



Cisco Nexus 5000 Series NX-OS Software Configuration Guide

First Published: October 09, 2009

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-16597-01



CONTENTS

[Trademark Notices](#) 41

Preface

[Preface](#) **xliii**

[Audience](#) **xliiii**

[Document Organization](#) **xliiii**

[Document Conventions](#) **xliv**

[Related Documentation](#) **xliv**

[Obtaining Documentation and Submitting a Service Request](#) **xlvi**

CHAPTER 1

[New and Changed Information](#) 1

[New and Changed Information for the Cisco Nexus 5000 Series](#) 1

CHAPTER 2

[Overview](#) 3

[Information About Cisco Nexus 5000 Series Switches](#) 3

[New Technologies in the Cisco Nexus 5000 Series](#) 3

[Fibre Channel over Ethernet](#) 3

[Data Center I/O Consolidation](#) 4

[Virtual Interfaces](#) 5

[Cisco Nexus 5000 Series Switch Software](#) 5

[Ethernet Switching](#) 5

[FCoE and Fibre Channel Switching](#) 5

[QoS](#) 6

[Virtual Port Channels](#) 6

[Serviceability](#) 6

[Switched Port Analyzer](#) 6

[Ethanalyzer](#) 6

[Call Home](#) 6

[Online Diagnostics](#) 7

Switch Management	7
Simple Network Management Protocol	7
Role-Based Access Control	7
Configuration Methods	7
Configuring with CLI, XML Management Interface, or SNMP	7
Configuring with Cisco Data Center Network Manager	7
Configuring with Cisco MDS Fabric Manager	7
Network Security Features	8
Virtual Device Contexts	8
Licensing	8
Typical Deployment Topologies	8
Ethernet TOR Switch Topology	8
Fabric Extender Deployment Topology	10
Data Center I/O Consolidation Topology	11
Supported Standards	11

PART I**Configuration Fundamentals 13****CHAPTER 3****Using the Command-Line Interface 15**

Accessing the Command Line Interface	15
Telnet Connection	15
SSH Connection	16
Using the CLI	17
Using CLI Command Modes	17
Changing Command Modes	18
Listing the Commands Used with Each Command Mode	18
CLI Command Hierarchy	18
EXEC Mode Commands	19
Configuration Mode Commands	21
Using Commands	23
Listing Commands and Syntax	23
Entering Command Sequences	23
Undoing or Reverting to Default Values or Conditions	23
Using Keyboard Shortcuts	24
Using CLI Variables	25

User-Defined Persistent CLI Variables	26
Using Command Aliases	27
Defining Command Aliases	27
Command Scripts	28
Executing Commands Specified in a Script	28
Using CLI Variables in Scripts	29
Setting the Delay Time	29

CHAPTER 4**Initial Switch Configuration 31**

Configuring the Switch	31
Image Files on the Switch	31
Starting the Switch	31
Boot Sequence	32
Console Settings	33
Upgrading the Switch Software	33
Downgrading from a Higher Release	36
Initial Configuration	37
Configuration Prerequisites	37
Initial Setup	38
Preparing to Configure the Switch	38
Default Login	39
Configuring the Switch	39
Changing the Initial Configuration	44
Management Interface Configuration	44
About the mgmt0 Interface	44
Configuring the Management Interface	44
Displaying Management Interface Configuration	46
Shutting Down the Management Interface	46

CHAPTER 5**Managing Licenses 47**

Licensing Terminology	47
Licensing Model	48
Licence Installation	49
Obtaining a Factory-Installed License	49
Performing a Manual Installation	50

Obtaining the License Key File	50
Installing the License Key File	51
Backing Up License Files	52
Identifying License Features in Use	53
Uninstalling Licenses	53
Updating Licenses	54
Grace Period Alerts	56
License Transfers Between Switches	57
Verifying the License Configuration	57

PART II
LAN Switching 59

CHAPTER 6
Configuring Ethernet Interfaces 61

Information About Ethernet Interfaces	61
About the Interface Command	61
About the Unidirectional Link Detection Parameter	62
UDLD Overview	62
Default UDLD Configuration	63
UDLD Aggressive and Nonaggressive Modes	63
About Interface Speed	64
About the Cisco Discovery Protocol	64
Default CDP Configuration	64
About the Debounce Timer Parameters	65
About MTU Configuration	65
Configuring Ethernet Interfaces	65
Configuring the UDLD Mode	65
Configuring Interface Speed	67
Configuring the Cisco Discovery Protocol	68
Configuring the CDP Characteristics	68
Enabling or Disabling CDP	69
Configuring the Debounce Timer	69
Configuring the Description Parameter	70
Disabling and Restarting Ethernet Interfaces	71
Displaying Interface Information	72
Default Physical Ethernet Settings	74

CHAPTER 7**Configuring VLANs 77**

Configuring VLANs 77

Information About VLANs 77

Understanding VLANs 77

Understanding VLAN Ranges 78

Creating, Deleting, and Modifying VLANs 79

Configuring a VLAN 80

Creating and Deleting a VLAN 80

Entering the VLAN Submode and Configuring the VLAN 81

Adding Ports to a VLAN 82

Verifying VLAN Configuration 83

CHAPTER 8**Configuring Private VLANs 85**

Information About Private VLANs 85

Primary and Secondary VLANs in Private VLANs 86

Private VLAN Ports 86

Primary, Isolated, and Community Private VLANs 87

Associating Primary and Secondary VLANs 88

Private VLAN Promiscuous Trunks 89

Private VLAN Isolated Trunks 89

Broadcast Traffic in Private VLANs 90

Private VLAN Port Isolation 90

Guidelines and Limitations for Private VLANs 90

Configuring a Private VLAN 90

Enabling Private VLANs 90

Configuring a VLAN as a Private VLAN 91

Associating Secondary VLANs with a Primary Private VLAN 92

Configuring an Interface as a Private VLAN Host Port 93

Configuring an Interface as a Private VLAN Promiscuous Port 94

Configuring a Promiscuous Trunk Port 95

Configuring an Isolated Trunk Port 97

Configuring the Allowed VLANs for PVLAN Trunking Ports 98

Configuring Native 802.1Q VLANs on Private VLANs 99

Verifying Private VLAN Configuration 99

CHAPTER 9**Configuring Access and Trunk Interfaces 101**

- Information About Access and Trunk Interfaces 101
 - Understanding Access and Trunk Interfaces 101
 - Understanding IEEE 802.1Q Encapsulation 103
 - Understanding Access VLANs 103
 - Understanding the Native VLAN ID for Trunk Ports 104
 - Understanding Allowed VLANs 104
 - Understanding Native 802.1Q VLANs 104
- Configuring Access and Trunk Interfaces 105
 - Configuring a LAN Interface as an Ethernet Access Port 105
 - Configuring Access Host Ports 106
 - Configuring Trunk Ports 107
 - Configuring the Native VLAN for 802.1Q Trunking Ports 107
 - Configuring the Allowed VLANs for Trunking Ports 108
 - Configuring Native 802.1Q VLANs 109
- Verifying Interface Configuration 110

CHAPTER 10**Configuring EtherChannels 111**

- Information About EtherChannels 111
 - Understanding EtherChannels 111
 - Compatibility Requirements 112
 - Load Balancing Using EtherChannels 113
 - Understanding LACP 115
 - LACP Overview 115
 - LACP ID Parameters 115
 - Channel Modes 116
 - LACP Marker Responders 117
 - LACP-Enabled and Static EtherChannels Differences 117
- Configuring EtherChannels 118
 - Creating an EtherChannel 118
 - Adding a Port to an EtherChannel 118
 - Configuring Load Balancing Using EtherChannels 120
 - Enabling LACP 121
 - Configuring Channel Mode for a Port 121

Configuring the LACP System Priority and System ID	122
Configuring the LACP Port Priority	123
Verifying EtherChannel Configuration	124

CHAPTER 11

Configuring Virtual Port Channels	125
Information About vPCs	125
vPC Overview	125
Terminology	127
vPC Terminology	127
Fabric Extender Terminology	127
Supported vPC Topologies	128
Cisco Nexus 5000 Series Switch vPC Topology	128
Single Homed Fabric Extender vPC Topology	129
Dual Homed Fabric Extender vPC Topology	130
vPC Domain	130
Peer-Keepalive Link and Messages	131
Compatibility Parameters for vPC Peer Links	131
Configuration Parameters That Must Be Identical	131
Configuration Parameters That Should Be Identical	133
vPC Peer Links	133
vPC Peer Link Overview	134
Manually Configured vPC Features	135
vPC Number	135
vPC Interactions with Other Features	136
vPC and LACP	136
vPC Peer Links and STP	136
CFSOE	137
vPC Guidelines and Limitations	137
Configuring vPCs	138
Enabling vPCs	138
Disabling vPCs	138
Creating a vPC Domain	139
Configuring a vPC Keepalive Link	140
Creating a vPC Peer Link	141
Checking the Configuration Compatibility	142

Creating an EtherChannel Host Interface	143
Moving Other EtherChannels into a vPC	144
Manually Configuring a vPC Domain MAC Address	145
Manually Configuring the System Priority	146
Manually Configuring a vPC Peer Switch Role	148
Verifying the vPC Configuration	149
vPC Example Configurations	150
Dual Homed Fabric Extender vPC Configuration Example	150
Single Homed Fabric Extender vPC Configuration Example	152
vPC Default Settings	154

CHAPTER 12**Configuring Rapid PVST+ 155**

Information About Rapid PVST+	155
Understanding STP	156
STP Overview	156
Understanding How a Topology is Created	156
Understanding the Bridge ID	156
Bridge Priority Value	157
Extended System ID	157
STP MAC Address Allocation	157
Understanding BPDUs	158
Election of the Root Bridge	159
Creating the Spanning Tree Topology	159
Understanding Rapid PVST+	160
Rapid PVST+ Overview	160
Rapid PVST+ BPDUs	162
Proposal and Agreement Handshake	163
Protocol Timers	164
Port Roles	164
Port States	165
Rapid PVST+ Port State Overview	165
Blocking State	166
Learning State	166
Forwarding State	166
Disabled State	167

Summary of Port States	167
Synchronization of Port Roles	167
Processing Superior BPDU Information	168
Processing Inferior BPDU Information	168
Detecting Unidirectional Link Failure	169
Port Cost	169
Port Priority	170
Rapid PVST+ and IEEE 802.1Q Trunks	170
Rapid PVST+ Interoperation with Legacy 802.1D STP	170
Rapid PVST+ Interoperation with 802.1s MST	171
Configuring Rapid PVST+	171
Enabling Rapid PVST+	171
Enabling Rapid PVST+ per VLAN	172
Configuring the Root Bridge ID	173
Configuring a Secondary Root Bridge	174
Configuring the Rapid PVST+ Port Priority	175
Configuring the Rapid PVST+ Pathcost Method and Port Cost	176
Configuring the Rapid PVST+ Bridge Priority of a VLAN	177
Configuring the Rapid PVST+ Hello Time for a VLAN	178
Configuring the Rapid PVST+ Forward Delay Time for a VLAN	178
Configuring the Rapid PVST+ Maximum Age Time for a VLAN	179
Specifying the Link Type	179
Restarting the Protocol	180
Verifying Rapid PVST+ Configurations	181

CHAPTER 13
Configuring Multiple Spanning Tree 183

Information About MST	183
MST Overview	183
MST Regions	184
MST BPDUs	184
MST Configuration Information	185
IST, CIST, and CST	185
IST, CIST, and CST Overview	185
Spanning Tree Operation Within an MST Region	186
Spanning Tree Operations Between MST Regions	186

MST Terminology	187
Hop Count	188
Boundary Ports	188
Detecting Unidirectional Link Failure	189
Port Cost and Port Priority	190
Interoperability with IEEE 802.1D	190
Interoperability with Rapid PVST+: Understanding PVST Simulation	191
Configuring MST	191
MST Configuration Guidelines	191
Enabling MST	191
Entering MST Configuration Mode	192
Specifying the MST Name	193
Specifying the MST Configuration Revision Number	194
Specifying the Configuration on an MST Region	195
Mapping and Unmapping VLANs to MST Instances	196
Mapping Secondary VLANs to Same MSTI as Primary VLANs for Private VLANs	197
Configuring the Root Bridge	198
Configuring a Secondary Root Bridge	199
Configuring the Port Priority	200
Configuring the Port Cost	201
Configuring the Switch Priority	202
Configuring the Hello Time	203
Configuring the Forwarding-Delay Time	204
Configuring the Maximum-Aging Time	204
Configuring the Maximum-Hop Count	205
Configuring PVST Simulation Globally	205
Configuring PVST Simulation Per Port	206
Specifying the Link Type	207
Restarting the Protocol	208
Verifying MST Configurations	208

CHAPTER 14**Configuring STP Extensions 211**

About STP Extensions 211

Information About STP Extensions 211

Understanding STP Port Types	211
Spanning Tree Edge Ports	211
Spanning Tree Network Ports	212
Spanning Tree Normal Ports	212
Understanding Bridge Assurance	212
Understanding BPDU Guard	212
Understanding BPDU Filtering	213
Understanding Loop Guard	214
Understanding Root Guard	214
Configuring STP Extensions	215
STP Extensions Configuration Guidelines	215
Configuring Spanning Tree Port Types Globally	215
Configuring Spanning Tree Edge Ports on Specified Interfaces	216
Configuring Spanning Tree Network Ports on Specified Interfaces	217
Enabling BPDU Guard Globally	218
Enabling BPDU Guard on Specified Interfaces	219
Enabling BPDU Filtering Globally	220
Enabling BPDU Filtering on Specified Interfaces	221
Enabling Loop Guard Globally	222
Enabling Loop Guard or Root Guard on Specified Interfaces	223
Verifying STP Extension Configuration	224

CHAPTER 15
Configuring the MAC Address Table 225

Information About MAC Addresses	225
Configuring MAC Addresses	225
Configuring a Static MAC Address	225
Configuring the Aging Time for the MAC Table	226
Clearing Dynamic Addresses from the MAC Table	227
Verifying the MAC Address Configuration	227

CHAPTER 16
Configuring IGMP Snooping 229

Information About IGMP Snooping	229
IGMPv1 and IGMPv2	230
IGMPv3	231
IGMP Snooping Querier	231

IGMP Forwarding	231
Configuring IGMP Snooping Parameters	232
Verifying IGMP Snooping Configuration	234

CHAPTER 17

Configuring Traffic Storm Control	237
Information About Traffic Storm Control	237
Traffic Storm Guidelines and Limitations	238
Configuring Traffic Storm Control	239
Verifying Traffic Storm Control Configuration	239
Traffic Storm Control Example Configuration	240
Default Traffic Storm Settings	240

PART III**Switch Security Features 241****CHAPTER 18**

Configuring Authentication, Authorization, and Accounting	243
Information About AAA	243
AAA Security Services	243
Benefits of Using AAA	244
Remote AAA Services	244
AAA Server Groups	245
AAA Service Configuration Options	245
Authentication and Authorization Process for User Login	246
Prerequisites for Remote AAA	247
Information about AAA Guidelines and Limitations	248
Configuring AAA	248
Configuring Console Login Authentication Methods	248
Configuring Default Login Authentication Methods	249
Enabling Login Authentication Failure Messages	251
Enabling MSCHAP Authentication	251
Configuring AAA Accounting Default Methods	253
Using AAA Server VSAs	254
About VSAs	254
VSA Format	254
Specifying Switch User Roles and SNMPv3 Parameters on AAA Servers	255
Displaying and Clearing the Local AAA Accounting Log	255

Verifying AAA Configuration	256
Example AAA Configuration	256
Default AAA Settings	256

CHAPTER 19**Configuring RADIUS 259**

Configuring RADIUS	259
Information About RADIUS	259
RADIUS Network Environments	259
RADIUS Operation	260
RADIUS Server Monitoring	260
Vendor-Specific Attributes	261
Prerequisites for RADIUS	262
Guidelines and Limitations for RADIUS	262
Configuring RADIUS Servers	262
Configuring RADIUS Server Hosts	263
Configuring RADIUS Global Preshared Keys	264
Configuring RADIUS Server Preshared Keys	264
Configuring RADIUS Server Groups	265
Allowing Users to Specify a RADIUS Server at Login	267
Configuring the Global RADIUS Transmission Retry Count and Timeout Interval	267
Configuring the RADIUS Transmission Retry Count and Timeout Interval for a Server	268
Configuring Accounting and Authentication Attributes for RADIUS Servers	270
Configuring Periodic RADIUS Server Monitoring	271
Configuring the Dead-Time Interval	272
Manually Monitoring RADIUS Servers or Groups	273
Verifying RADIUS Configuration	274
Displaying RADIUS Server Statistics	274
Example RADIUS Configuration	275
Default RADIUS Settings	275

CHAPTER 20**Configuring TACACS+ 277**

About Configuring TACACS+	277
Information About TACACS+	277
TACACS+ Advantages	277

User Login with TACACS+ 278	278
Default TACACS+ Server Encryption Type and Preshared Key 278	278
TACACS+ Server Monitoring 279	279
Prerequisites for TACACS+ 279	279
Guidelines and Limitations for TACACS+ 279	279
Configuring TACACS+ 280	280
TACACS+ Server Configuration Process 280	280
Enabling TACACS+ 280	280
Configuring TACACS+ Server Hosts 281	281
Configuring TACACS+ Global Preshared Keys 282	282
Configuring TACACS+ Server Preshared Keys 283	283
Configuring TACACS+ Server Groups 284	284
Specifying a TACACS+ Server at Login 285	285
Configuring the Global TACACS+ Timeout Interval 286	286
Configuring the Timeout Interval for a Server 287	287
Configuring TCP Ports 288	288
Configuring Periodic TACACS+ Server Monitoring 288	288
Configuring the Dead-Time Interval 290	290
Manually Monitoring TACACS+ Servers or Groups 291	291
Disabling TACACS+ 291	291
Displaying TACACS+ Statistics 292	292
Verifying TACACS+ Configuration 292	292
Example TACACS+ Configuration 293	293
Default TACACS+ Settings 293	293

CHAPTER 21
Configuring SSH and Telnet 295

Configuring SSH and Telnet 295	295
Information About SSH and Telnet 295	295
SSH Server 295	295
SSH Client 295	295
SSH Server Keys 295	295
Telnet Server 296	296
Guidelines and Limitations for SSH 296	296
Configuring SSH 296	296
Generating SSH Server Keys 296	296

Specifying the SSH Public Keys for User Accounts	297
Specifying the SSH Public Keys in Open SSH Format	297
Specifying the SSH Public Keys in IETF SECSH Format	298
Specifying the SSH Public Keys in PEM-Formatted Public Key Certificate Form	299
Starting SSH Sessions to Remote Devices	300
Clearing SSH Hosts	300
Disabling the SSH Server	301
Deleting SSH Server Keys	301
Clearing SSH Sessions	302
SSH Example Configuration	303
Configuring Telnet	304
Enabling the Telnet Server	304
Reenabling the Telnet Server	304
Starting Telnet Sessions to Remote Devices	304
Clearing Telnet Sessions	305
Verifying the SSH and Telnet Configuration	305
Default SSH Settings	306

CHAPTER 22

Configuring Access Control Lists	307
Information About ACLs	307
IP ACL Types and Applications	307
Application Order	308
Rules	308
Source and Destination	308
Protocols	308
Implicit Rules	309
Additional Filtering Options	309
Sequence Numbers	310
Logical Operators and Logical Operation Units	310
Configuring IP ACLs	311
Creating an IP ACL	311
Changing an IP ACL	312
Removing an IP ACL	313
Changing Sequence Numbers in an IP ACL	314

Applying an IP ACL as a Port ACL	314
Verifying IP ACL Configurations	315
Displaying and Clearing IP ACL Statistics	316
Configuring MAC ACLs	316
Creating a MAC ACL	316
Changing a MAC ACL	317
Removing a MAC ACL	318
Changing Sequence Numbers in a MAC ACL	319
Applying a MAC ACL as a Port ACL	320
Verifying MAC ACL Configurations	321
Displaying and Clearing MAC ACL Statistics	321
Example Configuration for MAC ACLs	321
Information About VLAN ACLs	321
VACLs and Access Maps	321
VACLs and Actions	321
Statistics	322
Configuring VACLs	322
Creating or Changing a VACL	322
Removing a VACL	323
Applying a VACL to a VLAN	324
Verifying VACL Configuration	324
Displaying and Clearing VACL Statistics	325
Example Configuration for VACL	325
Default ACL Settings	325

PART IV
System Management 327

CHAPTER 23
Using Cisco Fabric Services 329

Using Cisco Fabric Services	329
Information About CFS	329
CFS Distribution	330
CFS Distribution Modes	330
Uncoordinated Distribution	330
Coordinated Distribution	330
Unrestricted Uncoordinated Distributions	331

Disabling or Enabling CFS Distribution on a Switch	331
Verifying CFS Distribution Status	332
CFS Distribution over IP	332
CFS Distribution over Fibre Channel	333
CFS Distribution Scopes	333
CFS Merge Support	334
CFS Support for Applications	334
CFS Application Requirements	334
Enabling CFS for an Application	335
Verifying Application Registration Status	335
Locking the Network	335
Verifying CFS Lock Status	336
Committing Changes	336
Discarding Changes	336
Saving the Configuration	337
Clearing a Locked Session	337
CFS Regions	337
About CFS Regions	337
Example Scenario	337
Managing CFS Regions	338
Creating CFS Regions	338
Assigning Applications to CFS Regions	338
Moving an Application to a Different CFS Region	339
Removing an Application from a Region	339
Deleting CFS Regions	340
Configuring CFS over IP	340
Enabling CFS over IPv4	340
Enabling CFS over IPv6	341
Verifying the CFS Over IP Configuration	342
Configuring IP Multicast Address for CFS over IP	342
Configuring IPv4 Multicast Address for CFS	342
Configuring IPv6 Multicast Address for CFS	343
Verifying IP Multicast Address Configuration for CFS over IP	343
Displaying CFS Distribution Information	343
Default CFS Settings	345

CHAPTER 24**Configuring User Accounts and RBAC 347**

Configuring User Accounts and RBAC 347

Information About User Accounts and RBAC 347

About User Accounts 347

Characteristics of Strong Passwords 348

About User Roles 348

About Rules 349

About User Role Policies 349

Guidelines and Limitations for User Accounts 350

Configuring User Accounts 350

Configuring RBAC 351

Creating User Roles and Rules 351

Creating Feature Groups 353

Changing User Role Interface Policies 353

Changing User Role VLAN Policies 354

Changing User Role VSAN Policies 355

Verifying User Accounts and RBAC Configuration 356

Default User Account and RBAC Settings 356

CHAPTER 25**Configuring Session Manager 359**

Configuring Session Manager 359

Information About Session Manager 359

Configuration Guidelines and Limitations 359

Configuring Session Manager 360

Creating a Session 360

Configuring ACLs in a Session 360

Verifying a Session 361

Committing a Session 361

Saving a Session 361

Discarding a Session 362

Session Manager Example Configuration 362

Verifying Session Manager Configuration 362

CHAPTER 26**Configuring Online Diagnostics 363**

Information About Online Diagnostics	363
Online Diagnostics Overview	363
Bootup Diagnostics	363
Health Monitoring Diagnostics	364
Expansion Module Diagnostics	365
Configuring Online Diagnostics	366
Verifying Online Diagnostics Configuration	367
Default GOLD Settings	367

CHAPTER 27**Configuring System Message Logging 369**

Information About System Message Logging	369
syslog Servers	370
Configuring System Message Logging	370
Configuring System Message Logging to Terminal Sessions	370
Configuring System Message Logging to a File	373
Configuring Module and Facility Messages Logging	374
Configuring Logging Timestamps	376
Configuring syslog Servers	377
Configuring syslog on a UNIX or Linux System	378
Configuring syslog Server Configuration Distribution	379
Displaying and Clearing Log Files	381
Verifying System Message Logging Configuration	382
Default System Message Logging Settings	382

CHAPTER 28**Configuring Smart Call Home 385**

Configuring Smart Call Home	385
Information About Call Home	385
Call Home Overview	385
Destination Profiles	386
Call Home Alert Groups	386
Call Home Message Levels	388
Obtaining Smart Call Home	389
Prerequisites for Call Home	389
Configuration Guidelines and Limitations	390
Configuring Call Home	390

Procedures for Configuring Call Home	390
Configuring Contact Information	391
Creating a Destination Profile	392
Modifying a Destination Profile	394
Associating an Alert Group with a Destination Profile	395
Adding show Commands to an Alert Group	396
Configuring E-Mail	397
Configuring Periodic Inventory Notification	398
Disabling Duplicate Message Throttle	399
Enabling or Disabling Call Home	399
Testing Call Home Communications	400
Verifying Call Home Configuration	400
Default Call Home Settings	401
Additional References	402
Call Home Message Formats	402
Sample syslog Alert Notification in Full-Text Format	409
Sample syslog Alert Notification in XML Format	409

CHAPTER 29**Configuring SNMP 415**

Information About SNMP	415
SNMP Functional Overview	415
SNMP Notifications	416
SNMPv3	416
Security Models and Levels for SNMPv1, v2, v3	416
User-Based Security Model	417
CLI and SNMP User Synchronization	418
Group-Based SNMP Access	419
Configuration Guidelines and Limitations	419
Configuring SNMP	419
Configuring SNMP Users	419
Enforcing SNMP Message Encryption	420
Assigning SNMPv3 Users to Multiple Roles	420
Creating SNMP Communities	421
Configuring SNMP Notification Receivers	421
Configuring the Notification Target User	422

Enabling SNMP Notifications	422
Configuring Link Notifications	424
Disabling Link Notifications on an Interface	425
Enabling One-Time Authentication for SNMP over TCP	425
Assigning SNMP Switch Contact and Location Information	426
Configuring the Context to Network Entity Mapping	426
Verifying SNMP Configuration	427
Default SNMP Settings	427

CHAPTER 30
Configuring RMON 429

Configuring RMON	429
Information About RMON	429
RMON Alarms	429
RMON Events	430
Configuration Guidelines and Limitations	430
Configuring RMON	430
Configuring RMON Alarms	430
Configuring RMON Events	432
Verifying RMON Configuration	432
Default RMON Settings	433

PART V
Fibre Channel over Ethernet 435

CHAPTER 31
Configuring FCoE 437

Information About FCoE	437
Information About FCoE and FIP	437
FCoE Initiation Protocol	437
FIP Virtual Link Instantiation	438
FCoE Frame Format	438
VLAN Tagging for FCoE Frames	438
FIP Ethernet Frame Format	439
Pre-FIP Virtual Link Instantiation	439
Information About DCBX	439
Data Center Bridging Exchange Protocol	439
DCBX Feature Negotiation	440

Lossless Ethernet	441
Logical Link Up/Down	441
Converged Network Adapters	441
FCoE Topologies	442
Directly Connected CNA Topology	442
Remotely Connected CNA Topology	443
FCoE Best Practices	444
Directly Connected CNA Best Practice	444
Remotely Connected CNA Best Practice	446
Licensing Requirements for FCoE	447
Configuring FCoE	447
Enabling FCoE	447
Disabling FCoE	448
Disabling LAN Traffic on an FCoE Link	449
Configuring the FC-Map	450
Configuring the Fabric Priority	450
Setting the Advertisement Interval	451
Configuring LLDP	452
Configuring Global LLDP Commands	452
Configuring Interface LLDP Commands	452
Verifying FCoE Configuration	453
<hr/>	
CHAPTER 32	Configuring FCoE VLANs and Virtual Interfaces 457
	Information About Virtual Interfaces 457
	Guidelines and Limitations for FCoE VLANs and Virtual Interfaces 457
	Configuring Virtual Interfaces 459
	Mapping a VSAN to a VLAN 459
	Creating a Virtual Fibre Channel Interface 459
	Associating a Virtual Fibre Channel Interface to a VSAN 461
	Verifying the Virtual Interface 461
	Mapping VSANs to VLANs Example Configuration 463
<hr/>	
PART VI	Quality of Service 465
<hr/>	
CHAPTER 33	Configuring QoS 467

Information About QoS	467
MQC	468
System Classes	468
Default System Classes	469
Policy Types	469
Link-Level Flow Control	471
Priority Flow Control	472
MTU	472
Trust Boundaries	473
Ingress Queuing Policies	473
Ingress Classification Policies	473
Egress Queuing Policies	473
QoS for Multicast Traffic	474
Policy for Fibre Channel Interfaces	475
QoS for Traffic Directed to the CPU	475
QoS Configuration Guidelines and Limitations	475
Configuring System Classes	476
Configuring Class Maps	476
Configuring ACL Classification	476
Configuring CoS Classification	477
Configuring DSCP Classification	478
Configuring IP RTP Classification	480
Configuring Precedence Classification	481
Configuring Protocol Classification	483
Configuring QoS Group Classification	484
Configuring Policy Maps	485
Configuring Type Network QoS Policies	486
Configuring Type Queuing Policies	488
Configuring Type QoS Policies	489
Attaching the System Service Policy	490
Restoring the Default System Service Policies	491
Enabling the Jumbo MTU	493
Verifying the Jumbo MTU	493
Configuring QoS on Interfaces	494
Configuring Untagged CoS	494

Configuring Interface Service Policy	495
Configuring Priority Flow Control and Link-Level Flow Control	496
Configuring Priority Flow Control	496
Configuring Link-Level Flow Control	497
Verifying QoS Configuration	498
Example QoS Configurations	504
QoS Example 1	504
QoS Example 2	505
QoS Example 3	507

PART VII
SAN Switching 509

CHAPTER 34
Configuring Fibre Channel Interfaces 511

Configuring Fibre Channel Interfaces	511
Information About Fibre Channel Interfaces	511
Licensing Requirements for Fibre Channel	511
Physical Fibre Channel Interfaces	511
Virtual Fibre Channel Interfaces	512
Interface Modes	512
E Port	513
F Port	513
NP Port	513
TE Port	513
SD Port	514
Auto Mode	514
Interface States	514
Administrative States	514
Operational States	514
Reason Codes	515
Buffer-to-Buffer Credits	517
Configuring Fibre Channel Interfaces	518
Configuring a Fibre Channel Interface	518
Configuring a Range of Fibre Channel Interfaces	518
Setting the Interface Administrative State	519
Configuring Interface Modes	519

Configuring the Interface Description	520
Configuring Port Speeds	521
Autosensing	521
Configuring SD Port Frame Encapsulation	521
Configuring Receive Data Field Size	522
Understanding Bit Error Thresholds	522
Configuring Buffer-to-Buffer Credits	523
Configuring Global Attributes for Fibre Channel Interfaces	524
Configuring Switch Port Attribute Default Values	524
About N Port Identifier Virtualization	525
Enabling N Port Identifier Virtualization	526
Verifying Fibre Channel Interfaces	526
Verifying SFP Transmitter Types	526
Verifying Interface Information	526
Verifying BB_Credit Information	528
Default Fibre Channel Interface Settings	528

CHAPTER 35**Configuring Domain Parameters 531**

Configuring Domain Parameters	531
Information About Fibre Channel Domains	531
About Domain Restart	532
Restarting a Domain	533
About Domain Manager Fast Restart	533
Enabling Domain Manager Fast Restart	533
About Switch Priority	534
Configuring Switch Priority	534
About fcdomain Initiation	535
Disabling or Reenabling fcdomains	535
Configuring Fabric Names	535
About Incoming RCFs	536
Rejecting Incoming RCFs	536
About Autoreconfiguring Merged Fabrics	537
Enabling Autoreconfiguration	537
Domain IDs	538
About Domain IDs	538

Specifying Static or Preferred Domain IDs	540
About Allowed Domain ID Lists	541
Configuring Allowed Domain ID Lists	541
About CFS Distribution of Allowed Domain ID Lists	542
Enabling Distribution	542
Locking the Fabric	543
Committing Changes	543
Discarding Changes	544
Clearing a Fabric Lock	544
Displaying CFS Distribution Status	544
Displaying Pending Changes	545
Displaying Session Status	545
About Contiguous Domain ID Assignments	545
Enabling Contiguous Domain ID Assignments	545
FC IDs	546
About Persistent FC IDs	546
Enabling the Persistent FC ID Feature	547
Persistent FC ID Configuration Guidelines	547
Configuring Persistent FC IDs	548
About Unique Area FC IDs for HBAs	548
Configuring Unique Area FC IDs for an HBA	549
About Persistent FC ID Selective Purging	550
Purging Persistent FC IDs	550
Verifying fcdomain Information	551
Default Fibre Channel Domain Settings	552

CHAPTER 36

Configuring N Port Virtualization	553
Configuring N Port Virtualization	553
Information About NPV	553
NPV Overview	553
NPV Mode	554
Server Interfaces	554
NP Uplinks	555
FLOGI Operation	555
NPV Traffic Management	556

Automatic Uplink Selection	556
Traffic Maps	556
Disruptive Load Balancing	556
NPV Traffic Management Guidelines	557
NPV Guidelines and Limitations	557
Configuring NPV	558
Enabling NPV	558
Configuring NPV Interfaces	559
Configuring an NP Interface	559
Configuring a Server Interface	559
Configuring NPV Traffic Management	560
Configuring NPV Traffic Maps	560
Enabling Disruptive Load Balancing	560
Verifying NPV	561
Verifying NPV Examples	561
Verifying NPV Traffic Management	562

CHAPTER 37

Configuring VSAN Trunking	563
Configuring VSAN Trunking	563
Information About VSAN Trunking	563
VSAN Trunking Mismatches	564
VSAN Trunking Protocol	564
Configuring VSAN Trunking	565
Guidelines and Restrictions	565
Enabling or Disabling the VSAN Trunking Protocol	565
About Trunk Mode	565
Configuring Trunk Mode	566
About Trunk-Allowed VSAN Lists	567
Configuring an Allowed-Active List of VSANs	568
Displaying VSAN Trunking Information	569
Default Trunk Configuration Settings	569

CHAPTER 38

Configuring SAN Port Channel	571
Configuring SAN Port Channels	571
Information About SAN Port Channels	571

Understanding Port Channels and VSAN Trunking	572
Understanding Load Balancing	572
Configuring SAN Port Channels	574
SAN Port Channel Configuration Guidelines	575
Creating a SAN Port Channel	576
SAN Port Channel Configuration Guidelines	576
About SAN Port Channel Deletion	577
Configuring Active Mode SAN Port Channel	577
Deleting SAN Port Channels	578
Interfaces in a SAN Port Channel	579
About Interface Addition to a SAN Port Channel	579
Compatibility Check	579
Suspended and Isolated States	579
Adding an Interface to a SAN Port Channel	579
Forcing an Interface Addition	580
About Interface Deletion from a SAN Port Channel	581
Deleting an Interface from a SAN Port Channel	581
SAN Port Channel Protocol	582
About Channel Group Creation	582
Autocreation Guidelines	584
Enabling and Configuring Autocreation	585
About Manually Configured Channel Groups	585
Converting to Manually Configured Channel Groups	585
Verifying SAN Port Channel Configuration	586
Default Settings for SAN Port Channels	587

CHAPTER 39**Configuring and Managing VSANs 589**

Configuring and Managing VSANs	589
Information About VSANs	589
VSAN Topologies	589
VSAN Advantages	592
VSANs Versus Zones	592
Configuring VSANs	593
About VSAN Creation	594
Creating VSANs Statically	594

About Port VSAN Membership	595
Assigning Static Port VSAN Membership	595
Displaying VSAN Static Membership	596
About the Default VSAN	596
About the Isolated VSAN	597
Displaying Isolated VSAN Membership	597
Operational State of a VSAN	597
About Static VSAN Deletion	597
Deleting Static VSANs	598
About Load Balancing	599
Configuring Load Balancing	599
About Interop Mode	600
Displaying Static VSAN Configuration	600
Default VSAN Settings	600

CHAPTER 40

Configuring and Managing Zones	603
Configuring and Managing Zones	603
Information About Zoning	603
Zoning Features	603
Zoning Example	605
Zone Implementation	605
Active and Full Zone Set Configuration Guidelines	606
Configuring Zones	609
Configuring Zones Example	609
Zone Sets	611
Activating a Zone Set	611
About the Default Zone	612
Configuring the Default Zone Access Permission	612
About FC Alias Creation	613
Creating FC Aliases	613
Creating FC Aliases Example	614
Creating Zone Sets and Adding Member Zones	615
Zone Enforcement	616
Zone Set Distribution	616
Enabling Full Zone Set Distribution	616

Enabling a One-Time Distribution	617
About Recovering from Link Isolation	618
Importing and Exporting Zone Sets	618
Zone Set Duplication	619
Copying Zone Sets	619
Renaming Zones, Zone Sets, and Aliases	620
Cloning Zones, Zone Sets, FC Aliases, and Zone Attribute Groups	621
Clearing the Zone Server Database	621
Verifying Zone Information	622
Enhanced Zoning	622
About Enhanced Zoning	622
Changing from Basic Zoning to Enhanced Zoning	623
Changing from Enhanced Zoning to Basic Zoning	624
Enabling Enhanced Zoning	624
Modifying the Zone Database	625
Releasing Zone Database Locks	626
Merging the Database	626
Configuring Zone Merge Control Policies	627
Default Zone Policies	628
Configuring System Default Zoning Settings	628
Verifying Enhanced Zone Information	629
Compacting the Zone Database	629
Zone and Zone Set Analysis	630
Default Basic Zone Settings	630

CHAPTER 41
Distributing Device Alias Services 631

Distributing Device Alias Services	631
Information About Device Aliases	631
Device Alias Features	631
Device Alias Requirements	632
Zone Aliases Versus Device Aliases	632
Device Alias Databases	633
Creating Device Aliases	633
Device Alias Modes	634
Changing Device Alias Mode Guidelines	634

Configuring Device Alias Modes	635
About Device Alias Distribution	635
Locking the Fabric	636
Committing Changes	636
Discarding Changes	636
Fabric Lock Override	637
Disabling and Enabling Device Alias Distribution	637
About Legacy Zone Alias Configuration	638
Importing a Zone Alias	639
Device Alias Database Merge Guidelines	639
Verifying Device Alias Configuration	640
Default Device Alias Settings	641

CHAPTER 42
Configuring Fibre Channel Routing Services and Protocols 643

Configuring Fibre Channel Routing Services and Protocols	643
Information About FSPF	643
FSPF Examples	644
Fault Tolerant Fabric Example	644
Redundant Link Example	644
FSPF Global Configuration	645
About SPF Computational Hold Times	645
About Link State Records	645
Configuring FSPF on a VSAN	646
Resetting FSPF to the Default Configuration	647
Enabling or Disabling FSPF	647
Clearing FSPF Counters for the VSAN	648
FSPF Interface Configuration	648
About FSPF Link Cost	648
Configuring FSPF Link Cost	648
About Hello Time Intervals	649
Configuring Hello Time Intervals	649
About Dead Time Intervals	650
Configuring Dead Time Intervals	650
About Retransmitting Intervals	650
Configuring Retransmitting Intervals	651

About Disabling FSPF for Specific Interfaces	651
Disabling FSPF for Specific Interfaces	651
Clearing FSPF Counters for an Interface	652
FSPF Routes	652
About Fibre Channel Routes	653
Configuring Fibre Channel Routes	653
In-Order Delivery	654
About Reordering Network Frames	654
About Reordering SAN Port Channel Frames	655
About Enabling In-Order Delivery	655
Enabling In-Order Delivery Globally	656
Enabling In-Order Delivery for a VSAN	656
Displaying the In-Order Delivery Status	657
Configuring the Drop Latency Time	657
Displaying Latency Information	658
Flow Statistics Configuration	658
About Flow Statistics	658
Counting Aggregated Flow Statistics	658
Counting Individual Flow Statistics	659
Clearing FIB Statistics	659
Displaying Flow Statistics	660
Default FSPF Settings	660

CHAPTER 43

Managing FLOGI, Name Server, FDMI, and RSCN Databases	663
Managing FLOGI, Name Server, FDMI, and RSCN Databases	663
Information About Fabric Login	663
Name Server Proxy	664
About Registering Name Server Proxies	664
Registering Name Server Proxies	664
About Rejecting Duplicate pWWNs	664
Rejecting Duplicate pWWNs	664
About Name Server Database Entries	665
Displaying Name Server Database Entries	665
FDMI	666
Displaying FDMI	666

RSCN	666
About RSCN Information	667
Displaying RSCN Information	667
About the multi-pid Option	667
Configuring the multi-pid Option	667
Suppressing Domain Format SW-RSCNs	668
Clearing RSCN Statistics	668
Configuring the RSCN Timer	669
Verifying the RSCN Timer Configuration	670
RSCN Timer Configuration Distribution	670
Enabling RSCN Timer Configuration Distribution	670
Locking the Fabric	671
Committing the RSCN Timer Configuration Changes	671
Discarding the RSCN Timer Configuration Changes	671
Clearing a Locked Session	672
Displaying RSCN Configuration Distribution Information	672
Default RSCN Settings	672

CHAPTER 44
Discovering SCSI Targets 675

Discovering SCSI Targets	675
Information About SCSI LUN Discovery	675
About Starting SCSI LUN Discovery	675
Starting SCSI LUN Discovery	675
About Initiating Customized Discovery	676
Initiating Customized Discovery	676
Displaying SCSI LUN Information	677

CHAPTER 45
Advanced Fibre Channel Features and Concepts 679

Advanced Fibre Channel Features and Concepts	679
Fibre Channel Timeout Values	679
Timer Configuration Across All VSANs	679
Timer Configuration Per-VSAN	680
About fctimer Distribution	681
Enabling or Disabling fctimer Distribution	681
Committing fctimer Changes	682

Discarding fctimer Changes	682
Fabric Lock Override	683
Fabric Database Merge Guidelines	683
Verifying Configured fctimer Values	684
World Wide Names	684
Verifying WWN Information	685
Link Initialization WWN Usage	685
Configuring a Secondary MAC Address	685
FC ID Allocation for HBAs	686
Default Company ID List	686
Verifying the Company ID Configuration	687
Switch Interoperability	688
About Interop Mode	688
Configuring Interop Mode 1	691
Verifying Interoperating Status	692
Default Settings for Advanced Features	697

CHAPTER 46
Configuring FC-SP and DHCHAP 699

Configuring FC-SP and DHCHAP	699
Information About Fabric Authentication	699
DHCHAP	700
DHCHAP Compatibility with Fibre Channel Features	701
About Enabling DHCHAP	701
Enabling DHCHAP	702
About DHCHAP Authentication Modes	702
Configuring the DHCHAP Mode	703
About the DHCHAP Hash Algorithm	704
Configuring the DHCHAP Hash Algorithm	704
About the DHCHAP Group Settings	705
Configuring the DHCHAP Group Settings	705
About the DHCHAP Password	705
Configuring DHCHAP Passwords for the Local Switch	706
About Password Configuration for Remote Devices	706
Configuring DHCHAP Passwords for Remote Devices	706
About the DHCHAP Timeout Value	707

- Configuring the DHCPAP Timeout Value 707
- Configuring DHCPAP AAA Authentication 708
- Displaying Protocol Security Information 708
- Sample Configuration 708
- Default Fabric Security Settings 710

CHAPTER 47**Configuring Port Security 713**

- Configuring Port Security 713
 - Information About Port Security 713
 - Port Security Enforcement 713
 - About Auto-Learning 714
 - Port Security Activation 714
 - Configuring Port Security 715
 - Configuring Port Security with Auto-Learning and CFS Distribution 715
 - Configuring Port Security with Auto-Learning without CFS 716
 - Configuring Port Security with Manual Database Configuration 717
 - Enabling Port Security 717
 - Port Security Activation 718
 - Activating Port Security 718
 - Database Activation Rejection 718
 - Forcing Port Security Activation 719
 - Database Reactivation 719
 - Auto-Learning 720
 - About Enabling Auto-Learning 720
 - Enabling Auto-Learning 720
 - Disabling Auto-Learning 721
 - Auto-Learning Device Authorization 721
 - Authorization Scenario 722
 - Port Security Manual Configuration 724
 - WWN Identification Guidelines 724
 - Adding Authorized Port Pairs 724
 - Port Security Configuration Distribution 725
 - Enabling Port Security Distribution 726
 - Locking the Fabric 727
 - Committing the Changes 727

Discarding the Changes	727
Activation and Auto-Learning Configuration Distribution	728
Port Security Database Merge Guidelines	729
Database Interaction	730
Database Scenarios	731
Copying the Port Security Database	732
Deleting the Port Security Database	732
Clearing the Port Security Database	732
Displaying Port Security Configuration	733
Default Port Security Settings	733

CHAPTER 48**Configuring Fabric Binding 735**

Configuring Fabric Binding	735
Information About Fabric Binding	735
Licensing Requirements for Fabric Binding	735
Port Security Versus Fabric Binding	735
Fabric Binding Enforcement	736
Configuring Fabric Binding	736
Configuring Fabric Binding	736
Enabling Fabric Binding	737
About Switch WWN Lists	738
Configuring Switch WWN List	738
About Fabric Binding Activation and Deactivation	738
Activating Fabric Binding	739
Forcing Fabric Binding Activation	739
Copying Fabric Binding Configurations	740
Clearing the Fabric Binding Statistics	740
Deleting the Fabric Binding Database	740
Verifying Fabric Binding Information	741
Default Fabric Binding Settings	742

CHAPTER 49**Configuring Fabric Configuration Servers 743**

Configuring Fabric Configuration Servers	743
Information About FCS	743
FCS Characteristics	744

FCS Name Specification 745
 Displaying FCS Information 745
 Default FCS Settings 745

CHAPTER 50
Configuring Port Tracking 747

Configuring Port Tracking 747
 Information About Port Tracking 747
 Configuring Port Tracking 748
 Enabling Port Tracking 749
 About Configuring Linked Ports 749
 Operationally Binding a Tracked Port 749
 About Tracking Multiple Ports 750
 Tracking Multiple Ports 751
 About Monitoring Ports in a VSAN 751
 Monitoring Ports in a VSAN 751
 About Forceful Shutdown 752
 Forcefully Shutting Down a Tracked Port 752
 Displaying Port Tracking Information 753
 Default Port Tracking Settings 753

PART VIII
Troubleshooting 755

CHAPTER 51
Configuring SPAN 757

Configuring SPAN 757
 SPAN Sources 757
 Characteristics of Source Ports 757
 SPAN Destinations 758
 Characteristics of Destination Ports 758
 Configuring SPAN 759
 Creating and Deleting a SPAN Session 759
 Configuring the Destination Port 759
 Configuring an Ethernet Destination Port 759
 Configuring Fibre Channel Destination Port 760
 Configuring Source Ports 761
 Configuring Source Port Channels, VLANs, or VSANs 762

Configuring the Description of a SPAN Session	763
Activating a SPAN Session	763
Suspending a SPAN Session	763
Displaying SPAN Information	764

CHAPTER 52**Troubleshooting 767**

Troubleshooting	767
Recovering a Lost Password	767
Using the CLI with Network-Admin Privileges	767
Power Cycling the Switch	768
Using Ethalyzer	769
Troubleshooting Fibre Channel	773
fctrace	773
fcping	774
Verifying Switch Connectivity	775
show tech-support Command	776
show tech-support brief Command	778
show tech-support fc Command	781
show tech-support platform Command	783
Default Settings for Troubleshooting Features	785

APPENDIX A**Configuration Limits 787**

Cisco Nexus 5000 Series Configuration Limits	787
--	-----



Trademark Notices

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.

CCDE, CCENT, CCSI, Cisco Eos, Cisco HealthPresence, Cisco Ironport, the Cisco logo, Cisco Lumin, Cisco Nexus, Cisco Nurse Connect, Cisco Stackpower, Cisco StadiumVision, Cisco TelePresence, Cisco Unified Computing System, Cisco WebEx, DCE, Flip Channels, Flip for Good, Flip Mino, Flip Video, Flip Video (Design), Flipshare (Design), Flip Ultra, and Welcome to the Human Network are trademarks; Changing the Way We Work, Live, Play, and Learn, Cisco Store, and Flip Gift Card are service marks; and Access Registrar, Aironet, AsyncOS, Bringing the Meeting To You, Catalyst, CCDA, CCDP, CCIE, CCIP, CCNA, CCNP, CCSP, CCVP, Cisco, the Cisco Certified Internetwork Expert logo, Cisco IOS, Cisco Press, Cisco Systems, Cisco Systems Capital, the Cisco Systems logo, Cisco Unity, Collaboration Without Limitation, EtherFast, EtherSwitch, Event Center, Fast Step, Follow Me Browsing, FormShare, GigaDrive, HomeLink, Internet Quotient, IOS, iPhone, iQuick Study, IronPort, the IronPort logo, LightStream, Linksys, MediaTone, MeetingPlace, MeetingPlace Chime Sound, MGX, Networkers, Networking Academy, Network Registrar,

PCNow, PIX, PowerPanels, ProConnect, ScriptShare, SenderBase, SMARTnet, Spectrum Expert, StackWise, The Fastest Way to Increase Your Internet Quotient, TransPath, WebEx, and the WebEx logo are registered trademarks of Cisco Systems, Inc. and/or its affiliates in the United States and certain other countries.

All other trademarks mentioned in this document or website 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. (0907R)

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.

© 2008–2009 Cisco Systems, Inc. All rights reserved.



Preface

The preface describes the audience, organization, and conventions of the Cisco Nexus 5000 Series Switch NX-OS Software Configuration Guide. It also provides information on how to obtain related documentation.

- [Audience, page xliii](#)
- [Document Organization, page xliii](#)
- [Document Conventions, page xliv](#)
- [Related Documentation, page xlv](#)

Audience

This guide is for experienced network administrators who are responsible for configuring and maintaining n5k switches.

Document Organization

This document is organized as follows:

Part or Chapter	Description
Overview, on page 3	Presents an overview of the Cisco Nexus 5000 Series switch.
Configuration Fundamentals, on page 13	Describes how to use the CLI and initial switch configuration.
LAN Switching, on page 59	Describes how to configure Ethernet interfaces, VLANs, STP, EtherChannels, trunks, the MAC address table, and IGMP snooping.
Switch Security Features	Describes how to configure AAA, RADIUS, TACACS+, SSH/Telnet, and ACLs.

Part or Chapter	Description
System Management	Describes how to configure CFS, RBAC, System Message Logging, Call Home, SNMP, RMON, network management interfaces, storm control, and SPAN.
Fibre Channel over Ethernet	Describes how to configure FCoE and virtual interfaces.
Quality of Server	Describes how to configure QoS.
SAN Switching	Describes how to configure Fibre Channel interfaces and Fibre Channel capabilities including NPV, SAN port channels, zones, DDAS, FSPF, and security features.
Troubleshooting	Describes how to perform basic troubleshooting.

Document Conventions

Command descriptions use the following conventions:

Convention	Description
bold	Bold text indicates the commands and keywords that you enter literally as shown.
<i>Italic</i>	Italic text indicates arguments for which the user supplies the values.
[x]	Square brackets enclose an optional element(keyword or argument).
[x y]	Square brackets enclosing keywords or arguments separated by a vertical bar indicate an optional choice.
{x y}	Braces enclosing keywords or arguments separated by a vertical bar indicate a required choice.
[x {y z}]	Nested set of square brackets or braces indicate optional or required choices within optional or required elements. Braces and a vertical bar within square brackets indicate a required choice within an optional element.
<i>variable</i>	Indicates a variable for which you supply values, in context where italics cannot be used.
string	A nonquoted set of characters. Do not use quotation marks around the string or the string will include the quotation marks.

Screen examples use the following conventions:

Convention	Description
<code>screen font</code>	Terminal sessions and information the switch displays are in screen font.
boldface screen font	Information you must enter is in boldface screen font.
<i>italic screen font</i>	Arguments for which you supply values are in italic screen font.
<>	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.

This document uses the following conventions:



Note

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



Caution

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

Related Documentation

Documentation for Cisco Nexus 5000 Series Switches and Cisco Nexus 2000 Series Fabric Extender is available at the following URL: http://www.cisco.com/en/us/products/ps9670/tsd_products_support_series_home.html.

The following are related Cisco Nexus 5000 Series Switches and Cisco Nexus 2000 Series Fabric Extender documents:

Cisco Nexus 5000 Series Switch CLI Software Configuration Guide, Cisco NX-OS Release 4.1

Cisco Nexus 5000 Series Command Reference, Cisco NX-OS Release 4.1

Cisco Nexus 5000 Series Hardware Installation Guide

Cisco Nexus 5000 Series System Messages Reference

Cisco Nexus 5000 Series and Cisco Nexus 2000 Series Release Notes

Cisco Nexus 2000 Series Fabric Extender Software Configuration Guide, Cisco NX-OS Release 4.1

Cisco Nexus 2000 Series Fabric Extender Hardware Installation Guide

Cisco MDS 9000 and Nexus 5000 Series Fabric Manager Software Configuration Guide, Cisco Fabric Manager Release 4.1

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, at <http://www.cisco.com/en/US/docs/general/whatsnew/whatsnew.html>.

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.



CHAPTER

1

New and Changed Information

This chapter provides release-specific information for each new and changed feature in the *Cisco Nexus 5000 Series NX-OS Configuration Guide*.

- [New and Changed Information for the Cisco Nexus 5000 Series, page 1](#)

New and Changed Information for the Cisco Nexus 5000 Series

This chapter provides release specific information for each new and changed feature in the *Cisco Nexus 5000 Series NX-OS Configuration Guide*.

To check for additional information about Cisco NX-OS Release 4.1(3)N1(1), see the *Cisco Nexus 5000 Series and Cisco Nexus 2000 Series Release Notes, 31/July/2009* available at the following Cisco website: http://www.cisco.com/en/US/products/ps9670/prod_release_notes_list.html.

This table summarizes the new and changed features for the *Cisco Nexus 5000 Series NX-OS Configuration Guide* for Cisco NX-OS Release 4.1(3)N1(1), and tells you where they are documented.

Table 1: New and Changed Features for Release 4.1(3)N1(1)

Feature	Description	Changed in Release	Where Documented
PVLAN promiscuous and isolated trunk ports	Promiscuous and isolated trunk ports allow greater control of private VLANs.	4.1(3)N1(1)	Configuring Private VLANs
QoS enhancements	Three new class-map and policy-map object types which provide greater flexibility in QoS configuration.	4.1(3)N1(1)	Quality of Service
T11-compliant FCoE	The FCoE implementation on the Cisco Nexus 5000 Series is T11-compliant.	4.1(3)N1(1)	Fibre Channel over Ethernet

Feature	Description	Changed in Release	Where Documented
vPC	Added the configuration of virtual Port Channels (vPC).	4.1(3)N1(1)	Configuring Virtual Port Channels



CHAPTER 2

Overview

This chapter describes the Cisco Nexus 5000 Series switches. It includes the following sections:

- [Information About Cisco Nexus 5000 Series Switches, page 3](#)
- [New Technologies in the Cisco Nexus 5000 Series, page 3](#)
- [Cisco Nexus 5000 Series Switch Software, page 5](#)
- [Typical Deployment Topologies, page 8](#)
- [Supported Standards, page 11](#)

Information About Cisco Nexus 5000 Series Switches

The Cisco Nexus 5000 Series is a family of top-of-rack switches for the data center. The Cisco Nexus 5000 Series offers high-speed Ethernet switching and supports Fibre Channel over Ethernet (FCoE) to provide data center I/O consolidation.

Currently, the Cisco Nexus 5000 Series has two switches: the Cisco Nexus 5010 switch which provides 20 fixed Ethernet ports in a 1 RU switch, and the Cisco Nexus 5020 switch which provides 40 fixed Ethernet ports in a 2 RU switch. Optional expansion modules provide Fibre Channel ports and additional Ethernet ports.

New Technologies in the Cisco Nexus 5000 Series

Fibre Channel over Ethernet

Fibre Channel over Ethernet (FCoE) allows Fibre Channel traffic to be encapsulated over a physical Ethernet link. FCoE frames use a unique EtherType so that FCoE traffic and standard Ethernet traffic can be carried on the same link.

Classic Ethernet is a best-effort protocol; in the event of congestion, Ethernet will discard packets, relying on higher level protocols to provide retransmission and other reliability mechanisms. Fibre Channel traffic requires a lossless transport layer; as a data storage protocol, it is unacceptable to lose a single data packet. Native Fibre Channel implements a lossless service at the transport layer using a buffer-to-buffer credit system.

For FCoE traffic, the Ethernet link must provide a lossless service. Ethernet links on Cisco Nexus 5000 Series switches provide two mechanisms to ensure lossless transport for FCoE traffic: link-level flow control and priority flow control.

IEEE 802.3x link-level flow control allows a congested receiver to signal the far end to pause the data transmission for a short period of time. The pause functionality is applied to all the traffic on the link.

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).

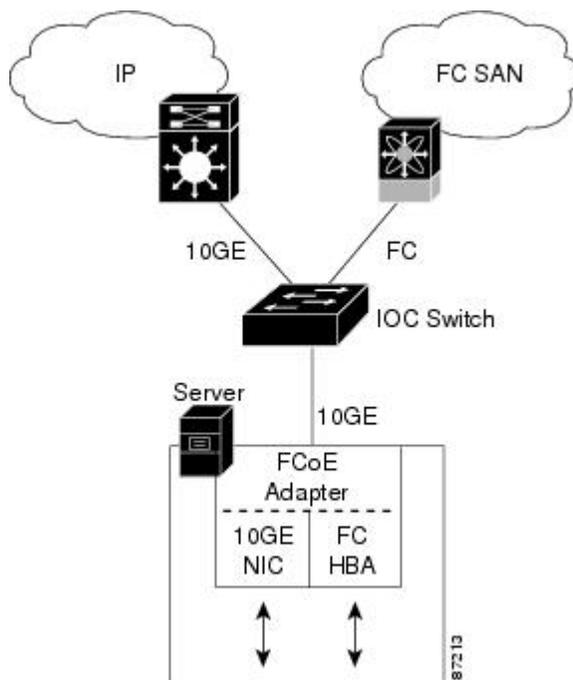
Data Center I/O Consolidation

I/O consolidation allows a single network technology to carry IP, SAN, and IPC traffic. FCoE is the single network technology that allows I/O consolidation. 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.

Cisco Nexus 5000 Series switches use FCoE to carry Fibre Channel and Ethernet traffic on the same physical Ethernet connection between the switch and the server. At the server, the connection terminates to a converged network adapter (CNA). The adapter presents two interfaces to the server's operating system (OS): one Ethernet NIC interface and one Fibre Channel host bus adapter (HBA) interface.

The server OS is not aware of the FCoE encapsulation (see the following figure). At the switch, the incoming Ethernet port separates the Ethernet and Fibre Channel traffic (using EtherType to differentiate the frames). Ethernet frames and Fibre Channel frames are switched to their respective network-side interfaces.

Figure 1: I/O Consolidation



Cisco Nexus 5000 Series switches provide quality of service (QoS) capabilities to ensure lossless or best-effort service across the switch. For Fibre Channel traffic (FCoE) you should apply the lossless QoS classes. By default, best-effort service is applied to all of the Ethernet traffic. You can configure different QoS levels for specific classes of Ethernet traffic.

Virtual Interfaces

When FCoE is enabled, a physical Ethernet cable carries traffic for a logical Fibre Channel connection.

The Cisco Nexus 5000 Series switch uses virtual interfaces to represent the logical Fibre Channel connections. For configuration purposes, virtual Fibre Channel interfaces are implemented as Layer 2 subinterfaces of the physical Ethernet interface.

Ethernet features (such as the link debounce timer and VLAN membership) are configured on the physical Ethernet interface. Logical Fibre Channel features (such as VSAN membership) are configured on the virtual Fibre Channel interfaces.

Cisco Nexus 5000 Series Switch Software

Ethernet Switching

Cisco Nexus 5000 Series switches are Layer 2 devices, which run Cisco NX-OS.

Cisco Nexus 5000 Series switches are designed to support high-density, high-performance Ethernet systems and provide the following Ethernet switching features:

- IEEE 802.1D-2004 Rapid and Multiple Spanning Tree Protocols (802.1w and 802.1s)
- IEEE 802.1Q VLANs and trunks
- IEEE 802.3ad link aggregation
- Private VLANs
- EtherChannels and virtual port channels (vPCs)
- Traffic suppression (unicast, multicast, and broadcast)

FCoE and Fibre Channel Switching

Cisco Nexus 5000 Series switches support data center I/O consolidation by providing FCoE interfaces (to the servers) and native Fibre Channel interfaces (to the SAN).

FCoE and Fibre Channel switching includes the following features:

- Cisco fabric services
- N-port virtualization
- VSANs and VSAN trunking
- Zoning

- Distributed device alias service
- SAN port channels

QoS

Cisco Nexus 5000 Series switches provide quality of service (QoS) capabilities such as traffic prioritization and bandwidth allocation on egress interfaces.

The default QoS configuration on the switch provides lossless service for Fibre Channel and FCoE traffic. QoS can be configured to provide additional classes of service for Ethernet traffic.

Virtual Port Channels

A virtual port channel (vPC) allows links that are physically connected to two different Cisco Nexus 5000 Series switches or Cisco Nexus 2000 Series Fabric Extenders to appear as a single port channel. A vPC can provide multipathing, which allows you to create redundancy by enabling multiple parallel paths between nodes and load balancing traffic where alternative paths exist.

Serviceability

The Cisco Nexus 5000 Series switch serviceability functions provide data for network planning and help to improve problem resolution time.

Switched Port Analyzer

The switched port analyzer (SPAN) feature allows an administrator to analyze all traffic between ports by nonintrusively directing the SPAN session traffic to a SPAN destination port that has an external analyzer attached to it.

Ethalyzer

Ethalyzer is a Cisco NX-OS protocol analyzer tool based on the Wireshark (formerly Ethereal) open source code. Ethalyzer is a command-line version of Wireshark for capturing and decoding packets. You can use Ethalyzer to troubleshoot your network and analyze the control-plane traffic.

Call Home

The Call Home feature continuously monitors hardware and software components to provide e-mail-based notification of critical system events. A versatile range of message formats is available for optimal compatibility with pager services, standard e-mail, and XML-based automated parsing applications. The feature offers alert grouping capabilities and customizable destination profiles. This feature can be used, for example, to directly page a network support engineer, send an e-mail message to a network operations center (NOC), and employ Cisco AutoNotify services to directly generate a case with the Cisco Technical Assistance Center (TAC). This feature is a step toward autonomous system operation, which enables networking devices to inform IT when a problem occurs and helps to ensure that the problem is resolved quickly.

Online Diagnostics

Cisco generic online diagnostics (GOLD) is a suite of diagnostic facilities to verify that hardware and internal data paths are operating as designed. Boot-time diagnostics, continuous monitoring, and on-demand and scheduled tests are part of the Cisco GOLD feature set. GOLD allows rapid fault isolation and continuous system monitoring.

Switch Management

Simple Network Management Protocol

Cisco NX-OS is compliant with Simple Network Management Protocol (SNMP) version 1, version 2, and version 3. A full set of Management Information Bases (MIBs) is supported.

Role-Based Access Control

With role-based access control (RBAC), you can limit access to switch operations by assigning roles to users. Administrators can customize access and restrict it to the users who require it.

Configuration Methods

Configuring with CLI, XML Management Interface, or SNMP

You can configure Cisco Nexus 5000 Series switches using the command-line interface (CLI), the XML management interface over SSH, or SNMP as follows:

- CLI—You can configure switches using the CLI from an SSH session, a Telnet session, or the console port. SSH provides a secure connection to the device.
- XML Management Interface over SSH—You can configure switches using the XML management interface, which is a programming interface based on the NETCONF protocol that complements the CLI functionality. For more information, see the *Cisco NX-OS XML Interfaces User Guide*.
- SNMP—SNMP allows you to configure switches using Management Information Bases (MIBs).

Configuring with Cisco Data Center Network Manager

You can configure Cisco Nexus 5000 Series switches using the Data Center Network Manager (DCNM) client, which runs on a local PC and uses the DCNM server.

For more information, see the *Cisco DCNM Configuration Guides*.

Configuring with Cisco MDS Fabric Manager

You can configure Cisco Nexus 5000 Series switches using the Fabric Manager client, which runs on a local PC and uses the Fabric Manager server.

For more information, see the *Cisco MDS 9000 and Nexus 5000 Series Fabric Manager Software Configuration Guide* for Cisco Fabric Manager Release 4.1.

Network Security Features

Cisco NX-OS Release 4.1 includes the following security features:

- Authentication, authorization, and accounting (AAA) and TACACS+
- RADIUS
- Secure Shell (SSH) Protocol Version 2
- Simple Network Management Protocol Version 3 (SNMPv3)
- MAC ACLs and IP ACLs, including port-based ACLs (PACLs) and VLAN-based ACLs (VACLs).

Virtual Device Contexts

Cisco NX-OS can segment operating system and hardware resources into virtual device contexts (VDC) that emulate virtual devices. The Cisco Nexus 5000 Series switch does not support multiple VDCs. All switch resources are managed in the default VDC.

For more information, see the *Cisco Nexus 7000 Series NX-OS Getting Started with Virtual Device Contexts*, Release 4.1.

Licensing

The Cisco Nexus 5000 Series switch is shipped with its licenses installed. The switch provides commands to manage the licenses and install additional licenses.

Typical Deployment Topologies

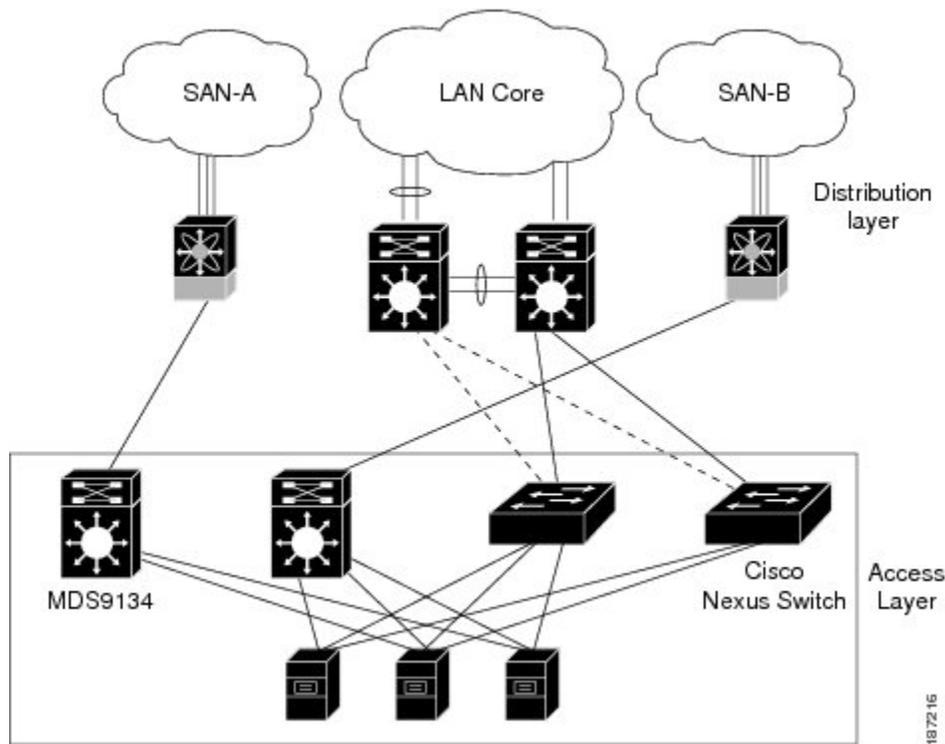
Ethernet TOR Switch Topology

The Cisco Nexus 5000 Series switch can be deployed as a 10-Gigabit Ethernet top-of-rack (TOR) switch, with uplinks to the data center LAN distribution layer switches. An example configuration is shown in the following figure.

In this example, the blade server rack incorporates blade switches that support 10-Gigabit Ethernet uplinks to the Cisco Nexus 5000 Series switch. The blade switches do not support FCoE, so there is no FCoE traffic and no Fibre Channel ports on the Cisco Nexus 5000 Series switch.

In the example configuration, the Cisco Nexus 5000 Series switch has Ethernet uplinks to two Catalyst switches. If STP is enabled in the data center LAN, the links to one of the switches will be STP active and the links to the other switch will be STP blocked.

