3) Argentina       20) French Guiana   37) Puerto Rico
4) Aruba            21) Greenland        38) St Kitts & Nevis
5) Bahamas         22) Grenada          39) St Lucia
6) Barbados        23) Guadeloupe      40) St Pierre & Miquelon
7) Belize          24) Guatemala       41) St Vincent
8) Bolivia         25) Guyana           42) Suriname
9) Brazil           26) Haiti            43) Trinidad & Tobago
10) Canada          27) Honduras        44) Turks & Caicos Is
11) Cayman Islands 28) Jamaica          45) United States
12) Chile           29) Martinique       46) Uruguay
13) Colombia       30) Mexico           47) Venezuela
14) Costa Rica     31) Montserrat      48) Virgin Islands (UK)
15) Cuba           32) Netherlands Antilles 49) Virgin Islands (US)
16) Dominica       33) Nicaragua
17) Dominican Republic 34) Panama
#? 45
Please select one of the following time zone regions.
1) Eastern Time
2) Eastern Time - Michigan - most locations
3) Eastern Time - Kentucky - Louisville area
4) Eastern Time - Kentucky - Wayne County
5) Eastern Standard Time - Indiana - most locations
6) Eastern Standard Time - Indiana - Crawford County
7) Eastern Standard Time - Indiana - Starke County
8) Eastern Standard Time - Indiana - Switzerland County
9) Central Time
10) Central Time - Michigan - Wisconsin border
11) Central Time - North Dakota - Oliver County
12) Mountain Time
13) Mountain Time - south Idaho & east Oregon
14) Mountain Time - Navajo
15) Mountain Standard Time - Arizona
16) Pacific Time
```

```

17) Alaska Time
18) Alaska Time - Alaska panhandle
19) Alaska Time - Alaska panhandle neck
20) Alaska Time - west Alaska
21) Aleutian Islands
22) Hawaii
#? 16

```

The following information has been given:

```

United States
Pacific Time

```

```

Therefore timezone 'America/Los_Angeles' will be set.
Local time is now:      Fri May 15 07:39:25 PDT 2009.
Universal Time is now:  Fri May 15 14:39:25 UTC 2009.
Is the above information OK?
1) Yes
2) No
#? 1
UCS-A /system/services* # commit-buffer
UCS-A /system/services # exit
UCS-A /system # exit
UCS-A# show timezone
Timezone: America/Los_Angeles (Pacific Time)
UCS-A#

```

Adding an NTP Server

Procedure

	Command or Action	Purpose
Step 1	UCS-A# scope system	Enters system mode.
Step 2	UCS-A /system # scope services	Enters system services mode.
Step 3	UCS-A /system/services # create ntp-server {hostname ip-addr}	Configures the system to use the NTP server with the specified hostname or IP address.
Step 4	UCS-A /system/services # commit-buffer	Commits the transaction to the system configuration.

The following example configures an NTP server with the IP address 192.168.200.101 and commits the transaction:

```

UCS-A# scope system
UCS-A /system # scope services
UCS-A /system/services # create ntp-server 192.168.200.101
UCS-A /system/services* # commit-buffer
UCS-A /system/services #

```

Deleting an NTP Server

Procedure

	Command or Action	Purpose
Step 1	UCS-A# scope system	Enters system mode.
Step 2	UCS-A /system # scope services	Enters system services mode.
Step 3	UCS-A /system/services # delete ntp-server <i>{hostname ip-addr}</i>	Deletes the NTP server with the specified IP address.

The following example deletes the NTP server with the IP address 192.168.200.101 and commits the transaction:

```
UCS-A# scope system
UCS-A /system # scope services
UCS-A /system/services # delete ntp-server 192.168.200.101
UCS-A /system/services* # commit-buffer
UCS-A /system/services #
```

Setting the System Clock Manually

System clock modifications take effect immediately.

Procedure

	Command or Action	Purpose
Step 1	UCS-A# scope system	Enters system mode.
Step 2	UCS-A /system # scope services	Enters system services mode.
Step 3	UCS-A /system/services # set clock <i>mon date</i> <i>year hour min sec</i>	Configures the system clock.

The following example configures the system clock and commits the transaction:

```
UCS-A# scope system
UCS-A /system # scope services
UCS-A /system/services # set clock apr 14 2010 15 27 00
UCS-A /system/services #
```



CHAPTER 34

Managing the Chassis

This chapter includes the following sections:

- [Guidelines for Removing and Decommissioning Chassis, page 507](#)
- [Acknowledging a Chassis, page 508](#)
- [Decommissioning a Chassis, page 508](#)
- [Removing a Chassis, page 509](#)
- [Recommissioning a Chassis, page 509](#)
- [Renumbering a Chassis, page 510](#)
- [Toggling the Locator LED, page 512](#)

Guidelines for Removing and Decommissioning Chassis

Consider the following guidelines when deciding whether to remove or decommission a chassis using Cisco UCS Manager:

Decommissioning a Chassis

Decommissioning is performed when a chassis is physically present and connected but you want to temporarily remove it from the configuration. Because it is expected that a decommissioned chassis will be eventually recommissioned, a portion of the chassis' information is retained by Cisco UCS Manager for future use.

Removing a Chassis

Removing is performed when you physically remove a chassis from the system. Once the physical removal of the chassis is completed, the configuration for that chassis can be removed in Cisco UCS Manager.



Note

You cannot remove a chassis from Cisco UCS Manager if it is physically present and connected.

If you need to add a removed chassis back to the configuration, it must be reconnected and then rediscovered. During rediscovery Cisco UCS Manager will assign the chassis a new ID that may be different from ID that it held before.

Acknowledging a Chassis

Perform the following procedure if you increase or decrease the number of links that connect the chassis to the fabric interconnect. Acknowledging the chassis ensures that Cisco UCS Manager is aware of the change in the number of links and that traffics flows along all available links.

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

Procedure

	Command or Action	Purpose
Step 1	UCS-A# acknowledge chassis <i>chassis-num</i>	Acknowledges the specified chassis.
Step 2	UCS-A# commit-buffer	Commits the transaction to the system configuration.

The following example acknowledges chassis 2 and commits the transaction:

```
UCS-A# acknowledge chassis 2
UCS-A* # commit-buffer
UCS-A #
```

Decommissioning a Chassis

Procedure

	Command or Action	Purpose
Step 1	UCS-A# decommission chassis <i>chassis-num</i>	Decommissions the specified chassis.
Step 2	UCS-A# commit-buffer	Commits the transaction to the system configuration.

The decommission may take several minutes to complete.

The following example decommissions chassis 2 and commits the transaction:

```
UCS-A# decommission chassis 2
UCS-A* # commit-buffer
UCS-A # show chassis

Chassis:
  Chassis  Overall Status  Admin State
  -----  -
          1 Operable          Acknowledged
          2 Accessibility Problem  Decommission
UCS-A #
```

Removing a Chassis

Before You Begin

Physically remove the chassis before performing the following procedure.

Procedure

	Command or Action	Purpose
Step 1	UCS-A# remove chassis <i>chassis-num</i>	Removes the specified chassis.
Step 2	UCS-A# commit-buffer	Commits the transaction to the system configuration.

The removal may take several minutes to complete.

The following example removes chassis 2 and commits the transaction:

```
UCS-A# remove chassis 2
UCS-A* # commit-buffer
UCS-A #
```

Recommissioning a Chassis

This procedure returns the chassis to the configuration and applies the chassis discovery policy to the chassis. After this procedure, you can access the chassis and any servers in it.

Before You Begin

Collect the following information about the chassis to be recommissioned:

- Vendor name
- Model name
- Serial number

Procedure

	Command or Action	Purpose
Step 1	UCS-A# recommission chassis <i>vendor-name</i> <i>model-name serial-num</i>	Recommissions the specified chassis.
Step 2	UCS-A# commit-buffer	Commits the transaction to the system configuration.

The following example recommissions a Cisco UCS 5108 chassis and commits the transaction:

```
UCS-A# show chassis

Chassis:
  Chassis      Overall Status      Admin State
  -----
  1 Accessibility Problem  Decommission

UCS-A# recommitment chassis "Cisco Systems Inc" "N20-C6508" FOX1252GNNN
UCS-A* # commit-buffer
UCS-A #
```

Renumbering a Chassis



Note

You cannot renumber a blade server through Cisco UCS Manager. The ID assigned to a blade server is determined by its physical slot in the chassis. To renumber a blade server, you must physically move the server to a different slot in the chassis.

Before You Begin

If you are swapping IDs between chassis, you must first decommission both chassis and then wait for the chassis decommission FSM to complete before proceeding with the renumbering steps.

Procedure

	Command or Action	Purpose
Step 1	UCS-A# show chassis inventory	Displays information about your chassis.
Step 2	Verify that the chassis inventory does not include the following:	<ul style="list-style-type: none"> • The chassis you want to renumber • A chassis with the number you want to use <p>If either of these chassis are listed in the chassis inventory, decommission those chassis. You must wait until the decommission FSM is complete and the chassis are not listed in the chassis inventory before continuing. This might take several minutes.</p> <p>To see which chassis have been decommissioned, issue the show chassis decommissioned command.</p>
Step 3	UCS-A# recommitment chassis <i>vendor-name model-name</i> <i>serial-num [chassis-num]</i>	Recommissions and renumbers the specified chassis.
Step 4	UCS-A# commit-buffer	Commits the transaction to the system configuration.

The following example decommissions two Cisco UCS chassis (chassis 7 and 8), switches their IDs, and commits the transaction:

```
UCS-A# show chassis inventory
```

Chassis	PID	Vendor	Serial (SN)	HW Revision
1	N20-C6508	Cisco Systems Inc	FOX1252GAAA	0
2	N20-C6508	Cisco Systems Inc	FOX1252BBBB	0
3	N20-C6508	Cisco Systems Inc	FOX1252CCCC	0
4	N20-C6508	Cisco Systems Inc	FOX1252DDDD	0
5	N20-C6508	Cisco Systems Inc	FOX1252EEEE	0
6	N20-C6508	Cisco Systems Inc	FOX1252FFFF	0
7	N20-C6508	Cisco Systems Inc	FOX1252GGGG	0
8	N20-C6508	Cisco Systems Inc	FOX1252HHHH	0
9	N20-C6508	Cisco Systems Inc	FOX1252IIII	0
10	N20-C6508	Cisco Systems Inc	FOX1252JJJJ	0
11	N20-C6508	Cisco Systems Inc	FOX1252KKKK	0
12	N20-C6508	Cisco Systems Inc	FOX1252LLLL	0
13	N20-C6508	Cisco Systems Inc	FOX1252MMMM	0
14	N20-C6508	Cisco Systems Inc	FOX1252NNNN	0

```
UCS-A# decommission chassis "Cisco Systems Inc" "N20-C6508" FOX1252HHHH
```

```
UCS-A*# commit-buffer
```

```
UCS-A# decommission chassis "Cisco Systems Inc" "N20-C6508" FOX1252IIII
```

```
UCS-A*# commit-buffer
```

```
UCS-A# show chassis inventory
```

Chassis	PID	Vendor	Serial (SN)	HW Revision
1	N20-C6508	Cisco Systems Inc	FOX1252GAAA	0
2	N20-C6508	Cisco Systems Inc	FOX1252BBBB	0
3	N20-C6508	Cisco Systems Inc	FOX1252CCCC	0
4	N20-C6508	Cisco Systems Inc	FOX1252DDDD	0
5	N20-C6508	Cisco Systems Inc	FOX1252EEEE	0
6	N20-C6508	Cisco Systems Inc	FOX1252FFFF	0
7	N20-C6508	Cisco Systems Inc	FOX1252GGGG	0
10	N20-C6508	Cisco Systems Inc	FOX1252JJJJ	0
11	N20-C6508	Cisco Systems Inc	FOX1252KKKK	0
12	N20-C6508	Cisco Systems Inc	FOX1252LLLL	0
13	N20-C6508	Cisco Systems Inc	FOX1252MMMM	0
14	N20-C6508	Cisco Systems Inc	FOX1252NNNN	0

```
UCS-A# show chassis decommissioned
```

Chassis	PID	Vendor	Serial (SN)	HW Revision
8	N20-C6508	Cisco Systems Inc	FOX1252HHHH	0
9	N20-C6508	Cisco Systems Inc	FOX1252IIII	0

```
UCS-A# recommission chassis "Cisco Systems Inc" "N20-C6508" FOX1252HHHH 9
```

```
UCS-A*# commit-buffer
```

```
UCS-A# recommission chassis "Cisco Systems Inc" "N20-C6508" FOX1252IIII 8
```

```
UCS-A*# commit-buffer
```

```
UCS-A # show chassis inventory
```

Chassis	PID	Vendor	Serial (SN)	HW Revision
1	N20-C6508	Cisco Systems Inc	FOX1252GAAA	0
2	N20-C6508	Cisco Systems Inc	FOX1252BBBB	0
3	N20-C6508	Cisco Systems Inc	FOX1252CCCC	0
4	N20-C6508	Cisco Systems Inc	FOX1252DDDD	0
5	N20-C6508	Cisco Systems Inc	FOX1252EEEE	0
6	N20-C6508	Cisco Systems Inc	FOX1252FFFF	0
7	N20-C6508	Cisco Systems Inc	FOX1252GGGG	0
8	N20-C6508	Cisco Systems Inc	FOX1252IIII	0
9	N20-C6508	Cisco Systems Inc	FOX1252HHHH	0
10	N20-C6508	Cisco Systems Inc	FOX1252JJJJ	0
11	N20-C6508	Cisco Systems Inc	FOX1252KKKK	0
12	N20-C6508	Cisco Systems Inc	FOX1252LLLL	0
13	N20-C6508	Cisco Systems Inc	FOX1252MMMM	0
14	N20-C6508	Cisco Systems Inc	FOX1252NNNN	0

Toggling the Locator LED

Turning On the Locator LED for a Chassis

Procedure

	Command or Action	Purpose
Step 1	UCS-A# scope chassis <i>chassis-num</i>	Enters chassis mode for the specified chassis.
Step 2	UCS-A /chassis # enable locator-led	Turns on the chassis locator LED.
Step 3	UCS-A /chassis # commit-buffer	Commits the transaction to the system configuration.

The following example turns on the locator LED for chassis 2 and commits the transaction:

```
UCS-A# scope chassis 2
UCS-A /chassis # enable locator-led
UCS-A /chassis* # commit-buffer
UCS-A /chassis #
```

Turning Off the Locator LED for a Chassis

Procedure

	Command or Action	Purpose
Step 1	UCS-A# scope chassis <i>chassis-num</i>	Enters chassis mode for the specified chassis.
Step 2	UCS-A /chassis # disable locator-led	Turns off the chassis locator LED.
Step 3	UCS-A /chassis # commit-buffer	Commits the transaction to the system configuration.

The following example turns off the locator LED for chassis 2 and commits the transaction:

```
UCS-A# scope chassis 2
UCS-A /chassis # disable locator-led
UCS-A /chassis* # commit-buffer
UCS-A /chassis #
```



CHAPTER **35**

Managing Blade Servers

This chapter includes the following sections:

- [Blade Server Management, page 513](#)
- [Guidelines for Removing and Decommissioning Blade Servers, page 514](#)
- [Booting a Blade Server, page 514](#)
- [Shutting Down a Blade Server, page 515](#)
- [Power Cycling a Blade Server, page 516](#)
- [Performing a Hard Reset on a Blade Server, page 516](#)
- [Avoiding Unexpected Server Power Changes, page 517](#)
- [Acknowledging a Blade Server, page 518](#)
- [Removing a Blade Server from a Chassis, page 519](#)
- [Decommissioning a Blade Server, page 519](#)
- [Turning On the Locator LED for a Blade Server, page 520](#)
- [Turning Off the Locator LED for a Blade Server, page 520](#)
- [Resetting the CMOS for a Blade Server, page 520](#)
- [Resetting the CIMC for a Blade Server, page 521](#)
- [Recovering the Corrupt BIOS on a Blade Server, page 522](#)
- [Issuing an NMI from a Blade Server, page 522](#)

Blade Server Management

You can manage and monitor all blade servers in a Cisco UCS domain through Cisco UCS Manager. Some blade server management tasks, such as changes to the power state, can be performed from the server and service profile.

The remaining management tasks can only be performed on the server.

If a blade server slot in a chassis is empty, Cisco UCS Manager provides information, errors, and faults for that slot. You can also reacknowledge the slot to resolve server mismatch errors and to have Cisco UCS Manager rediscover the blade server in the slot.

Guidelines for Removing and Decommissioning Blade Servers

Consider the following guidelines when deciding whether to remove or decommission a blade server using Cisco UCS Manager:

Decommissioning a Blade Server

Decommissioning is performed when a blade server is physically present and connected but you want to temporarily remove it from the configuration. Because it is expected that a decommissioned blade server will be eventually recommissioned, a portion of the server's information is retained by Cisco UCS Manager for future use.

Removing a Blade Server

Removing is performed when you physically remove a blade server from the server by disconnecting it from the chassis. You cannot remove a blade server from Cisco UCS Manager if it is physically present and connected to a chassis. Once the physical removal of the blade server is completed, the configuration for that blade server can be removed in Cisco UCS Manager.

During removal, active links to the blade server are disabled, all entries from databases are removed, and the server is automatically removed from any server pools that it was assigned to during discovery.



Note

Only those servers added to a server pool automatically during discovery will be removed automatically. Servers that have been manually added to a server pool have to be removed manually.

If you need to add a removed blade server back to the configuration, it must be reconnected and then rediscovered. When a server is reintroduced to Cisco UCS Manager it is treated like a new server and is subject to the deep discovery process. For this reason, it's possible that Cisco UCS Manager will assign the server a new ID that may be different from the ID that it held before.

Booting a Blade Server

Before You Begin

Associate a service profile with a blade server or server pool.

Procedure

	Command or Action	Purpose
Step 1	UCS-A# <code>scope org org-name</code>	Enters organization mode for the specified organization. To enter the root organization mode, type <code>/</code> as the <code>org-name</code> .

	Command or Action	Purpose
Step 2	UCS-A /org # scope service-profile <i>profile-name</i>	Enters organization service profile mode for the specified service profile.
Step 3	UCS-A /org/service-profile # power up	Boots the blade server associated with the service profile.
Step 4	UCS-A /org/service-profile # commit-buffer	Commits the transaction to the system configuration.

The following example boots the blade server associated with the service profile named ServProf34 and commits the transaction:

```
UCS-A# scope org /
UCS-A /org* # scope service-profile ServProf34
UCS-A /org/service-profile* # power up
UCS-A /org/service-profile* # commit-buffer
UCS-A /org/service-profile #
```

Shutting Down a Blade Server

When you use this procedure to shut down a server with an installed operating system, Cisco UCS Manager triggers the OS into a graceful shutdown sequence.

Before You Begin

Associate a service profile with a blade server or server pool.

Procedure

	Command or Action	Purpose
Step 1	UCS-A# scope org <i>org-name</i>	Enters organization mode for the specified organization. To enter the root organization mode, type / as the <i>org-name</i> .
Step 2	UCS-A /org # scope service-profile <i>profile-name</i>	Enters organization service profile mode for the specified service profile.
Step 3	UCS-A /org/service-profile # power down	Shuts down the blade server associated with the service profile.
Step 4	UCS-A /org/service-profile # commit-buffer	Commits the transaction to the system configuration.

The following example shuts down the blade server associated with the service profile named ServProf34 and commits the transaction:

```
UCS-A# scope org /
UCS-A /org # scope service-profile ServProf34
UCS-A /org/service-profile # power down
```

```
UCS-A /org/service-profile* # commit-buffer
UCS-A /org/service-profile #
```

Power Cycling a Blade Server

Procedure

	Command or Action	Purpose
Step 1	UCS-A# scope server <i>chassis-num</i> <i>/ server-num</i>	Enters chassis server mode for the specified blade server.
Step 2	UCS-A /chassis/server # cycle { cycle-immediate cycle-wait }	Power cycles the blade server. Use the cycle-immediate keyword to immediately begin power cycling the blade server; use the cycle-wait keyword to schedule the power cycle to begin after all pending management operations have completed.
Step 3	UCS-A# commit-buffer	Commits the transaction to the system configuration.

The following example immediately power cycles blade server 4 in chassis 2 and commits the transaction:

```
UCS-A# scope server 2/4
UCS-A /chassis/server # cycle cycle-immediate
UCS-A /chassis/server* # commit-buffer
UCS-A /chassis/server #
```

Performing a Hard Reset on a Blade Server

When you reset a server, Cisco UCS Manager sends a pulse on the reset line. You can choose to gracefully shut down the operating system. If the operating system does not support a graceful shut down, the server is power cycled. The option to have Cisco UCS Manager complete all management operations before it resets the server does not guarantee that these operations will be completed before the server is reset.



Note

If you are trying to boot a server from a power-down state, you should not use **Reset**.

If you continue the power-up with this process, the desired power state of the servers will become out of sync with the actual power state and the servers may unexpectedly shut down at a later time. To safely reboot the selected servers from a power-down state, click **Cancel** then select the **Boot Server** action.

Procedure

	Command or Action	Purpose
Step 1	UCS-A# scope server <i>chassis-num</i> / <i>server-num</i>	Enters chassis# server mode for the specified server.

	Command or Action	Purpose
Step 2	UCS-A /chassis/server # reset { hard-reset-immediate hard-reset-wait }	Performs a hard reset of the blade server. Use the hard-reset-immediate keyword to immediately begin hard resetting the server; use the hard-reset-wait keyword to schedule the hard reset to begin after all pending management operations have completed.
Step 3	UCS-A /server # commit-buffer	Commits the transaction to the system configuration.

The following example performs an immediate hard reset of blade server 4 in chassis 2 and commits the transaction:

```
UCS-A# scope server 2/4
UCS-A /chassis/server # reset hard-reset-immediate
UCS-A /chassis/server* # commit-buffer
UCS-A /chassis/server #
```

Avoiding Unexpected Server Power Changes

If a server is not associated with a service profile, you can use any available means to change the server power state, including the physical Power or Reset buttons on the server.

If a server is associated with, or assigned to, a service profile, you should only use the following methods to change the server power state:

- In Cisco UCS Manager GUI, go to the **General** tab for the server or the service profile associated with the server and select **Boot Server** or **Shutdown Server** from the **Actions** area.
- In Cisco UCS Manager CLI, scope to the server or the service profile associated with the server and use the **power up** or **power down** commands.



Important Do *not* use any of the following options on an associated server that is currently powered off:

- **Reset** in the GUI
- **cycle cycle-immediate** or **reset hard-reset-immediate** in the CLI
- The physical Power or Reset buttons on the server

If you reset, cycle, or use the physical power buttons on a server that is currently powered off, the server's actual power state may become out of sync with the desired power state setting in the service profile. If the communication between the server and Cisco UCS Manager is disrupted or if the service profile configuration changes, Cisco UCS Manager may apply the desired power state from the service profile to the server, causing an unexpected power change.

In Cisco UCS release 2.0.1, power synchronization issues can lead to unexpected server shut downs or restarts as shown below:

Desired Power State in Service Profile	Current Server Power State	Server Power State After Communication Is Disrupted
Up	Powered Off	Powered On
Down	Powered On	Powered Off

Starting in Cisco UCS release 2.0.2, power synchronization issues can lead to an unexpected server restart, as shown below:

Desired Power State in Service Profile	Current Server Power State	Server Power State After Communication Is Disrupted
Up	Powered Off	Powered On
Down	Powered On	Powered On Note Running servers are not shut down regardless of the desired power state in the service profile.

Acknowledging a Blade Server

Perform the following procedure if you need to have Cisco UCS Manager rediscover the server and all endpoints in the server. For example, you can use this procedure if a server is stuck in an unexpected state, such as the discovery state.

Procedure

	Command or Action	Purpose
Step 1	UCS-A# acknowledge server <i>chassis-num</i> / <i>server-num</i>	Acknowledges the specified blade server.
Step 2	UCS-A# commit-buffer	Commits the transaction to the system configuration.

The following example acknowledges server 4 in chassis 2 and commits the transaction:

```
UCS-A# acknowledge server 2/4
UCS-A* # commit-buffer
UCS-A #
```


Removing a Blade Server from a Chassis

Before You Begin

Physically remove the server from its chassis before performing the following procedure.

Procedure

	Command or Action	Purpose
Step 1	UCS-A# remove server <i>chassis-num</i> / <i>server-num</i>	Removes the specified blade server.
Step 2	UCS-A# commit-buffer	Commits the transaction to the system configuration.

The following example removes blade server 4 in chassis 2 and commits the transaction:

```
UCS-A# remove server 2/4
UCS-A* # commit-buffer
UCS-A #
```

What to Do Next

If you physically re-install the blade server, you must re-acknowledge the slot to have Cisco UCS Manager rediscover the server.

For more information, see [Acknowledging a Blade Server](#), on page 518.

Decommissioning a Blade Server

Procedure

	Command or Action	Purpose
Step 1	UCS-A# decommission server <i>chassis-num</i> / <i>server-num</i>	Decommissions the specified blade server.
Step 2	UCS-A# commit-buffer	Commits the transaction to the system configuration.

The following example decommissions blade server 4 in chassis 2 and commits the transaction:

```
UCS-A# decommission server 2/4
UCS-A* # commit-buffer
UCS-A #
```

Turning On the Locator LED for a Blade Server

Procedure

	Command or Action	Purpose
Step 1	UCS-A# scope server <i>chassis-num / server-num</i>	Enters chassis server mode for the specified chassis.
Step 2	UCS-A /chassis/server # enable locator-led	Turns on the blade server locator LED.
Step 3	UCS-A /chassis/server # commit-buffer	Commits the transaction to the system configuration.

The following example turns on the locator LED for blade server 4 in chassis 2 and commits the transaction:

```
UCS-A# scope server 2/4
UCS-A /chassis/server # enable locator-led
UCS-A /chassis/server* # commit-buffer
UCS-A /chassis/server #
```

Turning Off the Locator LED for a Blade Server

Procedure

	Command or Action	Purpose
Step 1	UCS-A# scope server <i>chassis-num / server-num</i>	Enters chassis mode for the specified chassis.
Step 2	UCS-A /chassis/server # disable locator-led	Turns off the blade server locator LED.
Step 3	UCS-A /chassis/server # commit-buffer	Commits the transaction to the system configuration.

The following example turns off the locator LED for blade server 4 in chassis 2 and commits the transaction:

```
UCS-A# scope chassis 2/4
UCS-A /chassis/server # disable locator-led
UCS-A /chassis/server* # commit-buffer
UCS-A /chassis/server #
```

Resetting the CMOS for a Blade Server

On rare occasions, troubleshooting a server may require you to reset the CMOS. This procedure is not part of the normal maintenance of a server.

Procedure

	Command or Action	Purpose
Step 1	UCS-A# scope server <i>chassis-num / server-num</i>	Enters chassis server mode for the specified chassis.
Step 2	UCS-A /chassis/server # reset-cmos	Resets the CMOS for the blade server.
Step 3	UCS-A /chassis/server # commit-buffer	Commits the transaction to the system configuration.

The following example resets the CMOS for blade server 4 in chassis 2 and commits the transaction:

```
UCS-A# scope server 2/4
UCS-A /chassis/server # reset-cmos
UCS-A /chassis/server* # commit-buffer
UCS-A /chassis/server #
```

Resetting the CIMC for a Blade Server

On rare occasions, such as an issue with the current running firmware, troubleshooting a server may require you to reset the CIMC. This procedure is not part of the normal maintenance of a server. After you reset the CIMC, the server boots with the running version of the firmware for that server.

If the CIMC is reset, the power monitoring functions of Cisco UCS become briefly unavailable for as long as it takes for the CIMC to reboot. While this usually only takes 20 seconds, there is a possibility that the peak power cap could be exceeded during that time. To avoid exceeding the configured power cap in a very low power-capped environment, consider staggering the rebooting or activation of CIMCs.

Procedure

	Command or Action	Purpose
Step 1	UCS-A# scope server <i>chassis-num / server-num</i>	Enters chassis server mode for the specified chassis.
Step 2	UCS-A /chassis/server # scope CIMC	Enters chassis server CIMC mode
Step 3	UCS-A /chassis/server/CIMC # reset	Resets the CIMC for the blade server.
Step 4	UCS-A /chassis/server/CIMC # commit-buffer	Commits the transaction to the system configuration.

The following example resets the CIMC for blade server 4 in chassis 2 and commits the transaction:

```
UCS-A# scope server 2/4
UCS-A /chassis/server # scope CIMC
UCS-A /chassis/server/cimc # reset
UCS-A /chassis/server/cimc* # commit-buffer
UCS-A /chassis/server/cimc #
```

Recovering the Corrupt BIOS on a Blade Server

On rare occasions, an issue with a blade server may require you to recover the corrupted BIOS. This procedure is not part of the normal maintenance of a server. After you recover the BIOS, the blade server boots with the running version of the firmware for that server.

Before You Begin



Important Remove all attached or mapped USB storage from a server before you attempt to recover the corrupt BIOS on that server. If an external USB drive is attached or mapped from vMedia to the server, BIOS recovery fails.

Procedure

	Command or Action	Purpose
Step 1	UCS-A# scope server <i>chassis-id / server-id</i>	Enters chassis server mode for the specified blade server in the specified chassis.
Step 2	UCS-A /chassis/server # recover-bios <i>version</i> [ignorecompcheck]	Loads and activates the specified BIOS version. To activate the firmware without making sure that it is compatible first, include the ignorecompcheck keyword. We recommend that you use this option only when explicitly directed to do so by a technical support representative.
Step 3	UCS-A /chassis/server # commit-buffer	Commits the transaction.

The following example shows how to recover the BIOS:

```
UCS-A# scope server 1/7
UCS-A /chassis/server # recover-bios S5500.0044.0.3.1.010620101125
UCS-A /chassis/server* # commit-buffer
UCS-A /chassis/server #
```

Issuing an NMI from a Blade Server

Perform the following procedure if the system remains unresponsive and you need Cisco UCS Manager to issue a Non Maskable Interrupt (NMI) to the BIOS or operating system from the CIMC. This action creates a core dump or stack trace, depending on the operating system installed on the server.

Procedure

	Command or Action	Purpose
Step 1	UCS-A# scope server [<i>chassis-num/server-num dynamic-uuid</i>]	Enters server mode for the specified server.
Step 2	UCS-A /chassis/server # diagnostic-interrupt	
Step 3	UCS-A /chassis/server* # commit-buffer	Commits any pending transactions.

The following example sends an NMI from server 4 in chassis 2 and commits the transaction:

```
UCS-A# scope server 2/4
UCS-A /chassis/server # diagnostic-interrupt
UCS-A /chassis/server* # commit-buffer
UCS-A /chassis/server #
```




CHAPTER **36**

Managing Rack-Mount Servers

This chapter includes the following sections:

- [Rack-Mount Server Management, page 525](#)
- [Guidelines for Removing and Decommissioning Rack-Mount Servers, page 526](#)
- [Booting a Rack-Mount Server, page 526](#)
- [Shutting Down a Rack-Mount Server, page 527](#)
- [Power Cycling a Rack-Mount Server, page 528](#)
- [Performing a Hard Reset on a Rack-Mount Server, page 528](#)
- [Avoiding Unexpected Server Power Changes, page 529](#)
- [Acknowledging a Rack-Mount Server, page 530](#)
- [Decommissioning a Rack-Mount Server, page 531](#)
- [Renumbering a Rack-Mount Server, page 531](#)
- [Removing a Rack-Mount Server, page 532](#)
- [Turning On the Locator LED for a Rack-Mount Server, page 533](#)
- [Turning Off the Locator LED for a Rack-Mount Server, page 534](#)
- [Resetting the CMOS for a Rack-Mount Server, page 534](#)
- [Resetting the CIMC for a Rack-Mount Server, page 535](#)
- [Recovering the Corrupt BIOS on a Rack-Mount Server, page 535](#)
- [Showing the Status for a Rack-Mount Server, page 536](#)
- [Issuing an NMI from a Rack-Mount Server, page 536](#)

Rack-Mount Server Management

You can manage and monitor all rack-mount servers that have been integrated with a Cisco UCS domain through Cisco UCS Manager. All management and monitoring features are supported for rack-mount servers except power capping. Some rack-mount server management tasks, such as changes to the power state, can

be performed from both the server and service profile. The remaining management tasks can only be performed on the server.

Cisco UCS Manager provides information, errors, and faults for each rack-mount server that it has discovered.

**Tip**

For information about how to integrate a supported Cisco UCS rack-mount server with Cisco UCS Manager, see the hardware installation guide for that server.

Guidelines for Removing and Decommissioning Rack-Mount Servers

Consider the following guidelines when deciding whether to remove or decommission a rack-mount server using Cisco UCS Manager:

Decommissioning a Rack-Mount server

Decommissioning is performed when a rack-mount server is physically present and connected but you want to temporarily remove it from the configuration. Because it is expected that a decommissioned rack-mount server will be eventually recommissioned, a portion of the server's information is retained by Cisco UCS Manager for future use.

Removing a Rack-Mount server

Removing is performed when you physically remove the server from the system by disconnecting the rack-mount server from the fabric extender. You cannot remove a rack-mount server from Cisco UCS Manager if it is physically present and connected to the fabric extender. Once the rack-mount server is disconnected, the configuration for that rack-mount server can be removed in Cisco UCS Manager.

During removal, management interfaces are disconnected, all entries from databases are removed, and the server is automatically removed from any server pools that it was assigned to during discovery.

**Note**

Only those servers added to a server pool automatically during discovery will be removed automatically. Servers that have been manually added to a server pool have to be removed manually.

If you need to add a removed rack-mount server back to the configuration, it must be reconnected and then rediscovered. When a server is reintroduced to Cisco UCS Manager it is treated like a new server and is subject to the deep discovery process. For this reason, it's possible that Cisco UCS Manager will assign the server a new ID that may be different from the ID that it held before.

Booting a Rack-Mount Server

Before You Begin

Associate a service profile with a rack-mount server.

Procedure

	Command or Action	Purpose
Step 1	UCS-A# scope org <i>org-name</i>	Enters organization mode for the specified organization. To enter the root organization mode, type / as the <i>org-name</i> .
Step 2	UCS-A /org # scope service-profile <i>profile-name</i>	Enters organization service profile mode for the specified service profile.
Step 3	UCS-A /org/service-profile # power up	Boots the rack-mount server associated with the service profile.
Step 4	UCS-A /org/service-profile # commit-buffer	Commits the transaction to the system configuration.

The following example boots the rack-mount server associated with the service profile named ServProf34 and commits the transaction:

```
UCS-A# scope org /
UCS-A /org* # scope service-profile ServProf34
UCS-A /org/service-profile # power up
UCS-A /org/service-profile* # commit-buffer
UCS-A /org/service-profile #
```

Shutting Down a Rack-Mount Server

When you use this procedure to shut down a server with an installed operating system, Cisco UCS Manager triggers the OS into a graceful shutdown sequence.

Before You Begin

Associate a service profile with a rack-mount server.

Procedure

	Command or Action	Purpose
Step 1	UCS-A# scope org <i>org-name</i>	Enters organization mode for the specified organization. To enter the root organization mode, type / as the <i>org-name</i> .
Step 2	UCS-A /org # scope service-profile <i>profile-name</i>	Enters organization service profile mode for the specified service profile.
Step 3	UCS-A /org/service-profile # power down	Shuts down the rack-mount server associated with the service profile.
Step 4	UCS-A /org/service-profile # commit-buffer	Commits the transaction to the system configuration.

The following example shuts down the rack-mount server associated with the service profile named ServProf34 and commits the transaction:

```
UCS-A# scope org /
UCS-A /org # scope service-profile ServProf34
UCS-A /org/service-profile # power down
UCS-A /org/service-profile* # commit-buffer
UCS-A /org/service-profile #
```

Power Cycling a Rack-Mount Server

Procedure

	Command or Action	Purpose
Step 1	UCS-A# <code>scope server server-num</code>	Enters server mode for the specified rack-mount server.
Step 2	UCS-A /server # <code>cycle</code> { <code>cycle-immediate</code> <code>cycle-wait</code> }	Power cycles the rack-mount server. Use the cycle-immediate keyword to immediately begin power cycling the rack-mount server; use the cycle-wait keyword to schedule the power cycle to begin after all pending management operations have completed.
Step 3	UCS-A# <code>commit-buffer</code>	Commits the transaction to the system configuration.

The following example immediately power cycles rack-mount server 2 and commits the transaction:

```
UCS-A# scope server 2
UCS-A /server # cycle cycle-immediate
UCS-A /server* # commit-buffer
UCS-A /server #
```

Performing a Hard Reset on a Rack-Mount Server

When you reset a server, Cisco UCS Manager sends a pulse on the reset line. You can choose to gracefully shut down the operating system. If the operating system does not support a graceful shut down, the server is power cycled. The option to have Cisco UCS Manager complete all management operations before it resets the server does not guarantee that these operations will be completed before the server is reset.



Note

If you are trying to boot a server from a power-down state, you should not use **Reset**.

If you continue the power-up with this process, the desired power state of the servers will become out of sync with the actual power state and the servers may unexpectedly shut down at a later time. To safely reboot the selected servers from a power-down state, click **Cancel** then select the **Boot Server** action.

Procedure

	Command or Action	Purpose
Step 1	UCS-A# scope server <i>server-num</i>	Enters server mode for the specified rack-mount server.
Step 2	UCS-A /server # reset { hard-reset-immediate hard-reset-wait }	Performs a hard reset of the rack-mount server. Use the hard-reset-immediate keyword to immediately begin hard resetting the rack-mount server; use the hard-reset-wait keyword to schedule the hard reset to begin after all pending management operations have completed.
Step 3	UCS-A /server # commit-buffer	Commits the transaction to the system configuration.

The following example performs an immediate hard reset of rack-mount server 2 and commits the transaction:

```
UCS-A# scope server 2
UCS-A /server # reset hard-reset-immediate
UCS-A /server* # commit-buffer
UCS-A /server #
```

Avoiding Unexpected Server Power Changes

If a server is not associated with a service profile, you can use any available means to change the server power state, including the physical Power or Reset buttons on the server.

If a server is associated with, or assigned to, a service profile, you should only use the following methods to change the server power state:

- In Cisco UCS Manager GUI, go to the **General** tab for the server or the service profile associated with the server and select **Boot Server** or **Shutdown Server** from the **Actions** area.
- In Cisco UCS Manager CLI, scope to the server or the service profile associated with the server and use the **power up** or **power down** commands.



Important

Do *not* use any of the following options on an associated server that is currently powered off:

- **Reset** in the GUI
- **cycle cycle-immediate** or **reset hard-reset-immediate** in the CLI
- The physical Power or Reset buttons on the server

If you reset, cycle, or use the physical power buttons on a server that is currently powered off, the server's actual power state may become out of sync with the desired power state setting in the service profile. If the communication between the server and Cisco UCS Manager is disrupted or if the service profile configuration changes, Cisco UCS Manager may apply the desired power state from the service profile to the server, causing an unexpected power change.

In Cisco UCS release 2.0.1, power synchronization issues can lead to unexpected server shut downs or restarts as shown below:

Desired Power State in Service Profile	Current Server Power State	Server Power State After Communication Is Disrupted
Up	Powered Off	Powered On
Down	Powered On	Powered Off

Starting in Cisco UCS release 2.0.2, power synchronization issues can lead to an unexpected server restart, as shown below:

Desired Power State in Service Profile	Current Server Power State	Server Power State After Communication Is Disrupted
Up	Powered Off	Powered On
Down	Powered On	Powered On Note Running servers are not shut down regardless of the desired power state in the service profile.

Acknowledging a Rack-Mount Server

Perform the following procedure if you need to have Cisco UCS Manager rediscover the server and all endpoints in the server. For example, you can use this procedure if a server is stuck in an unexpected state, such as the discovery state.

Procedure

	Command or Action	Purpose
Step 1	UCS-A# acknowledge server <i>server-num</i>	Acknowledges the specified rack-mount server.
Step 2	UCS-A# commit-buffer	Commits the transaction to the system configuration.

The following example acknowledges rack-mount server 2 and commits the transaction:

```
UCS-A# acknowledge server 2
UCS-A* # commit-buffer
UCS-A #
```

Decommissioning a Rack-Mount Server

Procedure

	Command or Action	Purpose
Step 1	UCS-A# decommission server <i>server-num</i>	Decommissions the specified rack-mount server.
Step 2	UCS-A# commit-buffer	Commits the transaction to the system configuration.

The following example decommissions rack-mount server 2 and commits the transaction:

```
UCS-A# decommission server 2
UCS-A* # commit-buffer
UCS-A #
```

Renumbering a Rack-Mount Server

Before You Begin

If you are swapping IDs between servers, you must first decommission both servers and then wait for the server decommission FSM to complete before proceeding with the renumbering steps.

Procedure

	Command or Action	Purpose
Step 1	UCS-A# show server inventory	Displays information about your servers.
Step 2	Verify that the server inventory does not include the following:	<ul style="list-style-type: none"> • The rack-mount server you want to renumber • A rack-mount server with the number you want to use <p>If either of these rack-mount servers are listed in the server inventory, decommission those servers. You must wait until the decommission FSM is complete and the rack-mount servers are not listed in the server inventory before continuing. This might take several minutes.</p> <p>To see which servers have been decommissioned, issue the show server decommissioned command.</p>
Step 3	UCS-A# recommission server <i>vendor-name model-name serial-numnew-id</i>	Recommissions and renumbers the specified rack-mount server.
Step 4	UCS-A# commit-buffer	Commits the transaction to the system configuration.

The following example decommissions a rack-mount server with ID 2, changes the ID to 3, recommissions that server, and commits the transaction:

UCS-A# **show server inventory**

Server Ackd Cores	Equipped	PID	Equipped VID	Equipped Serial (SN)	Slot	Status	Ackd Memory (MB)	
1/1 16	UCSB-B200-M3	V01		FCH1532718P		Equipped	131072	
1/2 16	UCSB-B200-M3	V01		FCH153271DF		Equipped	131072	
1/3 16	UCSB-B200-M3	V01		FCH153271DL		Equipped	114688	
1/4 1/5 1/6	UCSB-B200-M3	V01				Empty Empty Empty		
1/7 16 1/8	N20-B6730-1	V01		JAF1432CFDH		Equipped	65536	
1 12	R200-1120402W	V01		QCI1414A02J		N/A	49152	
2	R210-2121605W	V01		QCI1442AHFX		N/A	24576	8
4	UCSC-BSE-SFF-C200	V01		QCI1514A0J7		N/A	8192	8

UCS-A# **decommission server 2**

UCS-A*# **commit-buffer**

UCS-A# **show server decommissioned**

Vendor	Model	Serial (SN)	Server
Cisco Systems Inc	R210-2121605W	QCI1442AHFX	2

UCS-A# **recommission chassis "Cisco Systems Inc" "R210-2121605W" QCI1442AHFX 3**

UCS-A* # **commit-buffer**

UCS-A # **show server inventory**

Server Ackd Cores	Equipped	PID	Equipped VID	Equipped Serial (SN)	Slot	Status	Ackd Memory (MB)	
1/1 16	UCSB-B200-M3	V01		FCH1532718P		Equipped	131072	
1/2 16	UCSB-B200-M3	V01		FCH153271DF		Equipped	131072	
1/3 16	UCSB-B200-M3	V01		FCH153271DL		Equipped	114688	
1/4 1/5 1/6	UCSB-B200-M3	V01				Empty Empty Empty		
1/7 16 1/8	N20-B6730-1	V01		JAF1432CFDH		Equipped	65536	
1 12	R200-1120402W	V01		QCI1414A02J		N/A	49152	
3	R210-2121605W	V01		QCI1442AHFX		N/A	24576	8
4	UCSC-BSE-SFF-C200	V01		QCI1514A0J7		N/A	8192	8

Removing a Rack-Mount Server

Before You Begin

Physically disconnect the CIMC LOM cables that connect the rack-mount server to the fabric extender before performing the following procedure. For high availability setups, remove both cables.

Procedure

	Command or Action	Purpose
Step 1	UCS-A# remove server <i>server-num</i>	Removes the specified rack-mount server.
Step 2	UCS-A# commit-buffer	Commits the transaction to the system configuration.

The following example removes rack-mount server 4 and commits the transaction:

```
UCS-A# remove server 4
UCS-A* # commit-buffer
UCS-A #
```

What to Do Next

If you physically reconnect the rack-mount server, you must re-acknowledge it to have Cisco UCS Manager rediscover the server.

For more information, see [Acknowledging a Rack-Mount Server](#), on page 530.

Turning On the Locator LED for a Rack-Mount Server

Procedure

	Command or Action	Purpose
Step 1	UCS-A# scope server <i>server-num</i>	Enters server mode for the specified rack-mount server.
Step 2	UCS-A /server # enable locator-led	Turns on the rack-mount server locator LED.
Step 3	UCS-A /server # commit-buffer	Commits the transaction to the system configuration.

The following example turns on the locator LED for rack-mount server 2 and commits the transaction:

```
UCS-A# scope server 2
UCS-A /server # enable locator-led
UCS-A /server* # commit-buffer
UCS-A /server #
```

Turning Off the Locator LED for a Rack-Mount Server

Procedure

	Command or Action	Purpose
Step 1	UCS-A# scope server <i>server-num</i>	Enters server mode for the specified rack-mount server.
Step 2	UCS-A /server # disable locator-led	Turns off the rack-mount server locator LED.
Step 3	UCS-A /server # commit-buffer	Commits the transaction to the system configuration.

The following example turns off the locator LED for rack-mount server 2 and commits the transaction:

```
UCS-A# scope server 2
UCS-A /server # disable locator-led
UCS-A /server* # commit-buffer
UCS-A /server #
```

Resetting the CMOS for a Rack-Mount Server

On rare occasions, troubleshooting a server may require you to reset the CMOS. This procedure is not part of the normal maintenance of a server.

Procedure

	Command or Action	Purpose
Step 1	UCS-A# scope server <i>server-num</i>	Enters server mode for the rack-mount server.
Step 2	UCS-A /server # reset-cmos	Resets the CMOS for the rack-mount server.
Step 3	UCS-A /server # commit-buffer	Commits the transaction to the system configuration.

The following example resets the CMOS for rack-mount server 2 and commits the transaction:

```
UCS-A# scope server 2
UCS-A /server # reset-cmos
UCS-A /server* # commit-buffer
UCS-A /server #
```


Resetting the CIMC for a Rack-Mount Server

On rare occasions, such as an issue with the current running firmware, troubleshooting a server may require you to reset the CIMC. This procedure is not part of the normal maintenance of a server. After you reset the CIMC, the server boots with the running version of the firmware for that server.

Procedure

	Command or Action	Purpose
Step 1	UCS-A# scope server <i>server-num</i>	Enters server mode for the specified rack-mount server.
Step 2	UCS-A /server # scope CIMC	Enters server CIMC mode
Step 3	UCS-A /server/CIMC # reset	Resets the CIMC for the rack-mount server.
Step 4	UCS-A /server/CIMC # commit-buffer	Commits the transaction to the system configuration.

The following example resets the CIMC for rack-mount server 2 and commits the transaction:

```
UCS-A# scope server 2
UCS-A /server # scope CIMC
UCS-A /server/cimc # reset
UCS-A /server/cimc* # commit-buffer
UCS-A /server/cimc #
```

Recovering the Corrupt BIOS on a Rack-Mount Server

On rare occasions, an issue with a rack-mount server may require you to recover the corrupted BIOS. This procedure is not part of the normal maintenance of a rack-mount server. After you recover the BIOS, the rack-mount server boots with the running version of the firmware for that server.

Before You Begin



Important Remove all attached or mapped USB storage from a server before you attempt to recover the corrupt BIOS on that server. If an external USB drive is attached or mapped from vMedia to the server, BIOS recovery fails.

Procedure

	Command or Action	Purpose
Step 1	UCS-A# scope server <i>server-id</i>	Enters server mode for the specified rack-mount server.
Step 2	UCS-A /server # recover-bios <i>version [ignorecompcheck]</i>	Loads and activates the specified BIOS version.

	Command or Action	Purpose
		To activate the firmware without making sure that it is compatible first, include the ignorecompcheck keyword. We recommend that you use this option only when explicitly directed to do so by a technical support representative.
Step 3	UCS-A /server # commit-buffer	Commits the transaction.

The following example shows how to recover the BIOS:

```
UCS-A# scope server 1
UCS-A /server # recover-bios S5500.0044.0.3.1.010620101125
UCS-A /server* # commit-buffer
UCS-A /server #
```

Showing the Status for a Rack-Mount Server

Procedure

	Command or Action	Purpose
Step 1	UCS-A# show server status	Shows the status for all servers in the Cisco UCS domain.

The following example shows the status for all servers in the Cisco UCS domain. The servers numbered 1 and 2 do not have a slot listed in the table because they are rack-mount servers.

```
Server Slot  Status  Availability  Overall Status  Discovery
-----
1/1          Equipped  Unavailable  Ok              Complete
1/2          Equipped  Unavailable  Ok              Complete
1/3          Equipped  Unavailable  Ok              Complete
1/4          Empty     Unavailable  Ok              Complete
1/5          Equipped  Unavailable  Ok              Complete
1/6          Equipped  Unavailable  Ok              Complete
1/7          Empty     Unavailable  Ok              Complete
1/8          Empty     Unavailable  Ok              Complete
1            Equipped  Unavailable  Ok              Complete
2            Equipped  Unavailable  Ok              Complete
```

Issuing an NMI from a Rack-Mount Server

Perform the following procedure if the system remains unresponsive and you need Cisco UCS Manager to issue a Non Maskable Interrupt (NMI) to the BIOS or operating system from the CIMC. This action creates a core dump or stack trace, depending on the operating system installed on the server.

Procedure

	Command or Action	Purpose
Step 1	UCS-A# scope server [<i>chassis-num/server-num dynamic-uuid</i>]	Enters server mode for the specified server.
Step 2	UCS-A /chassis/server # diagnostic-interrupt	
Step 3	UCS-A /chassis/server* # commit-buffer	Commits any pending transactions.

The following example sends an NMI from server 4 in chassis 2 and commits the transaction:

```
UCS-A# scope server 2/4
UCS-A /chassis/server # diagnostic-interrupt
UCS-A /chassis/server* # commit-buffer
UCS-A /chassis/server #
```




CHAPTER 37

Managing the I/O Modules

This chapter includes the following sections:

- [I/O Module Management in Cisco UCS Manager GUI](#), page 539
- [Resetting the I/O Module](#), page 539

I/O Module Management in Cisco UCS Manager GUI

You can manage and monitor all I/O modules in a Cisco UCS domain through Cisco UCS Manager GUI.

Resetting the I/O Module

Procedure

	Command or Action	Purpose
Step 1	UCS-A# scope chassis <i>chassis-num</i>	Enters chassis mode for the specified chassis.
Step 2	UCS-A /chassis # scope iom {a b}	Enters chassis IOM mode for the specified IOM.
Step 3	UCS-A /chassis/iom # reset	Resets the IOM.
Step 4	UCS-A /chassis/iom # commit-buffer	Commits the transaction to the system configuration.

The following example resets the IOM on fabric A and commits the transaction:

```
UCS-A# scope chassis 1
UCS-A /chassis # scope iom a
UCS-A /chassis/iom # reset
UCS-A /chassis/iom* # commit-buffer
UCS-A /chassis/iom #
```




CHAPTER 38

Backing Up and Restoring the Configuration

This chapter includes the following sections:

- [Backup and Export Configuration, page 541](#)
- [Backup Types, page 541](#)
- [Considerations and Recommendations for Backup Operations, page 542](#)
- [Import Configuration, page 543](#)
- [Import Methods, page 543](#)
- [System Restore, page 543](#)
- [Required User Role for Backup and Import Operations, page 544](#)
- [Backup Operations, page 544](#)
- [Import Operations, page 548](#)
- [Restoring the Configuration for a Fabric Interconnect, page 552](#)
- [Erasing the Configuration, page 553](#)

Backup and Export Configuration

When you perform a backup through Cisco UCS Manager, you take a snapshot of all or part of the system configuration and export the file to a location on your network. You cannot use Cisco UCS Manager to back up data on the servers.

You can perform a backup while the system is up and running. The backup operation only saves information from the management plane. It does not have any impact on the server or network traffic.

Backup Types

You can perform one or more of the following types of backups through Cisco UCS Manager:

- **Full state**—A binary file that includes a snapshot of the entire system. You can use the file generated from this backup to restore the system during disaster recovery. This file can restore or rebuild the

configuration on the original fabric interconnect, or recreate the configuration on a different fabric interconnect. You cannot use this file for an import.

- **All configuration**—An XML file that includes all system and logical configuration settings. You can use the file generated from this backup to import these configuration settings to the original fabric interconnect or to a different fabric interconnect. You cannot use this file for a system restore. This file does not include passwords for locally authenticated users.
- **System configuration**—An XML file that includes all system configuration settings such as usernames, roles, and locales. You can use the file generated from this backup to import these configuration settings to the original fabric interconnect or to a different fabric interconnect. You cannot use this file for a system restore.
- **Logical configuration**—An XML file that includes all logical configuration settings such as service profiles, VLANs, VSANs, pools, and policies. You can use the file generated from this backup to import these configuration settings to the original fabric interconnect or to a different fabric interconnect. You cannot use this file for a system restore.

Considerations and Recommendations for Backup Operations

Before you create a backup operation, consider the following:

Backup Locations

The backup location is the destination or folder on the network where you want Cisco UCS Manager to export the backup file. You can maintain only one backup operation for each location where you plan to save a backup file.

Potential to Overwrite Backup Files

If you rerun a backup operation without changing the filename, Cisco UCS Manager overwrites the existing file on the server. To avoid overwriting existing backup files, change the filename in the backup operation or copy the existing file to another location.

Multiple Types of Backups

You can run and export more than one type of backup to the same location. You need to change the backup type before you rerun the backup operation. We recommend that you change the filename for easier identification of the backup type and to avoid overwriting the existing backup file.

Scheduled Backups

You cannot schedule a backup operation. You can, however, create a backup operation in advance and leave the admin state disabled until you are ready to run the backup. Cisco UCS Manager does not run the backup operation, save, or export the configuration file until you set the admin state of the backup operation to enabled.

Incremental Backups

You cannot perform incremental backups of the Cisco UCS Manager system configuration.

Backwards Compatibility

Starting with Release 1.1(1) of the Cisco UCS Manager, full state backups are encrypted so that passwords and other sensitive information are not exported as clear text. As a result, full state backups made from Release 1.1(1) or later cannot be restored to a Cisco UCS domain running an earlier software release.

Import Configuration

You can import any configuration file that was exported from Cisco UCS Manager. The file does not need to have been exported from the same Cisco UCS Manager.

The import function is available for all configuration, system configuration, and logical configuration files. You can perform an import while the system is up and running. An import operation modifies information on the management plane only. Some modifications caused by an import operation, such as a change to a vNIC assigned to a server, can cause a server reboot or other operations that disrupt traffic.

You cannot schedule an import operation. You can, however, create an import operation in advance and leave the admin state disabled until you are ready to run the import. Cisco UCS Manager will not run the import operation on the configuration file until you set the admin state to enabled.

You can maintain only one import operation for each location where you saved a configuration backup file.

Import Methods

You can use one of the following methods to import and update a system configuration through Cisco UCS Manager:

- **Merge**—The information in the imported configuration file is compared with the existing configuration information. If there are conflicts, the import operation overwrites the information on the Cisco UCS domain with the information in the import configuration file.
- **Replace**—The current configuration information is replaced with the information in the imported configuration file one object at a time.

System Restore

You can use the restore function for disaster recovery.

You can restore a system configuration from any full state backup file that was exported from Cisco UCS Manager. The file does not need to have been exported from the Cisco UCS Manager on the system that you are restoring. When restoring using a backup file that was exported from a different system, we strongly recommend that you use a system with the same or similar system configuration and hardware, including fabric interconnects, servers, adapters, and I/O module or FEX connectivity. Mismatched hardware and/or system configuration can lead to the restored system not fully functioning. If there is a mismatch between the I/O module links or servers on the two systems, acknowledge the chassis and/or servers after the restore operation.

The restore function is only available for a full state backup file. You cannot import a full state backup file. You perform a restore through the initial system setup.

Required User Role for Backup and Import Operations

You must have a user account that includes the admin role to create and run backup and import operations.

Backup Operations

Creating a Backup Operation

Before You Begin

Obtain the backup server IP address and authentication credentials.

Procedure

	Command or Action	Purpose
Step 1	UCS-A# scope system	Enters system mode.
Step 2	UCS-A /system # create backup URL backup-type {disabled enabled}	<p>Creates a backup operation. Specify the <i>URL</i> for the backup file using one of the following syntax:</p> <ul style="list-style-type: none"> • ftp:// username@hostname / path • scp:// username@hostname / path • sftp:// username@hostname / path • tftp:// hostname : port-num / path <p>The <i>backup-type</i> argument can be one of the following values:</p> <ul style="list-style-type: none"> • all-configuration —Backs up the server-, fabric-, and system-related configuration • logical-configuration —Backs up the fabric- and service profile-related configuration • system-configuration —Backs up the system-related configuration • full-state —Backs up the full state for disaster recovery <p>Note Full-state backup files cannot be imported using an import operation. They are used only to restore the configuration for a fabric interconnect.</p> <p>You can save multiple backup operations, but only one operation for each hostname is saved.</p> <p>If you use the enable keyword, the backup operation automatically runs as soon as you enter the commit-buffer command. If you use the disable keyword, the backup operation will not run until it is enabled. When enabling a backup operation, you must specify the hostname you used when creating the backup operation.</p>

	Command or Action	Purpose
Step 3	UCS-A /system # commit-buffer	Commits the transaction.

The following example creates a disabled all-configuration backup operation for hostname host35 and commits the transaction:

```
UCS-A# scope system
UCS-A /system* # create backup scp://user@host35/backups/all-config9.bak all-configuration
disabled
Password:
UCS-A /system* # commit-buffer
UCS-A /system #
```

Running a Backup Operation

Procedure

	Command or Action	Purpose
Step 1	UCS-A# scope system	Enters system mode.
Step 2	UCS-A /system # scope backup <i>hostname</i>	Enters system backup mode for the specified hostname.
Step 3	UCS-A /system/backup # enable	Enables the backup operation. Note For backup operations using FTP, SCP, SFTP, you are prompted for the password. Enter the password before committing the transaction.
Step 4	UCS-A /system/backup # commit-buffer	Commits the transaction.

The following example enables a backup operation named host35, enters the password for the SCP protocol, and commits the transaction:

```
UCS-A# scope system
UCS-A /system # scope backup host35
UCS-A /system/backup # enable
Password:
UCS-A /system/backup* # commit-buffer
UCS-A /system/backup #
```

Modifying a Backup Operation

You can modify a backup operation to save a file of another backup type to that location or to change the filename and avoid overwriting previous backup files.

Procedure

	Command or Action	Purpose
Step 1	UCS-A# scope system	Enters system mode.
Step 2	UCS-A /system # scope backup <i>hostname</i>	Enters system backup mode for the specified hostname.
Step 3	UCS-A /system/backup # disable	(Optional) Disables an enabled backup operation so that it does not automatically run when the transaction is committed.
Step 4	UCS-A /system/backup # enable	(Optional) Automatically runs the backup operation as soon as you commit the transaction.
Step 5	UCS-A /system/backup # set descr <i>description</i>	(Optional) Provides a description for the backup operation. Note If your description includes spaces, special characters, or punctuation, you must begin and end your description with quotation marks. The quotation marks will not appear in the description field of any show command output.
Step 6	UCS-A /system/backup # set protocol { ftp scp sftp tftp }	(Optional) Specifies the protocol to use when communicating with the remote server.
Step 7	UCS-A /system/backup # set remote-file <i>filename</i>	(Optional) Specifies the name of the configuration file that is being backed up.
Step 8	UCS-A /system/backup # set type <i>backup-type</i>	(Optional) Specifies the type of backup file to be made. The <i>backup-type</i> argument can be one of the following values: <ul style="list-style-type: none"> • all-configuration —Backs up the server, fabric, and system related configuration • logical-configuration —Backs up the fabric and service profile related configuration • system-configuration —Backs up the system related configuration • full-state —Backs up the full state for disaster recovery Note Full-state backup files cannot be imported using an import operation. They are used only to restore the configuration for a fabric interconnect.

	Command or Action	Purpose
Step 9	UCS-A /system/backup # set preserve-pooled-values {no yes}	(Optional) Specifies whether pool-derived identity values, such as vHBA WWPN, vNIC MAC, WWNN, and UUID, will be saved with the backup.
Step 10	UCS-A /system/backup # set user <i>username</i>	(Optional) Specifies the username the system should use to log in to the remote server. This step does not apply if the TFTP protocol is used.
Step 11	UCS-A /system/backup # set password <i>password</i>	(Optional) Specifies the password for the remote server username. This step does not apply if the TFTP protocol is used. Note Cisco UCS Manager does not store this password. Therefore, you do not need to enter this password unless you intend to enable and run the backup operation immediately.
Step 12	UCS-A /system/backup # commit-buffer	Commits the transaction.

The following example adds a description and changes the protocol, username, and password for the host35 backup operation and commits the transaction:

```
UCS-A# scope system
UCS-A /system # scope backup host35
UCS-A /system/backup # set descr "This is a backup operation for host35."
UCS-A /system/backup* # set protocol sftp
UCS-A /system/backup* # set user UserName32
UCS-A /system/backup* # set password
Password:
UCS-A /system/backup* # set preserve-pooled-values no
UCS-A /system/backup* # commit-buffer
UCS-A /system #
```

Deleting a Backup Operation

Procedure

	Command or Action	Purpose
Step 1	UCS-A# scope system	Enters system mode.
Step 2	UCS-A /system # delete backup <i>hostname</i>	Deletes the backup operation for the specified hostname.
Step 3	UCS-A /system # commit-buffer	Commits the transaction.

The following example deletes a backup operation for the host35 hostname and commits the transaction:

```
UCS-A# scope system
UCS-A /system # delete backup host35
UCS-A /system* # commit-buffer
UCS-A /system #
```

Import Operations

Creating an Import Operation

You cannot import a Full State configuration file. You can import any of the following configuration files:

- All configuration
- System configuration
- Logical configuration

Before You Begin

Collect the following information that you will need to import a configuration file:

- Backup server IP address and authentication credentials
- Fully qualified name of a backup file

Procedure

	Command or Action	Purpose
Step 1	UCS-A# scope system	Enters system mode.
Step 2	UCS-A /system # create import-config <i>URL</i> { disabled enabled } { merge replace }	<p>Creates an import operation. Specify the URL for the file being imported using one of the following syntax:</p> <ul style="list-style-type: none"> • ftp:// <i>username@hostname</i> / <i>path</i> • scp:// <i>username@hostname</i> / <i>path</i> • sftp:// <i>username@hostname</i> / <i>path</i> • tftp:// <i>hostname</i> : <i>port-num</i> / <i>path</i> <p>You can save multiple import operations, but only one operation for each hostname is saved.</p> <p>If you use the enable keyword, the import operation automatically runs as soon as you enter the commit-buffer command. If you use the disable keyword, the import operation will not run until it is enabled. When enabling an import operation, you must specify the hostname you used when creating the import operation.</p> <p>If you use the merge keyword, the configuration information is merged with the existing information. If there are conflicts, the system replaces the information on the current system with the information in</p>

	Command or Action	Purpose
		the import configuration file. If you use the replace keyword, the system takes each object in the import configuration file and overwrites the corresponding object in the current configuration.
Step 3	UCS-A /system/import-config# set descr <i>description</i>	(Optional) Provides a description for the import operation. Note If your description includes spaces, special characters, or punctuation, you must begin and end your description with quotation marks. The quotation marks will not appear in the description field of any show command output.
Step 4	UCS-A /system/import-config # commit-buffer	Commits the transaction.

The following example creates a disabled import operation for hostname host35 that replaces the existing configuration and commits the transaction:

```
UCS-A# scope system
UCS-A /system* # create import-config scp://user@host35/backups/all-config9.bak disabled
replace
Password:
UCS-A /system/import-config* # commit-buffer
UCS-A /system/import-config #
```

Running an Import Operation

You cannot import a Full State configuration file. You can import any of the following configuration files:

- All configuration
- System configuration
- Logical configuration

Procedure

	Command or Action	Purpose
Step 1	UCS-A# scope system	Enters system mode.
Step 2	UCS-A /system # scope import-config <i>hostname</i>	Enters system backup mode for the specified hostname.
Step 3	UCS-A /system/import-config # enable	Enables the import operation.
Step 4	UCS-A /system/import-config # commit-buffer	Commits the transaction.

The following example enables an import operation for the host35 hostname and commits the transaction:

```
UCS-A# scope system
UCS-A /system # scope import-config host35
UCS-A /system/import-config # enable
UCS-A /system/import-config* # commit-buffer
UCS-A /system/import-config #
```

Modifying an Import Operation

Procedure

	Command or Action	Purpose
Step 1	UCS-A# scope system	Enters system mode.
Step 2	UCS-A /system # scope import-config <i>hostname</i>	Enters system import configuration mode for the specified hostname.
Step 3	UCS-A /system/import-config # disable	(Optional) Disables an enabled import operation so that it does not automatically run when the transaction is committed.
Step 4	UCS-A /system/import-config # enable	(Optional) Automatically runs the import operation as soon as you commit the transaction.
Step 5	UCS-A /system/import-config # set action { merge replace }	(Optional) Specifies one of the following action types to use for the import operation: <ul style="list-style-type: none"> • Merge —The configuration information is merged with the existing information. If there are conflicts, the system replaces the information on the current system with the information in the import configuration file. • Replace —The system takes each object in the import configuration file and overwrites the corresponding object in the current configuration.
Step 6	UCS-A /system/import-config # set descr <i>description</i>	(Optional) Provides a description for the import operation. Note If your description includes spaces, special characters, or punctuation, you must begin and end your description with quotation marks. The quotation marks will not appear in the description field of any show command output.
Step 7	UCS-A /system/import-config # set password <i>password</i>	(Optional) Specifies the password for the remote server username. This step does not apply if the TFTP protocol is used.

	Command or Action	Purpose
		Note Cisco UCS Manager does not store this password. Therefore, you do not need to enter this password unless you intend to enable and run the import operation immediately.
Step 8	UCS-A /system/import-config # set protocol {ftp scp sftp tftp}	(Optional) Specifies the protocol to use when communicating with the remote server.
Step 9	UCS-A /system/import-config # set remote-file filename	(Optional) Specifies the name of the configuration file that is being imported.
Step 10	UCS-A /system/import-config # set user username	(Optional) Specifies the username the system should use to log in to the remote server. This step does not apply if the TFTP protocol is used.
Step 11	UCS-A /system/import-config # commit-buffer	Commits the transaction.

The following example adds a description, changes the password, protocol and username for the host35 import operation, and commits the transaction:

```
UCS-A# scope system
UCS-A /system # scope import-config host35
UCS-A /system/import-config # set descr "This is an import operation for host35."
UCS-A /system/import-config* # set password
Password:
UCS-A /system/import-config* # set protocol sftp
UCS-A /system/import-config* # set user jforlenz32
UCS-A /system/import-config* # commit-buffer
UCS-A /system/import-config #
```

Deleting an Import Operation

Procedure

	Command or Action	Purpose
Step 1	UCS-A# scope system	Enters system mode.
Step 2	UCS-A /system # delete import-config hostname	Deletes the import operation for the specified hostname.
Step 3	UCS-A /system # commit-buffer	Commits the transaction.

The following example deletes the import operation for the host35 hostname and commits the transaction:

```
UCS-A# scope system
UCS-A /system # delete import-config host35
```

```
UCS-A /system* # commit-buffer
UCS-A /system #
```

Restoring the Configuration for a Fabric Interconnect

Before You Begin

Collect the following information that you will need to restore the system configuration:

- Fabric interconnect management port IP address and subnet mask
- Default gateway IP address
- Backup server IP address and authentication credentials
- Fully qualified name of a Full State backup file



Note

You must have access to a Full State configuration file to perform a system restore. You cannot perform a system restore with any other type of configuration or backup file.

Procedure

-
- Step 1** Connect to the console port.
- Step 2** If the fabric interconnect is off, power on the fabric interconnect.
You will see the power on self-test message as the fabric interconnect boots.
- Step 3** At the installation method prompt, enter **console** .
- Step 4** Enter **restore** to restore the configuration from a full-state backup.
- Step 5** Enter **y** to confirm that you want to restore from a full-state backup.
- Step 6** Enter the IP address for the management port on the fabric interconnect.
- Step 7** Enter the subnet mask for the management port on the fabric interconnect.
- Step 8** Enter the IP address for the default gateway.
- Step 9** Enter one of the following protocols to use when retrieving the backup configuration file:
- **scp**
 - **ftp**
 - **tftp**
 - **sftp**
- Step 10** Enter the IP address of the backup server.
- Step 11** Enter the full path and filename of the Full State backup file.
- Step 12** Enter the username and password to access the backup server.
The fabric interconnect logs in to the backup server, retrieves a copy of the specified Full State backup file, and restores the system configuration. For a cluster configuration, you do not need to restore the secondary fabric interconnect. As soon as the secondary fabric interconnect reboots, Cisco UCS Manager synchronizes the configuration with the primary fabric interconnect.

The following example restores a system configuration from the Backup.bak file, which was retrieved from the 20.10.20.10 backup server using FTP:

```

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

Enter the setup mode; setup newly or restore from backup. (setup/restore) ? restore

NOTE:
  To configure Fabric interconnect using a backup file on a remote server,
  you will need to setup management interface.
  The management interface will be re-configured (if necessary),
  based on information stored in the backup file.

Continue to restore this Fabric interconnect from a backup file (yes/no) ? yes

Physical Switch Mgmt0 IPv4 address : 192.168.10.10

Physical Switch Mgmt0 IPv4 netmask : 255.255.255.0

IPv4 address of the default gateway : 192.168.10.1

Enter the protocol to get backup file (scp/ftp/tftp/sftp) ? scp
Enter the IP address of backup server: 20.10.20.10
Enter fully qualified backup file name: Backup.bak
Enter user ID: user
Enter password:
  Retrieved backup configuration file.
Configuration file - Ok

```

```

Cisco UCS 6100 Series Fabric Interconnect
UCS-A login:

```

Erasing the Configuration



Caution

You should erase the configuration only when it is necessary. Erasing the configuration completely removes the configuration and reboots the system in an unconfigured state. You must then either restore the configuration from a backup file or perform an initial system setup.

Procedure

	Command or Action	Purpose
Step 1	UCS-A# connect local-mgmt	Enters the local management CLI.
Step 2	UCS-A(local-mgmt)# erase configuration	Erases the configuration. You are prompted to confirm that you want to erase the configuration. Entering yes erases the configuration and reboots the system in an unconfigured state.

The following example erases the configuration:

```
UCS-A# connect local-mgmt  
UCS-A(local-mgmt)# erase configuration  
All UCS configurations will be erased and system will reboot. Are you sure? (yes/no): yes
```



CHAPTER 39

Recovering a Lost Password

This chapter includes the following sections:

- [Password Recovery for the Admin Account, page 555](#)
- [Determining the Leadership Role of a Fabric Interconnect, page 556](#)
- [Recovering the Admin Account Password in a Standalone Configuration, page 556](#)
- [Recovering the Admin Account Password in a Cluster Configuration, page 557](#)

Password Recovery for the Admin Account

The admin account is the system administrator or superuser account. If an administrator loses the password to this account, you can have a serious security issue. As a result, the procedure to recover the password for the admin account requires you to power cycle all fabric interconnects in a Cisco UCS domain.

When you recover the password for the admin account, you actually change the password for that account. You cannot retrieve the original password for that account.

You can reset the password for all other local accounts through Cisco UCS Manager. However, you must log in to Cisco UCS Manager with an account that includes aaa or admin privileges.



Caution

This procedure requires you to power down all fabric interconnects in a Cisco UCS domain. As a result, all data transmission in the Cisco UCS domain is stopped until you restart the fabric interconnects.

Determining the Leadership Role of a Fabric Interconnect

Procedure

	Command or Action	Purpose
Step 1	UCS-A# show cluster state	Displays the operational state and leadership role for both fabric interconnects in a cluster.

The following example displays the leadership role for both fabric interconnects in a cluster, where fabric interconnect A has the primary role and fabric interconnect B has the subordinate role:

```
UCS-A# show cluster state
Cluster Id: 0x4432f72a371511de-0xb97c000de1b1ada4

A: UP, PRIMARY
B: UP, SUBORDINATE

HA READY
```

Recovering the Admin Account Password in a Standalone Configuration

This procedure will help you to recover the password that you set for the admin account when you performed an initial system setup on the fabric interconnect. The admin account is the system administrator or superuser account.

Before You Begin

- 1 Physically connect the console port on the fabric interconnect to a computer terminal or console server
- 2 Determine the running versions of the following firmware:
 - The firmware kernel version on the fabric interconnect
 - The firmware system version



Tip

To find this information, you can log in with any user account on the Cisco UCS domain.

Procedure

- Step 1** Connect to the console port.
- Step 2** Power cycle the fabric interconnect:
 - a) Turn off the power to the fabric interconnect.
 - b) Turn on the power to the fabric interconnect.
- Step 3** In the console, press one of the following key combinations as it boots to get the `loader` prompt:

- Ctrl+l
- Ctrl+Shift+r

You may need to press the selected key combination multiple times before your screen displays the loader prompt.

Step 4 Boot the kernel firmware version on the fabric interconnect.

```
loader >
boot /installables/switch/
kernel_firmware_version
```

Example:

```
loader >
boot /installables/switch/ucs-6100-k9-kickstart.4.1.3.N2.1.0.11.gbin
```

Step 5 Enter config terminal mode.

```
Fabric(boot) #
config terminal
```

Step 6 Reset the admin password.

```
Fabric(boot) (config) #
admin-password
password
```

Choose a strong password that includes at least one capital letter and one number. The password cannot be blank.

The new password displays in clear text mode.

Step 7 Exit config terminal mode and return to the boot prompt.

Step 8 Boot the system firmware version on the fabric interconnect.

```
Fabric(boot) #
load /installables/switch/
system_firmware_version
```

Example:

```
Fabric(boot) #
load /installables/switch/ucs-6100-k9-system.4.1.3.N2.1.0.211.bin
```

Step 9 After the system image loads, log in to Cisco UCS Manager.

Recovering the Admin Account Password in a Cluster Configuration

This procedure will help you to recover the password that you set for the admin account when you performed an initial system setup on the fabric interconnects. The admin account is the system administrator or superuser account.

Before You Begin

- 1 Physically connect a console port on one of the fabric interconnects to a computer terminal or console server
- 2 Obtain the following information:
 - The firmware kernel version on the fabric interconnect
 - The firmware system version
 - Which fabric interconnect has the primary leadership role and which is the subordinate

**Tip**

To find this information, you can log in with any user account on the Cisco UCS domain.

Procedure

Step 1 Connect to the console port.

Step 2 For the subordinate fabric interconnect:

- a) Turn off the power to the fabric interconnect.
- b) Turn on the power to the fabric interconnect.
- c) In the console, press one of the following key combinations as it boots to get the `loader` prompt:
 - `Ctrl+l`
 - `Ctrl+Shift+r`

You may need to press the selected key combination multiple times before your screen displays the `loader` prompt.

Step 3 Power cycle the primary fabric interconnect:

- a) Turn off the power to the fabric interconnect.
- b) Turn on the power to the fabric interconnect.

Step 4 In the console, press one of the following key combinations as it boots to get the `loader` prompt:

- `Ctrl+l`
- `Ctrl+Shift+r`

You may need to press the selected key combination multiple times before your screen displays the `loader` prompt.

Step 5 Boot the kernel firmware version on the primary fabric interconnect.

```
loader > boot /installables/switch/
kernel_firmware_version
```


Example:

```
loader > boot /installables/switch/ucs-6100-k9-kickstart.4.1.3.N2.1.0.11.gbin
```

Step 6 Enter config terminal mode.

```
Fabric(boot) # config terminal
```

Step 7 Reset the admin password.

```
Fabric(boot) (config) # admin-password password
```

Choose a strong password that includes at least one capital letter and one number. The password cannot be blank.

The new password displays in clear text mode.

Step 8 Exit config terminal mode and return to the boot prompt.

Step 9 Boot the system firmware version on the primary fabric interconnect.

```
Fabric(boot) # load /installables/switch/
system_firmware_version
```

Example:

```
Fabric(boot) # load /installables/switch/ucs-6100-k9-system.4.1.3.N2.1.0.211.bin
```

Step 10 After the system image loads, log in to Cisco UCS Manager.

Step 11 In the console for the subordinate fabric interconnect, do the following to bring it up:

a) Boot the kernel firmware version on the subordinate fabric interconnect.

```
loader > boot /installables/switch/
kernel_firmware_version
```

b) Boot the system firmware version on the subordinate fabric interconnect.

```
Fabric(boot) # load /installables/switch/
system_firmware_version
```




PART **VII**

System Monitoring

- [Monitoring Traffic, page 563](#)
- [Monitoring Hardware, page 575](#)
- [Configuring Statistics-Related Policies, page 585](#)
- [Configuring Call Home, page 601](#)
- [Managing the System Event Log, page 623](#)
- [Configuring Settings for Faults, Events, and Logs, page 631](#)



CHAPTER 40

Monitoring Traffic

This chapter includes the following sections:

- [Traffic Monitoring, page 563](#)
- [Guidelines and Recommendations for Traffic Monitoring, page 564](#)
- [Creating an Ethernet Traffic Monitoring Session, page 565](#)
- [Creating a Fibre Channel Traffic Monitoring Session, page 566](#)
- [Adding Traffic Sources to a Monitoring Session, page 567](#)
- [Activating a Traffic Monitoring Session, page 572](#)
- [Deleting a Traffic Monitoring Session, page 573](#)

Traffic Monitoring

Traffic monitoring copies traffic from one or more sources and sends the copied traffic to a dedicated destination port for analysis by a network analyzer. This feature is also known as Switched Port Analyzer (SPAN).

Type of Session

When you create a traffic monitoring session, you can choose either an Ethernet or Fibre Channel destination port to receive the traffic. The type of destination port determines the type of session, which in turn determines the types of available traffic sources. For an Ethernet traffic monitoring session, the destination port must be an unconfigured physical port. For a Fibre Channel traffic monitoring session, the destination port must be a Fibre Channel uplink port.

Traffic Sources

An Ethernet traffic monitoring session can monitor any of the following traffic sources:

- Uplink Ethernet port
- Ethernet port channel
- VLAN
- Service profile vNIC

- Service profile vHBA
- FCoE port
- Port channels
- Server port

A Fibre Channel traffic monitoring session can monitor any of the following traffic sources:

- Uplink Fibre Channel port
- SAN port channel
- VSAN
- Service profile vHBA
- Fibre Channel storage port

Guidelines and Recommendations for Traffic Monitoring

When configuring or activating traffic monitoring, consider the following guidelines:

- You can create and store up to 16 traffic monitoring sessions, but only two can be active at the same time.
- A traffic monitoring session is disabled by default when created. To begin monitoring traffic, you must activate the session.
- To monitor traffic from a server, add all vNICs from the service profile corresponding to the server.
- You can monitor Fibre Channel traffic using either a Fibre Channel traffic analyzer or an Ethernet traffic analyzer. When Fibre Channel traffic is monitored using an Ethernet traffic monitoring session, with an Ethernet destination port, the destination traffic will be FCoE.
- Because a traffic monitoring destination is a single physical port, a traffic monitoring session can monitor only a single fabric. To monitor uninterrupted vNIC traffic across a fabric failover, you must create two sessions—one per fabric—and connect two analyzers. Add the vNIC as the traffic source for both sessions.
- All traffic sources must be located within the same switch as the destination port.
- A port configured as a destination port cannot also be configured as a source port.
- A member port of a port channel cannot be configured individually as a source. If the port channel is configured as a source, all member ports are source ports.
- A vHBA can be a source for either an Ethernet or Fibre Channel monitoring session, but it cannot be a source for both simultaneously.
- A server port can be a source only if it is a non-virtualized rack server adapter-facing port.
- A Fibre Channel port on a Cisco UCS 6248 fabric interconnect cannot be configured as a source port.
- If you change the port profile of a virtual machine, any associated vNICs being used as source ports are removed from monitoring, and you must reconfigure the monitoring session.

- If a traffic monitoring session was configured on a dynamic vNIC under a release earlier than Cisco UCS Manager Release 2.0, you must reconfigure the traffic monitoring session after upgrading.



Note Traffic monitoring can impose a significant load on your system resources. To minimize the load, select sources that carry as little unwanted traffic as possible and disable traffic monitoring when it is not needed.

Creating an Ethernet Traffic Monitoring Session



Note This procedure describes creating an Ethernet traffic monitoring session. To create a Fibre Channel traffic monitoring session, the following changes are required:

- Enter the **scope fc-traffic-mon** command instead of the **scope eth-traffic-mon** command in Step 1.
- Enter the **create fc-mon-session** command instead of the **create eth-mon-session** command in Step 3.

Procedure

	Command or Action	Purpose
Step 1	UCS-A# scope eth-traffic-mon	Enters Ethernet traffic monitoring command mode.
Step 2	UCS-A /eth-traffic-mon # scope fabric {a b}	Enters traffic monitoring command mode for the specified fabric.
Step 3	UCS-A /eth-traffic-mon/fabric # create eth-mon-session session-name	Creates a traffic monitoring session with the specified name.
Step 4	UCS-A /eth-traffic-mon/fabric/eth-mon-session # create dest-interface slot-num port-num	Configures the interface at the specified slot and port number to be the destination for the traffic monitoring session. Enters the command mode for the interface.
Step 5	UCS-A /eth-traffic-mon/fabric/eth-mon-session/dest-interface # set speedadmin-speed	Sets the data transfer rate of the port channel to be monitored. This can be: <ul style="list-style-type: none"> • 10gbps—10 Gbps • 1gbps—1 Gbps • 20gbps—20 Gbps • 40gbps—40 Gbps

	Command or Action	Purpose
Step 6	UCS-A /eth-traffic-mon/fabric/eth-mon-session/dest-interface # commit-buffer	Commits the transaction to the system configuration.

The following example creates an Ethernet traffic monitoring session to copy and forward traffic to the destination port at slot 2, port 12, sets the admin speed to 20 Gbps, and commits the transaction:

```
UCS-A# scope eth-traffic-mon
UCS-A /eth-traffic-mon # scope fabric a
UCS-A /eth-traffic-mon/fabric # create eth-mon-session EthMonitor33
UCS-A /eth-traffic-mon/fabric/eth-mon-session* # create dest-interface 2 12
UCS-A /eth-traffic-mon/fabric/eth-mon-session/dest-interface* # set speed 20gbps
UCS-A /eth-traffic-mon/fabric/eth-mon-session/dest-interface* # commit-buffer
UCS-A /eth-traffic-mon/fabric/eth-mon-session/dest-interface #
```

What to Do Next

- Add traffic sources to the traffic monitoring session.
- Activate the traffic monitoring session.

Creating a Fibre Channel Traffic Monitoring Session

Procedure

	Command or Action	Purpose
Step 1	UCS-A# scope fc-traffic-mon	Enters Fibre Channel traffic monitoring command mode.
Step 2	UCS-A /fc-traffic-mon # scope fabric {a b}	Enters Fibre Channel traffic monitoring command mode for the specified fabric.
Step 3	UCS-A /fc-traffic-mon/fabric # create fc-mon-session session-name	Creates a Fibre Channel traffic monitoring session with the specified name.
Step 4	UCS-A /fc-traffic-mon/fabric/fc-mon-session # create dest-interface slot-num port-num	Creates and enters the command mode of the destination slot and port for the Fibre Channel traffic monitoring session.
Step 5	UCS-A /fc-traffic-mon/fabric/fc-mon-session/dest-interface # set speedadmin-speed	Sets the data transfer rate of the port channel to be monitored. This can be: <ul style="list-style-type: none"> • 1gbps—1 Gbps • 2gbps—2 Gbps • 4gbps—4 Gbps • 8gbps—8 Gbps

	Command or Action	Purpose
		<ul style="list-style-type: none"> • auto—Cisco UCS determines the data transfer rate.
Step 6	UCS-A /fc-traffic-mon/fabric/fc-mon-session/dest-interface # commit-buffer	Commits the transaction to the system configuration.

The following example creates a Fibre channel traffic monitoring session to copy and forward traffic to the destination port at slot 1, port 10, sets the admin speed to 8 Gbps, and commits the transaction:

```
UCS-A# scope fc-traffic-mon
UCS-A /fc-traffic-mon # scope fabric a
UCS-A /fc-traffic-mon/fabric # create fc-mon-session FCMonitor
UCS-A /fc-traffic-mon/fabric/fc-mon-session* # create dest-interface 1 10
UCS-A /fc-traffic-mon/fabric/fc-mon-session/dest-interface* # set speed 8gbps
UCS-A /fc-traffic-mon/fabric/fc-mon-session/dest-interface* # commit-buffer
UCS-A /fc-traffic-mon/fabric/fc-mon-session/dest-interface #
```

What to Do Next

- Add traffic sources to the traffic monitoring session.
- Activate the traffic monitoring session.

Adding Traffic Sources to a Monitoring Session

Adding an Uplink Source Port to a Monitoring Session



Note

This procedure describes adding an Ethernet uplink port as a source for a traffic monitoring session. To add a Fibre Channel uplink port as a source, enter the **scope fc-uplink** command instead of the **scope eth-uplink** command in Step 1.

Before You Begin

A traffic monitoring session must be created.

Procedure

	Command or Action	Purpose
Step 1	UCS-A# scope eth-uplink	Enters Ethernet uplink command mode.
Step 2	UCS-A /eth-uplink # scope fabric {a b}	Enters uplink fabric mode for the specified fabric.

	Command or Action	Purpose
Step 3	UCS-A /eth-uplink/fabric # scope interface <i>slot-num port-num</i>	Enters the interface command mode for the specified uplink port.
Step 4	UCS-A /eth-uplink/fabric/interface # create mon-src <i>session-name</i>	Adds the uplink port as a source to the specified monitoring session.
Step 5	UCS-A /eth-uplink/fabric/interface/mon-src # set direction { both receive transmit }	(Optional) Specifies the traffic direction to be monitored.
Step 6	UCS-A /eth-uplink/fabric/interface/mon-src # commit-buffer	Commits the transaction to the system configuration.

The following example adds the ingress traffic on Ethernet uplink port 3 on slot 2 of fabric A as a source for a monitoring session and commits the transaction:

```
UCS-A# scope eth-uplink
UCS-A /eth-uplink # scope fabric a
UCS-A /eth-uplink/fabric # scope interface 2 3
UCS-A /eth-uplink/fabric/interface # create mon-src Monitor23
UCS-A /eth-uplink/fabric/interface/mon-src* # set direction receive
UCS-A /eth-uplink/fabric/interface/mon-src* # commit-buffer
UCS-A /eth-uplink/fabric/interface/mon-src #
```

What to Do Next

You can add additional sources to the traffic monitoring session.

Adding a vNIC or vHBA Source to a Monitoring Session



Note

This procedure describes adding a vNIC as a source for a traffic monitoring session. To add a vHBA as a source, enter the **scope vhba** command instead of the **scope vnic** command in Step 2.

Before You Begin

A traffic monitoring session must be created.

Procedure

	Command or Action	Purpose
Step 1	Switch-A# scope system	Enters system mode.
Step 2	Switch-A /system # scope vm-mgmt	Enters VM management mode.
Step 3	Switch-A /system/vm-mgmt # show virtual-machine	(Optional) Displays the running virtual machines.

	Command or Action	Purpose
Step 4	Switch-A /system/vm-mgmt # scope virtual-machine uuid	Enters command mode for the virtual machine that contains the dynamic vNIC.
Step 5	Switch-A /system/vm-mgmt/virtual-machine # show expand	(Optional) Displays the virtual machine details, including the vNIC MAC address.
Step 6	Switch-A /system/vm-mgmt/virtual-machine # scope vnic mac-address	Enters the command mode for the vNIC at the specified MAC address.
Step 7	Switch-A /system/vm-mgmt/virtual-machine/vnic # create mon-src session-name	Adds the vNIC as a source to the specified monitoring session.
Step 8	Switch-A /system/vm-mgmt/virtual-machine/vnic/mon-src # set direction {both receive transmit}	(Optional) Specifies the traffic direction to be monitored.
Step 9	Switch-A /system/vm-mgmt/virtual-machine/vnic/mon-src # commit-buffer	Commits the transaction to the system configuration.

The following example adds the ingress traffic on a dynamic vNIC as a source for a monitoring session and commits the transaction:

```
Switch-A# scope system
Switch-A /system # scope vm-mgmt
Switch-A /system/vm-mgmt # show virtual-machine
Virtual Machine:
  UUID: 42327c42-e00c-886f-e3f7-e615906f51e9
  Service Profile: org-root/ls-dsw-bld1-esx
  Server: sys/chassis-1/blade-1
  Status: Online
.
.
.
Switch-A /system/vm-mgmt # scope virtual-machine 42327c42-e00c-886f-e3f7-e615906f51e9
Switch-A /system/vm-mgmt/virtual-machine # show expand
Virtual Machine:
  UUID: 42327c42-e00c-886f-e3f7-e615906f51e9
  Service Profile: org-root/ls-dsw-bld1-esx
  Server: sys/chassis-1/blade-1
  Status: Online

vNIC:
  Name:
  Status: Online
  MAC Address: 00:50:56:B2:00:00

VIF:
  Vif Id: 32772
  Status: Online
  Phys Fabric ID: B
  Virtual Fabric:
Switch-A /system/vm-mgmt/virtual-machine # scope vnic 00:50:56:B2:00:00
Switch-A /system/vm-mgmt/virtual-machine/vnic # create mon-src Monitor23
Switch-A /system/vm-mgmt/virtual-machine/vnic/mon-src* # set direction receive
Switch-A /system/vm-mgmt/virtual-machine/vnic/mon-src* # commit-buffer

Switch-A /system/vm-mgmt/virtual-machine/vnic/mon-src #
```

What to Do Next

You can add additional sources to the traffic monitoring session.

Adding a VLAN or VSAN Source to a Monitoring Session



Note

This procedure describes adding a VLAN as a source for a traffic monitoring session. To add a VSAN as a source, the following changes are required:

- Enter the **scope fc-uplink** command instead of the **scope eth-uplink** command in Step 1.
- Enter the **create vsan** command instead of the **create vlan** command in Step 3.

Before You Begin

A traffic monitoring session must be created.

Procedure

	Command or Action	Purpose
Step 1	UCS-A# scope eth-uplink	Enters Ethernet uplink command mode.
Step 2	UCS-A /eth-uplink # scope fabric {a b}	Enters uplink fabric mode for the specified fabric. Note This step is required when adding a local VLAN as a source. To add a global VLAN as a source, omit this step.
Step 3	UCS-A /eth-uplink/fabric # create vlan vlan-name vlan-id	Creates a named VLAN, specifies the VLAN name and VLAN ID, and enters uplink VLAN mode.
Step 4	UCS-A /eth-uplink/fabric/vlan # create mon-src session-name	Adds the VLAN as a source to the specified monitoring session.
Step 5	UCS-A /eth-uplink/fabric/vlan/mon-src # commit-buffer	Commits the transaction to the system configuration.

The following example adds a local VLAN as a source for an Ethernet monitoring session and commits the transaction:

```
UCS-A# scope eth-uplink
UCS-A /eth-uplink # scope fabric a
UCS-A /eth-uplink/fabric # create vlan vlan23 23
UCS-A /eth-uplink/fabric/vlan # create mon-src Monitor23
UCS-A /eth-uplink/fabric/vlan/mon-src* # commit-buffer
UCS-A /eth-uplink/fabric/vlan/mon-src #
```

What to Do Next

You can add additional sources to the traffic monitoring session.

Adding a Storage Port Source to a Monitoring Session



Note

This procedure describes adding a Fibre Channel storage port as a source for a Fibre Channel traffic monitoring session. To add an FCoE storage port as a source for an Ethernet traffic monitoring session, enter the **create interface fcoe** command instead of the **create interface fc** command in Step 3.

Before You Begin

A traffic monitoring session must be created.

Procedure

	Command or Action	Purpose
Step 1	UCS-A# scope fc-storage	Enters Fibre Channel storage port command mode.
Step 2	UCS-A /fc-storage # scope fabric {a b}	Enters Fibre Channel storage port fabric mode for the specified fabric.
Step 3	UCS-A /fc-storage/fabric # create interface fc slot-num port-num	Creates a Fibre Channel storage port interface and enters the interface command mode.
Step 4	UCS-A /fc-storage/fabric/fc # create mon-src session-name	Adds the storage port as a source to the specified monitoring session.
Step 5	UCS-A /fc-storage/fabric/fc/mon-src # commit-buffer	Commits the transaction to the system configuration.

The following example adds a Fibre Channel storage port on port 3 of slot 2 as a source for a Fibre Channel monitoring session and commits the transaction:

```
UCS-A# scope fc-storage
UCS-A /fc-storage # scope fabric a
UCS-A /fc-storage/fabric # create interface fc 2 3
UCS-A /fc-storage/fabric/fc* # create mon-src Monitor23
UCS-A /fc-storage/fabric/fc/mon-src* # commit-buffer
UCS-A /fc-storage/fabric/fc/mon-src #
```

What to Do Next

You can add additional sources to the traffic monitoring session.

Activating a Traffic Monitoring Session



Note

This procedure describes activating an Ethernet traffic monitoring session. To activate a Fibre Channel traffic monitoring session, the following changes are required:

- Enter the **scope fc-traffic-mon** command instead of the **scope eth-traffic-mon** command in Step 1.
- Enter the **scope fc-mon-session** command instead of the **scope eth-mon-session** command in Step 3.

Before You Begin

Configure a traffic monitoring session.

Procedure

	Command or Action	Purpose
Step 1	UCS-A# scope eth-traffic-mon	Enters Ethernet traffic monitoring command mode.
Step 2	UCS-A /eth-traffic-mon # scope fabric {a b}	Enters traffic monitoring command mode for the specified fabric.
Step 3	UCS-A /eth-traffic-mon/fabric # scope eth-mon-session session-name	Enters the command mode of the traffic monitoring session with the specified name.
Step 4	UCS-A /eth-traffic-mon/fabric/eth-mon-session # disable enable	Disables or enables the traffic monitoring session.
Step 5	UCS-A /eth-traffic-mon/fabric/eth-mon-session # commit-buffer	Commits the transaction to the system configuration.

When activated, the traffic monitoring session begins forwarding traffic to the destination as soon as a traffic source is configured.

The following example activates an Ethernet traffic monitoring session and commits the transaction:

```
UCS-A# scope eth-traffic-mon
UCS-A /eth-traffic-mon # scope fabric a
UCS-A /eth-traffic-mon/fabric # scope eth-mon-session Monitor33
UCS-A /eth-traffic-mon/fabric/eth-mon-session # enable
UCS-A /eth-traffic-mon/fabric/eth-mon-session* # commit-buffer
UCS-A /eth-traffic-mon/fabric/eth-mon-session # show
```

```
Ether Traffic Monitoring Session:
  Name      Admin State   Oper State   Oper State Reason
  -----
  Monitor33 Enabled       Up           Active
```

```
UCS-A /eth-traffic-mon/fabric/eth-mon-session #
```

Deleting a Traffic Monitoring Session



Note

This procedure describes deleting an Ethernet traffic monitoring session. To delete a Fibre Channel traffic monitoring session, the following changes are required:

- Enter the **scope fc-traffic-mon** command instead of the **scope eth-traffic-mon** command in Step 1.
- Enter the **delete fc-mon-session** command instead of the **delete eth-mon-session** command in Step 3.

Procedure

	Command or Action	Purpose
Step 1	UCS-A# scope eth-traffic-mon	Enters Ethernet traffic monitoring command mode.
Step 2	UCS-A /eth-traffic-mon # scope fabric {a b}	Enters traffic monitoring command mode for the specified fabric.
Step 3	UCS-A /eth-traffic-mon/fabric # delete eth-mon-session session-name	Deletes the traffic monitoring session with the specified name.
Step 4	UCS-A /eth-traffic-mon/fabric # commit-buffer	Commits the transaction to the system configuration.

The following example deletes an Ethernet traffic monitoring session and commits the transaction:

```
UCS-A# scope eth-traffic-mon
UCS-A /eth-traffic-mon # scope fabric a
UCS-A /eth-traffic-mon/fabric # delete eth-mon-session Monitor33
UCS-A /eth-traffic-mon/fabric* # commit-buffer
UCS-A /eth-traffic-mon/fabric #
```




CHAPTER 41

Monitoring Hardware

This chapter includes the following sections:

- [Monitoring Fan Modules, page 575](#)
- [Monitoring Management Interfaces, page 577](#)
- [Server Disk Drive Monitoring, page 579](#)

Monitoring Fan Modules

Procedure

	Command or Action	Purpose
Step 1	UCS-A# scope chassis <i>chassis-num</i>	Enters chassis mode for the specified chassis.
Step 2	UCS-A /chassis # show environment fan	Displays the environment status for all fans within the chassis. This includes the following information: <ul style="list-style-type: none">• Overall status• Operability• Power state• Thermal status• Threshold status• Voltage status
Step 3	UCS-A /chassis # scope fan-module <i>tray-num module-num</i>	Enters fan module chassis mode for the specified fan module. Note Each chassis contains one tray, so the tray number in this command is always 1.

	Command or Action	Purpose
Step 4	UCS-A /chassis/fan-module # show [detail expand]	Displays the environment status for the specified fan module.

The following example displays information about the fan modules in chassis 1:

```
UCS-A# scope chassis 1
UCS-A /chassis # show environment fan
Chassis 1:
  Overall Status: Power Problem
  Operability: Operable
  Power State: Redundancy Failed
  Thermal Status: Upper Non Recoverable

Tray 1 Module 1:
  Threshold Status: OK
  Overall Status: Operable
  Operability: Operable
  Power State: On
  Thermal Status: OK
  Voltage Status: N/A

  Fan Module Stats:
    Ambient Temp (C): 25.000000

  Fan 1:
    Threshold Status: OK
    Overall Status: Operable
    Operability: Operable
    Power State: On
    Thermal Status: OK
    Voltage Status: N/A

  Fan 2:
    Threshold Status: OK
    Overall Status: Operable
    Operability: Operable
    Power State: On
    Thermal Status: OK
    Voltage Status: N/A

Tray 1 Module 2:
  Threshold Status: OK
  Overall Status: Operable
  Operability: Operable
  Power State: On
  Thermal Status: OK
  Voltage Status: N/A

  Fan Module Stats:
    Ambient Temp (C): 24.000000

  Fan 1:
    Threshold Status: OK
    Overall Status: Operable
    Operability: Operable
    Power State: On
    Thermal Status: OK
    Voltage Status: N/A

  Fan 2:
    Threshold Status: OK
    Overall Status: Operable
    Operability: Operable
    Power State: On
    Thermal Status: OK
```

Voltage Status: N/A

The following example displays information about fan module 2 in chassis 1:

```
UCS-A# scope chassis 1
UCS-A /chassis # scope fan-module 1 2
UCS-A /chassis/fan-module # show detail
Fan Module:
  Tray: 1
  Module: 2
  Overall Status: Operable
  Operability: Operable
  Threshold Status: OK
  Power State: On
  Presence: Equipped
  Thermal Status: OK
  Product Name: Fan Module for UCS 5108 Blade Server Chassis
  PID: N20-FAN5
  VID: V01
  Vendor: Cisco Systems Inc
  Serial (SN): NWG14350B6N
  HW Revision: 0
  Mfg Date: 1997-04-01T08:41:00.000
```

Monitoring Management Interfaces

Management Interfaces Monitoring Policy

This policy defines how the mgmt0 Ethernet interface on the fabric interconnect should be monitored. If Cisco UCS detects a management interface failure, a failure report is generated. If the configured number of failure reports is reached, the system assumes that the management interface is unavailable and generates a fault. By default, the management interfaces monitoring policy is disabled.

If the affected management interface belongs to a fabric interconnect which is the managing instance, Cisco UCS confirms that the subordinate fabric interconnect's status is up, that there are no current failure reports logged against it, and then modifies the managing instance for the end-points.

If the affected fabric interconnect is currently the primary inside of a high availability setup, a failover of the management plane is triggered. The data plane is not affected by this failover.

You can set the following properties related to monitoring the management interface:

- Type of mechanism used to monitor the management interface.
- Interval at which the management interface's status is monitored.
- Maximum number of monitoring attempts that can fail before the system assumes that the management is unavailable and generates a fault message.



Important

In the event of a management interface failure on a fabric interconnect, the managing instance may not change if one of the following occurs:

- A path to the end-point through the subordinate fabric interconnect does not exist.
- The management interface for the subordinate fabric interconnect has failed.
- The path to the end-point through the subordinate fabric interconnect has failed.

Configuring the Management Interfaces Monitoring Policy

Procedure

-
- Step 1** Enter monitoring mode.
UCS-A# **scope monitoring**
- Step 2** Enable or disable the management interfaces monitoring policy.
UCS-A /monitoring # **set mgmt-if-mon-policy admin-state {enabled | disabled}**
- Step 3** Specify the number of seconds that the system should wait between data recordings.
UCS-A /monitoring # **set mgmt-if-mon-policy poll-interval**
Enter an integer between 90 and 300.
- Step 4** Specify the maximum number of monitoring attempts that can fail before the system assumes that the management interface is unavailable and generates a fault message.
UCS-A /monitoring # **set mgmt-if-mon-policy max-fail-reports num-mon-attempts**
Enter an integer between 2 and 5.
- Step 5** Specify the monitoring mechanism that you want the system to use.
UCS-A /monitoring # **set mgmt-if-mon-policy monitor-mechanism {mii-status | ping-arp-targets | ping-gateway}**
- **mii-status** —The system monitors the availability of the Media Independent Interface (MII).
 - **ping-arp-targets** —The system pings designated targets using the Address Resolution Protocol (ARP).
 - **ping-gateway** —The system pings the default gateway address specified for this Cisco UCS domain in the management interface.
- Step 6** If you selected **mii-status** as your monitoring mechanism, configure the following properties:
- a) Specify the number of seconds that the system should wait before requesting another response from the MII if a previous attempt fails.
UCS-A /monitoring # **set mgmt-if-mon-policy mii-retry-interval num-seconds**
Enter an integer between 3 and 10.
 - b) Specify the number of times that the system polls the MII until the system assumes that the interface is unavailable.
UCS-A /monitoring # **set mgmt-if-mon-policy mii-retry-count num-retries**
Enter an integer between 1 and 3.
- Step 7** If you selected **ping-arp-targets** as your monitoring mechanism, configure the following properties:
- a) Specify the first IP address the system pings.
UCS-A /monitoring # **set mgmt-if-mon-policy arp-target1 ip-addr**
Type 0.0.0.0 to remove the ARP target.
 - b) Specify the second IP address the system pings.
UCS-A /monitoring # **set mgmt-if-mon-policy arp-target2 ip-addr**
Type 0.0.0.0 to remove the ARP target.

- c) Specify the third IP address the system pings.
UCS-A /monitoring # **set mgmt-if-mon-policy arp-target3** *ip-addr*
Type 0.0.0.0 to remove the ARP target.
- d) Specify the number of ARP requests to send to the target IP addresses.
UCS-A /monitoring # **set mgmt-if-mon-policy arp-requests** *num-requests*
Enter an integer between 1 and 5.
- e) Specify the number of seconds to wait for responses from the ARP targets before the system assumes that they are unavailable.
UCS-A /monitoring # **set mgmt-if-mon-policy arp-deadline** *num-seconds*
Enter a number between 5 and 15.

Step 8 If you selected **ping-gateway** as your monitoring mechanism, configure the following properties:

- a) Specify the number of times the system should ping the gateway.
UCS-A /monitoring # **set mgmt-if-mon-policy ping-requests**
Enter an integer between 1 and 5.
- b) Specify the number of seconds to wait for a response from the gateway until the system assumes that the address is unavailable.
UCS-A /monitoring # **set mgmt-if-mon-policy ping-deadline**
Enter an integer between 5 and 15.

Step 9 Commit the transaction to the system configuration.

UCS-A /monitoring # **commit-buffer**

The following example creates a monitoring interface management policy using the Media Independent Interface (MII) monitoring mechanism and commits the transaction:

```
UCS-A# scope monitoring
UCS-A /monitoring # set mgmt-if-mon-policy admin-state enabled
UCS-A /monitoring* # set mgmt-if-mon-policy poll-interval 250
UCS-A /monitoring* # set mgmt-if-mon-policy max-fail-reports 2
UCS-A /monitoring* # set mgmt-if-mon-policy monitor-mechanism set mii-status
UCS-A /monitoring* # set mgmt-if-mon-policy mii-retry-count 3
UCS-A /monitoring* # set mgmt-if-mon-policy mii-retry-interval 7
UCS-A /monitoring* # commit-buffer
UCS-A /monitoring #
```

Server Disk Drive Monitoring

The disk drive monitoring for Cisco UCS provides Cisco UCS Manager with blade-resident disk drive status for supported blade servers in a Cisco UCS domain. Disk drive monitoring provides a unidirectional fault signal from the LSI firmware to Cisco UCS Manager to provide status information.

The following server and firmware components gather, send, and aggregate information about the disk drive status in a server:

- Physical presence sensor—Determines whether the disk drive is inserted in the server drive bay.
- Physical fault sensor—Determines the operability status reported by the LSI storage controller firmware for the disk drive.

- IPMI disk drive fault and presence sensors—Sends the sensor results to Cisco UCS Manager.
- Disk drive fault LED control and associated IPMI sensors—Controls disk drive fault LED states (on/off) and relays the states to Cisco UCS Manager.

Support for Disk Drive Monitoring

Disk drive monitoring only supports certain blade servers and a specific LSI storage controller firmware level.

Supported Cisco UCS Servers

Through Cisco UCS Manager, you can monitor disk drives for the following servers:

- B-200 blade server
- B-230 blade server
- B-250 blade server
- B-440 blade server

Cisco UCS Manager cannot monitor disk drives in any other blade server or rack-mount server.

Storage Controller Firmware Level

The storage controller on a supported server must have LSI 1064E firmware.

Cisco UCS Manager cannot monitor disk drives in servers with a different level of storage controller firmware.

Prerequisites for Disk Drive Monitoring

In addition to the supported servers and storage controller firmware version, you must ensure that the following prerequisites have been met for disk drive monitoring to provide useful status information:

- The drive must be inserted in the server drive bay.
- The server must be powered on.
- The server must have completed discovery.
- The results of the BIOS POST complete must be TRUE.

Viewing the Status of a Disk Drive

Procedure

	Command or Action	Purpose
Step 1	UCS-A# scope chassis <i>chassis-num</i>	Enters chassis mode for the specified chassis.
Step 2	UCS-A /chassis # scope server <i>server-num</i>	Enters server chassis mode.

	Command or Action	Purpose						
Step 3	UCS-A /chassis/server # scope raid-controller <i>raid-contr-id</i> { sas sata }	Enters RAID controller server chassis mode.						
Step 4	UCS-A /chassis/server/raid-controller # show local-disk [<i>local-disk-id</i> detail expand]	<p>Displays the following local disk statistics:</p> <table border="1"> <thead> <tr> <th>Name</th> <th>Description</th> </tr> </thead> <tbody> <tr> <td>Operability field</td> <td> <p>The operational state of the disk drive. This can be the following:</p> <ul style="list-style-type: none"> • Operable—The disk drive is operable. • Inoperable—The disk drive is inoperable, possibly due to a hardware issue such as bad blocks. • N/A—The operability of the disk drive cannot be determined. This could be due to the server or firmware not being support for disk drive monitoring, or because the server is powered off. <p>Note The Operability field may show the incorrect status for several reasons, such as if the disk is part of a broken RAID set or if the BIOS POST (Power On Self Test) has not completed.</p> </td> </tr> <tr> <td>Presence field</td> <td> <p>The presence of the disk drive, and whether it can be detected in the server drive bay, regardless of its operational state. This can be the following:</p> <ul style="list-style-type: none"> • Equipped—A disk drive can be detected in the server drive bay. • Missing—No disk drive can be detected in the server drive bay. </td> </tr> </tbody> </table>	Name	Description	Operability field	<p>The operational state of the disk drive. This can be the following:</p> <ul style="list-style-type: none"> • Operable—The disk drive is operable. • Inoperable—The disk drive is inoperable, possibly due to a hardware issue such as bad blocks. • N/A—The operability of the disk drive cannot be determined. This could be due to the server or firmware not being support for disk drive monitoring, or because the server is powered off. <p>Note The Operability field may show the incorrect status for several reasons, such as if the disk is part of a broken RAID set or if the BIOS POST (Power On Self Test) has not completed.</p>	Presence field	<p>The presence of the disk drive, and whether it can be detected in the server drive bay, regardless of its operational state. This can be the following:</p> <ul style="list-style-type: none"> • Equipped—A disk drive can be detected in the server drive bay. • Missing—No disk drive can be detected in the server drive bay.
Name	Description							
Operability field	<p>The operational state of the disk drive. This can be the following:</p> <ul style="list-style-type: none"> • Operable—The disk drive is operable. • Inoperable—The disk drive is inoperable, possibly due to a hardware issue such as bad blocks. • N/A—The operability of the disk drive cannot be determined. This could be due to the server or firmware not being support for disk drive monitoring, or because the server is powered off. <p>Note The Operability field may show the incorrect status for several reasons, such as if the disk is part of a broken RAID set or if the BIOS POST (Power On Self Test) has not completed.</p>							
Presence field	<p>The presence of the disk drive, and whether it can be detected in the server drive bay, regardless of its operational state. This can be the following:</p> <ul style="list-style-type: none"> • Equipped—A disk drive can be detected in the server drive bay. • Missing—No disk drive can be detected in the server drive bay. 							

The following example shows the status of a disk drive:

```
UCS-A# scope chassis 1
UCS-A /chassis # scope server 6
UCS-A /chassis/server # scope raid-controller 1 sas
UCS-A /chassis/server/raid-controller # show local-disk 1
```

```

Local Disk:
  ID: 1
  Block Size: 512
  Blocks: 60545024
  Size (MB): 29563
  Operability: Operable
  Presence: Equipped

```

Interpreting the Status of a Monitored Disk Drive

Cisco UCS Manager displays the following properties for each monitored disk drive:

- **Operability**—The operational state of the disk drive.
- **Presence**—The presence of the disk drive, and whether it can be detected in the server drive bay, regardless of its operational state.

You need to look at both properties to determine the status of the monitored disk drive. The following table shows the likely interpretations of the property values.

Operability Status	Presence Status	Interpretation
Operable	Equipped	No fault condition. The disk drive is in the server and can be used.
Inoperable	Equipped	Fault condition. The disk drive is in the server, but one of the following could be causing an operability problem: <ul style="list-style-type: none"> • The disk drive is unusable due to a hardware issue such as bad blocks. • There is a problem with the IPMI link to the storage controller.
N/A	Missing	Fault condition. The server drive bay does not contain a disk drive.
N/A	Equipped	Fault condition. The disk drive is in the server, but one of the following could be causing an operability problem: <ul style="list-style-type: none"> • The server is powered off. • The storage controller firmware is the wrong version and does not support disk drive monitoring. • The server does not support disk drive monitoring.

**Note**

The **Operability** field may show the incorrect status for several reasons, such as if the disk is part of a broken RAID set or if the BIOS POST (Power On Self Test) has not completed.



CHAPTER 42

Configuring Statistics-Related Policies

This chapter includes the following sections:

- [Configuring Statistics Collection Policies, page 585](#)
- [Configuring Statistics Threshold Policies, page 586](#)

Configuring Statistics Collection Policies

Statistics Collection Policy

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

For NIC statistics, Cisco UCS Manager displays the average, minimum, and maximum of the change since the last collection of statistics. If the values are 0, there has been no change since the last collection.

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

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



Note

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

Configuring a Statistics Collection Policy

Procedure

	Command or Action	Purpose
Step 1	UCS-A# scope monitoring	Enters monitoring mode.
Step 2	UCS-A/monitoring# scope stats-collection-policy {adapter chassis host port server}	Enters statistics collection policy mode for the specified policy type.
Step 3	UCS-A /monitoring/stats-collection-policy # set collection-interval {1minute 2minutes 30seconds 5minutes}	Specifies the interval at which statistics are collected from the system.
Step 4	UCS-A /monitoring/stats-collection-policy # set reporting-interval {15minutes 30minutes 60minutes}	Specifies the interval at which collected statistics are reported.
Step 5	UCS-A /monitoring/stats-collection-policy # commit-buffer	Commits the transaction to the system configuration.

The following example creates a statistics collection policy for ports, sets the collection interval to one minute, the reporting interval to 30 minutes, and commits the transaction:

```
UCS-A# scope monitoring
UCS-A /monitoring # scope stats-collection-policy port
UCS-A /monitoring/stats-collection-policy* # set collection-interval 1minute
UCS-A /monitoring/stats-collection-policy* # set reporting-interval 30minutes
UCS-A /monitoring/stats-collection-policy* # commit-buffer
UCS-A /monitoring/stats-collection-policy #
```

Configuring Statistics Threshold Policies

Statistics Threshold Policy

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

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

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

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

**Note**

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

Server and Server Component Statistics Threshold Policy Configuration

Configuring a Server and Server Component Statistics Threshold Policy

Procedure

	Command or Action	Purpose
Step 1	UCS-A# scope org <i>org-name</i>	Enters organization mode for the specified organization. To enter the root organization mode, type / as the <i>org-name</i> .
Step 2	UCS-A /org # create stats-threshold-policy <i>policy-name</i>	Creates the specified statistics threshold policy and enters organization statistics threshold policy mode.
Step 3	UCS-A /org/stats-threshold-policy # set descr <i>description</i>	(Optional) Provides a description for the policy. Note If your description includes spaces, special characters, or punctuation, you must begin and end your description with quotation marks. The quotation marks will not appear in the description field of any show command output.
Step 4	UCS-A /org/stats-threshold-policy # commit-buffer	Commits the transaction to the system configuration.

The following example creates the server and server component statistics threshold policy named ServStatsPolicy, provides a description for the policy, and commits the transaction:

```
UCS-A# scope org /
UCS-A /org* # create stats-threshold-policy ServStatsPolicy
UCS-A /org/stats-threshold-policy* # set descr "Server stats threshold policy."
UCS-A /org/stats-threshold-policy* # commit-buffer
UCS-A /org/stats-threshold-policy #
```

What to Do Next

Configure one or more policy classes for the statistics threshold policy. For more information, see "[Configuring a Server and Server Component Statistics Threshold Policy Class](#), on page 588."

Deleting a Server and Server Component Statistics Threshold Policy

Procedure

	Command or Action	Purpose
Step 1	UCS-A# scope org <i>org-name</i>	Enters organization mode for the specified organization. To enter the root organization mode, type / as the <i>org-name</i> .
Step 2	UCS-A /org # delete stats-threshold-policy <i>policy-name</i>	Deletes the specified statistics threshold policy.
Step 3	UCS-A /org # commit-buffer	Commits the transaction to the system configuration.

The following example deletes the server and server component statistics threshold policy named ServStatsPolicy and commits the transaction:

```
UCS-A# scope org /
UCS-A /org* # delete stats-threshold-policy ServStatsPolicy
UCS-A /org* # commit-buffer
UCS-A /org #
```

Configuring a Server and Server Component Statistics Threshold Policy Class

Before You Begin

Configure or identify the server and server component statistics threshold policy that will contain the policy class. For more information, see "[Configuring a Server and Server Component Statistics Threshold Policy, on page 587.](#)"

Procedure

	Command or Action	Purpose
Step 1	UCS-A# scope org <i>org-name</i>	Enters organization mode for the specified organization. To enter the root organization mode, type / as the <i>org-name</i> .
Step 2	UCS-A /org # scope stats-threshold-policy <i>policy-name</i>	Enters organization statistics threshold policy mode.
Step 3	UCS-A /org/stats-threshold-policy # create class <i>class-name</i>	Creates the specified statistics threshold policy class and enters organization statistics threshold policy class mode. The <i>class-name</i> argument can be any of the class name keywords available for the particular statistics threshold policy being configured. To see a list of the available class name keywords, enter the create class ? command in organization statistics threshold policy mode. Note You can configure multiple classes for the statistics threshold policy.

	Command or Action	Purpose
Step 4	UCS-A /org/stats-threshold-policy /class # create property <i>property-name</i>	Creates the specified statistics threshold policy class property and enters organization statistics threshold policy class property mode. The <i>property-name</i> argument can be any of the property name keywords available for the particular policy class being configured. To see a list of the available property name keywords, enter the create property ? command in organization statistics threshold policy class mode. Note You can configure multiple properties for the policy class.
Step 5	UCS-A /org/stats-threshold-policy/class/property # set normal-value <i>value</i>	Specifies the normal value for the class property. The <i>value</i> format can vary depending on the class property being configured. To see the required format, enter the set normal-value ? command in organization statistics threshold policy class property mode.
Step 6	UCS-A /org/stats-threshold-policy /class/property # create threshold-value { above-normal below-normal } { cleared condition critical info major minor warning }	Creates the specified threshold value for the class property and enters organization statistics threshold policy class property threshold value mode. Note You can configure multiple threshold values for the class property.
Step 7	UCS-A /org/stats-threshold-policy /class/property/threshold-value # set { deescalating escalating } <i>value</i>	Specifies the de-escalating or escalating class property threshold value. The <i>value</i> format can vary depending on the class property threshold value being configured. To see the required format, enter the set deescalating ? or set escalating ? command in organization statistics threshold policy class property threshold value mode. Note You can specify both de-escalating and escalating class property threshold values.
Step 8	UCS-A /org/stats-threshold-policy /class/property/threshold-value # commit-buffer	Commits the transaction to the system configuration.

The following example creates the server and server component statistics threshold policy class for CPU statistics, creates a CPU temperature property, specifies that the normal CPU temperature is 48.5° C, creates an above normal warning threshold of 50° C, and commits the transaction:

```
UCS-A# scope org /
UCS-A /org* # scope stats-threshold-policy ServStatsPolicy
UCS-A /org/stats-threshold-policy* # create class cpu-stats
UCS-A /org/stats-threshold-policy/class* # create property cpu-temp
UCS-A /org/stats-threshold-policy/class/property* # set normal-value 48.5
UCS-A /org/stats-threshold-policy/class/property* # create threshold-value above-normal
warning
UCS-A /org/stats-threshold-policy/class/property/threshold-value* # set escalating 50.0
UCS-A /org/stats-threshold-policy/class/property/threshold-value* # commit-buffer
UCS-A /org/stats-threshold-policy/class/property/threshold-value #
```

Deleting a Server and Server Component Statistics Threshold Policy Class

Procedure

	Command or Action	Purpose
Step 1	UCS-A# scope org <i>org-name</i>	Enters organization mode for the specified organization. To enter the root organization mode, type / as the <i>org-name</i> .
Step 2	UCS-A /org # scope stats-threshold-policy <i>policy-name</i>	Enters the specified statistics threshold policy.
Step 3	UCS-A /org/stats-threshold-policy # delete class <i>class-name</i>	Deletes the specified statistics threshold policy class from the policy.
Step 4	UCS-A /org/stats-threshold-policy # commit-buffer	Commits the transaction to the system configuration.

The following example deletes the server and server component statistics threshold policy class for CPU statistics and commits the transaction:

```
UCS-A# scope org /
UCS-A /org* # scope stats-threshold-policy ServStatsPolicy
UCS-A /org/stats-threshold-policy* # delete class cpu-stats
UCS-A /org/stats-threshold-policy* # commit-buffer
UCS-A /org/stats-threshold-policy #
```

Uplink Ethernet Port Statistics Threshold Policy Configuration

Configuring an Uplink Ethernet Port Statistics Threshold Policy

Procedure

	Command or Action	Purpose
Step 1	UCS-A# scope eth-uplink	Enters Ethernet uplink mode.
Step 2	UCS-A /eth-uplink # scope stats-threshold-policy default	Enters Ethernet uplink statistics threshold policy mode. Note You cannot create (or delete) an uplink Ethernet port statistics threshold policy. You can only enter (scope to) the existing default policy.
Step 3	UCS-A /eth-uplink/stats-threshold-policy # set descr <i>description</i>	(Optional) Provides a description for the policy. Note If your description includes spaces, special characters, or punctuation, you must begin and end your description with quotation marks. The quotation marks will not appear in the description field of any show command output.

	Command or Action	Purpose
Step 4	UCS-A /eth-uplink/stats-threshold-policy # commit-buffer	Commits the transaction to the system configuration.

The following example enters the default uplink Ethernet port threshold policy, provides a description for the policy, and commits the transaction:

```
UCS-A# scope eth-uplink
UCS-A /eth-uplink* # scope stats-threshold-policy default
UCS-A /eth-uplink/stats-threshold-policy* # set descr "Uplink Ethernet port stats threshold
policy."
UCS-A /eth-uplink/stats-threshold-policy* # commit-buffer
UCS-A /eth-uplink/stats-threshold-policy #
```

What to Do Next

Configure one or more policy classes for the statistics threshold policy. For more information, see "[Configuring an Uplink Ethernet Port Statistics Threshold Policy Class](#), on page 591."

Configuring an Uplink Ethernet Port Statistics Threshold Policy Class

Procedure

	Command or Action	Purpose
Step 1	UCS-A# scope eth-uplink	Enters Ethernet uplink mode.
Step 2	UCS-A /eth-uplink # scope stats-threshold-policy default	Enters Ethernet uplink statistics threshold policy mode.
Step 3	UCS-A /eth-uplink/stats-threshold-policy # create class class-name	Creates the specified statistics threshold policy class and enters Ethernet uplink statistics threshold policy class mode. The <i>class-name</i> argument can be any of the class name keywords available for the particular statistics threshold policy being configured. To see a list of the available class name keywords, enter the create class ? command in Ethernet uplink statistics threshold policy mode. Note You can configure multiple classes for the statistics threshold policy.
Step 4	UCS-A /eth-uplink/stats-threshold-policy /class # create property property-name	Creates the specified statistics threshold policy class property and enters Ethernet uplink statistics threshold policy class property mode. The <i>property-name</i> argument can be any of the property name keywords available for the particular property class being configured. To see a list of the available property name keywords, enter the create property ? command in Ethernet uplink statistics threshold policy class mode. Note You can configure multiple properties for the policy class.

	Command or Action	Purpose
Step 5	UCS-A /eth-uplink/stats-threshold-policy /class/property # set normal-value <i>value</i>	Specifies the normal value for the class property. The <i>value</i> format can vary depending on the class property being configured. To see the required format, enter the set normal-value ? command in Ethernet uplink statistics threshold policy class property mode.
Step 6	UCS-A /eth-uplink/stats-threshold-policy /class/property # create threshold-value { above-normal below-normal } { cleared condition critical info major minor warning }	Creates the specified threshold value for the class property and enters Ethernet uplink statistics threshold policy class property threshold value mode. Note You can configure multiple threshold values for the class property.
Step 7	UCS-A /eth-uplink/stats-threshold-policy /class/property/threshold-value # set { deescalating escalating } <i>value</i>	Specifies the de-escalating or escalating class property threshold value. The <i>value</i> format can vary depending on the class property threshold value being configured. To see the required format, enter the set deescalating ? or set escalating ? command in Ethernet uplink statistics threshold policy class property threshold value mode. Note You can specify both de-escalating and escalating class property threshold values.
Step 8	UCS-A /eth-uplink/stats-threshold-policy /class/property/threshold-value # commit-buffer	Commits the transaction to the system configuration.

The following example creates the uplink Ethernet port statistics threshold policy class for Ethernet error statistics, creates a cyclic redundancy check (CRC) error count property, specifies that the normal CRC error count for each polling interval is 1000, creates an above normal warning threshold of 1250, and commits the transaction:

```
UCS-A# scope eth-uplink
UCS-A /eth-uplink* # scope stats-threshold-policy default
UCS-A /eth-uplink/stats-threshold-policy* # create class ether-error-stats
UCS-A /eth-uplink/stats-threshold-policy/class* # create property crc-delta
UCS-A /eth-uplink/stats-threshold-policy/class/property* # set normal-value 1000
UCS-A /eth-uplink/stats-threshold-policy/class/property* # create threshold-value above-normal
warning
UCS-A /eth-uplink/stats-threshold-policy/class/property/threshold-value* # set escalating
1250
UCS-A /eth-uplink/stats-threshold-policy/class/property/threshold-value* # commit-buffer
UCS-A /eth-uplink/stats-threshold-policy/class/property/threshold-value #
```

Deleting an Uplink Ethernet Port Statistics Threshold Policy Class

Procedure

	Command or Action	Purpose
Step 1	UCS-A# scope eth-uplink	Enters Ethernet uplink mode.
Step 2	UCS-A /eth-uplink # scope stats-threshold-policy default	Enters Ethernet uplink statistics threshold policy mode.
Step 3	UCS-A /eth-uplink/stats-threshold-policy # delete class class-name	Deletes the specified statistics threshold policy class from the policy.
Step 4	UCS-A /eth-uplink/stats-threshold-policy # commit-buffer	Commits the transaction to the system configuration.

The following example deletes the uplink Ethernet port statistics threshold policy class for Ethernet error statistics and commits the transaction:

```
UCS-A# scope eth-uplink
UCS-A /eth-uplink # scope stats-threshold-policy default
UCS-A /eth-uplink/stats-threshold-policy # delete class ether-error-stats
UCS-A /eth-uplink/stats-threshold-policy* # commit-buffer
UCS-A /eth-uplink/stats-threshold-policy #
```

Server Port, Chassis, and Fabric Interconnect Statistics Threshold Policy Configuration

Configuring a Server Port, Chassis, and Fabric Interconnect Statistics Threshold Policy

Procedure

	Command or Action	Purpose
Step 1	UCS-A# scope eth-server	Enters Ethernet server mode.
Step 2	UCS-A /eth-server # scope stats-threshold-policy default	Enters Ethernet server statistics threshold policy mode. Note You cannot create (or delete) a server port, chassis, and fabric interconnect statistics threshold policy. You can only enter (scope to) the existing default policy.
Step 3	UCS-A /eth-server/stats-threshold-policy # set descr description	(Optional) Provides a description for the policy. Note If your description includes spaces, special characters, or punctuation, you must begin and end your description with quotation marks. The quotation marks will not appear in the description field of any show command output.

	Command or Action	Purpose
Step 4	UCS-A /eth-server/stats-threshold-policy # commit-buffer	Commits the transaction to the system configuration.

The following example enters the default server port, chassis, and fabric interconnect statistics threshold policy, provides a description for the policy, and commits the transaction:

```
UCS-A# scope eth-server
UCS-A /eth-server* # scope stats-threshold-policy default
UCS-A /eth-server/stats-threshold-policy* # set descr "Server port, chassis, and fabric
interconnect stats threshold policy."
UCS-A /eth-server/stats-threshold-policy* # commit-buffer
UCS-A /eth-server/stats-threshold-policy #
```

What to Do Next

Configure one or more policy classes for the statistics threshold policy. For more information, see "[Configuring a Server Port, Chassis, and Fabric Interconnect Statistics Threshold Policy Class, on page 594.](#)"

Configuring a Server Port, Chassis, and Fabric Interconnect Statistics Threshold Policy Class

Procedure

	Command or Action	Purpose
Step 1	UCS-A# scope eth-server	Enters Ethernet server mode.
Step 2	UCS-A /eth-server # scope stats-threshold-policy default	Enters Ethernet server statistics threshold policy mode.
Step 3	UCS-A /eth-server/stats-threshold-policy # create class class-name	Creates the specified statistics threshold policy class and enters Ethernet server statistics threshold policy class mode. The <i>class-name</i> argument can be any of the class name keywords available for the particular statistics threshold policy being configured. To see a list of the available class name keywords, enter the create class ? command in Ethernet server statistics threshold policy mode. Note You can configure multiple classes for the statistics threshold policy.
Step 4	UCS-A /eth-server/stats-threshold-policy /class # create property property-name	Creates the specified statistics threshold policy class property and enters Ethernet server statistics threshold policy class property mode. The <i>property-name</i> argument can be any of the property name keywords available for the particular policy class being configured. To see a list of the available property name keywords, enter the create property ? command in Ethernet server statistics threshold policy class mode. Note You can configure multiple properties for the policy class.

	Command or Action	Purpose
Step 5	UCS-A /eth-server/stats-threshold-policy /class/property # set normal-value <i>value</i>	Specifies the normal value for the class property. The <i>value</i> format can vary depending on the class property being configured. To see the required format, enter the set normal-value ? command in Ethernet server statistics threshold policy class property mode.
Step 6	UCS-A /eth-server/stats-threshold-policy /class/property # create threshold-value { above-normal below-normal } { cleared condition critical info major minor warning }	Creates the specified threshold value for the class property and enters Ethernet server statistics threshold policy class property threshold value mode. Note You can configure multiple threshold values for the class property.
Step 7	UCS-A /eth-server/stats-threshold-policy /class/property/threshold-value # set { deescalating escalating } <i>value</i>	Specifies the de-escalating or escalating class property threshold value. The <i>value</i> format can vary depending on the class property threshold value being configured. To see the required format, enter the set deescalating ? or set escalating ? command in Ethernet server statistics threshold policy class property threshold value mode. Note You can specify both de-escalating and escalating class property threshold values.
Step 8	UCS-A /eth-server/stats-threshold-policy /class/property/threshold-value # commit-buffer	Commits the transaction to the system configuration.

The following example creates the server port, chassis, and fabric interconnect statistics threshold policy class for chassis statistics, creates an input power (Watts) property, specifies that the normal power is 8kW, creates an above normal warning threshold of 11kW, and commits the transaction:

```
UCS-A# scope eth-server
UCS-A /eth-server* # scope stats-threshold-policy default
UCS-A /eth-server/stats-threshold-policy* # create class chassis-stats
UCS-A /eth-server/stats-threshold-policy/class* # create property input-power
UCS-A /eth-server/stats-threshold-policy/class/property* # set normal-value 8000.0
UCS-A /eth-server/stats-threshold-policy/class/property* # create threshold-value above-normal
warning
UCS-A /eth-server/stats-threshold-policy/class/property/threshold-value* # set escalating
11000.0
UCS-A /eth-server/stats-threshold-policy/class/property/threshold-value* # commit-buffer
UCS-A /eth-server/stats-threshold-policy/class/property/threshold-value #
```

Deleting a Server Port, Chassis, and Fabric Interconnect Statistics Threshold Policy Class

Procedure

	Command or Action	Purpose
Step 1	UCS-A# scope eth-server	Enters Ethernet server mode.
Step 2	UCS-A /eth-server # scope stats-threshold-policy default	Enters Ethernet server statistics threshold policy mode.
Step 3	UCS-A /eth-server/stats-threshold-policy # delete class class-name	Deletes the specified statistics threshold policy class from the policy.
Step 4	UCS-A /eth-server/stats-threshold-policy # commit-buffer	Commits the transaction to the system configuration.

The following example deletes the server port, chassis, and fabric interconnect statistics threshold policy class for chassis statistics and commits the transaction:

```
UCS-A# scope eth-server
UCS-A /eth-server* # scope stats-threshold-policy default
UCS-A /eth-server/stats-threshold-policy* # delete class chassis-stats
UCS-A /eth-server/stats-threshold-policy* # commit-buffer
UCS-A /eth-server/stats-threshold-policy #
```

Fibre Channel Port Statistics Threshold Policy Configuration

Configuring a Fibre Channel Port Statistics Threshold Policy

Procedure

	Command or Action	Purpose
Step 1	UCS-A# scope fc-uplink	Enters Fibre Channel uplink mode.
Step 2	UCS-A /fc-uplink # scope stats-threshold-policy default	Enters Fibre Channel uplink statistics threshold policy mode. Note You cannot create (or delete) an uplink Fibre Channel port statistics threshold policy. You can only enter (scope to) the existing default policy.
Step 3	UCS-A /fc-uplink/stats-threshold-policy # set descr description	(Optional) Provides a description for the policy. Note If your description includes spaces, special characters, or punctuation, you must begin and end your description with quotation marks. The quotation marks will not appear in the description field of any show command output.

	Command or Action	Purpose
Step 4	UCS-A /fc-uplink/stats-threshold-policy # commit-buffer	Commits the transaction to the system configuration.

The following example enters the default uplink Fibre Channel port statistics threshold policy, provides a description for the policy, and commits the transaction:

```
UCS-A# scope fc-uplink
UCS-A /fc-uplink* # scope stats-threshold-policy default
UCS-A /fc-uplink/stats-threshold-policy* # set descr "Uplink Fibre Channel stats threshold
policy."
UCS-A /fc-uplink/stats-threshold-policy* # commit-buffer
UCS-A /fc-uplink/stats-threshold-policy #
```

What to Do Next

Configure one or more policy classes for the statistics threshold policy. For more information, see "[Configuring a Fibre Channel Port Statistics Threshold Policy Class, on page 597.](#)"

Configuring a Fibre Channel Port Statistics Threshold Policy Class

Procedure

	Command or Action	Purpose
Step 1	UCS-A# scope fc-uplink	Enters Fibre Channel uplink mode.
Step 2	UCS-A /fc-uplink # scope stats-threshold-policy default	Enters Fibre Channel uplink statistics threshold policy mode.
Step 3	UCS-A /fc-uplink/stats-threshold-policy # create class class-name	Creates the specified statistics threshold policy class and enters Fibre Channel uplink statistics threshold policy class mode. The <i>class-name</i> argument can be any of the class name keywords available for the particular statistics threshold policy being configured. To see a list of the available class name keywords, enter the create class ? command in Fibre Channel uplink statistics threshold policy mode. Note You can configure multiple classes for the statistics threshold policy.
Step 4	UCS-A /fc-uplink/stats-threshold-policy /class # create property property-name	Creates the specified statistics threshold policy class property and enters Fibre Channel uplink statistics threshold policy class property mode. The <i>property-name</i> argument can be any of the property name keywords available for the particular policy class being configured. To see a list of the available property name keywords, enter the create property ? command in Fibre Channel uplink statistics threshold policy class mode. Note You can configure multiple properties for the policy class.

	Command or Action	Purpose
Step 5	UCS-A /fc-uplink/stats-threshold-policy /class/property # set normal-value <i>value</i>	Specifies the normal value for the class property. The <i>value</i> format can vary depending on the class property being configured. To see the required format, enter the set normal-value ? command in Fibre Channel uplink statistics threshold policy class property mode.
Step 6	UCS-A /fc-uplink/stats-threshold-policy /class/property # create threshold-value { above-normal below-normal } { cleared condition critical info major minor warning }	Creates the specified threshold value for the class property and enters Fibre Channel uplink statistics threshold policy class property threshold value mode. Note You can configure multiple threshold values for the class property.
Step 7	UCS-A /fc-uplink/stats-threshold-policy /class/property/threshold-value # set { deescalating escalating } <i>value</i>	Specifies the de-escalating or escalating class property threshold value. The <i>value</i> format can vary depending on the class property threshold value being configured. To see the required format, enter the set deescalating ? or set escalating ? command in Fibre Channel uplink statistics threshold policy class property threshold value mode. Note You can specify both de-escalating and escalating class property threshold values.
Step 8	UCS-A /fc-uplink/stats-threshold-policy /class/property/threshold-value # commit-buffer	Commits the transaction to the system configuration.

The following example creates the uplink Fibre Channel port statistics threshold policy class for Fibre Channel statistics, creates an average bytes received property, specifies that the normal average number of bytes received for each polling interval is 150MB, creates an above normal warning threshold of 200MB, and commits the transaction:

```
UCS-A# scope fc-uplink
UCS-A /fc-uplink* # scope stats-threshold-policy default
UCS-A /fc-uplink/stats-threshold-policy* # create class fc-stats
UCS-A /fc-uplink/stats-threshold-policy/class* # create property bytes-rx-avg
UCS-A /fc-uplink/stats-threshold-policy/class/property* # set normal-value 150000000
UCS-A /fc-uplink/stats-threshold-policy/class/property* # create threshold-value above-normal
warning
UCS-A /fc-uplink/stats-threshold-policy/class/property/threshold-value* # set escalating
200000000
UCS-A /fc-uplink/stats-threshold-policy/class/property/threshold-value* # commit-buffer
UCS-A /fc-uplink/stats-threshold-policy/class/property/threshold-value #
```


Deleting an Uplink Fibre Channel Port Statistics Threshold Policy Class

Procedure

	Command or Action	Purpose
Step 1	UCS-A# scope fc-uplink	Enters Fibre Channel uplink mode.
Step 2	UCS-A /fc-uplink # scope stats-threshold-policy default	Enters Fibre Channel uplink statistics threshold policy mode.
Step 3	UCS-A /fc-uplink/stats-threshold-policy # delete class class-name	Deletes the specified statistics threshold policy class from the policy.
Step 4	UCS-A /fc-uplink/stats-threshold-policy # commit-buffer	Commits the transaction to the system configuration.

The following example deletes the uplink Fibre Channel port statistics threshold policy class for Fibre Channel statistics and commits the transaction:

```
UCS-A# scope fc-uplink
UCS-A /fc-uplink # scope stats-threshold-policy default
UCS-A /fc-uplink/stats-threshold-policy # delete class fc-stats
UCS-A /fc-uplink/stats-threshold-policy* # commit-buffer
UCS-A /fc-uplink/stats-threshold-policy #
```




CHAPTER 43

Configuring Call Home

This chapter includes the following sections:

- [Call Home, page 601](#)
- [Call Home Considerations and Guidelines, page 603](#)
- [Cisco UCS Faults and Call Home Severity Levels, page 604](#)
- [Cisco Smart Call Home, page 605](#)
- [Configuring Call Home, page 606](#)
- [Disabling Call Home, page 608](#)
- [Enabling Call Home, page 608](#)
- [Configuring System Inventory Messages, page 609](#)
- [Configuring Call Home Profiles, page 610](#)
- [Sending a Test Call Home Alert, page 613](#)
- [Configuring Call Home Policies, page 614](#)
- [Example: Configuring Call Home for Smart Call Home, page 616](#)

Call Home

Call Home provides an email-based notification for critical system policies. A range of message formats are available for compatibility with pager services or XML-based automated parsing applications. You can use this feature to page a network support engineer, email a Network Operations Center, or use Cisco Smart Call Home services to generate a case with the Technical Assistance Center.

The Call Home feature can deliver alert messages containing information about diagnostics and environmental faults and events.

The Call Home feature can deliver alerts to multiple recipients, referred to as Call Home destination profiles. Each profile includes configurable message formats and content categories. A predefined destination profile is provided for sending alerts to the Cisco TAC, but you also can define your own destination profiles.

When you configure Call Home to send messages, Cisco UCS Manager executes the appropriate CLI **show** command and attaches the command output to the message.


Cisco UCS delivers Call Home messages in the following formats:

- Short text format which provides a one or two line description of the fault that is suitable for pagers or printed reports.
- Full text format which provides fully formatted message with detailed information that is suitable for human reading.
- XML machine readable format that uses Extensible Markup Language (XML) and Adaptive Messaging Language (AML) XML schema definition (XSD). The AML XSD is published on the [Cisco.com website](#). The XML format enables communication with the Cisco Systems Technical Assistance Center.

For information about the faults that can trigger Call Home email alerts, see the *Cisco UCS Faults and Error Messages Reference*.

The following figure shows the flow of events after a Cisco UCS fault is triggered in a system with Call Home configured:

Figure 2: Flow of Events after a Fault is Triggered



Call Home Considerations and Guidelines

How you configure Call Home depends on how you intend to use the feature. The information you need to consider before you configure Call Home includes the following:

Destination Profile

You must configure at least one destination profile. The destination profile or profiles that you use depend upon whether the receiving entity is a pager, email, or automated service such as Cisco Smart Call Home.

If the destination profile uses email message delivery, you must specify a Simple Mail Transfer Protocol (SMTP) server when you configure Call Home.

Contact Information

The contact email, phone, and street address information should be configured so that the receiver can determine the origin of messages received from the Cisco UCS domain.

Cisco Smart Call Home sends the registration email to this email address after you send a system inventory to begin the registration process.

If an email address includes special characters, such as # (hash), spaces, or & (ampersand), the email server may not be able to deliver email messages to that address. Cisco recommends that you use email addresses which comply with RFC2821 and RFC2822 and include only 7bit ASCII characters.

IP Connectivity to Email Server or HTTP Server

The fabric interconnect must have IP connectivity to an email server or the destination HTTP server. In a cluster configuration, both fabric interconnects must have IP connectivity. This connectivity ensures that the current, active fabric interconnect can send Call Home email messages. The source of these email messages is always the IP address of a fabric interconnect. The virtual IP address assigned Cisco UCS Manager in a cluster configuration is never the source of the email.

Smart Call Home

If Cisco Smart Call Home is used, the following are required:

- An active service contract must cover the device being configured
- The customer ID associated with the Smart Call Home configuration in Cisco UCS must be the CCO (Cisco.com) account name associated with a support contract that includes Smart Call Home

Cisco UCS Faults and Call Home Severity Levels

Because Call Home is present across several Cisco product lines, Call Home has developed its own standardized severity levels. The following table describes how the underlying Cisco UCS fault levels map to the Call Home severity levels. You need to understand this mapping when you configure the Level setting for Call Home profiles.

Table 15: Mapping of Faults and Call Home Severity Levels

Call Home Severity	Cisco UCS Fault	Call Home Meaning
(9) Catastrophic	N/A	Network-wide catastrophic failure.
(8) Disaster	N/A	Significant network impact.
(7) Fatal	N/A	System is unusable.

Call Home Severity	Cisco UCS Fault	Call Home Meaning
(6) Critical	Critical	Critical conditions, immediate attention needed.
(5) Major	Major	Major conditions.
(4) Minor	Minor	Minor conditions.
(3) Warning	Warning	Warning conditions.
(2) Notification	Info	Basic notifications and informational messages. Possibly independently insignificant.
(1) Normal	Clear	Normal event, signifying a return to normal state.
(0) debug	N/A	Debugging messages.

Cisco Smart Call Home

Cisco Smart Call Home is a web application which leverages the Call Home feature of Cisco UCS. Smart Call Home offers proactive diagnostics and real-time email alerts of critical system events, which results in higher network availability and increased operational efficiency. Smart Call Home is a secure connected service offered by Cisco Unified Computing Support Service and Cisco Unified Computing Mission Critical Support Service for Cisco UCS.



Note Using Smart Call Home requires the following:

- A CCO ID associated with a corresponding Cisco Unified Computing Support Service or Cisco Unified Computing Mission Critical Support Service contract for your company.
- Cisco Unified Computing Support Service or Cisco Unified Computing Mission Critical Support Service for the device to be registered.

You can configure and register Cisco UCS Manager to send Smart Call Home email alerts to either the Smart Call Home System or the secure Transport Gateway. Email alerts sent to the secure Transport Gateway are forwarded to the Smart Call Home System using HTTPS.



Note For security reasons, we recommend using the Transport Gateway option. The Transport Gateway can be downloaded from Cisco.

To configure Smart Call Home, you must do the following:

- Enable the Smart Call Home feature.

- Configure the contact information.
- Configure the email information.
- Configure the SMTP server information.
- Configure the default CiscoTAC-1 profile.
- Send a Smart Call Home inventory message to start the registration process.
- Ensure that the CCO ID you plan to use as the Call Home Customer ID for the Cisco UCS domain has the contract numbers from the registration added to its entitlements. You can update the ID in the account properties under Additional Access in the Profile Manager on CCO.

Configuring Call Home

Procedure

	Command or Action	Purpose
Step 1	UCS-A# scope monitoring	Enters monitoring mode.
Step 2	UCS-A /monitoring # scope callhome	Enters monitoring call home mode.
Step 3	UCS-A /monitoring/callhome # enable	Enables Call Home.
Step 4	UCS-A /monitoring/callhome # set contact name	Specifies the name of the main Call Home contact person.
Step 5	UCS-A /monitoring/callhome # set email email-addr	Specifies the email address of the main Call Home contact person. Note If an email address includes special characters, such as # (hash), spaces, or & (ampersand), the email server may not be able to deliver email messages to that address. Cisco recommends that you use email addresses which comply with RFC2821 and RFC2822 and include only 7bit ASCII characters.
Step 6	UCS-A /monitoring/callhome # set phone-contact phone-num	Specifies the phone number of the main Call Home contact person. The phone number must be in international format, starting with a + (plus sign) and a country code.
Step 7	UCS-A /monitoring/callhome # set street-address street-addr	Specifies the street address of the main Call Home contact person. Enter up to 255 ASCII characters.
Step 8	UCS-A /monitoring/callhome # set customer-id id-num	Specifies the CCO identification number that includes the contract numbers for the support contract in its entitlements. The number can be up to 255 alphanumeric characters in free format.

	Command or Action	Purpose
Step 9	UCS-A /monitoring/callhome # set contract-id <i>id-num</i>	Specifies the contract identification number from the service agreement. The number can be up to 255 alphanumeric characters in free format.
Step 10	UCS-A /monitoring/callhome # set site-id <i>id-num</i>	Specifies the site identification number from the service agreement. The number can be up to 255 alphanumeric characters in free format.
Step 11	UCS-A /monitoring/callhome # set from-email <i>email-addr</i>	Specifies the email address to use for the From field in Call Home messages.
Step 12	UCS-A /monitoring/callhome # set reply-to-email <i>email-addr</i>	Specifies the email address to use for the Reply To field in Call Home messages.
Step 13	UCS-A /monitoring/callhome # set hostname { <i>hostname</i> <i>ip-addr</i> }	Specifies the hostname or IP address of the SMTP server that Call Home uses to send email messages.
Step 14	UCS-A /monitoring/callhome # set port <i>port-num</i>	Specifies the SMTP server port that Call Home uses to send email messages. Valid port numbers are 1 to 65535.
Step 15	UCS-A /monitoring/callhome # set throttling { <i>off</i> <i>on</i> }	Enables or disables Call Home throttling. When enabled, throttling prevents too many Call Home email messages from being sent for the same event. By default, throttling is enabled.
Step 16	UCS-A /monitoring/callhome # set urgency { <i>alerts</i> <i>critical</i> <i>debugging</i> <i>emergencies</i> <i>errors</i> <i>information</i> <i>notifications</i> <i>warnings</i> }	Specifies the urgency level for Call Home email messages. In the context of a large UCS deployment with several pairs of fabric interconnects, the urgency level potentially allows you to attach significance to Call Home messages from one particular Cisco UCS domain versus another. In the context of a small UCS deployment involving only two fabric interconnects, the urgency level holds little meaning.
Step 17	UCS-A /monitoring/callhome # commit-buffer	Commits the transaction to the system configuration.

The following example configures Call Home and commits the transaction:

```
UCS-A# scope monitoring
UCS-A /monitoring* # scope callhome
UCS-A /monitoring/callhome* # enable
UCS-A /monitoring/callhome* # set contact "Steve Jones"
UCS-A /monitoring/callhome* # set email admin@MyCompany.com
UCS-A /monitoring/callhome* # set phone-contact +1-001-408-555-1234
UCS-A /monitoring/callhome* # set street-address "123 N. Main Street, Anytown, CA, 99885"
UCS-A /monitoring/callhome* # set customer-id 1234567
UCS-A /monitoring/callhome* # set contract-id 99887766
UCS-A /monitoring/callhome* # set site-id 5432112
UCS-A /monitoring/callhome* # set from-email person@MyCompany.com
UCS-A /monitoring/callhome* # set reply-to-email person@MyCompany.com
UCS-A /monitoring/callhome* # set hostname 192.168.100.12
UCS-A /monitoring/callhome* # set port 25
UCS-A /monitoring/callhome* # set throttling on
```

```
UCS-A /monitoring/callhome* # set urgency information
UCS-A /monitoring/callhome* # commit-buffer
UCS-A /monitoring/callhome #
```

Disabling Call Home

Procedure

	Command or Action	Purpose
Step 1	UCS-A# scope monitoring	Enters monitoring mode.
Step 2	UCS-A /monitoring # scope callhome	Enters monitoring call home mode.
Step 3	UCS-A /monitoring/callhome # disable	Enables Call Home.
Step 4	UCS-A /monitoring/callhome # commit-buffer	Commits the transaction to the system configuration.

The following example disables Call Home and commits the transaction:

```
UCS-A# scope monitoring
UCS-A /monitoring # scope callhome
UCS-A /monitoring/callhome # disable
UCS-A /monitoring/callhome* # commit-buffer
UCS-A /monitoring/callhome #
```

Enabling Call Home

Procedure

	Command or Action	Purpose
Step 1	UCS-A# scope monitoring	Enters monitoring mode.
Step 2	UCS-A /monitoring # scope callhome	Enters monitoring call home mode.
Step 3	UCS-A /monitoring/callhome # enable	Enables Call Home.
Step 4	UCS-A /monitoring/callhome # commit-buffer	Commits the transaction to the system configuration.

The following example enables Call Home and commits the transaction:

```
UCS-A# scope monitoring
UCS-A /monitoring # scope callhome
UCS-A /monitoring/callhome # enable
UCS-A /monitoring/callhome* # commit-buffer
UCS-A /monitoring/callhome #
```

Configuring System Inventory Messages

Configuring System Inventory Messages

Procedure

	Command or Action	Purpose
Step 1	UCS-A# scope monitoring	Enters monitoring mode.
Step 2	UCS-A /monitoring # scope callhome	Enters monitoring call home mode.
Step 3	UCS-A /monitoring/callhome # scope inventory	Enters monitoring call home inventory mode.
Step 4	UCS-A /monitoring/callhome/inventory # set send-periodically {off on}	Enables or disables the sending of inventory messages. When the on keyword is specified, inventory messages are automatically sent to the Call Home database.
Step 5	UCS-A /monitoring/callhome/inventory # set interval-days interval-num	Specifies the time interval (in days) at which inventory messages will be sent.
Step 6	UCS-A /monitoring/callhome/inventory # set timeofday-hour hour	Specifies the hour (using 24-hour format) that inventory messages are sent.
Step 7	UCS-A /monitoring/callhome/inventory # set timeofday-minute minute	Specifies the number of minutes after the hour that inventory messages are sent.
Step 8	UCS-A /monitoring/callhome/inventory # commit-buffer	Commits the transaction to the system configuration.

The following example configures Call Home system inventory messages and commits the transaction:

```
UCS-A# scope monitoring
UCS-A /monitoring* # scope callhome
UCS-A /monitoring/callhome* # scope inventory
UCS-A /monitoring/callhome/inventory* # set send-periodically on
UCS-A /monitoring/callhome/inventory* # set interval-days 15
UCS-A /monitoring/callhome/inventory* # set timeofday-hour 21
UCS-A /monitoring/callhome/inventory* # set timeofday-minute 30
UCS-A /monitoring/callhome/inventory* # commit-buffer
UCS-A /monitoring/callhome/inventory #
```

Sending a System Inventory Message

Use this procedure if you need to manually send a system inventory message outside of the scheduled messages.



Note

The system inventory message is sent only to those recipients defined in CiscoTAC-1 profile.

Procedure

	Command or Action	Purpose
Step 1	UCS-A# scope monitoring	Enters monitoring mode.
Step 2	UCS-A /monitoring # scope callhome	Enters monitoring call home mode.
Step 3	UCS-A /monitoring/callhome # scope inventory	Enters monitoring call home inventory mode.
Step 4	UCS-A /monitoring/callhome/inventory # send	Sends the system inventory message to the Call Home database.

The following example sends the system inventory message to the Call Home database:

```
UCS-A# scope monitoring
UCS-A /monitoring # scope callhome
UCS-A /monitoring/callhome # scope inventory
UCS-A /monitoring/callhome/inventory* # send
```

Configuring Call Home Profiles

Call Home Profiles

Call Home profiles determine which alerts are sent to designated recipients. You can configure the profiles to send email alerts for events and faults at a desired severity level and for specific alert groups that represent categories of alerts. You can also use these profiles to specify the format of the alert for a specific set of recipients and alert groups.

Each alert that Cisco UCS generates fits into a category represented by an alert group, such as the following:

- Environmental alerts, including fans and power supplies
- Diagnostic alerts, such as POST completion failure on a server

Alert groups and Call Home profiles enable you to filter the alerts and ensure that a specific profile only receives certain categories of alerts. For example, a data center may have a hardware team that handles issues with fans and power supplies. This hardware team does not care about server POST failures or licensing issues. To ensure that the hardware team only receives relevant alerts, create a Call Home profile for the hardware team and check only the "environmental" alert group.

By default, you must configure the Cisco TAC-1 profile. However, you can also create additional profiles to send email alerts to one or more alert groups when events occur at the level that you specify and provide the recipients with the appropriate amount of information about those alerts.

For example, you may want to configure two profiles for faults with a major severity:

- A profile that sends an alert to the Supervisor alert group in the short text format. Members of this group receive a one- or two-line description of the fault that they can use to track the issue.

- A profile that sends an alert to the CiscoTAC alert group in the XML format. Members of this group receive a detailed message in the machine readable format preferred by the Cisco Systems Technical Assistance Center.

Configuring a Call Home Profile

By default, you must configure the Cisco TAC-1 profile. However, you can also create additional profiles to send email alerts to one or more specified groups when events occur at the level that you specify.

Procedure

	Command or Action	Purpose
Step 1	UCS-A# scope monitoring	Enters monitoring mode.
Step 2	UCS-A /monitoring # scope callhome	Enters monitoring call home mode.
Step 3	UCS-A /monitoring/callhome # create profile profile-name	Enters monitoring call home profile mode.
Step 4	UCS-A /monitoring/callhome/profile # set level {critical debug disaster fatal major minor normal notification warning}	Specifies the event level for the profile. Each profile can have its own unique event level. Cisco UCS faults that are greater than or equal to the event level will trigger this profile.
Step 5	UCS-A /monitoring/callhome/profile # set alertgroups group-name <ul style="list-style-type: none"> • ciscotac • diagnostic • environmental • inventory • license • lifecycle • linecard • supervisor • syslogport • system • test 	Specifies one or more groups that are alerted based on the profile. The <i>group-name</i> argument can be one or more of the following keywords entered on the same command line:
Step 6	UCS-A /monitoring/callhome/profile # add alertgroups group-names	(Optional) Adds one or more groups to the existing list of groups that are alerted based on the Call Home profile.

	Command or Action	Purpose
		Note You must use the add alertgroups command to add more alert groups to the existing alert group list. Using the set alertgroups command will replace any pre-existing alert groups with a new group list.
Step 7	UCS-A /monitoring/callhome/profile # set format {shorttxt xml}	Specifies the formatting method to use for the e-mail messages.
Step 8	UCS-A /monitoring/callhome/profile # set maxsize id-num	Specifies the maximum size (in characters) of the email message.
Step 9	UCS-A /monitoring/callhome/profile # create destination email-addr	Specifies the email address to which Call Home alerts should be sent. Use multiple create destination commands in monitoring call home profile mode to specify multiple email recipients. Use the delete destination command in monitoring call home profile mode to delete a specified email recipient.
Step 10	UCS-A /monitoring/callhome/profile/destination # commit-buffer	Commits the transaction to the system configuration.

The following example configures a Call Home profile and commits the transaction:

```
UCS-A# scope monitoring
UCS-A /monitoring* # scope callhome
UCS-A /monitoring/callhome* # create profile TestProfile
UCS-A /monitoring/callhome/profile* # set level normal
UCS-A /monitoring/callhome/profile* # set alertgroups test diagnostic
UCS-A /monitoring/callhome/profile* # set format xml
UCS-A /monitoring/callhome/profile* # set maxsize 100000
UCS-A /monitoring/callhome/profile* # create destination admin@MyCompany.com
UCS-A /monitoring/callhome/profile/destination* # commit-buffer
UCS-A /monitoring/callhome/profile/destination #
```

Deleting a Call Home Profile

Procedure

	Command or Action	Purpose
Step 1	UCS-A# scope monitoring	Enters monitoring mode.
Step 2	UCS-A /monitoring # scope callhome	Enters monitoring call home mode.
Step 3	UCS-A /monitoring/callhome # delete profile <i>profile-name</i>	Deletes the specified profile.

	Command or Action	Purpose
Step 4	UCS-A /monitoring/callhome # commit-buffer	Commits the transaction to the system configuration.

The following example deletes the Call Home profile named TestProfile and commits the transaction:

```
UCS-A# scope monitoring
UCS-A /monitoring # scope callhome
UCS-A /monitoring/callhome # delete profile TestProfile
UCS-A /monitoring/callhome* # commit-buffer
UCS-A /monitoring/callhome #
```

Sending a Test Call Home Alert

Before You Begin

Configure Call Home and a Call Home Profile

Procedure

	Command or Action	Purpose
Step 1	UCS-A# scope monitoring	Enters monitoring mode.
Step 2	UCS-A /monitoring # scope callhome	Enters monitoring call home mode.
Step 3	UCS-A /monitoring/callhome # send-test-alert {[alert-group { diagnostic environmental }] [alert-level { critical debug fatal major minor normal notify warning }] [alert-message-type { conf diag env inventory syslog test }] [alert-message-subtype { delta full goldmajor goldminor goldnormal major minor nosubtype test }] [alert-description <i>description</i>]}	Sends a test Call Home alert. The test Call Home alert must specify all alert-* parameters or Cisco UCS Manager cannot generate the test message. The alert-* parameters include the following : <ul style="list-style-type: none"> • alert-description—Alert description • alert-group—Alert group • alert-level—Event severity level • alert-message-type—Message type • alert-message-subtypeMessage subtype <p>When a test Call Home alert is sent, Call Home responds as it would to any other alert and delivers it to the configured destination email addresses.</p>

The following example sends a test Call Home alert to the configured destination email address of the environmental alert group:

```
UCS-A# scope monitoring
UCS-A /monitoring # scope callhome
UCS-A /monitoring/callhome # send-test-alert alert-description "This is a test alert"
alert-group diagnostic alert-level critical alert-message-subtype
```

```
major alert-message-type diag
```

Configuring Call Home Policies

Call Home Policies

Call Home policies determine whether or not Call Home alerts are sent for a specific type of fault or system event. By default, Call Home is enabled to send alerts for certain types of faults and system events. However, you can configure Cisco UCS not to process certain types.

To disable alerts for a type of fault or events, you must create a Call Home policy for that type, and you must first create a policy for that type and then disable the policy.

Configuring a Call Home Policy



Tip

By default, email alerts are sent for all critical system events. However, you can optionally configure Call Home policies to enable or disable sending email alerts for other critical system events.

Procedure

	Command or Action	Purpose
Step 1	UCS-A# scope monitoring	Enters monitoring mode.
Step 2	UCS-A /monitoring # scope callhome	Enters monitoring call home mode.
Step 3	UCS-A /monitoring/callhome # create policy { equipment-inoperable fru-problem identity-unestablishable thermal-problem voltage-problem }	Creates the specified policy and enters monitoring call home policy mode.
Step 4	UCS-A /monitoring/callhome/policy # { disabled enabled }	Disables or enables the sending of email alerts for the specified policy.
Step 5	UCS-A /monitoring/callhome/policy # commit-buffer	Commits the transaction to the system configuration.

The following example creates a Call Home policy that disables the sending of email alerts for system events pertaining to voltage problems and commits the transaction:

```
UCS-A# scope monitoring
UCS-A /monitoring* # scope callhome
UCS-A /monitoring/callhome* # create policy voltage-problem
UCS-A /monitoring/callhome/policy* # disabled
UCS-A /monitoring/callhome/policy* # commit-buffer
UCS-A /monitoring/callhome/policy #
```


Disabling a Call Home Policy

Procedure

	Command or Action	Purpose
Step 1	UCS-A# scope monitoring	Enters monitoring mode.
Step 2	UCS-A /monitoring # scope callhome	Enters monitoring call home mode.
Step 3	UCS-A /monitoring/callhome # scope policy { equipment-inoperable fru-problem identity-unestablishable thermal-problem voltage-problem }	Enters monitoring call home policy mode for the specified policy.
Step 4	UCS-A /monitoring/callhome/policy # disable	Disables the specified policy.
Step 5	UCS-A /monitoring/callhome/policy # commit-buffer	Commits the transaction to the system configuration.

The following example disables the Call Home policy named voltage-problem and commits the transaction:

```
UCS-A# scope monitoring
UCS-A /monitoring # scope callhome
UCS-A /monitoring/callhome # scope policy voltage-problem
UCS-A /monitoring/callhome/policy # disable
UCS-A /monitoring/callhome/policy* # commit-buffer
UCS-A /monitoring/callhome/policy #
```

Enabling a Call Home Policy

Procedure

	Command or Action	Purpose
Step 1	UCS-A# scope monitoring	Enters monitoring mode.
Step 2	UCS-A /monitoring # scope callhome	Enters monitoring call home mode.
Step 3	UCS-A /monitoring/callhome # scope policy { equipment-inoperable fru-problem identity-unestablishable thermal-problem voltage-problem }	Enters monitoring call home policy mode for the specified policy.
Step 4	UCS-A /monitoring/callhome/policy # enable	Enables the specified policy.
Step 5	UCS-A /monitoring/callhome/policy # commit-buffer	Commits the transaction to the system configuration.

The following example enables the Call Home policy named voltage-problem and commits the transaction:

```
UCS-A# scope monitoring
UCS-A /monitoring # scope callhome
UCS-A /monitoring/callhome # scope policy voltage-problem
UCS-A /monitoring/callhome/policy # enable
UCS-A /monitoring/callhome/policy* # commit-buffer
UCS-A /monitoring/callhome/policy #
```

Deleting a Call Home Policy

Procedure

	Command or Action	Purpose
Step 1	UCS-A# scope monitoring	Enters monitoring mode.
Step 2	UCS-A /monitoring # scope callhome	Enters monitoring call home mode.
Step 3	UCS-A /monitoring/callhome # delete policy { equipment-inoperable fru-problem identity-unestablishable thermal-problem voltage-problem }	Deletes the specified policy
Step 4	UCS-A /monitoring/callhome # commit-buffer	Commits the transaction to the system configuration.

The following example deletes the Call Home policy named voltage-problem and commits the transaction:

```
UCS-A# scope monitoring
UCS-A /monitoring # scope callhome
UCS-A /monitoring/callhome # delete policy voltage-problems
UCS-A /monitoring/callhome* # commit-buffer
UCS-A /monitoring/callhome #
```

Example: Configuring Call Home for Smart Call Home

Configuring Smart Call Home

Procedure

	Command or Action	Purpose
Step 1	UCS-A# scope monitoring	Enters monitoring mode.
Step 2	UCS-A /monitoring # scope callhome	Enters monitoring call home mode.
Step 3	UCS-A /monitoring/callhome # enable	Enables Call Home.
Step 4	UCS-A /monitoring/callhome # set contact <i>name</i>	Cisco Smart Call Home sends the registration email to this email address.

	Command or Action	Purpose
Step 5	UCS-A /monitoring/callhome # set email <i>email-addr</i>	Specifies the email address of the main Call Home contact person. Cisco Smart Call Home sends the registration email to this email address.
Step 6	UCS-A /monitoring/callhome # set phone-contact <i>phone-num</i>	Specifies the phone number of the main Call Home contact person. The phone number must be in international format, starting with a + (plus sign) and a country code.
Step 7	UCS-A /monitoring/callhome # set street-address <i>street-addr</i>	Specifies the street address of the main Call Home contact person.
Step 8	UCS-A /monitoring/callhome # set customer-id <i>id-num</i>	Specifies the CCO identification number that includes the contract numbers for the support contract in its entitlements. The number can be up to 255 alphanumeric characters in free format.
Step 9	UCS-A /monitoring/callhome # set contract-id <i>id-num</i>	Specifies the contract identification number from the service agreement. The number can be up to 255 alphanumeric characters in free format.
Step 10	UCS-A /monitoring/callhome # set site-id <i>id-num</i>	Specifies the site identification number from the service agreement. The number can be up to 255 alphanumeric characters in free format.
Step 11	UCS-A /monitoring/callhome # set from-email <i>email-addr</i>	Specifies the email address to use for the From field in Call Home messages.
Step 12	UCS-A /monitoring/callhome # set reply-to-email <i>email-addr</i>	Specifies the email address to use for the Reply To field in Call Home messages.
Step 13	UCS-A /monitoring/callhome # set hostname { <i>hostname</i> <i>ip-addr</i> }	Specifies the hostname or IP address of the SMTP server that Call Home uses to send email messages.
Step 14	UCS-A /monitoring/callhome # set port <i>port-num</i>	Specifies the SMTP server port that Call Home uses to send email messages. Valid port numbers are 1 to 65535.
Step 15	UCS-A /monitoring/callhome # set throttling { off on }	Enables or disables Call Home throttling. When enabled, throttling prevents too many Call Home email messages from being sent for the same event. By default, throttling is enabled.
Step 16	UCS-A /monitoring/callhome # set urgency { alerts critical debugging emergencies errors information notifications warnings }	Specifies the urgency level for Call Home email messages.
Step 17	UCS-A /monitoring/callhome # commit-buffer	Commits the transaction to the system configuration.

The following example configures Call Home and commits the transaction:

```
UCS-A# scope monitoring
UCS-A /monitoring* # scope callhome
UCS-A /monitoring/callhome* # enable
UCS-A /monitoring/callhome* # set contact "Steve Jones"
UCS-A /monitoring/callhome* # set email admin@MyCompany.com
UCS-A /monitoring/callhome* # set phone-contact +1-001-408-555-1234
UCS-A /monitoring/callhome* # set street-address "123 N. Main Street, Anytown, CA, 99885"
UCS-A /monitoring/callhome* # set customer-id 1234567
UCS-A /monitoring/callhome* # set contract-id 99887766
UCS-A /monitoring/callhome* # set site-id 5432112
UCS-A /monitoring/callhome* # set from-email person@MyCompany.com
UCS-A /monitoring/callhome* # set reply-to-email person@MyCompany.com
UCS-A /monitoring/callhome* # set hostname 192.168.100.12
UCS-A /monitoring/callhome* # set port 25
UCS-A /monitoring/callhome* # set throttling on
UCS-A /monitoring/callhome* # set urgency information
UCS-A /monitoring/callhome* # commit-buffer
UCS-A /monitoring/callhome #
```

What to Do Next

Continue to ["Configuring the Default Cisco TAC-1 Profile, on page 618"](#) to configure a Call Home profile for use with Smart Call Home.

Configuring the Default Cisco TAC-1 Profile

The following are the default settings for the CiscoTAC-1 profile:

- Level is normal
- Only the CiscoTAC alert group is selected
- Format is xml
- Maximum message size is 5000000

Before You Begin

Complete the ["Configuring Smart Call Home, on page 616"](#) section.

Procedure

	Command or Action	Purpose
Step 1	UCS-A /monitoring/callhome # scope profile CiscoTac-1	Enters monitoring call home profile mode for the default Cisco TAC-1 profile.
Step 2	UCS-A /monitoring/callhome/profile # set level normal	Specifies the normal event level for the profile.
Step 3	UCS-A /monitoring/callhome/profile # set alertgroups ciscotac	Specifies the ciscotac alert group for the profile.

	Command or Action	Purpose
Step 4	UCS-A /monitoring/callhome/profile # set format xml	Specifies the e-mail message format to xml .
Step 5	UCS-A /monitoring/callhome/profile # set maxsize 5000000	Specifies the maximum size of 5000000 for email messages.
Step 6	UCS-A /monitoring/callhome/profile # create destination callhome@cisco.com	Specifies the email recipient to callhome@cisco.com .
Step 7	UCS-A /monitoring/callhome/profile/destination # exit	Exits to monitoring call home profile mode.
Step 8	UCS-A /monitoring/callhome/profile # exit	Exits to monitoring call home mode.

The following example configures the default Cisco TAC-1 profile for use with Smart Call Home:

```
UCS-A /monitoring/callhome* # scope profile CiscoTac-1
UCS-A /monitoring/callhome/profile* # set level normal
UCS-A /monitoring/callhome/profile* # set alertgroups ciscotac
UCS-A /monitoring/callhome/profile* # set format xml
UCS-A /monitoring/callhome/profile* # set maxsize 5000000
UCS-A /monitoring/callhome/profile* # create destination callhome@cisco.com
UCS-A /monitoring/callhome/profile/destination* # exit
UCS-A /monitoring/callhome/profile* # exit
UCS-A /monitoring/callhome* #
```

What to Do Next

Continue to "[Configuring a System Inventory Message for Smart Call Home, on page 619](#)" to configure system inventory messages for use with Smart Call Home.

Configuring a System Inventory Message for Smart Call Home

Before You Begin

Complete the "[Configuring the Default Cisco TAC-1 Profile, on page 618](#)" section.

Procedure

	Command or Action	Purpose
Step 1	UCS-A /monitoring/callhome # scope inventory	Enters monitoring call home inventory mode.
Step 2	UCS-A /monitoring/callhome/inventory # set send-periodically {off on}	Enables or disables the sending of inventory messages. When the on keyword is specified, inventory messages are automatically sent to the Call Home database.

	Command or Action	Purpose
Step 3	UCS-A /monitoring/callhome/inventory # set interval-days <i>interval-num</i>	Specifies the the time interval (in days) at which inventory messages will be sent.
Step 4	UCS-A /monitoring/callhome/inventory # set timeofday-hour <i>hour</i>	Specifies the hour (using 24-hour format) that inventory messages are sent.
Step 5	UCS-A /monitoring/callhome/inventory # set timeofday-minute <i>minute</i>	Specifies the number of minutes after the hour that inventory messages are sent.
Step 6	UCS-A /monitoring/callhome/inventory # commit-buffer	Commits the transaction to the system configuration.

The following example configures Call Home system inventory messages and commits the transaction:

```
UCS-A /monitoring/callhome* # scope inventory
UCS-A /monitoring/callhome/inventory* # set send-periodically on
UCS-A /monitoring/callhome/inventory* # set interval-days 15
UCS-A /monitoring/callhome/inventory* # set timeofday-hour 21
UCS-A /monitoring/callhome/inventory* # set timeofday-minute 30
UCS-A /monitoring/callhome/inventory* # commit-buffer
UCS-A /monitoring/callhome/inventory #
```

What to Do Next

Continue to "[Registering Smart Call Home, on page 620](#)" to send an inventory message that starts the Smart Call Home registration process.

Registering Smart Call Home

Before You Begin

Complete the "[Configuring a System Inventory Message for Smart Call Home, on page 619](#)" section.

Procedure

	Command or Action	Purpose
Step 1	UCS-A /monitoring/callhome/inventory # send	Sends the system inventory message to the Smart Call Home database. When Cisco receives the system inventory, a Smart Call Home registration email is sent to the email address that you configured as the email address for the main Smart Call Home contact.

The following example sends the system inventory message to the Smart Call Home database:

```
UCS-A /monitoring/callhome/inventory # send
```

What to Do Next

When you receive the registration email from Cisco, do the following to complete registration for Smart Call Home:

- 1 Click the link in the email.
The link opens the [Cisco Smart Call Home portal](#) in your web browser.
- 2 Log into the Cisco Smart Call Home portal.
- 3 Follow the steps provided by Cisco Smart Call Home.

After you agree to the terms and conditions, the Cisco Smart Call Home registration for the Cisco UCS domain is complete.



CHAPTER 44

Managing the System Event Log

This chapter includes the following sections:

- [System Event Log, page 623](#)
- [Viewing the System Event Log for a Server, page 624](#)
- [Configuring the SEL Policy, page 625](#)
- [Backing Up the System Event Log for a Server, page 627](#)
- [Clearing the System Event Log for a Server, page 628](#)

System Event Log

The system event log (SEL) resides on the CIMC in NVRAM. It records most server-related events, such as over and under voltage, temperature events, fan events, and events from BIOS. The SEL is mainly used for troubleshooting purposes.

The SEL file is approximately 40KB in size, and no further events can be recorded when it is full. It must be cleared before additional events can be recorded.

You can use the SEL policy to backup the SEL to a remote server, and optionally clear the SEL after a backup operation occurs. Backup operations can be triggered based on specific actions, or they can occur at regular intervals. You can also manually backup or clear the SEL.

The backup file is automatically generated. The filename format is *sel-SystemName-ChassisID-ServerID-ServerSerialNumber-Timestamp*; for example, *sel-UCS-A-ch01-serv01-QCI12522939-20091121160736*.

Viewing the System Event Log for a Server

Viewing the System Event Log for an Individual Server

Procedure

	Command or Action	Purpose
Step 1	UCS-A# <code>show sel chassis-id / blade-id</code>	Displays the system event log for the specified server.

The following example displays the system event log for blade 3 in chassis 1.

```
UCS-A# show sel 1/3
 1 | 01/01/1970 01:23:27 | System Event 0x83 | Timestamp clock synch | SEL timestamp
clock updated, event is f
irst of pair | Asserted
 2 | 01/01/1970 01:23:28 | Drive slot(Bay) SAS0_LINK_STATUS | Transition to Degraded |
Asserted
 3 | 01/01/1970 01:23:28 | Drive slot(Bay) SAS0_LINK_STATUS | Transition to On Line |
Deasserted
 4 | 01/01/1970 01:23:28 | Platform alert LED_SAS0_FAULT | LED is blinking fast |
Asserted
 5 | 01/01/1970 01:23:28 | Platform alert LED_SAS0_FAULT | LED is on | Deasserted
 6 | 01/01/1970 01:23:28 | Platform alert LED_FPID | LED is on | Asserted
 7 | 01/01/1970 01:23:28 | Platform alert LED_FPID | LED is off | Deasserted
 8 | 01/01/1970 01:23:29 | Entity presence MAIN_POWER | Device Absent | Asserted
 9 | 01/01/1970 01:23:29 | Entity presence MAIN_POWER | Device Present | Deasserted
 a | 01/01/1970 01:23:29 | Platform alert LED_SAS0_FAULT | LED is on | Asserted
 b | 01/01/1970 01:23:29 | Platform alert LED_SAS0_FAULT | LED color is green | Asserted
 c | 01/01/1970 01:23:29 | Platform alert LED_SAS0_FAULT | LED is blinking fast |
Deasserted
 d | 01/01/1970 01:23:29 | Platform alert LED_SAS0_FAULT | LED color is amber | Deasserted
 e | 01/01/1970 00:00:22 | Drive slot(Bay) SAS0_LINK_STATUS | Transition to Degraded |
Asserted
 f | 01/01/1970 00:00:22 | Entity presence MEZZ_PRS | Device Present | Asserted
10 | 01/01/1970 00:00:22 | Entity presence HDD1_PRS | Device Absent | Asserted
```

Viewing the System Event Log for All of the Servers in a Chassis

Procedure

	Command or Action	Purpose
Step 1	UCS-A# <code>scope server chassis-id / blade-id</code>	Enters chassis server mode for the specified server.
Step 2	UCS-A /chassis/server # <code>show sel</code>	Displays the system event log.

The following example displays the system event log from chassis server mode for blade 3 in chassis 1.

```
UCS-A# scope server 1/3
UCS-A /chassis/server # show sel
 1 | 01/01/1970 01:23:27 | System Event 0x83 | Timestamp clock synch | SEL timestamp
clock updated, event is f
irst of pair | Asserted
 2 | 01/01/1970 01:23:28 | Drive slot(Bay) SAS0_LINK_STATUS | Transition to Degraded |
Asserted
 3 | 01/01/1970 01:23:28 | Drive slot(Bay) SAS0_LINK_STATUS | Transition to On Line |
Deasserted
 4 | 01/01/1970 01:23:28 | Platform alert LED_SAS0_FAULT | LED is blinking fast |
Asserted
 5 | 01/01/1970 01:23:28 | Platform alert LED_SAS0_FAULT | LED is on | Deasserted
 6 | 01/01/1970 01:23:28 | Platform alert LED_FPID | LED is on | Asserted
 7 | 01/01/1970 01:23:28 | Platform alert LED_FPID | LED is off | Deasserted
 8 | 01/01/1970 01:23:29 | Entity presence MAIN_POWER | Device Absent | Asserted
 9 | 01/01/1970 01:23:29 | Entity presence MAIN_POWER | Device Present | Deasserted
 a | 01/01/1970 01:23:29 | Platform alert LED_SAS0_FAULT | LED is on | Asserted
 b | 01/01/1970 01:23:29 | Platform alert LED_SAS0_FAULT | LED color is green | Asserted
 c | 01/01/1970 01:23:29 | Platform alert LED_SAS0_FAULT | LED is blinking fast |
Deasserted
 d | 01/01/1970 01:23:29 | Platform alert LED_SAS0_FAULT | LED color is amber | Deasserted
 e | 01/01/1970 00:00:22 | Drive slot(Bay) SAS0_LINK_STATUS | Transition to Degraded |
Asserted
 f | 01/01/1970 00:00:22 | Entity presence MEZZ_PRS | Device Present | Asserted
10 | 01/01/1970 00:00:22 | Entity presence HDD1_PRS | Device Absent | Asserted
```

Configuring the SEL Policy

Procedure

	Command or Action	Purpose
Step 1	UCS-A# scope org <i>org-name</i>	Enters organization mode for the specified organization. To enter the root organization mode, type / as the <i>org-name</i> .
Step 2	UCS-A /org # scope ep-log-policy sel	Enters organization endpoint log policy mode and scopes the SEL policy.
Step 3	UCS-A /org/ep-log-policy # set description <i>description</i>	(Optional) Provides a description for the policy. Note If your description includes spaces, special characters, or punctuation, you must begin and end your description with quotation marks. The quotation marks will not appear in the description field of any show command output.
Step 4	UCS-A /org/ep-log-policy # set backup action [log-full] [on-change-of-association] [on-clear] [timer] [none]	Specifies an action or actions that will trigger a backup operation.
Step 5	UCS-A /org/ep-log-policy # set backup clear-on-backup { no yes }	Specifies whether to clear the system event log after a backup operation occurs.

	Command or Action	Purpose
Step 6	UCS-A /org/ep-log-policy # set backup destination <i>URL</i>	<p>Specifies the protocol, user, password, remote hostname, and remote path for the backup operation. Depending on the protocol used, specify the URL using one of the following syntax:</p> <ul style="list-style-type: none"> • ftp:// <i>username@hostname / path</i> • scp:// <i>username @ hostname / path</i> • sftp:// <i>username @ hostname / path</i> • tftp:// <i>hostname : port-num / path</i> <p>Note You can also specify the backup destination by using the set backup hostname , set backup password , set backup protocol , set backup remote-path , set backup user commands, or by using the set backup destination command. Use either method to specify the backup destination.</p>
Step 7	UCS-A /org/ep-log-policy # set backup format { <i>ascii</i> <i>binary</i> }	Specifies the format for the backup file.
Step 8	UCS-A /org/ep-log-policy # set backup hostname { <i>hostname</i> <i>ip-addr</i> }	Specifies the hostname or IP address of the remote server.
Step 9	UCS-A /org/ep-log-policy # set backup interval { 1-hour 2-hours 4-hours 8-hours 24-hours never }	Specifies the time interval for the automatic backup operation. Specifying the never keyword means that automatic backups will not be made.
Step 10	UCS-A /org/ep-log-policy # set backup password <i>password</i>	Specifies the password for the username. This step does not apply if the TFTP protocol is used.
Step 11	UCS-A /org/ep-log-policy # set backup protocol { ftp scp sftp tftp }	Specifies the protocol to use when communicating with the remote server.
Step 12	UCS-A /org/ep-log-policy # set backup remote-path <i>path</i>	Specifies the path on the remote server where the backup file is to be saved.
Step 13	UCS-A /org/ep-log-policy # set backup user <i>username</i>	Specifies the username the system should use to log in to the remote server. This step does not apply if the TFTP protocol is used.
Step 14	UCS-A /org/ep-log-policy # commit-buffer	Commits the transaction.

The following example configures the SEL policy to back up the system event log (in ascii format) every 24 hours or when the log is full and clear the system event log after a backup operation occurs and commits the transaction:

```
UCS-A# scope org /
UCS-A /org # scope ep-log-policy sel
UCS-A /org/ep-log-policy # set backup destination scp://user@192.168.1.10/logs
Password:
UCS-A /org/ep-log-policy* # set backup action log-full
UCS-A /org/ep-log-policy* # set backup clear-on-backup yes
UCS-A /org/ep-log-policy* # set backup format ascii
UCS-A /org/ep-log-policy* # set backup interval 24-hours
UCS-A /org/ep-log-policy* # commit-buffer
UCS-A /org/ep-log-policy #
```

Backing Up the System Event Log for a Server

Backing Up the System Event Log for an Individual Server

Before You Begin

Configure the system event log policy. The manual backup operation uses the remote destination configured in the system event log policy.

Procedure

	Command or Action	Purpose
Step 1	UCS-A /chassis/server # backup sel chassis-id / blade-id	Clears the system event log.
Step 2	UCS-A# commit-buffer	Commits the transaction.

The following example backs up the system event log for blade 3 in chassis 1 and commits the transaction.

```
UCS-A# backup sel 1/3
UCS-A* # commit-buffer
UCS-A#
```

Backing Up the System Event Log for All of the Servers in a Chassis

Before You Begin

Configure the system event log policy. The manual backup operation uses the remote destination configured in the system event log policy.

Procedure

	Command or Action	Purpose
Step 1	UCS-A# scope server <i>chassis-id / blade-id</i>	Enters chassis server mode for the specified server.
Step 2	UCS-A /chassis/server # backup sel	Clears the system event log.
Step 3	UCS-A /chassis/server # commit-buffer	Commits the transaction.

The following example backs up the system event log from chassis server mode for blade 3 in chassis 1 and commits the transaction.

```
UCS-A# scope server 1/3
UCS-A /chassis/server # backup sel
UCS-A /chassis/server* # commit-buffer
UCS-A /chassis/server #
```

Clearing the System Event Log for a Server

Clearing the System Event Log for an Individual Server

Procedure

	Command or Action	Purpose
Step 1	UCS-A# clear sel <i>chassis-id / blade-id</i>	Clears the system event log.
Step 2	UCS-A# commit-buffer	Commits the transaction.

The following example clears the system event log for blade 3 in chassis 1 and commits the transaction:

```
UCS-A# clear sel 1/3
UCS-A* # commit-buffer
UCS-A#
```

Clearing the System Event Log for All of the Servers in a Chassis

Procedure

	Command or Action	Purpose
Step 1	UCS-A# scope server <i>chassis-id / blade-id</i>	Enters chassis server mode for the specified server.
Step 2	UCS-A /chassis/server # clear sel	Clears the system event log.

	Command or Action	Purpose
Step 3	UCS-A /chassis/server # commit-buffer	Commits the transaction.

The following example clears the system event log from chassis server mode for blade 3 in chassis 1 and commits the transaction:

```
UCS-A# scope server 1/3  
UCS-A /chassis/server # clear sel  
UCS-A /chassis/server* # commit-buffer  
UCS-A /chassis/server #
```




CHAPTER 45

Configuring Settings for Faults, Events, and Logs

This chapter includes the following sections:

- [Configuring Settings for the Fault Collection Policy, page 631](#)
- [Configuring Settings for the Core File Exporter, page 632](#)
- [Configuring the Syslog, page 634](#)

Configuring Settings for the Fault Collection Policy

Fault Collection Policy

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

A fault in Cisco UCS has the following lifecycle:

- 1 A condition occurs in the system and Cisco UCS Manager raises a fault. This is the active state.
- 2 When the fault is alleviated, it enters a flapping or soaking interval that is designed to prevent flapping. Flapping occurs when a fault is raised and cleared several times in rapid succession. During the flapping interval, the fault retains its severity for the length of time specified in the fault collection policy.
- 3 If the condition reoccurs during the flapping interval, the fault returns to the active state. If the condition does not reoccur during the flapping interval, the fault is cleared.
- 4 The cleared fault enters the retention interval. This interval ensures that the fault reaches the attention of an administrator even if the condition that caused the fault has been alleviated and the fault has not been deleted prematurely. The retention interval retains the cleared fault for the length of time specified in the fault collection policy.
- 5 If the condition reoccurs during the retention interval, the fault returns to the active state. If the condition does not reoccur, the fault is deleted.

Configuring the Fault Collection Policy

Procedure

	Command or Action	Purpose
Step 1	UCS-A# scope monitoring	Enters monitoring mode.
Step 2	UCS-A /monitoring # scope fault policy	Enters monitoring fault policy mode.
Step 3	UCS-A /monitoring/fault-policy # set clear-action {delete retain}	Specifies whether to retain or delete all cleared messages. If the retain option is specified, then the length of time that the messages are retained is determined by the set retention-interval command.
Step 4	UCS-A /monitoring/fault-policy # set flap-interval seconds	Specifies the time interval (in seconds) the system waits before changing a fault state. Flapping occurs when a fault is raised and cleared several times in rapid succession. To prevent this, the system does not allow a fault to change state until the flapping interval has elapsed after the last state change. If the fault is raised again during the flapping interval, it returns to the active state, otherwise, the fault is cleared.
Step 5	UCS-A /monitoring/fault-policy # set retention-interval {days hours minutes seconds forever}	Specifies the time interval the system retains all cleared fault messages before deleting them. The system can retain cleared fault messages forever, or for the specified number of days, hours, minutes, and seconds.
Step 6	UCS-A /monitoring/fault-policy # commit-buffer	Commits the transaction.

This example configures the fault collection policy to retain cleared fault messages for 30 days, sets the flapping interval to 10 seconds, and commits the transaction.

```
UCS-A# scope monitoring
UCS-A /monitoring # scope fault policy
UCS-A /monitoring/fault-policy # set clear-action retain
UCS-A /monitoring/fault-policy* # set flap-interval 10
UCS-A /monitoring/fault-policy* # set retention-interval 30 0 0 0
UCS-A /monitoring/fault-policy* # commit-buffer
UCS-A /monitoring/fault-policy #
```

Configuring Settings for the Core File Exporter

Core File Exporter

Cisco UCS Manager uses the Core File Exporter to export core files as soon as they occur to a specified location on the network through TFTP. This functionality allows you to export the tar file with the contents of the core file.

Configuring the Core File Exporter

Procedure

	Command or Action	Purpose
Step 1	UCS-A# scope monitoring	Enters monitoring mode.
Step 2	UCS-A /monitoring # scope sysdebug	Enters monitoring system debug mode.
Step 3	UCS-A /monitoring/sysdebug # enable core-export-target	Enables the core file exporter. When the core file exporter is enabled and an error causes the server to perform a core dump, the system exports the core file via TFTP to the specified remote server.
Step 4	UCS-A /monitoring/sysdebug # set core-export-target path path	Specifies the path to use when exporting the core file to the remote server.
Step 5	UCS-A /monitoring/sysdebug # set core-export-target port port-num	Specifies the port number to use when exporting the core file via TFTP. The range of valid values is 1 to 65,535.
Step 6	UCS-A /monitoring/sysdebug # set core-export-target server-description description	Provides a description for the remote server used to store the core file.
Step 7	UCS-A /monitoring/sysdebug # set core-export-target server-name hostname	Specifies the hostname of the remote server to connect with via TFTP.
Step 8	UCS-A /monitoring/sysdebug # commit-buffer	Commits the transaction.

The following example enables the core file exporter, specifies the path and port to use when sending the core file, specifies the remote server hostname, provides a description for the remote server, and commits the transaction.

```
UCS-A# scope monitoring
UCS-A /monitoring # scope sysdebug
UCS-A /monitoring/sysdebug # enable core-export-target
UCS-A /monitoring/sysdebug* # set core-export-target path /root/CoreFiles/core
UCS-A /monitoring/sysdebug* # set core-export-target port 45000
UCS-A /monitoring/sysdebug* # set core-export-target server-description CoreFile102.168.10.10
UCS-A /monitoring/sysdebug* # set core-export-target server-name 192.168.10.10
UCS-A /monitoring/sysdebug* # commit-buffer
UCS-A /monitoring/sysdebug #
```

Disabling the Core File Exporter

Procedure

	Command or Action	Purpose
Step 1	UCS-A# scope monitoring	Enters monitoring mode.
Step 2	UCS-A /monitoring # scope sysdebug	Enters monitoring system debug mode.
Step 3	UCS-A /monitoring/sysdebug # disable core-export-target	Disables the core file exporter. When the core file exporter is disabled core files are not automatically exported.
Step 4	UCS-A /monitoring/sysdebug # commit-buffer	Commits the transaction.

The following example disables the core file exporter and commits the transaction.