Figure 2: Ethernet TOR Switch Topology



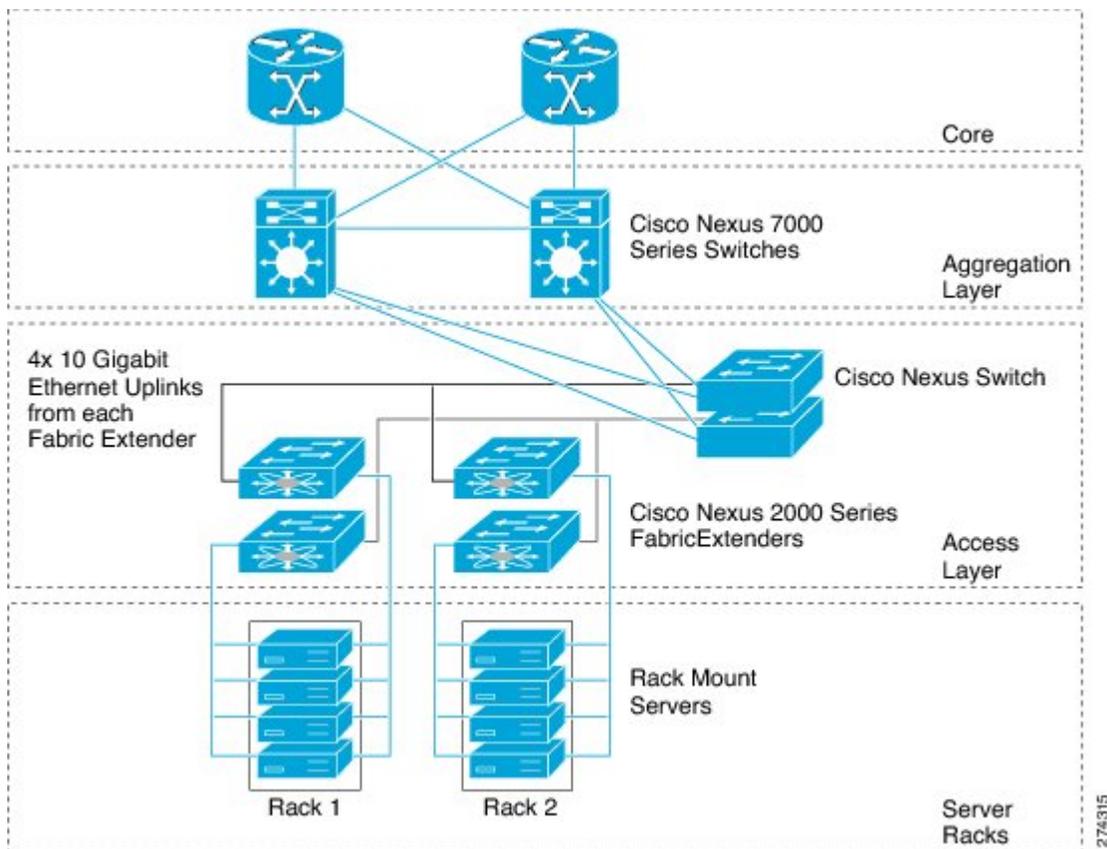
All of the server-side ports on the Cisco Nexus 5000 Series switch are running standard Ethernet. FCoE is not required, so the server ports are connected using 10-Gigabit Ethernet NICs.

The servers are connected to the data center SAN through MDS 9134 SAN switches. The server Fibre Channel ports require standard Fibre Channel HBAs.

Fabric Extender Deployment Topology

The following figure shows a simplified configuration using the Cisco Nexus 2000 Series Fabric Extender in combination with the Cisco Nexus 5000 Series switch to provide a simplified and cost-effective 1-Gigabit TOR solution.

Figure 3: Fabric Extender Deployment Topology



In the example configuration, the Fabric Extender top-of-rack units provide 1-Gigabit host interfaces connected to the servers. The Fabric Extender units are attached to their parent Cisco Nexus 5000 Series switches with 10-Gigabit fabric interfaces.

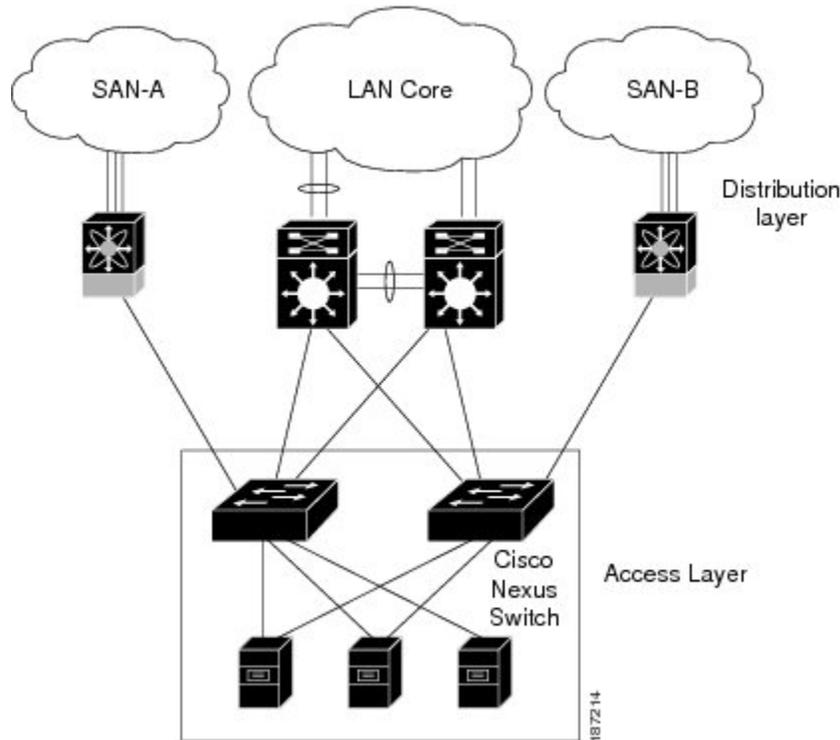
Each Fabric Extender acts as a Remote I/O Module on the parent Cisco Nexus 5000 Series switch. All device configurations are managed on the Cisco Nexus 5000 Series switch and configuration information is downloaded using inband communication to the Fabric Extender.

See the *Cisco Nexus 2000 Series Fabric Extender Software Configuration Guide* for an overview of the Fabric Extender and configuration details.

Data Center I/O Consolidation Topology

The following figure shows a typical I/O consolidation scenario for the Cisco Nexus 5000 Series switch.

Figure 4: I/O Consolidation Topology



The Cisco Nexus 5000 Series switch connects to the server ports using FCoE. Ports on the server require converged network adapters. For redundancy, each server connects to both switches. Dual-port CNA adapters can be used for this purpose. The CNA is configured in active-passive mode, and the server needs to support server-based failover.

On the Cisco Nexus 5000 Series switch, the Ethernet network-facing ports are connected to two Catalyst 6500 Series switches. Depending on required uplink traffic volume, there may be multiple ports connected to each Catalyst 6500 Series switch, configured as port channels. If STP is enabled in the data center LAN, the links to one of the switches will be STP active and the links to the other switch will be STP blocked.

The SAN network-facing ports on the Cisco Nexus 5000 Series switch are connected to Cisco MDS 9000 Family switches. Depending on the required traffic volume, there may be multiple Fibre Channel ports connected to each MDS 9000 Family switch, configured as SAN port channels.

Supported Standards

The following table lists the standards supported by the Cisco Nexus 5000 Series switches.

Table 2: IEEE Compliance

Standard	Description
802.1D	MAC Bridges
802.1s	Multiple Spanning Tree Protocol
802.1w	Rapid Spanning Tree Protocol
802.3ad	Link aggregation with LACP
802.3ae	10-Gigabit Ethernet
802.1Q	VLAN Tagging
802.1p	Class of Service Tagging for Ethernet frames



PART **I**

Configuration Fundamentals

- [Using the Command-Line Interface, page 15](#)
- [Initial Switch Configuration, page 31](#)
- [Managing Licenses, page 47](#)



Using the Command-Line Interface

This chapter describes how to use the command-line interface of the Cisco Nexus 5000 Series switch. It contains the following sections:

- [Accessing the Command Line Interface, page 15](#)
- [Using the CLI, page 17](#)
- [Using Commands, page 23](#)
- [Using CLI Variables, page 25](#)
- [Using Command Aliases , page 27](#)
- [Defining Command Aliases , page 27](#)
- [Command Scripts, page 28](#)

Accessing the Command Line Interface

You can connect to the switch using a terminal plugged into the console port.

You can also connect to the switch with Telnet or SSH. The switch supports up to eight simultaneous Telnet and SSH connections. To connect with Telnet or SSH, you need to know the hostname or IP address of the switch.

Telnet Connection

You can make a Telnet connection to a Cisco Nexus 5000 Series switch.

Before You Begin

Correctly set the console port parameters.

SUMMARY STEPS

1. Make a Telnet connection from your host to the switch that you want to access.
2. At the switch login prompt, enter your username and password.
3. Exit the session when finished.

DETAILED STEPS

Step 1 Make a Telnet connection from your host to the switch that you want to access.

```
host$ telnet {hostname | ip-addr}
```

Step 2 At the switch login prompt, enter your username and password.
The Cisco Nexus 5000 Series switch initiates authentication.

Note If no password has been configured, press
Return.

Step 3 Exit the session when finished.

```
switch# exit
```

This example shows how to make a Telnet connection to a switch:

```
host$ telnet 10.0.13.42
Trying 10.0.13.42...

Connected to 10.0.13.42
Escape character is '^]'.
switch Login: admin
Password: password
...
switch# exit
```

SSH Connection

You can make an SSH connection to a Cisco Nexus 5000 Series switch.

Before You Begin

Correctly set the console port parameters.

SUMMARY STEPS

1. Make an SSH connection from your host to the switch that you want to access.
2. At the switch login prompt, enter your username and password.
3. Exit the session when finished.

DETAILED STEPS

Step 1 Make an SSH connection from your host to the switch that you want to access.

```
host$ ssh [-D port] [{user@} hostname | ip_addr]
```

Step 2 At the switch login prompt, enter your username and password.
The Cisco Nexus 5000 Series switch initiates authentication.

Note If no password has been configured, press
Return.

- Step 3** Exit the session when finished.
switch# **exit**

This example shows how to make an SSH connection to a switch:

```
host$ ssh 10.0.13.42
The authenticity of host '(10.0.13.42)' can't be established.
RSA key fingerprint is 53:b4:ad:c8:51:17:99:4b:c9:08:ac:c1:b6:05:71:9b.
Are you sure you want to continue connecting (yes/no)? yes
Warning: Permanently added '10.0.13.42' (RSA) to the list of known hosts.
switch Login: admin
Password: password
...
switch# exit
```

Using the CLI

Using CLI Command Modes

Switches in the Cisco Nexus 5000 Series have two main command modes: user EXEC mode and configuration mode. The commands available to you depend on the mode you are in. To obtain a list of available commands in either mode, type a question mark (?) at the system prompt.

The following table lists and describes the two commonly used modes, how to enter the modes, and the resulting system prompts. The system prompt helps you identify which mode you are in and the commands that are available to you in that mode.

Table 3: Frequently Used Switch Command Modes

Mode	Description	How to Access	Prompt
EXEC	Enables you to temporarily change terminal settings, perform basic tests, and display system information. Note Changes made in this mode are generally not saved across system resets.	At the switch prompt, enter the required EXEC mode command.	switch#
Configuration mode	Enables you to configure features that affect the system as a whole. Note Changes made in this mode are saved across system resets if you save your configuration.	From EXEC mode, enter the configure terminal command.	switch(config)#

You can abbreviate commands and keywords by entering just enough characters to make the command unique from other commands. For example, you can abbreviate the **configure terminal** command to **conf t**.

Changing Command Modes

Configuration mode, also known as terminal configuration mode, has several submodes. Each of these submodes places you further down in the prompt hierarchy. When you type **exit**, the switch backs out of the current level and returns you to the previous level. When you type **end**, the switch backs out to the user EXEC level. You can also press **Ctrl-Z** in configuration mode as an alternative to typing **end**.

Listing the Commands Used with Each Command Mode

You can display the commands available in any command mode by typing a question mark (?) at the switch prompt.

CLI Command Hierarchy

CLI commands are organized hierarchically, with commands that perform similar functions grouped under the same level. For example, all commands that display information about the system, configuration, or hardware are grouped under the **show** command, and all commands that allow you to configure the switch are grouped under the **configure terminal** command.

To execute a command, you enter the command by starting at the top level of the hierarchy. For example, to configure an interface, use the **config terminal** command. Once you are in configuration mode, enter the **interface** command. When you are in the interface submenu, you can query the available commands.

The following example shows how to query the available command in the interface submenu:

```
switch# configure terminal
switch(config)# interface ethernet 1/1
switch(config-if)# ?

bandwidth          Set bandwidth informational parameter
cdp                Configure CDP interface parameters
channel-group      Add to/remove from a port-channel
delay              Specify interface throughput delay
description        Enter description of maximum 80 characters
exit               Exit from command interpreter
fcoe               Fibre channel over ethernet configuration
fex                Configure FEX fabric
flowcontrol        Configure interface flowcontrol
ip                 Configure IP features
ipv6               Configure IPv6 features
lacp               Configure LACP parameters
link               Configure link
lldp               Configure Interface LLDP parameters
logging            Configure logging for interface
mac                MAC configuration commands
no                 Negate a command or set its defaults
priority-flow-control
                  Configure interface priority-flowcontrol
service-policy     Configure QoS service policy
shutdown           Enable/disable an interface
snmp               Modify SNMP interface parameters
spanning-tree      Spanning Tree Subsystem
speed              Enter the port speed
storm-control       Configure Interface storm control
switchport         Configure switchport parameters
untagged           default to use for untagged packets on interface
```

EXEC Mode Commands

When you start a session on the switch, you begin in EXEC mode. From EXEC mode, you can enter configuration mode. Most of the EXEC commands are one-time commands, such as **show** commands, which display the current configuration status.

The following commands are available in EXEC mode:

```

switch# ?
  attach          Connect to a specific linecard
  callhome        callhome commands
  cd              Change current directory
  check           run consistency check on external storage device
  clear           Reset functions
  cli             CLI commands
  clock           Manage the system clock
  configure       Enter configuration mode
  copy            Copy from one file to another
  debug           Debugging functions
  debug-filter    Enable filtering for debugging functions
  delete          delete a file
  dir             list files in a directory
  discover        discover information
  echo            echo argument back to screen (usefull for run script)
  end             Exit configuration mode
  ethanalyzer     Configure cisco fabric analyzer
  exit            Exit from command interpreter
  fcping         Ping an N-Port
  fctrace        Trace the route for an N-Port.
  fex            FEX control commands
  find            Find a file below the current directory
  format         Format disks
  gunzip          Uncompresses LZ77 coded files
  gzip           Compresses file using LZ77 coding
  install        upgrade software
  license         Enter the license configuration mode
  mkdir          Create new directory
  move           Move files
  no             Negate a command or set its defaults
  ntp            Execute NTP commands
  ping           Test network reachability
  ping6          Test IPv6 network reachability
  purge          Deletes unused data
  pwd            View current directory
  reload         Reboot the entire box
  rmdir          Delete a directory
  routing-context Set the routing context
  run-script     Run shell scripts
  san-port-channel Port-Channel related commands
  send           Send message to open sessions
  session        Configure session preferences
  setup         Run the basic SETUP command facility
  show           Show running system information
  sleep          Sleep for the specified number of seconds
  ssh            SSH to another system
  ssh6           SSH to another system
  system         System management commands
  tac-pac        save tac information to a specific location
  tail           Display the last part of a file
  telnet         Telnet to another system
  telnet6        Telnet6 to another system
  terminal        Set terminal line parameters
  terminate      Terminates a config session
  test           test command
  traceroute     Traceroute to destination
  traceroute6    Traceroute6 to destination
  undebg         Disable Debugging functions (See also debug)
  unmount        unmount compact flash disk or usb drive
  update         Update license
  where          shows the cli context you are in
  write          Write current configuration
  xml            xml agent
  zone           Execute Zone Server commands
  zoneset        Execute zoneset commands

```

Configuration Mode Commands

Configuration mode allows you to make changes to the existing configuration. When you save the configuration, these commands are saved across switch reboots. Once you are in configuration mode, you can enter interface configuration mode, zone configuration mode, and a variety of protocol-specific modes. Configuration mode is the starting point for all configuration commands.

The following commands are available in configuration mode:

```

switch# configure terminal
switch(config)# ?
  aaa                Configure aaa functions
  banner             Configure banner message
  boot               Configure boot variables
  callhome           Enter the callhome configuration mode
  cdp                Configure CDP parameters
  cfs                CFS configuration commands
  class-map          Configure class-map
  cli                Configure CLI aliases
  clock              Configure time-of-day clock
  device-alias       Device-alias configuration commands
  diagnostic          Diagnostic commands
  end                Exit configuration mode
  exit               Exit from command interpreter
  fabric-binding     Fabric Binding configuration
  fcalias            Fcalias configuration commands
  fcdomain           Enter the fcdomain configuration mode
  fcdroplateness     configure switch or network latency
  fcflow             Configure fcflow
  fcid-allocation    Add/remove company id(or OUIs) from auto area list
  fcinterop          Interop commands
  fcns               name server configuration
  fcroute            Configure FC routes
  fcs                Configure Fabric Config Server
  fcsp               Config commands for FC-SP
  fctimer            configure fibre channel timers
  fdmi               config commands for FDMI
  feature            Command to enable/disable features
  fex                FEX configuration
  fspf               Configure fspf
  hostname            Configure system's host name
  hw-module          Enable/Disable OBFL information
  in-order-guarantee set in-order delivery guarantee
  interface           Configure interfaces
  ip                 Configure IP features
  ipv6               Configure IPv6 features
  lacp                Configure LACP parameters
  license             Modify license features
  line                Configure a terminal line
  lldp                Configure global LLDP parameters
  logging             Modify message logging facilities
  mac                MAC configuration commands
  mac-address-table  MAC Address Table
  monitor            Ethernet SPAN
  no                  Negate a command or set its defaults
  npiv                Nx port Id Virtualization (NPIV) feature enable
  npv                 Config commands for FC N_port Virtualizer
  ntp                 NTP Configuration
  policy-map          Configure policy-map
  port-channel        Configure port channel parameters
  port-security       Configure Port Security
  port-track          Configure Switch port track config
  privilege           Command privilege parameters
  radius-server       Configure RADIUS related parameters
  resequence          Resequenece a list with sequence numbers
  rib                 Configure RIB parameters
  rlir                config commands for RLIR
  rmon                Remote Monitoring
  role                Configure roles
  rscn                config commands for RSCN
  scsi-target         scsi-target configuration
  show                Show running system information
  snmp-server         Configure snmp server
  spanning-tree       Spanning Tree Subsystem
  ssh                 Configure SSH parameters
  switchname          Configure system's host name
  system              system config command
  system              System management commands
  tacacs+             Enable tacacs+
  telnet              Enable telnet

```

track	Object tracking configuration commands
trunk	Configure Switch wide trunk protocol
username	Configure user information.
vlan	Vlan commands
vrf	Configure VRF parameters
vsan	Enter the vsan configuration mode
wwn	Set secondary base MAC addr and range for additional WWNs
xml	xml agent
zone	Zone configuration commands
zoneset	Zoneset configuration commands

Using Commands

You can configure the CLI to function in two ways: configure it interactively by entering commands at the CLI prompt or create an ASCII file containing switch configuration information (use the CLI to edit and activate the file).

Listing Commands and Syntax

In any command mode, you can obtain a list of available commands by entering a question mark (?).

```
switch# ?
```

To see a list of commands that begin with a particular character sequence, type those characters followed by a question mark (?). Do not include a space before the question mark.

```
switch# co?
configure copy
```

To list keywords or arguments, enter a question mark in place of a keyword or argument. Include a space before the question mark. This form of help is called command syntax help because it reminds you which keywords or arguments are applicable based on the commands, keywords, and arguments you have already entered.

```
switch# # configure ?
<CR>
terminal Configure the system from terminal input
```

If you are having trouble entering a command, check the system prompt and enter the question mark (?) for a list of available commands. You might be in the wrong command mode or using incorrect syntax.

Entering Command Sequences

In any command mode, you can begin a particular command sequence, and then immediately press the **Tab** key to complete the rest of the command.

```
switch (config)# ro<Tab>
switch (config)# role <Tab>
switch (config)# role name
```

This form of help is called command completion because it completes a word for you. If several options are available for the typed letters, all options that match those letters are displayed.

Undoing or Reverting to Default Values or Conditions

You can enter the **no** form of any command to perform the following actions:

- Undo an incorrectly entered command.

If you enter the **zone member** command, you can undo the results:

```
switch(config)# zone name test vsan 1
switch(config-zone)# member pwn 12:12:12:12:12:12:12:12
switch(config-zone)# no member pwn 12:12:12:12:12:12:12:12
WARNING: Zone is empty. Deleting zone test. Exit the submode.
switch(config-zone)#
```

- Delete a created facility.

If you want to delete a zone that you created, enter the following commands:

```
switch(config)# zone name test vsan 1
switch(config-zone)# exit
switch(config)# no zone name test vsan 1
switch(config)#
```

You cannot delete a zone facility called test while still in zone configuration submode. You must first exit the zone submode and return to configuration mode.

- Revert to the default value.

If you enter the **zone merge-control restrict vsan** command, you can undo the results:

```
switch(config)# zone merge-control restrict vsan 10
switch(config)# no zone merge-control restrict vsan 10
switch(config)#
```

Using Keyboard Shortcuts

You can execute an EXEC mode command from a configuration mode or submode prompt. You can enter this command from any submode within the configuration mode. The command is executed at the EXEC level, and the prompt resumes its current mode level, as in the following example:

```
switch(config)# terminal session-timeout 0
```

In this example, **terminal session-timeout** is an EXEC mode command.

The following table lists some useful command keys that can be used in both EXEC and configuration modes.

Table 4: Useful Command Keys

Command	Description
Ctrl-P	Up history
Ctrl-N	Down history
Ctrl-X-H	List history
Alt-P	History search backwards Note The difference between Tab completion and Alt-P or Alt-N is that pressing Tab completes the current word, while Alt-P and Alt-N completes a previously entered command.

Command	Description
Alt-N	History search forwards
Ctrl-G	Exit
Ctrl-Z	End
Ctrl-L	Clear session

The following table describes the commonly used configuration submodes.

Table 5: Common Configuration Submodes

Submode Name	From Configuration Mode, Enter:	Submode Prompt
Call home	callhome	switch(config-callhome)#
FCS Registration	fcs register	switch(config-fcs-register)#
	From FCS registration submode: platform name name vsan vsan-id	switch(config-fcs-register-attrib)#
Fibre Channel alias	fcalias name name vsan vsan-id	switch(config-fcalias)#
FSPF	fspf config vsan vsan-id	switch(config-(fspf-config))#
Interface configuration	interface type slot/port	switch(config-if)#
Line console	line console	switch(config-console)
Virtual terminal line	line vty	switch(config-line)#
Role	role name	switch(config-role)#
VLAN	vlan	switch(config-vlan)#
VSAN database	vsan database	switch(config-vsan-db)#
Zone	zone name string vsan vsan-id	switch(config-zone)#
Zone set	zoneset name name vsan vsan-id	switch(config-zoneset)#

Using CLI Variables

The Cisco Nexus 5000 Series CLI parser supports the definition and use of variables in CLI commands. CLI variables can be used as follows:

- Entered directly on the command line.
- Passed to the child script initiated using the **run-script** command.
The variables defined in the parent shell are available for use in the child **run-script** command process.
- Passed as command line arguments to the **run-script** command.

CLI variables have the following characteristics:

- You cannot reference a variable through another variable using nested references.
- You can define persistent variables that are available across switch reloads.
- You can reference only one predefined system variable, which is the **TIMESTAMP** variable.

User-Defined Persistent CLI Variables

You can define CLI session variables to persist only for the duration of your CLI session using the **cli var name** command in EXEC mode. CLI session variables are useful for scripts that you execute periodically.

The following example shows how to create a user-defined CLI session variable:

```
switch# cli var name testinterface fc 1/1
```

You can reference a variable using the syntax **\$(variable)**. The following example shows how to reference a user-defined CLI session variable:

```
switch# show interface $(testinterface)
fc2/1 is up
Hardware is Fibre Channel, SFP is short wave laser w/o OFC (SN)
  Port WWN is 20:01:00:0d:ec:0e:1d:00
  Admin port mode is auto, trunk mode is on
  snmp traps are enabled
  Port mode is F, FCID is 0x01000b
  Port vsan is 1
  Speed is 2 Gbps
  Transmit B2B Credit is 7
  Receive B2B Credit is 16
  Receive data field Size is 2112
  Beacon is turned off
  5 minutes input rate 256 bits/sec, 32 bytes/sec, 1 frames/sec
  5 minutes output rate 256 bits/sec, 32 bytes/sec, 1 frames/sec
    232692 frames input, 7447280 bytes
      0 discards, 0 errors
      0 CRC, 0 unknown class
      0 too long, 0 too short
    232691 frames output, 7448692 bytes
      0 discards, 0 errors
      0 input OLS, 0 LRR, 0 NOS, 0 loop inits
      1 output OLS, 1 LRR, 0 NOS, 1 loop inits
      16 receive B2B credit remaining
      7 transmit B2B credit remaining
```

Use the **show cli variables** command to display user-defined CLI session variables. The following example displays user-defined CLI session variables:

```
switch# show cli variables
VSH Variable List
-----
TIMESTAMP="2005-10-24-21.29.33"
testinterface="fc 1/1"
```

Use the **cli no var name** command to remove user-defined CLI session variables. The following example removes a user-defined CLI session variable:

```
switch# cli no var name testinterface
```

Using Command Aliases

Command alias support has the following characteristics:

- Command aliases are global for all user sessions.
- Command aliases are saved across reboots.
- Commands being aliased must be typed in full without abbreviation.
- Command alias translation always takes precedence over any keyword in any configuration mode or submode.
- Command alias support is only available on the supervisor module, not the switching modules.
- Command alias configuration takes effect for other user sessions immediately.
- You cannot override the default command alias **alias**, which aliases the **show cli alias** command.
- Nesting of command aliases is permitted to a maximum depth of 1. One command alias can refer to another command alias that must refer to a valid command, not to another command alias.
- A command alias always replaces the first command keyword on the command line.
- You can define command aliases for commands in any configuration submode or the EXEC mode.

Defining Command Aliases

You can define command aliases using the **cli alias name** command in configuration mode.

This following example shows how to define command aliases:

```
switch# configure terminal
switch(config)# cli alias name eth interface ethernet
switch(config)# cli alias name shintbr show interface brief
switch(config)# cli alias name shfcintup shintbr | include up | include fc
```

You can display the command aliases defined on the switch using the **alias** default command alias.

The following example shows how to display the command aliases defined on the switch:

```
switch# alias
CLI alias commands
=====
alias      :show cli alias
gigint     :interface gigabitethernet
shintbr    :show interface brief
shfcintup :shintbr | include up | include fc
```

Command Scripts

Executing Commands Specified in a Script

The **run-script** command executes the commands specified in a file. To use this command, be sure to create the file and specify commands in the required order.



Note

You cannot create the script file at the switch prompt. You can create the script file on an external machine and copy it to the bootflash: directory. This section assumes that the script file resides in the bootflash: directory.

The syntax for this command is **run-script filename**.

This example displays the CLI commands specified in a test file that resides in the bootflash: directory.

```
switch# show file bootflash:testfile
configure terminal
interface fc 3/1
no shutdown
end
show interface fc 3/1
```

This file output is in response to the **run-script** command executing the contents in the test file:

```
switch# run-script bootflash:testfile
'configure terminal'
Enter configuration commands, one per line. End with CNTL/Z.
'interface fc 3/1'
'no shutdown'
'end'
'show interface fc 3/1'
fc3/1 is trunking
  Hardware is Fibre Channel, SFP is short wave laser w/o OFC (SN)
  Port WWN is 20:81:00:0d:ec:6b:cd:c0
  Peer port WWN is 20:01:00:0d:ec:0d:d0:00
  Admin port mode is auto, trunk mode is on
  snmp link state traps are enabled
  Port mode is TE
  Port vsan is 1
  Speed is 2 Gbps
  Transmit B2B Credit is 255
  Receive B2B Credit is 16
  Receive data field Size is 2112
  Beacon is turned off
  Trunk vsans (admin allowed and active) (1)
  Trunk vsans (up) (1)
  Trunk vsans (isolated) ()
  Trunk vsans (initializing) ()
  5 minutes input rate 96 bits/sec, 12 bytes/sec, 0 frames/sec
  5 minutes output rate 64 bits/sec, 8 bytes/sec, 0 frames/sec
    77423 frames input, 6708868 bytes
      0 discards, 0 errors
      0 CRC, 0 unknown class
      0 too long, 0 too short
    77302 frames output, 4184976 bytes
      0 discards, 0 errors
  1 input OLS, 2 LRR, 0 NOS, 0 loop inits
  1 output OLS, 0 LRR, 1 NOS, 0 loop inits
  16 receive B2B credit remaining
  255 transmit B2B credit remaining
```

Using CLI Variables in Scripts

You can use CLI variables defined by the **cli var** command or passed as arguments in the **run-script** command.

The following example shows how to use CLI session variables in a script file used by the **run-script** command:

```
switch# cli var name testinterface fc 1/1
switch# show file bootflash:test1.vsh
show interface $(testvar)
switch# run-script bootflash:test1.vsh
`show interface $(testvar)`
fc2/1 is down (SFP not present)
Hardware is Fibre Channel
Port WWN is 20:01:00:05:30:00:8e:1e
Admin port mode is auto, trunk mode is on
Port vsan is 1
Receive data field Size is 2112
Beacon is turned off
5 minutes input rate 0 bits/sec, 0 bytes/sec, 0 frames/sec
5 minutes output rate 0 bits/sec, 0 bytes/sec, 0 frames/sec
1 frames input, 128 bytes
0 discards, 0 errors
0 CRC, 0 unknown class
0 too long, 0 too short
1 frames output, 128 bytes
0 discards, 0 errors
0 input OLS, 0 LRR, 0 NOS, 0 loop inits
0 output OLS, 0 LRR, 0 NOS, 0 loop inits
0 receive B2B credit remaining
0 transmit B2B credit remaining
```

The following example shows how you can pass a CLI session variable as arguments to a child **run-script** command process:

```
switch# show file bootflash:test1.vsh
show interface $(var1) $(var2)
switch# run bootflash:test2.vsh var1="fc2/1" var2="brief"
`show interface $(var1) $(var2)`
-----
Interface  Vsan    Admin  Admin  Status          SFP    Oper  Oper  Port
          Mode    Mode   Trunk                               Mode  Speed Channel
                               Mode                                     (Gbps)
-----
fc2/1 1 auto on sfpAbsent -- -- -- \
```

Related Topics

[Using Command Aliases , on page 27](#)

Setting the Delay Time

The **sleep** command delays an action by a specified number of seconds.

The syntax for this command is **sleep seconds**.

```
switch# sleep 30
```

You will see the switch prompt return after 30 seconds. This command is useful within scripts. For example, if you create a command script called test-script and then you execute the script, the switch software executes

the **discover scsi-target remote** command, and then waits for 10 seconds before executing the **show scsi-target disk** command.

```
switch# show file bootflash:test-script
discover scsi-target remote
sleep 10
show scsi-target disk
switch# run-script bootflash:test-script
```



Initial Switch Configuration

This chapter describes the command-line interface (CLI) and CLI command modes of Cisco Nexus 5000 Series switches. It includes the following sections:

- [Configuring the Switch, page 31](#)

Configuring the Switch

Image Files on the Switch

The Cisco Nexus 5000 Series switches have the following images:

- BIOS and loader images combined in one file
- Kickstart image
- System image that includes a BIOS image that can be upgraded

The switch has flash memory that consists of two separate flash parts:

- A 2 MB flash part holds two BIOS and loader images.
- A 1 GB flash part holds configuration files, kickstart images, systems images, and other files.

The upgradeable BIOS and the golden BIOS are programmed onto the 2 MB flash part. You cannot upgrade the golden BIOS.

When you download a new pair of kickstart and system images, you also get a new BIOS image because it is included in the system image. You can use the **install all** command to upgrade the kickstart, system, and upgradeable BIOS images.

Starting the Switch

A Cisco Nexus 5000 Series switch starts its boot process as soon as its power cord is connected to an A/C source. The switch does not have a power switch.

Boot Sequence

When the switch boots, the golden BIOS validates the checksum of the upgradeable BIOS. If the checksum is valid, then control is transferred to the upgradeable BIOS image. The upgradeable BIOS launches the kickstart image, which then launches the system image. If the checksum of the upgradeable BIOS is not valid, then the golden BIOS launches the kickstart image, which then launches the system image.

You can force the switch to bypass the upgradeable BIOS and use the golden BIOS instead. If you press **Ctrl-Shift-6** within two seconds of when power is supplied to the switch, the golden BIOS will be used to launch the kickstart image, even if the checksum of the upgradeable BIOS is valid.

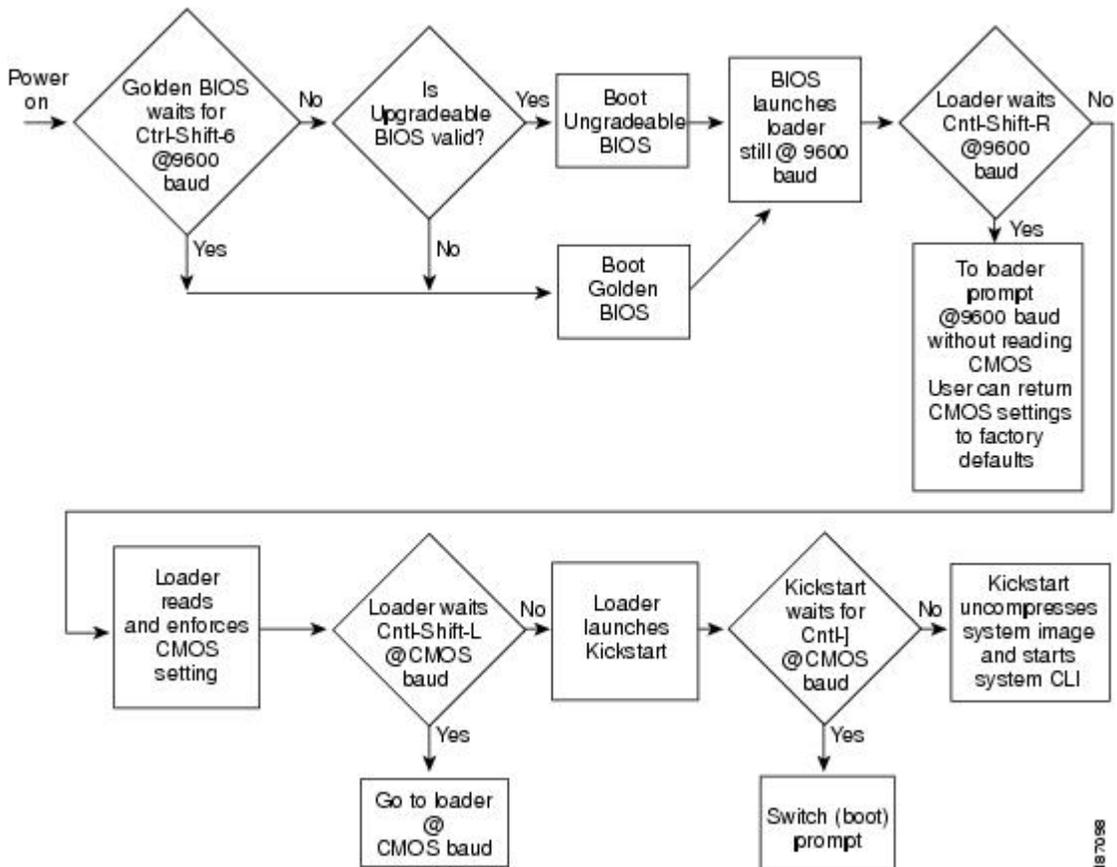


Note When you press **Ctrl-Shift-6**, the console settings must be set to their defaults: 9600 baud, 8 data bits, no parity, and 1 stop bit.

Before the boot sequence starts, the BIOS performs internal tests on the switch. If the tests fail, then the loader does not gain control. Instead, the BIOS image retains control and prints a message to the console at 9600 baud every 30 seconds that indicates a failure.

The following figure shows the normal and recovery boot sequence.

Figure 5: Boot Sequence



187096

For additional information, see *Troubleshooting*.

Console Settings

The loader, kickstart, and system images have the following factory default console settings:

- Speed—9600 baud
- Databits—8 bits per byte
- Stopbits—1 bit
- Parity—none

These settings are stored on the switch, and all three images use the stored console settings.

To change a console setting, use the **line console** command in configuration mode. The following example configures a line console and sets the options for that terminal line:

```
switch# configure terminal
switch(config)# line console
switch(config-console)# databits 7
switch(config-console)# exec-timeout 30
switch(config-console)# parity even
switch(config-console)# stopbits 2
```

You cannot change the BIOS console settings. These are the same as the default console settings.

Upgrading the Switch Software



Note

You must have the network-admin role before you can upgrade the software image on the switch.

You must log in to the switch on its console port connection.

To upgrade the software on the switch, follow these steps:

SUMMARY STEPS

1. Log in to Cisco.com to access the Software Download Center. To log in to Cisco.com, go to the URL <http://www.cisco.com/> and click **Log In** at the top of the page. Enter your Cisco username and password.
2. Access the Software Download Center using this URL: <http://www.cisco.com/kobayashi/sw-center/index.shtml>
3. Navigate to the software downloads for Cisco Nexus 5000 Series switches.
4. Read the release notes for the related image file.
5. Select and download the kickstart and system software files to a local server.
6. Ensure that the required space is available in the bootflash: directory for the image file(s) to be copied.
7. If you need more space on the active supervisor module bootflash, delete unnecessary files to make space available.
8. Copy the kickstart and system images to the switch bootflash using a transfer protocol. You can use **ftp**, **tftp**, **scp**, or **sftp**. The examples in this procedure use **scp**.
9. Install the new images, specifying the new image names that you downloaded in the previous step.
10. After the switch completes the installation, log in and verify that the switch is running the required software version.

DETAILED STEPS

Step 1 Log in to Cisco.com to access the Software Download Center. To log in to Cisco.com, go to the URL <http://www.cisco.com/> and click **Log In** at the top of the page. Enter your Cisco username and password.

Note Unregistered Cisco.com users cannot access the links provided in this document.

Step 2 Access the Software Download Center using this URL: <http://www.cisco.com/kobayashi/sw-center/index.shtml>

Step 3 Navigate to the software downloads for Cisco Nexus 5000 Series switches. You see links to the download images for the switch.

Step 4 Read the release notes for the related image file.

Step 5 Select and download the kickstart and system software files to a local server.

Step 6 Ensure that the required space is available in the bootflash: directory for the image file(s) to be copied.

Example:

```
switch# dir bootflash:
 4681      Nov 24 02:43:52 2008  config
13176836  Nov 24 07:19:36 2008  gdb.1
 49152     Jan 12 18:38:36 2009  lost+found/
 310556   Dec 23 02:53:28 2008  n1
20058112  Nov 07 02:35:22 2008  n5000-uk9-kickstart.4.0.1a.N1.0.62.bin
20217856  Jan 12 18:26:54 2009  n5000-uk9-kickstart.4.0.1a.N2.0.140.bin
76930262  Nov 07 02:35:22 2008  n5000-uk9.4.0.1a.N1.0.62.bin
103484727 Jan 12 18:29:08 2009  n5000-uk9.4.0.1a.N2.0.140.bin
```

```
Usage for bootflash://sup-local
74934272 bytes used
5550080 bytes free
80484352 bytes total
```

Caution We recommend that you keep the kickstart and system image files for at least one previous software release to use if the new image files do not load successfully.

Step 7 If you need more space on the active supervisor module bootflash, delete unnecessary files to make space available.

Example:

```
switch# delete bootflash:n5000-uk9-kickstart.4.0.1a.N1.0.62.bin
switch# delete bootflash:n5000-uk9.4.0.1a.N1.0.62.bin
```

- Step 8** Copy the kickstart and system images to the switch bootflash using a transfer protocol. You can use **ftp**, **tftp**, **scp**, or **sftp**. The examples in this procedure use **scp**.

Example:

```
switch# copy scp://user@scpserver.cisco.com/downloads/n5000-uk9.4.1.3.N1.0.96.bin
bootflash:n5000-uk9.4.1.3.N1.0.96.bin
switch# copy scp://user@scpserver.cisco.com/downloads/n5000-uk9-kickstart.4.1.3.N1.0.96.bin
bootflash:n5000-uk9-kickstart.4.1.3.N1.0.96.bin
```

- Step 9** Install the new images, specifying the new image names that you downloaded in the previous step.

Example:

```
switch# install all kickstart bootflash:n5000-uk9-kickstart.4.1.3.N1.0.96.bin system
bootflash:n5000-uk9.4.1.3.N1.0.96.bin
```

The **install all** command performs the following actions:

- Performs compatibility checks (equivalent to the **show incompatibility** command) for the images that you have specified. If there are compatibility issues, an error message is displayed and the installation does not proceed.
 - Displays the compatibility check results and displays whether the installation is disruptive.
 - Provides a prompt to allow you to continue or abort the installation.
- Caution** After completing the installation, all traffic through the switch is disrupted while the switch reboots.
- Updates the boot variables to reference the specified images and saves the configuration to the startup configuration file.

- Step 10** After the switch completes the installation, log in and verify that the switch is running the required software version.

Example:

```
switch# show version
Cisco Nexus Operating System (NX-OS) Software
TAC support: http://www.cisco.com/tac
Copyright (c) 2002-2009, Cisco Systems, Inc. All rights reserved.
The copyrights to certain works contained herein are owned by
other third parties and are used and distributed under license.
Some parts of this software are covered under the GNU Public
License. A copy of the license is available at
http://www.gnu.org/licenses/gpl.html.

Software
  BIOS:          version 1.2.0
  loader:        version N/A
  kickstart:     version 4.1(3)N1(1) [build 4.1(3)N1(0.96)]
  system:        version 4.1(3)N1(1) [build 4.1(3)N1(0.96)]
  BIOS compile time: 06/19/08
  kickstart image file is: bootflash:/n5000-uk9-kickstart.4.1.3.N1.0.96.bin
  kickstart compile time: 7/14/2009 4:00:00 [07/14/2009 04:27:38]
  system image file is: bootflash:/n5000-uk9.4.1.3.N1.0.96.bin
  system compile time: 7/14/2009 4:00:00 [07/14/2009 05:20:12]
```

```

Hardware
  cisco Nexus5020 Chassis ("40x10GE/Supervisor")
  Intel(R) Celeron(R) M CPU with 2074240 kB of memory.
  Processor Board ID JAB1232002F

  Device name: switch
  bootflash: 1003520 kB

Kernel uptime is 13 day(s), 23 hour(s), 25 minute(s), 5 second(s)

Last reset at 720833 usecs after Tue Jul 14 11:18:32 2009

Reason: Reset by installer
System version: 4.1(3)N1(0.96)
Service:

plugin
  Core Plugin, Ethernet Plugin

```

Downgrading from a Higher Release

The procedure to downgrade the switch is identical to a switch upgrade, except that the image files to be loaded are for an earlier release than the image currently running on the switch.



Note Prior to downgrading to a specific release, check the release notes for the current release installed on the switch, to ensure that your hardware is compatible with the specific release. There are special caveats you must be aware of before you downgrade the switch software to a 4.0(0)-based release. See the *Cisco Nexus 5000 Series and Cisco Nexus 2000 Series Release Notes* for details.

To downgrade the software on the switch, follow these steps:

SUMMARY STEPS

1. Locate the image files you will use for the downgrade by entering the **dir bootflash:** command.
2. Install the new images.
3. After the switch completes the installation, log in and verify that the switch is running the required software version.

DETAILED STEPS

Step 1

Locate the image files you will use for the downgrade by entering the **dir bootflash:** command.

If the image files are not stored on the bootflash memory, download the files from Cisco.com:

- a) Log in to Cisco.com to access the Software Download Center. To log in to Cisco.com, go to the URL <http://www.cisco.com/> and click **Log In** at the top of the page. Enter your Cisco username and password.

Note Unregistered Cisco.com users cannot access the links provided in this document.

- b) Access the Software Download Center using this URL: <http://www.cisco.com/kobayashi/sw-center/index.shtml>
- c) Navigate to the software downloads for Cisco Nexus 5000 Series switches.

You see links to the download images for the switch.

- d) Read the release notes for the related image file then select and download the kickstart and system software files to a local server
- e) Ensure that the required space is available in the bootflash: directory for the image file(s) to be copied.

Caution We recommend that you keep the kickstart and system image files for at least one previous software release to use if the new image files do not load successfully.
- f) Copy the kickstart and system images to the switch bootflash using a transfer protocol. You can use **ftp**, **tftp**, **scp**, or **sftp**.

Step 2 Install the new images.

Example:

```
switch# install all kickstart bootflash:n5000-uk9-kickstart.4.0.1a.N1.0.62.bin system
bootflash:n5000-uk9.4.0.1a.N1.0.62.bin
```

The **install all** command performs the following actions:

- Performs compatibility checks (equivalent to the **show incompatibility** command) for the images that you have specified. If there are compatibility issues, an error message is displayed and the installation does not proceed.
- Displays the compatibility check results and displays whether the installation is disruptive.
- Provides a prompt to allow you to continue or abort the installation.

Note A disruptive installation causes traffic disruption while the switch reboots.
- Updates the boot variables to reference the specified images and saves the configuration to the startup configuration file.

Step 3 After the switch completes the installation, log in and verify that the switch is running the required software version.

Example:

```
switch# show version
```

Initial Configuration

Configuration Prerequisites

The following procedure is a review of the tasks you should have completed during hardware installation. These tasks must be completed before you can configure the switch.

Before you can configure a switch, follow these steps:

SUMMARY STEPS

1. Verify the following physical connections for the new Cisco Nexus 5000 Series switch:
2. Verify that the default console port parameters are identical to those of the computer terminal (or terminal server) attached to the switch console port:

DETAILED STEPS

Step 1 Verify the following physical connections for the new Cisco Nexus 5000 Series switch:

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

Refer to the *Cisco Nexus 5000 Series Hardware Installation Guide* (for the required product) for more information.

Tip Save the host ID information for future use (for example, to enable licensed features). The host ID information is provided in the Proof of Purchase document that accompanies the switch.

Step 2 Verify that the default console port parameters are identical to those of the computer terminal (or terminal server) attached to the switch console port:

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

Initial Setup

The first time that you access a switch in the Cisco Nexus 5000 Series, it runs a setup program that prompts you for the IP address and other configuration information necessary for the switch to communicate over the Ethernet interface. This information is required to configure and manage the switch.



Note

The IP address can only be configured from the CLI. When the switch powers up for the first time, you should assign the IP address. After you perform this step, the Cisco MDS 9000 Family Fabric Manager can reach the switch through the console port.

Preparing to Configure the Switch

Before you configure Cisco Nexus 5000 Series switch for the first time, you need the following information:

- Administrator password.



Note

If a password is weak (short, easy-to-decipher), your password configuration is rejected. Be sure to configure a strong password.

- If you are using an IPv4 address for the management interface, you need the following information:
 - IPv4 subnet mask for the switch's management interface.

- IPv4 address of the default gateway (optional).

- SSH service on the switch (optional).

To enable this service, select the type of SSH key (dsa/rsa/rsa1) and number of SSH key bits (768 to 2048).

- NTP server IPv4 address (optional).
- SNMP community string (optional).
- Switch name (optional).

This is your switch prompt.

- An additional login account and password (optional).

**Note**

If you are using IPv4, be sure to configure the IPv4 route, the IPv4 default network address, and the IPv4 default gateway address to enable SNMP access.

Default Login

The switch has the network administrator as a default user (admin). You cannot change the default user at any time.

There is no default password so you must explicitly configure a strong password. If a password is trivial (short, easy-to-decipher), your password configuration is rejected. Be sure to configure a strong password. If you configure and subsequently forget this new password, you have the option to recover this password.

**Note**

If you enter the **write erase** command and reload the switch, you must reconfigure the default user (admin) password using the setup procedure.

Configuring the Switch

This section describes how to initially configure the switch.

**Note**

Press **Ctrl-C** at any prompt to skip the remaining configuration options and proceed with what you have configured up to that point. However, entering the new password for the administrator is a requirement and cannot be skipped.

**Tip**

If you do not want to answer a previously configured question, or if you want to skip answers to any questions, press **Enter**. If a default answer is not available (for example, switch name), the switch uses what was previously configured and skips to the next question.

To configure the switch for first time, follow these steps:

SUMMARY STEPS

1. Ensure that the switch is on. Switches in the Cisco Nexus 5000 Series boot automatically.
2. Enter the new password for the administrator.
3. Enter yes to enter the setup mode.
4. Enter the new password for the administrator (admin is the default).
5. Enter yes (no is the default) to create additional accounts.
6. Enter yes (yes is the default) to create an SNMP read-only community string.
7. Enter a name for the switch.
8. Enter yes (yes is the default) to configure out-of-band management and enter the mgmt0 IPv4 address.
9. Enter yes (yes is the default) to configure the IPv4 default gateway (recommended) and enter the IPv4 address for the default gateway.
10. Enter yes (yes is the default) to enable the Telnet service.
11. Enter yes (no is the default) to enable the SSH service.
12. Enter yes (no is the default) to configure the NTP server and enter the IPv4 address for the NTP server.
13. Enter yes (yes is the default) to configure basic Fibre Channel configurations.
14. Enter shut (shut is the default) to configure the default Fibre Channel switch port interface to the shut (disabled) state.
15. Enter on (on is the default) to configure the switch port trunk mode.
16. Enter permit (deny is the default) to deny a default zone policy configuration.
17. Enter yes (no is the default) to enable a full zone set distribution.
18. You see the new configuration. Review and edit the configuration that you have just entered. Enter no (no is the default) if you are satisfied with the configuration.
19. Enter yes (yes is default) to use and save this configuration:

DETAILED STEPS

Step 1 Ensure that the switch is on. Switches in the Cisco Nexus 5000 Series boot automatically.

Step 2 Enter the new password for the administrator.

Example:

Enter the password for admin: *<password>*

Note Clear text passwords cannot contain dollar signs (\$) or spaces anywhere in the password. Also, they cannot include these special characters at the beginning of the password: quotation marks (" or '), vertical bars (|), or right angle brackets (>).

Tip If a password is weak (short, easy-to-decipher), your password configuration is rejected. Be sure to configure a strong password. Passwords are case-sensitive.

Step 3 Enter yes to enter the setup mode.

Example:

This setup utility will guide you through the basic configuration of the system. Setup configures only enough connectivity for management of the system.

*Note: setup is mainly used for configuring the system initially, when no configuration is present. So setup always assumes system defaults and not the current system configuration values.

Press Enter at anytime to skip a dialog. Use ctrl-c at anytime to skip the remaining dialogs.

Would you like to enter the basic configuration dialog (yes/no): **yes**

The setup utility guides you through the basic configuration process. Press **Ctrl-C** at any prompt to end the configuration process.

Step 4 Enter the new password for the administrator (admin is the default).

Example:

Enter the password for admin: **admin**

Step 5 Enter yes (no is the default) to create additional accounts.

Example:

Create another login account (yes/no) [n]: **yes**

While configuring your initial setup, you can create an additional user account (in the network-admin role) besides the administrator's account.

a) Enter the user login ID.

Example:

Enter the user login ID: *user_name*

b) Enter the user password.

Example:

Enter the password for *user_name*: *user-password*

Step 6 Enter yes (yes is the default) to create an SNMP read-only community string.

Example:

Configure read-only SNMP community string (yes/no) [n]: **yes**

SNMP community string: *snmp_community*

Step 7 Enter a name for the switch.

Note The switch name is limited to 32 alphanumeric characters. The default name is "switch".

Example:

Enter the switch name: *switch_name*

Step 8 Enter yes (yes is the default) to configure out-of-band management and enter the mgmt0 IPv4 address.

Example:

Continue with Out-of-band (mgmt0) management configuration? [yes/no]: **yes**
 Mgmt0 IPv4 address: *ip_address*

Step 9 Enter yes (yes is the default) to configure the IPv4 default gateway (recommended) and enter the IPv4 address for the default gateway.

Example:

Configure the default-gateway: (yes/no) [y]: **yes**
 IPv4 address of the default-gateway: *default_gateway*

Step 10 Enter yes (yes is the default) to enable the Telnet service.

Example:

Enable the telnet service? (yes/no) [y]: **yes**

Step 11 Enter yes (no is the default) to enable the SSH service.

Example:

Enabled SSH service? (yes/no) [n]: **yes**

a) Enter the SSH key type that you would like to generate.

Example:

Type the SSH key you would like to generate (dsa/rsa/rsal)? **dsa**

b) Enter the number of key bits within the specified range.

Example:

Enter the number of key bits? (768 to 2048): **768**

Step 12 Enter yes (no is the default) to configure the NTP server and enter the IPv4 address for the NTP server.

Example:

Configure NTP server? (yes/no) [n]: **yes**
 NTP server IP address: *ntp_server_IP_address*

Step 13 Enter yes (yes is the default) to configure basic Fibre Channel configurations.

Example:

Enter basic FC configurations (yes/no) [n]: **yes**

Step 14 Enter shut (shut is the default) to configure the default Fibre Channel switch port interface to the shut (disabled) state.

Example:

Configure default physical FC switchport interface state (shut/noshut) [shut]: **shut**

Step 15 Enter on (on is the default) to configure the switch port trunk mode.

Example:

```
Configure default physical FC switchport trunk mode (on/off/auto) [on]: on
```

Step 16 Enter permit (deny is the default) to deny a default zone policy configuration.

Example:

```
Configure default zone policy (permit/deny) [deny]: permit
```

Permits traffic flow to all members of the default zone.

Note If you are executing the setup script after entering a **write erase** command, you explicitly must change the default zone policy to permit for VSAN 1 after finishing the script using the following command:

```
Configure read-only SNMP community string (yes/no) [n]: zone default-zone permit vsan 1
```

Step 17 Enter yes (no is the default) to enable a full zone set distribution.

Example:

```
Enable full zoneset distribution (yes/no) [n]: yes
```

Overrides the switch-wide default for the full zone set distribution feature.

Step 18 You see the new configuration. Review and edit the configuration that you have just entered. Enter no (no is the default) if you are satisfied with the configuration.

Example:

```
The following configuration will be applied:
username admin password <user-password> role network-admin
snmp-server community snmp_community ro
switchname switch
feature telnet
ssh key dsa 768 force
feature ssh
system default switchport shutdown san
system default switchport trunk mode on
system default zone default-zone permit
system default zone distribute full
Would you like to edit the configuration? (yes/no) [n]: no
```

Step 19 Enter yes (yes is default) to use and save this configuration:

Example:

```
Use this configuration and save it? (yes/no) [y]: yes
```

Caution If you do not save the configuration at this point, none of your changes are updated the next time the switch is rebooted. Type yes to save the new configuration. This operation ensures that the kickstart and system images are also automatically configured.

Related Topics

[Image Files on the Switch, on page 31](#)

Changing the Initial Configuration

To make changes to the initial configuration at a later time, enter the **setup** command in EXEC mode:

```
switch# setup
---- Basic System Configuration Dialog ----
This setup utility will guide you through the basic configuration of
the system. Setup configures only enough connectivity for management
of the system.
*Note: setup is mainly used for configuring the system initially,
when no configuration is present. So setup always assumes system
defaults and not the current system configuration values.
Press Enter at anytime to skip a dialog. Use ctrl-c at anytime
to skip the remaining dialogs.
Would you like to enter the basic configuration dialog (yes/no): yes
```

The setup utility guides you through the basic configuration process.

Management Interface Configuration

The management interface on the switch allows multiple simultaneous Telnet, SSH, or SNMP sessions. You can remotely configure the switch through the management interface (mgmt0), but first you must configure some IP parameters so that the switch is reachable. You can manually configure the management interface from the CLI through the console port.

About the mgmt0 Interface

The mgmt0 interface on a Cisco Nexus 5000 Series switch provides out-of-band management, which enables you to manage the switch by its IPv4 or IPv6 address. The mgmt0 interface is a 10/100/1000 Ethernet port.



Note

Before you begin to configure the management interface manually, obtain the switch's IP address and subnet mask. Also make sure that the console cable is connected to the console port.

Configuring the Management Interface

To configure the management (mgmt0) Ethernet interface to connect over IP, perform this task:

SUMMARY STEPS

1. switch# **configure terminal**
2. switch(config)# **interface mgmt 0**
3. Configure the IP address for IPv4 or IPv6:
4. switch(config-if)# **no shutdown**
5. switch(config-if)# **exit**
6. switch(config)# **vrf context management**
7. Configure the IP address (IPv4 or IPv6) for the next hop:
8. switch(config-vrf)# **exit**
9. (Optional) switch# **copy running-config startup-config**

DETAILED STEPS

- Step 1** switch# **configure terminal**
Enters configuration mode.
- Step 2** switch(config)# **interface mgmt 0**
Selects the management Ethernet interface on the switch and enters interface configuration submode.
- Step 3** Configure the IP address for IPv4 or IPv6:
- a) switch(config-if)# **ip address** *ipv4-address* [/ *length*]
Configures the IPv4 address and its subnet mask.
 - b) switch(config-if)# **ip address** *ipv4-address* [*subnet-mask*]
An alternative method that configures the IPv4 address and its subnet mask.
 - c) switch(config-if)# **ipv6 address** *ipv6-address* [/ *length*]
Configures the IPv6 address and its subnet mask.
- Step 4** switch(config-if)# **no shutdown**
Enables the interface.
- Step 5** switch(config-if)# **exit**
Returns to configuration mode.
- Step 6** switch(config)# **vrf context management**
Enters VRF context management configuration mode.
- Step 7** Configure the IP address (IPv4 or IPv6) for the next hop:
- a) switch(config-vrf)# **ip route** *ipv4-prefix* [/ *length*] *ipv4-nexthop-address*
Configures the IPv4 address of the next hop.
 - b) switch(config-vrf)# **ipv6 route** *ipv6-prefix* [/ *length*] *ipv6-nexthop-address*
Configures the IPv6 address of the next hop.
- Step 8** switch(config-vrf)# **exit**
Returns to EXEC mode.
- Step 9** (Optional) switch# **copy running-config startup-config**
Saves your configuration changes to the file system.
-

In some cases, a switch interface might be administratively shut down. You can check the status of an interface at any time by using the **show interface mgmt 0** command.

Displaying Management Interface Configuration

To display the management interface configuration, use the **show interface mgmt 0** command.

```
switch# show interface mgmt0
mgmt0 is up
  Hardware is GigabitEthernet, address is 000d.ec8f.cb00 (bia 000d.ec8f.cb00)
  Internet Address is 172.16.131.202/24
  MTU 1500 bytes, BW 0 Kbit, DLY 0 usec,
    reliability 255/255, txload 1/255, rxload 1/255
  Encapsulation ARPA
  full-duplex, 1000 Mb/s
  Input flow-control is off, output flow-control is off
  8540 packets input, 2835036 bytes
  5202 multicast frames, 0 compressed
  0 input errors, 0 frame, 0 overrun, 0 fifo
  570 packets output, 85555 bytes
  0 underrun, 0 output errors, 0 collisions
  0 fifo, 0 carrier errors
```

Shutting Down the Management Interface

To shut down the management interface (mgmt0), you use the **shutdown** command. A system prompt requests you confirm your action before it executes the command. You can use the force option to bypass this confirmation.

The following example shuts down the interface without using the force option:

```
switch# configure terminal
switch(config)# interface mgmt 0
switch(config-if)# shutdown
Shutting down this interface will drop all telnet sessions.
Do you wish to continue (y/n)? y
```

The following example shuts down the interface using the force option:

```
switch# configure terminal
switch(config)# interface mgmt 0
switch(config-if)# shutdown force
```



Managing Licenses

This chapter describes how to manage licenses on Cisco Nexus 5000 Series switches. It contains the following sections:

- [Licensing Terminology, page 47](#)
- [Licensing Model, page 48](#)
- [Licence Installation, page 49](#)
- [Obtaining the License Key File , page 50](#)
- [Installing the License Key File, page 51](#)
- [Backing Up License Files, page 52](#)
- [Identifying License Features in Use, page 53](#)
- [Uninstalling Licenses, page 53](#)
- [Updating Licenses, page 54](#)
- [Grace Period Alerts, page 56](#)
- [License Transfers Between Switches, page 57](#)
- [Verifying the License Configuration, page 57](#)

Licensing Terminology

Licensing allows you to access specified premium features on the switch after you install the appropriate license for that feature. This information describes licensing types, options, procedures, installation, and management for a Cisco Nexus 5000 Series switch.

The following terms are used:

- **Licensed feature**—Permission to use a particular feature through a license file, a hardware object, or a legal contract. This permission is limited to the number of users, number of instances, time span, and the implemented switch.
- **Licensed application**—A software feature that requires a license to be used.

- License enforcement—A mechanism that prevents a feature from being used without first obtaining a license.
- Node-locked license—A license that can only be used on a particular switch using the switch's unique host ID.
- Host IDs—A unique chassis serial number that is specific to each switch.
- Proof of purchase—A document entitling its rightful owner to use licensed features on one switch as described in that document. The proof of purchase document is also known as the claim certificate.
- Product Authorization Key (PAK)—The PAK allows you to obtain a license key from one of the sites listed in the proof of purchase document. After registering at the specified website, you will receive your license key file and installation instructions through e-mail.
- License key file—A switch-specific unique file that specifies the licensed features. Each file contains digital signatures to prevent tampering and modification. License keys are required to use a licensed feature. License keys are enforced within a specified time span.
- Missing license—If the bootflash has been corrupted or a supervisor module replaced after you have installed a license, that license shows as "missing." The feature still works, but the license count is inaccurate. You should reinstall the license as soon as possible.
- Incremental license—An additional licensed feature that was not in the initial license file. License keys are incremental. If you purchase some features now and others later, the license file and the software detect the sum of all features for the specified switch.
- Evaluation license—A temporary license. Evaluation licenses are time bound (valid for a specified number of days) and are not tied to a host ID (switch serial number).
- Permanent license—A license that is not time bound is called a permanent license.
- Grace period—The amount of time the features in a license package can continue functioning without a license.
- Support—If you purchased Cisco support through a Cisco reseller, contact the reseller directly. If you purchased support directly from Cisco Systems, contact Cisco Technical Support at this URL: <http://www.cisco.com/warp/public/687/Directory/DirTAC.shtml>

Licensing Model

The licensing model for the Cisco NX-OS software is feature-based. Feature-based licenses make features available to the entire physical switch. The following table lists the feature-based license packages.



Note

Any feature not included in the Storage Services license package is bundled with the Cisco NX-OS software and is provided with the switch hardware at no additional charge (See Base Services Package in the table).

Table 6: Feature-Based Licenses

Feature License	Features
Base Services Package N5000-AS	This package is included with the switch hardware at no additional charge. It includes all available Ethernet and system features, except features explicitly listed in the Storage Services Package.
Nexus 5010 Storage Protocols Services License N5010-SSK9	<ul style="list-style-type: none"> • N5010-SS includes the following services for one NX5010 system: • Native Fibre Channel • FCoE • NPV • FC Port Security • Fabric Binding
Nexus 5020 Storage Protocols Services License N5020-SS	<ul style="list-style-type: none"> • N5020-SS includes the following services for one NX5020 system: • Native Fibre Channel • FCoE • NPV • FC Port Security • Fabric Binding

Licence Installation

Obtaining a Factory-Installed License

You can obtain factory-installed licenses for a new Cisco Nexus 5000 Series switch.

To obtain a factory-installed license, perform this task:

SUMMARY STEPS

1. Contact your reseller or Cisco representative and request this service.
2. Your switch is shipped with the required licenses installed in the system. The proof of purchase document is sent along with the switch.
3. Obtain the host ID from the proof of purchase document for future use.

DETAILED STEPS

-
- Step 1** Contact your reseller or Cisco representative and request this service.
- Note** If you purchased Cisco support through a Cisco reseller, contact the reseller directly. If you purchased support directly from Cisco Systems, contact Cisco Technical Support at this URL: <http://www.cisco.com/warp/public/687/Directory/DirTAC.shtml>
- Step 2** Your switch is shipped with the required licenses installed in the system. The proof of purchase document is sent along with the switch.
- Step 3** Obtain the host ID from the proof of purchase document for future use.
-

Performing a Manual Installation

All Cisco Nexus 5000 Series licenses are factory-installed. Manual installation is not required.

Obtaining the License Key File

To obtain new or updated license key files, perform this task:

SUMMARY STEPS

1. Use the **show license host-id** command to obtain the serial number for your switch. The host ID is also referred to as the switch serial number.
2. Obtain either your claim certificate or your proof of purchase document. This document accompanies every Cisco Nexus 5000 Series switch.
3. Get the product authorization key (PAK) from either the claim certificate or the proof of purchase document.
4. Locate the website URL from either the claim certificate or the proof of purchase document.
5. Access the specified URL that applies to your switch and enter the switch serial number and the PAK.
6. Use the **copy licenses** command in EXEC mode to save your license file to one of two locations; either the bootflash or the volatile directory.

DETAILED STEPS

-
- Step 1** Use the **show license host-id** command to obtain the serial number for your switch. The host ID is also referred to as the switch serial number.

```
switch# show license host-id
License hostid: VDH=FOX064317SQ
```

Tip Use the entire ID that appears after the equals sign (=) sign. In this example, the host ID is FOX064317SQ.

- Step 2** Obtain either your claim certificate or your proof of purchase document. This document accompanies every Cisco Nexus 5000 Series switch.
- Step 3** Get the product authorization key (PAK) from either the claim certificate or the proof of purchase document.
- Step 4** Locate the website URL from either the claim certificate or the proof of purchase document.
- Step 5** Access the specified URL that applies to your switch and enter the switch serial number and the PAK. The license key file is sent to you by e-mail. The license key file is digitally signed to only authorize use on the requested switch. The requested features are also enabled once the Cisco NX-OS software on the specified switch accesses the license key file.
- Caution** Install the license key file in the specified Cisco Nexus 5000 Series switch without making any modifications. A license is either permanent or it expires on a fixed date. If you do not have a license, the grace period for using that feature starts from the first time you start using a feature offered by that license.
- Step 6** Use the **copy licenses** command in EXEC mode to save your license file to one of two locations; either the bootflash or the volatile directory.

Related Topics

[Backing Up License Files, on page 52](#)

[Grace Period Alerts, on page 56](#)

Installing the License Key File

To install a license key file in any switch, perform this task:

SUMMARY STEPS

1. Log into the switch through the console port.
2. Perform the installation by entering the **install license** command from the switch console.
3. Back up the license file to a .tar file on bootflash: using the **copy licenses** command.
4. Exit the switch console and open a new terminal session to view all license files installed on the switch using the **show license** command.

DETAILED STEPS

- Step 1** Log into the switch through the console port.
- Step 2** Perform the installation by entering the **install license** command from the switch console.

```
switch# install license bootflash:license_file.lic
Installing license ..done
```

Note If you provide a target name for the license key file, the file is installed with the specified name. Otherwise, the filename specified in the license key file is used to install the license.

Step 3 Back up the license file to a .tar file on bootflash: using the **copy licenses** command.

```
switch# copy licenses bootflash:/Enterprise.tar
Backing up license done
```

Step 4 Exit the switch console and open a new terminal session to view all license files installed on the switch using the **show license** command.

```
switch# show license
Enterprise.lic:
SERVER this_host ANY
VENDOR cisco
INCREMENT ENTERPRISE_PKG cisco 1.0 permanent uncounted \
  HOSTID=VDH=FOX0646S017 \
  NOTICE="<LicFileID></LicFileID><LicLineID>0</LicLineID> \
  <PAK>dummyPak</PAK>" SIGN=EE9F91EA4B64
```

Note If the license meets all guidelines when the **install license** command is entered, all features and modules continue functioning as configured.

You can use the **show license brief** command to display a list of license files installed on the switch.

```
switch# show license brief
Enterprise.lic
FibreChannel.lic
```

You can use the **show license file** command to display information about a specific license file installed on the switch.

```
switch# show license file Enterprise.lic
Enterprise.lic:
SERVER this_host ANY
VENDOR cisco
INCREMENT ENTERPRISE_PKG cisco 1.0 permanent uncounted \
  HOSTID=VDH=FOX0646S017 \
  NOTICE="<LicFileID></LicFileID><LicLineID>0</LicLineID> \
  <PAK>dummyPak</PAK>" SIGN=EE9F91EA4B64
```

Backing Up License Files

All installed license files can be backed up as a .tar file in the user specified location. Use the **copy licenses** command in EXEC mode to save your license file to one of two locations; bootflash: or volatile:. The following example saves all licenses to a file named Enterprise.tar:

```
switch# copy licenses bootflash:/Enterprise.tar
Backing up license done
```



Tip

We recommend backing up your license files immediately after installing them and just before running a **write erase** command.

**Caution**

If you erase any existing licenses, you can only install them using the **install license** command.

Identifying License Features in Use

When a Cisco NX-OS software feature is enabled, it can activate a license grace period. To identify the features active for a specific license, use the **show license usage** *license-name* command.

```
switch# show license usage FC_FEATURES_PKG
Application
-----
PFM
-----
```

Use the **show license usage** command to identify all of the active features on your switch.

```
switch# show license usage
Feature                               Ins  Lic  Status Expiry Date Comments
                               Count
-----
FM_SERVER_PKG                         No   -   Unused          -
ENTERPRISE_PKG                        No   -   In use          Grace 119D 23H
FC_FEATURES_PKG                        Yes  -   In use never    -
-----
```

Uninstalling Licenses

You can only uninstall a permanent license that is not in use. If you try to delete a permanent license that is currently being used, the software rejects the request with an error message. Uninstalling an unused license initiates the grace period. The grace period is measured from the first use of the feature without a license and is reset when a valid license file is installed.

**Note**

Permanent licenses cannot be uninstalled if they are currently being used. Features turned on by permanent licenses must first be disabled, before that license is uninstalled.

If you are using an evaluation license and would like to install a new permanent license, you can do so without service disruption and before the evaluation license expires. Removing an evaluation license immediately triggers a grace period without service disruption.

**Caution**

Disable related features before uninstalling a license. The delete procedure fails if the license is in use.

SUMMARY STEPS

1. Save your running configuration to a remote server using the **copy** command
2. Enter the **show license brief** command in EXEC mode to view a list of all installed license key files and identify the file to be uninstalled. In this example, the file to be uninstalled is the FibreChannel.lic file.
3. Disable the features provided by the license to be uninstalled. Enter the **show license usage package-name** command to view the enabled features for a specified package.
4. Uninstall the FibreChannel.lic file using the **clear license filename** command, where *filename* is the name of the installed license key file.
5. Enter yes (yes is the default) to continue with the license update.

DETAILED STEPS

-
- Step 1** Save your running configuration to a remote server using the **copy** command
- Step 2** Enter the **show license brief** command in EXEC mode to view a list of all installed license key files and identify the file to be uninstalled. In this example, the file to be uninstalled is the FibreChannel.lic file.
- ```
switch# show license brief
Enterprise.lic
FibreChannel.lic
```
- Step 3** Disable the features provided by the license to be uninstalled. Enter the **show license usage package-name** command to view the enabled features for a specified package.
- ```
switch# show license usage FC_FEATURES_PKG
Application
-----
PFM
-----
```
- Step 4** Uninstall the FibreChannel.lic file using the **clear license filename** command, where *filename* is the name of the installed license key file.
- ```
switch# clear license FibreChannel.lic
Clearing license FibreChannel.lic:
SERVER this_host ANY
VENDOR cisco
```
- Step 5** Enter yes (yes is the default) to continue with the license update.
- ```
Do you want to continue? (y/n) yes
Clearing license ..done
```
- The FibreChannel.lic license key file is now uninstalled.
-

Updating Licenses

If your license is time bound, you must obtain and install an updated license. Contact technical support to request an updated license.

**Note**

If you purchased Cisco support through a Cisco reseller, contact the reseller directly. If you purchased support directly from Cisco Systems, contact Cisco Technical Support at this URL: <http://www.cisco.com/warp/public/687/Directory/DirTAC.shtml>

To update a license, perform this task:

SUMMARY STEPS

1. Obtain the updated license file using the following procedure:
2. Save your running configuration to a remote server using the **copy** command.
3. Enter the **show license brief** command to verify the name of the file to be updated.
4. Update the license file using the **update license url** command, where *url* specifies the bootflash or volatile location of the updated license file.
5. Enter **yes** (yes is the default), to continue with the license update.

DETAILED STEPS**Step 1**

Obtain the updated license file using the following procedure:

- a) Use the **show license host-id** command to obtain the serial number for your switch. The host ID is also referred to as the switch serial number.
- b) Obtain either your claim certificate or your proof of purchase document. This document accompanies every Cisco Nexus 5000 Series switch.
- c) Get the product authorization key (PAK) from either the claim certificate or the proof of purchase document.
- d) Locate the website URL from either the claim certificate or the proof of purchase document.
- e) Access the specified URL that applies to your switch and enter the switch serial number and the PAK. The license key file is sent to you by e-mail. The license key file is digitally signed to only authorize use on the requested switch. The requested features are also enabled once the Cisco NX-OS software on the specified switch accesses the license key file.

Caution Install the license key file in the specified Cisco Nexus 5000 Series switch without making any modifications.

- f) Use the **copy licenses** command in EXEC mode to save your license file to one of two locations; either the bootflash or the volatile directory.

Step 2

Save your running configuration to a remote server using the **copy** command.

Step 3

Enter the **show license brief** command to verify the name of the file to be updated.

```
switch# show license brief
Enterprise.lic:
```

Step 4 Update the license file using the **update license url** command, where *url* specifies the bootflash or volatile location of the updated license file.

```
switch# update license bootflash:Advanced2.lic Advanced1.lic
Updating Advanced1.lic:
SERVER this_host ANY
VENDOR cisco
# An example fcports license
INCREMENT SAN_EXTN_OVER_IP cisco 1.000 permanent 1 HOSTID=VDH=ABCD \
    NOTICE=<LicFileID>Advanced1.lic</LicFileID><LicLineID>0</LicLineID> \
    SIGN=33088E76F668
with bootflash:/Advanced2.lic:
SERVER this_host ANY
VENDOR cisco
# An example fcports license
INCREMENT SAN_EXTN_OVER_IP cisco 1.000 permanent 1 HOSTID=VDH=ABCD \
    NOTICE=<LicFileID>Advanced2.lic</LicFileID><LicLineID>1</LicLineID> \
    SIGN=67CB2A8CCAC2
```

Step 5 Enter **yes** (yes is the default), to continue with the license update.

```
Do you want to continue? (y/n) y
Updating license ..done
```

The Enterprise.lic license key file is now updated.

Grace Period Alerts

Cisco NX-OS gives you a 120-day grace period. This grace period starts or continues when you are evaluating a feature for which you have not installed a license.

The grace period stops if you disable a feature you are evaluating, but if you enable that feature again without a valid license, the grace period countdown continues from when it had stopped.

The grace period operates across all features in a license package. License packages can contain several features. If you disable a feature during the grace period and there are other features in that license package that are still enabled, the countdown does not stop for that license package. To suspend the grace period countdown for a license package, you must disable every feature in that license package. Use the **show license usage license-name** command to determine which applications to disable.

```
switch# show license usage FC_FEATURES_PKG
Application
-----
PFM
-----
```

The Cisco NX-OS license counter keeps track of all licenses on a switch. If you are evaluating a feature and the grace period has started, you will receive console messages, SNMP traps, system messages, and Call Home messages on a daily basis.

The frequency of these messages become hourly during the last seven days of the grace period.

**Note**

You cannot modify the frequency of the grace period messages.

**Caution**

After the final seven days of the grace period, the feature is turned off and your network traffic may be disrupted. Any future upgrade will enforce license requirements and the 120-day grace period.

Use the **show license usage** command to display grace period information for a switch.

```
switch# show license usage
Feature                Installed    License Status  ExpiryDate  Comments
Count
-----
FM_SERVER_PKG          Yes         -               Unused      never       license missing
MAINFRAME_PKG          No          -               Unused      never       Grace Period 57days15hrs
ENTERPRISE_PKG         Yes         -               InUse       never       -
SAN_EXTN_OVER_IP      No          0               Unused      never       -
SAN_EXTN_OVER_IP_IPS4 No          0               Unused      never       -
-----
```

License Transfers Between Switches

A license is specific to the switch for which it is issued and is not valid on any other switch. If you need to transfer a license from one switch to another, contact your customer service representative.

**Note**

If you purchased Cisco support through a Cisco reseller, contact the reseller directly. If you purchased support directly from Cisco Systems, contact Cisco Technical Support at this URL: <http://www.cisco.com/warp/public/687/Directory/DirTAC.shtml>

Verifying the License Configuration

To display the license configuration information, perform one of the following tasks:

- **switch# show license [brief]**
Displays information for all installed license files.
- **switch# show license file**
Displays information for a specific license file.
- **switch# show license host-id**
Displays the host ID for the physical switch.
- **switch# show license usage**
Displays the usage information for installed licenses.



PART **II**

LAN Switching

- [Configuring Ethernet Interfaces, page 61](#)
- [Configuring VLANs, page 77](#)
- [Configuring Private VLANs, page 85](#)
- [Configuring Access and Trunk Interfaces, page 101](#)
- [Configuring EtherChannels, page 111](#)
- [Configuring Virtual Port Channels, page 125](#)
- [Configuring Rapid PVST+, page 155](#)
- [Configuring Multiple Spanning Tree, page 183](#)
- [Configuring STP Extensions, page 211](#)
- [Configuring the MAC Address Table, page 225](#)
- [Configuring IGMP Snooping, page 229](#)
- [Configuring Traffic Storm Control, page 237](#)



Configuring Ethernet Interfaces

This section describes the configuration of the Ethernet interfaces on a Cisco Nexus 5000 Series switch. It contains the following sections:

- [Information About Ethernet Interfaces, page 61](#)
- [Configuring Ethernet Interfaces, page 65](#)
- [Displaying Interface Information, page 72](#)

Information About Ethernet Interfaces

The Ethernet ports can operate as standard Ethernet interfaces connected to servers or to a LAN.

The Ethernet interfaces also support Fibre Channel over Ethernet (FCoE). FCoE allows the physical Ethernet link to carry both Ethernet and Fibre Channel traffic.

On a Cisco Nexus 5000 Series switch, the Ethernet interfaces are enabled by default.

About the Interface Command

You can enable the various capabilities of the Ethernet interfaces on a per-interface basis using the **interface** command. When you enter the **interface** command, you specify the following information:

- Interface type—All physical Ethernet interfaces use the **ethernet** keyword.
- Slot number
 - Slot 1 includes all the fixed ports.
 - Slot 2 includes the ports on the upper expansion module (if populated).
 - Slot 3 includes the ports on the lower expansion module (if populated).
- Port number
 - Port number within the group.

The interface numbering convention is extended to support use with a Cisco Nexus 2000 Series Fabric Extender as follows:

```
switch(config)# interface ethernet [chassis]/slot/port
```

- Chassis ID is an optional entry to address the ports of a connected Fabric Extender. The chassis ID is configured on a physical Ethernet or EtherChannel interface on the switch to identify the Fabric Extender discovered via the interface. The chassis ID ranges from 100 to 199.

About the Unidirectional Link Detection Parameter

UDLD Overview

The Cisco-proprietary Unidirectional Link Detection (UDLD) protocol allows ports that are connected through fiber optics or copper (for example, Category 5 cabling) Ethernet cables to monitor the physical configuration of the cables and detect when a unidirectional link exists. When the switch detects a unidirectional link, UDLD shuts down the affected LAN port and alerts the user. Unidirectional links can cause a variety of problems, including spanning tree topology loops.

UDLD is a Layer 2 protocol that works with the Layer 1 protocols to determine the physical status of a link. At Layer 1, autonegotiation takes care of physical signaling and fault detection. UDLD performs tasks that autonegotiation cannot perform, such as detecting the identities of neighbors and shutting down misconnected LAN ports. When you enable both autonegotiation and UDLD, Layer 1 and Layer 2 detections work together to prevent physical and logical unidirectional connections and the malfunctioning of other protocols.

A unidirectional link occurs whenever traffic transmitted by the local device over a link is received by the neighbor but traffic transmitted from the neighbor is not received by the local device. If one of the fiber strands in a pair is disconnected, as long as autonegotiation is active, the link does not stay up. In this case, the logical link is undetermined, and UDLD does not take any action. If both fibers are working normally at Layer 1, then UDLD at Layer 2 determines whether those fibers are connected correctly and whether traffic is flowing bidirectionally between the correct neighbors. This check cannot be performed by autonegotiation, because autonegotiation operates at Layer 1.

A Cisco Nexus 5000 Series switch periodically transmits UDLD frames to neighbor devices on LAN ports with UDLD enabled. If the frames are echoed back within a specific time frame and they lack a specific acknowledgment (echo), the link is flagged as unidirectional and the LAN port is shut down. Devices on both ends of the link must support UDLD in order for the protocol to successfully identify and disable unidirectional links.

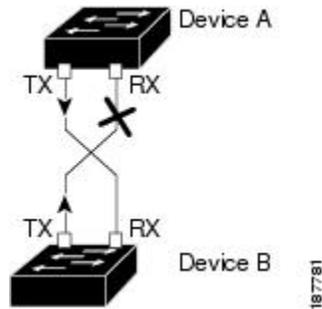


Note

By default, UDLD is locally disabled on copper LAN ports to avoid sending unnecessary control traffic on this type of media.

The following figure shows an example of a unidirectional link condition. Device B successfully receives traffic from Device A on the port. However, Device A does not receive traffic from Device B on the same port. UDLD detects the problem and disables the port.

Figure 6: Unidirectional Link



Default UDLD Configuration

The following table shows the default UDLD configuration.

Table 7: UDLD Default Configuration

Feature	Default Value
UDLD global enable state	Globally disabled
UDLD aggressive mode	Disabled
UDLD per-port enable state for fiber-optic media	Enabled on all Ethernet fiber-optic LAN ports
UDLD per-port enable state for twisted-pair (copper) media	Disabled on all Ethernet 10/100 and 1000BASE-TX LAN ports

Related Topics

[Configuring the UDLD Mode, on page 65](#)

UDLD Aggressive and Nonaggressive Modes

UDLD aggressive mode is disabled by default. You can configure UDLD aggressive mode only on point-to-point links between network devices that support UDLD aggressive mode. If UDLD aggressive mode is enabled, when a port on a bidirectional link that has a UDLD neighbor relationship established stops receiving UDLD frames, UDLD tries to reestablish the connection with the neighbor. After eight failed retries, the port is disabled.

To prevent spanning tree loops, nonaggressive UDLD with the default interval of 15 seconds is fast enough to shut down a unidirectional link before a blocking port transitions to the forwarding state (with default spanning tree parameters).

When you enable the UDLD aggressive mode, the following occurs:

- One side of a link has a port stuck (both transmission and receive)
- One side of a link remains up while the other side of the link is down

In these cases, the UDLD aggressive mode disables one of the ports on the link, which prevents traffic from being discarded.

About Interface Speed

A Cisco Nexus 5000 Series switch has a number of fixed 10-Gigabit ports, each equipped with SFP+ interface adapters. The Cisco Nexus 5010 switch has 20 fixed ports, the first 8 of which are switchable 1-Gigabit and 10-Gigabit ports. The Cisco Nexus 5020 switch has 40 fixed ports, the first 16 of which are switchable 1-Gigabit and 10-Gigabit ports.

About the Cisco Discovery Protocol

The Cisco Discovery Protocol (CDP) is a device discovery protocol that runs over Layer 2 (the data link layer) on all Cisco-manufactured devices (routers, bridges, access servers, and switches) and allows network management applications to discover Cisco devices that are neighbors of already known devices. With CDP, network management applications can learn the device type and the Simple Network Management Protocol (SNMP) agent address of neighboring devices running lower-layer, transparent protocols. This feature enables applications to send SNMP queries to neighboring devices.

CDP runs on all media that support Subnetwork Access Protocol (SNAP). Because CDP runs over the data-link layer only, two systems that support different network-layer protocols can learn about each other.

Each CDP-configured device sends periodic messages to a multicast address, advertising at least one address at which it can receive SNMP messages. The advertisements also contain time-to-live, or holdtime information, which is the length of time a receiving device holds CDP information before discarding it. Each device also listens to the messages sent by other devices to learn about neighboring devices.

The switch supports both CDP Version 1 and Version 2.

Default CDP Configuration

The following table shows the default CDP configuration.

Table 8: Default CDP Configuration

Feature	Default Setting
CDP interface state	Enabled
CDP timer (packet update frequency)	60 seconds
CDP holdtime (before discarding)	180 seconds
CDP Version-2 advertisements	Enabled

About the Debounce Timer Parameters

The port debounce time is the amount of time that an interface waits to notify the supervisor of a link going down. During this time, the interface waits to see if the link comes back up. The wait period is a time when traffic is stopped.

You can enable the debounce timer for each interface and specify the delay time in milliseconds.

**Caution**

When you enable the port debounce timer the link up and link down detections are delayed, resulting in a loss of traffic during the debounce period. This situation might affect the convergence and reconvergence of some protocols.

About MTU Configuration

The Cisco Nexus 5000 Series switch is a Layer 2 device. This means it does not fragment frames. As a result, the switch cannot have two ports in the same Layer 2 domain with different maximum transmission units (MTUs). A per-physical Ethernet interface MTU is not supported. Instead, the MTU is set according to the QoS classes. You modify the MTU by setting Class and Policy maps.

**Note**

When you show the interface settings, a default MTU of 1500 is displayed for physical Ethernet interfaces and a receive data field size of 2112 is displayed for Fibre Channel interfaces.

Configuring Ethernet Interfaces

Configuring the UDLD Mode

You can configure normal or aggressive unidirectional link detection (UDLD) modes for Ethernet interfaces on devices configured to run UDLD. Before you can enable a UDLD mode for an interface, you must make sure that UDLD is already enabled on the device that includes the interface. UDLD must also be enabled on the other linked interface and its device.

To use the normal UDLD mode, you must configure one of the ports for normal mode and configure the other port for the normal or aggressive mode. To use the aggressive UDLD mode, you must configure both ports for the aggressive mode.

**Note**

Before you begin, UDLD must be enabled for the other linked port and its device.

To configure the UDLD mode, perform this task:

SUMMARY STEPS

1. switch# **configure terminal**
2. switch(config)# **feature uddld**
3. switch(config)# **no feature uddld**
4. switch(config)# **show uddld global**
5. switch(config)# **interface type slot/port**
6. switch(config-if)# **uddld {enable | disable | aggressive}**
7. switch(config-if)# **show uddld interface**

DETAILED STEPS

	Command or Action	Purpose
Step 1	switch# configure terminal	Enters configuration mode.
Step 2	switch(config)# feature uddld	Enables UDLD for the device.
Step 3	switch(config)# no feature uddld	Disables UDLD for the device.
Step 4	switch(config)# show uddld global	Displays the UDLD status for the device.
Step 5	switch(config)# interface type slot/port	Specifies an interface to configure, and enters interface configuration mode.
Step 6	switch(config-if)# uddld {enable disable aggressive}	Enables the normal UDLD mode, disables UDLD, or enables the aggressive UDLD mode.
Step 7	switch(config-if)# show uddld interface	Displays the UDLD status for the interface.

This example shows how to enable the UDLD for the switch:

```
switch# configure terminal
switch(config)# feature uddld
```

This example shows how to enable the normal UDLD mode for an Ethernet port:

```
switch# configure terminal
switch(config)# interface ethernet 1/4
switch(config-if)# uddld enable
```

This example shows how to enable the aggressive UDLD mode for an Ethernet port:

```
switch# configure terminal
switch(config)# interface ethernet 1/4
switch(config-if)# uddld aggressive
```

This example shows how to disable UDLD for an Ethernet port:

```
switch# configure terminal
switch(config)# interface ethernet 1/4
switch(config-if)# uddld disable
```

This example shows how to disable UDLD for the switch:

```
switch# configure terminal
switch(config)# no feature udld
```

Configuring Interface Speed

The first 8 ports of a Cisco Nexus 5010 switch and the first 16 ports of a Cisco Nexus 5020 switch are switchable 1-Gigabit and 10-Gigabit ports. The default interface speed is 10-Gigabit. To configure these ports for 1-Gigabit Ethernet, insert a 1-Gigabit Ethernet SFP transceiver into the applicable port and then set its speed with the **speed** command.

To configure a 1-Gigabit Ethernet port, perform this task:

SUMMARY STEPS

1. switch# **configure terminal**
2. switch(config)# **interface** *type slot/port*
3. switch(config-if)# **speed** *speed*

DETAILED STEPS

	Command or Action	Purpose
Step 1	switch# configure terminal	Enters configuration mode.
Step 2	switch(config)# interface <i>type slot/port</i>	Enters interface configuration mode for the specified interface. This interface must have a 1-Gigabit Ethernet SFP transceiver inserted into it.
Step 3	switch(config-if)# speed <i>speed</i>	Sets the speed on the interface.

The following example shows how to set the speed for a 1-Gigabit Ethernet port:

```
switch# configure terminal
switch(config)# interface ethernet 1/4
switch(config-if)# speed 1000
```

This command can only be applied to a physical Ethernet interface.



Note

If the interface and transceiver speed is mismatched, the SFP validation failed message is displayed when you enter the **show interface ethernet slot/port** command. For example, if you insert a 1-Gigabit SFP transceiver into a port without configuring the speed 1000 command, you will get this error. By default, all ports are 10 Gigabits.

Configuring the Cisco Discovery Protocol

Configuring the CDP Characteristics

You can configure the frequency of Cisco Discovery Protocol (CDP) updates, the amount of time to hold the information before discarding it, and whether or not to send Version-2 advertisements.

To configure CDP characteristics for an interface, perform this task:

SUMMARY STEPS

1. switch# **configure terminal**
2. (Optional) switch(config)# **[no] cdp advertise {v1 | v2 }**
3. (Optional) switch(config)# **[no] cdp format device-id {mac-address | serial-number | system-name}**
4. (Optional) switch(config)# **[no] cdp holdtime seconds**
5. (Optional) switch(config)# **[no] cdp timer seconds**

DETAILED STEPS

	Command or Action	Purpose
Step 1	switch# configure terminal	Enters configuration mode.
Step 2	switch(config)# [no] cdp advertise {v1 v2 }	(Optional) Configures the version to use to send CDP advertisements. Version-2 is the default state. Use the no form of the command to return to its default setting.
Step 3	switch(config)# [no] cdp format device-id {mac-address serial-number system-name}	(Optional) Configures the format of the CDP device ID. The default is the system name, which can be expressed as a fully qualified domain name. Use the no form of the command to return to its default setting.
Step 4	switch(config)# [no] cdp holdtime seconds	(Optional) Specifies the amount of time a receiving device should hold the information sent by your device before discarding it. The range is 10 to 255 seconds; the default is 180 seconds. Use the no form of the command to return to its default setting.
Step 5	switch(config)# [no] cdp timer seconds	(Optional) Sets the transmission frequency of CDP updates in seconds. The range is 5 to 254; the default is 60 seconds. Use the no form of the command to return to its default setting.

This example shows how to configure CDP characteristics:

```
switch# configure terminal
switch(config)# cdp timer 50
switch(config)# cdp holdtime 120
switch(config)# cdp advertise v2
```

Enabling or Disabling CDP

You can enable or disable CDP for Ethernet interfaces. This protocol works only when you have it enabled on both interfaces on the same link.

To enable or disable CDP for an interface, perform this task:

SUMMARY STEPS

1. switch# **configure terminal**
2. switch(config)# **interface** *type slot/port*
3. switch(config-if)# **cdp enable**
4. switch(config-if)# **no cdp enable**

DETAILED STEPS

	Command or Action	Purpose
Step 1	switch# configure terminal	Enters configuration mode.
Step 2	switch(config)# interface <i>type slot/port</i>	Enters interface configuration mode for the specified interface.
Step 3	switch(config-if)# cdp enable	Enables CDP for the interface. To work correctly, this parameter must be enabled for both interfaces on the same link.
Step 4	switch(config-if)# no cdp enable	Disables CDP for the interface.

The following example shows how to enable CDP for an Ethernet port:

```
switch# configure terminal
switch(config)# interface ethernet 1/4
switch(config-if)# cdp enable
```

This command can only be applied to a physical Ethernet interface.

Configuring the Debounce Timer

You can enable the debounce timer for Ethernet ports by specifying a debounce time (in milliseconds) or disable the timer by specifying a debounce time of 0.

You can show the debounce times for all of the Ethernet ports by using the **show interface debounce** command.

To enable or disable the debounce timer, perform this task:

SUMMARY STEPS

1. switch# **configure terminal**
2. switch(config)# **interface** *type slot/port*
3. switch(config-if)# **link debounce time** *milliseconds*

DETAILED STEPS

	Command or Action	Purpose
Step 1	switch# configure terminal	Enters configuration mode.
Step 2	switch(config)# interface <i>type slot/port</i>	Enters interface configuration mode for the specified interface.
Step 3	switch(config-if)# link debounce time <i>milliseconds</i>	Enables the debounce timer for the amount of time (1 to 5000 milliseconds) specified. Disables the debounce timer if you specify 0 milliseconds.

This example shows how to enable the debounce timer and set the debounce time to 1000 milliseconds for an Ethernet interface:

```
switch# configure terminal
switch(config)# interface ethernet 1/4
switch(config-if)# link debounce time 1000
```

This example shows how to disable the debounce timer for an Ethernet interface:

```
switch# configure terminal
switch(config)# interface ethernet 1/4
switch(config-if)# link debounce time 0
```

This command can only be applied to a physical Ethernet interface.

Configuring the Description Parameter

To provide textual interface descriptions for the Ethernet ports, perform this task:

SUMMARY STEPS

1. switch# **configure terminal**
2. switch(config)# **interface** *type slot/port*
3. switch(config-if)# **description** *test*

DETAILED STEPS

	Command or Action	Purpose
Step 1	switch# configure terminal	Enters configuration mode.

	Command or Action	Purpose
Step 2	switch(config)# interface <i>type slot/port</i>	Enters interface configuration mode for the specified interface.
Step 3	switch(config-if)# description <i>test</i>	Specifies the description for the interface.

This example shows how to set the interface description to "Server 3 Interface."

```
switch# configure terminal
switch(config)# interface ethernet 1/3
switch(config-if)# description Server 3 Interface
```

Disabling and Restarting Ethernet Interfaces

You can shut down and restart an Ethernet interface. This action disables all of the interface functions and marks the interface as being down on all monitoring displays. This information is communicated to other network servers through all dynamic routing protocols. When shut down, the interface is not included in any routing updates.

To disable an interface, perform this task:

SUMMARY STEPS

1. switch# **configure terminal**
2. switch(config)# **interface** *type slot/port*
3. switch(config-if)# **shutdown**
4. switch(config-if)# **no shutdown**

DETAILED STEPS

	Command or Action	Purpose
Step 1	switch# configure terminal	Enters configuration mode.
Step 2	switch(config)# interface <i>type slot/port</i>	Enters interface configuration mode for the specified interface.
Step 3	switch(config-if)# shutdown	Disables the interface.
Step 4	switch(config-if)# no shutdown	Restarts the interface.

The following example shows how to disable an Ethernet port:

```
switch# configure terminal
switch(config)# interface ethernet 1/4
switch(config-if)# shutdown
```

The following example shows how to restart an Ethernet interface:

```
switch# configure terminal
switch(config)# interface ethernet 1/4
switch(config-if)# no shutdown
```

Displaying Interface Information

To view configuration information about the defined interfaces, perform one of these tasks:

Command	Purpose
switch# show interface <i>type slot/port</i>	Displays the detailed configuration of the specified interface.
switch# show interface <i>type slot/port capabilities</i>	Displays detailed information about the capabilities of the specified interface. This option is only available for physical interfaces
switch# show interface <i>type slot/port transceiver</i>	Displays detailed information about the transceiver connected to the specified interface. This option is only available for physical interfaces.
switch# show interface brief	Displays the status of all interfaces.
switch# show interface debounce	Displays the debounce status of all interfaces.
switch# show interface flowcontrol	Displays the detailed listing of the flow control settings on all interfaces.

The **show interface** command is invoked from EXEC mode and displays the interface configurations. Without any arguments, this command displays the information for all the configured interfaces in the switch.

The following example shows how to display the physical Ethernet interface:

```
switch# show interface ethernet 1/1
Ethernet1/1 is up
  Hardware is 1000/10000 Ethernet, address is 000d.eca3.5f08 (bia 000d.eca3.5f08)
  MTU 1500 bytes, BW 10000000 Kbit, DLY 10 usec,
    reliability 255/255, txload 190/255, rxload 192/255
  Encapsulation ARPA
  Port mode is trunk
  full-duplex, 10 Gb/s, media type is 1/10g
  Input flow-control is off, output flow-control is off
  Auto-mdix is turned on
  Rate mode is dedicated
  Switchport monitor is off
  Last clearing of "show interface" counters never
  5 minute input rate 942201806 bytes/sec, 14721892 packets/sec
  5 minute output rate 935840313 bytes/sec, 14622492 packets/sec
Rx
  129141483840 input packets 0 unicast packets 129141483847 multicast packets
  0 broadcast packets 0 jumbo packets 0 storm suppression packets
  8265054965824 bytes
  0 No buffer 0 runt 0 Overrun
  0 crc 0 Ignored 0 Bad etype drop
  0 Bad proto drop
Tx
  119038487241 output packets 119038487245 multicast packets
  0 broadcast packets 0 jumbo packets
  7618463256471 bytes
  0 output CRC 0 ecc
  0 underrun 0 if down drop      0 output error 0 collision 0 deferred
  0 late collision 0 lost carrier 0 no carrier
  0 babble
  0 Rx pause 8031547972 Tx pause 0 reset
```

The following example shows how to display the physical Ethernet capabilities:

```
switch# show interface ethernet 1/1 capabilities
Ethernet1/1
  Model:                734510033
  Type:                 10Gbase-(unknown)
  Speed:                1000,10000
  Duplex:               full
  Trunk encap. type:   802.1Q
  Channel:              yes
  Broadcast suppression: percentage(0-100)
  Flowcontrol:         rx-(off/on),tx-(off/on)
  Rate mode:           none
  QoS scheduling:      rx-(6q1t),tx-(1p6q0t)
  CoS rewrite:         no
  ToS rewrite:         no
  SPAN:                yes
  UDLD:                yes
  Link Debounce:       yes
  Link Debounce Time:  yes
  MDIX:                no
  FEX Fabric:          yes
```

The following example shows how to display the physical Ethernet transceiver:

```
switch# show interface ethernet 1/1 transceiver
Ethernet1/1
  sfp is present
  name is CISCO-EXCELIGHT
  part number is SPP5101SR-C1
  revision is A
  serial number is ECL120901AV
  nominal bitrate is 10300 Mbits/sec
  Link length supported for 50/125mm fiber is 82 m(s)
  Link length supported for 62.5/125mm fiber is 26 m(s)
  cisco id is --
  cisco extended id number is 4
```

The following example shows how to display a brief interface status (some of the output has been removed for brevity):

```
switch# show interface brief
```

```
-----
Ethernet      VLAN   Type Mode   Status Reason                               Speed   Port
Interface                                           Ch #
-----
Eth1/1        200    eth trunk up      none                               10G(D) --
Eth1/2        1      eth trunk up      none                               10G(D) --
Eth1/3        300    eth access down   SFP not inserted                 10G(D) --
Eth1/4        300    eth access down   SFP not inserted                 10G(D) --
Eth1/5        300    eth access down   Link not connected               1000(D) --
Eth1/6        20     eth access down   Link not connected               10G(D) --
Eth1/7        300    eth access down   SFP not inserted                 10G(D) --
...

```

The following example shows how to display the link debounce status (some of the output has been removed for brevity):

```
switch# show interface debounce
```

```
-----
Port          Debounce time  Value(ms)
-----
...
Eth1/1        enable         100
Eth1/2        enable         100
Eth1/3        enable         100
...

```

The following example shows how to display the CDP neighbors:

```
switch# show cdp neighbors
```

```
Capability Codes: R - Router, T - Trans-Bridge, B - Source-Route-Bridge
                  S - Switch, H - Host, I - IGMP, r - Repeater,
                  V - VoIP-Phone, D - Remotely-Managed-Device,
                  s - Supports-STP-Dispute
```

```
Device ID          Local Intrfce  Hldtme  Capability  Platform      Port ID
d13-dist-1         mgmt0         148     S I         WS-C2960-24TC Fas0/9
n5k(FLC12080012)  Eth1/5        8       S I s       N5K-C5020P-BA Eth1/5
```



Note

From Cisco NX-OS Release 4.0(1a)N1(1), the default value of the device ID field for CDP advertisement has been changed from the chassis serial number to the hostname and serial number, as in the example above.

Default Physical Ethernet Settings

The following table lists the default settings for all physical Ethernet interfaces:

Parameter	Default Setting
Debounce	Enable, 100 milliseconds
Duplex	Auto (full-duplex)
Encapsulation	ARPA

Parameter	Default Setting
MTU ¹	1500 bytes
Port Mode	Access
Speed	Auto (10000)

¹ MTU cannot be changed per-physical Ethernet interface. You modify MTU by selecting maps of QoS classes.



Configuring VLANs

This chapter describes how to configure VLANs on the Cisco Nexus 5000 Series switch. It contains the following sections:

- [Configuring VLANs, page 77](#)

Configuring VLANs

You can use virtual LANs (VLANs) to divide the network into separate logical areas. VLANs can also be considered as broadcast domains.

Any switch port can belong to a VLAN, and unicast, broadcast, and multicast packets are forwarded and flooded only to end stations in that VLAN. Each VLAN is considered a logical network, and packets destined for stations that do not belong to the VLAN must be forwarded through a router.

Information About VLANs

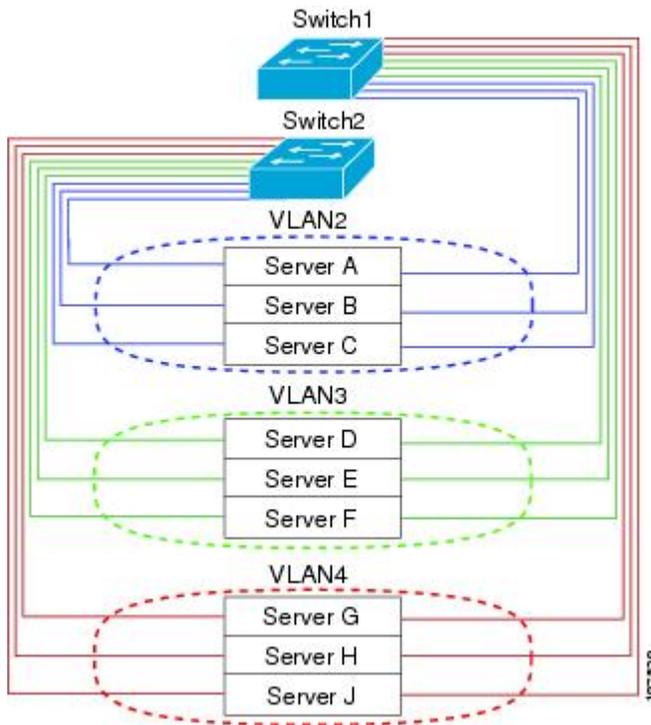
Understanding VLANs

A VLAN is a group of end stations in a switched network that is logically segmented by function, project team, or application, without regard to the physical locations of the users. VLANs have the same attributes as physical LANs, but you can group end stations even if they are not physically located on the same LAN segment.

Any port can belong to a VLAN, and unicast, broadcast, and multicast packets are forwarded and flooded only to end stations in that VLAN. Each VLAN is considered a logical network. Packets destined for stations that do not belong to the VLAN must be forwarded through a router.

The following figure shows VLANs as logical networks. In this diagram, the stations in the engineering department are assigned to one VLAN, the stations in the marketing department are assigned to another VLAN, and the stations in the accounting department are assigned to yet another VLAN.

Figure 7: VLANs as Logically Defined Networks



VLANs are usually associated with IP subnetworks. For example, all the end stations in a particular IP subnet belong to the same VLAN. To communicate between VLANs, you must route the traffic.

By default, a newly created VLAN is operational. To disable the VLAN use the **shutdown** command. Additionally, you can configure VLANs to be in the active state, which is passing traffic, or the suspended state, in which the VLANs are not passing packets. By default, the VLANs are in the active state and pass traffic.



Note The VLAN Trunking Protocol (VTP) mode is OFF. VTP BPDUs are dropped on all interfaces of a Cisco Nexus 5000 Series switch. This has the effect of partitioning VTP domains if other switches have VTP turned on.

Understanding VLAN Ranges

The Cisco Nexus 5000 Series switch supports VLAN numbers 1 to 4094 in accordance with the IEEE 802.1Q standard. These VLANs are organized into ranges. You use each range slightly differently. The switch is physically limited in the number of VLANs it can support. The hardware also shares this available range with its VSANs

The following table describes the details of the VLAN ranges.

Table 9: VLAN Ranges

VLANs Numbers	Range	Usage
1	Normal	Cisco default. You can use this VLAN, but you cannot modify or delete it.
2—1005	Normal	You can create, use, modify, and delete these VLANs.
1006—4094	Extended	You can create, name, and use these VLANs. You cannot change the following parameters: <ul style="list-style-type: none"> • State is always active. • VLAN is always enabled. You cannot shut down these VLANs.
3968—4047 and 4094	Internally allocated	These 80 VLANs, plus VLAN 4094, are allocated for internal use. You cannot create, delete, or modify any VLANs within the block reserved for internal use.

**Note**

VLANs 3968 to 4047 and 4094 are reserved for internal use; these VLANs cannot be changed or used.

Cisco NX-OS allocates a group of 80 VLAN numbers for those features, such as multicast and diagnostics, that need to use internal VLANs for their operation. By default, the system allocates VLANs numbered 3968 to 4047 for internal use. VLAN 4094 is also reserved for internal use by the switch.

You cannot use, modify, or delete any of the VLANs in the reserved group. You can display the VLANs that are allocated internally and their associated use.

Creating, Deleting, and Modifying VLANs

VLANs are numbered from 1 to 4094. All configured ports belong to the default VLAN when you first bring up the switch. The default VLAN (VLAN1) uses only default values. You cannot create, delete, or suspend activity in the default VLAN.

You create a VLAN by assigning a number to it. You can delete VLANs as well as move them from the active operational state to the suspended operational state. If you attempt to create a VLAN with an existing VLAN ID, the switch goes into the VLAN submode but does not create the same VLAN again.

Newly created VLANs remain unused until ports are assigned to the specific VLAN. All the ports are assigned to VLAN1 by default.

Depending on the range of the VLAN, you can configure the following parameters for VLANs (except the default VLAN):

- VLAN name
- Shutdown or not shutdown

When you delete a specified VLAN, the ports associated to that VLAN are shut down and no traffic flows. However, the system retains all the VLAN-to-port mapping for that VLAN, and when you reenables, or recreates, the specified VLAN, the system automatically reinstates all the original ports to that VLAN.

**Note**

Commands entered in the VLAN configuration submode are immediately executed.

VLANs 3968 to 4047 and 4094 are reserved for internal use; these VLANs cannot be changed or used.

Configuring a VLAN

Creating and Deleting a VLAN

You can create or delete all VLANs except the default VLAN and those VLANs that are internally allocated for use by the switch. Once a VLAN is created, it is automatically in the active state.

**Note**

When you delete a VLAN, ports associated to that VLAN shut down. The traffic does not flow and the packets are dropped.

SUMMARY STEPS

1. switch# **configure terminal**
2. switch(config)# **vlan** {vlan-id | vlan-range}
3. switch(config-vlan)# **no vlan** {vlan-id | vlan-range}

DETAILED STEPS

	Command or Action	Purpose
Step 1	switch# configure terminal	Enters configuration mode.
Step 2	switch(config)# vlan {vlan-id vlan-range}	Creates a VLAN or a range of VLANs. If you enter a number that is already assigned to a VLAN, the switch puts you into the VLAN configuration submode for that VLAN. If you enter a number that is assigned to an internally allocated VLAN, the system returns an error message. However, if you enter a range of VLANs and one or more of the specified VLANs is outside the range of internally allocated VLANs, the command takes effect on <i>only</i> those VLANs outside the range. The range is from 2 to 4094; VLAN1 is the default VLAN and cannot be created or deleted. You cannot create or delete those VLANs that are reserved for internal use.

	Command or Action	Purpose
Step 3	switch(config-vlan)# no vlan {vlan-id vlan-range}	Deletes the specified VLAN or range of VLANs and removes you from the VLAN configuration submode. You cannot delete VLAN1 or the internally allocated VLANs.

This example shows how to create a range of VLANs from 15 to 20:

```
switch# configure terminal
switch(config)# vlan 15-20
```



Note You can also create and delete VLANs in the VLAN configuration submode.

Entering the VLAN Submode and Configuring the VLAN

To configure or modify the VLAN for the following parameters, you must be in the VLAN configuration submode:

- Name
- Shut down



Note You cannot create, delete, or modify the default VLAN or the internally allocated VLANs. Additionally, some of these parameters cannot be modified on some VLANs.

SUMMARY STEPS

1. switch# **configure terminal**
2. switch(config)# **vlan** {vlan-id | vlan-range}
3. switch(config-vlan)# **name** vlan-name
4. switch(config-vlan)# **state** {active | suspend}
5. (Optional) switch(config-vlan)# **no shutdown**

DETAILED STEPS

	Command or Action	Purpose
Step 1	switch# configure terminal	Enters configuration mode.
Step 2	switch(config)# vlan {vlan-id vlan-range}	Enters VLAN configuration submode. If the VLAN does not exist, the system first creates the specified VLAN.
Step 3	switch(config-vlan)# name vlan-name	Names the VLAN. You can enter up to 32 alphanumeric characters to name the VLAN. You cannot change the name of VLAN1 or the internally allocated

	Command or Action	Purpose
		VLANs. The default value is VLANxxxx where xxxx represent four numeric digits (including leading zeroes) equal to the VLAN ID number.
Step 4	switch(config-vlan)# state { active suspend }	Sets the state of the VLAN to active or suspend. While the VLAN state is suspended, the ports associated with this VLAN are shut down, and that VLAN does not pass any traffic. The default state is active. You cannot suspend the state for the default VLAN or VLANs 1006 to 4094.
Step 5	switch(config-vlan)# no shutdown	(Optional) Enables the VLAN. The default value is no shutdown (or enabled). You cannot shut down the default VLAN, VLAN1, or VLANs 1006 to 4094.

This example shows how to configure optional parameters for VLAN 5:

```
switch# configure terminal
switch(config)# vlan 5
switch(config-vlan)# name accounting
switch(config-vlan)# state active
switch(config-vlan)# no shutdown
```

Adding Ports to a VLAN

After you have completed the configuration of a VLAN, assign ports to it. To add ports, perform this task:

SUMMARY STEPS

1. switch# **configure terminal**
2. switch(config)# **interface** {**ethernet slot/port** | **port-channel number**}
3. switch(config-if)# **switchport access vlan vlan-id**

DETAILED STEPS

	Command or Action	Purpose
Step 1	switch# configure terminal	Enters configuration mode.
Step 2	switch(config)# interface { ethernet slot/port port-channel number }	Specifies the interface to configure, and enters the interface configuration mode. The interface can be a physical Ethernet port or an EtherChannel.
Step 3	switch(config-if)# switchport access vlan vlan-id	Sets the access mode of the interface to the specified VLAN.

This example shows how to configure an Ethernet interface to join VLAN 5:

```
switch# configure terminal
switch(config)# interface ethernet 1/13
switch(config-if)# switchport access vlan 5
```

Verifying VLAN Configuration

To display VLAN configuration information, perform one of these tasks:

Command	Purpose
switch# show running-config vlan [<i>vlan_id</i> <i>vlan_range</i>]	Displays VLAN information.
switch# show vlan [brief id [<i>vlan_id</i> <i>vlan_range</i>] name <i>name</i> summary]	Displays selected configuration information for the defined VLAN(s).

The following example shows all VLANs defined in the range of 1 to 21.

```
switch# show running-config vlan 1-21
version 4.0(1a)N1(1)
vlan 1
vlan 5
```

The following example shows the VLANs created on the switch and their status:

```
switch# show vlan
```

VLAN Name	Status	Ports
1 default	active	Eth1/1, Eth1/2, Eth1/3, Eth1/4 Eth1/5, Eth1/6, Eth1/7, Eth1/8 Eth1/9, Eth1/10, Eth1/11 Eth1/12, Eth1/15, Eth1/16 Eth1/17, Eth1/18, Eth1/19 Eth1/20, Eth1/21, Eth1/22 Eth1/23, Eth1/24, Eth1/25 Eth1/26, Eth1/27, Eth1/28 Eth1/29, Eth1/30, Eth1/31 Eth1/32, Eth1/33, Eth1/34 Eth1/35, Eth1/36, Eth1/37 Eth1/38, Eth1/39, Eth1/40 Eth3/1, Eth3/2, Eth3/3, Eth3/4 veth1/1
13 VLAN0005	active	Eth1/13, Eth1/14

The following example shows the details of VLAN 13 including its member ports:

```
switch# show vlan id 13
```

VLAN Name	Status	Ports
13 VLAN0005	active	Eth1/13, Eth1/14

```
VLAN Type  MTU
-----
13  enet  576

Remote SPAN VLAN
-----
Disabled
```

Primary	Secondary	Type	Ports
-----	-----	-----	-----

The following example shows the VLAN settings summary:

```
switch# show vlan summary
Number of existing VLANs           : 2
Number of existing user VLANs      : 2
Number of existing extended VLANs  : 0
```



Configuring Private VLANs

This chapter describes how to configure private VLANs on the Cisco Nexus 5000 Series switch. It contains the following sections:

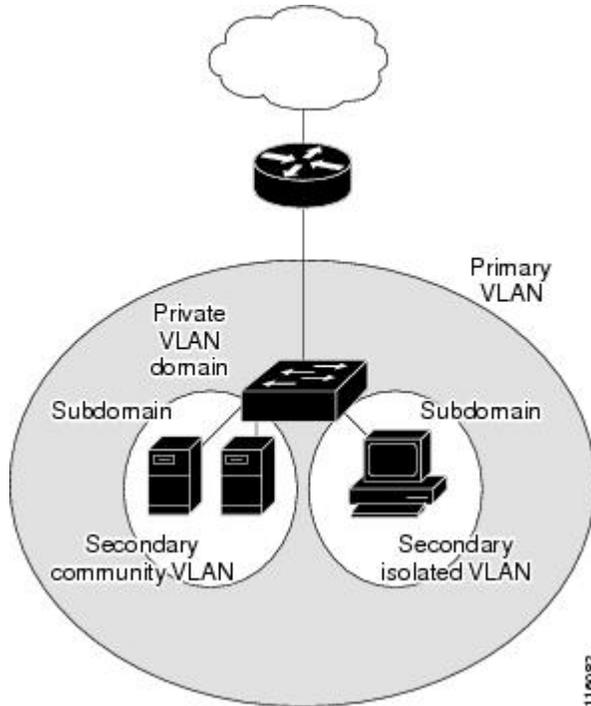
- [Information About Private VLANs, page 85](#)
- [Guidelines and Limitations for Private VLANs, page 90](#)
- [Configuring a Private VLAN, page 90](#)
- [Verifying Private VLAN Configuration, page 99](#)

Information About Private VLANs

A private VLAN partitions the Ethernet broadcast domain of a VLAN into subdomains, allowing you to isolate the ports on the switch from each other. A subdomain consists of a primary VLAN and one or more secondary VLANs (see the following figure). All VLANs in a private VLAN domain share the same primary VLAN. The secondary VLAN ID differentiates one subdomain from another. The secondary VLANs may either be isolated VLANs or community VLANs. A host on an isolated VLAN can only communicate with the associated

promiscuous port in its primary VLAN. Hosts on community VLANs can communicate among themselves and with their associated promiscuous port but not with ports in other community VLANs.

Figure 8: Private VLAN Domain



Note You must first create the VLAN before you can convert it to a private VLAN, either primary or secondary.

Primary and Secondary VLANs in Private VLANs

A private VLAN domain has only one primary VLAN. Each port in a private VLAN domain is a member of the primary VLAN; the primary VLAN is the entire private VLAN domain.

Secondary VLANs provide isolation between ports within the same private VLAN domain. The following two types are secondary VLANs within a primary VLAN:

- Isolated VLANs—Ports within an isolated VLAN cannot communicate directly with each other at the Layer 2 level.
- Community VLANs—Ports within a community VLAN can communicate with each other but cannot communicate with ports in other community VLANs or in any isolated VLANs at the Layer 2 level.

Private VLAN Ports

The three types of private VLAN ports are as follows:

- **Promiscuous**—A promiscuous port belongs to the primary VLAN. The promiscuous port can communicate with all interfaces, including the community and isolated host ports, that belong to those secondary VLANs associated to the promiscuous port and associated with the primary VLAN. You can have several promiscuous ports in a primary VLAN. Each promiscuous port can have several secondary VLANs or no secondary VLANs that are associated to that port. You can associate a secondary VLAN to more than one promiscuous port, as long as the promiscuous port and secondary VLANs are within the same primary VLAN. You may want to do this for load-balancing or redundancy purposes. You can also have secondary VLANs that are not associated to any promiscuous port.

A promiscuous port can be configured either as an access port or as a trunk port.

- **Isolated**—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, except that it can communicate with associated promiscuous ports. Private VLANs 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.

An isolated port can be configured as either an access port or a trunk port.

- **Community**—A community port is a host port that belongs to a community secondary VLAN. Community ports communicate with other ports in the same community VLAN and with associated promiscuous ports. These interfaces are isolated from all other interfaces in other communities and from all isolated ports within the private VLAN domain.

A community port must be configured as an access port. A community VLAN must not be enabled on an isolated trunk.


Note

Because trunks can support the VLANs that carry traffic between promiscuous, isolated, and community ports, the isolated and community port traffic might enter or leave the switch through a trunk interface.

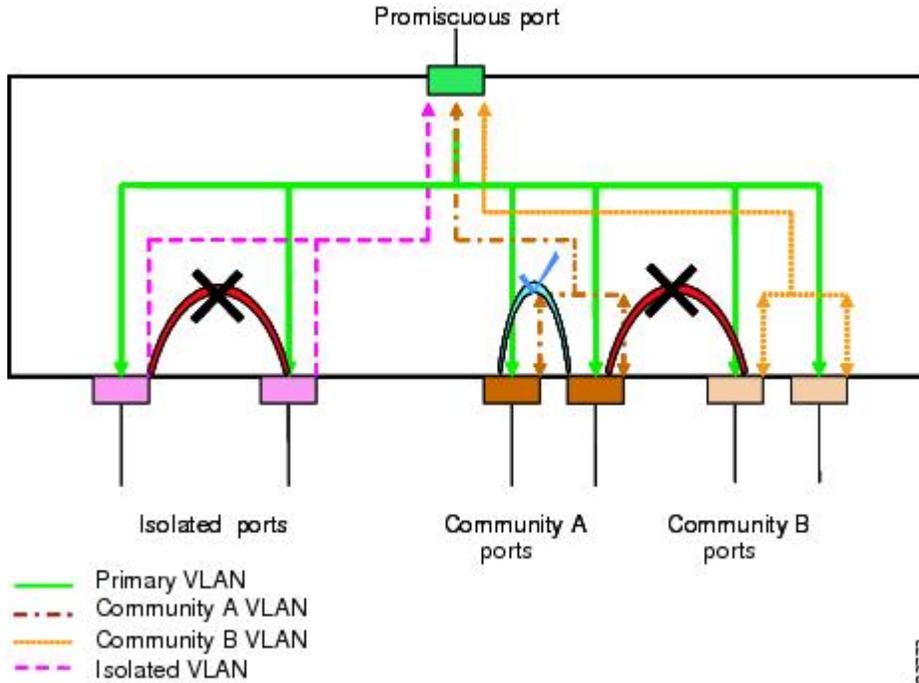
Primary, Isolated, and Community Private VLANs

Primary VLANs and the two types of secondary VLANs (isolated and community) have these characteristics:

- **Primary VLAN**— The primary VLAN carries traffic from the promiscuous ports to the host ports, both isolated and community, and to other promiscuous ports.
- **Isolated VLAN** —An isolated VLAN is a secondary VLAN that carries unidirectional traffic upstream from the hosts toward the promiscuous ports. You can only configure one isolated VLAN in a private VLAN domain. An isolated VLAN can have several isolated ports. The traffic from each isolated port also remains completely separate.
- **Community VLAN**—A community VLAN is a secondary VLAN that carries upstream traffic from the community ports to the promiscuous port and to other host ports in the same community. You can configure multiple community VLANs in a private VLAN domain. The ports within one community can communicate, but these ports cannot communicate with ports in any other community or isolated VLAN in the private VLAN.

The following figure shows the traffic flows within a private VLAN, along with the types of VLANs and types of ports.

Figure 9: Private VLAN Traffic Flows



Note

The private VLAN traffic flows are unidirectional from the host ports to the promiscuous ports. Traffic received on primary VLAN enforces no separation and forwarding is done as in normal VLAN.

A promiscuous access port can serve only one primary VLAN and multiple secondary VLANs (community and isolated VLANs). A promiscuous trunk port can carry traffic for several primary VLANs. Multiple secondary VLANs under a given primary VLAN can be mapped to promiscuous trunk ports. With a promiscuous port, you can connect a wide range of devices as access points to a private VLAN. For example, you can use a promiscuous port to monitor or back up all the private VLAN servers from an administration workstation.

In a switched environment, you can assign an individual private VLAN and associated IP subnet to each individual or common group of end stations. The end stations need to communicate only with a default gateway to communicate outside the private VLAN.

Associating Primary and Secondary VLANs

To allow host ports in secondary VLANs to communicate outside the private VLAN, you associate secondary VLANs to the primary VLAN. If the association is not operational, the host ports (community and isolated ports) in the secondary VLAN are brought down.



Note You can associate a secondary VLAN with only one primary VLAN.

For an association to be operational, the following conditions must be met:

- The primary VLAN must exist and be configured as a primary VLAN.
- The secondary VLAN must exist and be configured as either an isolated or community VLAN.



Note Use the **show vlan private-vlan** command to verify that the association is operational. The switch does not display an error message when the association is nonoperational.

If you delete either the primary or secondary VLAN, the ports that are associated with the VLAN become inactive. Use the **no private-vlan** command to return the VLAN to the normal mode. All primary and secondary associations on that VLAN are suspended, but the interfaces remain in private VLAN mode. When you convert the VLAN back to private VLAN mode, the original associations are reinstated.

If you enter the **no vlan** command for the primary VLAN, all private VLAN associations with that VLAN are deleted. However, if you enter the **no vlan** command for a secondary VLAN, the private VLAN associations with that VLAN are suspended and are restored when you recreate the specified VLAN and configure it as the previous secondary VLAN.

In order to change the association between a secondary and primary VLAN, you must first remove the current association and then add the desired association.

Private VLAN Promiscuous Trunks

A promiscuous trunk port can carry traffic for several primary VLANs. Multiple secondary VLANs under a given primary VLAN can be mapped to promiscuous trunk port. Traffic on the promiscuous port is received and transmitted with a primary VLAN tag.

Private VLAN Isolated Trunks

An isolated trunk port can carry traffic for multiple isolated private VLANs. Traffic for a community VLAN is not carried by isolated trunk ports. Traffic on isolated trunk ports is received and transmitted with an isolated VLAN tag. Isolated trunk ports are intended to be connected to host servers.

To support isolated private VLAN ports on a Cisco Nexus 2000 Series Fabric Extender, the Cisco Nexus 5000 Series switch must prevent communication between the isolated ports on the Fabric Extender; all forwarding occurs through the Cisco Nexus 5000 Series switch.

For unicast traffic, it is simple to prevent such a communication without any side effects.

For multicast traffic, the Fabric Extender provides replication of the frames. To prevent communication between isolated private VLAN ports on the Fabric Extender, the Cisco Nexus 5000 Series switch prevents multicast frames from being sent back through the fabric ports. This restriction prevents communication between an isolated VLAN and a promiscuous port on the Fabric Extender. However as its host interfaces are not intended to be connected to another switch or router, you cannot enable a promiscuous port on Fabric Extender.

Broadcast Traffic in Private VLANs

Broadcast traffic from ports in a private VLAN flows in the following ways:

- The broadcast traffic flows from a promiscuous port to all ports in the primary VLAN (which includes all the ports in the community and isolated VLANs). This broadcast traffic is distributed to all ports within the primary VLAN, including those ports that are not configured with private VLAN parameters.
- The broadcast traffic from an isolated port is distributed only to those promiscuous ports in the primary VLAN that are associated to that isolated port.
- The broadcast traffic from community ports is distributed to all ports within the port's community and to all promiscuous ports that are associated to the community port. The broadcast packets are not distributed to any other communities within the primary VLAN, or to any isolated ports.

Private VLAN Port Isolation

You can use private VLANs to control access to end stations as follows:

- Configure selected interfaces connected to end stations as isolated ports to prevent any communication. For example, if the end stations are servers, this configuration prevents communication between the servers.
- Configure interfaces connected to default gateways and selected end stations (for example, backup servers) as promiscuous ports to allow all end stations access to a default gateway.

Guidelines and Limitations for Private VLANs

When configuring private VLANs, follow these guidelines:

- You must have already created the VLAN before you can assign the specified VLAN as a private VLAN.
- You must enable private VLANs before the switch can apply the private VLAN functionality.
- You cannot disable private VLANs if the switch has any operational ports in a private VLAN mode.
- Enter the **private-vlan synchronize** command to map the secondary VLANs to the same Multiple Spanning Tree (MST) instance as the primary VLAN.

Related Topics

[Mapping Secondary VLANs to Same MSTI as Primary VLANs for Private VLANs](#), on page 197

Configuring a Private VLAN

Enabling Private VLANs

You must enable private VLANs on the switch to use the private VLAN functionality.



Note The private VLAN commands do not appear until you enable the private VLAN feature.

SUMMARY STEPS

1. switch# **configure terminal**
2. switch(config)# **feature private-vlan**
3. (Optional) switch(config)# **no feature private-vlan**

DETAILED STEPS

	Command or Action	Purpose
Step 1	switch# configure terminal	Enters configuration mode.
Step 2	switch(config)# feature private-vlan	Enables the private VLAN feature on the switch.
Step 3	switch(config)# no feature private-vlan	(Optional) Disables the private VLAN feature on the switch. Note You cannot disable private VLANs if there are operational ports on the switch that are in private VLAN mode.

This example shows how to enable the private VLAN feature on the switch:

```
switch# configure terminal
switch(config)# feature private-vlan
```

Configuring a VLAN as a Private VLAN

To create a private VLAN, you first create a VLAN, and then configure that VLAN to be a private VLAN.

Before You Begin

Ensure that the private VLAN feature is enabled.

SUMMARY STEPS

1. switch# **configure terminal**
2. switch(config)# **vlan** {vlan-id | vlan-range}
3. switch(config-vlan)# **private-vlan** {community | isolated | primary}
4. (Optional) switch(config-vlan)# **no private-vlan** {community | isolated | primary}

DETAILED STEPS

	Command or Action	Purpose
Step 1	switch# configure terminal	Enters configuration mode.
Step 2	switch(config)# vlan { <i>vlan-id</i> <i>vlan-range</i> }	Places you into the VLAN configuration submenu.
Step 3	switch(config-vlan)# private-vlan { community isolated primary }	Configures the VLAN as either a community, isolated, or primary private VLAN. In a private VLAN, you must have one primary VLAN. You can have multiple community and isolated VLANs.
Step 4	switch(config-vlan)# no private-vlan { community isolated primary }	(Optional) Removes the private VLAN configuration from the specified VLAN(s) and returns it to normal VLAN mode. If you delete either the primary or secondary VLAN, the ports that are associated with the VLAN become inactive.

This example shows how to assign VLAN 5 to a private VLAN as the primary VLAN:

```
switch# configure terminal
switch(config)# vlan 5
switch(config-vlan)# private-vlan primary
```

This example shows how to assign VLAN 100 to a private VLAN as a community VLAN:

```
switch# configure terminal
switch(config)# vlan 100
switch(config-vlan)# private-vlan community
```

This example shows how to assign VLAN 200 to a private VLAN as an isolated VLAN:

```
switch# configure terminal
switch(config)# vlan 200
switch(config-vlan)# private-vlan isolated
```

Associating Secondary VLANs with a Primary Private VLAN

When you associate secondary VLANs with a primary VLAN, follow these guidelines:

- The *secondary-vlan-list* parameter cannot contain spaces. It can contain multiple comma-separated items. Each item can be a single secondary VLAN ID or a hyphenated range of secondary VLAN IDs.
- The *secondary-vlan-list* parameter can contain multiple community VLAN IDs and one isolated VLAN ID.
- Enter a *secondary-vlan-list* or use the add keyword with a *secondary-vlan-list* to associate secondary VLANs with a primary VLAN.
- Use the remove keyword with a *secondary-vlan-list* to clear the association between secondary VLANs and a primary VLAN.
- You change the association between a secondary and primary VLAN by removing the existing association and then adding the desired association.

If you delete either the primary or secondary VLAN, the VLAN becomes inactive on the port where the association is configured. When you enter the **no private-vlan** command, the VLAN returns to the normal VLAN mode. All primary and secondary associations on that VLAN are suspended, but the interfaces remain in private VLAN mode. If you again convert the specified VLAN to private VLAN mode, the original associations are reinstated.

If you enter the **no vlan** command for the primary VLAN, all private VLAN associations with that VLAN are lost. However, if you enter the **no vlan** command for a secondary VLAN, the private VLAN associations with that VLAN are suspended and are reinstated when you recreate the specified VLAN and configure it as the previous secondary VLAN.

Before You Begin

Ensure that the private VLAN feature is enabled.

SUMMARY STEPS

1. switch# **configure terminal**
2. switch(config)# **vlan primary-vlan-id**
3. switch(config-vlan)# **private-vlan association** {[add] secondary-vlan-list | remove secondary-vlan-list}
4. (Optional) switch(config-vlan)# **no private-vlan association**

DETAILED STEPS

	Command or Action	Purpose
Step 1	switch# configure terminal	Enters configuration mode.
Step 2	switch(config)# vlan primary-vlan-id	Enters the number of the primary VLAN that you are working in for the private VLAN configuration.
Step 3	switch(config-vlan)# private-vlan association {[add] secondary-vlan-list remove secondary-vlan-list}	Associates the secondary VLANs with the primary VLAN.
Step 4	switch(config-vlan)# no private-vlan association	(Optional) Removes all associations from the primary VLAN and returns it to normal VLAN mode.

This example shows how to associate community VLANs 100 through 110 and isolated VLAN 200 with primary VLAN 5:

```
switch# configure terminal
switch(config)# vlan 5
switch(config-vlan)# private-vlan association 100-110, 200
```

Configuring an Interface as a Private VLAN Host Port

In private VLANs, host ports are part of the secondary VLANs, which are either community VLANs or isolated VLANs. Configuring a private VLAN host port involves two steps. First, you define the port as a private VLAN host port and then you configure a host association between the primary and secondary VLANs.

**Note**

We recommend that you enable BPDU Guard on all interfaces configured as a host ports.

Before You Begin

Ensure that the private VLAN feature is enabled.

SUMMARY STEPS

1. switch# **configure terminal**
2. switch(config)# **interface** *type* [*chassis*/]*slot/port*
3. switch(config-if)# **switchport mode private-vlan host**
4. switch(config-if)# **switchport private-vlan host-association** {*primary-vlan-id*} {*secondary-vlan-id*}
5. (Optional) switch(config-if)# **no switchport private-vlan host-association**

DETAILED STEPS

	Command or Action	Purpose
Step 1	switch# configure terminal	Enters configuration mode.
Step 2	switch(config)# interface <i>type</i> [<i>chassis</i> /] <i>slot/port</i>	Selects the port to configure as a private VLAN host port. This port can be on a Fabric Extender (identified by the chassis option).
Step 3	switch(config-if)# switchport mode private-vlan host	Configures the port as a host port for a private VLAN.
Step 4	switch(config-if)# switchport private-vlan host-association { <i>primary-vlan-id</i> } { <i>secondary-vlan-id</i> }	Associates the port with the primary and secondary VLANs of a private VLAN. The secondary VLAN can be either an isolated or community VLAN.
Step 5	switch(config-if)# no switchport private-vlan host-association	(Optional) Removes the private VLAN association from the port.

This example shows how to configure Ethernet port 1/12 as a host port for a private VLAN and associate it to primary VLAN 5 and secondary VLAN 101:

```
switch# configure terminal
switch(config)# interface ethernet 1/12
switch(config-if)# switchport mode private-vlan host
switch(config-if)# switchport private-vlan host-association 5 101
```

Configuring an Interface as a Private VLAN Promiscuous Port

In a private VLAN domain, promiscuous ports are part of the primary VLAN. Configuring a promiscuous port involves two steps. First, you define the port as a promiscuous port and then you configure the mapping between a secondary VLAN and the primary VLAN.

Before You Begin

Ensure that the private VLAN feature is enabled.

SUMMARY STEPS

1. switch# **configure terminal**
2. switch(config)# **interface** *type slot/port*
3. switch(config-if)# **switchport mode private-vlan promiscuous**
4. switch(config-if)# **switchport private-vlan mapping** *{primary-vlan-id}* *{secondary-vlan-list | add secondary-vlan-list | remove secondary-vlan-list}*
5. (Optional) switch(config-if)# **no switchport private-vlan mapping**

DETAILED STEPS

	Command or Action	Purpose
Step 1	switch# configure terminal	Enters configuration mode.
Step 2	switch(config)# interface <i>type slot/port</i>	Selects the port to configure as a private VLAN promiscuous port. A physical interface is required. This port cannot be on a Fabric Extender.
Step 3	switch(config-if)# switchport mode private-vlan promiscuous	Configures the port as a promiscuous port for a private VLAN. You can only enable a physical Ethernet port as the promiscuous port.
Step 4	switch(config-if)# switchport private-vlan mapping <i>{primary-vlan-id}</i> <i>{secondary-vlan-list add secondary-vlan-list remove secondary-vlan-list}</i>	Configures the port as a promiscuous port and associates the specified port with a primary VLAN and a selected list of secondary VLANs. The secondary VLAN can be either an isolated or community VLAN.
Step 5	switch(config-if)# no switchport private-vlan mapping	(Optional) Clears the mapping from the private VLAN.

This example shows how to configure Ethernet interface 1/4 as a promiscuous port associated with primary VLAN 5 and secondary isolated VLAN 200:

```
switch# configure terminal
switch(config)# interface ethernet 1/4
switch(config-if)# switchport mode private-vlan promiscuous
switch(config-if)# switchport private-vlan mapping 5 200
```

Configuring a Promiscuous Trunk Port

In a private VLAN domain, promiscuous trunks are part of the primary VLAN. Promiscuous trunk ports can carry multiple primary VLANs. Multiple secondary VLANs under a given primary VLAN can be mapped to a promiscuous trunk port.

Configuring a promiscuous port involves two steps. First, you define the port as a promiscuous port and then you configure the mapping between a secondary VLAN and the primary VLAN. Multiple primary VLANs can be enabled by configuring multiple mappings.

Before You Begin

Ensure that the private VLAN feature is enabled.

SUMMARY STEPS

1. switch# **configure terminal**
2. switch(config)# **interface** *type slot/port*
3. switch(config-if)# **switchport mode private-vlan trunk promiscuous**
4. switch(config-if)# **switchport private-vlan mapping trunk** {*primary-vlan-id*} {*secondary-vlan-id*}
5. (Optional) switch(config-if)# **no switchport private-vlan mapping trunk** [*primary-vlan-id*]

DETAILED STEPS

	Command or Action	Purpose
Step 1	switch# configure terminal	Enters configuration mode.
Step 2	switch(config)# interface <i>type slot/port</i>	Selects the port to configure as a private VLAN promiscuous trunk port.
Step 3	switch(config-if)# switchport mode private-vlan trunk promiscuous	Configures the port as a promiscuous trunk port for a private VLAN. You can only enable a physical Ethernet port as the promiscuous port. This port cannot be on a Fabric Extender.
Step 4	switch(config-if)# switchport private-vlan mapping trunk { <i>primary-vlan-id</i> } { <i>secondary-vlan-id</i> }	Maps the trunk port with the primary and secondary VLANs of a private VLAN. The secondary VLAN can be either an isolated or community VLAN.
Step 5	switch(config-if)# no switchport private-vlan mapping trunk [<i>primary-vlan-id</i>]	(Optional) Removes the private VLAN mapping from the port. If the <i>primary-vlan-id</i> is not supplied, all private VLAN mappings are removed from the port.

This example shows how to configure Ethernet interface 1/1 as a promiscuous trunk port for a private VLAN and then map the secondary VLANs to the primary VLAN:

```
switch# configure terminal
switch(config)# interface ethernet 1/1
switch(config-if)# switchport mode private-vlan trunk promiscuous
switch(config-if)# switchport private-vlan mapping trunk 5 100
switch(config-if)# switchport private-vlan mapping trunk 5 200
switch(config-if)# switchport private-vlan mapping trunk 6 300
```

Configuring an Isolated Trunk Port

In a private VLAN domain, isolated trunks are part of a secondary VLAN. Isolated trunk ports can carry multiple isolated VLANs. Only one isolated VLAN under a given primary VLAN can be associated to an isolated trunk port. Configuring an isolated trunk port involves two steps. First, you define the port as an isolated trunk port and then you configure the association between the isolated and primary VLANs. Multiple isolated VLANs can be enabled by configuring multiple associations.

Before You Begin

Ensure that the private VLAN feature is enabled.