```
UCS-A# scope monitoring
UCS-A /monitoring # scope sysdebug
UCS-A /monitoring/sysdebug # disable core-export-target
UCS-A /monitoring/sysdebug* # commit-buffer
UCS-A /monitoring/sysdebug #
```

Configuring the Syslog

Procedure

	Command or Action	Purpose
Step 1	UCS-A# scope monitoring	Enters monitoring mode.
Step 2	UCS-A /monitoring # { enable disable } syslog console	Enables or disables the sending of syslogs to the console.
Step 3	UCS-A /monitoring # set syslog console level { emergencies alerts critical }	(Optional) Select the lowest message level that you want displayed. If syslogs are enabled, the system displays that level and above on the console. The level options are listed in order of decreasing urgency. The default level is Critical.
Step 4	UCS-A /monitoring # { enable disable } syslog monitor	Enables or disables the monitoring of syslog information by the operating system.
Step 5	UCS-A /monitoring # set syslog monitor level { emergencies alerts critical errors warnings notifications information debugging }	(Optional) Select the lowest message level that you want displayed. If the monitor state is enabled, the system displays that level and above. The level options are

	Command or Action	Purpose
		listed in order of decreasing urgency. The default level is Critical. Note Messages at levels below Critical are displayed on the terminal monitor only if you have entered the terminal monitor command.
Step 6	UCS-A /monitoring # {enable disable} syslog file	Enables or disables the writing of syslog information to a syslog file.
Step 7	UCS-A /monitoring # set syslog file name <i>filename</i>	The name of the file in which the messages are logged. Up to 16 characters are allowed in the file name.
Step 8	UCS-A /monitoring # set syslog file level {emergencies alerts critical errors warnings notifications information debugging}	(Optional) Select the lowest message level that you want stored to a file. If the file state is enabled, the system stores that level and above in the syslog file. The level options are listed in order of decreasing urgency. The default level is Critical.
Step 9	UCS-A /monitoring # set syslog file size <i>filesize</i>	(Optional) The maximum file size, in bytes, before the system begins to write over the oldest messages with the newest ones. The range is 4096 to 4194304 bytes.
Step 10	UCS-A /monitoring # {enable disable} syslog remote-destination {server-1 server-2 server-3}	Enables or disables the sending of syslog messages to up to three external syslog servers.
Step 11	UCS-A /monitoring # set syslog remote-destination {server-1 server-2 server-3} level {emergencies alerts critical errors warnings notifications information debugging}	(Optional) Select the lowest message level that you want stored to the external log. If the remote-destination is enabled, the system sends that level and above to the external server. The level options are listed in order of decreasing urgency. The default level is Critical.
Step 12	UCS-A /monitoring # set syslog remote-destination {server-1 server-2 server-3} hostname <i>hostname</i>	The hostname or IP address of the specified remote syslog server. Up to 256 characters are allowed in the hostname.
Step 13	UCS-A /monitoring # set syslog remote-destination {server-1 server-2 server-3} facility {local0 local1 local2 local3 local4 local5 local6 local7}	(Optional) The facility level contained in the syslog messages sent to the specified remote syslog server.
Step 14	UCS-A /monitoring # commit-buffer	Commits the transaction.

This example shows how to enable the storage of syslog messages in a local file and commits the transaction:

```
UCS-A# scope monitoring
UCS-A /monitoring # disable syslog console
```

```
UCS-A /monitoring* # disable syslog monitor
UCS-A /monitoring* # enable syslog file
UCS-A /monitoring* # set syslog file name SysMsgsUCSA
UCS-A /monitoring* # set syslog file level notifications
UCS-A /monitoring* # set syslog file size 4194304
UCS-A /monitoring* # disable syslog remote-destination server-1
UCS-A /monitoring* # disable syslog remote-destination server-2
UCS-A /monitoring* # disable syslog remote-destination server-3
UCS-A /monitoring* # commit-buffer
UCS-A /monitoring #
```



INDEX

A

- accounts [159, 160, 162, 178, 179, 180, 181](#)
 - admin [160](#)
 - expiration [160](#)
 - locally authenticated [160, 178, 179, 180, 181](#)
 - remotely authenticated [160](#)
 - user [159, 162](#)
 - username guidelines [160](#)
- acknowledging [508, 518, 522, 530, 536](#)
 - blade server [518, 522, 536](#)
 - chassis [508](#)
 - rack server [530](#)
- activate firmware [188](#)
- activating [214, 216, 218, 220, 222, 223, 224, 233, 239](#)
 - adapters [214](#)
 - BIOS firmware [216](#)
 - board controller firmware [222](#)
 - capability catalog [233](#)
 - CIMC [218](#)
 - fabric interconnects [224](#)
 - I/O modules [220](#)
 - management extensions [239](#)
 - UCS manager [223](#)
- activities [455, 462, 463](#)
 - pending [455, 462, 463](#)
- adapter [98](#)
 - port channels [98](#)
- adapter port channels [98](#)
 - viewing [98](#)
- adapter qualification [395, 396](#)
 - creating [395](#)
 - deleting [396](#)
- adapters [41, 205, 214, 406](#)
 - NIC [41](#)
 - updating and activating [214](#)
 - vCon placement [406](#)
 - verifying status [205](#)
 - VIC [41](#)
 - virtualization [41](#)
- Address Resolution Protocol [578](#)
 - management interfaces monitoring policy [578](#)
 - configuring [578](#)
- admin account [160](#)
- administration [43](#)
- aging time [250](#)
 - MAC address table [250](#)
- aging time, Mac address table [250](#)
 - configuring [250](#)
- alert groups [610](#)
 - profiles [610](#)
- alert, call home test [613](#)
- all configuration [541](#)
- appliance port channel, NAS port channel [83](#)
 - assigning MAC address [83](#)
- appliance port channels [94](#)
 - member ports [94](#)
 - deleting [94](#)
- appliance port, NAS port [83](#)
 - assigning MAC address [83](#)
- appliance ports [92, 93, 94](#)
 - port channels [92, 93, 94](#)
 - member ports, adding [93](#)
 - member ports, deleting [94](#)
 - unconfiguring [92](#)
- appliance, NAS [81](#)
 - ports [81](#)
 - configuring [81](#)
- architectural simplification [9](#)
- assigning port channels to vlans [305](#)
- assigning ports to vlans [305](#)
- authentication [125](#)
 - primary [125](#)
 - remote [125](#)
- authentication domains [146, 147](#)
 - about [146](#)
- authentication profile [424, 425](#)
- authentication service [149](#)
 - selecting default [149](#)
- authentication services [125](#)
 - about [125](#)
- authNoPriv [117](#)

authPriv [117](#)
 autoconfiguration policy [25, 385](#)
 about [25, 385](#)

B

backing up [202, 541, 542, 544, 545, 547](#)
 about [541](#)
 all-configuration [202](#)
 considerations [542](#)
 creating operations [544](#)
 deleting operations [547](#)
 modifying operations [545](#)
 running operations [545](#)
 types [541](#)
 user role [544](#)
 backup operations [544, 545, 547](#)
 creating [544](#)
 deleting [547](#)
 modifying [545](#)
 running [545](#)
 banner [52, 53, 54](#)
 pre-login [52, 53, 54](#)
 beacon leds [74, 78](#)
 configuring [78](#)
 best effort system class [37, 282](#)
 BIOS [216, 347, 348, 350, 356, 358, 360, 361, 362, 363, 368, 371, 381](#)
 actual settings [371](#)
 default settings [368](#)
 policy [368](#)
 scrub policy [381](#)
 settings [347, 348, 350, 356, 358, 360, 361, 362, 363](#)
 about [347](#)
 boot options [362](#)
 Intel Directed I/O [356](#)
 main [348](#)
 PCI configuration [361](#)
 processor [350](#)
 RAS memory [358](#)
 serial port [360](#)
 server management [363](#)
 USB [360](#)
 updating and activating [216](#)
 BIOS defaults [370](#)
 modifying [370](#)
 BIOS policy [369](#)
 creating [369](#)
 BIOS, recovering [522, 535](#)
 blade [498](#)
 viewing power cap [498](#)
 blade server [514, 515, 516, 518, 519, 520, 521, 522, 536](#)
 acknowledging [518, 522, 536](#)

blade server (*continued*)
 booting [514](#)
 CIMC [521](#)
 resetting [521](#)
 decommissioning [519](#)
 power cycling [516](#)
 removing [519](#)
 resetting [516](#)
 resetting CMOS [520](#)
 shutting down [515](#)
 turning off locator LED [520](#)
 turning on locator LED [520](#)
 blade servers [513, 517, 522, 529](#)
 managing [513](#)
 recovering BIOS [522](#)
 unexpected power changes [517, 529](#)
 blade-level power cap [498](#)
 setting for server [498](#)
 board controllers, activating firmware [222](#)
 boot [415, 448, 449, 450](#)
 LAN [448](#)
 local disk [449](#)
 SAN [415](#)
 virtual media [450](#)
 boot definitions [480, 481, 482, 483, 484](#)
 configuring [480](#)
 deleting [484](#)
 LAN boot [481](#)
 storage boot [482](#)
 virtual media boot [483](#)
 boot options, BIOS settings [362](#)
 boot policies [15, 413, 414, 415, 416, 448, 449, 450, 451](#)
 about [15, 413](#)
 configuring [414](#)
 deleting [451](#)
 LAN boot [448](#)
 local disk boot [449](#)
 SAN boot [415, 416](#)
 virtual media boot [450](#)
 boot process [418](#)
 iSCSI [418](#)
 bootflash, available space [212](#)
 bronze system class [37, 282](#)
 bundles, firmware [186](#)
 burned in values [14, 466](#)

C

Cabling Considerations for Port Channel Mode} [99](#)
 About [99](#)
 call home [606, 609, 611, 612, 613, 614, 615, 616, 618](#)
 configuring [606](#)

- call home (*continued*)
 - inventory messages, configuring [609](#)
 - inventory messages, sending [609](#)
 - policies, configuring [614](#)
 - policies, deleting [616](#)
 - policies, disabling [615](#)
 - policies, enabling [615](#)
 - profiles, configuring [611](#)
 - profiles, deleting [612](#)
 - sending test alert [613](#)
 - smart call home, configuring [616](#)
 - TAC-1 profile, configuring [618](#)
- Call Home [601, 603, 604, 605, 610, 614](#)
 - about [601](#)
 - considerations [603](#)
 - policies [614](#)
 - profiles [610](#)
 - severity levels [604](#)
 - Smart Call Home [605](#)
- canceling image downloads [209](#)
- capability catalog [232, 233, 234, 235, 237](#)
 - about [232](#)
 - activating [233](#)
 - contents [232](#)
 - restarting an update [234](#)
 - updates [233](#)
 - updating [237](#)
 - verifying version [234](#)
 - viewing provider [235](#)
- capping server power usage [498](#)
- catalog [187, 232, 233](#)
 - capability [232, 233](#)
 - firmware images [187](#)
- certificate [106](#)
 - about [106](#)
- changing [132](#)
 - LDAP group rule [132](#)
- chassis [16, 243, 508, 509, 510, 512](#)
 - acknowledging [508](#)
 - decommissioning [508](#)
 - discovery policy [16, 243](#)
 - recommissioning [509](#)
 - removing [509](#)
 - renumbering [510](#)
 - turning off locator LED [512](#)
 - turning on locator LED [512](#)
- chassis connectivity policy [248](#)
 - system-related policies [248](#)
 - chassis [248](#)
- chassis discovery policy [16, 243, 246](#)
 - about [16, 243](#)
 - discovery policies [246](#)
 - chassis [246](#)
- chassis management [509](#)
 - removing [509](#)
- chassis qualification [396, 397](#)
 - configuring [396](#)
 - deleting [397](#)
- CIM XML [104](#)
 - configuring [104](#)
- CIMC [218, 339, 521, 535](#)
 - IP address [339](#)
 - resetting [521, 535](#)
 - blade server [521](#)
 - rack server [535](#)
 - updating and activating [218](#)
- Cisco Discovery Protocol [23, 297](#)
- Cisco UCS Manager [43, 198](#)
 - about [43](#)
 - impact of firmware upgrade [198](#)
- Cisco VM-FEX [41](#)
- cisco-av-pair [126](#)
- CiscoAVPair [126](#)
- CLI session limits [51](#)
- clock [506](#)
 - setting manually [506](#)
- commands for object management [49](#)
- communication services [51, 103, 107, 108, 109, 111, 112, 119, 122, 123, 124](#)
 - about [103](#)
 - disabling [124](#)
 - HTTPS [107, 108, 109, 111, 112](#)
 - SNMP [119, 122, 123](#)
 - web session limits [51](#)
- community, SNMP [119](#)
- component, firmware [186](#)
- configuration [202, 543, 544, 548, 552, 553](#)
 - backing up [202, 544](#)
 - erasing [553](#)
 - import methods [543](#)
 - importing [543, 548](#)
 - restoring [543, 552](#)
- configuring [81, 84, 107, 108, 109, 111, 112, 340, 341, 342, 493](#)
 - blade server [340, 341](#)
 - management IP pool [341](#)
 - static IP address [340](#)
 - Fibre Channel [84](#)
 - storage ports [84](#)
 - global cap policy [493](#)
 - HTTPS [107, 108, 109, 111, 112](#)
 - ports [81](#)
 - appliance, NAS [81](#)
 - rack server [341, 342](#)
 - management IP pool [342](#)
 - static IP address [341](#)
- considerations [542, 603](#)
 - backup operations [542](#)

- considerations (*continued*)
 - Call Home [603](#)
 - console authentication service [148](#)
 - selecting [148](#)
 - converged network adapters [41](#)
 - virtualization [41](#)
 - core file exporter [633, 634](#)
 - configuring [633](#)
 - disabling [634](#)
 - Core File Exporter [632](#)
 - about [632](#)
 - corrupt BIOS [522, 535](#)
 - cpu qualification [398, 399](#)
 - creating [398](#)
 - deleting [399](#)
 - create [49](#)
 - creating [134, 141, 266, 267, 494, 496](#)
 - LDAP group map [134](#)
 - LDAP provider group [141](#)
 - named VLANs [266, 267](#)
 - storage Ethernet mode [267](#)
 - uplink Ethernet mode [266](#)
 - power groups [494](#)
 - power power control policy [496](#)
 - custom roles [164](#)
 - reserved words [164](#)
- D**
- database [541, 543](#)
 - backing up [541](#)
 - restoring [543](#)
 - decommissioning [508, 519, 531](#)
 - blade server [519](#)
 - chassis [508](#)
 - rack server [531](#)
 - decommissioning chassis, guidelines [507](#)
 - decommissioning rack-mount servers [526](#)
 - decommissioning servers [514](#)
 - default service profiles [14, 466](#)
 - default zoning [86, 87](#)
 - about [86](#)
 - disabling [87](#)
 - enabling [86](#)
 - deferring deployment [29, 453, 454, 455, 456, 457, 459, 460, 461, 462, 463](#)
 - guidelines [455](#)
 - maintenance policies [29, 454, 461, 462](#)
 - one time occurrences [456, 459](#)
 - pending activities [455, 462, 463](#)
 - deploying [463](#)
 - viewing [462](#)
 - deferring deployment (*continued*)
 - recurring occurrences [457, 459](#)
 - schedules [454, 456, 460](#)
 - service profiles [453](#)
 - delete [49](#)
 - deleting [135, 142, 144, 146, 495, 497](#)
 - LDAP group map [135](#)
 - LDAP provider group [142](#)
 - power groups [495](#)
 - power power control policy [497](#)
 - RADIUS provider group [144](#)
 - TACACS provider group [146](#)
 - deleting a RADIUS provider [138](#)
 - deleting a TACACS provider [140](#)
 - deleting an LDAP provider [133](#)
 - DHCP [436, 437](#)
 - iSCSI initiator [436, 437](#)
 - deleting DHCP boot parameter [437](#)
 - setting [436](#)
 - disabling [87](#)
 - default zoning [87](#)
 - disaster recovery [541, 543](#)
 - discovery policy [16, 25, 243, 248, 249, 387](#)
 - chassis [16, 243](#)
 - rack server [25, 248, 249](#)
 - server [25, 387](#)
 - disjoint L2 networks [301, 302, 303, 304](#)
 - about [301](#)
 - configuring [304](#)
 - guidelines [302](#)
 - pinning considerations [303](#)
 - disk drive monitoring [579, 580, 582](#)
 - about [579](#)
 - interpreting results [580, 582](#)
 - limitations [580](#)
 - prerequisites [580](#)
 - support [580](#)
 - disk scrub policy [381](#)
 - DNS servers [241, 242](#)
 - about [241](#)
 - configuring [241](#)
 - deleting [242](#)
 - domains, disjoint L2 [301](#)
 - downgrading [201](#)
 - firmware [201](#)
 - prerequisites [201](#)
 - download firmware [188](#)
 - downloading [207, 209](#)
 - canceling [209](#)
 - images [207](#)
 - dynamic vNIC connection policy [19](#)
 - about [19](#)

E

- enabling [86, 116, 119](#)
 - default zoning [86](#)
 - http to https redirection [116](#)
 - SNMP [119](#)
- endpoints [196, 198, 199](#)
 - direct firmware upgrade [196, 198](#)
 - service profile upgrade [199](#)
- enforcing password strength [173](#)
- enter [49](#)
- Ethernet [11, 29, 38, 71, 79, 80, 81, 87, 88, 89, 90, 287, 567](#)
 - Fibre Channel over [11](#)
 - flow control policies [29, 38, 287](#)
 - server ports [79](#)
 - configuring [79](#)
 - unconfiguring [79](#)
 - uplink port channels [87, 88, 89, 90](#)
 - configuring [87](#)
 - member ports, adding [89](#)
 - member ports, deleting [90](#)
 - unconfiguring [88](#)
 - uplink ports [71, 80, 81, 567](#)
 - configuring [80](#)
 - monitoring [567](#)
 - unconfiguring [81](#)
- Ethernet adapter policies [19, 294, 329](#)
 - about [19, 294, 329](#)
- Ethernet adapter policy [295, 297](#)
 - configuring [295](#)
 - deleting [297](#)
- Ethernet switching mode [66, 67](#)
 - about [66](#)
- expiration, accounts [160](#)
- explicit assignment [407, 410](#)
- exporting [541, 544, 633](#)
 - backup types [541](#)
 - configuration [541](#)
 - core file [633](#)
 - user role [544](#)

F

- fabric [99](#)
 - port channels [99](#)
- fabric interconnects [46, 57, 58, 64, 66, 68, 198, 203, 212, 224, 251, 252, 254, 255, 257, 312, 555, 556, 557](#)
 - activating [224](#)
 - admin password recover [556, 557](#)
 - admin password recovery [555](#)
 - available space [212](#)
 - determining leadership role [556](#)
 - enabling standalone for cluster [64](#)

- fabric interconnects (*continued*)
 - Ethernet switching mode [66](#)
 - failover [58](#)
 - FC uplink trunking [312](#)
 - Fibre Channel switching mode [68](#)
 - high availability [46](#)
 - host ID [252](#)
 - impact of firmware upgrade [198](#)
 - initial setup [57, 58](#)
 - about [57](#)
 - management port [58](#)
 - setup mode [58](#)
 - licenses [251, 254, 255, 257](#)
 - installing [254](#)
 - uninstalling [257](#)
 - viewing [254](#)
 - viewing usage [255](#)
 - system configuration type [58](#)
 - verifying high availability status and roles [203](#)
 - verifying operability [203](#)
- fabric port channels [100, 101](#)
 - member ports [101](#)
 - deleting [101](#)
 - viewing [100](#)
- fabric ports [101](#)
 - port channels [101](#)
 - member ports, deleting [101](#)
- fan environment [575](#)
- fan modules [575](#)
- fault collection policy [28, 631, 632](#)
 - about [28, 631](#)
 - configuring [632](#)
- faults [28, 604, 631, 632](#)
 - Call Home severity levels [604](#)
 - collection policy [28, 631](#)
 - Core File Exporter [632](#)
 - lifecycle [28, 631](#)
- FC end-host mode [312](#)
 - VSAN ID restrictions [312](#)
- FC switch mode [312](#)
 - VSAN ID restrictions [312](#)
- FC uplinks [312](#)
 - trunking [312](#)
 - about [312](#)
- FCoE [11](#)
- FCoE ports [85](#)
 - unconfiguring [85](#)
- FCoE VLAN ID [318, 319](#)
 - changing [318, 319](#)
- feature [251](#)
 - licenses [251](#)
- features [38](#)
 - opt-in [38](#)
 - stateless computing [38](#)

Fibre Channel **11, 71, 93, 94, 95, 96, 97, 596, 597**
 link-level flow control **11**
 over Ethernet **11**
 port channels **94, 95, 96, 97**
 configuring **95**
 enabling **96**
 member ports, adding **97**
 unconfiguring **96**
 priority flow control **11**
 statistics threshold policies **596**
 statistics threshold policy classes **597**
 uplink port channels **93**
 enabling **93**
 uplink ports **71**

Fibre Channel adapter policies **19, 294, 329**
 about **19, 294, 329**

Fibre Channel adapter policy **330, 332**
 configuring **330**
 deleting **332**

Fibre Channel ports **97**
 port channels **97**
 member ports, adding **97**

Fibre Channel storage ports **84, 85**
 configuring **84**
 unconfiguring **85**

Fibre Channel switching **86**
 default zoning **86**

Fibre channel switching mode **68**

Fibre Channel switching mode **68**
 about **68**

Fibre Channel system class **37, 282**

firmware **21, 22, 185, 186, 187, 188, 189, 190, 192, 194, 195, 196, 197, 198, 199, 200, 201, 206, 207, 209, 210, 211, 213, 214, 216, 218, 220, 222, 223, 224, 225, 226, 229, 230, 238**
 about **185**
 activating board controller **222**
 adapters **214**
 BIOS **216**
 bundles **186**
 canceling image download **209**
 CIMC **218**
 deleting images **213**
 deleting packages **213**
 direct upgrade **196**
 displaying download status **209**
 downgrades **201**
 downloading packages **207**
 fabric interconnects **224**
 guidelines **189, 190, 192**
 host package **21, 199, 225, 229**
 I/O modules **220**
 image headers **187**
 images **187, 210**
 displaying **210**

firmware (*continued*)
 management **188**
 management extensions **238**
 management package **22, 200, 226, 230**
 obtaining packages **206**
 outage impacts **198**
 package contents **211**
 packages **211**
 displaying **211**
 prerequisites **201**
 service profiles **199**
 UCS manager **223**
 upgrade order **194, 195**
 upgrade stages **197, 200**
 upgrades **189**

firmwareadapters **205**
 verifying status **204, 205**

firmwarecluster configuration **203**
 verifying overall status **203**

firmwarefabric interconnects **203**
 verifying overall status **203**

firmwareI/O modules **204**
 verifying status **204, 205**

firmwareservers **205**
 verifying status **204, 205**

flexibility **10**

flow control **11**
 link-level **11**
 priority **11**

flow control policies **287, 288**
 configuring **287**
 deleting **288**

flow control policy **29, 38, 287**
 about **29, 38, 287**

full state **541**

G

global cap policy **20, 493**
 configuring **493**

gold system class **37, 282**

group map **134, 135**
 creating **134**
 LDAP **134**
 deleting **135**
 LDAP **135**

group rule **132**
 changing **132**
 LDAP **132**

guidelines **34, 36, 160, 162, 189, 190, 192, 302, 312, 376, 377, 455, 466, 564**
 deferred deployment **455**

guidelines (*continued*)

- disjoint L2 networks [302](#)
- firmware upgrades [189, 190, 192](#)
- local disk configuration policy [376, 377](#)
- named VSANs [312](#)
- oversubscription [34](#)
- passwords [162](#)
- pinning [36](#)
- service profiles [466](#)
- traffic monitoring [564](#)
- usernames [160](#)

H

- hardware based service profiles [471](#)
 - creating [471](#)
- hardware-based service profiles [14, 466](#)
- hardware, stateless [38](#)
- headers, images [187](#)
- high availability [10, 46, 58](#)
 - about [46](#)
 - fabric interconnect failover [58](#)
- high availability configuration [46](#)
 - about [46](#)
- history, passwords [179](#)
- host firmware package [21, 199, 225, 229](#)
 - about [21, 199, 225](#)
 - creating and updating [229](#)
- host ID, obtaining [252](#)
- host port channels [98](#)
 - viewing [98](#)
- HTTP [51, 105, 106](#)
 - configuring [105](#)
 - unconfiguring [106](#)
 - web session limits [51](#)
- HTTPS [51, 107, 108, 109, 111, 112, 113, 115](#)
 - certificate request [108, 109](#)
 - configuring [113](#)
 - creating key ring [107](#)
 - importing certificate [112](#)
 - regenerating key ring [107](#)
 - trusted point [111](#)
 - unconfiguring [115](#)
 - web session limits [51](#)

I

- I/O module [539](#)
 - management [539](#)
- I/O modules [204, 220](#)
 - updating and activating [220](#)

I/O modules (*continued*)

- verifying status [204](#)
- IEEE 802.1Qbh [41](#)
- IEEE 802.3x link-level flow control [11](#)
- images [185, 186, 187, 206, 207, 213](#)
 - bundle [186](#)
 - contents [187](#)
 - deleting [213](#)
 - downloading [207](#)
 - headers [187](#)
 - obtaining [206](#)
 - packages, deleting with [213](#)
- implicit assignment [407](#)
- import operations [548, 549, 550, 551](#)
 - creating [548](#)
 - deleting [551](#)
 - modifying [550](#)
 - running [549](#)
- importing [543, 544, 548, 549, 550, 551](#)
 - about [543](#)
 - creating operations [548](#)
 - deleting operations [551](#)
 - modifying operations [550](#)
 - restore methods [543](#)
 - running operations [549](#)
 - user role [544](#)
- informs [117](#)
 - about [117](#)
- inheritance, servers [26, 389](#)
- inherited values [14, 466](#)
- initial setup [57, 58](#)
 - about [57](#)
 - management port IP address [58](#)
 - setup mode [58](#)
- initial templates [15, 467](#)
- Intel Directed I/O, BIOS settings [356](#)
- inventory messages, call home [609](#)
 - configuring [609](#)
 - sending [609](#)
- inventory messages, smart call home [619](#)
 - configuring [619](#)
- IOM [539](#)
 - resetting [539](#)
- IP address pools [425](#)
 - adding iSCSI initiator address block [425](#)
- IP addresses [33, 58, 339, 344](#)
 - CIMC [339](#)
 - management IP pool [33, 344](#)
 - management port [58](#)
- IP pools [33, 344](#)
 - management [33, 344](#)
- IPMI access profile [372, 374, 375](#)
 - configuring [372, 374](#)
 - deleting [374, 375](#)

- IPMI access profiles [21, 372](#)
 - about [21, 372](#)
 - IQN pools [438, 440, 441, 442](#)
 - about [438](#)
 - adding a block [440](#)
 - creating [438](#)
 - deleting [441](#)
 - deleting a block [440](#)
 - viewing usage [442](#)
 - iSCSI adapter policies [422, 424](#)
 - creating [422](#)
 - deleting [424](#)
 - iSCSI auto target [445, 447](#)
 - creating [445](#)
 - deleting [447](#)
 - iSCSI boot [418, 420, 424, 425](#)
 - authentication profile [424, 425](#)
 - configuring [420](#)
 - high-level tasks [420](#)
 - overview [418](#)
 - prerequisites [418](#)
 - iscsi boot policies [429](#)
 - deleting [429](#)
 - iSCSI boot policies [427](#)
 - creating [427](#)
 - iSCSI boot process [418](#)
 - iSCSI initiator [432, 433, 434, 435, 436, 437](#)
 - boot parameters [434, 435](#)
 - iSCSI initiator [434, 435](#)
 - deleting boot from IP pool parameter [435](#)
 - setting to boot from IP pool [434](#)
 - deleting boot from DHCP parameter [437](#)
 - deleting boot from IP Pool parameter [435](#)
 - deleting static IP address boot parameters [433](#)
 - IP pools [434, 435](#)
 - iSCSI initiator [434, 435](#)
 - service profile [432, 433, 434, 435, 436, 437](#)
 - iSCSI initiator [432, 433, 434, 435, 436, 437](#)
 - deleting boot from IP pool parameter [435](#)
 - deleting static IP address boot parameters [433](#)
 - deleting the boot from DHCP parameter [437](#)
 - setting to boot from DHCP [436](#)
 - setting to boot from IP pool [434](#)
 - setting to boot from static IP address [432](#)
 - setting to boot from DHCP [436](#)
 - setting to boot from IP Pool [434](#)
 - setting to boot from static IP address [432](#)
- iSCSI static target [442, 445](#)
 - creating [442](#)
 - deleting [445](#)
- iSCSI vNIC [429, 431](#)
 - creating [429](#)
 - deleting from service profile [431](#)
- ## K
- key ring [106, 107, 108, 109, 111, 112, 114](#)
 - about [106](#)
 - certificate request [108, 109](#)
 - creating [107](#)
 - deleting [114](#)
 - importing certificate [112](#)
 - regenerating [107](#)
 - trusted point [111](#)
 - KVM Console [339](#)
 - IP address [339](#)
- ## L
- L2 networks, disjoint [301](#)
 - LAN [27, 263, 264, 275, 291, 301](#)
 - disjoint L2 networks [301](#)
 - pin groups [275](#)
 - PVLANS [264](#)
 - VLANs [263](#)
 - vNIC policy [27, 291](#)
 - LAN boot [448, 481](#)
 - about [448](#)
 - LAN boot, boot policies [448](#)
 - lanes, virtual [37, 281](#)
 - LDAP [129, 132, 134, 135, 141, 142](#)
 - creating a provider [129](#)
 - group map [134, 135](#)
 - creating [134](#)
 - deleting [135](#)
 - group rule [132](#)
 - changing [132](#)
 - provider group [141, 142](#)
 - creating [141](#)
 - deleting [142](#)
 - LDAP group mapping [133](#)
 - LDAP group rule [128](#)
 - LDAP provider [125, 126, 128](#)
 - about [125](#)
 - configuring properties [128](#)
 - user attribute [126](#)
 - LDAP providers [133](#)
 - deleting [133](#)
 - licenses [251, 252, 253, 254, 255, 257](#)
 - about [251](#)
 - installing [254](#)
 - obtaining [253](#)
 - obtaining host ID [252](#)
 - uninstalling [257](#)
 - viewing [254](#)
 - viewing usage [255](#)
 - lifecycle, faults [28, 631](#)

- link-level flow control [11](#)
 - local disk boot [449](#)
 - about [449](#)
 - configuring [449](#)
 - local disk configuration policy [22, 376, 377](#)
 - about [22, 376](#)
 - guidelines [376](#)
 - RAID configuration [377](#)
 - local disks [379, 380, 478](#)
 - policies [379, 380](#)
 - service profiles [478](#)
 - locales [166, 169, 170, 171, 175, 176](#)
 - about [166](#)
 - assigning an organization [170](#)
 - assigning to user accounts [175](#)
 - creating [169](#)
 - deleting [171](#)
 - deleting an organization from [170](#)
 - removing from user accounts [176](#)
 - locally authenticated users [160, 161, 178, 179, 180, 181](#)
 - accounts [160](#)
 - change interval [180](#)
 - clearing password history [178](#)
 - no change interval [181](#)
 - password history count [181](#)
 - password profile [179](#)
 - reserved words for accounts [161](#)
 - log, system [634](#)
 - log, system event [623, 624, 625, 627, 628](#)
 - about [623](#)
 - backing up [627](#)
 - all servers in chassis [627](#)
 - individual server [627](#)
 - clearing [628](#)
 - all servers in chassis [628](#)
 - individual server [628](#)
 - policy [625](#)
 - viewing [624](#)
 - chassis server mode [624](#)
 - one server [624](#)
 - logical configuration [541](#)
- M**
- MAC address table [250](#)
 - aging time, about [250](#)
 - MAC address table aging time [250](#)
 - configuring [250](#)
 - MAC addresses [31, 277](#)
 - pools [31, 277](#)
 - MAC pools [277, 278](#)
 - creating [277](#)
 - MAC pools *(continued)*
 - deleting [278](#)
 - MAC sync [58](#)
 - main, BIOS settings [348](#)
 - maintenance policies [29, 454, 456, 460, 461, 462](#)
 - about [29, 454](#)
 - creating [461](#)
 - deleting [462](#)
 - schedules [456, 460](#)
 - management [513, 525, 539](#)
 - blade servers [513](#)
 - I/O modules [539](#)
 - rack-mount servers [525](#)
 - management extensions [238, 239](#)
 - about [238](#)
 - activating [239](#)
 - management firmware package [22, 200, 226, 230](#)
 - about [22, 200, 226](#)
 - creating and updating [230](#)
 - management interfaces monitoring policy [23, 577, 578](#)
 - about [23, 577](#)
 - configuring [578](#)
 - management IP address [343](#)
 - service profile [343](#)
 - setting [343](#)
 - management IP addresses [339](#)
 - management IP pool [341, 342](#)
 - configuring [341, 342](#)
 - blade server [341](#)
 - rack server [342](#)
 - management IP pools [33, 344, 345](#)
 - about [33, 344](#)
 - configuring [344](#)
 - deleting [345](#)
 - management port IP address [58, 65](#)
 - changing [65](#)
 - manual blade-level power capping [497](#)
 - Media Independent Interface [578](#)
 - management interfaces monitoring policy [578](#)
 - configuring [578](#)
 - member ports, port channel [89, 90, 97](#)
 - adding [89, 97](#)
 - deleting [90](#)
 - memory qualification [401, 402](#)
 - creating [401](#)
 - deleting [402](#)
 - merging configuration [543](#)
 - mobility [38](#)
 - mode [58, 66, 68](#)
 - end-host [66, 68](#)
 - Ethernet switching [66](#)
 - Fibre Channel switching [68](#)
 - setup [58](#)

monitoring [578, 579, 580, 582](#)
 disk drive [579, 580](#)
 disk drives [580, 582](#)
 interface management [578](#)
 multi-tenancy [39, 40, 153, 154](#)
 about [39](#)
 name resolution [154](#)
 opt-in [40](#)
 opt-out [40](#)
 organizations [153](#)
 multiple authentication systems [141](#)
 mutual inclusion [301](#)

N

name resolution [154, 241](#)
 named VLANs [263, 266, 267, 268, 269, 270](#)
 about [263](#)
 creating for dual fabric interconnects [266, 267](#)
 storage Ethernet mode [267](#)
 uplink Ethernet mode [266](#)
 creating for single fabric interconnect [268, 269](#)
 deleting [270](#)
 named VSANs [311, 312, 314, 315, 316, 317, 318](#)
 about [311](#)
 creating for dual fabric interconnects [314, 315, 317](#)
 creating for single fabric interconnect [316](#)
 deleting [318](#)
 FC uplink trunking [312](#)
 ID range restrictions [312](#)
 named VSANS [318](#)
 FCoE VLAN ID [318](#)
 network [12, 263, 264, 311](#)
 connectivity [12](#)
 named VLANs [263](#)
 named VSANs [311](#)
 private VLANs [264](#)
 network control policies [298, 299](#)
 configuring [298](#)
 deleting [299](#)
 network control policy [23, 297](#)
 networks, disjoint L2 [301](#)
 NIC adapters [41](#)
 virtualization [41](#)
 noAuthNoPriv [117](#)
 NTP servers [503, 505, 506](#)
 about [503](#)
 configuring [505](#)
 deleting [506](#)

O

obtaining [206, 237](#)
 capability catalog updates [237](#)
 firmware image bundles [206](#)
 occurrences [454, 456, 457, 459](#)
 one time [454, 456, 459](#)
 about [454](#)
 creating [456](#)
 deleting [459](#)
 recurring [454, 457, 459](#)
 about [454](#)
 creating [457](#)
 deleting [459](#)
 one time occurrences [454, 456, 459](#)
 about [454](#)
 creating [456](#)
 deleting [459](#)
 opt-in [38, 39, 40](#)
 about [38](#)
 multi-tenancy [40](#)
 stateless computing [39](#)
 opt-out [38, 39, 40](#)
 multi-tenancy [40](#)
 stateless computing [39](#)
 organizations [39, 153, 154, 156, 157, 166](#)
 about [153](#)
 configuring under non-root [156](#)
 configuring under root [156](#)
 deleting [157](#)
 locales [166](#)
 multi-tenancy [39](#)
 name resolution [154](#)
 outage impacts [198](#)
 Cisco UCS Manager [198](#)
 fabric interconnects [198](#)
 firmware upgrade [198](#)
 overriding server identity [13, 465](#)
 oversubscription [33, 34](#)
 about [33](#)
 considerations [33](#)
 guidelines [34](#)
 overview [9](#)

P

packages [206, 207, 211](#)
 contents [211](#)
 downloading [207](#)
 obtaining [206](#)
 packs [21, 22, 199, 200, 225, 226](#)
 host firmware [21, 199, 225](#)
 management firmware [22, 200, 226](#)

- password profile [178, 179, 180, 181](#)
 - about [179](#)
 - change interval [180](#)
 - clearing password history [178](#)
 - no change interval [181](#)
 - password history count [181](#)
- passwords [173, 179](#)
 - change interval [179](#)
 - history count [179](#)
 - strength check [173](#)
- passwords, guidelines [162](#)
- passwords, recovering admin [555, 556, 557](#)
- PCI configuration, BIOS settings [361](#)
- pending activities [455, 462, 463](#)
 - about [455](#)
 - deploying [463](#)
 - viewing [462](#)
- pending commands [50](#)
- PFC [11](#)
- physical qualification [402, 403](#)
 - creating [402](#)
 - deleting [403](#)
- pin groups [35, 275, 321](#)
 - about [35](#)
 - LAN [275](#)
 - SAN [321](#)
- ping gateway [578](#)
 - management interfaces monitoring policy [578](#)
 - configuring [578](#)
- pinning [35, 36, 303](#)
 - about [35](#)
 - disjoint L2 networks [303](#)
 - guidelines [36](#)
 - servers to server ports [35](#)
- PKI [106](#)
- placement policies, vNIC/VHBA [409](#)
 - configuring [409](#)
 - vcons [409](#)
- placement profiles, vNIC/VHBA [410](#)
 - deleting [410](#)
- platinum system class [37, 282](#)
- policies [15, 16, 19, 20, 21, 22, 23, 24, 25, 26, 27, 28, 29, 30, 38, 150, 199, 200, 225, 226, 243, 248, 249, 284, 287, 291, 294, 297, 327, 329, 368, 369, 372, 376, 379, 380, 381, 383, 385, 387, 389, 391, 393, 405, 413, 422, 424, 427, 454, 461, 462, 492, 493, 496, 577, 585, 586, 614, 631, 632](#)
 - about [15](#)
 - autoconfiguration [25, 385](#)
 - BIOS [368, 369](#)
 - boot [15, 413](#)
 - Call Home [614](#)
 - chassis discovery [16, 243](#)
 - dynamic vNIC connection [19](#)
 - about [19](#)
- policies (*continued*)
 - Ethernet [19, 294, 329](#)
 - fault collection [28, 631, 632](#)
 - Fibre Channel adapter [19, 294, 329](#)
 - flow control [29, 38, 287](#)
 - global cap policy [20, 493](#)
 - host firmware [21, 199, 225](#)
 - IPMI access [21, 372](#)
 - iSCSI adapter [422, 424](#)
 - iSCSI boot [427](#)
 - local disk configuration [22, 376](#)
 - local disks [379, 380](#)
 - maintenance [29, 454, 461, 462](#)
 - management firmware [22, 200, 226](#)
 - management interfaces monitoring [23, 577](#)
 - network control [23, 297](#)
 - power [24, 492](#)
 - power control [24, 496](#)
 - PSU [24, 492](#)
 - QoS [25, 38, 284](#)
 - rack server discovery [25, 248, 249](#)
 - role for remote users [150](#)
 - scrub [29, 381](#)
 - serial over LAN [30, 383](#)
 - about [30, 383](#)
 - server discovery [25, 387](#)
 - server inheritance [26, 389](#)
 - about [26, 389](#)
 - server pool [26, 391](#)
 - server pool qualification [26, 393](#)
 - statistics collection [30, 585](#)
 - threshold [30, 586](#)
 - vHBA [27, 327](#)
 - VM lifecycle [27](#)
 - vNIC [27, 291](#)
 - vNIC/vHBA placement [28, 405](#)
- policies, call home [614, 615, 616](#)
 - configuring [614](#)
 - deleting [616](#)
 - disabling [615](#)
 - enabling [615](#)
- policy classes [588, 590, 591, 593, 594, 596, 597, 599](#)
 - Fibre Channel port statistics, configuring [597](#)
 - server port statistics, configuring [594](#)
 - server port statistics, deleting [596](#)
 - server statistics, configuring [588](#)
 - server statistics, deleting [590](#)
 - uplink Ethernet port statistics, configuring [591](#)
 - uplink Ethernet port statistics, deleting [593, 599](#)
- policy-driven chassis group power capping [493](#)
- pools [31, 32, 33, 277, 278, 323, 324, 325, 335, 337, 338, 344, 425, 438, 441, 442](#)
 - about [31](#)
 - adding iSCSI initiator address block [425](#)

- pools (*continued*)
 - IQN [438, 441, 442](#)
 - MAC [31, 277, 278](#)
 - management IP [33, 344](#)
 - server [335](#)
 - servers [31, 335](#)
 - UUID suffix [337, 338](#)
 - UUID suffixes [32, 337](#)
 - WWN [32, 323, 324, 325](#)
 - port channels [87, 88, 89, 90, 92, 93, 94, 95, 96, 97, 98, 99](#)
 - adapter [98](#)
 - appliance [90](#)
 - configuring [87, 90, 95](#)
 - appliance [90](#)
 - enabling [93, 96](#)
 - fabric [99](#)
 - Fibre Channel [94](#)
 - member ports [89, 90, 93, 97](#)
 - adding [89, 93, 97](#)
 - deleting [90](#)
 - unconfiguring [88, 92, 96](#)
 - uplink Ethernet [87](#)
 - port modes [72, 75](#)
 - about [72](#)
 - impact on data traffic [75](#)
 - port types [73](#)
 - ports [35, 58, 71, 72, 74, 79, 80, 81, 84, 87, 88, 90, 92, 93, 94, 95, 96, 98, 99, 251, 275, 321, 567](#)
 - appliance [81](#)
 - appliance port channels [93](#)
 - member ports, adding [93](#)
 - appliance port, NAS port [84](#)
 - unconfiguring [84](#)
 - appliance, NAS [71](#)
 - configuring [81](#)
 - appliance, NAS [81](#)
 - fabric interconnect [71](#)
 - Fibre Channel storage [84](#)
 - configuring [84](#)
 - licenses [251](#)
 - management [58](#)
 - pin groups [275, 321](#)
 - pinning server traffic [35](#)
 - port channels [87, 88, 90, 92, 93, 94, 95, 96, 98, 99](#)
 - configuring [87, 90, 95](#)
 - enabling [93, 96](#)
 - Fibre Channel [94](#)
 - unconfiguring [88, 92, 96](#)
 - server [71, 79](#)
 - configuring [79](#)
 - unconfiguring [79](#)
 - unified [72, 74](#)
 - uplink [71](#)
 - ports (*continued*)
 - uplink Ethernet [80, 81, 567](#)
 - configuring [80](#)
 - monitoring [567](#)
 - unconfiguring [81](#)
 - ports, [89, 90, 94, 97, 101](#)
 - appliance port channels [94](#)
 - member ports, deleting [94](#)
 - fabric port channels [101](#)
 - member ports, deleting [101](#)
 - port channels [89, 90, 97](#)
 - member ports, adding [89, 97](#)
 - member ports, deleting [90](#)
 - power cap [498](#)
 - viewing [498](#)
 - power capping [493, 497](#)
 - manual blade-level [497](#)
 - policy-driven chassis group [493](#)
 - power control policy [24, 496, 497](#)
 - creating [496](#)
 - deleting [497](#)
 - power cycling [516, 528](#)
 - blade server [516](#)
 - rack server [528](#)
 - power group qualification [400](#)
 - configuring [400](#)
 - deleting [400](#)
 - power groups [494, 495](#)
 - creating [494](#)
 - deleting [495](#)
 - power management [24, 491, 494, 495, 496, 497](#)
 - creating [494, 496](#)
 - power control policy [496](#)
 - power groups [494](#)
 - deleting [495, 497](#)
 - power control policy [497](#)
 - power groups [495](#)
 - policies [24, 496](#)
 - power control [24, 496](#)
 - power groups [494](#)
 - rack server [491](#)
 - power policy [24, 492](#)
 - about [24, 492](#)
 - power state, synchronizing [517, 529](#)
 - pre-login banner [52, 53, 54](#)
 - creating [52](#)
 - deleting [54](#)
 - modifying [53](#)
 - primary authentication [125](#)
 - remote [125](#)
 - primary VLAN [271, 272](#)
 - priority flow control [11](#)
 - private VLANs [264, 271, 272, 273](#)
 - about [264](#)

private VLANs (*continued*)
 creating primary 271, 272
 creating secondary 272, 273

privileges 164, 167, 168
 about 164
 adding to user roles 167
 removing from user roles 168

processor, BIOS settings 350

profiles 12, 179, 610
 Call Home alert groups 610
 password 179

profiles, call home 611, 612
 configuring 611
 deleting 612

profiles, TAC-1, smart call home 618
 configuring 618

provider group 141, 142, 143, 144, 145, 146
 creating 141, 143, 145
 LDAP 141
 RADIUS 143
 TACACS 145
 deleting 142, 144, 146
 LDAP 142, 146
 RADIUS 144

provider groups 141, 146
 authentication domains 146

provider, capability catalog 232, 235

PSU policy 24, 492

PVLANs 264, 271, 272, 273
 about 264
 creating primary 271, 272
 creating secondary 272, 273

Q

QoS policies 25, 38, 284, 286
 about 25, 38, 284
 configuring 284
 deleting 286

qualification 395, 398, 401, 402, 403
 creating adapter 395
 creating cpu 398
 creating memory 401
 creating physical 402
 creating storage 403

quality of service 25, 29, 37, 38, 281, 282, 283, 284, 287
 about 37, 281
 flow control policies 29, 38, 287
 policies 25, 38, 284
 system classes 37, 281, 282, 283
 configuring 282
 disabling 283

R

rack server 526, 527, 528, 530, 531, 533, 534, 535
 acknowledging 530
 booting 526
 CIMC 535
 resetting 535
 decommissioning 531
 power cycling 528
 resetting 528
 resetting CMOS 534
 shutting down 527
 turning off locator LED 534
 turning on locator LED 533

rack server discovery policy 25, 248, 249
 about 25, 248
 configuring 249

rack server power management 491

rack servers 535
 recovering BIOS 535

rack-mount servers 25, 195, 248, 249, 466, 525
 discovery policy 25, 248, 249
 guidelines for service profiles 466
 integrating 195
 managing 525

RADIUS 136, 143, 144
 creating provider 136
 provider group 143, 144
 creating 143
 deleting 144

RADIUS provider 125, 126, 135
 about 125
 configuring properties 135
 user attribute 126

RADIUS provider group 143

RADIUS providers 138
 deleting 138

RAID configurations, local disk configuration policies 377

range restrictions, VSAN IDs 312

RAS memory, BIOS settings 358

recommendations 542
 backup operations 542

recommissioning 509
 chassis 509

recovering admin password 555, 556, 557

recovering BIOS 522, 535

recurring occurrences 454, 457, 459
 about 454
 creating 457
 deleting 459

remote authentication 126
 user accounts 126
 user roles 126

- remotely authenticated users [160](#)
 - accounts [160](#)
- removing [509, 519](#)
 - blade server [519](#)
 - chassis [509](#)
- removing chassis, guidelines [507](#)
- removing rack-mount servers [526](#)
- removing servers [514](#)
- renumbering [510, 531](#)
 - blade servers [510](#)
 - chassis [510](#)
 - rack-mount servers [531](#)
- replacing configuration [543](#)
- reserved words [161, 164](#)
 - custom roles [164](#)
 - locally authenticated user accounts [161](#)
- resetting [516, 528](#)
 - blade server [516](#)
 - rack server [528](#)
- resetting CMOS [520, 534](#)
 - blade server [520](#)
 - rack server [534](#)
- resolution, name [241](#)
- restoring [543, 544, 552](#)
 - about [543](#)
 - configuration [552](#)
 - user role [544](#)
- Restoring [85](#)
 - ports [85](#)
 - Uplink Fibre Channel port [85](#)
- role policy for remote users [150, 151](#)
 - about [150](#)
- role-based access control [159, 177](#)
 - user account [177](#)
 - enabling, disabling [177](#)
- roles [162, 163, 164, 174, 176, 544](#)
 - about [162](#)
 - assigning to user accounts [174](#)
 - backing up [544](#)
 - default [163](#)
 - privileges [164](#)
 - removing from user accounts [176](#)
- RSA [106](#)

S

- SAN [27, 311, 321, 327](#)
 - pin groups [321](#)
 - vHBA policy [27, 327](#)
 - VSANs [311](#)
- SAN boot [415](#)
 - about [415](#)
- SAN boot, boot policies [416](#)
- scalability [10](#)
- schedules [454, 456, 457, 459, 460](#)
 - about [454](#)
 - creating [456](#)
 - deleting [460](#)
 - one time occurrences [456, 459](#)
 - creating [456](#)
 - deleting [459](#)
 - recurring occurrences [457, 459](#)
 - creating [457](#)
 - deleting [459](#)
- scope [49](#)
- scrub policies [381, 382](#)
 - creating [381](#)
 - deleting [382](#)
- scrub policy [29, 381](#)
 - about [29, 381](#)
- secondary VLAN [272, 273](#)
- SEL [623, 624, 625, 627, 628](#)
 - about [623](#)
 - backing up [627](#)
 - all servers in chassis [627](#)
 - individual server [627](#)
 - clearing [628](#)
 - all servers in chassis [628](#)
 - individual server [628](#)
 - policy [625](#)
 - viewing [624](#)
 - chassis server mode [624](#)
 - one server [624](#)
- selective exclusion [301](#)
- serial number, obtaining [252](#)
- serial over LAN [383, 479](#)
 - policies [383](#)
 - service profiles [479](#)
- serial over LAN policies [385](#)
 - deleting [385](#)
- serial over LAN policy [30, 383, 384](#)
 - about [30, 383](#)
 - viewing [384](#)
- serial port, BIOS settings [360](#)
- server [498](#)
 - setting power blade-level power cap [498](#)
- server autoconfiguration policies [385, 387](#)
 - configuring [385](#)
 - deleting [387](#)
- server autoconfiguration policy [25, 385](#)
 - about [25, 385](#)
- server discovery policies [387, 389](#)
 - configuring [387](#)
 - deleting [389](#)
 - discovery policies [387, 389](#)
 - server, configuring [387](#)

- server discovery policies (*continued*)
 - discovery policies (*continued*)
 - server, deleting [389](#)
- server discovery policy [25, 387](#)
 - about [25, 387](#)
- server inheritance policies [389, 391](#)
 - configuring [389](#)
 - deleting [391](#)
- server inheritance policy [26, 389](#)
 - about [26, 389](#)
- server management [513, 525](#)
- server management, BIOS settings [363](#)
- server pool policies [391, 392](#)
 - configuring [391](#)
 - deleting [392](#)
- server pool policy [26, 391](#)
 - about [26, 391](#)
- server pool policy qualification [26, 393, 394](#)
 - about [26, 393](#)
 - creating [393](#)
 - deleting [394](#)
- server pools [335, 336](#)
 - creating [335](#)
 - deleting [336](#)
- server ports [71, 79](#)
 - about [71](#)
 - configuring [79](#)
 - unconfiguring [79](#)
- server virtualization [10](#)
- servers [12, 13, 15, 21, 22, 25, 26, 31, 35, 38, 39, 195, 205, 214, 216, 218, 241, 335, 347, 348, 350, 356, 358, 360, 361, 362, 363, 368, 369, 370, 371, 372, 376, 387, 389, 391, 393, 413, 422, 424, 427, 453, 465, 487, 513, 517, 525, 529, 531, 579, 580, 582](#)
 - actual BIOS settings [371](#)
 - adapters [214](#)
 - updating and activating [214](#)
 - adding previously unsupported [195](#)
 - BIOS [216](#)
 - updating and activating [216](#)
 - BIOS defaults [368, 370](#)
 - BIOS policies [368](#)
 - BIOS policy [369](#)
 - BIOS settings [347, 348, 350, 356, 358, 360, 361, 362, 363](#)
 - blade [513](#)
 - boot policies [15, 413](#)
 - CIMC [218](#)
 - updating and activating [218](#)
 - configuration [12](#)
 - discovery policy [25, 387](#)
 - disk drive monitoring [579, 580](#)
 - disk drive status [580, 582](#)
 - disk drive support [580](#)
 - DNS [241](#)
 - inheritance policy [26, 389](#)
- servers (*continued*)
 - IPMI access [21, 372](#)
 - iSCSI adapter policies [422, 424](#)
 - iSCSI boot policies [427](#)
 - local disk configuration [22, 376](#)
 - multi-tenancy [39](#)
 - pinning [35](#)
 - pool policy [26, 391](#)
 - pool qualifications [26, 393](#)
 - pools [31, 335](#)
 - rack-mount [525](#)
 - renumbering [531](#)
 - resetting UUID [487](#)
 - service profiles [12, 13, 453, 465](#)
 - stateless [38](#)
 - unexpected power changes [517, 529](#)
 - verifying status [205](#)
- service profile [429, 431](#)
 - iSCSI vNIC [429, 431](#)
 - creating [429](#)
 - deleting [431](#)
- service profiles [12, 13, 14, 15, 199, 453, 465, 466, 467, 470, 471, 474, 476, 478, 479, 480, 481, 482, 483, 484, 485, 486, 487, 488, 489, 517, 529](#)
 - about [12](#)
 - associating [485, 486](#)
 - blade server [485](#)
 - rack server [486](#)
 - server pool [485](#)
 - boot definitions [480, 484](#)
 - configuration [12](#)
 - deferring deployment [453](#)
 - disassociating [486](#)
 - blade server [486](#)
 - server pool [486](#)
 - firmware upgrades [199](#)
 - guidelines [466](#)
 - hardware based [471](#)
 - creating [471](#)
 - inherited values [14, 466](#)
 - instance, creating from template [470](#)
 - LAN boot [481](#)
 - local disks [478](#)
 - network connectivity [12](#)
 - override identity [13, 465](#)
 - resetting MAC address [488](#)
 - resetting UUID [487](#)
 - resetting WWPN [489](#)
 - serial over LAN [479](#)
 - storage boot [482](#)
 - template, creating [467](#)
 - templates [15, 467](#)
 - unexpected power changes [517, 529](#)
 - vHBAs [476](#)
 - virtual media boot [483](#)

- service profiles (*continued*)
 - vNICs [474](#)
- setting [343](#)
 - management IP address [343](#)
 - service profile [343](#)
- setup mode [58](#)
- severity levels, Call Home [604](#)
- shutting down [515, 527](#)
 - blade server [515](#)
 - rack server [527](#)
- silver system class [37, 282](#)
- smart call home [616, 618, 619, 620](#)
 - configuring [616](#)
 - inventory messages, configuring [619](#)
 - registering [620](#)
 - TAC-1 profile, configuring [618](#)
- Smart Call Home [603, 604, 605](#)
 - about [605](#)
 - considerations [603](#)
 - severity levels [604](#)
- SNMP [116, 117, 119, 120, 121, 122, 123](#)
 - about [116](#)
 - community [119](#)
 - enabling [119](#)
 - notifications [117](#)
 - privileges [117](#)
 - security levels [117](#)
 - SNMPv3 users [123](#)
 - support [116, 119](#)
 - trap [121](#)
 - deleting [121](#)
 - traps [120](#)
 - creating [120](#)
 - users [122, 123](#)
 - creating [122](#)
 - deleting [123](#)
 - Version 3 security features [119](#)
- SNMPv3 [119](#)
 - security features [119](#)
- software [185](#)
- SOL policies [384, 385](#)
 - deleting [385](#)
 - viewing [384](#)
- SPAN, See [traffic monitoring](#)
- stages, firmware upgrades [197, 200](#)
- stateless computing [38, 39](#)
 - about [38](#)
 - opt-in [39](#)
 - opt-out [39](#)
- statelessness [38](#)
- static address [432, 433](#)
 - iSCSI initiator [432, 433](#)
 - deleting [433](#)
 - setting [432](#)
 - static IP address [340, 341](#)
 - blade server [340](#)
 - configuring [340](#)
 - rack server [341](#)
 - configuring [341](#)
- statistics [30, 586](#)
 - threshold policies [30, 586](#)
- statistics collection policies [30, 585](#)
 - about [30, 585](#)
- statistics collection policy [586](#)
 - configuring [586](#)
- statistics threshold policies [587, 588, 590, 591, 593, 594, 596, 597, 599](#)
 - Fibre channel port, classes, configuring [597](#)
 - Fibre Channel port, configuring [596](#)
 - server classes, deleting [590](#)
 - server port classes, configuring [594](#)
 - server port classes, deleting [596](#)
 - server port, configuring [593](#)
 - server, classes, configuring [588](#)
 - server, configuring [587](#)
 - server, deleting [588](#)
 - uplink Ethernet port classes, deleting [593, 599](#)
 - uplink Ethernet port, classes, configuring [591](#)
 - uplink Ethernet port, configuring [590](#)
- statistics threshold policy classes [588, 590, 591, 593, 594, 596, 597, 599](#)
 - Fibre Channel port, configuring [597](#)
 - server port, configuring [594](#)
 - server port, deleting [596](#)
 - server, configuring [588](#)
 - server, deleting [590](#)
 - uplink Ethernet port, configuring [591](#)
 - uplink Ethernet port, deleting [593, 599](#)
- storage boot [482](#)
- storage qualification [403, 405](#)
 - creating [403](#)
 - deleting [405](#)
- storage VSANS [319](#)
 - FCoE VLAN ID [319](#)
- supported tasks [44](#)
- switching mode [67, 68](#)
- syslog [634](#)
- system class [282, 283](#)
 - configuring [282](#)
 - disabling [283](#)
- system classes [37, 281, 282](#)
 - best effort [37, 282](#)
 - bronze [37, 282](#)
 - Fibre Channel [37, 282](#)
 - gold [37, 282](#)
 - platinum [37, 282](#)
 - silver [37, 282](#)
- system configuration [541](#)

- system event log [623, 624, 625, 627, 628](#)
 - about [623](#)
 - backing up [627](#)
 - all servers in chassis [627](#)
 - individual server [627](#)
 - clearing [628](#)
 - all servers in chassis [628](#)
 - individual server [628](#)
 - policy [625](#)
 - viewing [624](#)
 - chassis server mode [624](#)
 - one server [624](#)
- system management [513, 525, 539](#)
 - blade servers [513](#)
 - I/O module [539](#)
 - rack-mount servers [525](#)
- system name [65](#)
 - changing [65](#)

T

- TACACS [145, 146](#)
 - provider group [145, 146](#)
 - creating [145](#)
 - deleting [146](#)
- TACACS provider group [145](#)
- TACACS providers [140](#)
 - deleting [140](#)
- TACACS+ [139](#)
 - creating provider [139](#)
- TACACS+ provider [125, 126, 138](#)
 - about [125](#)
 - configuring properties [138](#)
 - user attribute [126](#)
- tasks [44, 46](#)
 - supported [44](#)
 - unsupported [46](#)
- telnet [123](#)
 - enabling [123](#)
- templates [15, 467](#)
 - service profiles [15, 467](#)
- test alert, call home [613](#)
- TFTP Core Exporter [632](#)
- threshold policies [30, 586](#)
 - about [30, 586](#)
- time zones [503, 505, 506](#)
 - about [503](#)
 - configuring NTP servers [505](#)
 - deleting NTP servers [506](#)
 - setting [503](#)
 - setting clock manually [506](#)
 - viewing [503](#)

- traffic management [33, 34, 37, 281](#)
 - oversubscription [33, 34](#)
 - quality of service [37, 281](#)
 - system classes [37, 281](#)
 - virtual lanes [37, 281](#)
- traffic monitoring [563, 564, 565, 566, 567, 568, 570, 571, 572, 573](#)
 - about [563](#)
 - activating a session [572](#)
 - adding a storage port source [571](#)
 - adding a VLAN or VSAN source [570](#)
 - adding a vNIC or vHBA source [568](#)
 - adding an uplink source port [567](#)
 - deleting a session [573](#)
 - Ethernet session [565](#)
 - Fibre Channel session [566](#)
 - guidelines [564](#)
 - storage ports [571](#)
 - monitoring [571](#)
- trap [121](#)
 - deleting [121](#)
- traps [117, 120](#)
 - about [117](#)
 - creating [120](#)
- tray [575](#)
- trunking [312, 319](#)
 - Fibre Channel [312, 319](#)
 - uplink [312, 319](#)
- trunking, named VSANs [312](#)
- trusted points [106, 111, 115](#)
 - about [106](#)
 - creating [111](#)
 - deleting [115](#)
- turning off locator LED [512, 520, 534](#)
 - blade server [520](#)
 - chassis [512](#)
 - rack server [534](#)
- turning on locator LED [512, 520, 533](#)
 - blade server [520](#)
 - chassis [512](#)
 - rack server [533](#)

U

- UCS manager [223](#)
 - activating [223](#)
- unconfiguring [84](#)
 - ports [84](#)
 - appliance, NAS port [84](#)
- unexpected power changes, avoiding [517, 529](#)
- unified fabric [10, 11](#)
 - about [10](#)
 - Fibre Channel [11](#)

- unified ports [72, 73, 74, 75](#)
 - guidelines [74](#)
 - port modes [72, 75](#)
 - port types [73](#)
 - ports [74](#)
 - unified [74](#)
 - guidelines [74](#)
- unsupported tasks [46](#)
- updating [194, 195, 214, 216, 218, 220, 237, 453](#)
 - adapters [214](#)
 - BIOS firmware [216](#)
 - capability catalog [237](#)
 - CIMC [218](#)
 - firmware order [194, 195](#)
 - I/O modules [220](#)
 - service profiles [453](#)
- updating templates [15, 467](#)
- upgrading [189, 190, 192, 196, 197, 199, 200, 201, 232, 233](#)
 - capability catalog [232, 233](#)
 - firmware [189, 197, 200](#)
 - firmware, direct [196](#)
 - firmware, guidelines [189, 190, 192](#)
 - firmware, service profiles [199](#)
 - prerequisites [201](#)
- upgrading firmware [206, 207, 211](#)
 - downloading images [207, 211](#)
 - obtaining packages [206](#)
- Uplink Fibre Channel port [85](#)
 - restoring [85](#)
- uplink ports [29, 38, 71, 80, 81, 87, 88, 89, 90, 93, 275, 287, 321, 567](#)
 - about [71](#)
 - Ethernet [80, 81, 567](#)
 - configuring [80](#)
 - monitoring [567](#)
 - unconfiguring [81](#)
 - flow control policies [29, 38, 287](#)
 - pin groups [275, 321](#)
 - port channels [87, 88, 89, 90, 93](#)
 - configuring [87](#)
 - enabling [93](#)
 - member ports, adding [89](#)
 - member ports, deleting [90](#)
 - unconfiguring [88](#)
 - uplink Ethernet [87](#)
- uplink trunking [312, 319](#)
 - Fibre Channel [312, 319](#)
 - about [312](#)
 - enabling, disabling [319](#)
- upstream disjoint L2 networks, See [disjoint L2 networks](#)
- usage, licenses [255](#)
- USB, BIOS settings [360](#)
- user account [177](#)
 - enabling, disabling [177](#)
- user accounts [159, 160, 162, 171, 174, 175, 176, 178, 179, 180, 181, 182](#)
 - about [159, 162](#)
 - creating [171](#)
 - deleting [178](#)
 - locales [175, 176](#)
 - assigning [175](#)
 - removing [176](#)
 - monitoring [182](#)
 - password profile [178, 179, 180, 181](#)
 - roles [174, 176](#)
 - assigning [174](#)
 - removing [176](#)
 - username guidelines [160](#)
 - web session limits [162, 174](#)
- user attributes [126](#)
 - LDAP [126](#)
 - RADIUS [126](#)
 - TACACS+ [126](#)
- user roles [162, 163, 164, 167, 169](#)
 - about [162](#)
 - creating [167](#)
 - default [163](#)
 - deleting [169](#)
 - privileges [164](#)
- usernames, guidelines [160](#)
- users [51, 122, 123, 125, 126, 150, 159, 160, 162, 163, 164, 166, 173, 178, 179, 180, 181, 555, 556, 557](#)
 - access control [159](#)
 - accounts [159, 162](#)
 - authentication [125](#)
 - CLI session limits [51](#)
 - default roles [163](#)
 - guidelines [160](#)
 - locales [166](#)
 - about [166](#)
 - locally authenticated [160, 178, 179, 180, 181](#)
 - password strength check [173](#)
 - privileges [164](#)
 - recovering admin password [555, 556, 557](#)
 - remote authentication [126](#)
 - remote, role policy [150](#)
 - remotely authenticated [160](#)
 - roles [162](#)
 - SNMPv3 [122, 123](#)
 - web session limits [51, 162](#)
- UUID [487](#)
 - resetting [487](#)
- UUID suffix pools [32, 337, 338](#)
 - about [32, 337](#)
 - creating [337](#)
 - deleting [338](#)

V

- vcons [409](#)
 - vNIC/vHBA placement policies [409](#)
- vCons [28, 405, 406, 407, 410](#)
 - about [28, 405](#)
 - adapter placement [406](#)
 - vNIC/vHBA assignment [407, 410](#)
- vHBA SAN Connectivity policies [27, 327](#)
 - about [27, 327](#)
- vHBA templates [27, 327, 329](#)
 - about [27, 327](#)
 - configuring [327](#)
 - deleting [329](#)
- vHBAs [407, 476, 489, 568](#)
 - assignment to vCon [407](#)
 - monitoring [568](#)
 - resetting WWPN [489](#)
 - service profiles [476](#)
- VIC adapters [41](#)
 - virtualization [41](#)
- viewing [384, 498](#)
 - blade-level power cap [498](#)
 - serial over LAN policies [384](#)
- virtual lanes [37, 281](#)
- virtual machines [40](#)
- virtual media boot [450, 483](#)
 - about [450](#)
- virtual media boot, boot policies [450](#)
- virtualization [27, 40, 41](#)
 - about [40](#)
 - converged network adapters [41](#)
 - NIC adapters [41](#)
 - VIC adapter [41](#)
 - VM lifecycle policy [27](#)
 - VM-FEX [41](#)
 - about [41](#)
- VLAN port count [274](#)
- VLAN port limitations [265](#)
- VLANs [263, 264, 271, 272, 273, 301, 302, 303, 304, 570](#)
 - disjoint L2 networks [301, 302, 303, 304](#)
 - monitoring [570](#)
 - named [263](#)
 - about [263](#)
 - private [264, 271, 272, 273](#)
 - about [264](#)
 - creating primary [271, 272](#)
 - creating secondary [272, 273](#)

- VM lifecycle policy [27](#)
 - about [27](#)
- VM-FEX [41](#)
 - about [41](#)
- vNIC [27, 291](#)
 - policy [27, 291](#)
- vNIC LAN Connectivity policies [27, 291](#)
 - about [27, 291](#)
- vNIC templates [27, 291, 292, 294](#)
 - about [27, 291](#)
 - configuring [292](#)
 - deleting [294](#)
- vNIC/vHBA placement policies [28, 405, 406, 407, 409, 410](#)
 - about [28, 405](#)
 - configuring [409](#)
 - deleting [410](#)
 - vcons [409](#)
 - vCons [28, 405, 406](#)
 - vNIC/vHBA assignment [407](#)
- vNICs [19, 407, 410, 429, 431, 474, 488, 568](#)
 - assignment to vCon [407, 410](#)
 - dynamic vNIC connection policy [19](#)
 - iSCSI [429, 431](#)
 - creating [429](#)
 - deleting for service profile [431](#)
 - monitoring [568](#)
 - resetting MAC address [488](#)
 - service profiles [474](#)
- VSANs [311, 312, 570](#)
 - monitoring [570](#)
 - named [311, 312](#)

W

- web session limits [51, 162, 174](#)
 - user accounts [174](#)
- WWN pools [32, 323, 324, 325](#)
 - about [32, 323](#)
 - creating [324](#)
 - deleting [325](#)
- WWNN pools [32, 323](#)
 - about [32, 323](#)
- WWPN pools [32, 323](#)
 - about [32, 323](#)