SUMMARY STEPS

1. switch# **configure terminal**
2. switch(config)# **interface** *type* [*chassis*]/*slot*/*port*
3. switch(config-if)# **switchport mode private-vlan trunk** [**secondary**]
4. switch(config-if)# **switchport private-vlan association trunk** {*primary-vlan-id*} {*secondary-vlan-id*}
5. (Optional) switch(config-if)# **no switchport private-vlan association trunk** [*primary-vlan-id*]

DETAILED STEPS

	Command or Action	Purpose
Step 1	switch# configure terminal	Enters configuration mode.
Step 2	switch(config)# interface <i>type</i> [<i>chassis</i>]/ <i>slot</i> / <i>port</i>	Selects the port to configure as a private VLAN isolated trunk port. This port can be on a Fabric Extender (identified by the <i>chassis</i> option).
Step 3	switch(config-if)# switchport mode private-vlan trunk [secondary]	Configures the port as a secondary trunk port for a private VLAN. Note The secondary keyword is assumed if it is not present.
Step 4	switch(config-if)# switchport private-vlan association trunk { <i>primary-vlan-id</i> } { <i>secondary-vlan-id</i> }	Associates the isolated trunk port with the primary and secondary VLANs of a private VLAN. The secondary VLAN should be an isolated VLAN. Only one isolated VLAN can be mapped under a given primary VLAN.
Step 5	switch(config-if)# no switchport private-vlan association trunk [<i>primary-vlan-id</i>]	(Optional) Removes the private VLAN association from the port. If the <i>primary-vlan-id</i> is not supplied, all private VLAN associations are removed from the port.

This example shows how to configure Ethernet interface 1/1 as a promiscuous trunk port for a private VLAN and then map the secondary VLANs to the primary VLAN:

```
switch# configure terminal
switch(config)# interface ethernet 1/1
switch(config-if)# switchport mode private-vlan trunk secondary
switch(config-if)# switchport private-vlan association 5 100
switch(config-if)# switchport private-vlan association 6 200
```

Configuring the Allowed VLANs for PVLAN Trunking Ports

Isolated trunk and promiscuous trunk ports can carry traffic from regular VLANs along with private VLANs.

Before You Begin

Ensure that the private VLAN feature is enabled.

SUMMARY STEPS

1. switch# **configure terminal**
2. switch(config)# **interface** *type* [*chassis*]/*slot*/*port*
3. switch(config-if)# **switchport private-vlan trunk allowed vlan** {*vlan-list* | **all** | **none** [**add** | **except** | **none** | **remove** {*vlan-list*}]}

DETAILED STEPS

	Command or Action	Purpose
Step 1	switch# configure terminal	Enters configuration mode.
Step 2	switch(config)# interface <i>type</i> [<i>chassis</i>]/ <i>slot</i> / <i>port</i>	Selects the port to configure as a private VLAN host port. This port can be on a Fabric Extender (identified by the chassis option).
Step 3	switch(config-if)# switchport private-vlan trunk allowed vlan { <i>vlan-list</i> all none [add except none remove { <i>vlan-list</i> }]}	<p>Sets the allowed VLANs for the private trunk interface. The default is to allow only mapped/associated VLANs on the private VLAN trunk interface.</p> <p>Note The primary VLANs do not need to be explicitly added to the allowed VLAN list. They are added automatically once there is a mapping between primary and secondary VLANs.</p>

This example shows how to add VLANs to the list of allowed VLANs on an Ethernet private VLAN trunk port:

```
switch# configure terminal
switch(config)# interface ethernet 1/3
switch(config-if)# switchport private-vlan trunk allowed vlan 15-20
```

Configuring Native 802.1Q VLANs on Private VLANs

Typically, you configure 802.1Q trunks with a native VLAN ID, which strips tagging from all packets on that VLAN. This configuration allows untagged traffic and control traffic to transit the Cisco Nexus 5000 Series switch. Secondary VLANs cannot be configured with a native VLAN ID on promiscuous trunk ports. Primary VLANs cannot be configured with a native VLAN ID on isolated trunk ports.



Note

A trunk can carry the traffic of multiple VLANs. Traffic belonging to the native VLAN is not encapsulated to transit the trunk. Traffic for other VLANs is encapsulated with tags which identify the VLAN the traffic belongs to.

Before You Begin

Ensure that the private VLAN feature is enabled.

SUMMARY STEPS

1. switch# **configure terminal**
2. switch(config)# **interface** *type* [*chassis*]/*slot*/*port*
3. switch(config-if)# **switchport private-vlan trunk native** {*vlan vlan-id*}
4. (Optional) switch(config-if)# **no switchport private-vlan trunk native** {*vlan vlan-id*}

DETAILED STEPS

	Command or Action	Purpose
Step 1	switch# configure terminal	Enters configuration mode.
Step 2	switch(config)# interface <i>type</i> [<i>chassis</i>]/ <i>slot</i> / <i>port</i>	Selects the port to configure as a private VLAN host port. This port can be on a Fabric Extender (identified by the chassis option).
Step 3	switch(config-if)# switchport private-vlan trunk native { <i>vlan vlan-id</i> }	Sets the native VLAN ID for the private VLAN trunk. The default is VLAN 1.
Step 4	switch(config-if)# no switchport private-vlan trunk native { <i>vlan vlan-id</i> }	(Optional) Removes the native VLAN ID from the private VLAN trunk.

Verifying Private VLAN Configuration

To display private VLAN configuration information, use the following commands:

Command	Purpose
switch# show feature	Displays the features enabled on the switch.

Command	Purpose
switch# show interface switchport	Displays information on all interfaces configured as switch ports.
switch# show vlan private-vlan [type]	Displays the status of the private VLAN.

The following example shows how to display the private VLAN configuration:

```
switch# show vlan private-vlan
Primary Secondary Type Ports
-----
5 100 community
5 101 community Eth1/12, Eth100/1/1
5 102 community
5 110 community
5 200 isolated Eth1/2
switch# show vlan private-vlan type
Vlan Type
-----
5 primary
100 community
101 community
102 community
110 community
200 isolated
```

The following example shows how to display enabled features (some of the output has been removed for brevity):

```
switch# show feature
Feature Name Instance State
-----
fcsp 1 enabled
...
interface-vlan 1 enabled
private-vlan 1 enabled
udld 1 disabled
...
```



Configuring Access and Trunk Interfaces

Ethernet interfaces can be configured either as access ports or trunk ports. Trunks carry the traffic of multiple VLANs over a single link and allow you to extend VLANs across the network.



Note

Cisco NX-OS supports only IEEE 802.1Q-type VLAN trunk encapsulation.

This chapter describes the configuration of access or trunk ports on Cisco Nexus 5000 Series switches. It includes the following sections:

- [Information About Access and Trunk Interfaces, page 101](#)
- [Configuring Access and Trunk Interfaces, page 105](#)
- [Verifying Interface Configuration, page 110](#)

Information About Access and Trunk Interfaces

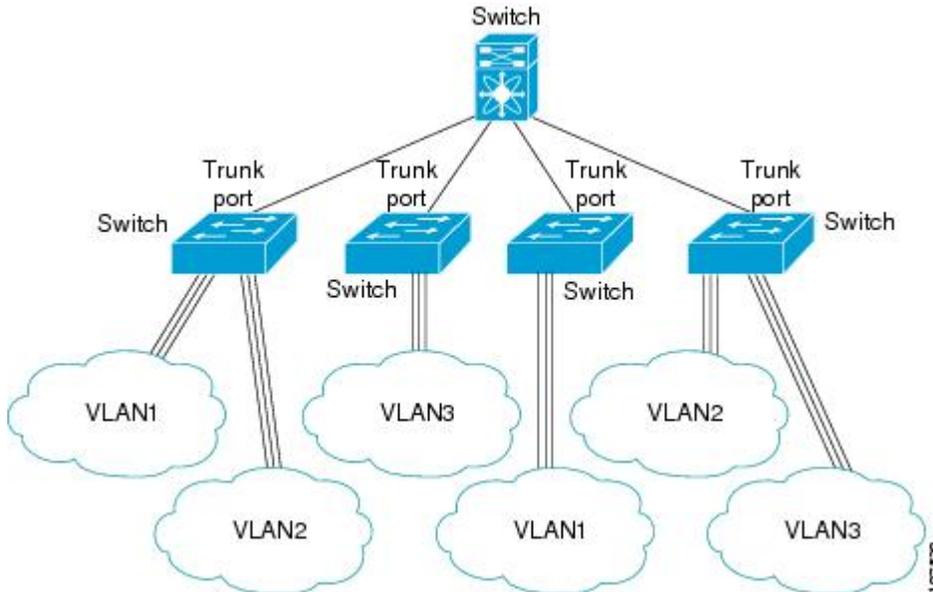
Understanding Access and Trunk Interfaces

Ethernet interfaces can be configured either as access ports or a trunk ports, as follows:

- An access port can have only one VLAN configured on the interface; it can carry traffic for only one VLAN.
- A trunk port can have two or more VLANs configured on the interface; it can carry traffic for several VLANs simultaneously.

The following figure shows how you can use trunk ports in the network. The trunk port carries traffic for two or more VLANs.

Figure 10: Devices in a Trunking Environment



In order to correctly deliver the traffic on a trunk port with several VLANs, the device uses the IEEE 802.1Q encapsulation or tagging method.

To optimize the performance on access ports, you can configure the port as a host port. Once the port is configured as a host port, it is automatically set as an access port, and channel grouping is disabled. Use the host designation to decrease the time it takes the designated port to begin to forward packets.



Note Only an end station can be set as a host port; you will receive an error message if you attempt to configure other ports as hosts.

If an access port receives a packet with an 802.1Q tag in the header other than the access VLAN value, that port drops the packet without learning its MAC source address.



Note An Ethernet interface can function as either an access port or a trunk port; it cannot function as both port types simultaneously.

Related Topics

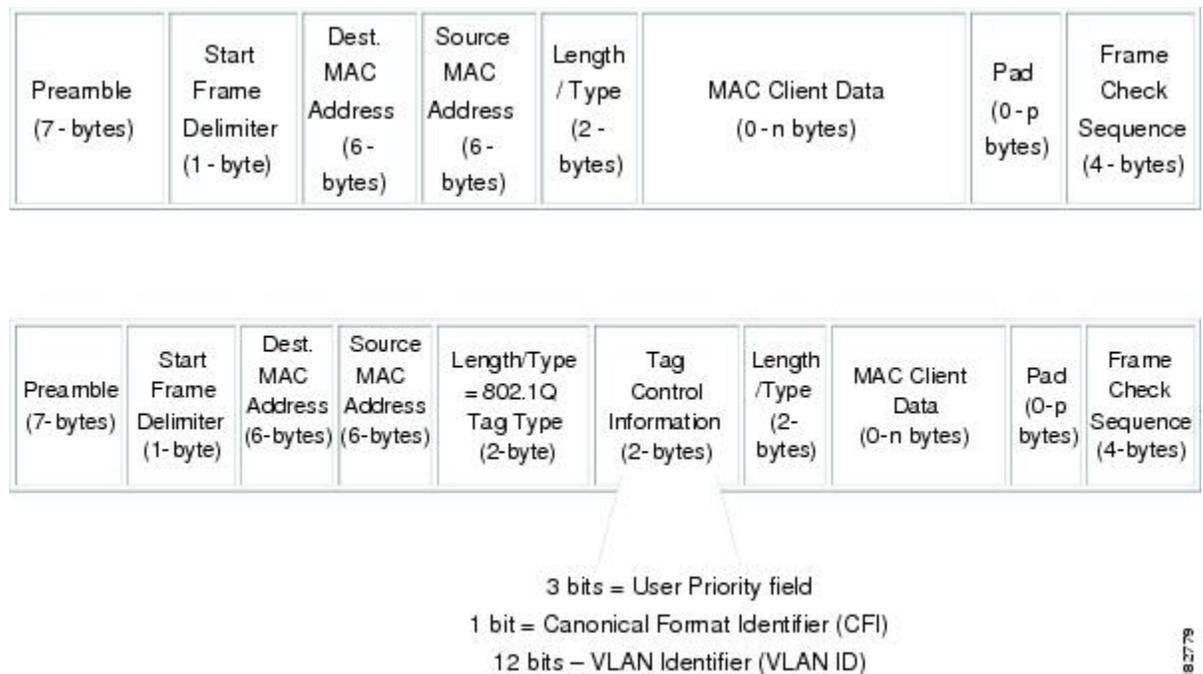
[Understanding IEEE 802.1Q Encapsulation, on page 103](#)

Understanding IEEE 802.1Q Encapsulation

A trunk is a point-to-point link between the device and another networking device. Trunks carry the traffic of multiple VLANs over a single link and allow you to extend VLANs across an entire network.

To correctly deliver the traffic on a trunk port with several VLANs, the device uses the IEEE 802.1Q encapsulation (tagging) method that uses a tag that is inserted into the frame header. This tag carries information about the specific VLAN to which the frame and packet belong. This method allows packets that are encapsulated for several different VLANs to traverse the same port and maintain traffic separation between the VLANs. The encapsulated VLAN tag also allows the trunk to move traffic end-to-end through the network on the same VLAN.

Figure 11: Header without and with 802.1Q Tag Included



Understanding Access VLANs

When you configure a port in access mode, you can specify which VLAN will carry the traffic for that interface. If you do not configure the VLAN for a port in access mode, or an access port, the interface carries traffic for the default VLAN (VLAN1).

You can change the access port membership in a VLAN by specifying the new VLAN. You must create the VLAN before you can assign it as an access VLAN for an access port. If you change the access VLAN on an access port to a VLAN that is not yet created, the system will shut that access port down.

If an access port receives a packet with an 802.1Q tag in the header other than the access VLAN value, that port drops the packet without learning its MAC source address.

**Note**

If you assign an access VLAN that is also a primary VLAN for a private VLAN, all access ports with that access VLAN will also receive all the broadcast traffic for the primary VLAN in the private VLAN mode.

Understanding the Native VLAN ID for Trunk Ports

A trunk port can carry untagged packets simultaneously with the 802.1Q tagged packets. When you assign a default port VLAN ID to the trunk port, all untagged traffic travels on the default port VLAN ID for the trunk port, and all untagged traffic is assumed to belong to this VLAN. This VLAN is referred to as the native VLAN ID for a trunk port. The native VLAN ID is the VLAN that carries untagged traffic on trunk ports.

The trunk port sends an egressing packet with a VLAN that is equal to the default port VLAN ID as untagged; all the other egressing packets are tagged by the trunk port. If you do not configure a native VLAN ID, the trunk port uses the default VLAN.

**Note**

Native VLAN ID numbers *must* match on both ends of the trunk.

Understanding Allowed VLANs

By default, a trunk port sends traffic to and receives traffic from all VLANs. All VLAN IDs are allowed on each trunk. However, you can remove VLANs from this inclusive list to prevent traffic from the specified VLANs from passing over the trunk. You can add any specific VLANs later that you may want the trunk to carry traffic for back to the list.

To partition spanning tree protocol (STP) topology for the default VLAN, you can remove VLAN1 from the list of allowed VLANs. Otherwise, VLAN1, which is enabled on all ports by default, will have a very big STP topology, which can result in problems during STP convergence. When you remove VLAN1, all data traffic for VLAN1 on this port is blocked, but the control traffic continues to move on the port.

Understanding Native 802.1Q VLANs

To provide additional security for traffic passing through an 802.1Q trunk port, the **vlan dot1q tag native** command was introduced. This feature provides a means to ensure that all packets going out of a 802.1Q trunk port are tagged and to prevent reception of untagged packets on the 802.1Q trunk port.

Without this feature, all tagged ingress frames received on a 802.1Q trunk port are accepted as long as they fall inside the allowed VLAN list and their tags are preserved. Untagged frames are tagged with the native VLAN ID of the trunk port before further processing. Only those egress frames whose VLAN tags are inside the allowed range for that 802.1Q trunk port are received. If the VLAN tag on a frame happens to match that of the native VLAN on the trunk port, the tag is stripped off and the frame is sent untagged.

This behavior will cause "VLAN hopping" in which it is possible for frames to jump to a different VLAN. It is also possible for traffic to become part of the native VLAN by sending untagged packets into an 802.1Q trunk port.

To address the above issues, the **vlan dot1q tag native** command performs the following functions:

- On the ingress side, all untagged data traffic is dropped.

- On the egress side, all traffic is tagged. If traffic belongs to native VLAN then it is tagged with the native VLAN ID.

This feature is supported on all the directly connected Ethernet and EtherChannel interfaces of the Cisco Nexus 5000 Series switch. It is also supported on all the host interface ports of any attached Cisco Nexus 2000 Series Fabric Extender.



Note The `vlan dot1q tag native` command is enabled on global basis.

Configuring Access and Trunk Interfaces

Configuring a LAN Interface as an Ethernet Access Port

You can configure an Ethernet interface as an access port. An access port transmits packets on only one, untagged VLAN. You specify which VLAN traffic that the interface carries. If you do not specify a VLAN for an access port, the interface carries traffic only on the default VLAN. The default VLAN is VLAN1.

The VLAN must exist before you can specify that VLAN as an access VLAN. The system shuts down an access port that is assigned to an access VLAN that does not exist.

SUMMARY STEPS

1. `switch# configure terminal`
2. `switch(config)# interface {{type slot/port} | {port-channel number}}`
3. `switch(config-if)# switchport mode {access | trunk}`
4. `switch(config-if)# switchport access vlan vlan-id`

DETAILED STEPS

	Command or Action	Purpose
Step 1	<code>switch# configure terminal</code>	Enters configuration mode.
Step 2	<code>switch(config)# interface {{type slot/port} {port-channel number}}</code>	Specifies an interface to configure, and enters interface configuration mode.
Step 3	<code>switch(config-if)# switchport mode {access trunk}</code>	Sets the interface as a nontrunking nontagged single-VLAN Ethernet interface. An access port can carry traffic in one VLAN only. By default, an access port carries traffic for VLAN1; to set the access port to carry traffic for a different VLAN, use the switchport access vlan command.
Step 4	<code>switch(config-if)# switchport access vlan vlan-id</code>	Specifies the VLAN for which this access port will carry traffic. If you do not enter this command, the access port carries traffic on VLAN1 only; use this command to change the VLAN for which the access port carries traffic.

This example shows how to set an interface as an Ethernet access port that carries traffic for a specific VLAN only:

```
switch# configure terminal
switch(config)# interface ethernet 1/10
switch(config-if)# switchport mode access
switch(config-if)# switchport access vlan 5
```

Configuring Access Host Ports

You can optimize performance on access ports that are connected to end stations by simultaneously setting that port as an access port. An access host port handles the Spanning Tree Protocol (STP) like an edge port and immediately moves to the forwarding state without passing through the blocking and learning states. Configuring an interface as an access host port also disables EtherChannel on that interface.

Before You Begin

Ensure that you are configuring the correct interface; it must be an interface that is connected to an end station.

SUMMARY STEPS

1. switch# **configure terminal**
2. switch(config)# **interface type slot/port**
3. switch(config-if)# **switchport host**

DETAILED STEPS

	Command or Action	Purpose
Step 1	switch# configure terminal	Enters configuration mode.
Step 2	switch(config)# interface type slot/port	Specifies an interface to configure, and enters interface configuration mode.
Step 3	switch(config-if)# switchport host	Sets the interface to be an access host port, which immediately moves to the spanning tree forwarding state and disables port channeling on this interface. Note Apply this command only to end stations.

This example shows how to set an interface as an Ethernet access host port with EtherChannel disabled:

```
switch# configure terminal
switch(config)# interface ethernet 1/10
switch(config-if)# switchport host
```

Configuring Trunk Ports

You can configure an Ethernet port as a trunk port; a trunk port transmits untagged packets for the native VLAN plus encapsulated, tagged, packets for multiple VLANs.



Note Cisco NX-OS supports only 802.1Q encapsulation.

To configure a trunk port, perform this task:

SUMMARY STEPS

1. switch# **configure terminal**
2. switch(config)# **interface** {*type slot/port* | **port-channel number**}
3. switch(config-if)# **switchport mode** {**access** | **trunk**}

DETAILED STEPS

	Command or Action	Purpose
Step 1	switch# configure terminal	Enters configuration mode.
Step 2	switch(config)# interface { <i>type slot/port</i> port-channel number }	Specifies an interface to configure, and enters interface configuration mode.
Step 3	switch(config-if)# switchport mode { access trunk }	Sets the interface as an Ethernet trunk port. A trunk port can carry traffic in one or more VLANs on the same physical link (VLANs are based on the trunk-allowed VLANs list). By default, a trunk interface can carry traffic for all VLANs. To specify that only certain VLANs are allowed on the specified trunk, use the switchport trunk allowed vlan command.

This example shows how to set an interface as an Ethernet trunk port:

```
switch# configure terminal
switch(config)# interface ethernet 1/3
switch(config-if)# switchport mode trunk
```

Related Topics

[Understanding IEEE 802.1Q Encapsulation, on page 103](#)

Configuring the Native VLAN for 802.1Q Trunking Ports

If you do not configure this parameter, the trunk port uses the default VLAN as the native VLAN ID.

SUMMARY STEPS

1. switch# **configure terminal**
2. switch(config)# **interface** {*type slot/port* | **port-channel number**}
3. switch(config-if)# **switchport trunk native vlan** *vlan-id*

DETAILED STEPS

	Command or Action	Purpose
Step 1	switch# configure terminal	Enters configuration mode.
Step 2	switch(config)# interface { <i>type slot/port</i> port-channel number }	Specifies an interface to configure, and enters interface configuration mode.
Step 3	switch(config-if)# switchport trunk native vlan <i>vlan-id</i>	Sets the native VLAN for the 802.1Q trunk. Valid values are from 1 to 4094, except those VLANs reserved for internal use. The default value is VLAN1.

This example shows how to set the native VLAN for an Ethernet trunk port:

```
switch# configure terminal
switch(config)# interface ethernet 1/3
switch(config-if)# switchport trunk native vlan 5
```

Configuring the Allowed VLANs for Trunking Ports

You can specify the IDs for the VLANs that are allowed on the specific trunk port.

Before you configure the allowed VLANs for the specified trunk ports, ensure that you are configuring the correct interfaces and that the interfaces are trunks.

SUMMARY STEPS

1. switch# **configure terminal**
2. switch(config)# **interface** {*type slot/port* | **port-channel number**}
3. switch(config-if)# **switchport trunk allowed vlan** {*vlan-list all* | **none** [**add** |**except** | **none** | **remove** {*vlan-list*}]}

DETAILED STEPS

	Command or Action	Purpose
Step 1	switch# configure terminal	Enters configuration mode.

	Command or Action	Purpose
Step 2	switch(config)# interface { <i>type slot/port</i> port-channel <i>number</i> }	Specifies an interface to configure, and enters interface configuration mode.
Step 3	switch(config-if)# switchport trunk allowed vlan { <i>vlan-list all</i> none [add except none remove { <i>vlan-list</i> }]}	<p>Sets allowed VLANs for the trunk interface. The default is to allow all VLANs on the trunk interface: 1 to 3967 and 4048 to 4094. VLANs 3968 to 4047 are the default VLANs reserved for internal use by default; this group of VLANs is configurable. By default, all VLANs are allowed on all trunk interfaces.</p> <p>Note You cannot add internally allocated VLANs as allowed VLANs on trunk ports. The system returns a message if you attempt to list an internally allocated VLAN as an allowed VLAN.</p>

This example shows how to add VLANs to the list of allowed VLANs on an Ethernet trunk port:

```
switch# configure terminal
switch(config)# interface ethernet 1/3
switch(config-if)# switchport trunk allow vlan 15-20
```

Configuring Native 802.1Q VLANs

Typically, you configure 802.1Q trunks with a native VLAN ID, which strips tagging from all packets on that VLAN. This configuration allows all untagged traffic and control traffic to transit the Cisco Nexus 5000 Series switch. Packets that enter the switch with 802.1Q tags that match the native VLAN ID value are similarly stripped of tagging.

To maintain the tagging on the native VLAN and drop untagged traffic, enter the **vlan dot1q tag native** command. The switch will tag the traffic received on the native VLAN and admit only 802.1Q-tagged frames, dropping any untagged traffic, including untagged traffic in the native VLAN.

Control traffic continues to be accepted untagged on the native VLAN on a trunked port, even when the **vlan dot1q tag native** command is enabled.



Note The **vlan dot1q tag native** command is enabled on global basis.

SUMMARY STEPS

1. switch# **configure terminal**
2. switch(config)# **vlan dot1q tag native**
3. (Optional) switch(config)# **no vlan dot1q tag native**
4. (Optional) switch# **show vlan dot1q tag native**

DETAILED STEPS

	Command or Action	Purpose
Step 1	switch# configure terminal	Enters configuration mode.
Step 2	switch(config)# vlan dot1q tag native	Enables dot1q (IEEE 802.1Q) tagging for all native VLANs on all trunked ports on the Cisco Nexus 5000 Series switch. By default, this feature is disabled.
Step 3	switch(config)# no vlan dot1q tag native	(Optional) Disables dot1q (IEEE 802.1Q) tagging for all native VLANs on all trunked ports on the switch.
Step 4	switch# show vlan dot1q tag native	(Optional) Displays the status of tagging on the native VLANs.

The following example shows how to enable 802.1Q tagging on the switch:

```
switch# configure terminal
switch(config)# vlan dot1q tag native
switch(config)# exit
switch# show vlan dot1q tag native
vlan dot1q native tag is enabled
```

Verifying Interface Configuration

To display access and trunk interface configuration information, perform one of these tasks:

Command	Purpose
switch# show interface	Displays the interface configuration
switch# show interface switchport	Displays information for all Ethernet interfaces, including access and trunk interfaces.
switch# show interface brief	Displays interface configuration information.



Configuring EtherChannels

This chapter describes how to configure EtherChannels and to apply and configure the Link Aggregation Control Protocol (LACP) for more efficient use of EtherChannels in Cisco NX-OS. It contains the following sections:

- [Information About EtherChannels, page 111](#)
- [Configuring EtherChannels, page 118](#)
- [Verifying EtherChannel Configuration, page 124](#)

Information About EtherChannels

An EtherChannel bundles up to 16 individual interfaces into a group to provide increased bandwidth and redundancy. Port channeling also load balances traffic across these physical interfaces. The EtherChannel stays operational as long as at least one physical interface within the EtherChannel is operational.

You create an EtherChannel by bundling compatible interfaces. You can configure and run either static EtherChannels or EtherChannels running the Link Aggregation Control Protocol (LACP).

Any configuration changes that you apply to the EtherChannel are applied to each member interface of that EtherChannel. For example, if you configure Spanning Tree Protocol (STP) parameters on the EtherChannel, the Cisco NX-OS applies those parameters to each interface in the EtherChannel.

You can use static EtherChannels, with no associated protocol, for a simplified configuration. For more efficient use of the EtherChannel, you can use the Link Aggregation Control Protocol (LACP), which is defined in IEEE 802.3ad. When you use LACP, the link passes protocol packets.

Related Topics

[LACP Overview, on page 115](#)

Understanding EtherChannels

Using EtherChannels, Cisco NX-OS provides wider bandwidth, redundancy, and load balancing across the channels.

You can collect up to 16 ports into a static EtherChannel or you can enable the Link Aggregation Control Protocol (LACP). Configuring EtherChannels with LACP requires slightly different steps than configuring static EtherChannels.



Note Cisco NX-OS does not support Port Aggregation Protocol (PAgP) for EtherChannels.

An EtherChannel bundles individual links into a channel group to create a single logical link that provides the aggregate bandwidth of up to 16 physical links. If a member port within an EtherChannel fails, traffic previously carried over the failed link switches to the remaining member ports within the EtherChannel.

Each port can be in only one EtherChannel. All the ports in an EtherChannel must be compatible; they must use the same speed and operate in full-duplex mode. When you are running static EtherChannels, without LACP, the individual links are all in the on channel mode; you cannot change this mode without enabling LACP.



Note You cannot change the mode from ON to Active or from ON to Passive.

You can create an EtherChannel directly by creating the port-channel interface, or you can create a channel group that acts to aggregate individual ports into a bundle. When you associate an interface with a channel group, Cisco NX-OS creates a matching EtherChannel automatically if the EtherChannel does not already exist. You can also create the EtherChannel first. In this instance, Cisco NX-OS creates an empty channel group with the same channel number as the EtherChannel and takes the default configuration.



Note An EtherChannel is operationally up when at least one of the member ports is up and that port's status is channeling. The EtherChannel is operationally down when all member ports are operationally down.

Compatibility Requirements

When you add an interface to a channel group, Cisco NX-OS checks certain interface attributes to ensure that the interface is compatible with the channel group. Cisco NX-OS also checks a number of operational attributes for an interface before allowing that interface to participate in the port-channel aggregation.

The compatibility check includes the following operational attributes:

- Port mode
- Access VLAN
- Trunk native VLAN
- Allowed VLAN list
- Speed
- 802.3x flow control setting
- MTU

The Cisco Nexus 5000 Series switch only supports system level MTU. This attribute cannot be changed on an individual port basis.

- Broadcast/Unicast/Multicast Storm Control setting
- Priority-Flow-Control
- Untagged CoS

Use the **show port-channel compatibility-parameters** command to see the full list of compatibility checks that Cisco NX-OS uses.

You can only add interfaces configured with the channel mode set to on to static EtherChannels. You can also only add interfaces configured with the channel mode as active or passive to EtherChannels that are running LACP. You can configure these attributes on an individual member port.

When the interface joins an EtherChannel, the following individual parameters are replaced with the values on the EtherChannel:

- Bandwidth
- MAC address
- Spanning Tree Protocol

The following interface parameters remain unaffected when the interface joins an EtherChannel:

- Description
- CDP
- LACP port priority
- Debounce

Related Topics

[Channel Modes](#), on page 116

Load Balancing Using EtherChannels

Cisco NX-OS load balances traffic across all operational interfaces in an EtherChannel by reducing part of the binary pattern formed from the addresses in the frame to a numerical value that selects one of the links in the channel. EtherChannels provide load balancing by default and the basic configuration uses the following criteria to select the link:

- For a Layer 2 frame, it uses the source and destination MAC addresses.
- For a Layer 3 frame, it uses the source and destination MAC addresses and the source and destination IP addresses.
- For a Layer 4 frame, it uses the source and destination MAC addresses, the source and destination IP addresses, and the source and destination port number.

You can configure the switch to use one of the following methods to load balance across the EtherChannel:

- Destination MAC address
- Source MAC address
- Source and destination MAC address

- Destination IP address
- Source IP address
- Source and destination IP address
- Destination TCP/UDP port number
- Source TCP/UDP port number
- Source and destination TCP/UDP port number

The following table shows the criteria used for each configuration:

Table 10: EtherChannel Load-Balancing Criteria

Configuration	Layer 2 Criteria	Layer 3 Criteria	Layer 4 Criteria
Destination MAC	Destination MAC	Destination MAC	Destination MAC
Source MAC	Source MAC	Source MAC	Source MAC
Source and destination MAC	Source and destination MAC	Source and destination MAC	Source and destination MAC
Destination IP	Destination MAC	Destination MAC, destination IP	Destination MAC, destination IP
Source IP	Source MAC	Source MAC, source IP	Source MAC, source IP
Source and destination IP	Source and destination MAC	Source and destination MAC, source and destination IP	Source and destination MAC, source and destination IP
Destination TCP/UDP port	Destination MAC	Destination MAC, destination IP	Destination MAC, destination IP, destination port
Source TCP/UDP port	Source MAC	Source MAC, source IP	Source MAC, source IP, source port
Source and destination TCP/UDP port	Source and destination MAC	Source and destination MAC, source and destination IP	Source and destination MAC, source and destination IP, source and destination port

Use the option that provides the balance criteria with the greatest variety in your configuration. For example, if the traffic on an EtherChannel is going only to a single MAC address and you use the destination MAC address as the basis of port-channel load balancing, the EtherChannel always chooses the same link in that EtherChannel; using source addresses or IP addresses might result in better load balancing.

Understanding LACP

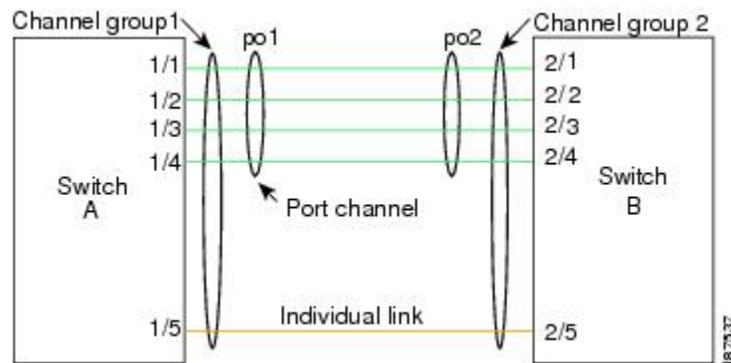
LACP Overview



Note You must enable the LACP feature before you can configure and use LACP functions.

The following figure shows how individual links can be combined into LACP EtherChannels and channel groups as well as function as individual links.

Figure 12: Individual Links Combined into an EtherChannel



With LACP, you can bundle up to 16 interfaces in a channel group.



Note When you delete the EtherChannel, Cisco NX-OS automatically deletes the associated channel group. All member interfaces revert to their previous configuration.

You cannot disable LACP while any LACP configurations are present.

LACP ID Parameters

LACP uses the following parameters:

- LACP system priority—Each system that runs LACP has an LACP system priority value. You can accept the default value of 32768 for this parameter, or you can configure a value between 1 and 65535. LACP uses the system priority with the MAC address to form the system ID and also uses the system priority during negotiation with other devices. A higher system priority value means a lower priority.



Note The LACP system ID is the combination of the LACP system priority value and the MAC address.

- LACP port priority—Each port configured to use LACP has an LACP port priority. You can accept the default value of 32768 for the LACP port priority, or you can configure a value between 1 and 65535. LACP uses the port priority with the port number to form the port identifier. LACP uses the port priority

to decide which ports should be put in standby mode when there is a limitation that prevents all compatible ports from aggregating and which ports should be put into active mode. A higher port priority value means a lower priority for LACP. You can configure the port priority so that specified ports have a lower priority for LACP and are most likely to be chosen as active links, rather than hot-standby links.

- LACP administrative key—LACP automatically configures an administrative key value equal to the channel-group number on each port configured to use LACP. The administrative key defines the ability of a port to aggregate with other ports. A port's ability to aggregate with other ports is determined by these factors:
 - Port physical characteristics, such as the data rate, the duplex capability, and the point-to-point or shared medium state
 - Configuration restrictions that you establish

Channel Modes

Individual interfaces in EtherChannels are configured with channel modes. When you run static EtherChannels, with no protocol, the channel mode is always set to on. After you enable LACP globally on the device, you enable LACP for each channel by setting the channel mode for each interface to active or passive. You can configure either channel mode for individual links in the LACP channel group.



Note

You must enable LACP globally before you can configure an interface in either the active or passive channel mode.

The following table describes the channel modes.

Table 11: Channel Modes for Individual Links in an EtherChannel

Channel Mode	Description
passive	LACP mode that places a port into a passive negotiating state, in which the port responds to LACP packets that it receives but does not initiate LACP negotiation.
active	LACP mode that places a port into an active negotiating state, in which the port initiates negotiations with other ports by sending LACP packets.

Channel Mode	Description
on	<p>All static EtherChannels, that is, that are not running LACP, remain in this mode. If you attempt to change the channel mode to active or passive before enabling LACP, the device returns an error message.</p> <p>You enable LACP on each channel by configuring the interface in that channel for the channel mode as either active or passive. When an LACP attempts to negotiate with an interface in the on state, it does not receive any LACP packets and becomes an individual link with that interface; it does not join the LACP channel group.</p>

Both the passive and active modes allow LACP to negotiate between ports to determine if they can form an EtherChannel, based on criteria such as the port speed and the trunking state. The passive mode is useful when you do not know whether the remote system, or partner, supports LACP.

Ports can form an LACP EtherChannel when they are in different LACP modes as long as the modes are compatible as in the following examples:

- A port in active mode can form an EtherChannel successfully with another port that is in active mode.
- A port in active mode can form an EtherChannel with another port in passive mode.
- A port in passive mode cannot form an EtherChannel with another port that is also in passive mode because neither port will initiate negotiation.
- A port in on mode is not running LACP.

LACP Marker Responders

Using EtherChannels, data traffic may be dynamically redistributed due to either a link failure or load balancing. LACP uses the Marker Protocol to ensure that frames are not duplicated or reordered because of this redistribution. Cisco NX-OS supports only Marker Responders.

LACP-Enabled and Static EtherChannels Differences

The following table provides a brief summary of major differences between EtherChannels with LACP enabled and static EtherChannels.

Table 12: EtherChannels with LACP Enabled and Static EtherChannels

Configurations	EtherChannels with LACP Enabled	Static EtherChannels
Protocol applied	Enable globally.	Not applicable.
Channel mode of links	Can be either: <ul style="list-style-type: none"> • Active • Passive 	Can only be On.

Configurations	EtherChannels with LACP Enabled	Static EtherChannels
Maximum number of links in channel	16	16

Configuring EtherChannels

Creating an EtherChannel

You can create an EtherChannel before creating a channel group. Cisco NX-OS automatically creates the associated channel group.



Note

If you want LACP-based EtherChannels, you need to enable LACP.

SUMMARY STEPS

1. switch# **configure terminal**
2. switch(config)# **interface port-channel** *channel-number*
3. switch(config)# **no interface port-channel** *channel-number*

DETAILED STEPS

	Command or Action	Purpose
Step 1	switch# configure terminal	Enters configuration mode.
Step 2	switch(config)# interface port-channel <i>channel-number</i>	Specifies the port-channel interface to configure, and enters the interface configuration mode. The range is from 1 to 4096. Cisco NX-OS automatically creates the channel group if it does not already exist.
Step 3	switch(config)# no interface port-channel <i>channel-number</i>	Removes the EtherChannel and deletes the associated channel group.

This example shows how to create an EtherChannel:

```
switch# configure terminal
switch (config)# interface port-channel 1
```

Adding a Port to an EtherChannel

You can add a port to a new channel group or to a channel group that already contains ports. Cisco NX-OS creates the EtherChannel associated with this channel group if the EtherChannel does not already exist.



Note If you want LACP-based EtherChannels, you need to enable LACP.

SUMMARY STEPS

1. switch# **configure terminal**
2. switch(config)# **interface** *type slot/port*
3. (Optional) switch(config-if)# **switchport mode trunk**
4. (Optional) switch(config-if)# **switchport trunk** {**allowed vlan** *vlan-id* | **native vlan** *vlan-id*}
5. switch(config-if)# **channel-group** *channel-number*
6. (Optional) switch(config-if)# **no channel-group**

DETAILED STEPS

	Command or Action	Purpose
Step 1	switch# configure terminal	Enters configuration mode.
Step 2	switch(config)# interface <i>type slot/port</i>	Specifies the interface that you want to add to a channel group and enters the interface configuration mode.
Step 3	switch(config-if)# switchport mode trunk	(Optional) Configures the interface as a trunk port.
Step 4	switch(config-if)# switchport trunk { allowed vlan <i>vlan-id</i> native vlan <i>vlan-id</i> }	(Optional) Configures necessary parameters for a trunk port.
Step 5	switch(config-if)# channel-group <i>channel-number</i>	Configures the port in a channel group and sets the mode. The channel-number range is from 1 to 4096. Cisco NX-OS creates the EtherChannel associated with this channel group if the EtherChannel does not already exist. This is called implicit EtherChannel creation.
Step 6	switch(config-if)# no channel-group	(Optional) Removes the port from the channel group. The port reverts to its original configuration.

This example shows how to add an Ethernet interface 1/4 to channel group 1:

```
switch# configure terminal
switch (config)# interface ethernet 1/4
switch(config-if)# switchport mode trunk
switch(config-if)# channel-group 1
```

Related Topics

[Enabling LACP, on page 121](#)

Configuring Load Balancing Using EtherChannels

You can configure the load-balancing algorithm for EtherChannels that applies to the entire device.


Note

If you want LACP-based EtherChannels, you need to enable LACP.

SUMMARY STEPS

1. switch# **configure terminal**
2. switch(config)# **port-channel load-balance ethernet {destination-ip | destination-mac | destination-port | source-dest-ip | source-dest-mac | source-dest-port | source-ip | source-mac | source-port}**
3. (Optional) switch(config)# **no port-channel load-balance ethernet**
4. (Optional) switch(config-router)# **show port-channel load-balance**

DETAILED STEPS

	Command or Action	Purpose
Step 1	switch# configure terminal	Enters configuration mode.
Step 2	switch(config)# port-channel load-balance ethernet {destination-ip destination-mac destination-port source-dest-ip source-dest-mac source-dest-port source-ip source-mac source-port}	Specifies the load-balancing algorithm for the device. The range depends on the device. The default is source-dest-mac.
Step 3	switch(config)# no port-channel load-balance ethernet	(Optional) Restores the default load-balancing algorithm of source-dest-mac.
Step 4	switch(config-router)# show port-channel load-balance	(Optional) Displays the port-channel load-balancing algorithm.

This example shows how to configure source IP load balancing for EtherChannels:

```
switch# configure terminal
switch (config)# port-channel load-balance ethernet source-ip
```


Note

Before Release 4.0(1a)N1 of Cisco NX-OS, the source-dest-ip, source-dest-mac, and source-dest-port keywords were source-destination-ip, source-destination-mac, and source-destination-port, respectively.

Related Topics

[Enabling LACP, on page 121](#)

Enabling LACP

LACP is disabled by default; you must enable LACP before you begin LACP configuration. You cannot disable LACP while any LACP configuration is present.

LACP learns the capabilities of LAN port groups dynamically and informs the other LAN ports. Once LACP identifies correctly matched Ethernet links, it facilitates grouping the links into an EtherChannel. The EtherChannel is then added to the spanning tree as a single bridge port.

SUMMARY STEPS

1. switch# **configure terminal**
2. switch(config)# **feature lacp**
3. (Optional) switch(config)# **show feature**

DETAILED STEPS

	Command or Action	Purpose
Step 1	switch# configure terminal	Enters configuration mode.
Step 2	switch(config)# feature lacp	Enables LACP on the switch.
Step 3	switch(config)# show feature	(Optional) Displays enabled features.

This example shows how to enable LACP:

```
switch# configure terminal
switch(config)# feature lacp
```

Configuring Channel Mode for a Port

You can configure the channel mode for each individual link in the LACP EtherChannel as active or passive. This channel configuration mode allows the link to operate with LACP.

When you configure EtherChannels with no associated protocol, all interfaces on both sides of the link remain in the on channel mode.

Before You Begin

Ensure that you have enabled the LACP feature.

SUMMARY STEPS

1. switch# **configure terminal**
2. switch(config)# **interface type slot/port**
3. switch(config-if)# **channel-group number mode {active | on | passive}**
4. switch(config-if)# **no channel-group number mode**

DETAILED STEPS

	Command or Action	Purpose
Step 1	switch# configure terminal	Enters configuration mode.
Step 2	switch(config)# interface <i>type slot/port</i>	Specifies the interface to configure, and enters the interface configuration mode.
Step 3	switch(config-if)# channel-group <i>number</i> mode { active on passive }	Specifies the port mode for the link in an EtherChannel. After LACP is enabled, you configure each link or the entire channel as active or passive. When you run EtherChannels with no associated protocol, the channel mode is always on. The default channel mode is on.
Step 4	switch(config-if)# no channel-group <i>number</i> mode	Returns the port mode to on for the specified interface.

This example shows how to set the LACP-enabled interface to active port-channel mode for Ethernet interface 1/4 in channel group 5:

```
switch# configure terminal
switch (config)# interface ethernet 1/4
switch(config-if)# channel-group 5 mode active
```

Configuring the LACP System Priority and System ID

The LACP system ID is the combination of the LACP system priority value and the MAC address.

Before You Begin

Ensure that you have enabled the LACP feature.

SUMMARY STEPS

1. switch# **configure terminal**
2. switch(config)# **lACP system-priority** *priority*
3. (Optional) switch# **show lACP system-identifier**

DETAILED STEPS

	Command or Action	Purpose
Step 1	switch# configure terminal	Enters configuration mode.

	Command or Action	Purpose
Step 2	switch(config)# lACP system-priority <i>priority</i>	Configures the system priority for use with LACP. Valid values are 1 through 65535, and higher numbers have lower priority. The default value is 32768.
Step 3	switch# show lACP system-identifier	(Optional) Displays the LACP system identifier.

This example shows how to set the LACP system priority to 2500:

```
switch# configure terminal
switch(config)# lACP system-priority 2500
```

Configuring the LACP Port Priority

You can configure each link in the LACP EtherChannel for the port priority.

Before You Begin

Ensure that you have enabled the LACP feature.

SUMMARY STEPS

1. switch# **configure terminal**
2. switch(config)# **interface** *type slot/port*
3. switch(config-if)# **lACP port-priority** *priority*

DETAILED STEPS

	Command or Action	Purpose
Step 1	switch# configure terminal	Enters configuration mode.
Step 2	switch(config)# interface <i>type slot/port</i>	Specifies the interface to configure, and enters the interface configuration mode.
Step 3	switch(config-if)# lACP port-priority <i>priority</i>	Configures the port priority for use with LACP. Valid values are 1 through 65535, and higher numbers have lower priority. The default value is 32768.

This example shows how to set the LACP port priority for Ethernet interface 1/4 to 40000:

```
switch# configure terminal
switch (config)# interface ethernet 1/4
switch(config-if)# lACP port priority 40000
```

Verifying EtherChannel Configuration

To display EtherChannel configuration information, perform one of the following tasks:

Command	Purpose
switch# show interface port-channel <i>channel-number</i>	Displays the status of a EtherChannel interface.
switch# show feature	Displays enabled features.
switch# show resource	Displays the number of resources currently available in the system.
switch# show lacp {counters interface type <i>slot/port</i> neighbor port-channel system-identifier}	Displays LACP information.
switch# show port-channel compatibility-parameters	Displays the parameters that must be the same among the member ports in order to join an EtherChannel.
switch# show port-channel database [interface port-channel <i>channel-number</i>]	Displays the aggregation state for one or more port-channel interfaces.
switch# show port-channel summary	Displays a summary for the EtherChannel interfaces.
switch# show port-channel traffic	Displays the traffic statistics for EtherChannels.
switch# show port-channel usage	Displays the range of used and unused channel numbers.
switch# show port-channel database	Displays information on current running of the EtherChannel feature.
switch# show port-channel load-balance	Displays information about load-balancing using EtherChannels.



CHAPTER 11

Configuring Virtual Port Channels

This chapter describes how to configure virtual port channels (vPCs) on Cisco Nexus 5000 Series switches. It contains the following sections:

- [Information About vPCs, page 125](#)
- [vPC Guidelines and Limitations, page 137](#)
- [Configuring vPCs, page 138](#)
- [Verifying the vPC Configuration, page 149](#)
- [vPC Example Configurations, page 150](#)
- [vPC Default Settings, page 154](#)

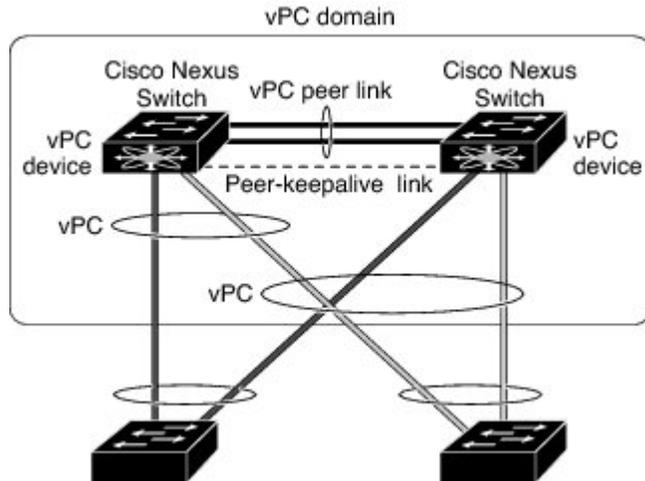
Information About vPCs

vPC Overview

A virtual port channel (vPC) allows links that are physically connected to two different Cisco Nexus 5000 Series switches or Cisco Nexus 2000 Series Fabric Extenders to appear as a single port channel by a third device (see the following figure). The third device can be a switch, server, or any other networking device. Beginning with Cisco NX-OS Release 4.1(3)N1(1), you can configure vPCs in topologies that include Cisco Nexus 5000 Series switches connected to the Fabric Extender. A vPC can provide multipathing, which allows

you to create redundancy by enabling multiple parallel paths between nodes and load balancing traffic where alternative paths exist.

Figure 13: vPC Architecture



You configure the EtherChannels by using one of the following:

- No protocol
- Link Aggregation Control Protocol (LACP)

When you configure the EtherChannels in a vPC—including the vPC peer link channel—each switch can have up to 16 active links in a single EtherChannel. When you configure a vPC on a Fabric Extender, only one port is allowed in an EtherChannel.



Note

You must enable the vPC feature before you can configure or run the vPC functionality.

To enable the vPC functionality, you must create a peer-keepalive link under a vPC domain to send heartbeat messages between the two vPC peer devices.

To create a vPC peer link you configure an EtherChannel on one Cisco Nexus 5000 Series switch by using two or more Ethernet ports. On the other switch, you configure another EtherChannel again using two or more Ethernet ports. Connecting these two EtherChannels together creates a vPC peer link.



Note

We recommend that you configure the vPC peer-link EtherChannels as trunks.

The vPC domain includes both vPC peer devices, the vPC peer-keepalive link, the vPC peer link, and all of the EtherChannels in the vPC domain connected to the downstream device. You can have only one vPC domain ID on each vPC peer device.



Note

Always attach all vPC devices using EtherChannels to both vPC peer devices.

A vPC provides the following benefits:

- Allows a single device to use an EtherChannel across two upstream devices
- Eliminates Spanning Tree Protocol (STP) blocked ports
- Provides a loop-free topology
- Uses all available uplink bandwidth
- Provides fast convergence if either the link or a switch fails
- Provides link-level resiliency
- Assures high availability

Terminology

vPC Terminology

The terminology used in vPCs is as follows:

- vPC—The combined EtherChannel between the vPC peer devices and the downstream device.
- vPC peer device—One of a pair of devices that are connected with the special EtherChannel known as the vPC peer link.
- vPC peer link—The link used to synchronize states between the vPC peer devices.
- vPC member port—Interfaces that belong to the vPCs.
- Host vPC port—Fabric Extender host interfaces that belong to a vPC.
- vPC domain—This domain includes both vPC peer devices, the vPC peer-keepalive link, and all of the port channels in the vPC connected to the downstream devices. It is also associated to the configuration mode that you must use to assign vPC global parameters. The vPC domain ID must be the same on both switches.
- vPC peer-keepalive link—The peer-keepalive link monitors the vitality of a vPC peer Cisco Nexus 5000 Series device. The peer-keepalive link sends configurable, periodic keepalive messages between vPC peer devices.

No data or synchronization traffic moves over the vPC peer-keepalive link; the only traffic on this link is a message that indicates that the originating switch is operating and running vPCs.

Fabric Extender Terminology

The terminology used for the Cisco Nexus 2000 Series Fabric Extender is as follows:

- Fabric interface—A 10-Gigabit Ethernet uplink port designated for connection from the Fabric Extender to its parent switch. A fabric interface cannot be used for any other purpose. It must be directly connected to the parent switch.
- EtherChannel fabric interface—An EtherChannel uplink connection from the Fabric Extender to its parent switch. This connection consists of fabric interfaces bundled into a single logical channel.

- Host interface—An Ethernet interface for server or host connectivity. These ports are 1-Gigabit Ethernet interfaces.
- EtherChannel host interface—An EtherChannel downlink connection from the Fabric Extender host interface to a server port.



Note In Release 4.1(3)N1(1), an EtherChannel host interface consists of only one host interface and can be configured either as a Link Aggregation Control Protocol (LACP) or non-LACP EtherChannel.

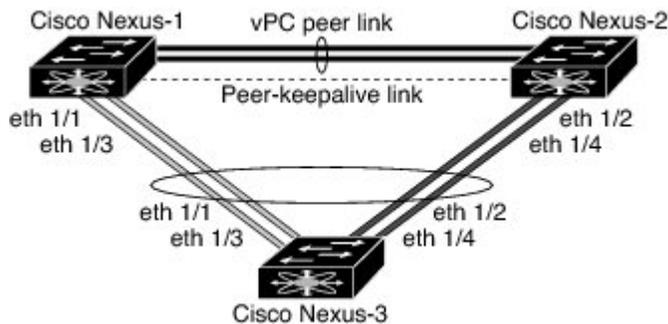
For further information about the Fabric Extender, refer to the *Cisco Nexus 2000 Series Fabric Extender Software Configuration Guide*.

Supported vPC Topologies

Cisco Nexus 5000 Series Switch vPC Topology

You can connect a pair of Cisco Nexus 5000 Series switches configured in a vPC directly to another switch or to a server. Up to 8 interfaces could be connected to each Cisco Nexus 5000 Series switch providing 16 interfaces bundled for the vPC pair. The topology that is shown in the following figure provides the vPC functionality to dual connected switches or servers with 10-Gigabit or 1-Gigabit Ethernet uplink interfaces.

Figure 14: Cisco Nexus 5000 Series Switch-to-Switch vPC Topology



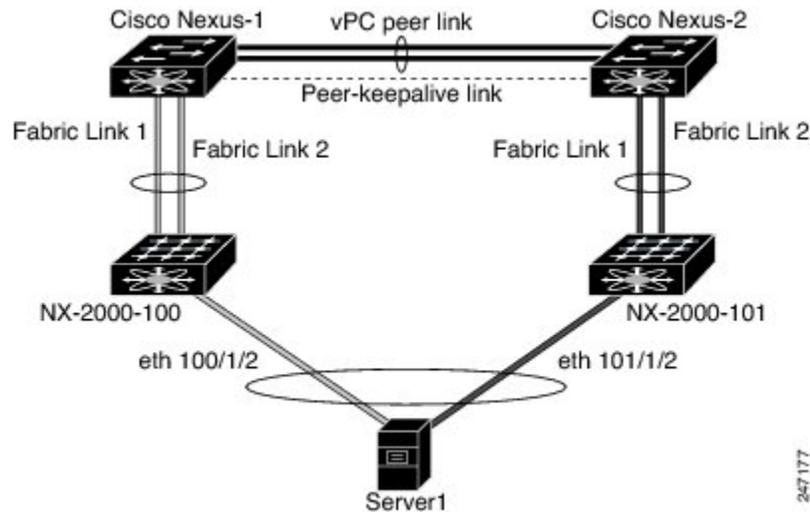
Note The first 8 ports on the Cisco Nexus 5010 switch and the first 16 ports on the Cisco Nexus 5020 switch are switchable 1-Gigabit and 10-Gigabit ports. You can enable vPC functionality on these ports in 1-Gigabit mode.

The switch connected to the pair of Cisco Nexus 5000 Series switches can be any standards-based Ethernet switch. Common environments to use this configuration include Blade Chassis with dual switches connected to the pair of Cisco Nexus 5000 Series switches through vPC or Unified Computing Systems connected to the pair of Cisco Nexus 5000 Series switches.

Single Homed Fabric Extender vPC Topology

You can connect a server with dual interfaces that are configured in a vPC to a pair of Cisco Nexus 2000 Series Fabric Extenders that you then connect single homed to the Cisco Nexus 5000 Series switches. The topology that is shown in the following figure provides the vPC functionality to dual homed servers with 1-Gigabit Ethernet uplink interfaces.

Figure 15: Single Homed Fabric Extender vPC Topology



The Cisco Nexus 5000 Series switch can support up to 12 configured single homed Fabric Extenders (576 ports) with this topology however only 480 dual homed host servers can be configured in a vPCs with this configuration.



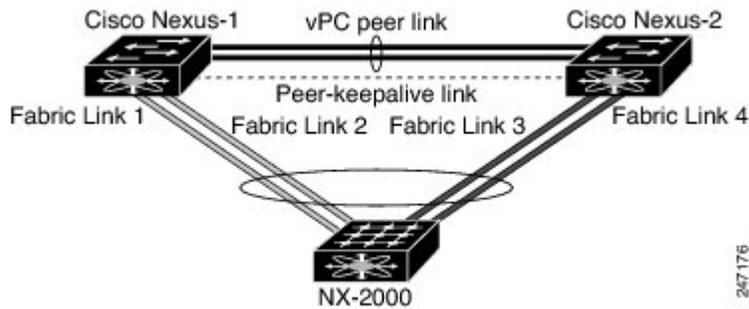
Note

The current generation of Cisco Nexus 2000 Series Fabric Extender does not support EtherChannels on its host interfaces. Therefore a maximum of two links can be configured in an EtherChannel from the server where each link is connected to a separate Fabric Extender.

Dual Homed Fabric Extender vPC Topology

You can connect the Cisco Nexus 2000 Series Fabric Extender to two upstream Cisco Nexus 5000 Series switches and downstream to a number of single homed servers. The topology shown in the following figure provides the vPC functionality to singly connected servers with 1-Gigabit Ethernet uplink interfaces.

Figure 16: Dual Homed Fabric Extender vPC Topology



The Cisco Nexus 5000 Series switch can support up to 12 configured dual homed Fabric Extenders with this topology. A maximum of 480 single homed servers can be connected to this configuration.

vPC Domain

You can use the vPC domain ID to identify the vPC peer links and the ports that are connected to the vPC downstream switches.

The vPC domain is a configuration mode that you use to configure the keepalive messages and also configure other vPC peer link parameters rather than accept the default values.

To create a vPC domain, you must first create a vPC domain ID on each vPC peer switch using a number from 1 to 1000. This ID must be the same on a set of vPC peer devices. Within this domain, the system provides a loop-free topology and multipathing.

You can configure the EtherChannels and vPC peer links by using LACP or no protocol. When possible, we recommend that you use LACP with the interfaces in active mode to configure EtherChannels in each vPC to ensure an optimized, graceful recovery in a port-channel failover scenario and provides configuration checks against a configuration mismatch among the EtherChannels.

The vPC peer switches use the vPC domain ID that you configure to automatically assign a unique vPC system MAC address. Each vPC domain has a unique MAC address that is used as a unique identifier for the specific vPC-related operations, although the switches use the vPC system MAC addresses only for link-scope operations, such as LACP. We recommend that you create each vPC domain within the contiguous network with a unique domain ID. You can also configure a specific MAC address for the vPC domain, rather than having the Cisco NX-OS software assign the address.



Note

You must set a unique vPC domain ID to avoid system ID issues with LACP vPCs.

After you create a vPC domain, the Cisco NX-OS software automatically creates a system priority for the vPC domain. You can also manually configure a specific system priority for the vPC domain.

**Note**

If you manually configure the system priority, you must ensure that you assign the same priority value on both vPC peer switches. If the vPC peer switches have different system priority values, the vPC will not come up.

Peer-Keepalive Link and Messages

The Cisco NX-OS software uses a peer-keepalive link between the vPC peers to transmit periodic, configurable keepalive messages. You must have Layer 3 connectivity between the peer switches to transmit these messages; the system cannot bring up the vPC peer link unless a peer-keepalive link is already up and running.

If one of the vPC peer switches fails, the vPC peer switch on the other side of the vPC peer link senses the failure when it does not receive any peer-keepalive messages. The default interval time for the vPC peer-keepalive message is 1 second. You can configure the interval between 400 milliseconds and 10 seconds. You can also configure a timeout value with a range of 3 to 20 seconds; the default timeout value is 5 seconds. The peer-keepalive status is checked only when the peer-link goes down.

The vPC peer-keepalive can be carried either in the management or default VRF on the Cisco Nexus 5000 Series switch. When you configure the switches to use the management VRF, the source and destination for the keepalive messages are the mgmt 0 interface IP addresses. When you configure the switches to use the default VRF, an SVI must be created to act as the source and destination addresses for the vPC peer-keepalive messages. Ensure that both the source and destination IP addresses used for the peer-keepalive messages are unique in your network and these IP addresses are reachable from the VRF associated with the vPC peer-keepalive link.

**Note**

We recommend that you configure the vPC peer-keepalive link on the Cisco Nexus 5000 Series switch to run in the management VRF using the mgmt 0 interfaces. If you configure the default VRF, ensure that the vPC peer link is not used to carry the vPC peer-keepalive messages.

Compatibility Parameters for vPC Peer Links

Many configuration and operational parameters must be identical on all interfaces in the vPC. After you enable the vPC feature and configure the peer link on both vPC peer switches, Cisco Fabric Services (CFS) messages provide a copy of the configuration on the local vPC peer switch configuration to the remote vPC peer switch. The system then determines whether any of the crucial configuration parameters differ on the two switches.

Enter the **show vpc consistency-parameters** command to display the configured values on all interfaces in the vPC. The displayed configurations are only those configurations that would limit the vPC peer link and vPC from coming up.

The compatibility check process for vPCs differs from the compatibility check for regular EtherChannels.

Configuration Parameters That Must Be Identical

The configuration parameters in this section must be configured identically on both switches at either end of the vPC peer link or the vPC is moved into suspend mode.

**Note**

You must ensure that all interfaces in the vPC have the identical operational and configuration parameters listed in this section.

Enter the **show vpc consistency-parameters** command to display the configured values on all interfaces in the vPC. The displayed configurations are only those configurations that would limit the vPC peer link and vPC from coming up.

The switch automatically check for compatibility of these parameters on the vPC interfaces. The per-interface parameters must be consistent per interface, and the global parameters must be consistent globally.

- Port-channel mode: on, off, or active
- Link speed per channel
- Duplex mode per channel
- Trunk mode per channel:
 - Native VLAN
 - VLANs allowed on trunk
 - Tagging of native VLAN traffic
- Spanning Tree Protocol (STP) mode
- STP region configuration for Multiple Spanning Tree (MST)
- Enable or disable state per VLAN
- STP global settings:
 - Bridge Assurance setting
 - Port type setting—We recommend that you set all vPC interfaces as network ports
 - Loop Guard settings
- STP interface settings:
 - Port type setting
 - Loop Guard
 - Root Guard
- Maximum Transmission Unit (MTU)
- Quality of Service global settings
 - System QoS policy
 - System Network-QoS policy
 - System Input Queuing policy
 - System Output Queuing policy

- For the Fabric Extender vPC topology, all the interface level parameters mentioned above should be identically configured for host interface from both the switches.
- Fabric Extender FEX number configured on an EtherChannel fabric interface; for the Fabric Extender vPC topology.

If any of these parameters are not enabled or defined on either switch, the vPC consistency check ignores those parameters.


Note

To ensure that none of the vPC interfaces are in the suspend mode, enter the **show vpc brief** and **show vpc consistency-parameters** commands and check the syslog messages.

Configuration Parameters That Should Be Identical

When any of the following parameters are not configured identically on both vPC peer switches, a misconfiguration may cause undesirable behavior in the traffic flow:

- MAC aging timers
- Static MAC entries
- VLAN interface—Each switch on the end of the vPC peer link must have a VLAN interface configured for the same VLAN on both ends and they must be in the same administrative and operational mode. Those VLANs configured on only one switch of the peer link do not pass traffic using the vPC or peer link. You must create all VLANs on both the primary and secondary vPC switches, or the VLAN will be suspended.
- Private VLAN configuration
- All ACL configurations and parameters
- Quality of service (QoS) configuration and parameters—Local parameters; global parameters must be identical
- STP interface settings:
 - BPDU Filter
 - BPDU Guard
 - Cost
 - Link type
 - Priority
 - VLANs (Rapid PVST+)

To ensure that all the configuration parameters are compatible, we recommend that you display the configurations for each vPC peer switch once you configure the vPC.

vPC Peer Links

A vPC peer link is the link that is used to synchronize the states between the vPC peer devices.

**Note**

You must configure the peer-keepalive link before you configure the vPC peer link or the peer link will not come up.

vPC Peer Link Overview

You can have only two switches as vPC peers; each switch can serve as a vPC peer to only one other vPC peer. The vPC peer switches can also have non-vPC links to other switches.

To make a valid configuration, you configure an EtherChannel on each switch and then configure the vPC domain. You assign the EtherChannel on each switch as a peer link. For redundancy, we recommend that you should configure at least two dedicated ports into the EtherChannel; if one of the interfaces in the vPC peer link fails, the switch automatically falls back to use another interface in the peer link.

**Note**

We recommend that you configure the EtherChannels in trunk mode.

Many operational parameters and configuration parameters must be the same in each switch connected by a vPC peer link. Because each switch is completely independent on the management plane, you must ensure that the switches are compatible on the critical parameters. vPC peer switches have separate control planes. After configuring the vPC peer link, you should display the configuration on each vPC peer switch to ensure that the configurations are compatible.

**Note**

You must ensure that the two switches connected by the vPC peer link have certain identical operational and configuration parameters.

When you configure the vPC peer link, the vPC peer switches negotiate that one of the connected switches is the primary switch and the other connected switch is the secondary switch. By default, the Cisco NX-OS software uses the lowest MAC address to elect the primary switch. The software takes different actions on each switch—that is, the primary and secondary—only in certain failover conditions. If the primary switch fails, the secondary switch becomes the operational primary switch when the system recovers, and the previously primary switch is now the secondary switch.

You can also configure which of the vPC switches is the primary switch. If you want to configure the role priority again to make one vPC switch the primary switch, configure the role priority on both the primary and secondary vPC switches with the appropriate values, shut down the EtherChannel that is the vPC peer link on both switches by entering the **shutdown** command, and reenab the EtherChannel on both switches by entering the **no shutdown** command.

For known unicast traffic, you should use the local links of the vPC because you cannot load balance traffic across the peer link. Unknown unicast, multicast, and broadcast traffic are flooded across the vPC peer link. The software keeps the multicast forwarding state synchronized between the two peers for groups learned over the vPC link. If the multicast forwarding states are learned over isolated ports on either peer, the states are not synchronized, but since the peer link is a router port, the traffic is forwarded to the peer; the peer then handles the forwarding if there are other isolated ports on that switch.

MAC addresses that are learned over vPC links are also synchronized between the peers.

Configuration information flows across the vPC peer links using the Cisco Fabric Services over Ethernet (CFS over Ethernet) protocol. All MAC addresses for those VLANs configured on both switches are synchronized between vPC peer switches. The software uses CFS over Ethernet for this synchronization.

If the vPC peer link fails, the software checks the status of the remote vPC peer switch using the peer-keepalive link, which is a link between vPC peer switches, to ensure that both switches are up. If the vPC peer switch is up, the secondary vPC switch disables all vPC ports on its switch to prevent loops and disappearing or flooding traffic. The data then forwards down the remaining active links of the EtherChannel.

The software learns of a vPC peer switch failure when the keepalive messages are not returned over the peer-keepalive link.

Use a separate link (vPC peer-keepalive link) to send configurable keepalive messages between the vPC peer switches. The keepalive messages on the vPC peer-keepalive link determines whether a failure is on the vPC peer link only or on the vPC peer switch. The keepalive messages are used only when all the links in the peer link fail.

Manually Configured vPC Features

You must manually configure the following features to conform to the primary/secondary mapping of each of the vPC peer devices:

- STP root—Configure the primary vPC peer device as the highest STP root priority, and configure the secondary device with a lower root priority.

We recommend that you configure the vPC peer link interfaces as STP network ports so that Bridge Assurance is enabled on all vPC peer links

We recommend that you configure Rapid PVST+ so that the primary device is the root for all VLANs and configure MST so that the primary device is the root for all instances.

- We recommend that you configure Unidirectional Link Detection (UDLD) on both sides of the vPC peer link.

vPC Number

Once you have created the vPC domain ID and the vPC peer link, you can create EtherChannels to attach the downstream switch to each vPC peer switch. That is, you create one EtherChannel from the downstream switch to the primary vPC peer switch and you create another EtherChannel from the downstream switch to the secondary peer switch.

On each vPC peer switch, you assign the same vPC number to the EtherChannel that connects to the downstream switch. You will experience minimal traffic disruption when you are creating vPCs. To simplify the configuration, you can assign the vPC ID number for each EtherChannel to be the same as the EtherChannel itself (that is, vPC ID 10 for EtherChannel 10).



Note

The vPC number that you assign to the EtherChannel connecting to the downstream switch from the vPC peer switch must be identical on both vPC peer switches.

vPC Interactions with Other Features

vPC and LACP

The Link Aggregation Control Protocol (LACP) uses the system MAC address of the vPC domain to form the LACP Aggregation Group (LAG) ID for the vPC.

You can use LACP on all the vPC EtherChannels, including those channels from the downstream switch. We recommend that you configure LACP with active mode on the interfaces on each EtherChannel on the vPC peer switches. This configuration allows you to more easily detect compatibility between switches, unidirectional links, and multihop connections, and provides dynamic reaction to run-time changes and link failures.

The vPC peer link supports 16 EtherChannel LACP interfaces. You should manually configure the system priority on the vPC peer-link switches to ensure that the vPC peer-link switches have a higher LACP priority than the downstream connected switches. A lower numerical value system priority means a higher LACP priority.



Note

When manually configuring the system priority, you must ensure that you assign the same priority value on both vPC peer switches. If the vPC peer switches have different system priority values, vPC will not come up.

vPC Peer Links and STP

When you first bring up the vPC functionality, STP reconverges. STP treats the vPC peer link as a special link and always includes the vPC peer link in the STP active topology.

We recommend that you set all the vPC peer link interfaces to the STP network port type so that Bridge Assurance is automatically enabled on all vPC peer links. We also recommend that you do not enable any of the STP enhancement features on VPC peer links.

You must configure a list of parameters to be identical on the vPC peer switches on both sides of the vPC peer link.

STP is distributed; that is, the protocol continues running on both vPC peer switches. However, the configuration on the vPC peer switch elected as the primary switch controls the STP process for the vPC interfaces on the secondary vPC peer switch.

The primary vPC switch synchronizes the STP state on the vPC secondary peer switch using Cisco Fabric Services over Ethernet (CFSoE).

The vPC manager performs a proposal/handshake agreement between the vPC peer switches that sets the primary and secondary switches and coordinates the two switches for STP. The primary vPC peer switch then controls the STP protocol for vPC interfaces on both the primary and secondary switches. We recommend that you configure the primary vPC peer switch as the highest STP root priority, and configure the secondary switch with a lower root priority.

If the primary vPC peer switch fails over to the secondary vPC peer switch, there is no change in the STP topology.

The Bridge Protocol Data Units (BPDUs) use the MAC address set for the vPC for the STP bridge ID in the designated bridge ID field. The vPC primary switch sends these BPDUs on the vPC interfaces.

**Note**

Display the configuration on both sides of the vPC peer link to ensure that the settings are identical. Use the **show spanning-tree** command to display information about the vPC.

CFSOE

The Cisco Fabric Services over Ethernet (CFSOE) is a reliable state transport mechanism that you can use to synchronize the actions of the vPC peer devices. CFSOE carries messages and packets for many features linked with vPC, such as STP and IGMP. Information is carried in CFS/CFSOE protocol data units (PDUs).

When you enable the vPC feature, the device automatically enables CFSOE, and you do not have to configure anything. CFSOE distributions for vPCs do not need the capabilities to distribute over IP or the CFS regions. You do not need to configure anything for the CFSOE feature to work correctly on vPCs.

You can use the **show mac address-table** command to display the MAC addresses that CFSOE synchronizes for the vPC peer link.

**Note**

Do not enter the **no cfs eth distribute** or the **no cfs distribute** command. CFSOE must be enabled for vPC functionality. If you do enter either of these commands when vPC is enabled, the system displays an error message.

When you enter the **show cfs application** command, the output displays "Physical-eth," which shows the applications that are using CFSOE.

vPC Guidelines and Limitations

vPC has the following configuration guidelines and limitations:

- You must enable the vPC feature before you can configure vPC peer-link and vPC interfaces.
- You must configure the peer-keepalive link before the system can form the vPC peer link.
- Only EtherChannels can be in vPCs. A vPC can be configured on a normal EtherChannel (switch-to-switch vPC topology), on an EtherChannel fabric interface (fabric extender vPC topology), and on an EtherChannel host interface (host interface vPC topology).

**Note**

Refer to the Cisco Nexus 2000 Series Fabric Extender Software Configuration Guide for information about Fabric Extender host and fabric interfaces.

- A Fabric Extender can be a member of a Host Interface vPC topology or a Fabric Extender vPC topology but not both simultaneously.
- You must configure both vPC peer switches; the configuration is not automatically synchronized between the vPC peer devices.
- Check that the necessary configuration parameters are compatible on both sides of the vPC peer link.
- You may experience minimal traffic disruption while configuring vPCs.

- You should configure all the EtherChannels in the vPC using LACP with the interfaces in active mode.

Configuring vPCs

Enabling vPCs

You must enable the vPC feature before you can configure and use vPCs.

SUMMARY STEPS

1. switch# **configure terminal**
2. switch(config)# **feature vpc**
3. (Optional) switch# **show feature**
4. (Optional) switch# **copy running-config startup-config**

DETAILED STEPS

	Command or Action	Purpose
Step 1	switch# configure terminal	Enters configuration mode.
Step 2	switch(config)# feature vpc	Enables vPCs on the switch.
Step 3	switch# show feature	(Optional) Displays which features are enabled on the switch.
Step 4	switch# copy running-config startup-config	(Optional) Copies the running configuration to the startup configuration.

This example shows how to enable the vPC feature:

```
switch# configure terminal
switch(config)# feature vpc
```

Disabling vPCs

You can disable the vPC feature.



Note

When you disable the vPC feature, the Cisco Nexus 5000 Series switch clears all the vPC configurations.

SUMMARY STEPS

1. switch# **configure terminal**
2. switch(config)# **no feature vpc**
3. (Optional) switch# **show feature**
4. (Optional) switch# **copy running-config startup-config**

DETAILED STEPS

	Command or Action	Purpose
Step 1	switch# configure terminal	Enters configuration mode.
Step 2	switch(config)# no feature vpc	Disables vPCs on the switch.
Step 3	switch# show feature	(Optional) Displays which features are enabled on the switch.
Step 4	switch# copy running-config startup-config	(Optional) Copies the running configuration to the startup configuration.

This example shows how to disable the vPC feature:

```
switch# configure terminal
switch(config)# no feature vpc
```

Creating a vPC Domain

You must create identical vPC domain IDs on both the vPC peer devices. This domain ID is used to automatically form the vPC system MAC address.

Before You Begin

Ensure that you have enabled the vPC feature.

You must configure both switches on either side of the vPC peer link with the following procedure.

SUMMARY STEPS

1. switch# **configure terminal**
2. switch(config)# **vpc domain domain-id**
3. (Optional) switch# **show vpc brief**
4. (Optional) switch# **copy running-config startup-config**

DETAILED STEPS

	Command or Action	Purpose
Step 1	switch# configure terminal	Enters configuration mode.
Step 2	switch(config)# vpc domain <i>domain-id</i>	Creates a vPC domain on the switch, and enters the vpc-domain configuration mode. There is no default <i>domain-id</i> ; the range is from 1 to 1000. Note You can also use the vpc domain command to enter the vpc-domain configuration mode for an existing vPC domain.
Step 3	switch# show vpc brief	(Optional) Displays brief information about each vPC domain.
Step 4	switch# copy running-config startup-config	(Optional) Copies the running configuration to the startup configuration.

This example shows how to create a vPC domain:

```
switch# configure terminal
switch(config)# vpc domain 5
```

Configuring a vPC Keepalive Link

You can configure the destination IP for the peer-keepalive link that carries the keepalive messages. Optionally, you can configure other parameters for the keepalive messages.

Before You Begin

Ensure that you have enabled the vPC feature.

You must configure the vPC peer-keepalive link before the system can form the vPC peer link.

You must configure both switches on either side of the vPC peer link with the following procedure.

SUMMARY STEPS

1. switch# **configure terminal**
2. switch(config)# **vpc domain** *domain-id*
3. switch(config-vpc-domain)# **peer-keepalive destination** *ipaddress* [**hold-timeout** *secs* | **interval** *msecs* {**timeout** *secs*} | **precedence** {*prec-value* | **network** | **internet** | **critical** | **flash-override** | **flash** | **immediate** | **priority** | **routine**} | **tos** {*tos-value* | **max-reliability** | **max-throughput** | **min-delay** | **min-monetary-cost** | **normal**} | **tos-byte** *tos-byte-value*} | **source** *ipaddress* | **vrf** {*name* | **management vpc-keepalive**}]
4. (Optional) switch# **show vpc peer-keepalive**
5. (Optional) switch# **copy running-config startup-config**

DETAILED STEPS

	Command or Action	Purpose
Step 1	switch# configure terminal	Enters configuration mode.
Step 2	switch(config)# vpc domain <i>domain-id</i>	Creates a vPC domain on the switch if it does not already exist, and enters the vpc-domain configuration mode.
Step 3	switch(config-vpc-domain)# peer-keepalive destination <i>ipaddress</i> [hold-timeout <i>secs</i> interval <i>msecs</i> { timeout <i>secs</i> } precedence { <i>prec-value</i> network internet critical flash-override flash immediate priority routine } tos { <i>tos-value</i> max-reliability max-throughput min-delay min-monetary-cost normal } tos-byte <i>tos-byte-value</i> } source <i>ipaddress</i> vrf { <i>name</i> management vpc-keepalive }]	Configures the IPv4 address for the remote end of the vPC peer-keepalive link. Note The system does not form the vPC peer link until you configure a vPC peer-keepalive link. The management ports and VRF are the defaults
Step 4	switch# show vpc peer-keepalive	(Optional) Displays information about the configuration for the keepalive messages.
Step 5	switch# copy running-config startup-config	(Optional) Copies the running configuration to the startup configuration.

This example shows how to configure the destination IP address for the vPC-peer-keepalive link:

```
switch# configure terminal
switch(config)# vpc domain 5
switch(config-vpc-domain)# peer-keepalive destination 10.10.10.42
```

Creating a vPC Peer Link

You can create a vPC peer link by designating the EtherChannel that you want on each switch as the peer link for the specified vPC domain. We recommend that you configure the EtherChannels that you are designating as the vPC peer link in trunk mode and that you use two ports on separate modules on each vPC peer switch for redundancy.

Before You Begin

Ensure that you have enabled the vPC feature.

You must configure both switches on either side of the vPC peer link with the following procedures

SUMMARY STEPS

1. switch# **configure terminal**
2. switch(config)# **interface port-channel** *channel-number*
3. switch(config-if)# **vpc peer-link**
4. (Optional) switch# **show vpc brief**
5. (Optional) switch# **copy running-config startup-config**

DETAILED STEPS

	Command or Action	Purpose
Step 1	switch# configure terminal	Enters configuration mode.
Step 2	switch(config)# interface port-channel <i>channel-number</i>	Selects the EtherChannel that you want to use as the vPC peer link for this switch, and enters the interface configuration mode.
Step 3	switch(config-if)# vpc peer-link	Configures the selected EtherChannel as the vPC peer link, and enters the vpc-domain configuration mode.
Step 4	switch# show vpc brief	(Optional) Displays information about each vPC, including information about the vPC peer link.
Step 5	switch# copy running-config startup-config	(Optional) Copies the running configuration to the startup configuration.

This example shows how to configure a vPC peer link:

```
switch# configure terminal
switch(config)# interface port-channel 20
switch(config-if)# vpc peer-link
```

Checking the Configuration Compatibility

After you have configured the vPC peer link on both vPC peer switches, check that the configurations are consistent on all vPC interfaces.

Parameter	Default Setting
switch# show vpc consistency-parameters { global interface port-channel <i>channel-number</i> }	Displays the status of those parameters that must be consistent across all vPC interfaces.

This example shows how to check that the required configurations are compatible across all the vPC interfaces:

```
switch# show vpc consistency-parameters global
```

Legend:

Type 1 : vPC will be suspended in case of mismatch

Name	Type	Local Value	Peer Value
QoS	1	([], [3], [0,7], [2], [4], [6])	([], [3], [0,7], [2], [4], [6])
Network QoS (MTU)	1	(1538, 2240, 0, 0, 0, 0)	(1538, 2240, 0, 0, 0, 0)
Network QoS (Pause)	1	(F, T, F, F, F, F)	(F, T, F, F, F, F)
Input Queuing (Bandwidth)	1	(50, 50, 0, 0, 0, 0)	(50, 50, 0, 0, 0, 0)
Input Queuing (Absolute Priority)	1	(F, F, F, F, F, F)	(F, F, F, F, F, F)
Output Queuing (Bandwidth)	1	(50, 50, 0, 0, 0, 0)	(50, 50, 0, 0, 0, 0)
Output Queuing (Absolute Priority)	1	(F, F, F, F, F, F)	(F, F, F, F, F, F)
STP Mode	1	MST	MST
STP Disabled	1	None	None
STP MST Region Name	1	""	""
STP MST Region Revision	1	0	0
STP MST Region Instance to VLAN Mapping	1		
STP Loopguard	1	Disabled	Disabled
STP Bridge Assurance	1	Enabled	Enabled
STP Port Type	1	Normal	Normal
STP MST Simulate PVST	1	Enabled	Enabled
Allowed VLANs	-	-	-

This example shows how to check that the required configurations are compatible for an EtherChannel interface:

```
switch# show vpc consistency-parameters interface port-channel 20
```

Legend:

Type 1 : vPC will be suspended in case of mismatch

Name	Type	Local Value	Peer Value
Fex id	1	20	20
STP Port Type	1	Default	Default
STP Port Guard	1	None	None
STP MST Simulate PVST	1	Default	Default
mode	1	on	on
Speed	1	10 Gb/s	10 Gb/s
Duplex	1	full	full
Port Mode	1	fex-fabric	fex-fabric
Shut Lan	1	No	No
Allowed VLANs	-	1,3-3967,4048-4093	1-3967,4048-4093

Creating an EtherChannel Host Interface

To connect to a downstream server from a Cisco Nexus 2000 Series Fabric Extender you can create a EtherChannel host interface. An EtherChannel host interface can have only one host interface as a member. You need to create an EtherChannel host interface to configure a vPC on it that uses the Fabric Extender topology.



Note

See the *Cisco Nexus 2000 Series Fabric Extender Software Configuration Guide* for information on attaching a Fabric Extender to a Cisco Nexus 5000 Series switch.

Before You Begin

Ensure that you have enabled the vPC feature.

Ensure that the connected Fabric Extender is online.

You must configure both switches on either side of the vPC peer link with the following procedure.

SUMMARY STEPS

1. switch# **configure terminal**
2. switch(config)# **interface ethernet** *chassis/slot/port*
3. switch(config-if)# **channel-group** *channel-number* **mode** {**active** | **passive** | **on**}
4. (Optional) switch# **show port-channel summary**
5. (Optional) switch# **copy running-config startup-config**

DETAILED STEPS

	Command or Action	Purpose
Step 1	switch# configure terminal	Enters configuration mode.
Step 2	switch(config)# interface ethernet <i>chassis/slot/port</i>	Specifies an interface to configure, and enters interface configuration mode.
Step 3	switch(config-if)# channel-group <i>channel-number</i> mode { active passive on }	Creates an EtherChannel host interface on the selected host interface.
Step 4	switch# show port-channel summary	(Optional) Displays information about each EtherChannel host interface.
Step 5	switch# copy running-config startup-config	(Optional) Copies the running configuration to the startup configuration.

This example shows how to configure an EtherChannel host interface:

```
switch# configure terminal
switch(config)# interface ethernet 101/1/20
switch(config-if)# channel-group 7 mode active
```

Moving Other EtherChannels into a vPC

To connect to the downstream switch, you create an EtherChannel from the downstream switch to the primary vPC peer switch and you create another EtherChannel from the downstream switch to the secondary peer switch. Finally, working on each vPC peer switch, you assign a vPC number to the EtherChannel that connects to the downstream switch. You will experience minimal traffic disruption when you are creating vPCs.

Before You Begin

Ensure that you have enabled the vPC feature.

You must configure both switches on either side of the vPC peer link with the following procedure.

SUMMARY STEPS

1. switch# **configure terminal**
2. switch(config)# **interface port-channel** *channel-number*
3. switch(config-if)# **vpc** *number*
4. (Optional) switch# **show vpc brief**
5. (Optional) switch# **copy running-config startup-config**

DETAILED STEPS

	Command or Action	Purpose
Step 1	switch# configure terminal	Enters configuration mode.
Step 2	switch(config)# interface port-channel <i>channel-number</i>	Selects the EtherChannel that you want to put into the vPC to connect to the downstream switch, and enters the interface configuration mode. Note A vPC can be configured on a normal EtherChannel (physical vPC topology), on an EtherChannel fabric interface (fabric extender vPC topology), and on an EtherChannel host interface (host interface vPC topology)
Step 3	switch(config-if)# vpc <i>number</i>	Configures the selected EtherChannel into the vPC to connect to the downstream switch. The range is from 1 to 4096. The vPC <i>number</i> that you assign to the EtherChannel connecting to the downstream switch from the vPC peer switch must be identical on both vPC peer switches.
Step 4	switch# show vpc brief	(Optional) Displays information about each vPC.
Step 5	switch# copy running-config startup-config	(Optional) Copies the running configuration to the startup configuration.

This example shows how to configure an EtherChannel that will connect to the downstream device:

```
switch# configure terminal
switch(config)# interface port-channel 20
switch(config-if)# vpc 5
```

Manually Configuring a vPC Domain MAC Address

You create the vPC peer link by designating the EtherChannel that you want on each switch as the peer link for the specified vPC domain. We recommend that you configure the EtherChannels that you are designating as the vPC peer link in trunk mode and that you use two ports on separate modules on each vPC peer switch for redundancy.

Before You Begin

Ensure that you have enabled the vPC feature.

You must configure both switches on either side of the vPC peer link with the following procedure.

SUMMARY STEPS

1. switch# **configure terminal**
2. switch(config)# **vpc domain** *domain-id*
3. switch(config-vpc-domain)# **system-mac** *mac-address*
4. (Optional) switch# **show vpc role**
5. (Optional) switch# **copy running-config startup-config**

DETAILED STEPS

	Command or Action	Purpose
Step 1	switch# configure terminal	Enters configuration mode.
Step 2	switch(config)# vpc domain <i>domain-id</i>	Selects an existing vPC domain on the switch, or creates a new vPC domain, and enters the vpc-domain configuration mode. There is no default <i>domain-id</i> ; the range is from 1 to 1000.
Step 3	switch(config-vpc-domain)# system-mac <i>mac-address</i>	Enters the MAC address that you want for the specified vPC domain in the following format: aaaa.bbbb.cccc.
Step 4	switch# show vpc role	(Optional) Displays the vPC system MAC address.
Step 5	switch# copy running-config startup-config	(Optional) Copies the running configuration to the startup configuration.

This example shows how to configure a vPC domain MAC address:

```
switch# configure terminal
switch(config)# vpc domain 5
switch(config-if)# system-mac 23fb.4ab5.4c4e
```

Manually Configuring the System Priority

When you create a vPC domain, the system automatically creates a vPC system priority. However, you can also manually configure a system priority for the vPC domain.

**Note**

We recommend that you manually configure the vPC system priority when you are running LACP to ensure that the vPC peer switches are the primary switches on LACP. When you manually configure the system priority, ensure that you configure the same priority value on both vPC peer switches. If these values do not match, vPC will not come up.

Before You Begin

Ensure that you have enabled the vPC feature.

You must configure both switches on either side of the vPC peer link with the following procedure.

SUMMARY STEPS

1. switch# **configure terminal**
2. switch(config)# **vpc domain** *domain-id*
3. switch(config-vpc-domain)# **system-priority** *priority*
4. (Optional) switch# **show vpc brief**
5. (Optional) switch# **copy running-config startup-config**

DETAILED STEPS

	Command or Action	Purpose
Step 1	switch# configure terminal	Enters configuration mode.
Step 2	switch(config)# vpc domain <i>domain-id</i>	Selects an existing vPC domain on the switch, or creates a new vPC domain, and enters the vpc-domain configuration mode. There is no default <i>domain-id</i> ; the range is from 1 to 1000.
Step 3	switch(config-vpc-domain)# system-priority <i>priority</i>	Enters the system priority that you want for the specified vPC domain. The range of values is from 1 to 65535. The default value is 32667.
Step 4	switch# show vpc brief	(Optional) Displays information about each vPC, including information about the vPC peer link.
Step 5	switch# copy running-config startup-config	(Optional) Copies the running configuration to the startup configuration.

This example shows how to configure a vPC peer link:

```
switch# configure terminal
switch(config)# vpc domain 5
switch(config-if)# system-priority 4000
```

Manually Configuring a vPC Peer Switch Role

By default, the Cisco NX-OS software elects a primary and secondary vPC peer switch after you configure the vPC domain and both sides of the vPC peer link. However, you may want to elect a specific vPC peer switch as the primary switch for the vPC. Then, you would manually configure the role value for the vPC peer switch that you want as the primary switch to be lower than the other vPC peer switch.

vPC does not support role preemption. If the primary vPC peer switch fails, the secondary vPC peer switch takes over to become operationally the vPC primary switch. However, the original operational roles are not restored when the formerly primary vPC comes up again.

Before You Begin

Ensure that you have enabled the vPC feature.

You must configure both switches on either side of the vPC peer link with the following procedure.

SUMMARY STEPS

1. switch# **configure terminal**
2. switch(config)# **vpc domain** *domain-id*
3. switch(config-vpc-domain)# **role priority** *priority*
4. (Optional) switch# **show vpc brief**
5. (Optional) switch# **copy running-config startup-config**

DETAILED STEPS

	Command or Action	Purpose
Step 1	switch# configure terminal	Enters configuration mode.
Step 2	switch(config)# vpc domain <i>domain-id</i>	Selects an existing vPC domain on the switch, or creates a new vPC domain, and enters the vpc-domain configuration mode. There is no default <i>domain-id</i> ; the range is from 1 to 1000.
Step 3	switch(config-vpc-domain)# role priority <i>priority</i>	Enters the role priority that you want for the vPC system priority. The range of values is from 1 to 65535. The default value is 32667.
Step 4	switch# show vpc brief	(Optional) Displays information about each vPC, including information about the vPC peer link.
Step 5	switch# copy running-config startup-config	(Optional) Copies the running configuration to the startup configuration.

This example shows how to configure a vPC peer link:

```
switch# configure terminal
switch(config)# vpc domain 5
switch(config-if)# role priority 4000
```

Verifying the vPC Configuration

Use the following commands to display vPC configuration information:

Command	Purpose
switch# show feature	Displays whether vPC is enabled or not.
switch# show port-channel capacity	Displays how many EtherChannels are configured and how many are still available on the switch.
switch# show running-config vpc	Displays running configuration information for vPCs.
switch# show vpc brief	Displays brief information on the vPCs.
switch# show vpc consistency-parameters	Displays the status of those parameters that must be consistent across all vPC interfaces.
switch# show vpc peer-keepalive	Displays information on the peer-keepalive messages.
switch# show vpc role	Displays the peer status, the role of the local switch, the vPC system MAC address and system priority, and the MAC address and priority for the local vPC switch.
switch# show vpc statistics	Displays statistics on the vPCs. Note This command displays the vPC statistics only for the vPC peer device that you are working on.

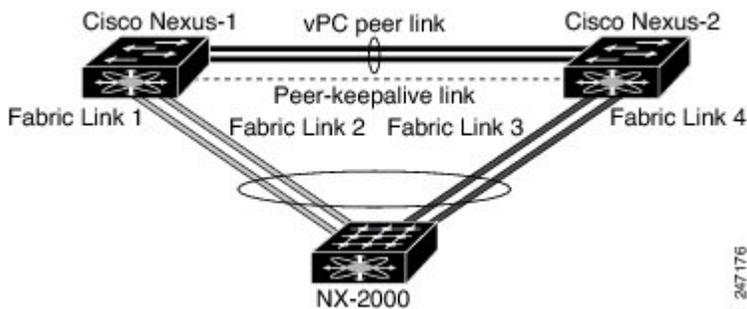
For detailed information about the fields in the output from these commands, see the *Cisco Nexus 5000 Series Command Reference*.

vPC Example Configurations

Dual Homed Fabric Extender vPC Configuration Example

The following example shows how to configure the dual homed Fabric Extender vPC topology using the management VRF to carry the peer-keepalive messages on switch NX-5000-1 as shown in following figure:

Figure 17: vPC Configuration Example



Before You Begin

Ensure that the Cisco Nexus 2000 Series Fabric Extender NX-2000-100 is attached and online.

SUMMARY STEPS

1. Enable vPC and LACP.
2. Create the vPC domain and add the vPC peer-keepalive link.
3. Configure the vPC peer link as a two port Etherchannel.
4. Configure the Fabric Extender NX-2000-100.
5. Configure the fabric EtherChannel links for the Fabric Extender NX-2000-100.
6. Configure the host interface ports on the Fabric Extender NX-2000-100.
7. Save the configuration.

DETAILED STEPS

Step 1

Enable vPC and LACP.

```
NX-5000-1# configure terminal
NX-5000-1(config)# feature lACP
NX-5000-1(config)# feature vpc
```

Step 2

Create the vPC domain and add the vPC peer-keepalive link.

```
NX-5000-1(config)# vpc domain 1
NX-5000-1(config-vpc-domain)# peer-keepalive destination 10.10.10.237
NX-5000-1(config-vpc-domain)# exit
```

Step 3 Configure the vPC peer link as a two port Etherchannel.

```
NX-5000-1(config)# interface ethernet 1/1-2
NX-5000-1(config-if-range)# switchport mode trunk
NX-5000-1(config-if-range)# switchport trunk allowed vlan 20-50
NX-5000-1(config-if-range)# switchport trunk native vlan 20
NX-5000-1(config-if-range)# channel-group 20 mode active
NX-5000-1(config-if-range)# exit
NX-5000-1(config)# interface port-channel 20
NX-5000-1(config-if)# vpc peer-link
NX-5000-1(config-if)# exit
```

Step 4 Configure the Fabric Extender NX-2000-100.

```
NX-5000-1(config)# fex 100
NX-5000-1(config-fex)# pinning max-links 1
NX-5000-1(fex)# exit
```

Step 5 Configure the fabric EtherChannel links for the Fabric Extender NX-2000-100.

```
NX-5000-1(config)# interface ethernet 1/20
NX-5000-1(config-if)# channel-group 100
NX-5000-1(config-if)# exit
NX-5000-1(config)# interface port-channel 100
NX-5000-1(config-if)# switchport mode fex-fabric
NX-5000-1(config-if)# vpc 100
NX-5000-1(config-if)# fex associate 100
NX-5000-1(config-if)# exit
```

Step 6 Configure the host interface ports on the Fabric Extender NX-2000-100.

```
switch(config)# interface ethernet 100/1/1-48
switch(config-if)# switchport mode access
switch(config-if)# switchport access vlan 50
switch(config-if)# no shutdown
switch(config-if)# exit
```

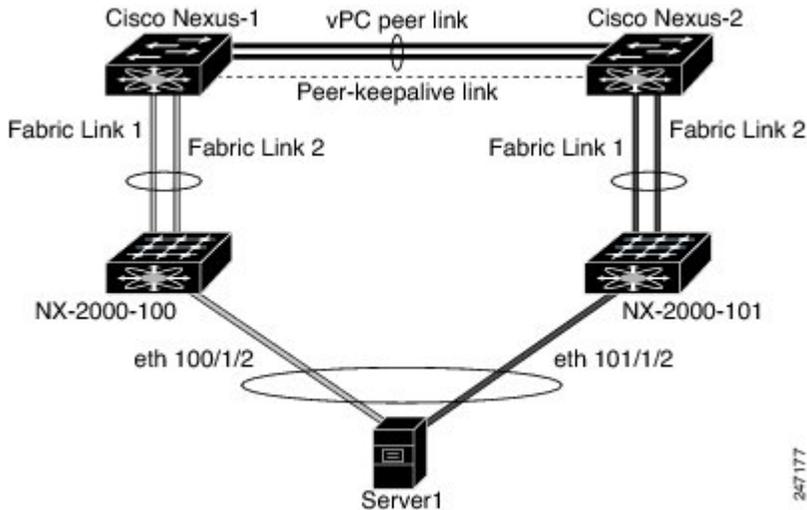
Step 7 Save the configuration.

```
switch(config)# copy running-config startup-config
```

Single Homed Fabric Extender vPC Configuration Example

The following example shows how to configure the single homed Fabric Extender vPC topology using the default VRF to carry the peer-keepalive messages on switch NX-5000-1 as shown in following figure:

Figure 18: vPC Configuration Example



Note

The following example only shows the configuration of NX-5000-1 which is connected to the Fabric Extender NX-2000-100. You must repeat these steps on its vPC peer, NX-5000-2, which is connected to the Fabric Extender NX-2000-101.

Before You Begin

Ensure that the Cisco Nexus 2000 Series Fabric Extenders NX-2000-100 and NX-2000-101 are attached and online.

SUMMARY STEPS

1. Enable vPC and LACP.
2. Enable SVI interfaces, create the VLAN and SVI to be used by the vPC peer-keepalive link.
3. Create the vPC domain and add the vPC peer-keepalive link in the default VRF.
4. Configure the vPC peer link as a two port Etherchannel.
5. Configure the Fabric Extender NX-2000-100.
6. Configure the fabric EtherChannel links for the Fabric Extender NX-2000-100.
7. Configure a vPC server port on on the Fabric Extender NX-2000-100.
8. Save the configuration.

DETAILED STEPS

Step 1 Enable vPC and LACP.

```
NX-5000-1# configure terminal
NX-5000-1(config)# feature lacp
NX-5000-1(config)# feature vpc
```

Step 2 Enable SVI interfaces, create the VLAN and SVI to be used by the vPC peer-keepalive link.

```
NX-5000-1(config)# feature interface-vlan
NX-5000-1(config)# vlan 900
NX-5000-1(config-vlan)# int vlan 900
NX-5000-1(config-if)# ip address 10.10.10.236 255.255.255.0
NX-5000-1(config-if)# no shutdown
NX-5000-1(config-if)# exit
```

Step 3 Create the vPC domain and add the vPC peer-keepalive link in the default VRF.

```
NX-5000-1(config)# vpc domain 30
NX-5000-1(config-vpc-domain)# peer-keepalive destination 10.10.10.237 source 10.10.10.236 vrf default
NX-5000-1(config-vpc-domain)# exit
```

Note VLAN 900 must **not** be trunked across the vPC peer-link because it carries the vPC peer-keepalive messages. There must be an alternative path between switches NX-5000-1 and NX-5000-2 for the vPC peer-keepalive messages.

Step 4 Configure the vPC peer link as a two port Etherchannel.

```
NX-5000-1(config)# interface ethernet 1/1-2
NX-5000-1(config-if-range)# switchport mode trunk
NX-5000-1(config-if-range)# switchport trunk allowed vlan 20-50
NX-5000-1(config-if-range)# switchport trunk native vlan 20
NX-5000-1(config-if-range)# channel-group 30 mode active
NX-5000-1(config-if-range)# exit
NX-5000-1(config)# interface port-channel 30
NX-5000-1(config-if)# vpc peer-link
NX-5000-1(config-if)# exit
```

Step 5 Configure the Fabric Extender NX-2000-100.

```
NX-5000-1(config)# fex 100
NX-5000-1(config-fex)# pinning max-links 1
NX-5000-1(fex)# exit
```

Step 6 Configure the fabric EtherChannel links for the Fabric Extender NX-2000-100.

```
NX-5000-1(config)# interface ethernet 1/20-21
NX-5000-1(config-if)# channel-group 100
NX-5000-1(config-if)# exit
NX-5000-1(config)# interface port-channel 100
NX-5000-1(config-if)# switchport mode fex-fabric
NX-5000-1(config-if)# fex associate 100
NX-5000-1(config-if)# exit
```

Step 7 Configure a vPC server port on the Fabric Extender NX-2000-100.

```
NX-5000-1(config-if)# interface ethernet 100/1/1
NX-5000-1(config-if)# switchport mode trunk
NX-5000-1(config-if)# switchport trunk native vlan 100
NX-5000-1(config-if)# switchport trunk allowed vlan 100-105
NX-5000-1(config-if)# channel-group 600
NX-5000-1(config-if)# no shutdown
NX-5000-1(config-if)# exit
NX-5000-1(config)# interface port-channel 600
NX-5000-1(config-if)# vpc 600
NX-5000-1(config-if)# no shutdown
NX-5000-1(config-if)# exit
```

Step 8 Save the configuration.

```
NX-5000-1(config)# copy running-config startup-config
```

vPC Default Settings

The following table lists the default settings for vPC parameters.

Table 13: Default vPC Parameters

Parameters	Default
vPC system priority	32667
vPC peer-keepalive message	Disabled
vPC peer-keepalive interval	1 second
vPC peer-keepalive timeout	5 seconds
vPC peer-keepalive UDP port	3200



Configuring Rapid PVST+

Rapid per VLAN Spanning Tree (Rapid PVST+) is an updated implementation of STP that allows you to create one spanning tree topology for each VLAN. Rapid PVST+ is the default Spanning Tree Protocol (STP) mode on the switch.



Note

Spanning tree is used to refer to IEEE 802.1w and IEEE 802.1s. If the text is discussing the IEEE 802.1D Spanning Tree Protocol, 802.1D is stated specifically.

This chapter describes the configuration of Rapid PVST+ on Cisco Nexus 5000 Series switches. It includes the following sections:

- [Information About Rapid PVST+, page 155](#)
- [Configuring Rapid PVST+, page 171](#)
- [Verifying Rapid PVST+ Configurations, page 181](#)

Information About Rapid PVST+

The Rapid PVST+ protocol is the IEEE 802.1w standard, Rapid Spanning Tree Protocol (RSTP), implemented on a per VLAN basis. Rapid PVST+ interoperates with the IEEE 802.1D standard, which mandates a single STP instance for all VLANs, rather than per VLAN.

Rapid PVST+ is enabled by default on the default VLAN (VLAN1) and on all newly created VLANs in software. Rapid PVST+ interoperates with switches that run legacy IEEE 802.1D STP.

RSTP is an improvement on the original STP standard, 802.1D, which allows faster convergence.

Related Topics

[Rapid PVST+ and IEEE 802.1Q Trunks, on page 170](#)

[Rapid PVST+ Interoperation with Legacy 802.1D STP, on page 170](#)

Understanding STP

STP Overview

For an Ethernet network to function properly, only one active path can exist between any two stations. STP operation is transparent to end stations, which cannot detect whether they are connected to a single LAN segment or a switched LAN of multiple segments.

When you create fault-tolerant internetworks, you must have a loop-free path between all nodes in a network. The STP algorithm calculates the best loop-free path throughout a switched network. LAN ports send and receive STP frames, which are called Bridge Protocol Data Units (BPDUs), at regular intervals. Switches do not forward these frames, but use the frames to construct a loop-free path.

Multiple active paths between end stations cause loops in the network. If a loop exists in the network, end stations might receive duplicate messages and switches might learn end station MAC addresses on multiple LAN ports. These conditions result in a broadcast storm, which creates an unstable network.

STP defines a tree with a root bridge and a loop-free path from the root to all switches in the network. STP forces redundant data paths into a blocked state. If a network segment in the spanning tree fails and a redundant path exists, the STP algorithm recalculates the spanning tree topology and activates the blocked path.

When two LAN ports on a switch are part of a loop, the STP port priority and port path cost setting determine which port on the switch is put in the forwarding state and which port is put in the blocking state.

Understanding How a Topology is Created

All switches in an extended LAN that participate in a spanning tree gather information about other switches in the network by exchanging of BPDUs. This exchange of BPDUs results in the following actions:

- The system elects a unique root switch for the spanning tree network topology.
- The system elects a designated switch for each LAN segment.
- The system eliminates any loops in the switched network by placing redundant interfaces in a backup state; all paths that are not needed to reach the root switch from anywhere in the switched network are placed in an STP-blocked state.

The topology on an active switched network is determined by the following:

- The unique switch identifier Media Access Control (MAC) address of the switch that is associated with each switch
- The path cost to the root that is associated with each interface
- The port identifier that is associated with each interface

In a switched network, the root switch is the logical center of the spanning tree topology. STP uses BPDUs to elect the root switch and root port for the switched network, as well as the root port and designated port for each switched segment.

Understanding the Bridge ID

Each VLAN on each switch has a unique 64-bit bridge ID consisting of a bridge priority value, an extended system ID (IEEE 802.1t), and an STP MAC address allocation.

Bridge Priority Value

The bridge priority is a 4-bit value when the extended system ID is enabled.



Note In Cisco NX-OS, the extended system ID is always enabled; you cannot be disable the extended system ID.

Related Topics

[Configuring the Rapid PVST+ Bridge Priority of a VLAN, on page 177](#)

Extended System ID

A 12-bit extended system ID field is part of the bridge ID.

Figure 19: Bridge ID with Extended System ID



The switches always use the 12-bit extended system ID.

Combined with the bridge ID, the system ID extension functions as the unique identifier for a VLAN.

Table 14: Bridge Priority Value and Extended System ID with the Extended System ID Enabled

Bridge Priority Value				Extended System ID (Set Equal to the VLAN ID)											
Bit 16	Bit 15	Bit 14	Bit 13	Bit 12	Bit 11	Bit 10	Bit 9	Bit 8	Bit 7	Bit 6	Bit 5	Bit 4	Bit 3	Bit 2	Bit 1
32768	16384	8192	4096	2048	1024	512	256	128	64	32	16	8	4	2	1

STP MAC Address Allocation



Note Extended system ID and MAC address reduction is always enabled on the software.

With MAC address reduction enabled on any switch, you should also enable MAC address reduction on all other connected switches to avoid undesirable root bridge election and spanning tree topology issues.

When MAC address reduction is enabled, the root bridge priority becomes a multiple of 4096 plus the VLAN ID. You can only specify a switch bridge ID (used by the spanning tree algorithm to determine the identity of the root bridge, the lowest being preferred) as a multiple of 4096. Only the following values are possible:

- 0
- 4096
- 8192

- 12288
- 16384
- 20480
- 24576
- 28672
- 32768
- 36864
- 40960
- 45056
- 49152
- 53248
- 57344
- 61440

STP uses the extended system ID plus a MAC address to make the bridge ID unique for each VLAN.



Note

If another bridge in the same spanning tree domain does not run the MAC address reduction feature, it could achieve root bridge ownership because its bridge ID may fall between the values specified by the MAC address reduction feature.

Understanding BPDUs

Switches transmit bridge protocol data units (BPDUs) throughout the STP instance. Each switch sends configuration BPDUs to communicate and compute the spanning tree topology. Each configuration BPDU contains the following minimal information:

- The unique bridge ID of the switch that the transmitting switch determines is the root bridge
- The STP path cost to the root
- The bridge ID of the transmitting bridge
- Message age
- The identifier of the transmitting port
- Values for the hello, forward delay, and max-age protocol timer
- Additional information for STP extension protocols

When a switch transmits a Rapid PVST+ BPDU frame, all switches connected to the VLAN on which the frame is transmitted receive the BPDU. When a switch receives a BPDU, it does not forward the frame but instead uses the information in the frame to calculate a BPDU, and, if the topology changes, initiate a BPDU transmission.

A BPDU exchange results in the following:

- One switch is elected as the root bridge.
- The shortest distance to the root bridge is calculated for each switch based on the path cost.
- A designated bridge for each LAN segment is selected. This is the switch closest to the root bridge through which frames are forwarded to the root.
- A root port is selected. This is the port providing the best path from the bridge to the root bridge.
- Ports included in the spanning tree are selected.

Related Topics

[Rapid PVST+ BPDUs, on page 162](#)

Election of the Root Bridge

For each VLAN, the switch with the highest bridge ID (that is, the lowest numerical ID value) is elected as the root bridge. If all switches are configured with the default priority (32768), the switch with the lowest MAC address in the VLAN becomes the root bridge. The bridge priority value occupies the most significant bits of the bridge ID.

When you change the bridge priority value, you change the probability that the switch will be elected as the root bridge. Configuring a lower value increases the probability; a higher value decreases the probability.

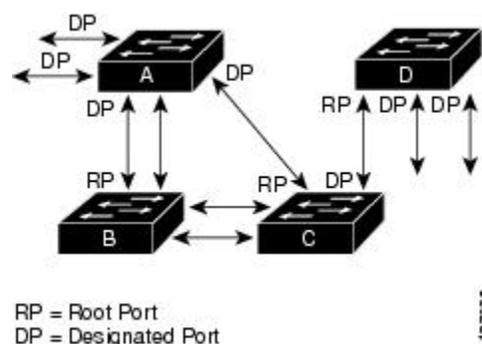
The STP root bridge is the logical center of each spanning tree topology in a network. All paths that are not needed to reach the root bridge from anywhere in the network are placed in STP blocking mode.

BPDUs contain information about the transmitting bridge and its ports, including bridge and MAC addresses, bridge priority, port priority, and path cost. STP uses this information to elect the root bridge for the STP instance, to elect the root port leading to the root bridge, and to determine the designated port for each segment.

Creating the Spanning Tree Topology

In the following figure, Switch A is elected as the root bridge because the bridge priority of all the switches is set to the default (32768) and Switch A has the lowest MAC address. However, due to traffic patterns, number of forwarding ports, or link types, Switch A might not be the ideal root bridge. By increasing the priority (lowering the numerical value) of the ideal switch so that it becomes the root bridge, you force an STP recalculation to form a new spanning tree topology with the ideal switch as the root.

Figure 20: Spanning Tree Topology



When the spanning tree topology is calculated based on default parameters, the path between source and destination end stations in a switched network might not be ideal. For instance, connecting higher-speed links to a port that has a higher number than the current root port can cause a root-port change. The goal is to make the fastest link the root port.

For example, assume that one port on Switch B is a fiber-optic link, and another port on Switch B (an unshielded twisted-pair [UTP] link) is the root port. Network traffic might be more efficient over the high-speed fiber-optic link. By changing the STP port priority on the fiber-optic port to a higher priority (lower numerical value) than the root port, the fiber-optic port becomes the new root port.

Understanding Rapid PVST+

Rapid PVST+ Overview

Rapid PVST+ is the IEEE 802.1w (RSTP) standard implemented per VLAN. A single instance of STP runs on each configured VLAN (if you do not manually disable STP). Each Rapid PVST+ instance on a VLAN has a single root switch. You can enable and disable STP on a per-VLAN basis when you are running Rapid PVST+.



Note Rapid PVST+ is the default STP mode for the switch.

Rapid PVST+ uses point-to-point wiring to provide rapid convergence of the spanning tree. The spanning tree reconfiguration can occur in less than 1 second with Rapid PVST+ (in contrast to 50 seconds with the default settings in the 802.1D STP).



Note Rapid PVST+ supports one STP instance for each VLAN.

Using Rapid PVST+, STP convergence occurs rapidly. Each designated or root port in the STP sends out a BPDU every 2 seconds by default. On a designated or root port in the topology, if hello messages are missed three consecutive times, or if the maximum age expires, the port immediately flushes all protocol information in the table. A port considers that it loses connectivity to its direct neighbor root or designated port if it misses three BPDUs or if the maximum age expires. This rapid aging of the protocol information allows quick failure detection. The switch automatically checks the PVID.

Rapid PVST+ provides for rapid recovery of connectivity following the failure of a network device, a switch port, or a LAN. It provides rapid convergence for edge ports, new root ports, and ports connected through point-to-point links as follows:

- Edge ports—When you configure a port as an edge port on an RSTP switch, the edge port immediately transitions to the forwarding state. (This immediate transition was previously a Cisco-proprietary feature named PortFast.) You should only configure on ports that connect to a single end station as edge ports. Edge ports do not generate topology changes when the link changes.

Enter the **spanning-tree port type** interface configuration command to configure a port as an STP edge port.



Note We recommend that you configure all ports connected to a host as edge ports.

- Root ports—If Rapid PVST+ selects a new root port, it blocks the old root port and immediately transitions the new root port to the forwarding state.
- Point-to-point links—If you connect a port to another port through a point-to-point link and the local port becomes a designated port, it negotiates a rapid transition with the other port by using the proposal-agreement handshake to ensure a loop-free topology.

Rapid PVST+ achieves rapid transition to the forwarding state only on edge ports and point-to-point links. Although the link type is configurable, the system automatically derives the link type information from the duplex setting of the port. Full-duplex ports are assumed to be point-to-point ports, while half-duplex ports are assumed to be shared ports.

Edge ports do not generate topology changes, but all other designated and root ports generate a topology change (TC) BPDU when they either fail to receive three consecutive BPDUs from the directly connected neighbor or the maximum age times out. At this point, the designated or root port sends out a BPDU with the TC flag set. The BPDUs continue to set the TC flag as long as the TC While timer runs on that port. The value of the TC While timer is the value set for the hello time plus 1 second. The initial detector of the topology change immediately floods this information throughout the entire topology.

When Rapid PVST+ detects a topology change, the protocol does the following:

- Starts the TC While timer with a value equal to twice the hello time for all the non-edge root and designated ports, if necessary.
- Flushes the MAC addresses associated with all these ports.

The topology change notification floods quickly across the entire topology. The system flushes dynamic entries immediately on a per-port basis when it receives a topology change.


Note

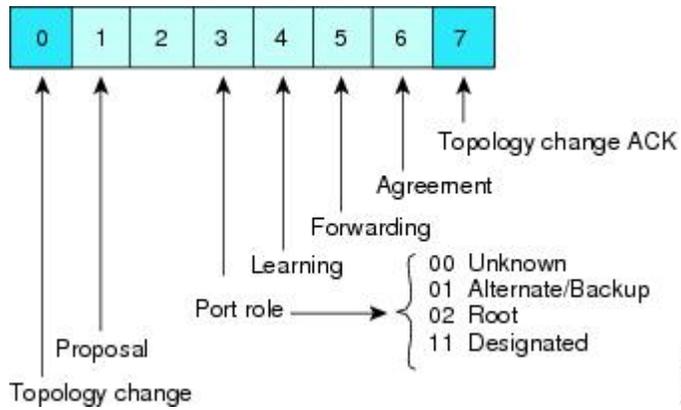
The TCA flag is used only when the switch is interacting with switches that are running legacy 802.1D STP.

The proposal and agreement sequence then quickly propagates toward the edge of the network and quickly restores connectivity after a topology change.

Rapid PVST+ BPDUs

Rapid PVST+ and 802.1w use all six bits of the flag byte to add the role and state of the port that originates the BPDU, and the proposal and agreement handshake. The following figure shows the use of the BPDU flags in Rapid PVST+.

Figure 21: Rapid PVST+ Flag Byte in BPDU

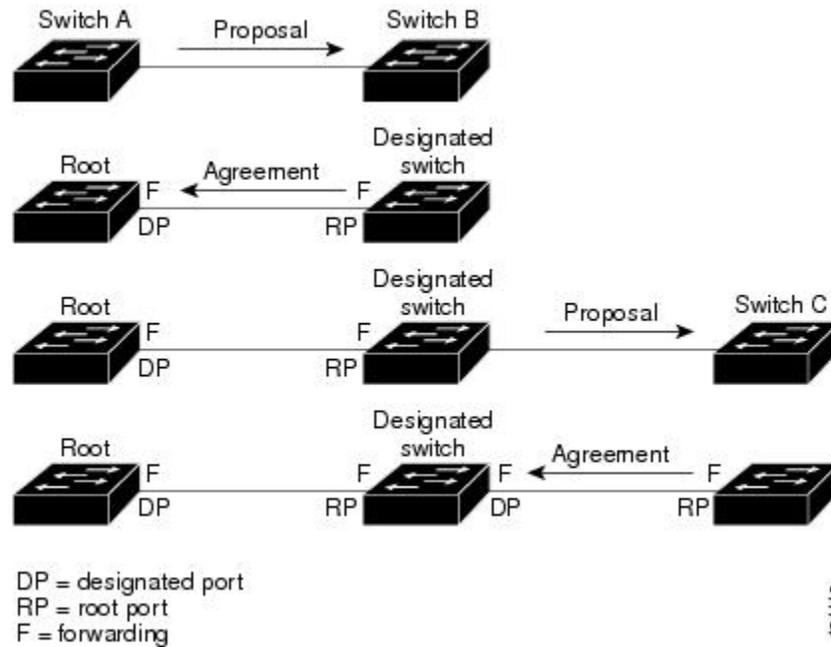


Another important change is that the Rapid PVST+ BPDU is type 2, version 2, which makes it possible for the switch to detect connected legacy (802.1D) bridges. The BPDU for 802.1D is version 0.

Proposal and Agreement Handshake

As shown in the following figure, switch A is connected to switch B through a point-to-point link, and all of the ports are in the blocking state. Assume that the priority of switch A is a smaller numerical value than the priority of switch B.

Figure 22: Proposal and Agreement Handshaking for Rapid Convergence



Switch A sends a proposal message (a configuration BPDU with the proposal flag set) to switch B, proposing itself as the designated switch.

After receiving the proposal message, switch B selects as its new root port the port from which the proposal message was received, forces all non-edge ports to the blocking state, and sends an agreement message (a BPDU with the agreement flag set) through its new root port.

After receiving the agreement message from switch B, switch A also immediately transitions its designated port to the forwarding state. No loops in the network can form because switch B blocked all of its non-edge ports and because there is a point-to-point link between switches A and B.

When switch C connects to switch B, a similar set of handshaking messages are exchanged. Switch C selects the port connected to switch B as its root port, and both ends of the link immediately transition to the forwarding state. With each iteration of this handshaking process, one more network device joins the active topology. As the network converges, this proposal-agreement handshaking progresses from the root toward the leaves of the spanning tree.

The switch learns the link type from the port duplex mode: a full-duplex port is considered to have a point-to-point connection and a half-duplex port is considered to have a shared connection. You can override the default setting that is controlled by the duplex setting by entering the **spanning-tree link-type** interface configuration command.

This proposal/agreement handshake is initiated only when a non-edge port moves from the blocking to the forwarding state. The handshaking process then proliferates step-by-step throughout the topology.

Related Topics

[Summary of Port States, on page 167](#)

Protocol Timers

The following table describes the protocol timers that affect the Rapid PVST+ performance.

Table 15: Rapid PVST+ Protocol Timers

Variable	Description
Hello timer	Determines how often each switch broadcasts BPDUs to other switches. The default is 2 seconds, and the range is from 1 to 10.
Forward delay timer	Determines how long each of the listening and learning states last before the port begins forwarding. This timer is generally not used by the protocol but is used as a backup. The default is 15 seconds, and the range is from 4 to 30 seconds.
Maximum age timer	Determines the amount of time protocol information received on an port is stored by the switch. This timer is generally not used by the protocol, but it is used when interoperating with 802.1D spanning tree. The default is 20 seconds; the range is from 6 to 40 seconds.

Port Roles

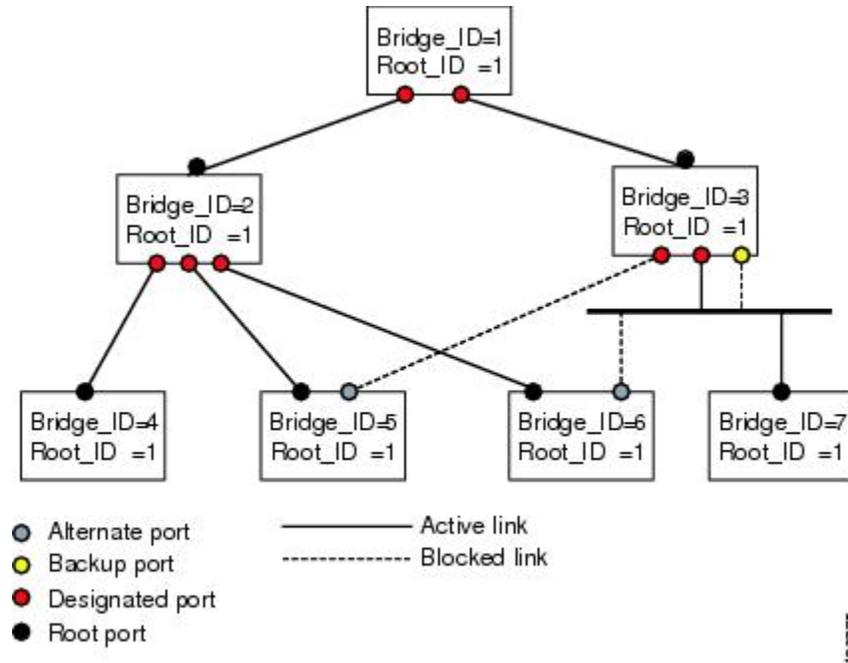
Rapid PVST+ provides rapid convergence of the spanning tree by assigning port roles and learning the active topology. Rapid PVST+ builds upon the 802.1D STP to select the switch with the highest priority (lowest numerical priority value) as the root bridge. Rapid PVST+ then assigns one of these port roles to individual ports:

- Root port—Provides the best path (lowest cost) when the switch forwards packets to the root bridge.
- Designated port—Connects to the designated switch, which incurs the lowest path cost when forwarding packets from that LAN to the root bridge. The port through which the designated switch is attached to the LAN is called the designated port.
- Alternate port—Offers an alternate path toward the root bridge to the path provided by the current root port. An alternate port provides a path to another switch in the topology.
- Backup port—Acts as a backup for the path provided by a designated port toward the leaves of the spanning tree. A backup port can exist only when two ports are connected in a loopback by a point-to-point link or when a switch has two or more connections to a shared LAN segment. A backup port provides another path in the topology to the switch.
- Disabled port—Has no role within the operation of the spanning tree.

In a stable topology with consistent port roles throughout the network, Rapid PVST+ ensures that every root port and designated port immediately transition to the forwarding state while all alternate and backup ports are always in the blocking state. Designated ports start in the blocking state. The port state controls the operation of the forwarding and learning processes.

A port with the root or a designated port role is included in the active topology. A port with the alternate or backup port role is excluded from the active topology (see the following figure).

Figure 23: Sample Topology Demonstrating Port Roles



Related Topics

[Election of the Root Bridge, on page 159](#)

Port States

Rapid PVST+ Port State Overview

Propagation delays can occur when protocol information passes through a switched LAN. As a result, topology changes can take place at different times and at different places in a switched network. When a LAN port transitions directly from nonparticipation in the spanning tree topology to the forwarding state, it can create temporary data loops. Ports must wait for new topology information to propagate through the switched LAN before starting to forward frames.

Each LAN port on a software using Rapid PVST+ or MST exists in one of the following four states:

- Blocking—The LAN port does not participate in frame forwarding.
- Learning—The LAN port prepares to participate in frame forwarding.
- Forwarding—The LAN port forwards frames.
- Disabled—The LAN port does not participate in STP and is not forwarding frames.

When you enable Rapid PVST+, every port in the software, VLAN, and network goes through the blocking state and the transitory states of learning at power up. If properly configured, each LAN port stabilizes to the forwarding or blocking state.

When the STP algorithm places a LAN port in the forwarding state, the following process occurs:

- The LAN port is put into the blocking state while it waits for protocol information that suggests it should go to the learning state.
- The LAN port waits for the forward delay timer to expire, moves the LAN port to the learning state, and restarts the forward delay timer.
- In the learning state, the LAN port continues to block frame forwarding as it learns the end station location information for the forwarding database.
- The LAN port waits for the forward delay timer to expire and then moves the LAN port to the forwarding state, where both learning and frame forwarding are enabled.

Blocking State

A LAN port in the blocking state does not participate in frame forwarding.

A LAN port in the blocking state performs as follows:

- Discards frames received from the attached segment.
- Discards frames switched from another port for forwarding.
- Does not incorporate the end station location into its address database. (There is no learning on a blocking LAN port, so there is no address database update.)
- Receives BPDUs and directs them to the system module.
- Receives, processes, and transmits BPDUs received from the system module.
- Receives and responds to network management messages.

Learning State

A LAN port in the learning state prepares to participate in frame forwarding by learning the MAC addresses for the frames. The LAN port enters the learning state from the blocking state.

A LAN port in the learning state performs as follows:

- Discards frames received from the attached segment.
- Discards frames switched from another port for forwarding.
- Incorporates the end station location into its address database.
- Receives BPDUs and directs them to the system module.
- Receives, processes, and transmits BPDUs received from the system module.
- Receives and responds to network management messages.

Forwarding State

A LAN port in the forwarding state forwards frames. The LAN port enters the forwarding state from the learning state.

A LAN port in the forwarding state performs as follows:

- Forwards frames received from the attached segment.
- Forwards frames switched from another port for forwarding.
- Incorporates the end station location information into its address database.
- Receives BPDUs and directs them to the system module.
- Processes BPDUs received from the system module.
- Receives and responds to network management messages.

Disabled State

A LAN port in the disabled state does not participate in frame forwarding or STP. A LAN port in the disabled state is virtually nonoperational.

A disabled LAN port performs as follows:

- Discards frames received from the attached segment.
- Discards frames switched from another port for forwarding.
- Does not incorporate the end station location into its address database. (There is no learning, so there is no address database update.)
- Does not receive BPDUs from neighbors.
- Does not receive BPDUs for transmission from the system module.

Summary of Port States

The following table lists the possible operational and Rapid PVST+ states for ports and the corresponding inclusion in the active topology.

Table 16: Port State Active Topology

Operational Status	Port State	Is Port Included in the Active Topology?
Enabled	Blocking	No
Enabled	Learning	Yes
Enabled	Forwarding	Yes
Disabled	Disabled	No

Synchronization of Port Roles

When the switch receives a proposal message on one of its ports and that port is selected as the new root port, Rapid PVST+ forces all other ports to synchronize with the new root information.

The switch is synchronized with superior root information received on the root port if all other ports are synchronized. An individual port on the switch is synchronized if either of the following applies:

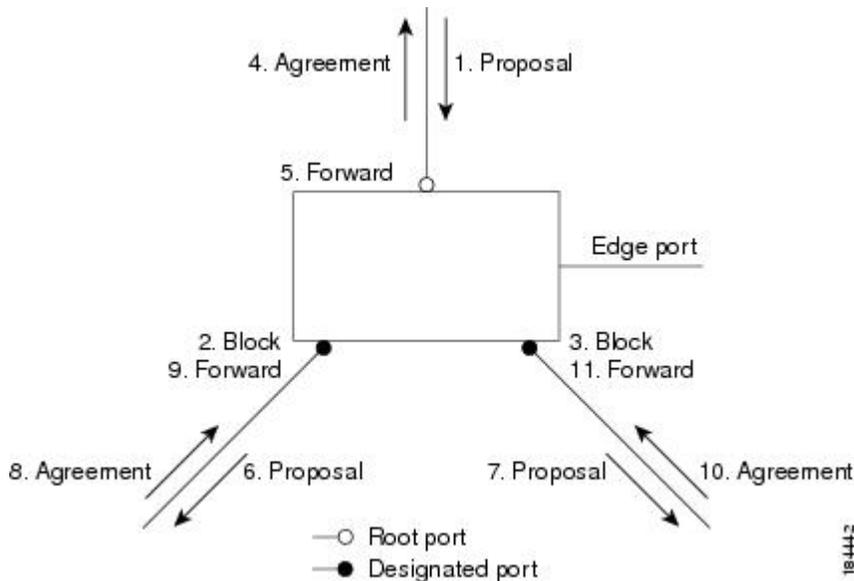
- That port is in the blocking state.

- It is an edge port (a port configured to be at the edge of the network).

If a designated port is in the forwarding state and is not configured as an edge port, it transitions to the blocking state when the Rapid PVST+ forces it to synchronize with new root information. In general, when the Rapid PVST+ forces a port to synchronize with root information and the port does not satisfy any of the above conditions, its port state is set to blocking.

After ensuring that all of the ports are synchronized, the switch sends an agreement message to the designated switch that corresponds to its root port. When the switches connected by a point-to-point link are in agreement about their port roles, Rapid PVST+ immediately transitions the port states to the forwarding state. The sequence of events is shown in the following figure.

Figure 24: Sequence of Events During Rapid Convergence



Processing Superior BPDUs Information

A superior BPDUs is a BPDUs with root information (such as a lower switch ID or lower path cost) that is superior to what is currently stored for the port.

If a port receives a superior BPDUs, Rapid PVST+ triggers a reconfiguration. If the port is proposed and is selected as the new root port, Rapid PVST+ forces all the other ports to synchronize.

If the received BPDUs is a Rapid PVST+ BPDUs with the proposal flag set, the switch sends an agreement message after all of the other ports are synchronized. The new root port transitions to the forwarding state as soon as the previous port reaches the blocking state.

If the superior information received on the port causes the port to become a backup port or an alternate port, Rapid PVST+ sets the port to the blocking state and sends an agreement message. The designated port continues sending BPDUs with the proposal flag set until the forward-delay timer expires. At that time, the port transitions to the forwarding state.

Processing Inferior BPDUs Information

An inferior BPDUs is a BPDUs with root information (such as a higher switch ID or higher path cost) that is inferior to what is currently stored for the port.

If a designated port receives an inferior BPDUs, it immediately replies with its own information.

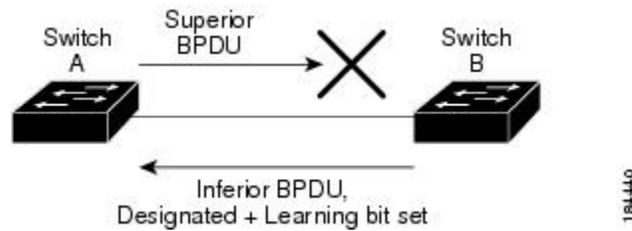
Detecting Unidirectional Link Failure

The software checks the consistency of the port role and state in the received BPDUs to detect unidirectional link failures that could cause bridging loops.

When a designated port detects a conflict, it keeps its role, but reverts to a discarding state because disrupting connectivity in case of inconsistency is preferable to opening a bridging loop.

The following figure illustrates a unidirectional link failure that typically creates a bridging loop. Switch A is the root bridge, and its BPDUs are lost on the link leading to switch B. The 802.1w-standard BPDUs include the role and state of the sending port. With this information, switch A can detect that switch B does not react to the superior BPDUs it sends and that switch B is the designated, not root port. As a result, switch A blocks (or keeps blocking) its port, thus preventing the bridging loop. The block is shown as an STP dispute.

Figure 25: Detecting Unidirectional Link Failure



Port Cost



Note

Rapid PVST+ uses the short (16-bit) pathcost method to calculate the cost by default. With the short pathcost method, you can assign any value in the range of 1 to 65535. However, you can configure the switch to use the long (32-bit) pathcost method, which allows you to assign any value in the range of 1 to 200,000,000. You configure the pathcost calculation method globally.

The STP port path-cost default value is determined from the media speed and path-cost calculation method of a LAN interface. If a loop occurs, STP considers the port cost when selecting a LAN interface to put into the forwarding state.

Table 17: Default Port Cost

Bandwidth	Short Path-cost Method of Port Cost	Long Path-cost Method of Port Cost
10 Mbps	100	2,000,000
100 Mbps	19	200,000
1 Gigabit Ethernet	4	20,000
10 Gigabit Ethernet	2	2,000

You can assign lower cost values to LAN interfaces that you want STP to select first and higher cost values to LAN interfaces that you want STP to select last. If all LAN interfaces have the same cost value, STP puts the LAN interface with the lowest LAN interface number in the forwarding state and blocks other LAN interfaces.

On access ports, you assign port cost by the port. On trunk ports, you assign the port cost by the VLAN; you can configure the same port cost to all the VLANs on a trunk port.

Port Priority

If a loop occurs and multiple ports have the same path cost, Rapid PVST+ considers the port priority when selecting which LAN port to put into the forwarding state. You can assign lower priority values to LAN ports that you want Rapid PVST+ to select first and higher priority values to LAN ports that you want Rapid PVST+ to select last.

If all LAN ports have the same priority value, Rapid PVST+ puts the LAN port with the lowest LAN port number in the forwarding state and blocks other LAN ports. The possible priority range is from 0 through 224 (the default is 128), configurable in increments of 32. software uses the port priority value when the LAN port is configured as an access port and uses VLAN port priority values when the LAN port is configured as a trunk port.

Rapid PVST+ and IEEE 802.1Q Trunks

802.1Q trunks impose some limitations on the STP strategy for a network. In a network of Cisco switches connected through 802.1Q trunks, the switches maintain one instance of STP for each VLAN allowed on the trunks. However, non-Cisco 802.1Q switches maintain only one instance of STP for all VLANs allowed on the trunks.

When you connect a Cisco switch to a non-Cisco switch through an 802.1Q trunk, the Cisco switch combines the STP instance of the 802.1Q VLAN of the trunk with the STP instance of the non-Cisco 802.1Q switch. However, all per-VLAN STP information that is maintained by Cisco switches is separated by a cloud of non-Cisco 802.1Q switches. The non-Cisco 802.1Q cloud that separates the Cisco switches is treated as a single trunk link between the switches.

Rapid PVST+ Interoperation with Legacy 802.1D STP

Rapid PVST+ can interoperate with switches that are running the legacy 802.1D protocol. The switch knows that it is interoperating with equipment running 802.1D when it receives a BPDU version 0. The BPDUs for Rapid PVST+ are version 2. If the BPDU received is an 802.1w BPDU version 2 with the proposal flag set, the switch sends an agreement message after all of the other ports are synchronized. If the BPDU is an 802.1D BPDU version 0, the switch does not set the proposal flag and starts the forward-delay timer for the port. The new root port requires twice the forward-delay time to transition to the forwarding state.

The switch interoperates with legacy 802.1D switches as follows:

- **Notification**—Unlike 802.1D BPDUs, 802.1w does not use TCN BPDUs. However, for interoperability with 802.1D switches, Cisco NX-OS processes and generates TCN BPDUs.
- **Acknowledgement**—When an 802.1w switch receives a TCN message on a designated port from an 802.1D switch, it replies with an 802.1D configuration BPDU with the TCA bit set. However, if the TC-while timer (the same as the TC timer in 802.1D) is active on a root port connected to an 802.1D switch and a configuration BPDU with the TCA set is received, the TC-while timer is reset.

This method of operation is required only for 802.1D switches. The 802.1w BPDUs do not have the TCA bit set.

- Protocol migration—For backward compatibility with 802.1D switches, 802.1w selectively sends 802.1D configuration BPDUs and TCN BPDUs on a per-port basis.

When a port is initialized, the migrate-delay timer is started (specifies the minimum time during which 802.1w BPDUs are sent), and 802.1w BPDUs are sent. While this timer is active, the switch processes all BPDUs received on that port and ignores the protocol type.

If the switch receives an 802.1D BPDU after the port migration-delay timer has expired, it assumes that it is connected to an 802.1D switch and starts using only 802.1D BPDUs. However, if the 802.1w switch is using 802.1D BPDUs on a port and receives an 802.1w BPDU after the timer has expired, it restarts the timer and starts using 802.1w BPDUs on that port.



Note If you want all switches to renegotiate the protocol, you must restart Rapid PVST+.

Related Topics

[Restarting the Protocol](#) , on page 208

Rapid PVST+ Interoperation with 802.1s MST

Rapid PVST+ interoperates seamlessly with the IEEE 802.1s Multiple Spanning Tree (MST) standard. No user configuration is needed.

Configuring Rapid PVST+

Rapid PVST+, which has the 802.1w standard applied to the Rapid PVST+ protocol, is the default STP setting in the software.

You enable Rapid PVST+ on a per-VLAN basis. The software maintains a separate instance of STP for each VLAN (except on those VLANs on which you disable STP). By default, Rapid PVST+ is enabled on the default VLAN and on each VLAN that you create.

Enabling Rapid PVST+

Once you enable Rapid PVST+ on the switch, you must enable Rapid PVST+ on the specified VLANs.

Rapid PVST+ is the default STP mode. You cannot simultaneously run MST and Rapid PVST+.



Note Changing the spanning tree mode disrupts traffic because all spanning tree instances are stopped for the previous mode and started for the new mode.

SUMMARY STEPS

1. switch# **configure terminal**
2. switch(config)# **spanning-tree mode rapid-pvst**

DETAILED STEPS

	Command or Action	Purpose
Step 1	switch# configure terminal	Enters configuration mode.
Step 2	switch(config)# spanning-tree mode rapid-pvst	Enables Rapid PVST+ on the switch. Rapid PVST+ is the default spanning tree mode. Note Changing the spanning tree mode disrupts traffic because all spanning tree instances are stopped for the previous mode and started for the new mode.

This example shows how to enable Rapid PVST+ on the switch:

```
switch# configure terminal
switch(config)# spanning-tree mode rapid-pvst
```



Note Because STP is enabled by default, entering the **show running-config** command to view the resulting configuration does not display the command that you entered to enable Rapid PVST+.

Enabling Rapid PVST+ per VLAN

You can enable or disable Rapid PVST+ on each VLAN.



Note Rapid PVST+ is enabled by default on the default VLAN and on all VLANs that you create.

SUMMARY STEPS

1. switch# **configure terminal**
2. switch(config)# **spanning-tree vlan-range**
3. (Optional) switch(config)# **no spanning-tree vlan-range**

DETAILED STEPS

	Command or Action	Purpose
Step 1	switch# configure terminal	Enters configuration mode.
Step 2	switch(config)# spanning-tree <i>vlan-range</i>	Enables Rapid PVST+ (default STP) on a per VLAN basis. The <i>vlan-range</i> value can be 2 through 4094 (except reserved VLAN values).
Step 3	switch(config)# no spanning-tree <i>vlan-range</i>	(Optional) Disables Rapid PVST+ on the specified VLAN. Caution Do not disable spanning tree on a VLAN unless all switches and bridges in the VLAN have spanning tree disabled. You cannot disable spanning tree on some of the switches and bridges in a VLAN and leave it enabled on other switches and bridges. This action can have unexpected results because switches and bridges with spanning tree enabled will have incomplete information regarding the physical topology of the network. Do not disable spanning tree in a VLAN without ensuring that there are no physical loops present in the VLAN. Spanning tree serves as a safeguard against misconfigurations and cabling errors.

This example shows how to enable STP on a VLAN:

```
switch# configure terminal
switch(config)# spanning-tree vlan 5
```

Configuring the Root Bridge ID

The software maintains a separate instance of STP for each active VLAN in Rapid PVST+. For each VLAN, the switch with the lowest bridge ID becomes the root bridge for that VLAN.

To configure a VLAN instance to become the root bridge, modify the bridge priority from the default value (32768) to a significantly lower value.

When you enter the **spanning-tree vlan *vlan_ID* root** command, the switch checks the bridge priority of the current root bridges for each VLAN. The switch sets the bridge priority for the specified VLANs to 24576 if this value will cause the switch to become the root for the specified VLANs. If any root bridge for the specified VLANs has a bridge priority lower than 24576, the switch sets the bridge priority for the specified VLANs to 4096 less than the lowest bridge priority.


Note

The **spanning-tree vlan *vlan_ID* root** command fails if the value required to be the root bridge is less than 1.


Caution

The root bridge for each instance of STP should be a backbone or distribution switch. Do not configure an access switch as the STP primary root.

Enter the **diameter** keyword to specify the network diameter (that is, the maximum number of bridge hops between any two end stations in the network). When you specify the network diameter, the software automatically selects an optimal hello time, forward delay time, and maximum age time for a network of that diameter, which can significantly reduce the STP convergence time. You can enter the **hello-time** keyword to override the automatically calculated hello time.



Note With the switch configured as the root bridge, do not manually configure the hello time, forward-delay time, and maximum-age time using the **spanning-tree mst hello-time**, **spanning-tree mst forward-time**, and **spanning-tree mst max-age** configuration commands.

SUMMARY STEPS

1. switch# **configure terminal**
2. switch(config)# **spanning-tree vlan** *vlan-range* **root primary** [**diameter** *dia* [**hello-time** *hello-time*]]

DETAILED STEPS

	Command or Action	Purpose
Step 1	switch# configure terminal	Enters configuration mode.
Step 2	switch(config)# spanning-tree vlan <i>vlan-range</i> root primary [diameter <i>dia</i> [hello-time <i>hello-time</i>]]	Configures a software switch as the primary root bridge. The <i>vlan-range</i> value can be 2 through 4094 (except reserved VLAN values.) The <i>dia</i> default is 7. The <i>hello-time</i> can be from 1 to 10 seconds, and the default value is 2 seconds.

This example shows how to configure the switch as the root bridge for a VLAN:

```
switch# configure terminal
switch(config)# spanning-tree vlan 5 root primary diameter 4
```

Configuring a Secondary Root Bridge

When you configure a software switch as the secondary root, the STP bridge priority is modified from the default value (32768) so that the switch is likely to become the root bridge for the specified VLANs if the primary root bridge fails (assuming the other switches in the network use the default bridge priority of 32768). STP sets the bridge priority to 28672.

Enter the **diameter** keyword to specify the network diameter (that is, the maximum number of bridge hops between any two end stations in the network). When you specify the network diameter, the software automatically selects an optimal hello time, forward delay time, and maximum age time for a network of that diameter, which can significantly reduce the STP convergence time. You can enter the **hello-time** keyword to override the automatically calculated hello time.

You configure more than one switch in this manner to have multiple backup root bridges. Enter the same network diameter and hello time values that you used when configuring the primary root bridge.

**Note**

With the switch configured as the root bridge, do not manually configure the hello time, forward-delay time, and maximum-age time using the **spanning-tree mst hello-time**, **spanning-tree mst forward-time**, and **spanning-tree mst max-age** global configuration commands.

SUMMARY STEPS

1. switch# **configure terminal**
2. switch(config)# **spanning-tree vlan *vlan-range* root secondary [diameter *dia* [hello-time *hello-time*]]**

DETAILED STEPS

	Command or Action	Purpose
Step 1	switch# configure terminal	Enters configuration mode.
Step 2	switch(config)# spanning-tree vlan <i>vlan-range</i> root secondary [diameter <i>dia</i> [hello-time <i>hello-time</i>]]	Configures a software switch as the secondary root bridge. The <i>vlan-range</i> value can be 2 through 4094 (except reserved VLAN values.) The <i>dia</i> default is 7. The <i>hello-time</i> can be from 1 to 10 seconds, and the default value is 2 seconds.

This example shows how to configure the switch as the secondary root bridge for a VLAN:

```
switch# configure terminal
switch(config)# spanning-tree vlan 5 root secondary diameter 4
```

Configuring the Rapid PVST+ Port Priority

You can assign lower priority values to LAN ports that you want Rapid PVST+ to select first and higher priority values to LAN ports that you want Rapid PVST+ to select last. If all LAN ports have the same priority value, Rapid PVST+ puts the LAN port with the lowest LAN port number in the forwarding state and blocks other LAN ports.

The software uses the port priority value when the LAN port is configured as an access port and uses VLAN port priority values when the LAN port is configured as a trunk port.

SUMMARY STEPS

1. switch# **configure terminal**
2. switch(config)# **interface *type slot/port***
3. switch(config-if)# **spanning-tree [vlan *vlan-list*] port-priority *priority***

DETAILED STEPS

	Command or Action	Purpose
Step 1	switch# configure terminal	Enters configuration mode.
Step 2	switch(config)# interface <i>type slot/port</i>	Specifies the interface to configure, and enters interface configuration mode.
Step 3	switch(config-if)# spanning-tree [vlan <i>vlan-list</i>] port-priority <i>priority</i>	Configures the port priority for the LAN interface. The <i>priority</i> value can be from 0 to 224. The lower the value, the higher the priority. The priority values are 0, 32, 64, 96, 128, 160, 192, and 224. All other values are rejected. The default value is 128.

This example shows how to configure the access port priority of an Ethernet interface:

```
switch# configure terminal
switch(config)# interface ethernet 1/4
switch(config-if)# spanning-tree port-priority 160
```

You can only apply this command to a physical Ethernet interface.

Configuring the Rapid PVST+ Pathcost Method and Port Cost

On access ports, you assign port cost by the port. On trunk ports, you assign the port cost by VLAN; you can configure the same port cost on all the VLANs on a trunk.


Note

In Rapid PVST+ mode, you can use either the short or long pathcost method, and you can configure the method in either the interface or configuration submenu. The default pathcost method is short.

SUMMARY STEPS

1. switch# **configure terminal**
2. switch(config)# **spanning-tree pathcost method** {**long** | **short**}
3. switch(config)# **interface** *type slot/port*
4. switch(config-if)# **spanning-tree** [**vlan** *vlan-id*] **cost** [*value* | **auto**]

DETAILED STEPS

	Command or Action	Purpose
Step 1	switch# configure terminal	Enters configuration mode.
Step 2	switch(config)# spanning-tree pathcost method { long short }	Selects the method used for Rapid PVST+ pathcost calculations. The default method is the short method.

	Command or Action	Purpose
Step 3	switch(config)# interface <i>type slot/port</i>	Specifies the interface to configure, and enters interface configuration mode.
Step 4	switch(config-if)# spanning-tree [vlan <i>vlan-id</i>] cost [<i>value</i> auto]	Configures the port cost for the LAN interface. The cost value, depending on the pathcost calculation method, can be as follows: <ul style="list-style-type: none"> • short—1 to 65535 • long—1 to 200000000 <p>Note You configure this parameter per interface on access ports and per VLAN on trunk ports. The default is auto, which sets the port cost on both the pathcost calculation method and the media speed.</p>

This example shows how to configure the access port cost of an Ethernet interface:

```
switch# configure terminal
switch (config)# spanning-tree pathcost method long
switch (config)# interface ethernet 1/4
switch(config-if)# spanning-tree cost 1000
```

You can only apply this command to a physical Ethernet interface.

Configuring the Rapid PVST+ Bridge Priority of a VLAN

You can configure the Rapid PVST+ bridge priority of a VLAN.



Note

Be careful when using this configuration. For most situations, we recommend that you configure the primary root and secondary root to modify the bridge priority.

SUMMARY STEPS

1. switch# **configure terminal**
2. switch(config)# **spanning-tree vlan** *vlan-range* **priority** *value*

DETAILED STEPS

	Command or Action	Purpose
Step 1	switch# configure terminal	Enters configuration mode.
Step 2	switch(config)# spanning-tree vlan <i>vlan-range</i> priority <i>value</i>	Configures the bridge priority of a VLAN. Valid values are 0, 4096, 8192, 12288, 16384, 20480, 24576, 28672, 32768, 36864, 40960, 45056, 49152, 53248, 57344, and 61440. All other values are rejected. The default value is 32768.

Command or Action	Purpose
-------------------	---------

This example shows how to configure the bridge priority of a VLAN:

```
switch# configure terminal
switch(config)# spanning-tree vlan 5 priority 8192
```

Configuring the Rapid PVST+ Hello Time for a VLAN

You can configure the Rapid PVST+ hello time for a VLAN.



Note

Be careful when using this configuration. For most situations, we recommend that you configure the primary root and secondary root to modify the hello time.

SUMMARY STEPS

1. switch# **configure terminal**
2. switch(config)# **spanning-tree vlan *vlan-range* hello-time *hello-time***

DETAILED STEPS

	Command or Action	Purpose
Step 1	switch# configure terminal	Enters configuration mode.
Step 2	switch(config)# spanning-tree vlan <i>vlan-range</i> hello-time <i>hello-time</i>	Configures the hello time of a VLAN. The hello time value can be from 1 to 10 seconds. The default is 2 seconds.

This example shows how to configure the hello time for a VLAN:

```
switch# configure terminal
switch(config)# spanning-tree vlan 5 hello-time 7
```

Configuring the Rapid PVST+ Forward Delay Time for a VLAN

You can configure the forward delay time per VLAN when using Rapid PVST+.

SUMMARY STEPS

1. switch# **configure terminal**
2. switch(config)# **spanning-tree vlan *vlan-range* forward-time *forward-time***

DETAILED STEPS

	Command or Action	Purpose
Step 1	switch# configure terminal	Enters configuration mode.
Step 2	switch(config)# spanning-tree vlan <i>vlan-range</i> forward-time <i>forward-time</i>	Configures the forward delay time of a VLAN. The forward delay time value can be from 4 to 30 seconds, and the default is 15 seconds.

This example shows how to configure the forward delay time for a VLAN:

```
switch# configure terminal
switch(config)# spanning-tree vlan 5 forward-time 21
```

Configuring the Rapid PVST+ Maximum Age Time for a VLAN

You can configure the maximum age time per VLAN when using Rapid PVST+.

SUMMARY STEPS

1. switch# **configure terminal**
2. switch(config)# **spanning-tree vlan *vlan-range* max-age *max-age***

DETAILED STEPS

	Command or Action	Purpose
Step 1	switch# configure terminal	Enters configuration mode.
Step 2	switch(config)# spanning-tree vlan <i>vlan-range</i> max-age <i>max-age</i>	Configures the maximum aging time of a VLAN. The maximum aging time value can be from 6 to 40 seconds, and the default is 20 seconds.

This example shows how to configure the maximum aging time for a VLAN:

```
switch# configure terminal
switch(config)# spanning-tree vlan 5 max-age 36
```

Specifying the Link Type

Rapid connectivity (802.1w standard) is established only on point-to-point links. By default, the link type is controlled from the duplex mode of the interface. A full-duplex port is considered to have a point-to-point connection; a half-duplex port is considered to have a shared connection.

If you have a half-duplex link physically connected point-to-point to a single port on a remote switch, you can override the default setting on the link type and enable rapid transitions.

If you set the link to shared, STP moves back to 802.1D.

SUMMARY STEPS

1. switch# **configure terminal**
2. switch(config)# **interface** *type slot/port*
3. switch(config-if)# **spanning-tree link-type** {**auto** | **point-to-point** | **shared**}

DETAILED STEPS

	Command or Action	Purpose
Step 1	switch# configure terminal	Enters configuration mode.
Step 2	switch(config)# interface <i>type slot/port</i>	Specifies the interface to configure, and enters the interface configuration mode.
Step 3	switch(config-if)# spanning-tree link-type { auto point-to-point shared }	Configures the link type to be either a point-to-point link or shared link. The system reads the default value from the switch connection, as follows: half duplex links are shared and full-duplex links are point-to-point. If the link type is shared, the STP reverts to 802.1D. The default is auto, which sets the link type based on the duplex setting of the interface.

This example shows how to configure the link type as a point-to-point link:

```
switch# configure terminal
switch (config)# interface ethernet 1/4
switch(config-if)# spanning-tree link-type point-to-point
```

You can only apply this command to a physical Ethernet interface.

Restarting the Protocol

A bridge running Rapid PVST+ can send 802.1D BPDUs on one of its ports when it is connected to a legacy bridge. However, the STP protocol migration cannot determine whether the legacy switch has been removed from the link unless the legacy switch is the designated switch. You can restart the protocol negotiation (force the renegotiation with neighboring switches) on the entire switch or on specified interfaces.

Command	Purpose
switch# clear spanning-tree detected-protocol [interface <i>interface</i> [<i>interface-num</i> <i>port-channel</i>]]	Restarts Rapid PVST+ on all interfaces on the switch or specified interfaces.

The following example shows how to restart Rapid PVST+ on an Ethernet interface:

```
switch# clear spanning-tree detected-protocol interface ethernet 1/8
```

Verifying Rapid PVST+ Configurations

To display Rapid PVST+ configuration information, perform one of these tasks:

Command	Purpose
switch# show running-config spanning-tree [all]	Displays the current spanning tree configuration.
switch# show spanning-tree [options]	Displays selected detailed information for the current spanning tree configuration.

This example shows how to display spanning tree status:

```
switch# show spanning-tree brief

VLAN0001
  Spanning tree enabled protocol rstp
  Root ID    Priority    32768
             Address     001c.b05a.5447
             Cost        2
             Port        131 (Ethernet1/3)
             Hello Time  2 sec  Max Age 20 sec  Forward Delay 15 sec
  Bridge ID  Priority    32769 (priority 32768 sys-id-ext 1)
             Address     000d.ec6d.7841
             Hello Time  2 sec  Max Age 20 sec  Forward Delay 15 sec

Interface    Role Sts Cost      Prio.Nbr Type
-----
Eth1/3       Root FWD 2         128.131 P2p Peer (STP)
veth1/1     Desg FWD 2         128.129 Edge P2p
```




Configuring Multiple Spanning Tree

Multiple Spanning Tree (MST), which is the IEEE 802.1s standard, allows you to assign two or more VLANs to a spanning tree instance. MST is not the default spanning tree mode; Rapid per VLAN Spanning Tree (Rapid PVST+) is the default mode. MST instances with the same name, revision number, and VLAN-to-instance mapping combine to form an MST region. The MST region appears as a single bridge to spanning tree configurations outside the region. MST fails over to IEEE 802.1D Spanning Tree Protocol (STP) when it receives an 802.1D message from a neighboring switch.

**Note**

Spanning tree is used to refer to IEEE 802.1w and IEEE 802.1s. If the text is discussing the IEEE 802.1D Spanning Tree Protocol, 802.1D is stated specifically.

This chapter describes how to configure MST on Cisco Nexus 5000 Series switches. It contains the following sections:

- [Information About MST, page 183](#)
- [Configuring MST, page 191](#)
- [Verifying MST Configurations, page 208](#)

Information About MST

MST Overview

**Note**

You must enable MST; Rapid PVST+ is the default spanning tree mode.

MST maps multiple VLANs into a spanning tree instance, with each instance having a spanning tree topology independent of other spanning tree instances. This architecture provides multiple forwarding paths for data traffic, enables load balancing, and reduces the number of STP instances required to support a large number of VLANs. MST improves the fault tolerance of the network because a failure in one instance (forwarding path) does not affect other instances (forwarding paths).

MST provides rapid convergence through explicit handshaking as each MST instance uses the IEEE 802.1w standard, which eliminates the 802.1D forwarding delay and quickly transitions root bridge ports and designated ports to the forwarding state.

MAC address reduction is always enabled while you are using MST. You cannot disable this feature.

MST improves spanning tree operation and maintains backward compatibility with these STP versions:

- Original 802.1D spanning tree
- Rapid per-VLAN spanning tree (Rapid PVST+)
 - IEEE 802.1w defined the Rapid Spanning Tree Protocol (RSTP) and was incorporated into IEEE 802.1D.
- IEEE 802.1s defined MST and was incorporated into IEEE 802.1Q.

MST Regions

To allow switches to participate in MST instances, you must consistently configure the switches with the same MST configuration information.

A collection of interconnected switches that have the same MST configuration is an MST region. An MST region is a linked group of MST bridges with the same MST configuration.

The MST configuration controls the MST region to which each switch belongs. The configuration includes the name of the region, the revision number, and the MST VLAN-to-instance assignment map.

A region can have one or multiple members with the same MST configuration. Each member must be capable of processing 802.1w bridge protocol data units (BPDUs). There is no limit to the number of MST regions in a network.

Each region can support up to 65 MST instances (MSTIs). Instances are identified by any number in the range from 1 to 4094. The system reserves Instance 0 for a special instance, which is the IST. You can assign a VLAN to only one MST instance at a time.

The MST region appears as a single bridge to adjacent MST regions and to other Rapid PVST+ regions and 802.1D spanning tree protocols.



Note

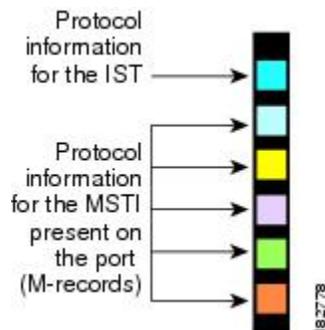
We recommend that you do not partition the network into a large number of regions.

MST BPDUs

Each region has only one MST BPDU, and that BPDU carries an M-record for each MSTI within the region (see the following figure). Only the IST sends BPDUs for the MST region; all M-records are encapsulated in

that one BPDU that the IST sends. Because the MST BPDU carries information for all instances, the number of BPDUs that need to be processed to support MSTIs is significantly reduced.

Figure 26: MST BPDU with M-Records for MSTIs



MST Configuration Information

The MST configuration that must be identical on all switches within a single MST region is configured by the user.

You can configure the following three parameters of the MST configuration:

- Name—32-character string, null padded and null terminated, identifying the MST region
- Revision number—Unsigned 16-bit number that identifies the revision of the current MST configuration



Note

You must set the revision number when required as part of the MST configuration. The revision number is *not* incremented automatically each time that the MST configuration is committed.

- MST configuration table—4096-element table that associates each of the potential 4094 VLANs supported to a given instance with the first (0) and last element (4095) set to 0. The value of element number X represents the instance to which VLAN X is mapped.



Caution

When you change the VLAN-to-MSTI mapping, the system restarts MST.

MST BPDUs contain these three configuration parameters. An MST bridge accepts an MST BPDU into its own region only if these three configuration parameters match exactly. If one configuration attribute differs, the MST bridge considers the BPDU to be from another MST region.

IST, CIST, and CST

IST, CIST, and CST Overview

Unlike Rapid PVST+, in which all the STP instances are independent, MST establishes and maintains IST, CIST, and CST spanning trees, as follows:

- An IST is the spanning tree that runs in an MST region.

MST establishes and maintains additional spanning trees within each MST region; these spanning trees are called, multiple spanning tree instances (MSTIs).

Instance 0 is a special instance for a region, known as the IST. The IST always exists on all ports; you cannot delete the IST, or Instance 0. By default, all VLANs are assigned to the IST. All other MST instances are numbered from 1 to 4094.

The IST is the only STP instance that sends and receives BPDUs. All of the other MSTI information is contained in MST records (M-records), which are encapsulated within MST BPDUs.

All MSTIs within the same region share the same protocol timers, but each MSTI has its own topology parameters, such as the root bridge ID, the root path cost, and so forth.

An MSTI is local to the region; for example, MSTI 9 in region A is independent of MSTI 9 in region B, even if regions A and B are interconnected.

- The CST interconnects the MST regions and any instance of 802.1D and 802.1w STP that may be running on the network. The CST is the one STP instance for the entire bridged network and encompasses all MST regions and 802.1w and 802.1D instances.
- A CIST is a collection of the ISTs in each MST region. The CIST is the same as an IST inside an MST region, and the same as a CST outside an MST region.

The spanning tree computed in an MST region appears as a subtree in the CST that encompasses the entire switched domain. The CIST is formed by the spanning tree algorithm running among switches that support the 802.1w, 802.1s, and 802.1D standards. The CIST inside an MST region is the same as the CST outside a region.

Spanning Tree Operation Within an MST Region

The IST connects all the MST switches in a region. When the IST converges, the root of the IST becomes the CIST regional root. The CIST regional root is also the CIST root if there is only one region in the network. If the CIST root is outside the region, the protocol selects one of the MST switches at the boundary of the region as the CIST regional root.

When an MST switch initializes, it sends BPDUs that identify itself as the root of the CIST and the CIST regional root, with both the path costs to the CIST root and to the CIST regional root set to zero. The switch also initializes all of its MSTIs and claims to be the root for all of them. If the switch receives superior MST root information (lower switch ID, lower path cost, and so forth) than the information that is currently stored for the port, it relinquishes its claim as the CIST regional root.

During initialization, an MST region might have many subregions, each with its own CIST regional root. As switches receive superior IST information from a neighbor in the same region, they leave their old subregions and join the new subregion that contains the true CIST regional root. This action causes all subregions to shrink except for the subregion that contains the true CIST regional root.

All switches in the MST region must agree on the same CIST regional root. Any two switches in the region will only synchronize their port roles for an MSTI if they converge to a common CIST regional root.

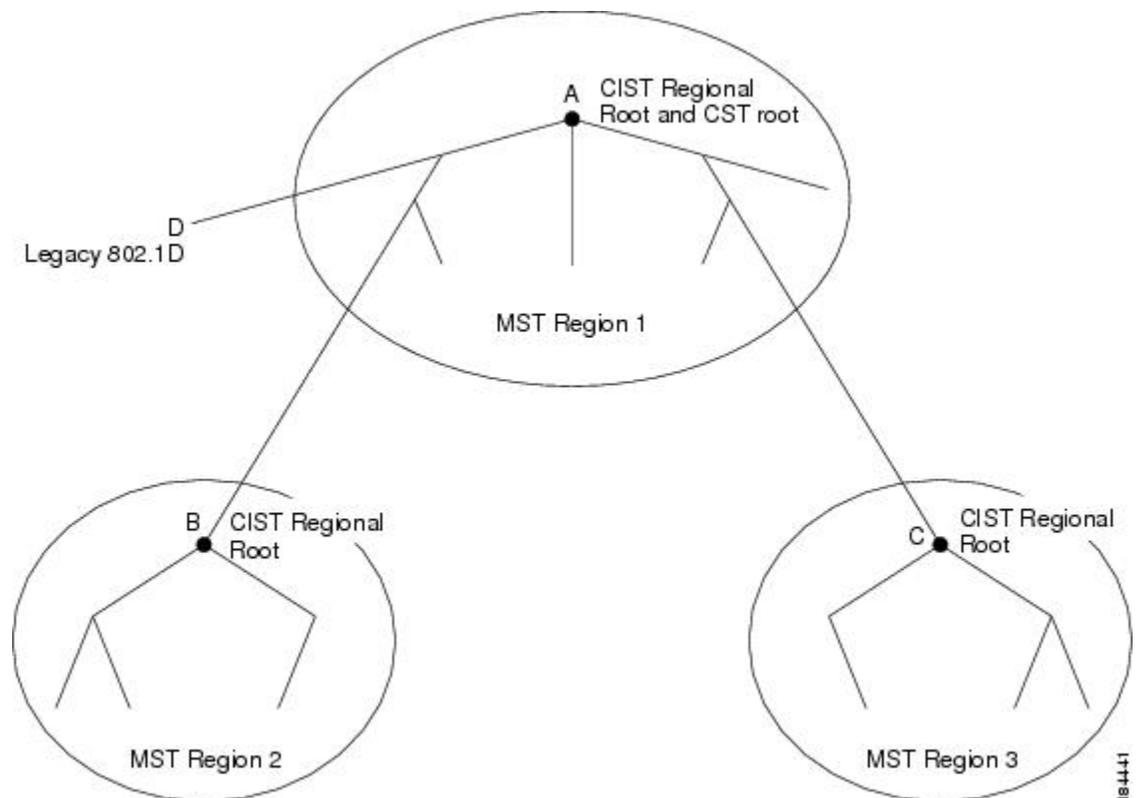
Spanning Tree Operations Between MST Regions

If you have multiple regions or 802.1w or 802.1D STP instances within a network, MST establishes and maintains the CST, which includes all MST regions and all 802.1w and 802.1D STP switches in the network. The MSTIs combine with the IST at the boundary of the region to become the CST.

The IST connects all the MST switches in the region and appears as a subtree in the CIST that encompasses the entire switched domain. The root of the subtree is the CIST regional root. The MST region appears as a virtual switch to adjacent STP switches and MST regions.

The following figure shows a network with three MST regions and an 802.1D switch (D). The CIST regional root for region 1 (A) is also the CIST root. The CIST regional root for region 2 (B) and the CIST regional root for region 3 (C) are the roots for their respective subtrees within the CIST.

Figure 27: MST Regions, CIST Regional Roots, and CST Root



Only the CST instance sends and receives BPDUs. MSTIs add their spanning tree information into the BPDUs (as M-records) to interact with neighboring switches and compute the final spanning tree topology. Because of this, the spanning tree parameters related to the BPDU transmission (for example, hello time, forward time, max-age, and max-hops) are configured only on the CST instance but affect all MSTIs. You can configure the parameters related to the spanning tree topology (for example, the switch priority, the port VLAN cost, and the port VLAN priority) on both the CST instance and the MSTI.

MST switches use Version 3 BPDUs or 802.1D STP BPDUs to communicate with 802.1D-only switches. MST switches use MST BPDUs to communicate with MST switches.

MST Terminology

MST naming conventions include identification of some internal or regional parameters. These parameters are used only within an MST region, compared to external parameters that are used throughout the whole network. Because the CIST is the only spanning tree instance that spans the whole network, only the CIST

parameters require the external qualifiers and not the internal or regional qualifiers. The MST terminology is as follows:

- The CIST root is the root bridge for the CIST, which is the unique instance that spans the whole network.
- The CIST external root path cost is the cost to the CIST root. This cost is left unchanged within an MST region. An MST region looks like a single switch to the CIST. The CIST external root path cost is the root path cost calculated between these virtual switches and switches that do not belong to any region.
- If the CIST root is in the region, the CIST regional root is the CIST root. Otherwise, the CIST regional root is the closest switch to the CIST root in the region. The CIST regional root acts as a root bridge for the IST.
- The CIST internal root path cost is the cost to the CIST regional root in a region. This cost is only relevant to the IST, instance 0.

Hop Count

MST does not use the message-age and maximum-age information in the configuration BPDU to compute the STP topology inside the MST region. Instead, the protocol uses the path cost to the root and a hop-count mechanism similar to the IP time-to-live (TTL) mechanism.

By using the **spanning-tree mst max-hops** global configuration command, you can configure the maximum hops inside the region and apply it to the IST and all MST instances in that region.

The hop count achieves the same result as the message-age information (triggers a reconfiguration). The root bridge of the instance always sends a BPDU (or M-record) with a cost of 0 and the hop count set to the maximum value. When a switch receives this BPDU, it decrements the received remaining hop count by one and propagates this value as the remaining hop count in the BPDUs that it generates. When the count reaches zero, the switch discards the BPDU and ages the information held for the port.

The message-age and maximum-age information in the 802.1w portion of the BPDU remain the same throughout the region (only on the IST), and the same values are propagated by the region-designated ports at the boundary.

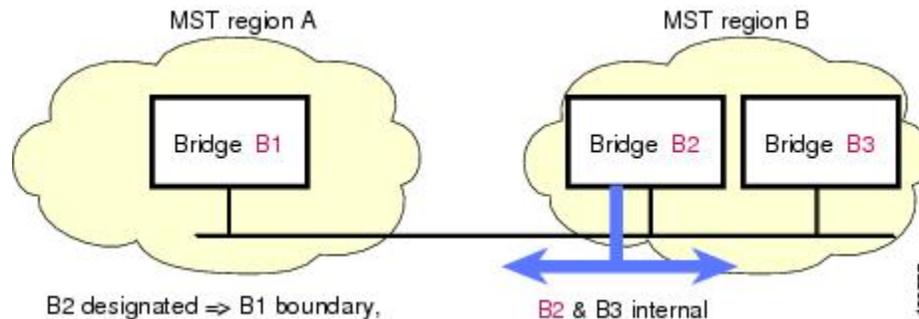
You configure a maximum aging time as the number of seconds that a switch waits without receiving spanning tree configuration messages before attempting a reconfiguration.

Boundary Ports

A boundary port is a port that connects to a LAN, the designated bridge of which is either a bridge with a different MST configuration (and so, a separate MST region) or a Rapid PVST+ or 802.1D STP bridge. A designated port knows that it is on the boundary if it detects an STP bridge or receives an agreement proposal from an MST bridge with a different configuration or a Rapid PVST+ bridge. This definition allows two ports

that are internal to a region to share a segment with a port that belongs to a different region, creating the possibility of receiving both internal and external messages on a port (see the following figure).

Figure 28: MST Boundary Ports



At the boundary, the roles of MST ports do not matter; the system forces their state to be the same as the IST port state. If the boundary flag is set for the port, the MST port-role selection process assigns a port role to the boundary and assigns the same state as the state of the IST port. The IST port at the boundary can take up any port role except a backup port role.

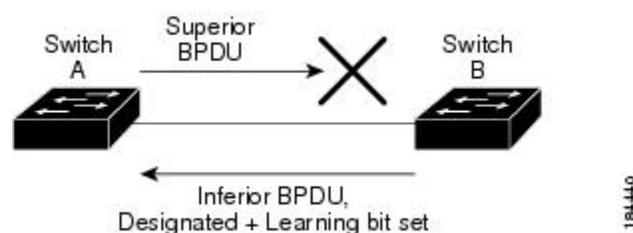
Detecting Unidirectional Link Failure

Currently, this feature is not present in the IEEE MST standard, but it is included in the standard-compliant implementation. The software checks the consistency of the port role and state in the received BPDUs to detect unidirectional link failures that could cause bridging loops.

When a designated port detects a conflict, it keeps its role, but reverts to a discarding state because disrupting connectivity in case of inconsistency is preferable to opening a bridging loop.

The following figure shows a unidirectional link failure that typically creates a bridging loop. Switch A is the root bridge, and its BPDUs are lost on the link leading to switch B. Rapid PVST+ (802.1w) and MST BPDUs include the role and state of the sending port. With this information, switch A can detect that switch B does not react to the superior BPDUs that it sends and that switch B is the designated, not root port. As a result, switch A blocks (or keeps blocking) its port, which prevents the bridging loop. The block is shown as an STP dispute.

Figure 29: Detecting a Unidirectional Link Failure



Port Cost and Port Priority

Spanning tree uses port costs to break a tie for the designated port. Lower values indicate lower port costs, and spanning tree chooses the least costly path. Default port costs are taken from the bandwidth of the interface, as follows:

- 10 Mbps—2,000,000
- 100 Mbps—200,000
- 1 Gigabit Ethernet—20,000
- 10 Gigabit Ethernet—2,000

You can configure the port costs in order to influence which port is chosen.



Note

MST always uses the long path cost calculation method, so the range of valid values is between 1 and 200,000,000.

The system uses port priorities to break ties among ports with the same cost. A lower number indicates a higher priority. The default port priority is 128. You can configure the priority to values between 0 and 224, in increments of 32.

Interoperability with IEEE 802.1D

A switch that runs MST supports a built-in protocol migration feature that enables it to interoperate with 802.1D STP switches. If this switch receives an 802.1D configuration BPDU (a BPDU with the protocol version set to 0), it sends only 802.1D BPDUs on that port. In addition, an MST switch can detect that a port is at the boundary of a region when it receives an 802.1D BPDU, an MST BPDU (Version 3) associated with a different region, or an 802.1w BPDU (Version 2).

However, the switch does not automatically revert to the MST mode if it no longer receives 802.1D BPDUs because it cannot detect whether the 802.1D switch has been removed from the link unless the 802.1D switch is the designated switch. A switch might also continue to assign a boundary role to a port when the switch to which this switch is connected has joined the region.

To restart the protocol migration process (force the renegotiation with neighboring switches), enter the **clear spanning-tree detected-protocols** command.

All Rapid PVST+ switches (and all 802.1D STP switches) on the link can process MST BPDUs as if they are 802.1w BPDUs. MST switches can send either Version 0 configuration and topology change notification (TCN) BPDUs or Version 3 MST BPDUs on a boundary port. A boundary port connects to a LAN, the designated switch of which is either a single spanning tree switch or a switch with a different MST configuration.



Note

MST interoperates with the Cisco prestandard MSTP whenever it receives prestandard MSTP on an MST port; no explicit configuration is necessary.

Interoperability with Rapid PVST+: Understanding PVST Simulation

MST interoperates with Rapid PVST+ with no need for user configuration. The PVST simulation feature enables this seamless interoperability.



Note

PVST simulation is enabled by default. That is, by default, all interfaces on the switch interoperate between MST and Rapid PVST+.

However, you may want to control the connection between MST and Rapid PVST+ to protect against accidentally connecting an MST-enabled port to a Rapid PVST+-enabled port. Because Rapid PVST+ is the default STP mode, you may encounter many Rapid PVST+-enabled connections.

Disabling Rapid PVST+ simulation, which can be done per port or globally for the entire switch, moves the MST-enabled port to the blocking state once it detects it is connected to a Rapid PVST+-enabled port. This port remains in the inconsistent state until the port stops receiving Rapid PVST+/SSTP BPDUs, and then the port resumes the normal STP transition process.

Configuring MST

MST Configuration Guidelines

When configuring MST, follow these guidelines:

- When you work with private VLANs, enter the **private-vlan synchronize** command to map the secondary VLANs to the same MST instance as the primary VLAN.
- When you are in the MST configuration mode, the following guidelines apply:
 - Each command reference line creates its pending regional configuration.
 - The pending region configuration starts with the current region configuration.
 - To leave the MST configuration mode without committing any changes, enter the **abort** command.
 - To leave the MST configuration mode and commit all the changes that you made before you left the mode, enter the **exit** command.

Enabling MST

You must enable MST; Rapid PVST+ is the default.



Caution

Changing the spanning tree mode disrupts traffic because all spanning tree instances are stopped for the previous mode and started for the new mode.

SUMMARY STEPS

1. switch# **configure terminal**
2. switch(config)# **spanning-tree mode mst**
3. (Optional) switch(config)# **no spanning-tree mode mst**

DETAILED STEPS

	Command or Action	Purpose
Step 1	switch# configure terminal	Enters configuration mode.
Step 2	switch(config)# spanning-tree mode mst	Enables MST on the switch.
Step 3	switch(config)# no spanning-tree mode mst	(Optional) Disables MST on the switch and returns you to Rapid PVST+.

This example shows how to enable MST on the switch:

```
switch# configure terminal
switch(config)# spanning-tree mode mst
```



Note

Because STP is enabled by default, entering a **show running-config** command to view the resulting configuration does not display the command that you entered to enable STP.

Entering MST Configuration Mode

You enter MST configuration mode to configure the MST name, VLAN-to-instance mapping, and MST revision number on the switch.

For two or more switches to be in the same MST region, they must have the identical MST name, VLAN-to-instance mapping, and MST revision number.



Note

Each command reference line creates its pending regional configuration in MST configuration mode. In addition, the pending region configuration starts with the current region configuration.

When you are working in MST configuration mode, note the difference between the **exit** and **abort** commands.

SUMMARY STEPS

1. switch# **configure terminal**
2. switch(config)# **spanning-tree mst configuration**
3. switch(config-mst)# **exit** or switch(config-mst)# **abort**
4. (Optional) switch(config)# **no spanning-tree mst configuration**

DETAILED STEPS

	Command or Action	Purpose
Step 1	switch# configure terminal	Enters configuration mode.
Step 2	switch(config)# spanning-tree mst configuration	Enters MST configuration mode on the system. You must be in the MST configuration mode to assign the MST configuration parameters, as follows: <ul style="list-style-type: none"> • MST name • Instance-to-VLAN mapping • MST revision number • Synchronize primary and secondary VLANs in private VLANs
Step 3	switch(config-mst)# exit or switch(config-mst)# abort	<ul style="list-style-type: none"> • The first form commits all the changes and exits MST configuration mode. • The second form exits the MST configuration mode without committing any of the changes.
Step 4	switch(config)# no spanning-tree mst configuration	(Optional) Returns the MST region configuration to the following default values: <ul style="list-style-type: none"> • The region name is an empty string. • No VLANs are mapped to any MST instance (all VLANs are mapped to the CIST instance). • The revision number is 0.

Specifying the MST Name

You configure a region name on the bridge. For two or more bridges to be in the same MST region, they must have the identical MST name, VLAN-to-instance mapping, and MST revision number.

SUMMARY STEPS

1. switch# **configure terminal**
2. switch(config)# **spanning-tree mst configuration**
3. switch(config-mst)# **name name**

DETAILED STEPS

	Command or Action	Purpose
Step 1	switch# configure terminal	Enters configuration mode.
Step 2	switch(config)# spanning-tree mst configuration	Enters MST configuration submode.
Step 3	switch(config-mst)# name name	Specifies the name for MST region. The <i>name</i> string has a maximum length of 32 characters and is case-sensitive. The default is an empty string.

This example shows how to set the name of the MST region:

```
switch# configure terminal
switch(config)# spanning-tree mst configuration
switch(config-mst)# name accounting
```

Specifying the MST Configuration Revision Number

You configure the revision number on the bridge. For two or more bridges to be in the same MST region, they must have the identical MST name, VLAN-to-instance mapping, and MST revision number.

SUMMARY STEPS

1. switch# **configure terminal**
2. switch(config)# **spanning-tree mst configuration**
3. switch(config-mst)# **revision version**

DETAILED STEPS

	Command or Action	Purpose
Step 1	switch# configure terminal	Enters configuration mode.
Step 2	switch(config)# spanning-tree mst configuration	Enters MST configuration submode.
Step 3	switch(config-mst)# revision version	Specifies the revision number for the MST region. The range is from 0 to 65535, and the default value is 0.

This example shows how to configure the revision number of the MSTI region for 5:

```
switch# configure terminal
switch(config)# spanning-tree mst configuration
switch(config-mst)# revision 5
```

Specifying the Configuration on an MST Region

For two or more switches to be in the same MST region, they must have the same VLAN-to-instance mapping, the same configuration revision number, and the same MST name.

A region can have one member or multiple members with the same MST configuration; each member must be capable of processing IEEE 802.1w RSTP BPDUs. There is no limit to the number of MST regions in a network, but each region can support only up to 65 MST instances. You can assign a VLAN to only one MST instance at a time.

SUMMARY STEPS

1. switch# **configure terminal**
2. switch(config)# **spanning-tree mst configuration**
3. switch(config-mst)# **instance** *instance-id* **vlan** *vlan-range*
4. switch(config-mst)# **name** *name*
5. switch(config-mst)# **revision** *version*

DETAILED STEPS

	Command or Action	Purpose
Step 1	switch# configure terminal	Enters configuration mode.
Step 2	switch(config)# spanning-tree mst configuration	Enters MST configuration submode.
Step 3	switch(config-mst)# instance <i>instance-id</i> vlan <i>vlan-range</i>	<p>Maps VLANs to an MST instance as follows:</p> <ul style="list-style-type: none"> • For <i>instance-id</i> , the range is from 1 to 4094. • For vlan <i>vlan-range</i> , the range is from 1 to 4094. <p>When you map VLANs to an MST instance, the mapping is incremental, and the VLANs specified in the command are added to or removed from the VLANs that were previously mapped.</p> <p>To specify a VLAN range, enter a hyphen; for example, enter the instance 1 vlan 1-63 command to map VLANs 1 through 63 to MST instance 1.</p> <p>To specify a VLAN series, enter a comma; for example, enter the instance 1 vlan 10, 20, 30 command to map VLANs 10, 20, and 30 to MST instance 1.</p>

	Command or Action	Purpose
Step 4	switch(config-mst)# name <i>name</i>	Specifies the instance name. The <i>name</i> string has a maximum length of 32 characters and is case sensitive.
Step 5	switch(config-mst)# revision <i>version</i>	Specifies the configuration revision number. The range is from 0 to 65535.

To return to defaults, do the following:

- To return to the default MST region configuration settings, enter the **no spanning-tree mst configuration** configuration command.
- To return to the default VLAN-to-instance map, enter the **no instance** *instance-id* **vlan** *vlan-range* MST configuration command.
- To return to the default name, enter the **no name** MST configuration command.
- To return to the default revision number, enter the **no revision** MST configuration command.
- To reenables Rapid PVST+, enter the **no spanning-tree mode** or the **spanning-tree mode rapid-pvst** global configuration command.

This example shows how to enter MST configuration mode, map VLANs 10 to 20 to MST instance 1, name the region *region1*, set the configuration revision to 1, display the pending configuration, apply the changes, and return to global configuration mode:

```
switch(config)# spanning-tree mst configuration
switch(config-mst)# instance 1 vlan 10-20
switch(config-mst)# name region1
switch(config-mst)# revision 1
switch(config-mst)# show pending
Pending MST configuration
Name      [region1]
Revision  1
Instances configured 2
Instance  Vlans Mapped
-----
0          1-9,21-4094
1          10-20
-----
```

Mapping and Unmapping VLANs to MST Instances



Caution

When you change the VLAN-to-MSTI mapping, the system restarts MST.



Note

You cannot disable an MSTI.

For two or more bridges to be in the same MST region, they must have the identical MST name, VLAN-to-instance mapping, and MST revision number.

SUMMARY STEPS

1. switch# **configure terminal**
2. switch(config)# **spanning-tree mst configuration**
3. switch(config-mst)# **instance** *instance-id* **vlan** *vlan-range*
4. switch(config-mst)# **no instance** *instance-id* **vlan** *vlan-range*

DETAILED STEPS

	Command or Action	Purpose
Step 1	switch# configure terminal	Enters configuration mode.
Step 2	switch(config)# spanning-tree mst configuration	Enters MST configuration submenu.
Step 3	switch(config-mst)# instance <i>instance-id</i> vlan <i>vlan-range</i>	Maps VLANs to an MST instance, as follows: <ul style="list-style-type: none"> • For <i>instance-id</i> the range is from 1 to 4094. Instance 0 is reserved for the IST for each MST region. • For <i>vlan-range</i> the range is from 1 to 4094. When you map VLANs to an MSTI, the mapping is incremental, and the VLANs specified in the command are added to or removed from the VLANs that were previously mapped.
Step 4	switch(config-mst)# no instance <i>instance-id</i> vlan <i>vlan-range</i>	Deletes the specified instance and returns the VLANs to the default MSTI, which is the CIST.

This example shows how to map VLAN 200 to MSTI 3:

```
switch# configure terminal
switch(config)# spanning-tree mst configuration
switch(config-mst)# instance 3 vlan 200
```

Mapping Secondary VLANs to Same MSTI as Primary VLANs for Private VLANs

When you are working with private VLANs on the system, all secondary VLANs must be in the same MSTI and their associated primary VLAN.

SUMMARY STEPS

1. switch# **configure terminal**
2. switch(config)# **spanning-tree mst configuration**
3. switch(config-mst)# **private-vlan synchronize**

DETAILED STEPS

	Command or Action	Purpose
Step 1	switch# configure terminal	Enters configuration mode.
Step 2	switch(config)# spanning-tree mst configuration	Enters MST configuration submenu.
Step 3	switch(config-mst)# private-vlan synchronize	Automatically maps all secondary VLANs to the same MSTI and their associated primary VLAN for all private VLANs.

This example shows how to automatically map all the secondary VLANs to the same MSTI as their associated primary VLANs in all private VLANs:

```
switch# configure terminal
switch(config)# spanning-tree mst configuration
switch(config-mst)# private-vlan synchronize
```

Configuring the Root Bridge

You can configure the switch to become the root bridge.



Note

The root bridge for each MSTI should be a backbone or distribution switch. Do not configure an access switch as the spanning tree primary root bridge.

Enter the **diameter** keyword, which is available only for MSTI 0 (or the IST), to specify the network diameter (that is, the maximum number of hops between any two end stations in the network). When you specify the network diameter, the switch automatically sets an optimal hello time, forward-delay time, and maximum-age time for a network of that diameter, which can significantly reduce the convergence time. You can enter the **hello** keyword to override the automatically calculated hello time.



Note

With the switch configured as the root bridge, do not manually configure the hello time, forward-delay time, and maximum-age time using the **spanning-tree mst hello-time**, **spanning-tree mst forward-time**, and **spanning-tree mst max-age** global configuration commands.

SUMMARY STEPS

1. switch# **configure terminal**
2. switch(config)# **spanning-tree mst instance-id root {primary | secondary} [diameter dia [hello-time hello-time]]**
3. (Optional) switch(config)# **no spanning-tree mst instance-id root**

DETAILED STEPS

	Command or Action	Purpose
Step 1	switch# configure terminal	Enters configuration mode.
Step 2	switch(config)# spanning-tree mst <i>instance-id</i> root { primary secondary } [diameter <i>dia</i> [hello-time <i>hello-time</i>]]	Configures a switch as the root bridge as follows: <ul style="list-style-type: none"> • For <i>instance-id</i>, you can specify a single instance, a range of instances separated by a hyphen, or a series of instances separated by a comma. The range is from 1 to 4094. • For diameter <i>net-diameter</i>, specify the maximum number of hops between any two end stations. The default is 7. This keyword is available only for MST instance 0. • For hello-time <i>seconds</i>, specify the interval in seconds between the generation of configuration messages by the root bridge. The range is from 1 to 10 seconds; the default is 2 seconds.
Step 3	switch(config)# no spanning-tree mst <i>instance-id</i> root	(Optional) Returns the switch priority, diameter, and hello time to default values.

This example shows how to configure the switch as the root switch for MSTI 5:

```
switch# configure terminal
switch(config)# spanning-tree mst 5 root primary
```

Configuring a Secondary Root Bridge

You can execute this command on more than one switch to configure multiple backup root bridges. Enter the same network diameter and hello-time values that you used when you configured the primary root bridge with the **spanning-tree mst root primary** configuration command.

SUMMARY STEPS

1. switch# **configure terminal**
2. switch(config)# **spanning-tree mst** *instance-id* **root** {**primary** | **secondary**} [**diameter** *dia* [**hello-time** *hello-time*]]
3. (Optional) switch(config)# **no spanning-tree mst** *instance-id* **root**

DETAILED STEPS

	Command or Action	Purpose
Step 1	switch# configure terminal	Enters configuration mode.

	Command or Action	Purpose
Step 2	switch(config)# spanning-tree mst <i>instance-id</i> root { primary secondary } [diameter <i>dia</i> [hello-time <i>hello-time</i>]]	Configures a switch as the secondary root bridge as follows: <ul style="list-style-type: none"> • For <i>instance-id</i>, you can specify a single instance, a range of instances separated by a hyphen, or a series of instances separated by a comma. The range is from 1 to 4094. • For diameter <i>net-diameter</i>, specify the maximum number of hops between any two end stations. The default is 7. This keyword is available only for MST instance 0. • For hello-time <i>seconds</i>, specify the interval in seconds between the generation of configuration messages by the root bridge. The range is from 1 to 10 seconds; the default is 2 seconds.
Step 3	switch(config)# no spanning-tree mst <i>instance-id</i> root	(Optional) Returns the switch priority, diameter, and hello-time to default values.

This example shows how to configure the switch as the secondary root switch for MSTI 5:

```
switch# configure terminal
switch(config)# spanning-tree mst 5 root secondary
```

Configuring the Port Priority

If a loop occurs, MST uses the port priority when selecting an interface to put into the forwarding state. You can assign lower priority values to interfaces that you want selected first and higher priority values to the interface that you want selected last. If all interfaces have the same priority value, MST puts the interface with the lowest interface number in the forwarding state and blocks the other interfaces.

SUMMARY STEPS

1. switch# **configure terminal**
2. switch(config)# **interface** {{*type slot/port*} | {**port-channel** *number*}}
3. switch(config-if)# **spanning-tree mst** *instance-id* **port-priority** *priority*

DETAILED STEPS

	Command or Action	Purpose
Step 1	switch# configure terminal	Enters configuration mode.
Step 2	switch(config)# interface {{ <i>type slot/port</i> } { port-channel <i>number</i> }}	Specifies an interface to configure, and enters interface configuration mode.
Step 3	switch(config-if)# spanning-tree mst <i>instance-id</i> port-priority <i>priority</i>	Configures the port priority as follows:

	Command or Action	Purpose
		<ul style="list-style-type: none"> • For <i>instance-id</i>, you can specify a single MSTI, a range of MSTIs separated by a hyphen, or a series of MSTIs separated by a comma. The range is from 1 to 4094. • For <i>priority</i>, the range is 0 to 224 in increments of 32. The default is 128. A lower number indicates a higher priority. <p>The priority values are 0, 32, 64, 96, 128, 160, 192, and 224. The system rejects all other values.</p>

This example shows how to set the MST interface port priority for MSTI 3 on Ethernet port 3/1 to 64:

```
switch# configure terminal
switch(config)# interface ethernet 3/1
switch(config-if)# spanning-tree mst 3 port-priority 64
```

You can only apply this command to a physical Ethernet interface.

Configuring the Port Cost

The MST path cost default value is derived from the media speed of an interface. If a loop occurs, MST uses the cost when selecting an interface to put in the forwarding state. You can assign lower cost values to interfaces that you want selected first and higher cost to interfaces values that you want selected last. If all interfaces have the same cost value, MST puts the interface with the lowest interface number in the forwarding state and blocks the other interfaces.



Note MST uses the long pathcost calculation method.

SUMMARY STEPS

1. switch# **configure terminal**
2. switch(config)# **interface** *{{type slot/port} | {port-channel number}}*
3. switch(config-if)# **spanning-tree mst** *instance-id* **cost** [*cost* | **auto**]

DETAILED STEPS

	Command or Action	Purpose
Step 1	switch# configure terminal	Enters configuration mode.
Step 2	switch(config)# interface <i>{{type slot/port} {port-channel number}}</i>	Specifies an interface to configure, and enters interface configuration mode.
Step 3	switch(config-if)# spanning-tree mst <i>instance-id</i> cost [<i>cost</i> auto]	Configures the cost.

	Command or Action	Purpose
		<p>If a loop occurs, MST uses the path cost when selecting an interface to place into the forwarding state. A lower path cost represents higher-speed transmission as follows:</p> <ul style="list-style-type: none"> • For <i>instance-id</i>, you can specify a single instance, a range of instances separated by a hyphen, or a series of instances separated by a comma. The range is from 1 to 4094. • For <i>cost</i>, the range is from 1 to 200000000. The default value is auto, which is derived from the media speed of the interface.

This example shows how to set the MST interface port cost on Ethernet 3/1 for MSTI 4:

```
switch# configure terminal
switch(config)# interface ethernet 3/1
switch(config-if)# spanning-tree mst 4 cost 17031970
```

Configuring the Switch Priority

You can configure the switch priority for an MST instance so that it is more likely that the specified switch is chosen as the root bridge.



Note

Exercise care when using this command. For most situations, we recommend that you enter the **spanning-tree mst root primary** and the **spanning-tree mst root secondary** global configuration commands to modify the switch priority.

SUMMARY STEPS

1. switch# **configure terminal**
2. switch(config)# **spanning-tree mst instance-id priority priority-value**

DETAILED STEPS

	Command or Action	Purpose
Step 1	switch# configure terminal	Enters configuration mode.
Step 2	switch(config)# spanning-tree mst instance-id priority priority-value	<p>Configures a switch priority as follows:</p> <ul style="list-style-type: none"> • For <i>instance-id</i>, you can specify a single instance, a range of instances separated by a hyphen, or a series of instances separated by a comma. The range is from 1 to 4094.

	Command or Action	Purpose
		<ul style="list-style-type: none"> For priority, the range is from 0 to 61440 in increments of 4096; the default is 32768. A lower number indicates that the switch will most likely be chosen as the root bridge. <p>Priority values are 0, 4096, 8192, 12288, 16384, 20480, 24576, 28672, 32768, 36864, 40960, 45056, 49152, 53248, 57344, and 61440. The system rejects all other values.</p>

This example shows how to configure the priority of the bridge to 4096 for MSTI 5:

```
switch# configure terminal
switch(config)# spanning-tree mst 5 priority 4096
```

Configuring the Hello Time

You can configure the interval between the generation of configuration messages by the root bridge for all instances on the switch by changing the hello time.



Note

Exercise care when using this command. For most situations, we recommend that you enter the **spanning-tree mst *instance-id* root primary** and the **spanning-tree mst *instance-id* root secondary** configuration commands to modify the hello time.

SUMMARY STEPS

1. switch# **configure terminal**
2. switch(config)# **spanning-tree mst hello-time *seconds***

DETAILED STEPS

	Command or Action	Purpose
Step 1	switch# configure terminal	Enters configuration mode.
Step 2	switch(config)# spanning-tree mst hello-time <i>seconds</i>	Configures the hello time for all MST instances. The hello time is the interval between the generation of configuration messages by the root bridge. These messages mean that the switch is alive. For <i>seconds</i> , the range is from 1 to 10, and the default is 2 seconds.

This example shows how to configure the hello time of the switch to 1 second:

```
switch# configure terminal
switch(config)# spanning-tree mst hello-time 1
```

Configuring the Forwarding-Delay Time

You can set the forward delay timer for all MST instances on the switch with one command.

SUMMARY STEPS

1. switch# **configure terminal**
2. switch(config)# **spanning-tree mst forward-time** *seconds*

DETAILED STEPS

	Command or Action	Purpose
Step 1	switch# configure terminal	Enters configuration mode.
Step 2	switch(config)# spanning-tree mst forward-time <i>seconds</i>	Configures the forward time for all MST instances. The forward delay is the number of seconds that a port waits before changing from its spanning tree blocking and learning states to the forwarding state. For <i>seconds</i> , the range is from 4 to 30, and the default is 15 seconds.

This example shows how to configure the forward-delay time of the switch to 10 seconds:

```
switch# configure terminal
switch(config)# spanning-tree mst forward-time 10
```

Configuring the Maximum-Aging Time

The maximum-aging timer is the number of seconds that a switch waits without receiving spanning tree configuration messages before attempting a reconfiguration.

You set the maximum-aging timer for all MST instances on the switch with one command (the maximum age time only applies to the IST).

SUMMARY STEPS

1. switch# **configure terminal**
2. switch(config)# **spanning-tree mst max-age** *seconds*

DETAILED STEPS

	Command or Action	Purpose
Step 1	switch# configure terminal	Enters configuration mode.
Step 2	switch(config)# spanning-tree mst max-age <i>seconds</i>	Configures the maximum-aging time for all MST instances. The maximum-aging time is the number of seconds that a switch waits without receiving spanning tree configuration messages before attempting a

	Command or Action	Purpose
		reconfiguration. For <i>seconds</i> , the range is from 6 to 40, and the default is 20 seconds.

This example shows how to configure the maximum-aging timer of the switch to 40 seconds:

```
switch# configure terminal
switch(config)# spanning-tree mst max-age 40
```

Configuring the Maximum-Hop Count

MST uses the path cost to the IST regional root and a hop-count mechanism similar to the IP time-to-live (TTL) mechanism. You configure the maximum hops inside the region and apply it to the IST and all MST instances in that region. The hop count achieves the same result as the message-age information (triggers a reconfiguration).

SUMMARY STEPS

1. switch# **configure terminal**
2. switch(config)# **spanning-tree mst max-hops *hop-count***

DETAILED STEPS

	Command or Action	Purpose
Step 1	switch# configure terminal	Enters configuration mode.
Step 2	switch(config)# spanning-tree mst max-hops <i>hop-count</i>	Specifies the number of hops in a region before the BPDU is discarded, and the information held for a port is aged. For <i>hop-count</i> , the range is from 1 to 255, and the default value is 20 hops.

This example shows how to set the maximum hops to 40:

```
switch# configure terminal
switch(config)# spanning-tree mst max-hops 40
```

Configuring PVST Simulation Globally

You can block this automatic feature either globally or per port. You can enter the global command, and change the PVST simulation setting for the entire switch while you are in interface command mode.

SUMMARY STEPS

1. switch# **configure terminal**
2. switch(config)# **no spanning-tree mst simulate pvst global**

DETAILED STEPS

	Command or Action	Purpose
Step 1	switch# configure terminal	Enters configuration mode.
Step 2	switch(config)# no spanning-tree mst simulate pvst global	Disables all interfaces on the switch from automatically interoperating with connected switch that is running in Rapid PVST+ mode. The default for this is enabled; that is, by default, all interfaces on the switch operate seamlessly between Rapid PVST+ and MST.

This example shows how to prevent the switch from automatically interoperating with a connecting switch that is running Rapid PVST+:

```
switch# configure terminal
switch(config)# no spanning-tree mst simulate pvst global
```

Configuring PVST Simulation Per Port

MST interoperates seamlessly with Rapid PVST+. However, to prevent an accidental connection to a switch that does not run MST as the default STP mode, you may want to disable this automatic feature. If you disable PVST simulation, the MST-enabled port moves to the blocking state once it detects it is connected to a Rapid PVST+-enabled port. This port remains in the inconsistent state until the port stops receiving BPDUs, and then the port resumes the normal STP transition process.

You can block this automatic feature either globally or per port.

SUMMARY STEPS

1. switch# **configure terminal**
2. switch(config)# **interface** *{{type slot/port} | {port-channel number}}*
3. switch(config-if)# **spanning-tree mst simulate pvst disable**
4. switch(config-if)# **spanning-tree mst simulate pvst**
5. switch(config-if)# **no spanning-tree mst simulate pvst**

DETAILED STEPS

	Command or Action	Purpose
Step 1	switch# configure terminal	Enters configuration mode.

	Command or Action	Purpose
Step 2	switch(config)# interface <i>{{type slot/port}}</i> {port-channel number}	Specifies an interface to configure, and enters interface configuration mode.
Step 3	switch(config-if)# spanning-tree mst simulate pvst disable	Disables specified interfaces from automatically interoperating with connected switch that is running in Rapid PVST+ mode. By default, all interfaces on the switch operate seamlessly between Rapid PVST+ and MST.
Step 4	switch(config-if)# spanning-tree mst simulate pvst	Re-enables seamless operation between MST and Rapid PVST+ on specified interfaces.
Step 5	switch(config-if)# no spanning-tree mst simulate pvst	Sets the interface to the switch-wide MST and Rapid PVST+ interoperation that you configured using the spanning-tree mst simulate pvst global command.

This example shows how to prevent the specified interfaces from automatically interoperating with a connecting switch that is not running MST:

```
switch# configure terminal
switch(config)# interface ethernet 1/4
switch(config-if)# spanning-tree mst simulate pvst disable
```

Specifying the Link Type

Rapid connectivity (802.1w standard) is established only on point-to-point links. By default, the link type is controlled from the duplex mode of the interface. A full-duplex port is considered to have a point-to-point connection; a half-duplex port is considered to have a shared connection.

If you have a half-duplex link physically connected point-to-point to a single port on a remote switch, you can override the default setting on the link type and enable rapid transitions.

If you set the link to shared, STP reverts to 802.1D.

SUMMARY STEPS

1. switch# **configure terminal**
2. switch(config)# **interface** *type slot/port*
3. switch(config-if)# **spanning-tree link-type** {auto | point-to-point | shared}

DETAILED STEPS

	Command or Action	Purpose
Step 1	switch# configure terminal	Enters configuration mode.

	Command or Action	Purpose
Step 2	switch(config)# interface <i>type slot/port</i>	Specifies the interface to configure, and enters interface configuration mode.
Step 3	switch(config-if)# spanning-tree link-type { auto point-to-point shared }	Configures the link type to be either point to point or shared. The system reads the default value from the switch connection. Half-duplex links are shared and full-duplex links are point to point. If the link type is shared, the STP reverts to 802.1D. The default is auto, which sets the link type based on the duplex setting of the interface.

This example shows how to configure the link type as point to point:

```
switch# configure terminal
switch (config)# interface ethernet 1/4
switch(config-if)# spanning-tree link-type point-to-point
```

Restarting the Protocol

An MST bridge can detect that a port is at the boundary of a region when it receives a legacy BPDU or an MST BPDU that is associated with a different region. However, the STP protocol migration cannot determine whether the legacy switch, which is a switch that runs only IEEE 802.1D, has been removed from the link unless the legacy switch is the designated switch. Enter this command to restart the protocol negotiation (force the renegotiation with neighboring switches) on the entire switch or on specified interfaces.

SUMMARY STEPS

1. switch# **clear spanning-tree detected-protocol** [**interface** *interface* [*interface-num* | *port-channel*]]

DETAILED STEPS

	Command or Action	Purpose
Step 1	switch# clear spanning-tree detected-protocol [interface <i>interface</i> [<i>interface-num</i> <i>port-channel</i>]]	Restarts MST on entire switch or specified interfaces.

This example shows how to restart MST on the Ethernet interface on slot 2, port 8:

```
switch# clear spanning-tree detected-protocol interface ethernet 2/8
```

Verifying MST Configurations

To display MST configuration information, perform one of the following tasks:

Command	Purpose
switch# show running-config spanning-tree [all]	Displays the current spanning tree configuration.
switch# show spanning-tree mst [options]	Displays detailed information for the current MST configuration.

The following example shows how to display current MST configuration:

```
switch# show spanning-tree mst configuration
% Switch is not in mst mode
Name      [mist-attempt]
Revision  1      Instances configured 2
Instance  Vlans mapped
-----
0         1-12,14-41,43-4094
1         13,42
```




Configuring STP Extensions

This chapter describes the configuration of extensions to the Spanning Tree Protocol (STP) on Cisco Nexus 5000 Series switches. It includes the following sections:

- [About STP Extensions, page 211](#)

About STP Extensions

Cisco has added extensions to STP that make convergence more efficient. In some cases, even though similar functionality may be incorporated into the IEEE 802.1w Rapid Spanning Tree Protocol (RSTP) standard, we recommend using these extensions. All of these extensions can be used with both RPVST+ and MST.

The available extensions are spanning tree port types, Bridge Assurance, BPDU Guard, BPDU Filtering, Loop Guard, and Root Guard. Many of these features can be applied either globally or on specified interfaces.



Note

Spanning tree is used to refer to IEEE 802.1w and IEEE 802.1s. If the text is discussing the IEEE 802.1D Spanning Tree Protocol, 802.1D is stated specifically.

Information About STP Extensions

Understanding STP Port Types

You can configure a spanning tree port as an edge port, a network port, or a normal port. A port can be in only one of these states at a given time. The default spanning tree port type is normal. Depending on the type of device to which the interface is connected, you can configure a spanning tree port as one of these port types.

Spanning Tree Edge Ports

Edge ports, which are connected to hosts, can be either an access port or a trunk port. The edge port interface immediately transitions to the forwarding state, without moving through the blocking or learning states. (This immediate transition was previously configured as the Cisco-proprietary feature PortFast.)

Interfaces that are connected to hosts should not receive STP Bridge Protocol Data Units (BPDUs).

**Note**

If you configure a port connected to another switch as an edge port, you might create a bridging loop.

Spanning Tree Network Ports

Network ports are connected only to switches or bridges. Bridge Assurance is enabled only on network ports.

**Note**

If you mistakenly configure ports that are connected to hosts or other edge devices, as spanning tree network ports, those ports will automatically move into the blocking state.

Spanning Tree Normal Ports

Normal ports can be connected to either hosts, switches, or bridges. These ports function as normal spanning tree ports.

The default spanning tree interface is a normal port.

Understanding Bridge Assurance

You can use Bridge Assurance to protect against certain problems that can cause bridging loops in the network. Specifically, you use Bridge Assurance to protect against a unidirectional link failure and a device that continues to forward data traffic when it is no longer running the spanning tree algorithm.

**Note**

Bridge Assurance is supported only by Rapid PVST+ and MST. Legacy 802.1D spanning tree does not support Bridge Assurance.

Bridge Assurance is enabled by default and can only be disabled globally. Also, Bridge Assurance can be enabled only on spanning tree network ports that are point-to-point links. Finally, both ends of the link must have Bridge Assurance enabled.

With Bridge Assurance enabled, BPDUs are sent out on all operational network ports, including alternate and backup ports, for each hello time period. If the port does not receive a BPDU for a specified period, the port moves into the blocking state and is not used in the root port calculation. Once that port receives a BPDU, it resumes the normal spanning tree transitions.

Understanding BPDU Guard

Enabling BPDU Guard shuts down that interface if a BPDU is received.

You can configure BPDU Guard at the interface level. When configured at the interface level, BPDU Guard shuts the port down as soon as the port receives a BPDU, regardless of the port type configuration.

When you configure BPDU Guard globally, it is effective only on operational spanning tree edge ports. In a valid configuration, LAN edge interfaces do not receive BPDUs. A BPDU that is received by an edge LAN interface signals an invalid configuration, such as the connection of an unauthorized host or switch. BPDU Guard, when enabled globally, shuts down all spanning tree edge ports when they receive a BPDU.

BPDU Guard provides a secure response to invalid configurations, because you must manually put the LAN interface back in service after an invalid configuration.

**Note**

When enabled globally, BPDU Guard applies to all operational spanning tree edge interfaces.

Understanding BPDU Filtering

You can use BPDU Filtering to prevent the switch from sending or even receiving BPDUs on specified ports.

When configured globally, BPDU Filtering applies to all operational spanning tree edge ports. You should connect edge ports only to hosts, which typically drop BPDUs. If an operational spanning tree edge port receives a BPDU, it immediately returns to a normal spanning tree port type and moves through the regular transitions. In that case, BPDU Filtering is disabled on this port, and spanning tree resumes sending BPDUs on this port.

In addition, you can configure BPDU Filtering by the individual interface. When you explicitly configure BPDU Filtering on a port, that port does not send any BPDUs and drops all BPDUs that it receives. You can effectively override the global BPDU Filtering setting on individual ports by configuring the specific interface. This BPDU Filtering command on the interface applies to the entire interface, whether the interface is trunking or not.

**Caution**

Use care when configuring BPDU Filtering per interface. If you explicitly configuring BPDU Filtering on a port that is not connected to a host, it can result in bridging loops because the port will ignore any BPDU that it receives and go to forwarding.

If the port configuration is not set to default BPDU Filtering, then the edge configuration will not affect BPDU Filtering. The following table lists all the BPDU Filtering combinations.

Table 18: BPDU Filtering Configurations

BPDU Filtering Per Port Configuration	BPDU Filtering Global Configuration	STP Edge Port Configuration	BPDU Filtering State
Default	Enable	Enable	Enable The port transmits at least 10 BPDUs. If this port receives any BPDUs, the port returns to the spanning tree normal port state and BPDU Filtering is disabled.
Default	Enable	Disable	Disable
Default	Disable	Not applicable	Disable
Disable	Not applicable	Not applicable	Disable
Enable	Not applicable	Not applicable	Enable

Understanding Loop Guard

Loop Guard protects networks from loops that are caused by the following:

- Network interfaces that malfunction
- Busy CPUs
- Anything that prevents the normal forwarding of BPDUs

An STP loop occurs when a blocking port in a redundant topology erroneously transitions to the forwarding state. This transition usually happens because one of the ports in a physically redundant topology (not necessarily the blocking port) stops receiving BPDUs.

Loop Guard is only useful in switched networks where devices are connected by point-to-point links. On a point-to-point link, a designated bridge cannot disappear unless it sends an inferior BPDU or brings the link down.



Note

Loop Guard can be enabled only on network and normal spanning tree port types.

You can use Loop Guard to determine if a root port or an alternate/backup root port receives BPDUs. If the port does not receive BPDUs, Loop Guard puts the port into an inconsistent state (blocking) until the port starts to receive BPDUs again. A port in the inconsistent state does not transmit BPDUs. If the port receives BPDUs again, the protocol removes its loop-inconsistent condition, and the STP determines the port state because such recovery is automatic.

Loop Guard isolates the failure and allows STP to converge to a stable topology without the failed link or bridge. Disabling Loop Guard moves all loop-inconsistent ports to the listening state.

You can enable Loop Guard on a per-port basis. When you enable Loop Guard on a port, it is automatically applied to all of the active instances or VLANs to which that port belongs. When you disable Loop Guard, it is disabled for the specified ports.

Understanding Root Guard

When you enable Root Guard on a port, Root Guard does not allow that port to become a root port. If a received BPDU triggers an STP convergence that makes that designated port become a root port, that port is put into a root-inconsistent (blocked) state. After the port stops send superior BPDUs, the port is unblocked again. Through STP, the port moves to the forwarding state. Recovery is automatic.

Root Guard enabled on an interface applies this functionality to all VLANs to which that interface belongs.

You can use Root Guard to enforce the root bridge placement in the network. Root Guard ensures that the port on which Root Guard is enabled is the designated port. Normally, root bridge ports are all designated ports, unless two or more of the ports of the root bridge are connected. If the bridge receives superior BPDUs on a Root Guard-enabled port, the bridge moves this port to a root-inconsistent STP state. In this way, Root Guard enforces the position of the root bridge.

You cannot configure Root Guard globally.



Note

You can enable Root Guard on all spanning tree port types: normal, edge, and network ports.

Configuring STP Extensions

STP Extensions Configuration Guidelines

When configuring STP extensions, follow these guidelines:

- Configure all access and trunk ports connected to hosts as edge ports.
- Bridge Assurance runs only on point-to-point spanning tree network ports. You must configure each side of the link for this feature.
- Loop Guard does not run on spanning tree edge ports.
- Enabling Loop Guard on ports that are not connected to a point-to-point link will not work.
- You cannot enable Loop Guard if Root Guard is enabled.

Configuring Spanning Tree Port Types Globally

The spanning tree port type designation depends on the type of device the port is connected to, as follows:

- Edge—Edge ports are connected to hosts and can be either an access port or a trunk port.
- Network—Network ports are connected only to switches or bridges.
- Normal—Normal ports are neither edge ports nor network ports; they are normal spanning tree ports. These ports can be connected to any type of device.

You can configure the port type either globally or per interface. By default, the spanning tree port type is normal.

Before You Begin

Ensure that STP is configured.

Ensure that you are configuring the ports correctly for the type of device to which the interface is connected.

SUMMARY STEPS

1. switch# **configure terminal**
2. switch(config)# **spanning-tree port type edge default**
3. switch(config)# **spanning-tree port type network default**

DETAILED STEPS

	Command or Action	Purpose
Step 1	switch# configure terminal	Enters configuration mode.
Step 2	switch(config)# spanning-tree port type edge default	Configures all interfaces as edge ports. This assumes all ports are connected to hosts/servers. Edge ports immediately transition to the forwarding state without passing through the blocking or learning state at linkup. By default, spanning tree ports are normal port types.

	Command or Action	Purpose
Step 3	<code>switch(config)# spanning-tree port type network default</code>	Configures all interfaces as spanning tree network ports. This assumes all ports are connected to switches and bridges. If you enable Bridge Assurance, it automatically runs on network ports. By default, spanning tree ports are normal port types. Note If you configure interfaces connected to hosts as network ports, those ports automatically move into the blocking state.

This example shows how to configure all access and trunk ports connected to hosts as spanning tree edge ports:

```
switch# configure terminal
switch(config)# spanning-tree port type edge default
```

This example shows how to configure all ports connected to switches or bridges as spanning tree network ports:

```
switch# configure terminal
switch(config)# spanning-tree port type network default
```

Configuring Spanning Tree Edge Ports on Specified Interfaces

You can configure spanning tree edge ports on specified interfaces. Interfaces configured as spanning tree edge ports immediately transition to the forwarding state, without passing through the blocking or learning states, on linkup.

This command has four states:

- **spanning-tree port type edge**—This command explicitly enables edge behavior on the access port.
- **spanning-tree port type edge trunk**—This command explicitly enables edge behavior on the trunk port.



Note If you enter the **spanning-tree port type edge trunk** command, the port is configured as an edge port even in the access mode.

- **spanning-tree port type normal**—This command explicitly configures the port as a normal spanning tree port and the immediate transition to the forwarding state is not enabled.
- **no spanning-tree port type**—This command implicitly enables edge behavior if you define the **spanning-tree port type edge default** command in global configuration mode. If you do not configure the edge ports globally, the **no spanning-tree port type** command is equivalent to the **spanning-tree port type disable** command.

Before You Begin

Ensure that STP is configured.

Ensure that the interface is connected to hosts.

SUMMARY STEPS

1. switch# **configure terminal**
2. switch(config)# **interface type slot/port**
3. switch(config-if)# **spanning-tree port type edge**

DETAILED STEPS

	Command or Action	Purpose
Step 1	switch# configure terminal	Enters configuration mode.
Step 2	switch(config)# interface type slot/port	Specifies the interface to configure, and enters the interface configuration mode.
Step 3	switch(config-if)# spanning-tree port type edge	Configures the specified access interfaces to be spanning edge ports. Edge ports immediately transition to the forwarding state without passing through the blocking or learning state at linkup. By default, spanning tree ports are normal port types.

This example shows how to configure the Ethernet access interface 1/4 to be a spanning tree edge port:

```
switch# configure terminal
switch(config)# interface ethernet 1/4
switch(config-if)# spanning-tree port type edge
```

Configuring Spanning Tree Network Ports on Specified Interfaces

You can configure spanning tree network ports on specified interfaces.

Bridge Assurance runs only on spanning tree network ports.

This command has three states:

- **spanning-tree port type network**—This command explicitly configures the port as a network port. If you enable Bridge Assurance globally, it automatically runs on a spanning tree network port.
- **spanning-tree port type normal**—This command explicitly configures the port as a normal spanning tree port and Bridge Assurance cannot run on this interface.
- **no spanning-tree port type**—This command implicitly enables the port as a spanning tree network port if you define the **spanning-tree port type network default** command in global configuration mode. If you enable Bridge Assurance globally, it automatically runs on this port.



Note

A port connected to a host that is configured as a network port automatically moves into the blocking state.

Before You Begin

Ensure that STP is configured.

Ensure that the interface is connected to switches or routers.

SUMMARY STEPS

1. switch# **configure terminal**
2. switch(config)# **interface** type slot/port
3. switch(config-if)# **spanning-tree port type network**

DETAILED STEPS

	Command or Action	Purpose
Step 1	switch# configure terminal	Enters configuration mode.
Step 2	switch(config)# interface type slot/port	Specifies the interface to configure, and enters the interface configuration mode. The interface can be a physical Ethernet port.
Step 3	switch(config-if)# spanning-tree port type network	Configures the specified interfaces to be spanning network ports. If you enable Bridge Assurance, it automatically runs on network ports. By default, spanning tree ports are normal port types.

This example shows how to configure the Ethernet interface 1/4 to be a spanning tree network port:

```
switch# configure terminal
switch(config)# interface ethernet 1/4
switch(config-if)# spanning-tree port type network
```

Enabling BPDU Guard Globally

You can enable BPDU Guard globally by default. In this condition, the system shuts down an edge port that receives a BPDU.

**Note**

We recommend that you enable BPDU Guard on all edge ports.

Before You Begin

Ensure that STP is configured.

Ensure that you have configured some spanning tree edge ports.

SUMMARY STEPS

1. switch# **configure terminal**
2. switch(config)# **spanning-tree port type edge bpduguard default**

DETAILED STEPS

	Command or Action	Purpose
Step 1	switch# configure terminal	Enters configuration mode.
Step 2	switch(config)# spanning-tree port type edge bpduguard default	Enables BPDU Guard by default on all spanning tree edge ports. By default, global BPDU Guard is disabled.

This example shows how to enable BPDU Guard on all spanning tree edge ports:

```
switch# configure terminal
switch(config)# spanning-tree port type edge bpduguard default
```

Enabling BPDU Guard on Specified Interfaces

You can enable BPDU Guard on specified interfaces. Enabling BPDU Guard shuts down the port if it receives a BPDU.

You can configure BPDU Guard on specified interfaces as follows:

- **spanning-tree bpduguard enable**—Unconditionally enables BPDU Guard on the interface.
- **spanning-tree bpduguard disable**—Unconditionally disables BPDU Guard on the interface.
- **no spanning-tree bpduguard**—Enables BPDU Guard on the interface if it is an operational edge port and if the **spanning-tree port type edge bpduguard default** command is configured.

Before You Begin

Ensure that STP is configured.

SUMMARY STEPS

1. switch# **configure terminal**
2. switch(config)# **interface** *type slot/port*
3. switch(config-if)# **spanning-tree bpduguard** {enable | disable}
4. (Optional) switch(config-if)# **no spanning-tree bpduguard**

DETAILED STEPS

	Command or Action	Purpose
Step 1	switch# configure terminal	Enters configuration mode.
Step 2	switch(config)# interface <i>type slot/port</i>	Specifies the interface to configure, and enters the interface configuration mode.

	Command or Action	Purpose
Step 3	switch(config-if)# spanning-tree bpduguard {enable disable}	Enables or disables BPDU Guard for the specified spanning tree edge interface. By default, BPDU Guard is disabled on physical Ethernet interfaces.
Step 4	switch(config-if)# no spanning-tree bpduguard	(Optional) Disables BPDU Guard on the interface. Note Enables BPDU Guard on the interface if it is an operational edge port and if you enter the spanning-tree port type edge bpduguard default command.

This example shows how to explicitly enable BPDU Guard on the Ethernet edge port 1/4:

```
switch# configure terminal
switch (config)# interface ethernet 1/4
switch(config-if)# spanning-tree bpduguard enable

switch(config-if)# no spanning-tree bpduguard
```

Enabling BPDU Filtering Globally

You can enable BPDU Filtering globally by default on spanning tree edge ports.

If an edge port with BPDU Filtering enabled receives a BPDU, it loses its operation status and as edge port and resumes the regular STP transitions. However, this port maintains its configuration as an edge port.



Caution

Be careful when using this command: using it incorrectly can cause bridging loops.



Note

When enabled globally, BPDU Filtering is applied *only* on ports that are operational edge ports. Ports send a few BPDUs at linkup before they effectively filter outbound BPDUs. If a BPDU is received on an edge port, it immediately loses its operational edge port status and BPDU Filtering is disabled.

Before You Begin

Ensure that STP is configured.

Ensure that you have configured some spanning tree edge ports.

SUMMARY STEPS

1. switch# **configure terminal**
2. switch(config)# **spanning-tree port type edge bpdudfilter default**

DETAILED STEPS

	Command or Action	Purpose
Step 1	switch# configure terminal	Enters configuration mode.
Step 2	switch(config)# spanning-tree port type edge bpdudfilter default	Enables BPDU Filtering by default on all operational spanning tree edge ports. Global BPDU Filtering is disabled by default.

This example shows how to enable BPDU Filtering on all operational spanning tree edge ports:

```
switch# configure terminal
switch(config)# spanning-tree port type edge bpdudfilter default
```

Enabling BPDU Filtering on Specified Interfaces

You can apply BPDU Filtering to specified interfaces. When enabled on an interface, that interface does not send any BPDUs and drops all BPDUs that it receives. This BPDU Filtering functionality applies to the entire interface, whether trunking or not.



Caution

Be careful when you enter the **spanning-tree bpdudfilter enable** command on specified interfaces. Explicitly configuring BPDU Filtering on a port that is not connected to a host can result in bridging loops as the port will ignore any BPDU it receives and go to forwarding.

You can enter this command to override the port configuration on specified interfaces.

This command has three states:

- **spanning-tree bpdudfilter enable**—Unconditionally enables BPDU Filtering on the interface.
- **spanning-tree bpdudfilter disable**—Unconditionally disables BPDU Filtering on the interface.
- **no spanning-tree bpdudfilter**—Enables BPDU Filtering on the interface if the interface is in operational edge port and if you configure the **spanning-tree port type edge bpdudfilter default** command.



Note

When you enable BPDU Filtering locally on a port, this feature prevents the device from receiving or sending BPDUs on this port.

Before You Begin

Ensure that STP is configured.

SUMMARY STEPS

1. switch# **configure terminal**
2. switch(config)# **interface** *type slot/port*
3. switch(config-if)# **spanning-tree bpdufilter** {enable | disable}
4. (Optional) switch(config-if)# **no spanning-tree bpdufilter**

DETAILED STEPS

	Command or Action	Purpose
Step 1	switch# configure terminal	Enters configuration mode.
Step 2	switch(config)# interface <i>type slot/port</i>	Specifies the interface to configure, and enters the interface configuration mode.
Step 3	switch(config-if)# spanning-tree bpdufilter {enable disable}	Enables or disables BPDU Filtering for the specified spanning tree edge interface. By default, BPDU Filtering is disabled.
Step 4	switch(config-if)# no spanning-tree bpdufilter	(Optional) Disables BPDU Filtering on the interface. Note Enables BPDU Filtering on the interface if the interface is an operational spanning tree edge port and if you enter the spanning-tree port type edge bpdufilter default command.

This example shows how to explicitly enable BPDU Filtering on the Ethernet spanning tree edge port 1/4:

```
switch# configure terminal
switch (config)# interface ethernet 1/4
switch(config-if)# spanning-tree bpdufilter enable
```

Enabling Loop Guard Globally

You can enable Loop Guard globally by default on all point-to-point spanning tree normal and network ports. Loop Guard does not run on edge ports.

Loop Guard provides additional security in the bridge network. Loop Guard prevents alternate or root ports from becoming the designated port because of a failure that could lead to a unidirectional link.



Note Entering the Loop Guard command for the specified interface overrides the global Loop Guard command.

Before You Begin

Ensure that STP is configured.

Ensure that you have spanning tree normal ports or have configured some network ports.

SUMMARY STEPS

1. switch# **configure terminal**
2. switch(config)# **spanning-tree loopguard default**

DETAILED STEPS

	Command or Action	Purpose
Step 1	switch# configure terminal	Enters configuration mode.
Step 2	switch(config)# spanning-tree loopguard default	Enables Loop Guard by default on all spanning tree normal and network ports. By default, global Loop Guard is disabled.

This example shows how to enable Loop Guard on all spanning tree normal or network ports:

```
switch# configure terminal
switch(config)# spanning-tree loopguard default
```

Enabling Loop Guard or Root Guard on Specified Interfaces

You can enable either Loop Guard or Root Guard on specified interfaces.

Enabling Root Guard on a port means that port cannot become a root port, and LoopGuard prevents alternate or root ports from becoming the designated port because of a failure that could lead to a unidirectional link.

Both Loop Guard and Root Guard enabled on an interface apply to all VLANs to which that interface belongs.

**Note**

Entering the Loop Guard command for the specified interface overrides the global Loop Guard command.

Before You Begin

Ensure that STP is configured.

Ensure that you are configuring Loop Guard on spanning tree normal or network ports.

SUMMARY STEPS

1. switch# **configure terminal**
2. switch(config)# **interface type slot/port**
3. switch(config-if)# **spanning-tree guard {loop | root | none}**

DETAILED STEPS

	Command or Action	Purpose
Step 1	switch# configure terminal	Enters configuration mode.
Step 2	switch(config)# interface <i>type slot/port</i>	Specifies the interface to configure, and enters the interface configuration mode.
Step 3	switch(config-if)# spanning-tree guard { loop root none }	Enables or disables either Loop Guard or Root Guard for the specified interface. By default, Root Guard is disabled by default, and Loop Guard on specified ports is also disabled. Note Loop Guard runs only on spanning tree normal and network interfaces.

This example shows how to enable Root Guard on Ethernet port 1/4:

```
switch# configure terminal
switch (config)# interface ethernet 1/4
switch(config-if)# spanning-tree guard root
```

Verifying STP Extension Configuration

To display the configuration information for the STP extensions, perform one of the following tasks:

Command	Purpose
switch# show running-config spanning-tree [all]	Displays the current status of spanning tree on the switch
switch# show spanning-tree [<i>options</i>]	Displays selected detailed information for the current spanning tree configuration.



Configuring the MAC Address Table

All Ethernet interfaces on Cisco Nexus 5000 Series switches maintain media access control (MAC) address tables. This chapter describes the configuration of the MAC address tables. It includes the following sections:

- [Information About MAC Addresses, page 225](#)
- [Configuring MAC Addresses, page 225](#)
- [Verifying the MAC Address Configuration, page 227](#)

Information About MAC Addresses

To switch frames between LAN ports efficiently, the switch maintains an address table. When the switch receives a frame, it associates the media access control (MAC) address of the sending network device with the LAN port on which it was received.

The switch dynamically builds the address table by using the MAC source address of the frames received. When the switch receives a frame for a MAC destination address not listed in its address table, it floods the frame to all LAN ports of the same VLAN except the port that received the frame. When the destination station replies, the switch adds its relevant MAC source address and port ID to the address table. The switch then forwards subsequent frames to a single LAN port without flooding all LAN ports.

You can also enter a MAC address, which is termed a static MAC address, into the table. These static MAC entries are retained across a reboot of the switch.

In addition, you can enter a multicast address as a statically configured MAC address. A multicast address can accept more than one interface as its destination.

The address table can store a number of unicast and multicast address entries without flooding any frames. The switch uses an aging mechanism, defined by a configurable aging timer, so if an address remains inactive for a specified number of seconds, it is removed from the address table.

Configuring MAC Addresses

Configuring a Static MAC Address

You can configure MAC addresses for the switch. These addresses are static MAC addresses.

**Note**

You can also configure a static MAC address in interface configuration mode or VLAN configuration mode.

SUMMARY STEPS

1. switch# **configure terminal**
2. switch(config)# **mac-address-table static** *mac_address* **vlan** *vlan-id* **{drop | interface** *{type slot/port}* **| port-channel** *number* **} [auto-learn]**
3. (Optional) switch(config)# **no mac-address-table static** *mac_address* **vlan** *vlan-id*

DETAILED STEPS

	Command or Action	Purpose
Step 1	switch# configure terminal	Enters configuration mode.
Step 2	switch(config)# mac-address-table static <i>mac_address</i> vlan <i>vlan-id</i> {drop interface <i>{type slot/port}</i> port-channel <i>number</i> } [auto-learn]	Specifies a static address to add to the MAC address table. If you enable the auto-learn option, the switch will update the entry if the same MAC address is seen on a different port.
Step 3	switch(config)# no mac-address-table static <i>mac_address</i> vlan <i>vlan-id</i>	(Optional) Deletes the static entry from the MAC address table.

This example shows how to put a static entry in the MAC address table:

```
switch# configure terminal
switch(config)# mac-address-table static 12ab.47dd.ff89 vlan 3 interface ethernet 2/1
```

You can use the **mac-address-table static** command to assign a static MAC address to a virtual interface.

Configuring the Aging Time for the MAC Table

You can configure the amount of time that an entry (the packet source MAC address and port that packet ingresses) remain in the MAC table.

**Note**

You can also configure MAC aging time in interface configuration mode or VLAN configuration mode.

SUMMARY STEPS

1. switch# **configure terminal**
2. switch(config)# **mac-address-table aging-time** *seconds* [**vlan** *vlan_id*]

DETAILED STEPS

	Command or Action	Purpose
Step 1	switch# configure terminal	Enters configuration mode.
Step 2	switch(config)# mac-address-table aging-time <i>seconds</i> [<i>vlan vlan_id</i>]	Specifies the time before an entry ages out and is discarded from the MAC address table. The range is from 0 to 1000000; the default is 300 seconds. Entering the value 0 disables the MAC aging. If a VLAN is not specified, the aging specification applies to all VLANs.

This example shows how to set the aging time for entries in the MAC address table to 600 seconds (10 minutes):

```
switch# configure terminal
switch(config)# mac-address-table aging-time 600
```

Clearing Dynamic Addresses from the MAC Table

You can clear all dynamic entries in the MAC address table.

Command	Purpose
switch(config)# clear mac-address-table dynamic { <i>address mac-addr</i> } { <i>interface [type slot/port port-channel number]</i> } { <i>vlan vlan-id</i> }	Clears the dynamic address entries from the MAC address table.

This example shows how to clear the dynamic entries in the MAC address table:

```
switch# clear mac-address-table dynamic
```

Verifying the MAC Address Configuration

To display MAC address configuration information, perform one of these tasks:

Command	Purpose
switch# show mac-address-table aging-time	Displays the MAC address aging time for all VLANs defined in the switch.
switch# show mac-address-table	Displays the contents of the MAC address table.

This example shows how to display the MAC address table:

```
switch# show mac-address-table
VLAN      MAC Address      Type      Age      Port
-----+-----+-----+-----+-----
1         0018.b967.3cd0   dynamic  10      Eth1/3
1         001c.b05a.5380   dynamic  200     Eth1/3
Total MAC Addresses: 2
```

This example shows how to display the current aging time:

```
switch# show mac-address-table aging-time
Vlan      Aging Time
-----
1         300
13        300
42        300
```



Configuring IGMP Snooping

Internet Group Management Protocol (IGMP) snooping streamlines multicast traffic handling for VLANs. By examining (snooping) IGMP membership report messages from interested hosts, multicast traffic is limited to the subset of VLAN interfaces on which the hosts reside.

This chapter describes the configuration of IGMP snooping on Cisco Nexus 5000 Series switches. It includes the following sections:

- [Information About IGMP Snooping, page 229](#)
- [Configuring IGMP Snooping Parameters, page 232](#)
- [Verifying IGMP Snooping Configuration, page 234](#)

Information About IGMP Snooping

The IGMP snooping software examines IGMP protocol messages within a VLAN to discover which interfaces are connected to hosts or other devices interested in receiving this traffic. Using the interface information, IGMP snooping can reduce bandwidth consumption in a multi-access LAN environment to avoid flooding the entire VLAN. The IGMP snooping feature tracks which ports are attached to multicast-capable routers to help it manage the forwarding of IGMP membership reports. The IGMP snooping software responds to topology change notifications.



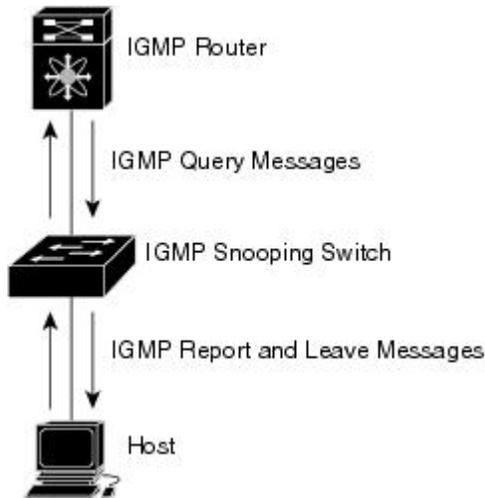
Note

IGMP snooping is supported on all Ethernet interfaces. The term *snooping* is used because Layer 3 control plane packets are intercepted and influence Layer 2 forwarding decisions.

Cisco NX-OS supports IGMPv2 and IGMPv3. IGMPv2 supports IGMPv1, and IGMPv3 supports IGMPv2. Although not all features of an earlier version of IGMP are supported, the features related to membership query and membership report messages are supported for all IGMP versions.

The following figure shows an IGMP snooping switch that is located between the host and the IGMP router. The IGMP snooping switch snoops the IGMP membership reports and leave messages and forwards them only when necessary to the connected IGMP routers.

Figure 30: IGMP Snooping Switch



Note

The switch supports IGMPv3 snooping based only on the destination multicast MAC address. It does not support snooping based on the source MAC address or on proxy reports.

The Cisco NX-OS IGMP snooping software supports optimized multicast flooding (OMF) that forwards unknown traffic to routers only and performs no data driven state creation. For more information about IGMP snooping, see <http://tools.ietf.org/wg/magma/draft-ietf-magma-snoop/rfc4541.txt>.

IGMPv1 and IGMPv2

Both IGMPv1 and IGMPv2 support membership report suppression, which means that if two hosts on the same subnet want to receive multicast data for the same group, then the host that receives a member report from the other host suppresses sending its report. Membership report suppression occurs for hosts that share a port.

If no more than one host is attached to each VLAN switch port, then you can configure the fast leave feature in IGMPv2. The fast leave feature does not send last member query messages to hosts. As soon as the software receives an IGMP leave message, the software stops forwarding multicast data to that port.

IGMPv1 does not provide an explicit IGMP leave message, so the software must rely on the membership message timeout to indicate that no hosts remain that want to receive multicast data for a particular group.



Note

Cisco NX-OS ignores the configuration of last member query interval when you enable the fast leave feature because it does not check for remaining hosts.

IGMPv3

The IGMPv3 snooping implementation on the switch forwards IGMPv3 reports to allow the upstream multicast router do source-based filtering.

By default, the software tracks hosts on each VLAN port. The explicit tracking feature provides a fast leave mechanism. Because every IGMPv3 host sends membership reports, a report suppression feature limits the amount of traffic the switch sends to other multicast capable routers. When report suppression is enabled, and no IGMPv1 or IGMPv2 hosts requested the same group, the software provides proxy reporting. The proxy feature builds group state from membership reports from the downstream hosts and generates membership reports in response to queries from upstream queriers.

Even though the IGMPv3 membership reports provide a full accounting of group members on a LAN segment, when the last host leaves, the software sends a membership query. You can configure the parameter last member query interval. If no host responds before the timeout, the software removes the group state.

IGMP Snooping Querier

When there is no multicast router in the VLAN to originate the queries, you must configure an IGMP snooping querier to send membership queries.

When an IGMP snooping querier is enabled, it sends out periodic IGMP queries that trigger IGMP report messages from hosts that want to receive IP multicast traffic. IGMP snooping listens to these IGMP reports to establish appropriate forwarding.

IGMP Forwarding

The control plane of the Cisco Nexus 5000 Series switch is able to detect IP addresses but forwarding occurs using the MAC address only.

When a host connected to the switch wants to join an IP multicast group, it sends an unsolicited IGMP join message, specifying the IP multicast group to join. Alternatively, when the switch receives a general query from a connected router, it forwards the query to all interfaces, physical and virtual, in the VLAN. Hosts wanting to join the multicast group respond by sending a join message to the switch. The switch CPU creates a multicast forwarding table entry for the group if it is not already present. The CPU also adds the interface where the join message was received to the forwarding table entry. The host associated with that interface receives multicast traffic for that multicast group.

The router sends periodic multicast general queries and the switch forwards these queries through all ports in the VLAN. Interested hosts respond to the queries. If at least one host in the VLAN wants to receive multicast traffic, the router continues forwarding the multicast traffic to the VLAN. The switch forwards multicast group traffic to only those hosts listed in the forwarding table for that multicast group.

When hosts want to leave a multicast group, they can either silently leave, or they can send a leave message. When the switch receives a leave message from a host, it sends a group-specific query to determine if any other devices connected to that interface are interested in traffic for the specific multicast group. The switch then updates the forwarding table for that MAC group so that only those hosts interested in receiving multicast traffic for the group are listed in the forwarding table. If the router receives no reports from a VLAN, it removes the group for the VLAN from its IGMP cache.

Configuring IGMP Snooping Parameters

To manage the operation of the IGMP snooping process, you can configure the optional IGMP snooping parameters described in the following table.

Table 19: IGMP Snooping Parameters

Parameter	Description
IGMP snooping	Enables IGMP snooping on a per-VLAN basis. The default is enabled. Note If the global setting is disabled, then all VLANs are treated as disabled, whether they are enabled or not.
Explicit tracking	Tracks IGMPv3 membership reports from individual hosts for each port on a per-VLAN basis. The default is enabled.
Fast leave	Enables the software to remove the group state when it receives an IGMP Leave report without sending an IGMP query message. This parameter is used for IGMPv2 hosts when no more than one host is present on each VLAN port. The default is disabled.
Last member query interval	Sets the interval that the software waits after sending an IGMP query to verify that no hosts that want to receive a particular multicast group remain on a network segment. If no hosts respond before the last member query interval expires, the software removes the group from the associated VLAN port. Values range from 1 to 25 seconds. The default is 1 second.
Snooping querier	Configures a snooping querier on an interface when there is no multicast router in the VLAN to generate queries. The default is disabled.
Report suppression	Limits the membership report traffic sent to multicast-capable routers. When you disable report suppression, all IGMP reports are sent as is to multicast-capable routers. The default is enabled.
Multicast router	Configures a static connection to a multicast router. The interface to the router must be in the selected VLAN.
Static group	Configures an interface belonging to a VLAN as a static member of a multicast group.

You can disable IGMP snooping either globally or for a specific VLAN.

SUMMARY STEPS

1. switch# **configure terminal**
2. switch(config)# **ip igmp snooping**
3. switch(config)# **vlan *vlan-id***
4. switch(config-vlan)# **ip igmp snooping**
5. switch(config-vlan)# **ip igmp snooping explicit-tracking**
6. switch(config-vlan)# **ip igmp snooping fast-leave**
7. switch(config-vlan)# **ip igmp snooping last-member-query-interval *seconds***
8. switch(config-vlan)# **ip igmp snooping querier *IP-address***
9. switch(config-vlan)# **ip igmp snooping report-suppression**
10. switch(config-vlan)# **ip igmp snooping mrouter interface *interface***
11. switch(config-vlan)# **ip igmp snooping static-group *group-ip-addr* [*source source-ip-addr*] interface *interface***

DETAILED STEPS

	Command or Action	Purpose
Step 1	switch# configure terminal	Enters configuration mode.
Step 2	switch(config)# ip igmp snooping	Globally enables IGMP snooping. The default is enabled. Note If the global setting is disabled, then all VLANs are treated as disabled, whether they are enabled or not.
Step 3	switch(config)# vlan <i>vlan-id</i>	Enters VLAN configuration mode.
Step 4	switch(config-vlan)# ip igmp snooping	Enables IGMP snooping for the current VLAN. The default is enabled. Note If IGMP snooping is enabled globally, this command is not required.
Step 5	switch(config-vlan)# ip igmp snooping explicit-tracking	Tracks IGMPv3 membership reports from individual hosts for each port on a per-VLAN basis. The default is enabled on all VLANs.
Step 6	switch(config-vlan)# ip igmp snooping fast-leave	Supports IGMPv2 hosts that cannot be explicitly tracked because of the host report suppression mechanism of the IGMPv2 protocol. When you enable fast leave, the IGMP software assumes that no more than one host is present on each VLAN port. The default is disabled for all VLANs.
Step 7	switch(config-vlan)# ip igmp snooping last-member-query-interval <i>seconds</i>	Removes the group from the associated VLAN port if no hosts respond to an IGMP query message before the last member query interval expires. Values range from 1 to 25 seconds. The default is 1 second.
Step 8	switch(config-vlan)# ip igmp snooping querier <i>IP-address</i>	Configures a snooping querier when you do not enable PIM because multicast traffic does not need to be routed. The IP address is used as the source in messages. The default is disabled.

	Command or Action	Purpose
Step 9	switch(config-vlan)# ip igmp snooping report-suppression	Limits the membership report traffic sent to multicast-capable routers. When you disable report suppression, all IGMP reports are sent as is to multicast-capable routers. The default is enabled.
Step 10	switch(config-vlan)# ip igmp snooping mrouter interface interface	Configures a static connection to a multicast router. The interface to the router must be in the selected VLAN. You can specify the interface by type and number.
Step 11	switch(config-vlan)# ip igmp snooping static-group group-ip-addr [source source-ip-addr] interface interface	Configures an interface belonging to a VLAN as a static member of a multicast group. You can specify the interface by type and number.

The following example shows configuring IGMP snooping parameters for a VLAN:

```
switch# configure terminal
switch(config)# vlan 5
switch(config-vlan)# ip igmp snooping last-member-query-interval 3
switch(config-vlan)# ip igmp snooping querier 172.20.52.106
switch(config-vlan)# ip igmp snooping explicit-tracking
switch(config-vlan)# ip igmp snooping fast-leave
switch(config-vlan)# ip igmp snooping report-suppression
switch(config-vlan)# ip igmp snooping mrouter interface ethernet 1/10
switch(config-vlan)# ip igmp snooping static-group 230.0.0.1 interface ethernet 1/10
switch(config-vlan)# end
```

Verifying IGMP Snooping Configuration

To verify the IGMP snooping configuration, perform one of these tasks:

Command	Description
switch# show ip igmp snooping [[vlan] <i>vlan-id</i>]	IGMP snooping configuration by VLAN.
switch# show ip igmp snooping groups [[vlan] <i>vlan-id</i>] [detail]	IGMP snooping information about groups by VLAN.
switch# show ip igmp snooping querier [[vlan] <i>vlan-id</i>]	IGMP snooping queriers by VLAN.
switch# show ip igmp snooping mrouter [[vlan] <i>vlan-id</i>]	Multicast router ports by VLAN.
switch# show ip igmp snooping explicit-tracking vlan <i>vlan-id</i>	IGMP snooping explicit tracking information by VLAN.

The following example shows how to verify the IGMP snooping parameters:

```
switch# show ip igmp snooping
Global IGMP Snooping Information:
  IGMP Snooping enabled
```

```
IGMP Snooping information for vlan 1
  IGMP snooping enabled
  IGMP querier none
  Switch-querier disabled
  Explicit tracking enabled
  Fast leave disabled
  Report suppression enabled
  Router port detection using PIM Hellos, IGMP Queries
  Number of router-ports: 0
  Number of groups: 0
IGMP Snooping information for vlan 5
IGMP snooping enabled
  IGMP querier present, address: 172.16.24.1, version: 3
  Querier interval: 125 secs
  Querier last member query interval: 10 secs
  Querier robustness: 2
  Switch-querier enabled, address 172.16.24.1, currently running
  Explicit tracking enabled
  Fast leave enabled
  Report suppression enabled
  Router port detection using PIM Hellos, IGMP Queries
  Number of router-ports: 1
  Number of groups: 1
```




CHAPTER 17

Configuring Traffic Storm Control

This chapter describes how to configure traffic storm control on Cisco Nexus 5000 Series switches. It contains the following sections:

- [Information About Traffic Storm Control, page 237](#)
- [Traffic Storm Guidelines and Limitations, page 238](#)
- [Configuring Traffic Storm Control, page 239](#)
- [Traffic Storm Control Example Configuration, page 240](#)
- [Default Traffic Storm Settings, page 240](#)

Information About Traffic Storm Control

A traffic storm occurs when packets flood the LAN, creating excessive traffic and degrading network performance. You can use the traffic storm control feature to prevent disruptions on Ethernet interfaces by a broadcast, multicast, or unknown unicast traffic storm.

Traffic storm control (also called traffic suppression) allows you to monitor the levels of the incoming broadcast, multicast, and unicast traffic over a 10-microsecond interval. During this interval, the traffic level, which is a percentage of the total available bandwidth of the port, is compared with the traffic storm control level that you configured. When the ingress traffic reaches the traffic storm control level that is configured on the port, traffic storm control drops the traffic until the interval ends.

The following figure shows the broadcast traffic patterns on an Ethernet interface during a specified time interval. In this example, traffic storm control occurs between times T1 and T2 and between T4 and T5. During those intervals, the amount of broadcast traffic exceeded the configured threshold.

The traffic storm control threshold numbers and the time interval allow the traffic storm control algorithm to work with different levels of packet granularity. For example, a higher threshold allows more packets to pass through.

Traffic storm control on the Cisco Nexus 5000 Series switch is implemented in the hardware. The traffic storm control circuitry monitors packets that pass from an Ethernet interface to the switching bus. Using the Individual/Group bit in the packet destination address, the circuitry determines if the packet is unicast or broadcast, tracks the current count of packets within the 10-microsecond interval, and filters out subsequent packets when a threshold is reached.

Traffic storm control uses a bandwidth-based method to measure traffic. You set the percentage of total available bandwidth that the controlled traffic can use. Because packets do not arrive at uniform intervals, the 10-microsecond interval can affect the operation of traffic storm control.

The following are examples of how traffic storm control operation is affected:

- If you enable broadcast traffic storm control, and broadcast traffic exceeds the level within the 10-microsecond interval, traffic storm control drops all broadcast traffic until the end of the interval.
- If you enable multicast traffic storm control, and the multicast traffic exceeds the level within the 10-microsecond interval, traffic storm control drops all multicast traffic until the end of the interval.
- If you enable broadcast and multicast traffic storm control, and broadcast traffic exceeds the level within the 10-microsecond interval, traffic storm control drops all broadcast traffic until the end of the interval.
- If you enable broadcast and multicast traffic storm control, and multicast traffic exceeds the level within the 10-microsecond interval, traffic storm control drops all multicast traffic until the end of the interval.

By default, Cisco NX-OS takes no corrective action when the traffic exceeds the configured level.

Traffic Storm Guidelines and Limitations

When configuring the traffic storm control level, follow these guidelines and limitations:

- You can configure traffic storm control on a port-channel interface.
- Specify the level as a percentage of the total interface bandwidth:
 - The level can be from 0 to 100.
 - The optional fraction of a level can be from 0 to 99.
 - 100 percent means no traffic storm control.
 - 0.0 percent suppresses all traffic.
- In the Cisco Nexus 5000 switch, storm-control does not distinguish between IP, non-IP, registered, or unregistered multicast traffic. All multicast traffic is subject to a single-multicast storm control policer when configured.
- In the Cisco Nexus 5500 switch, storm-control is applied only to unregistered or unknown multicast MAC address.
- The link-level control protocols (LACP, LLDP and so on) are not affected in case of a traffic storm. The storm control is applied to data plane traffic only.
- The burst size values are:
 - For a 10G port, 48.68 Mbytes/390Mbits
 - For a 1G port, 25 Mbytes/200Mbits

Because of hardware limitations and the method by which packets of different sizes are counted, the level percentage is an approximation. Depending on the sizes of the frames that make up the incoming traffic, the actual enforced level might differ from the configured level by several percentage points.

Configuring Traffic Storm Control

You can set the percentage of total available bandwidth that the controlled traffic can use.



Note Traffic storm control uses a 10-microsecond interval that can affect the operation of traffic storm control.

SUMMARY STEPS

1. switch# **configure terminal**
2. switch(config)# **interface** {*ethernet slot/port* | **port-channel number**}
3. switch(config-if)# **storm-control** {**broadcast** | **multicast** | **unicast**} **level** *percentage*[*fraction*]

DETAILED STEPS

	Command or Action	Purpose
Step 1	switch# configure terminal	Enters configuration mode.
Step 2	switch(config)# interface { <i>ethernet slot/port</i> port-channel number }	Enters interface configuration mode.
Step 3	switch(config-if)# storm-control { broadcast multicast unicast } level <i>percentage</i> [<i>fraction</i>]	Configures traffic storm control for traffic on the interface. The default state is disabled.

This example shows how to configure unicast traffic storm control for Ethernet interface 1/4:

```
switch# configure terminal
switch(config)# interface ethernet 1/4
switch(config-if)# storm-control unicast level 40
```

Verifying Traffic Storm Control Configuration

To display traffic storm control configuration information, perform one of these tasks:

Command	Purpose
switch# show interface [<i>ethernet slot/port</i> port-channel number] counters storm-control	Displays the traffic storm control configuration for the interfaces. Note Traffic storm control uses a 10-microsecond interval that can affect the operation of traffic storm control.
switch# show running-config interface	Displays the traffic storm control configuration.

Traffic Storm Control Example Configuration

The following example shows how to configure traffic storm control:

```
switch# configure terminal
switch(config)# interface ethernet 1/4
switch(config-if)# storm-control broadcast level 40
switch(config-if)# storm-control multicast level 40
switch(config-if)# storm-control unicast level 40
```

Default Traffic Storm Settings

The following table lists the default settings for traffic storm control parameters.

Table 20: Default Traffic Storm Control Parameters

Parameters	Default
Traffic storm control	Disabled
Threshold percentage	100



PART

Switch Security Features

- [Configuring Authentication, Authorization, and Accounting, page 243](#)
- [Configuring RADIUS, page 259](#)
- [Configuring TACACS+, page 277](#)
- [Configuring SSH and Telnet, page 295](#)
- [Configuring Access Control Lists, page 307](#)



Configuring Authentication, Authorization, and Accounting

This chapter describes how to configure authentication, authorization, and accounting (AAA) on Cisco Nexus 5000 Series switches. It contains the following sections:

- [Information About AAA, page 243](#)
- [Prerequisites for Remote AAA, page 247](#)
- [Information about AAA Guidelines and Limitations, page 248](#)
- [Configuring AAA, page 248](#)
- [Displaying and Clearing the Local AAA Accounting Log , page 255](#)
- [Verifying AAA Configuration, page 256](#)
- [Example AAA Configuration, page 256](#)
- [Default AAA Settings, page 256](#)

Information About AAA

AAA Security Services

The authentication, authorization, and accounting (AAA) features allows you to verify the identity of, grant access to, and track the actions of users managing Cisco Nexus 5000 Series switches. The Cisco Nexus 5000 Series switches support Remote Access Dial-In User Service (RADIUS) or Terminal Access Controller Access Control device Plus (TACACS+) protocols.

Based on the user ID and password combination that you provide, the Cisco Nexus 5000 Series switches perform local authentication or authorization using the local database or remote authentication or authorization using one or more AAA servers. A preshared secret key provides security for communication between the switch and AAA servers. You can configure a common secret key for all AAA servers or for only a specific AAA server.

AAA security provides the following services:

- **Authentication**—Identifies users, including login and password dialog, challenge and response, messaging support, and, encryption depending on the security protocol that you select.

Authentication is the process of verifying the identity of the person or device accessing the Cisco Nexus 5000 Series switches. This process is based on the user ID and password combination provided by the entity trying to access the switch. The Cisco Nexus 5000 Series switches allow you to perform local authentication (using the local lookup database) or remote authentication (using one or more RADIUS or TACACS+ servers).

- **Authorization**—Provides access control.

AAA authorization is the process of assembling a set of attributes that describe what the user is authorized to perform. Authorization in Cisco Nexus 5000 Series switches is provided by attributes that are downloaded from AAA servers. Remote security servers, such as RADIUS and TACACS+, authorize users for specific rights by associating attribute-value (AV) pairs, which define those rights with the appropriate user.

- **Accounting**—Provides the method for collecting information, logging the information locally, and sending the information to the AAA server for billing, auditing, and reporting.

The accounting feature tracks and maintains a log of every management session used to access the Cisco Nexus 5000 Series switches. You can use this information to generate reports for troubleshooting and auditing purposes. You can store accounting logs locally or send them to remote AAA servers.

**Note**

The Cisco NX-OS software supports authentication, authorization, and accounting independently. For example, you can configure authentication and authorization without configuring accounting.

Benefits of Using AAA

AAA provides the following benefits:

- Increased flexibility and control of access configuration
- Scalability
- Standardized authentication methods, such as RADIUS and TACACS+
- Multiple backup devices

Remote AAA Services

Remote AAA services provided through RADIUS and TACACS+ protocols have the following advantages over local AAA services:

- User password lists for each Cisco Nexus 5000 Series switch in the fabric are easier to manage.
- AAA servers are already deployed widely across enterprises and can be easily used for AAA services.
- The accounting log for all switches in the fabric can be centrally managed.
- User attributes for each switch in the fabric than using the local databases on the switches are easier to manage.

AAA Server Groups

You can specify remote AAA servers for authentication, authorization, and accounting using server groups. A server group is a set of remote AAA servers that implement the same AAA protocol. The purpose of a server group is to provide for failover servers in case a remote AAA server fails to respond. If the first remote server in the group fails to respond, the next remote server in the group is tried until one of the servers sends a response. If all the AAA servers in the server group fail to respond, then that server group option is considered a failure. If required, you can specify multiple server groups. If a Cisco Nexus 5000 Series switch encounters errors from the servers in the first group, it tries the servers in the next server group.

AAA Service Configuration Options

On Cisco Nexus 5000 Series switches, you can have separate AAA configurations for the following services:

- User Telnet or Secure Shell (SSH) login authentication
- Console login authentication
- User management session accounting

The following table lists the CLI commands for each AAA service configuration option.

Table 21: AAA Service Configuration Commands

AAA Service Configuration Option	Related Command
Telnet or SSH login	aaa authentication login default
Console login	aaa authentication login console
User session accounting	aaa accounting default

You can specify the following authentication methods for the AAA services:

- RADIUS server groups—Uses the global pool of RADIUS servers for authentication.
- Specified server groups—Uses specified RADIUS or TACACS+ server groups for authentication.
- Local—Uses the local username or password database for authentication.
- None—Uses only the user name.



Note

If the method is for all RADIUS servers, instead of a specific server group, the Nexus 5000 Series switches choose the RADIUS server from the global pool of configured RADIUS servers in the order of configuration. Servers from this global pool are the servers that can be selectively configured in a RADIUS server group on the Nexus 5000 Series switches.

The following table describes the AAA authentication methods that you can configure for the AAA services.

Table 22: AAA Authentication Methods for AAA Services

AAA Service	AAA Methods
Console login authentication	Server groups, local, and none
User login authentication	Server groups, local, and none
User management session accounting	Server groups and local

**Note**

For console login authentication, user login authentication, and user management session accounting, the Cisco Nexus 5000 Series switches try each option in the order specified. The local option is the default method when other configured options fail.

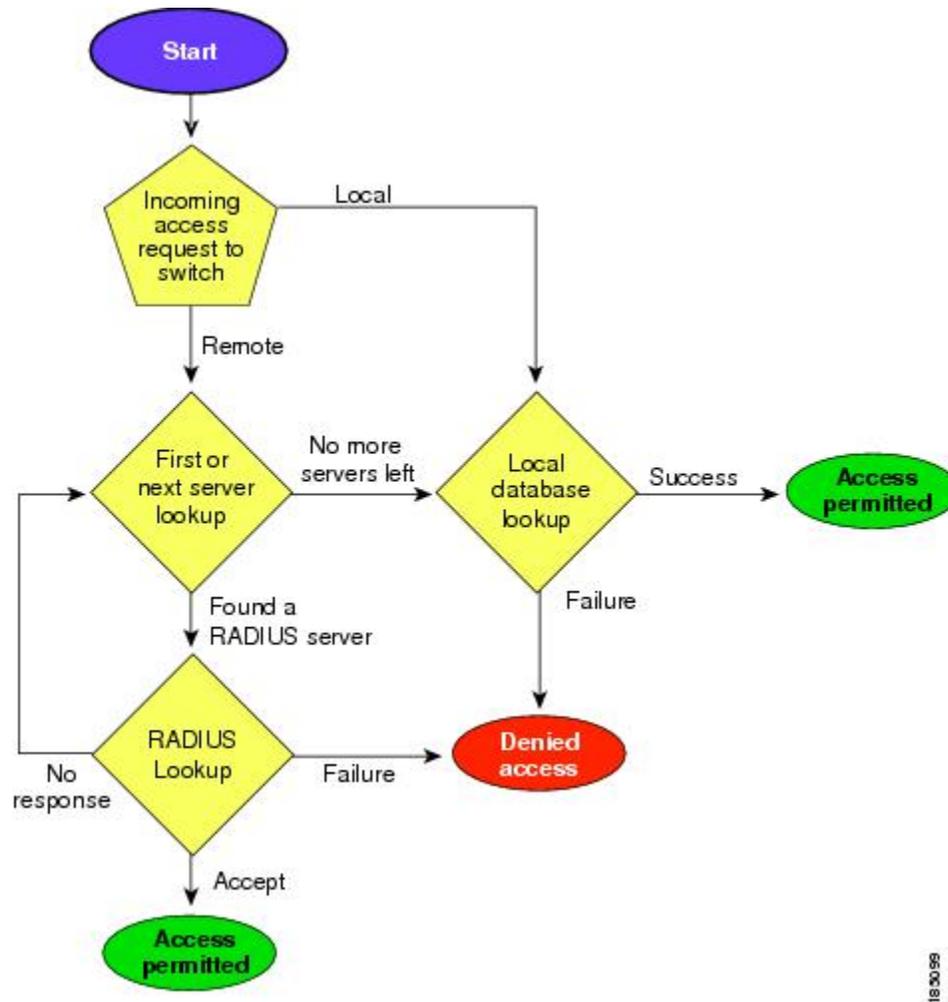
Authentication and Authorization Process for User Login

The figure below shows a flowchart of the authentication and authorization process for user login. The following process occurs:

- When you log in to the required Cisco Nexus 5000 Series switch, you can use the Telnet, SSH, Fabric Manager or Device Manager, or console login options.
- When you have configured the AAA server groups using the server group authentication method, the Cisco Nexus 5000 Series switch sends an authentication request to the first AAA server in the group as follows:
 - If the AAA server fails to respond, then the next AAA server is tried and so on until the remote server responds to the authentication request.
 - If all AAA servers in the server group fail to respond, then the servers in the next server group are tried.
 - If all configured methods fail, then the local database is used for authentication.
- If the Cisco Nexus 5000 Series switches successfully authenticate you through a remote AAA server, then the following possibilities apply:
 - If the AAA server protocol is RADIUS, then user roles specified in the cisco-av-pair attribute are downloaded with an authentication response.
 - If the AAA server protocol is TACACS+, then another request is sent to the same server to get the user roles specified as custom attributes for the shell.

- If your username and password are successfully authenticated locally, the Cisco Nexus 5000 Series switch logs you in and assigns you the roles configured in the local database.

Figure 31: Authorization and Authentication Flow for User Login



Note

"No more server groups left" means that there is no response from any server in all server groups.

"No more servers left" means that there is no response from any server within this server group.

Prerequisites for Remote AAA

Remote AAA servers have the following prerequisites:

- At least one RADIUS or TACACS+ server must be IP reachable.
- The Cisco Nexus 5000 Series switch is configured as a client of the AAA servers.

- The preshared secret key is configured on the Cisco Nexus 5000 Series switch and on the remote AAA servers.
- The remote server responds to AAA requests from the Cisco Nexus 5000 Series switch.

Related Topics

[Configuring RADIUS Server Hosts, on page 263](#)

[Configuring TACACS+ Server Hosts, on page 281](#)

[Manually Monitoring RADIUS Servers or Groups, on page 273](#)

[Manually Monitoring TACACS+ Servers or Groups, on page 291](#)

Information about AAA Guidelines and Limitations

The Cisco Nexus 5000 Series switches do not support all numeric usernames, whether created with TACACS+ or RADIUS, or created locally. If an all numeric username exists on an AAA server and is entered during login, the Cisco Nexus 5000 Series switch will still log in the user.



Caution

You should not create user accounts with usernames that are all numeric.

Configuring AAA

Configuring Console Login Authentication Methods

This section describes how to configure the authentication methods for the console login.

The authentication methods include the following:

- Global pool of RADIUS servers
- Named subset of RADIUS or TACACS+ servers
- Local database on the Nexus 5000 Series switch
- Username only (**none**)

The default method is local.



Note

The **group radius** and **group *server-name*** forms of the **aaa authentication** command are used for a set of previously defined RADIUS servers. Use the **radius server-host** command to configure the host servers. Use the **aaa group server radius** command to create a named group of servers.

Before you configure console login authentication methods, configure RADIUS or TACACS+ server groups as needed. To configure console login authentication methods, perform this task:

SUMMARY STEPS

1. switch# **configure terminal**
2. switch(config)# **aaa authentication login console {group group-list [none] | local | none}**
3. switch(config)# **exit**
4. (Optional) switch# **show aaa authentication**
5. (Optional) switch# **copy running-config startup-config**

DETAILED STEPS

	Command or Action	Purpose
Step 1	switch# configure terminal	Enters configuration mode.
Step 2	switch(config)# aaa authentication login console {group group-list [none] local none}	<p>Configures login authentication methods for the console.</p> <p>The <i>group-list</i> argument consists of a space-delimited list of group names. The group names are the following:</p> <ul style="list-style-type: none"> • radius —Uses the global pool of RADIUS servers for authentication. • named-group —Uses a named subset of TACACS+ or RADIUS servers for authentication. <p>The local method uses the local database for authentication. The none method uses the username only.</p> <p>The default console login method is local, which is used when no methods are configured or when all of the configured methods fail to respond.</p>
Step 3	switch(config)# exit	Exits configuration mode.
Step 4	switch# show aaa authentication	(Optional) Displays the configuration of the console login authentication methods.
Step 5	switch# copy running-config startup-config	(Optional) Copies the running configuration to the startup configuration.

The following example shows how to configure authentication methods for the console login:

```
switch# configure terminal
switch(config)# aaa authentication login console group radius
switch(config)# exit
switch# show aaa authentication
switch# copy running-config startup-config
```

Configuring Default Login Authentication Methods

The authentication methods include the following:

- Global pool of RADIUS servers

- Named subset of RADIUS or TACACS+ servers
- Local database on the Nexus 5000 Series switch
- Username only

The default method is local.

Before you configure default login authentication methods, configure RADIUS or TACACS+ server groups as needed. To configure default login authentication methods, perform this task:

SUMMARY STEPS

1. switch# **configure terminal**
2. switch(config)# **aaa authentication login default {group *group-list* [none] | local | none}**
3. switch(config)# **exit**
4. (Optional) switch# **show aaa authentication**
5. (Optional) switch# **copy running-config startup-config**

DETAILED STEPS

	Command or Action	Purpose
Step 1	switch# configure terminal	Enters configuration mode.
Step 2	switch(config)# aaa authentication login default {group <i>group-list</i> [none] local none}	<p>Configures the default authentication methods.</p> <p>The <i>group-list</i> argument consists of a space-delimited list of group names. The group names are the following:</p> <ul style="list-style-type: none"> • radius —Uses the global pool of RADIUS servers for authentication. • named-group —Uses a named subset of TACACS+ or RADIUS servers for authentication. <p>The local method uses the local database for authentication. The none method uses the username only.</p> <p>The default login method is local, which is used when no methods are configured or when all of the configured methods do not respond.</p>
Step 3	switch(config)# exit	Exits configuration mode.
Step 4	switch# show aaa authentication	(Optional) Displays the configuration of the default login authentication methods.
Step 5	switch# copy running-config startup-config	(Optional) Copies the running configuration to the startup configuration.

Enabling Login Authentication Failure Messages

When you log in, the login is processed by the local user database if the remote AAA servers do not respond. If you have enabled the displaying of login failure messages, the following message is displayed :

```
Remote AAA servers unreachable; local authentication done.
Remote AAA servers unreachable; local authentication failed.
```

To enable login authentication failure messages, perform this task:

SUMMARY STEPS

1. switch# **configure terminal**
2. switch(config)# **aaa authentication login error-enable**
3. switch(config)# **exit**
4. (Optional) switch# **show aaa authentication**
5. (Optional) switch# **copy running-config startup-config**

DETAILED STEPS

	Command or Action	Purpose
Step 1	switch# configure terminal	Enters configuration mode.
Step 2	switch(config)# aaa authentication login error-enable	Enables login authentication failure messages. The default is disabled.
Step 3	switch(config)# exit	Exits configuration mode.
Step 4	switch# show aaa authentication	(Optional) Displays the login failure message configuration.
Step 5	switch# copy running-config startup-config	(Optional) Copies the running configuration to the startup configuration.

Enabling MSCHAP Authentication

Microsoft Challenge Handshake Authentication Protocol (MSCHAP) is the Microsoft version of CHAP. You can use MSCHAP for user logins to a Cisco Nexus 5000 Series switch through a remote authentication server (RADIUS or TACACS+).

By default, the Cisco Nexus 5000 Series switch uses Password Authentication Protocol (PAP) authentication between the switch and the remote server. If you enable MSCHAP, you need to configure your RADIUS server to recognize the MSCHAP vendor-specific attributes (VSAs).

The following table describes the RADIUS VSAs required for MSCHAP.

Table 23: MSCHAP RADIUS VSAs

Vendor-ID Number	Vendor-Type Number	VSA	Description
311	11	MSCHAP-Challenge	Contains the challenge sent by an AAA server to an MSCHAP user. It can be used in both Access-Request and Access-Challenge packets.
211	11	MSCHAP-Response	Contains the response value provided by an MSCHAP user in response to the challenge. It is only used in Access-Request packets.

To enable MSCHAP authentication, perform this task:

SUMMARY STEPS

1. switch# **configure terminal**
2. switch(config)# **aaa authentication login mschap enable**
3. switch(config)# **exit**
4. (Optional) switch# **show aaa authentication login mschap**
5. (Optional) switch# **copy running-config startup-config**

DETAILED STEPS

	Command or Action	Purpose
Step 1	switch# configure terminal	Enters configuration mode.
Step 2	switch(config)# aaa authentication login mschap enable	Enables MS-CHAP authentication. The default is disabled.
Step 3	switch(config)# exit	Exits configuration mode.
Step 4	switch# show aaa authentication login mschap	(Optional) Displays the MS-CHAP configuration.
Step 5	switch# copy running-config startup-config	(Optional) Copies the running configuration to the startup configuration.

Related Topics

[About VSAs, on page 254](#)

Configuring AAA Accounting Default Methods

The Cisco Nexus 5000 Series switch supports TACACS+ and RADIUS methods for accounting. The switches report user activity to TACACS+ or RADIUS security servers in the form of accounting records. Each accounting record contains accounting attribute-value (AV) pairs and is stored on the AAA server.

When you activate AAA accounting, the Cisco Nexus 5000 Series switch reports these attributes as accounting records, which are then stored in an accounting log on the security server.

You can create default method lists defining specific accounting methods, which include the following:

- RADIUS server group—Uses the global pool of RADIUS servers for accounting.
- Specified server group—Uses a specified RADIUS or TACACS+ server group for accounting.
- Local—Uses the local username or password database for accounting.



Note

If you have configured server groups and the server groups do not respond, by default the local database is used for authentication.

Before you configure AAA accounting default methods, configure RADIUS or TACACS+ server groups as needed.

To configure AAA accounting default methods, perform this task:

SUMMARY STEPS

1. switch# **configure terminal**
2. switch(config)# **aaa accounting default {group group-list | local}**
3. switch(config)# **exit**
4. (Optional) switch# **show aaa accounting**
5. (Optional) switch# **copy running-config startup-config**

DETAILED STEPS

	Command or Action	Purpose
Step 1	switch# configure terminal	Enters configuration mode.
Step 2	switch(config)# aaa accounting default {group group-list local}	Configures default accounting method. One or more server group names can be specified in a space separated list. The <i>group-list</i> argument consists of a space-delimited list of group names. The group names are of the following: <ul style="list-style-type: none"> • radius —Uses the global pool of RADIUS servers for accounting.

	Command or Action	Purpose
		<ul style="list-style-type: none"> • <i>named-group</i>—Uses a named subset of TACACS+ or RADIUS servers for accounting. <p>The local method uses the local database for accounting.</p> <p>The default method is local , which is used when no server groups are configured or when all the configured server group do not respond.</p>
Step 3	switch(config)# exit	Exits configuration mode.
Step 4	switch# show aaa accounting	(Optional) Displays the configuration AAA accounting default methods.
Step 5	switch# copy running-config startup-config	(Optional) Copies the running configuration to the startup configuration.

Using AAA Server VSAs

About VSAs

You can use vendor-specific attributes (VSAs) to specify the Cisco Nexus 5000 Series user roles and SNMPv3 parameters on AAA servers.

The Internet Engineering Task Force (IETF) draft standard specifies a method for communicating VSAs between the network access server and the RADIUS server. The IETF uses attribute 26. VSAs allow vendors to support their own extended attributes that are not suitable for general use. The Cisco RADIUS implementation supports one vendor-specific option using the format recommended in the specification. The Cisco vendor ID is 9, and the supported option is vendor type 1, which is named `cisco-av-pair`. The value is a string with the following format:

```
protocol : attribute separator value *
```

The protocol is a Cisco attribute for a particular type of authorization, separator is an equal sign (=) for mandatory attributes, and an asterisk (*) indicates optional attributes.

When you use RADIUS servers for authentication on a Nexus 5000 Series switch, the RADIUS protocol directs the RADIUS server to return user attributes, such as authorization information, along with authentication results. This authorization information is specified through VSAs.

VSA Format

The following VSA protocol options are supported by the Cisco Nexus 5000 Series switches:

- Shell—Used in access-accept packets to provide user profile information.
- Accounting—Used in accounting-request packets. If a value contains any white spaces, put it within double quotation marks.

The following attributes are supported by the Cisco Nexus 5000 Series switches:

- **roles**—Lists all the roles assigned to the user. The value field is a string that stores the list of group names delimited by white space.
- **accountinginfo**—Stores additional accounting information in addition to the attributes covered by a standard RADIUS accounting protocol. This attribute is sent only in the VSA portion of the Account-Request frames from the RADIUS client on the switch, and it can only be used with the accounting protocol-related PDUs.

Specifying Switch User Roles and SNMPv3 Parameters on AAA Servers

You can use the VSA `cisco-av-pair` on AAA servers to specify user role mapping for the Cisco Nexus 5000 Series switch using this format:

```
shell:roles="roleA roleB ..."
```

If you do not specify the `role` option in the `cisco-av-pair` attribute, the default user role is `network-operator`.

You can also specify your SNMPv3 authentication and privacy protocol attributes as follows:

```
shell:roles="roleA roleB..." snmpv3:auth=SHA priv=AES-128
```

The SNMPv3 authentication protocol options are SHA and MD5. The privacy protocol options are AES-128 and DES. If you do not specify these options in the `cisco-av-pair` attribute, MD5 and DES are the default authentication protocols.

For additional information, see the Configuring User Accounts and RBAC chapter in the *Cisco Nexus 5000 Series NX-OS System Management Configuration Guide*.

Displaying and Clearing the Local AAA Accounting Log

The Cisco Nexus 5000 Series switch maintains a local log for the AAA accounting activity. To display this log and clear it, perform this task:

SUMMARY STEPS

1. `switch# show accounting log [size] [start-time year month day hh : mm : ss]`
2. (Optional) `switch# clear accounting log`

DETAILED STEPS

	Command or Action	Purpose
Step 1	<code>switch# show accounting log [size] [start-time year month day hh : mm : ss]</code>	Displays the accounting log contents. By default, the command output contains up to 250,000 bytes of the accounting log. You can use the <code>size</code> argument to limit command output. The range is from 0 to 250000 bytes. You can also specify a start time for the log output.
Step 2	<code>switch# clear accounting log</code>	(Optional) Clears the accounting log contents.

Verifying AAA Configuration

To display AAA configuration information, perform one of the following tasks:

SUMMARY STEPS

1. `show aaa accounting`
2. `show aaa authentication [login {error-enable | mschap}]`
3. `show aaa groups`
4. `show running-config aaa [all]`
5. `show startup-config aaa`

DETAILED STEPS

	Command or Action	Purpose
Step 1	<code>show aaa accounting</code>	Displays AAA accounting configuration.
Step 2	<code>show aaa authentication [login {error-enable mschap}]</code>	Displays AAA authentication information.
Step 3	<code>show aaa groups</code>	Displays the AAA server group configuration.
Step 4	<code>show running-config aaa [all]</code>	Displays the AAA configuration in the running configuration.
Step 5	<code>show startup-config aaa</code>	Displays the AAA configuration in the startup configuration.

Example AAA Configuration

The following example shows how to configure AAA:

```
switch(config)# aaa authentication login default group radius
switch(config)# aaa authentication login console group radius
switch(config)# aaa accounting default group radius
```

Default AAA Settings

The following table lists the default settings for AAA parameters.

Table 24: Default AAA Parameters

Parameters	Default
Console authentication method	local
Default authentication method	local

Parameters	Default
Login authentication failure messages	Disabled
MSCHAP authentication	Disabled
Default accounting method	local
Accounting log display length	250 KB



Configuring RADIUS

This chapter contains the following sections:

- [Configuring RADIUS, page 259](#)

Configuring RADIUS

Information About RADIUS

The Remote Access Dial-In User Service (RADIUS) distributed client/server system allows you to secure networks against unauthorized access. In the Cisco implementation, RADIUS clients run on Cisco Nexus 5000 Series switches and send authentication and accounting requests to a central RADIUS server that contains all user authentication and network service access information.

RADIUS Network Environments

RADIUS can be implemented in a variety of network environments that require high levels of security while maintaining network access for remote users.

You can use RADIUS in the following network environments that require access security:

- Networks with multiple-vendor network devices, each supporting RADIUS.

For example, network devices from several vendors can use a single RADIUS server-based security database.

- Networks already using RADIUS.

You can add a Nexus 5000 Series switch with RADIUS to the network. This action might be the first step when you make a transition to a AAA server.

- Networks that require resource accounting.

You can use RADIUS accounting independent of RADIUS authentication or authorization. The RADIUS accounting functions allow data to be sent at the start and end of services, indicating the amount of resources (such as time, packets, bytes, and so on) used during the session. An Internet service provider (ISP) might use a freeware-based version of the RADIUS access control and accounting software to meet special security and billing needs.

- Networks that support authentication profiles.

Using the RADIUS server in your network, you can configure AAA authentication and set up per-user profiles. Per-user profiles enable the Nexus 5000 Series switch to better manage ports using their existing RADIUS solutions and to efficiently manage shared resources to offer different service-level agreements.

RADIUS Operation

When a user attempts to log in and authenticate to a Cisco Nexus 5000 Series switch using RADIUS, the following process occurs:

- 1 The user is prompted for and enters a username and password.
- 2 The username and encrypted password are sent over the network to the RADIUS server.
- 3 The user receives one of the following responses from the RADIUS server:
 - ACCEPT—The user is authenticated.
 - REJECT—The user is not authenticated and is prompted to reenter the username and password, or access is denied.
 - CHALLENGE—A challenge is issued by the RADIUS server. The challenge collects additional data from the user.
 - CHANGE PASSWORD—A request is issued by the RADIUS server, asking the user to select a new password.

The ACCEPT or REJECT response is bundled with additional data that is used for EXEC or network authorization. You must first complete RADIUS authentication before using RADIUS authorization. The additional data included with the ACCEPT or REJECT packets consists of the following:

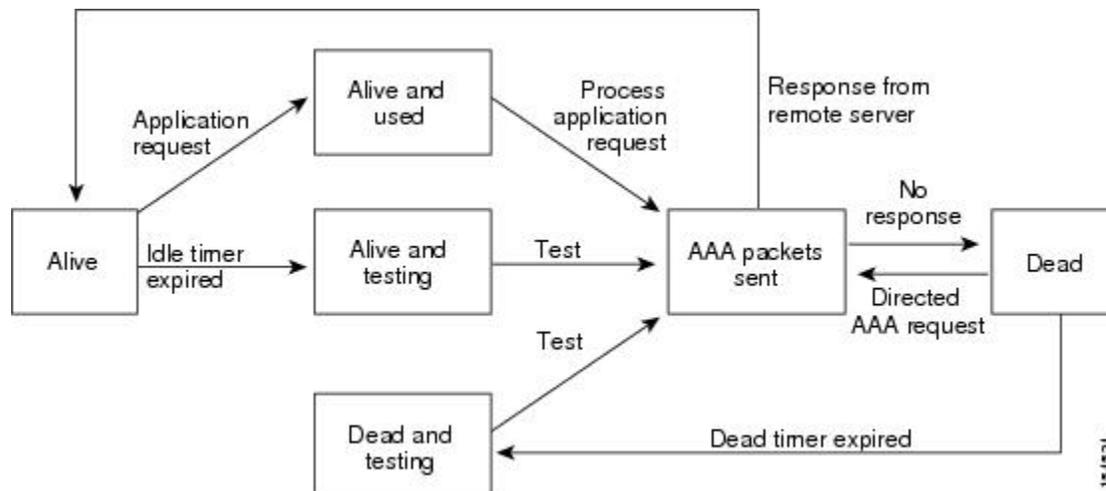
- Services that the user can access, including Telnet, rlogin, or local-area transport (LAT) connections, and Point-to-Point Protocol (PPP), Serial Line Internet Protocol (SLIP), or EXEC services.
- Connection parameters, including the host or client IPv4 or IPv6 address, access list, and user timeouts.

RADIUS Server Monitoring

An unresponsive RADIUS server can cause delay in processing of AAA requests. You can configure the Cisco Nexus 5000 Series switch to periodically monitor a RADIUS server to check whether it is responding (or alive) to save time in processing AAA requests. The Cisco Nexus 5000 Series switch marks unresponsive RADIUS servers as dead and does not send AAA requests to any dead RADIUS servers. The switch periodically monitors the dead RADIUS servers and brings them to the alive state once they respond. This monitoring process verifies that a RADIUS server is in a working state before real AAA requests are sent its way. Whenever

a RADIUS server changes to the dead or alive state, a Simple Network Management Protocol (SNMP) trap is generated and the Cisco Nexus 5000 Series switch displays an error message that a failure is taking place.

Figure 32: RADIUS Server States



Note

The monitoring interval for alive servers and dead servers are different and can be configured by the user. The RADIUS server monitoring is performed by sending a test authentication request to the RADIUS server.

Vendor-Specific Attributes

The Internet Engineering Task Force (IETF) draft standard specifies a method for communicating vendor-specific attributes (VSAs) between the network access server and the RADIUS server. The IETF uses attribute 26. VSAs allow vendors to support their own extended attributes that are not suitable for general use. The Cisco RADIUS implementation supports one vendor-specific option using the format recommended in the specification. The Cisco vendor ID is 9, and the supported option is vendor type 1, which is named cisco-av-pair. The value is a string with the following format:

```
protocol : attribute separator value *
```

The protocol is a Cisco attribute for a particular type of authorization, the separator is an equal sign (=) for mandatory attributes, and an asterisk (*) indicates optional attributes.

When you use RADIUS servers for authentication on a Cisco Nexus 5000 Series switch, the RADIUS protocol directs the RADIUS server to return user attributes, such as authorization information, along with authentication results. This authorization information is specified through VSAs.

The following VSA protocol options are supported by the Cisco Nexus 5000 Series switch:

- Shell— Used in access-accept packets to provide user profile information.
- Accounting— Used in accounting-request packets. If a value contains any white spaces, you should enclose the value within double quotation marks.

The Nexus 5000 Series switch supports the following attributes:

- **roles**—Lists all the roles to which the user belongs. The value field is a string that lists the role names delimited by white space.
- **accountinginfo**—Stores accounting information in addition to the attributes covered by a standard RADIUS accounting protocol. This attribute is sent only in the VSA portion of the Account-Request frames from the RADIUS client on the switch. It can be used only with the accounting protocol data units (PDUs).

Prerequisites for RADIUS

RADIUS has the following prerequisites:

- Obtain IPv4 or IPv6 addresses or host names for the RADIUS servers.
- Obtain preshared keys from the RADIUS servers.
- Ensure that the Cisco Nexus 5000 Series switch is configured as a RADIUS client of the AAA servers.

Guidelines and Limitations for RADIUS

RADIUS has the following guidelines and limitations:

- You can configure a maximum of 64 RADIUS servers on the Cisco Nexus 5000 Series switch.

Configuring RADIUS Servers

To configure RADIUS servers, perform this task:

SUMMARY STEPS

1. Establish the RADIUS server connections to the Cisco Nexus 5000 Series switch.
2. Configure the preshared secret keys for the RADIUS servers.
3. If needed, configure RADIUS server groups with subsets of the RADIUS servers for AAA authentication methods.
4. If needed, configure any of the following optional parameters:
5. If needed, configure periodic RADIUS server monitoring.

DETAILED STEPS

-
- Step 1** Establish the RADIUS server connections to the Cisco Nexus 5000 Series switch.
- Step 2** Configure the preshared secret keys for the RADIUS servers.
- Step 3** If needed, configure RADIUS server groups with subsets of the RADIUS servers for AAA authentication methods.
- Step 4** If needed, configure any of the following optional parameters:
- Dead-time interval.
 - Allow specification of a RADIUS server at login.
 - Transmission retry count and timeout interval.

- Accounting and authentication attributes.

Step 5 If needed, configure periodic RADIUS server monitoring.

Configuring RADIUS Server Hosts

You must configure the IPv4 or IPv6 address or the host name for each RADIUS server that you want to use for authentication. All RADIUS server hosts are added to the default RADIUS server group. You can configure up to 64 RADIUS servers.

To configure a RADIUS server host, perform this task:

SUMMARY STEPS

1. switch# **configure terminal**
2. switch(config) #**radius-server host** {*ipv4-address* | *ipv6-address* | *host-name*}
3. switch(config)# **exit**
4. (Optional) switch# **show radius-server**
5. (Optional) switch# **copy running-config startup-config**

DETAILED STEPS

	Command or Action	Purpose
Step 1	switch# configure terminal	Enters configuration mode.
Step 2	switch(config) # radius-server host { <i>ipv4-address</i> <i>ipv6-address</i> <i>host-name</i> }	Specifies the IPv4 or IPv6 address or hostname for a RADIUS server.
Step 3	switch(config)# exit	Exits configuration mode.
Step 4	switch# show radius-server	(Optional) Displays the RADIUS server configuration.
Step 5	switch# copy running-config startup-config	(Optional) Copies the running configuration to the startup configuration.

The following example shows how to configure a RADIUS server host:

```
switch# configure terminal
switch(config)# radius-server host 10.10.1.1
switch(config)# exit
switch# show radius-server
switch# copy running-config startup-config
```

Configuring RADIUS Global Preshared Keys

You can configure preshared keys at the global level for all servers used by the Cisco Nexus 5000 Series switch. A preshared key is a shared secret text string between the switch and the RADIUS server hosts.

To configure global preshared keys, obtain the preshared key values for the remote RADIUS servers and perform this task:

SUMMARY STEPS

1. switch# **configure terminal**
2. switch(config)# **radius-server key [0 | 7] key-value**
3. switch(config)# **exit**
4. (Optional) switch# **show radius-server**
5. (Optional) switch# **copy running-config startup-config**

DETAILED STEPS

	Command or Action	Purpose
Step 1	switch# configure terminal	Enters configuration mode.
Step 2	switch(config)# radius-server key [0 7] key-value	Specifies a preshared key for all RADIUS servers. You can specify a clear text (0) or encrypted (7) preshared key. The default format is clear text. The maximum length is 63 characters. By default, no preshared key is configured.
Step 3	switch(config)# exit	Exits configuration mode.
Step 4	switch# show radius-server	(Optional) Displays the RADIUS server configuration. Note The preshared keys are saved in encrypted form in the running configuration. Use the show running-config command to display the encrypted preshared keys.
Step 5	switch# copy running-config startup-config	(Optional) Copies the running configuration to the startup configuration.

The following example shows how to configure the preshared key values for a remote RADIUS server:

```
switch# configure terminal
switch(config)# radius-server key 0 QsEfThUkO
switch(config)# exit
switch# show radius-server
switch# copy running-config startup-config
```

Configuring RADIUS Server Preshared Keys

You can configure preshared keys for a RADIUS server. A preshared key is a shared secret text string between the Cisco Nexus 5000 Series switch and the RADIUS server host.

To configure radius server preshared keys, obtain the preshared key values for the remote RADIUS servers and perform this task:

SUMMARY STEPS

1. switch# **configure terminal**
2. switch(config)# **radius-server host** {*ipv4-address* | *ipv6-address* | *host-name*} **key** [0 | 7] *key-value*
3. switch(config)# **exit**
4. (Optional) switch# **show radius-server**
5. (Optional) switch# **copy running-config startup-config**

DETAILED STEPS

	Command or Action	Purpose
Step 1	switch# configure terminal	Enters configuration mode.
Step 2	switch(config)# radius-server host { <i>ipv4-address</i> <i>ipv6-address</i> <i>host-name</i> } key [0 7] <i>key-value</i>	Specifies a preshared key for a specific RADIUS server. You can specify a clear text (0) or encrypted (7) preshared key. The default format is clear text. The maximum length is 63 characters. This preshared key is used instead of the global preshared key.
Step 3	switch(config)# exit	Exits configuration mode.
Step 4	switch# show radius-server	(Optional) Displays the RADIUS server configuration. Note The preshared keys are saved in encrypted form in the running configuration. Use the show running-config command to display the encrypted preshared keys.
Step 5	switch# copy running-config startup-config	(Optional) Copies the running configuration to the startup configuration.

The following example shows how to configure a preshared keys for a RADIUS server:

```
switch# configure terminal
switch(config)# radius-server host 10.10.1.1 key 0 P1IjUhYg
switch(config)# exit
switch# show radius-server
switch# copy running-config startup-config
```

Configuring RADIUS Server Groups

You can specify one or more remote AAA servers for authentication using server groups. All members of a group must belong to the RADIUS protocol. The servers are tried in the same order in which you configure them.

You can configure these server groups at any time but they only take effect when you apply them to an AAA service.

To configure radius server groups, perform this task:

SUMMARY STEPS

1. switch# **configure terminal**
2. switch(config)# **aaa group server radius** *group-name*
3. switch(config-radius)# **server** {*ipv4-address* | *ipv6-address* | *server-name*}
4. (Optional) switch(config-radius)# **deadtime** *minutes*
5. switch(config-radius)# **exit**
6. (Optional) switch(config) #**show radius-server group** [*group-name*]
7. (Optional) switch(config)# **copy running-config startup-config**

DETAILED STEPS

	Command or Action	Purpose
Step 1	switch# configure terminal	Enters configuration mode.
Step 2	switch(config)# aaa group server radius <i>group-name</i>	Creates a RADIUS server group and enters the RADIUS server group configuration submode for that group. The <i>group-name</i> argument is a case-sensitive alphanumeric string with a maximum length of 127 characters.
Step 3	switch(config-radius)# server { <i>ipv4-address</i> <i>ipv6-address</i> <i>server-name</i> }	Configures the RADIUS server as a member of the RADIUS server group. If the specified RADIUS server is not found, configure it using the radius-server host command and retry this command.
Step 4	switch(config-radius)# deadtime <i>minutes</i>	(Optional) Configures the monitoring dead time. The default is 0 minutes. The range is from 1 through 1440. Note If the dead-time interval for a RADIUS server group is greater than zero (0), that value takes precedence over the global dead-time value. See the example that shows how to configure periodic RADIUS server monitoring.
Step 5	switch(config-radius)# exit	Exits configuration mode.
Step 6	switch(config) # show radius-server group [<i>group-name</i>]	(Optional) Displays the RADIUS server group configuration.
Step 7	switch(config)# copy running-config startup-config	(Optional) Copies the running configuration to the startup configuration.

The following example shows how to configure a RADIUS server group:

```
switch# configure terminal
switch(config)# aaa group server radius RadServer
switch(config-radius)# server 10.10.1.1
switch(config-radius)# deadtime 30
switch(config-radius)# use-vrf management
switch(config-radius)# exit
switch(config)# show radius-server group
switch(config)# copy running-config startup-config
```

Allowing Users to Specify a RADIUS Server at Login

To allow users to specify a RADIUS server at login, perform this task:

SUMMARY STEPS

1. switch# **configure terminal**
2. switch(config)# **radius-server directed-request**
3. switch(config)# **exit**
4. (Optional) switch# **show radius-server directed-request**
5. (Optional) switch# **copy running-config startup-config**

DETAILED STEPS

	Command or Action	Purpose
Step 1	switch# configure terminal	Enters configuration mode.
Step 2	switch(config)# radius-server directed-request	Allows users to specify a RADIUS server to send the authentication request when logging in. The default is disabled.
Step 3	switch(config)# exit	Exits configuration mode.
Step 4	switch# show radius-server directed-request	(Optional) Displays the directed request configuration.
Step 5	switch# copy running-config startup-config	(Optional) Copies the running configuration to the startup configuration.

Configuring the Global RADIUS Transmission Retry Count and Timeout Interval

You can configure a global retransmission retry count and timeout interval for all RADIUS servers. By default, a switch retries transmission to a RADIUS server only once before reverting to local authentication. You can increase this number up to a maximum of five retries per server. The timeout interval determines how long the Cisco Nexus 5000 Series switch waits for responses from RADIUS servers before declaring a timeout failure.

To configure the global RADIUS transmission retry count and timeout interval, perform this task:

SUMMARY STEPS

1. switch# **configure terminal**
2. switch(config)# **radius-server retransmit** *count*
3. switch(config)# **radius-server timeout** *seconds*
4. switch(config)# **exit**
5. (Optional) switch# **show radius-server**
6. (Optional) switch# **copy running-config startup-config**

DETAILED STEPS

	Command or Action	Purpose
Step 1	switch# configure terminal	Enters configuration mode.
Step 2	switch(config)# radius-server retransmit <i>count</i>	Specifies the retransmission count for all RADIUS servers. The default retransmission count is 1 and the range is from 0 to 5.
Step 3	switch(config)# radius-server timeout <i>seconds</i>	Specifies the transmission timeout interval for RADIUS servers. The default timeout interval is 5 seconds and the range is from 1 to 60 seconds.
Step 4	switch(config)# exit	Exits configuration mode.
Step 5	switch# show radius-server	(Optional) Displays the RADIUS server configuration.
Step 6	switch# copy running-config startup-config	(Optional) Copies the running configuration to the startup configuration.

Configuring the RADIUS Transmission Retry Count and Timeout Interval for a Server

By default, a Cisco Nexus 5000 Series switch retries transmission to a RADIUS server only once before reverting to local authentication. You can increase this number up to a maximum of five retries per server. You can also set a timeout interval that the switch waits for responses from RADIUS servers before declaring a timeout failure.

To configure RADIUS transmission retry count and timeout interval for a server, perform this task:

SUMMARY STEPS

1. switch# **configure terminal**
2. #switch(config)# **radius-server host** {*ipv4-address* |*ipv6-address* | *host-name*} **retransmit** *count*
3. switch(config)# switch(config)# **radius-server host** {*ipv4-address* |*ipv6-address* | *host-name*} **timeout** *seconds*
4. switch(config)# **exit**
5. (Optional) switch# **show radius-server**
6. (Optional) switch# **copy running-config startup-config**

DETAILED STEPS

	Command or Action	Purpose
Step 1	switch# configure terminal	Enters configuration mode.
Step 2	#switch(config)# radius-server host { <i>ipv4-address</i> <i>ipv6-address</i> <i>host-name</i> } retransmit <i>count</i>	Specifies the retransmission count for a specific server. The default is the global value. Note The retransmission count value specified for a RADIUS server overrides the count specified for all RADIUS servers.
Step 3	switch(config)# switch(config)# radius-server host { <i>ipv4-address</i> <i>ipv6-address</i> <i>host-name</i> } timeout <i>seconds</i>	Specifies the transmission timeout interval for a specific server. The default is the global value. Note The timeout interval value specified for a RADIUS server overrides the interval value specified for all RADIUS servers.
Step 4	switch(config)# exit	Exits configuration mode.
Step 5	switch# show radius-server	(Optional) Displays the RADIUS server configuration.
Step 6	switch# copy running-config startup-config	(Optional) Copies the running configuration to the startup configuration.

The following example shows how to configure RADIUS transmission retry count and timeout interval for a server:

```
switch# configure terminal
switch(config)# radius-server host server1 retransmit 3
switch(config)# radius-server host server1 timeout 10
switch(config)# exit
switch# show radius-server
switch# copy running-config startup-config
```

Configuring Accounting and Authentication Attributes for RADIUS Servers

You can specify that a RADIUS server is to be used only for accounting purposes or only for authentication purposes. By default, RADIUS servers are used for both accounting and authentication. You can also specify the destination UDP port numbers where RADIUS accounting and authentication messages should be sent.

To configure the accounting and authentication attributes for RADIUS servers, perform this task:

SUMMARY STEPS

1. switch# **configure terminal**
2. (Optional) switch(config) #**radius-server host** {*ipv4-address* | *ipv6-address* | *host-name*} **acct-port** *udp-port*
3. (Optional) switch(config)# **radius-server host** {*ipv4-address* | *ipv6-address* | *host-name*} **accounting**
4. (Optional) switch(config)# **radius-server host** {*ipv4-address* | *ipv6-address* | *host-name*} **auth-port** *udp-port*
5. (Optional) switch(config)# **radius-server host** {*ipv4-address* | *ipv6-address* | *host-name*} **authentication**
6. switch(config)# **exit**
7. (Optional) switch(config)# **show radius-server**
8. (Optional) switch# **copy running-config startup-config**

DETAILED STEPS

	Command or Action	Purpose
Step 1	switch# configure terminal	Enters configuration mode.
Step 2	switch(config) # radius-server host { <i>ipv4-address</i> <i>ipv6-address</i> <i>host-name</i> } acct-port <i>udp-port</i>	(Optional) Specifies a UDP port to use for RADIUS accounting messages. The default UDP port is 1812. The range is from 0 to 65535.
Step 3	switch(config)# radius-server host { <i>ipv4-address</i> <i>ipv6-address</i> <i>host-name</i> } accounting	(Optional) Specifies that the specified RADIUS server it to be used only for accounting purposes. The default is both accounting and authentication.
Step 4	switch(config)# radius-server host { <i>ipv4-address</i> <i>ipv6-address</i> <i>host-name</i> } auth-port <i>udp-port</i>	(Optional) Specifies a UDP port to use for RADIUS authentication messages. The default UDP port is 1812. The range is from 0 to 65535.
Step 5	switch(config)# radius-server host { <i>ipv4-address</i> <i>ipv6-address</i> <i>host-name</i> } authentication	(Optional) Specifies that the specified RADIUS server only be used for authentication purposes. The default is both accounting and authentication.
Step 6	switch(config)# exit	Exits configuration mode.

	Command or Action	Purpose
Step 7	switch(config)# show radius-server	(Optional) Displays the RADIUS server configuration.
Step 8	switch# copy running-config startup-config	(Optional) Copies the running configuration to the startup configuration.

The following example shows how to configure the accounting and authentication attributes for a RADIUS server:

```
switch# configure terminal
switch(config)# radius-server host 10.10.1.1 acct-port 2004
switch(config)# radius-server host 10.10.1.1 accounting
switch(config)# radius-server host 10.10.2.2 auth-port 2005
switch(config)# radius-server host 10.10.2.2 authentication
switch(config)# exit
switch# show radius-server
switch# copy running-config startup-config
```

Configuring Periodic RADIUS Server Monitoring

You can monitor the availability of RADIUS servers. These parameters include the username and password to use for the server and an idle timer. The idle timer specifies the interval during which a RADIUS server receives no requests before the Cisco Nexus 5000 Series switch sends out a test packet. You can configure this option to test servers periodically.



Note For security reasons, we recommend that you do not configure a test username that is the same as an existing user in the RADIUS database.

The test idle timer specifies the interval during which a RADIUS server receives no requests before the Cisco Nexus 5000 Series switch sends out a test packet.



Note The default idle timer value is 0 minutes. When the idle time interval is 0 minutes, the Cisco Nexus 5000 Series switch does not perform periodic RADIUS server monitoring.

SUMMARY STEPS

1. switch# **configure terminal**
2. switch(config)# **radius-server host** {ipv4-address | ipv6-address | host-name} **test** {idle-time minutes | password password [idle-time minutes] | username name [password password [idle-time minutes]]}
3. switch(config)# **radius-server deadtime** minutes
4. switch(config)# **exit**
5. (Optional) switch# **show radius-server**
6. (Optional) switch# **copy running-config startup-config**

DETAILED STEPS

	Command or Action	Purpose
Step 1	switch# configure terminal	Enters configuration mode.
Step 2	switch(config)# radius-server host { <i>ipv4-address</i> <i>ipv6-address</i> <i>host-name</i> } test { <i>idle-time minutes</i> password <i>password</i> [<i>idle-time minutes</i>] username <i>name</i> [password <i>password</i> [<i>idle-time minutes</i>]]}	Specifies parameters for server monitoring. The default username is test and the default password is test. The default value for the idle timer is 0 minutes. The valid range is 0 to 1440 minutes. Note For periodic RADIUS server monitoring, you must set the idle timer to a value greater than 0.
Step 3	switch(config)# radius-server deadtime <i>minutes</i>	Specifies the number of minutes before the Cisco Nexus 5000 Series switch checks a RADIUS server that was previously unresponsive. The default value is 0 minutes. The valid range is 1 to 1440 minutes.
Step 4	switch(config)# exit	Exits configuration mode.
Step 5	switch# show radius-server	(Optional) Displays the RADIUS server configuration.
Step 6	switch# copy running-config startup-config	(Optional) Copies the running configuration to the startup configuration.

To configure periodic RADIUS server monitoring, perform this task:

```
switch# configure terminal
switch(config)# radius-server host 10.10.1.1 test username user1 password Ur2Gd2BH idle-time
3
switch(config)# radius-server deadtime 5
switch(config)# exit
switch# show radius-server
switch# copy running-config startup-config
```

Configuring the Dead-Time Interval

You can configure the dead-time interval for all RADIUS servers. The dead-time interval specifies the time that the Cisco Nexus 5000 Series switch waits after declaring a RADIUS server is dead, before sending out a test packet to determine if the server is now alive. The default value is 0 minutes.


Note

When the dead-time interval is 0 minutes, RADIUS servers are not marked as dead even if they are not responding. You can configure the dead-time interval for a RADIUS server group.

To configure dead time interval, perform this task:

SUMMARY STEPS

1. switch# **configure terminal**
2. #switch(config)# **radius-server deadtime**
3. switch(config)# **exit**
4. (Optional) switch# **show radius-server**
5. (Optional) switch# **copy running-config startup-config**

DETAILED STEPS

	Command or Action	Purpose
Step 1	switch# configure terminal	Enters configuration mode.
Step 2	#switch(config)# radius-server deadtime	Configures the dead-time interval. The default value is 0 minutes. The range is from 1 to 1440 minutes.
Step 3	switch(config)# exit	Exits configuration mode.
Step 4	switch# show radius-server	(Optional) Displays the RADIUS server configuration.
Step 5	switch# copy running-config startup-config	(Optional) Copies the running configuration to the startup configuration.

Manually Monitoring RADIUS Servers or Groups

To manually send a test message to a RADIUS server or to a server group, perform this task:

SUMMARY STEPS

1. switch# **test aaa server radius** {*ipv4-address* | *ipv6-address* | *server-name*} [**vrf** *vrf-name*] *username password*
2. switch# **test aaa group** *group-name username password*

DETAILED STEPS

	Command or Action	Purpose
Step 1	switch# test aaa server radius { <i>ipv4-address</i> <i>ipv6-address</i> <i>server-name</i> } [vrf <i>vrf-name</i>] <i>username password</i>	Sends a test message to a RADIUS server to confirm availability.
Step 2	switch# test aaa group <i>group-name username password</i>	Sends a test message to a RADIUS server group to confirm availability.

The following example shows how to manually send a test message to a RADIUS server:

```
switch# test aaa server radius 10.10.1.1 user1 Ur2Gd2BH
switch# test aaa group RadGroup user2 As3He3CI
```

Verifying RADIUS Configuration

To display RADIUS configuration information, perform one of the following tasks:

SUMMARY STEPS

1. switch# **show running-config radius** [all]
2. switch# **show startup-config radius**
3. switch# **show radius-server** [server-name | ipv4-address | ipv6-address] [directed-request | groups | sorted | statistics]

DETAILED STEPS

	Command or Action	Purpose
Step 1	switch# show running-config radius [all]	Displays the RADIUS configuration in the running configuration.
Step 2	switch# show startup-config radius	Displays the RADIUS configuration in the startup configuration.
Step 3	switch# show radius-server [server-name ipv4-address ipv6-address] [directed-request groups sorted statistics]	Displays all configured RADIUS server parameters.

For detailed information about the fields in the output from this command, refer to the *Cisco Nexus 5000 Series Command Reference*.

Displaying RADIUS Server Statistics

To display the statistics the Cisco Nexus 5000 Series switch maintains for RADIUS server activity, perform this task:

SUMMARY STEPS

1. switch# **show radius-server statistics** {hostname | ipv4-address | ipv6-address}

DETAILED STEPS

	Command or Action	Purpose
Step 1	switch# show radius-server statistics {hostname ipv4-address ipv6-address}	Displays the RADIUS statistics.

The following example shows how to display statistics:

```
switch# show radius-server statistics 10.10.1.1
```

Example RADIUS Configuration

The following example shows how to configure RADIUS:

```
switch# configure terminal
switch(config)# radius-server key 7 "ToIkLhPpG"
switch(config)# radius-server host 10.10.1.1 key 7 "ShMoMhTl" authentication accounting
switch(config)# aaa group server radius RadServer
switch(config-radius)# server 10.10.1.1
switch(config-radius)# exit
switch(config-radius)# use-vrf management
```

Default RADIUS Settings

The following table lists the default settings for RADIUS parameters.

Table 25: Default RADIUS Parameters

Parameters	Default
Server roles	Authentication and accounting
Dead timer interval	0 minutes
Retransmission count	1
Retransmission timer interval	5 seconds
Idle timer interval	0 minutes
Periodic server monitoring username	test
Periodic server monitoring password	test



Configuring TACACS+

This chapter contains the following sections:

- [About Configuring TACACS+, page 277](#)

About Configuring TACACS+

Information About TACACS+

The Terminal Access Controller Access Control System Plus (TACACS+) security protocol provides centralized validation of users attempting to gain access to a Cisco Nexus 5000 Series switch. TACACS+ services are maintained in a database on a TACACS+ daemon typically running on a UNIX or Windows NT workstation. You must have access to and must configure a TACACS+ server before the configured TACACS+ features on your Cisco Nexus 5000 Series switch are available.

TACACS+ provides for separate authentication, authorization, and accounting facilities. TACACS+ allows for a single access control server (the TACACS+ daemon) to provide each service (authentication, authorization, and accounting) independently. Each service is associated with its own database to take advantage of other services available on that server or on the network, depending on the capabilities of the daemon.

The TACACS+ client/server protocol uses TCP (TCP port 49) for transport requirements. Cisco Nexus 5000 Series switches provide centralized authentication using the TACACS+ protocol.

TACACS+ Advantages

TACACS+ has the following advantages over RADIUS authentication:

- Provides independent AAA facilities. For example, the Cisco Nexus 5000 Series switch can authorize access without authenticating.
- Uses the TCP transport protocol to send data between the AAA client and server, making reliable transfers with a connection-oriented protocol.
- Encrypts the entire protocol payload between the switch and the AAA server to ensure higher data confidentiality. The RADIUS protocol only encrypts passwords.

User Login with TACACS+

When a user attempts a Password Authentication Protocol (PAP) login to a Cisco Nexus 5000 Series switch using TACACS+, the following actions occur:

- 1 When the Cisco Nexus 5000 Series switch establishes a connection, it contacts the TACACS+ daemon to obtain the username and password.



Note

TACACS+ allows an arbitrary conversation between the daemon and the user until the daemon receives enough information to authenticate the user. This action is usually done by prompting for a username and password combination, but may include prompts for other items, such as the user's mother's maiden name.

- 2 The Cisco Nexus 5000 Series switch will receive one of the following responses from the TACACS+ daemon:
 - ACCEPT—User authentication succeeds and service begins. If the Cisco Nexus 5000 Series switch requires user authorization, authorization begins.
 - REJECT—User authentication failed. The TACACS+ daemon either denies further access to the user or prompts the user to retry the login sequence.
 - ERROR—An error occurred at some time during authentication either at the daemon or in the network connection between the daemon and the Cisco Nexus 5000 Series switch. If the Cisco Nexus 5000 Series switch receives an ERROR response, the switch tries to use an alternative method for authenticating the user.

The user also undergoes an additional authorization phase, if authorization has been enabled on the Cisco Nexus 5000 Series switch. Users must first successfully complete TACACS+ authentication before proceeding to TACACS+ authorization.

- 3 If TACACS+ authorization is required, the Cisco Nexus 5000 Series switch again contacts the TACACS+ daemon and it returns an ACCEPT or REJECT authorization response. An ACCEPT response contains attributes that are used to direct the EXEC or NETWORK session for that user and determines the services that the user can access.

Services include the following:

- Telnet, rlogin, Point-to-Point Protocol (PPP), Serial Line Internet Protocol (SLIP), or EXEC services
- Connection parameters, including the host or client IP address (IPv4 or IPv6), access list, and user timeouts

Default TACACS+ Server Encryption Type and Preshared Key

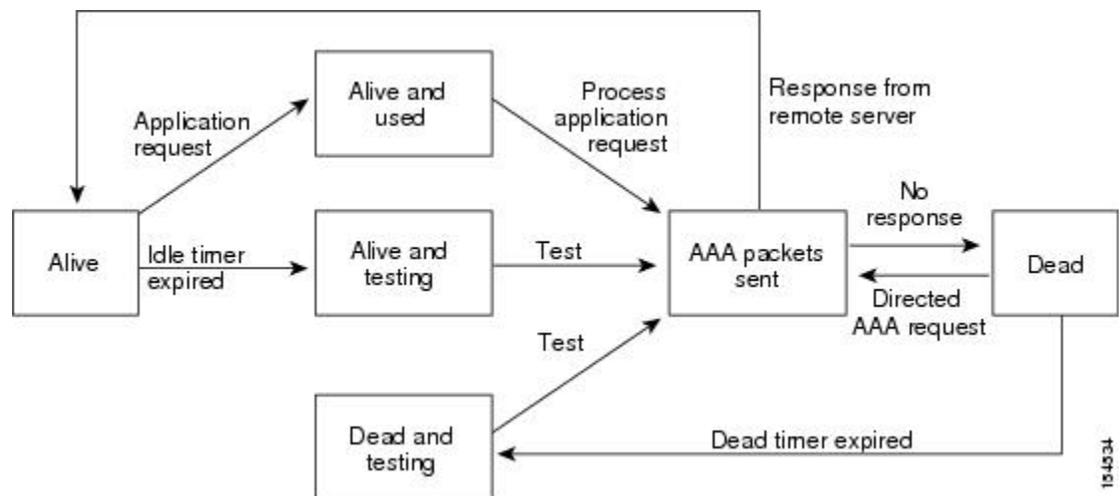
You must configure the TACACS+ preshared key to authenticate the switch to the TACACS+ server. A preshared key is a secret text string shared between the Cisco Nexus 5000 Series switch and the TACACS+ server host. The length of the key is restricted to 63 characters and can include any printable ASCII characters (white spaces are not allowed). You can configure a global preshared secret key for all TACACS+ server configurations on the Cisco Nexus 5000 Series switch to use.

You can override the global preshared key assignment by explicitly using the **key** option when configuring an individual TACACS+ server.

TACACS+ Server Monitoring

An unresponsive TACACS+ server can delay the processing of AAA requests. A Cisco Nexus 5000 Series switch can periodically monitor an TACACS+ server to check whether it is responding (or alive) to save time in processing AAA requests. The Cisco Nexus 5000 Series switch marks unresponsive TACACS+ servers as dead and does not send AAA requests to any dead TACACS+ servers. A Cisco Nexus 5000 Series switch periodically monitors dead TACACS+ servers and brings them to the alive state once they are responding. This process verifies that a TACACS+ server is in a working state before real AAA requests are sent its way. Whenever an TACACS+ server changes to the dead or alive state, a Simple Network Management Protocol (SNMP) trap is generated and the Cisco Nexus 5000 Series switch displays an error message that a failure is taking place before it can impact performance.

Figure 33: TACACS+ Server States



Note

The monitoring interval for alive servers and dead servers are different and can be configured by the user. The TACACS+ server monitoring is performed by sending a test authentication request to the TACACS+ server.

Prerequisites for TACACS+

TACACS+ has the following prerequisites:

- Obtain the IPv4 or IPv6 addresses or host names for the TACACS+ servers.
- Obtain the preshared keys from the TACACS+ servers, if any.
- Ensure that the Cisco Nexus 5000 Series switch is configured as a TACACS+ client of the AAA servers.

Guidelines and Limitations for TACACS+

TACACS+ has the following guidelines and limitations:

- You can configure a maximum of 64 TACACS+ servers on the Cisco Nexus 5000 Series switch.

Configuring TACACS+

TACACS+ Server Configuration Process

To configure TACACS+ servers, perform this task:

SUMMARY STEPS

1. Enable TACACS+.
2. Establish the TACACS+ server connections to the Cisco Nexus 5000 Series switch.
3. Configure the preshared secret keys for the TACACS+ servers.
4. If needed, configure TACACS+ server groups with subsets of the TACACS+ servers for AAA authentication methods.
5. If needed, configure any of the following optional parameters:
6. If needed, configure periodic TACACS+ server monitoring.

DETAILED STEPS

-
- Step 1** Enable TACACS+.
- Step 2** Establish the TACACS+ server connections to the Cisco Nexus 5000 Series switch.
- Step 3** Configure the preshared secret keys for the TACACS+ servers.
- Step 4** If needed, configure TACACS+ server groups with subsets of the TACACS+ servers for AAA authentication methods.
- Step 5** If needed, configure any of the following optional parameters:
- Dead-time interval
 - Allow TACACS+ server specification at login
 - Timeout interval
 - TCP port
- Step 6** If needed, configure periodic TACACS+ server monitoring.
-

Enabling TACACS+

By default, the TACACS+ feature is disabled on the Cisco Nexus 5000 Series switch. To explicitly enable the TACACS+ feature to access the configuration and verification commands for authentication, perform this task:

SUMMARY STEPS

1. switch# **configure terminal**
2. switch(config)# **feature tacacs+**
3. switch(config)# **exit**
4. (Optional) switch# **copy running-config startup-config**

DETAILED STEPS

	Command or Action	Purpose
Step 1	switch# configure terminal	Enters configuration mode.
Step 2	switch(config)# feature tacacs+	Enables TACACS+.
Step 3	switch(config)# exit	Exits configuration mode.
Step 4	switch# copy running-config startup-config	(Optional) Copies the running configuration to the startup configuration.

Configuring TACACS+ Server Hosts

To access a remote TACACS+ server, you must configure the IPv4 or IPv6 address or the hostname for the TACACS+ server on the Cisco Nexus 5000 Series switch. All TACACS+ server hosts are added to the default TACACS+ server group. You can configure up to 64 TACACS+ servers.

If a preshared key is not configured for a configured TACACS+ server, a warning message is issued if a global key is not configured. If a TACACS+ server key is not configured, the global key (if configured) is used for that server.

Before you configure TACACS+ server hosts, you should do the following:

- Enable TACACS+.
- Obtain the IPv4 or IPv6 addresses or the hostnames for the remote TACACS+ servers.

To configure TACACS+ server hosts, perform this task:

SUMMARY STEPS

1. switch# **configure terminal**
2. switch(config)# **tacacs-server host** {*ipv4-address* | *ipv6-address* | *host-name*}
3. switch(config)# **exit**
4. (Optional) switch# **show tacacs-server**
5. (Optional) switch# **copy running-config startup-config**

DETAILED STEPS

	Command or Action	Purpose
Step 1	switch# configure terminal	Enters configuration mode.
Step 2	switch(config)# tacacs-server host { <i>ipv4-address</i> <i>ipv6-address</i> <i>host-name</i> }	Specifies the IPv4 or IPv6 address or hostname for a TACACS+ server.
Step 3	switch(config)# exit	Exits configuration mode.
Step 4	switch# show tacacs-server	(Optional) Displays the TACACS+ server configuration.
Step 5	switch# copy running-config startup-config	(Optional) Copies the running configuration to the startup configuration.

You can delete a TACACS+ server host from a server group.

Configuring TACACS+ Global Preshared Keys

You can configure preshared keys at the global level for all servers used by the Cisco Nexus 5000 Series switch. A preshared key is a shared secret text string between the Cisco Nexus 5000 Series switch and the TACACS+ server hosts.

Before you configure preshared keys, you should do the following:

- Enable TACACS+.
- Obtain the preshared key values for the remote TACACS+ servers.

To configure global preshared keys, perform this task:

SUMMARY STEPS

1. switch# **configure terminal**
2. switch(config)# **tacacs-server key** [0 | 7] *key-value*
3. switch(config)# **exit**
4. (Optional) switch# **show tacacs-server**
5. (Optional) switch# **copy running-config startup-config**

DETAILED STEPS

	Command or Action	Purpose
Step 1	switch# configure terminal	Enters configuration mode.
Step 2	switch(config)# tacacs-server key [0 7] <i>key-value</i>	Specifies a preshared key for all TACACS+ servers. You can specify a clear text (0) or encrypted (7) preshared key. The default format is clear text. The maximum length is 63 characters.

	Command or Action	Purpose
		By default, no preshared key is configured.
Step 3	switch(config)# exit	Exits configuration mode.
Step 4	switch# show tacacs-server	(Optional) Displays the TACACS+ server configuration. Note The preshared keys are saved in encrypted form in the running configuration. Use the show running-config command to display the encrypted preshared keys.
Step 5	switch# copy running-config startup-config	(Optional) Copies the running configuration to the startup configuration.

The following example shows how to configure global preshared keys:

```
switch# configure terminal
switch(config)# tacacs-server key 0 QsEfThUkO
switch(config)# exit
switch# show tacacs-server
switch# copy running-config startup-config
```

Related Topics

[Enabling TACACS+ , on page 280](#)

Configuring TACACS+ Server Preshared Keys

You can configure preshared keys for a TACACS+ server. A preshared key is a shared secret text string between the Cisco Nexus 5000 Series switch and the TACACS+ server host.

To configure the TACACS+ preshared keys, perform this task:

SUMMARY STEPS

1. switch# **configure terminal**
2. switch(config)# **tacacs-server host** {*ipv4-address* | *ipv6-address* | *host-name*} **key** [0 | 7] *key-value*
3. switch(config)# **exit**
4. (Optional) switch# **show tacacs-server**
5. (Optional) switch# **copy running-config startup-config**

DETAILED STEPS

	Command or Action	Purpose
Step 1	switch# configure terminal	Enters configuration mode.
Step 2	switch(config)# tacacs-server host { <i>ipv4-address</i> <i>ipv6-address</i> <i>host-name</i> } key [0 7] <i>key-value</i>	Specifies a preshared key for a specific TACACS+ server. You can specify a clear text (0) or encrypted (7) preshared key. The default format is clear text. The maximum length is 63 characters.

	Command or Action	Purpose
		This preshared key is used instead of the global preshared key.
Step 3	switch(config)# exit	Exits configuration mode.
Step 4	switch# show tacacs-server	(Optional) Displays the TACACS+ server configuration. Note The preshared keys are saved in encrypted form in the running configuration. Use the show running-config command to display the encrypted preshared keys.
Step 5	switch# copy running-config startup-config	(Optional) Copies the running configuration to the startup configuration.

The following example shows how to configure the TACACS+ preshared keys:

```
switch# configure terminal
switch(config)# tacacs-server host 10.10.1.1 key 0 P1IjUhYg
switch(config)# exit
switch# show tacacs-server
switch# copy running-config startup-config
```

Configuring TACACS+ Server Groups

You can specify one or more remote AAA servers to authenticate users using server groups. All members of a group must belong to the TACACS+ protocol. The servers are tried in the same order in which you configure them.

You can configure these server groups at any time but they only take effect when you apply them to an AAA service.

To configure TACACS+ server groups, perform this task:

SUMMARY STEPS

1. switch# **configure terminal**
2. switch(config)# **aaa group server tacacs+ group-name**
3. switch(config-tacacs+)# **server {ipv4-address | ipv6-address | host-name}**
4. (Optional) switch(config-tacacs+)# **deadtime minutes**
5. switch(config-tacacs+)# **exit**
6. (Optional) switch(config)# **show tacacs-server groups**
7. (Optional) switch(config)# **copy running-config startup-config**

DETAILED STEPS

	Command or Action	Purpose
Step 1	switch# configure terminal	Enters configuration mode.

	Command or Action	Purpose
Step 2	switch(config)# aaa group server tacacs+ group-name	Creates a TACACS+ server group and enters the TACACS+ server group configuration mode for that group.
Step 3	switch(config-tacacs+)# server {ipv4-address ipv6-address host-name}	Configures the TACACS+ server as a member of the TACACS+ server group. If the specified TACACS+ server is not found, configure it using the tacacs-server host command and retry this command.
Step 4	switch(config-tacacs+)# deadtime minutes	(Optional) Configures the monitoring dead time. The default is 0 minutes. The range is from 0 through 1440. Note If the dead-time interval for a TACACS+ server group is greater than zero (0), that value takes precedence over the global dead-time value.
Step 5	switch(config-tacacs+)# exit	Exits configuration mode.
Step 6	switch(config)# show tacacs-server groups	(Optional) Displays the TACACS+ server group configuration.
Step 7	switch(config)# copy running-config startup-config	(Optional) Copies the running configuration to the startup configuration.

The following example shows how to configure a TACACS+ server group:

```
switch# configure terminal
switch(config)# aaa group server tacacs+ TacServer
switch(config-tacacs+)# server 10.10.2.2
switch(config-tacacs+)# deadtime 30
switch(config-tacacs+)# exit
switch(config)# show tacacs-server groups
switch(config)# copy running-config startup-config
```

Specifying a TACACS+ Server at Login

You can configure the switch to allow the user to specify which TACACS+ server to send the authenticate request by enabling the directed-request option. By default, a Cisco Nexus 5000 Series switch forwards an authentication request based on the default AAA authentication method. If you enable this option, the user can log in as `username@hostname`, where `hostname` is the name of a configured RADIUS server.



Note User specified logins are only supported for Telnet sessions.

To specify a TACACS+ server at login, perform this task:

SUMMARY STEPS

1. switch# **configure terminal**
2. switch(config)# **tacacs-server directed-request**
3. switch(config)# **exit**
4. (Optional) switch# **show tacacs-server directed-request**
5. (Optional) switch# **copy running-config startup-config**

DETAILED STEPS

	Command or Action	Purpose
Step 1	switch# configure terminal	Enters configuration mode.
Step 2	switch(config)# tacacs-server directed-request	Allows users to specify a TACACS+ server to send the authentication request when logging in. The default is disabled.
Step 3	switch(config)# exit	Exits configuration mode.
Step 4	switch# show tacacs-server directed-request	(Optional) Displays the TACACS+ directed request configuration.
Step 5	switch# copy running-config startup-config	(Optional) Copies the running configuration to the startup configuration.

Configuring the Global TACACS+ Timeout Interval

You can set a global timeout interval that the Cisco Nexus 5000 Series switch waits for responses from all TACACS+ servers before declaring a timeout failure. The timeout interval determines how long the switch waits for responses from TACACS+ servers before declaring a timeout failure.

To specify a TACACS+ global timeout interval, perform this task:

SUMMARY STEPS

1. switch# **configure terminal**
2. switch(config)# **tacacs-server timeout seconds**
3. switch(config)# **exit**
4. (Optional) switch# **show tacacs-server**
5. (Optional) switch# **copy running-config startup-config**

DETAILED STEPS

	Command or Action	Purpose
Step 1	switch# configure terminal	Enters configuration mode.

	Command or Action	Purpose
Step 2	switch(config)# tacacs-server timeout <i>seconds</i>	Specifies the timeout interval for TACACS+ servers. The default timeout interval is 5 second and the range is from 1 to 60 seconds.
Step 3	switch(config)# exit	Exits configuration mode.
Step 4	switch# show tacacs-server	(Optional) Displays the TACACS+ server configuration.
Step 5	switch# copy running-config startup-config	(Optional) Copies the running configuration to the startup configuration.

Configuring the Timeout Interval for a Server

You can set a timeout interval that the Cisco Nexus 5000 Series switch waits for responses from a TACACS+ server before declaring a timeout failure. The timeout interval determines how long the switch waits for responses from a TACACS+ server before declaring a timeout failure.

To configure the timeout interval for a server, perform this task:

SUMMARY STEPS

1. switch# **configure terminal**
2. switch(config)# switch(config)# **tacacs-server host** {*ipv4-address* | *ipv6-address* | *host-name*} **timeout** *seconds*
3. switch(config)# **exit**
4. (Optional) switch# **show tacacs-server**
5. (Optional) switch# **copy running-config startup-config**

DETAILED STEPS

	Command or Action	Purpose
Step 1	switch# configure terminal	Enters configuration mode.
Step 2	switch(config)# switch(config)# tacacs-server host { <i>ipv4-address</i> <i>ipv6-address</i> <i>host-name</i> } timeout <i>seconds</i>	Specifies the timeout interval for a specific server. The default is the global value. Note The timeout interval value specified for a TACACS+ server overrides the global timeout interval value specified for all TACACS+ servers.
Step 3	switch(config)# exit	Exits configuration mode.
Step 4	switch# show tacacs-server	(Optional) Displays the TACACS+ server configuration.

	Command or Action	Purpose
Step 5	switch# copy running-config startup-config	(Optional) Copies the running configuration to the startup configuration.

Configuring TCP Ports

You can configure another TCP port for the TACACS+ servers if there are conflicts with another application. By default, Cisco Nexus 5000 Series switches use port 49 for all TACACS+ requests.

To configure TCP ports, perform this task:

SUMMARY STEPS

1. switch# **configure terminal**
2. switch(config)# **tacacs-server host** {*ipv4-address* | *ipv6-address* | *host-name*} **port** *tcp-port*
3. switch(config)# **exit**
4. (Optional) switch# **show tacacs-server**
5. (Optional) switch# **copy running-config startup-config**

DETAILED STEPS

	Command or Action	Purpose
Step 1	switch# configure terminal	Enters configuration mode.
Step 2	switch(config)# tacacs-server host { <i>ipv4-address</i> <i>ipv6-address</i> <i>host-name</i> } port <i>tcp-port</i>	Specifies the UDP port to use for TACACS+ accounting messages. The default TCP port is 49. The range is from 1 to 65535.
Step 3	switch(config)# exit	Exits configuration mode.
Step 4	switch# show tacacs-server	(Optional) Displays the TACACS+ server configuration.
Step 5	switch# copy running-config startup-config	(Optional) Copies the running configuration to the startup configuration.

The following example shows how to configure TCP ports:

```
switch# configure terminal
switch(config)# tacacs-server host 10.10.1.1 port 2
switch(config)# exit
switch# show tacacs-server
switch# copy running-config startup-config
```

Configuring Periodic TACACS+ Server Monitoring

You can monitor the availability of TACACS+ servers. These parameters include the username and password to use for the server and an idle timer. The idle timer specifies the interval in which a TACACS+ server

receives no requests before the Cisco Nexus 5000 Series switch sends out a test packet. You can configure this option to test servers periodically, or you can run a one-time only test.



Note To protect network security, we recommend that you use a user name that is not the same as an existing username in the TACACS+ database.

The test idle timer specifies the interval in which a TACACS+ server receives no requests before the Cisco Nexus 5000 Series switch sends out a test packet.



Note The default idle timer value is 0 minutes. When the idle time interval is 0 minutes, periodic TACACS+ server monitoring is not performed.

To configure periodic TACACS+ server monitoring, perform this task:

SUMMARY STEPS

1. switch# **configure terminal**
2. switch(config)# **tacacs-server host** {*ipv4-address* | *ipv6-address* | *host-name*} **test** {*idle-time minutes* | **password** *password* [*idle-time minutes*] | **username** *name* [**password** *password* [*idle-time minutes*]]}
3. switch(config)# **tacacs-server dead-time** *minutes*
4. switch(config)# **exit**
5. (Optional) switch# **show tacacs-server**
6. (Optional) switch# **copy running-config startup-config**

DETAILED STEPS

	Command or Action	Purpose
Step 1	switch# configure terminal	Enters configuration mode.
Step 2	switch(config)# tacacs-server host { <i>ipv4-address</i> <i>ipv6-address</i> <i>host-name</i> } test { <i>idle-time minutes</i> password <i>password</i> [<i>idle-time minutes</i>] username <i>name</i> [password <i>password</i> [<i>idle-time minutes</i>]]}	Specifies parameters for server monitoring. The default username is test and the default password is test. The default value for the idle timer is 0 minutes and the valid range is 0 to 1440 minutes. Note For periodic TACACS+ server monitoring, the idle timer value must be greater than 0.
Step 3	switch(config)# tacacs-server dead-time <i>minutes</i>	Specifies the number minutes before the Cisco Nexus 5000 Series switch checks a TACACS+ server that was previously unresponsive. The default value is 0 minutes and the valid range is 0 to 1440 minutes.
Step 4	switch(config)# exit	Exits configuration mode.
Step 5	switch# show tacacs-server	(Optional) Displays the TACACS+ server configuration.

	Command or Action	Purpose
Step 6	switch# copy running-config startup-config	(Optional) Copies the running configuration to the startup configuration.

The following example shows how to configure periodic TACACS+ server monitoring:

```
switch# configure terminal
switch(config)# tacacs-server host 10.10.1.1 test username user1 password Ur2Gd2BH idle-time 3
switch(config)# tacacs-server dead-time 5
switch(config)# exit
switch# show tacacs-server
switch# copy running-config startup-config
```

Configuring the Dead-Time Interval

You can configure the dead-time interval for all TACACS+ servers. The dead-time interval specifies the time that the Cisco Nexus 5000 Series switch waits, after declaring a TACACS+ server is dead, before sending out a test packet to determine if the server is now alive.



Note

When the dead-timer interval is 0 minutes, TACACS+ servers are not marked as dead even if they are not responding. You can configure the dead-timer per group.

To configure the dead-time interval for all TACACS+ servers, perform this task:

SUMMARY STEPS

1. switch# **configure terminal**
2. switch(config)# **tacacs-server deadtime *minutes***
3. switch(config)# **exit**
4. (Optional) switch# **show tacacs-server**
5. (Optional) switch# **copy running-config startup-config**

DETAILED STEPS

	Command or Action	Purpose
Step 1	switch# configure terminal	Enters configuration mode.
Step 2	switch(config)# tacacs-server deadtime <i>minutes</i>	Configures the global dead-time interval. The default value is 0 minutes. The range is from 1 to 1440 minutes.
Step 3	switch(config)# exit	Exits configuration mode.
Step 4	switch# show tacacs-server	(Optional) Displays the TACACS+ server configuration.

	Command or Action	Purpose
Step 5	switch# copy running-config startup-config	(Optional) Copies the running configuration to the startup configuration.

Manually Monitoring TACACS+ Servers or Groups

To manually issue a test message to a TACACS+ server or to a server group, perform this task:

SUMMARY STEPS

1. switch# **test aaa server tacacs+** {*ipv4-address* | *ipv6-address* | *host-name*} [**vrf** *vrf-name*] *username password*
2. switch# **test aaa group** *group-name username password*

DETAILED STEPS

	Command or Action	Purpose
Step 1	switch# test aaa server tacacs+ { <i>ipv4-address</i> <i>ipv6-address</i> <i>host-name</i> } [vrf <i>vrf-name</i>] <i>username password</i>	Sends a test message to a TACACS+ server to confirm availability.
Step 2	switch# test aaa group <i>group-name username password</i>	Sends a test message to a TACACS+ server group to confirm availability.

The following example shows how to manually issue a test message:

```
switch# test aaa server tacacs+ 10.10.1.1 user1 Ur2Gd2BH
switch# test aaa group TacGroup user2 As3He3CI
```

Disabling TACACS+

You can disable TACACS+.



Caution

When you disable TACACS+, all related configurations are automatically discarded.

To disable TACACS+, perform this task:

SUMMARY STEPS

1. switch# **configure terminal**
2. switch(config)# **no feature tacacs+**
3. switch(config)# **exit**
4. (Optional) switch# **copy running-config startup-config**

DETAILED STEPS

	Command or Action	Purpose
Step 1	switch# configure terminal	Enters configuration mode.
Step 2	switch(config)# no feature tacacs+	Disables TACACS+.
Step 3	switch(config)# exit	Exits configuration mode.
Step 4	switch# copy running-config startup-config	(Optional) Copies the running configuration to the startup configuration.

Displaying TACACS+ Statistics

To display the statistics the Cisco Nexus 5000 Series switch maintains for TACACS+ activity, perform this task:

SUMMARY STEPS

1. switch# **show tacacs-server statistics** {hostname | ipv4-address | ipv6-address}

DETAILED STEPS

	Command or Action	Purpose
Step 1	switch# show tacacs-server statistics {hostname ipv4-address ipv6-address}	Displays the TACACS+ statistics.

For detailed information about the fields in the output from this command, see the *Cisco Nexus 5000 Series Command Reference*.

Verifying TACACS+ Configuration

To display TACACS+ configuration information, perform one of the following tasks:

SUMMARY STEPS

1. switch# **show tacacs+** {status | pending | pending-diff}
2. switch# **show running-config tacacs** [all]
3. switch# **show startup-config tacacs**
4. switch# **show tacacs-serve** [host-name | ipv4-address | ipv6-address] [directed-request | groups | sorted | statistics]

DETAILED STEPS

	Command or Action	Purpose
Step 1	switch# show tacacs+ {status pending pending-diff}	Displays the TACACS+ Cisco Fabric Services distribution status and other details.
Step 2	switch# show running-config tacacs [all]	Displays the TACACS+ configuration in the running configuration.
Step 3	switch# show startup-config tacacs	Displays the TACACS+ configuration in the startup configuration.
Step 4	switch# show tacacs-serve [host-name ipv4-address ipv6-address] [directed-request groups sorted statistics]	Displays all configured TACACS+ server parameters.

Example TACACS+ Configuration

The following example shows how to configure TACACS+:

```
switch# configure terminal
switch(config)# feature tacacs+
switch(config)# tacacs-server key 7 "ToIkLhPpG"
switch(config)# tacacs-server host 10.10.2.2 key 7 "ShMoMhTl"
switch(config)# aaa group server tacacs+ TacServer
switch(config-tacacs+)# server 10.10.2.2
switch(config-tacacs+)# use-vrf management
```

Default TACACS+ Settings

The following table lists the default settings for TACACS+ parameters.

Table 26: Default TACACS+ Parameters

Parameters	Default
TACACS+	Disabled
Dead timer interval	0 minutes
Timeout interval	5 seconds
Idle timer interval	0 minutes
Periodic server monitoring username	test
Periodic server monitoring password	test



Configuring SSH and Telnet

This chapter contains the following sections:

- [Configuring SSH and Telnet, page 295](#)

Configuring SSH and Telnet

Information About SSH and Telnet

SSH Server

The Secure Shell Protocol (SSH) server feature enables a SSH client to make a secure, encrypted connection to a Cisco Nexus 5000 Series switch. SSH uses strong encryption for authentication. The SSH server in the Cisco Nexus 5000 Series switch will interoperate with publicly and commercially available SSH clients.

The user authentication mechanisms supported for SSH are RADIUS, TACACS+, and the use of locally stored user names and passwords.

SSH Client

The SSH client feature is an application running over the SSH protocol to provide device authentication and encryption. The SSH client enables a Cisco Nexus 5000 Series switch to make a secure, encrypted connection to another Cisco Nexus 5000 Series switch or to any other device running an SSH server. This connection provides an outbound connection that is encrypted. With authentication and encryption, the SSH client allows for a secure communication over an insecure network.

The SSH client in the Cisco Nexus 5000 Series switch works with publicly and commercially available SSH servers.

SSH Server Keys

SSH requires server keys for secure communications to the Cisco Nexus 5000 Series switch. You can use SSH keys for the following SSH options:

- SSH version 2 using Rivest, Shamir, and Adelman (RSA) public-key cryptography

- SSH version 2 using the Digital System Algorithm (DSA)

Be sure to have an SSH server key-pair with the appropriate version before enabling the SSH service. You can generate the SSH server key-pair according to the SSH client version used. The SSH service accepts three types of key-pairs for use by SSH version 2:

- The `dsa` option generates the DSA key-pair for the SSH version 2 protocol.
- The `rsa` option generates the RSA key-pair for the SSH version 2 protocol.

By default, the Cisco Nexus 5000 Series switch generates an RSA key using 1024 bits.

SSH supports the following public key formats:

- OpenSSH
- IETF Secure Shell (SECSH)



Caution

If you delete all of the SSH keys, you cannot start the SSH services.

Telnet Server

The Telnet protocol enables TCP/IP connections to a host. Telnet allows a user at one site to establish a TCP connection to a login server at another site, and then passes the keystrokes from one system to the other. Telnet can accept either an IP address or a domain name as the remote system address.

The Telnet server is enabled by default on the Cisco Nexus 5000 Series switch.

Guidelines and Limitations for SSH

SSH has the following configuration guidelines and limitations:

- The Cisco Nexus 5000 Series switch supports only SSH version 2 (SSHv2).

Configuring SSH

Generating SSH Server Keys

You can generate an SSH server key based on your security requirements. The default SSH server key is an RSA key generated using 1024 bits. To generate SSH server keys, perform this task:

SUMMARY STEPS

1. `switch# configure terminal`
2. `switch(config)# ssh key {dsa [force] | rsa [bits [force]]}`
3. `switch(config)# exit`
4. (Optional) `switch# show ssh key`
5. (Optional) `switch# copy running-config startup-config`

DETAILED STEPS

	Command or Action	Purpose
Step 1	switch# configure terminal	Enters configuration mode.
Step 2	switch(config)# ssh key {dsa [force] rsa [bits] [force]}	Generates the SSH server key. The <i>bits</i> argument is the number of bits used to generate the key. The range is 768 to 2048 and the default value is 1024. Use the force keyword to replace an existing key.
Step 3	switch(config)# exit	Exits global configuration mode.
Step 4	switch# show ssh key	(Optional) Displays the SSH server keys.
Step 5	switch# copy running-config startup-config	(Optional) Copies the running configuration to the startup configuration.

The following example shows how to generate an SSH server key:

```
switch# configure terminal
switch(config)# ssh key rsa 2048
switch(config)# exit
switch# show ssh key
switch# copy running-config startup-config
```

Specifying the SSH Public Keys for User Accounts

You can configure an SSH public key to log in using the SSH client without being prompted for a password. You can specify the SSH public key in one of three different formats:

- Open SSH format
- IETF SECSH format
- Public Key Certificate in PEM format

Specifying the SSH Public Keys in Open SSH Format

You can specify the SSH public keys in SSH format for user accounts.

To specify the SSH public keys in open SSH format, generate an SSH public key in open SSH format and perform this task:

SUMMARY STEPS

1. switch# **configure terminal**
2. switch(config)# **username username sshkey ssh-key**
3. switch(config)# **exit**
4. (Optional) switch# **show user-account**
5. (Optional) switch# **copy running-config startup-config**

DETAILED STEPS

	Command or Action	Purpose
Step 1	switch# configure terminal	Enters configuration mode.
Step 2	switch(config)# username <i>username</i> sshkey <i>ssh-key</i>	Configures the SSH public key in SSH format.
Step 3	switch(config)# exit	Exits global configuration mode.
Step 4	switch# show user-account	(Optional) Displays the user account configuration.
Step 5	switch# copy running-config startup-config	(Optional) Copies the running configuration to the startup configuration.

The following example shows how to specify an SSH public keys in open SSH format:

```
switch# configure terminal
switch(config)# username User1 sshkey ssh-rsa
AAAAB3NzaC1yc2EAAAABIwAAAIEAri3mQy4W1AV9Y2t2hrEWgbUEYz
CfTPO5B8LRkedn56BEy2N9ZcdpqE6aqJLZwFZcTFEzaAAZp9AS86dgBAjsKGS7UxnhGySr8ZELv+DQBsDQH6rz0KR+2Da8hJD4Z
XIeccWk0gS1DQUNZ300xstQsYZUtqnx1bvm5Ninn0McNinn0Mc=
switch(config)# exit
switch# show user-account
switch# copy running-config startup-config
```

**Note**

The **username** command example above is a single line that has been broken for legibility.

Specifying the SSH Public Keys in IETF SECSH Format

You can specify the SSH public keys in IETF SECSH format for user accounts.

To specify the SSH public keys in IETF SECSH format, generate an SSH public key in IETF SCHSH format, and perform this task:

SUMMARY STEPS

1. switch# **copy server-file bootflash:** *filename*
2. switch# **configure terminal**
3. switch(config)# **username** *username* **sshkey file** *filename*
4. switch(config)# **exit**
5. (Optional) switch# **show user-account**
6. (Optional) switch# **copy running-config startup-config**

DETAILED STEPS

	Command or Action	Purpose
Step 1	switch# copy server-file bootflash: filename	Downloads the file containing the SSH key in IETF SECSH format from a server. The server can be FTP, SCP, SFTP, or TFTP.
Step 2	switch# configure terminal	Enters configuration mode.
Step 3	switch(config)# username username sshkey file filename	Configures the SSH public key in SSH format.
Step 4	switch(config)# exit	Exits global configuration mode.
Step 5	switch# show user-account	(Optional) Displays the user account configuration.
Step 6	switch# copy running-config startup-config	(Optional) Copies the running configuration to the startup configuration.

The following example shows how to specify the SSH public keys in the IETF SECSH format:

```
switch#copy tftp://10.10.1.1/secsh_file.pub bootflash:secsh_file.pub
switch# configure terminal
switch(config)# username User1 sshkey file bootflash:secsh_file.pub
switch(config)# exit
switch# show user-account
switch# copy running-config startup-config
```

Specifying the SSH Public Keys in PEM-Formatted Public Key Certificate Form

You can specify the SSH public keys in PEM-formatted Public Key Certificate form for user accounts.

To specify the SSH public keys in PEM-formatted Public Key Certificate form, generate an SSH public key in PEM-Formatted Public Key Certificate form and perform this task:

SUMMARY STEPS

1. switch# **copy server-file bootflash: filename**
2. switch# **configure terminal**
3. (Optional) switch# **show user-account**
4. (Optional) switch# **copy running-config startup-config**

DETAILED STEPS

	Command or Action	Purpose
Step 1	switch# copy server-file bootflash: filename	Downloads the file containing the SSH key in PEM-formatted Public Key Certificate form from a server. The server can be FTP, SCP, SFTP, or TFTP

	Command or Action	Purpose
Step 2	switch# configure terminal	Enters configuration mode.
Step 3	switch# show user-account	(Optional) Displays the user account configuration.
Step 4	switch# copy running-config startup-config	(Optional) Copies the running configuration to the startup configuration.

The following example shows how to specify the SSH public keys in PEM-formatted public key certificate form:

```
switch# copy tftp://10.10.1.1/cert.pem bootflash:cert.pem
switch# configure terminal
switch# show user-account
switch# copy running-config startup-config
```

Starting SSH Sessions to Remote Devices

To start SSH sessions to connect to remote devices from your Cisco Nexus 5000 Series switch, perform this task:

SUMMARY STEPS

1. switch# **ssh** {hostname | username@hostname} [**vrf** vrf-name]

DETAILED STEPS

	Command or Action	Purpose
Step 1	switch# ssh {hostname username@hostname} [vrf vrf-name]	Creates an SSH session to a remote device. The <i>hostname</i> argument can be an IPv4 address, an IPv6 address, or a host name.

Clearing SSH Hosts

When you download a file from a server using SCP or SFTP, you establish a trusted SSH relationship with that server. To clear the list of trusted SSH servers for your user account, perform this task:

SUMMARY STEPS

1. switch# **clear ssh hosts**

DETAILED STEPS

	Command or Action	Purpose
Step 1	switch# clear ssh hosts	Clears the SSH host sessions.

Disabling the SSH Server

By default, the SSH server is enabled on the Cisco Nexus 5000 Series switch.

To disable the SSH server to prevent SSH access to the switch, perform this task:

SUMMARY STEPS

1. switch# **configure terminal**
2. switch(config)# **no feature ssh**
3. switch(config)# **exit**
4. (Optional) switch# **show ssh server**
5. (Optional) switch# **copy running-config startup-config**

DETAILED STEPS

	Command or Action	Purpose
Step 1	switch# configure terminal	Enters configuration mode.
Step 2	switch(config)# no feature ssh	Disables the SSH server. The default is enabled.
Step 3	switch(config)# exit	Exits global configuration mode.
Step 4	switch# show ssh server	(Optional) Displays the SSH server configuration.
Step 5	switch# copy running-config startup-config	(Optional) Copies the running configuration to the startup configuration.

Deleting SSH Server Keys

You can delete SSH server keys after you disable the SSH server.

**Note**

To reenabling SSH, you must first generate an SSH server key.

To delete the SSH server keys, perform this task:

SUMMARY STEPS

1. switch# **configure terminal**
2. switch(config)# **no feature ssh**
3. switch(config)# **no ssh key [dsa | rsa]**
4. switch(config)# **exit**
5. (Optional) switch# **show ssh key**
6. (Optional) switch# **copy running-config startup-config**

DETAILED STEPS

	Command or Action	Purpose
Step 1	switch# configure terminal	Enters configuration mode.
Step 2	switch(config)# no feature ssh	Disables the SSH server.
Step 3	switch(config)# no ssh key [dsa rsa]	Deletes the SSH server key. The default is to delete all the SSH keys.
Step 4	switch(config)# exit	Exits global configuration mode.
Step 5	switch# show ssh key	(Optional) Displays the SSH server configuration.
Step 6	switch# copy running-config startup-config	(Optional) Copies the running configuration to the startup configuration.

Clearing SSH Sessions

To clear SSH sessions from the Cisco Nexus 5000 Series switch, perform this task:

SUMMARY STEPS

1. switch# **show users**
2. switch# **clear line vty-line**

DETAILED STEPS

	Command or Action	Purpose
Step 1	switch# show users	Displays user session information.
Step 2	switch# clear line vty-line	Clears a user SSH session.

SSH Example Configuration

The following example shows how to configure SSH:

SUMMARY STEPS

1. Generate an SSH server key.
2. Enable the SSH server.
3. Display the SSH server key.
4. Specify the SSH public key in Open SSH format.
5. Save the configuration.

DETAILED STEPS

Step 1 Generate an SSH server key.

```
switch(config)# ssh key rsa
generating rsa key(1024 bits).....
.
generated rsa key
```

Step 2 Enable the SSH server.

```
switch# configure terminal
switch(config)# feature ssh
```

Note This step should not be required as the SSH server is enabled by default.

Step 3 Display the SSH server key.

```
switch(config)# show ssh key
rsa Keys generated:Fri May 8 22:09:47 2009
```

```
ssh-rsa AAAAB3NzaC1yc2EAAAABIwAAAIEAri3mQy4W1AV9Y2t2hrEWgbUEYzCfTPO5B8LRkedn56BEy2N9ZcdpqE6aqJLZwfZ/
cTFEzaAAZp9AS86dgBAjsKGS7UxnhGySr8ZELv+DQBsDQH6rZt0KR+2Da8hJD4ZXIeccWk0gS1DQUNZ300xstQsYZUtqnX1bvm5/
Ninn0Mc=
```

```
bitcount:1024
```

```
fingerprint:
4b:4d:f6:b9:42:e9:d9:71:3c:bd:09:94:4a:93:ac:ca
*****
could not retrieve dsa key information
*****
```

Step 4 Specify the SSH public key in Open SSH format.

```
switch(config)# username User1 sshkey ssh-rsa AAAAB3NzaC1yc2EAAAABIwAAAIEAri3mQy4W1AV9Y2t2hrEWgbUEYz
CfTPO5B8LRkedn56BEy2N9ZcdpqE6aqJLZwfZcTFEzaAAZp9AS86dgBAjsKGS7UxnhGySr8ZELv+DQBsDQH6rZt0KR+2Da8hJD4Z
XIeccWk0gS1DQUNZ300xstQsYZUtqnX1bvm5Ninn0McNinn0Mc=
```

Step 5 Save the configuration.

```
switch(config)# copy running-config startup-config
```

Configuring Telnet

Enabling the Telnet Server

By default, the Telnet server is enabled. You can disable the Telnet server on your Cisco Nexus 5000 Series switch.

SUMMARY STEPS

1. switch# **configure terminal**
2. switch(config)# **feature telnet**

DETAILED STEPS

	Command or Action	Purpose
Step 1	switch# configure terminal	Enters configuration mode.
Step 2	switch(config)# feature telnet	Disables the Telnet server. The default is enabled.

Reenabling the Telnet Server

If the Telnet server on your Cisco Nexus 5000 Series switch has been disabled, you can reenabling it.

SUMMARY STEPS

1. switch(config)# **feature telnet**

DETAILED STEPS

	Command or Action	Purpose
Step 1	switch(config)# feature telnet	Reenables the Telnet server.

Starting Telnet Sessions to Remote Devices

Before you start a Telnet session to connect to remote devices, you should do the following:

- Obtain the hostname for the remote device and, if needed, the user name on the remote device.
- Enable the Telnet server on the Cisco Nexus 5000 Series switch.
- Enable the Telnet server on the remote device.

To start Telnet sessions to connect to remote devices from your Cisco Nexus 5000 Series switch, perform this task:

SUMMARY STEPS

1. switch# **telnet** *hostname*

DETAILED STEPS

	Command or Action	Purpose
Step 1	switch# telnet <i>hostname</i>	Creates a Telnet session to a remote device. The <i>hostname</i> argument can be an IPv4 address, an IPv6 address, or a device name.

The following example shows starting a Telnet session to connect to a remote device:

```
switch# telnet 10.10.1.1
Trying 10.10.1.1...
Connected to 10.10.1.1.
Escape character is '^]'.
switch login:
```

Clearing Telnet Sessions

To clear Telnet sessions from the Cisco Nexus 5000 Series switch, perform this task:

SUMMARY STEPS

1. switch# **show users**
2. switch# **clear line** *vtty-line*

DETAILED STEPS

	Command or Action	Purpose
Step 1	switch# show users	Displays user session information.
Step 2	switch# clear line <i>vtty-line</i>	Clears a user Telnet session.

Verifying the SSH and Telnet Configuration

To display the SSH configuration information, perform one of the following tasks:

- switch# **show ssh key** [*dsa* | *rsa*]
Displays SSH server key-pair information.
- switch# **show running-config security** [*all*]

Displays the SSH and user account configuration in the running configuration. The **all** keyword displays the default values for the SSH and user accounts.

- switch# **show ssh server**
Displays the SSH server configuration.
- switch# **show user-account**
Displays user account information.

Default SSH Settings

The following table lists the default settings for SSH parameters.

Table 27: Default SSH Parameters

Parameters	Default
SSH server	Enabled
SSH server key	RSA key generated with 1024 bits
RSA key bits for generation	1024
Telnet server	Enabled



Configuring Access Control Lists

This chapter contains the following sections:

- [Information About ACLs, page 307](#)
- [Configuring IP ACLs, page 311](#)
- [Configuring MAC ACLs, page 316](#)
- [Example Configuration for MAC ACLs, page 321](#)
- [Information About VLAN ACLs, page 321](#)
- [Configuring VACLs, page 322](#)
- [Example Configuration for VACL, page 325](#)
- [Default ACL Settings, page 325](#)

Information About ACLs

An access control list (ACL) is an ordered set of rules that you can use to filter traffic. Each rule specifies a set of conditions that a packet must satisfy to match the rule. When the switch determines that an ACL applies to a packet, it tests the packet against the conditions of all rules. The first match determines whether the packet is permitted or denied. If there is no match, the switch applies the applicable default rule. The switch continues processing packets that are permitted and drops packets that are denied.

You can use ACLs to protect networks and specific hosts from unnecessary or unwanted traffic. For example, you could use ACLs to disallow HTTP traffic from a high-security network to the Internet. You could also use ACLs to allow HTTP traffic but only to specific sites, using the IP address of the site to identify it in an IP ACL.

IP ACL Types and Applications

The Cisco Nexus 5000 Series switch supports IPv4, IPv6, and MAC ACLs for security traffic filtering. The switch allows you to use IP ACLs as port ACLs and VLAN ACLs, as shown in the following table.

Table 28: Security ACL Applications

Application	Supported Interfaces	Types of ACLs Supported
Port ACL	<p>An ACL is considered a port ACL when you apply it to one of the following:</p> <ul style="list-style-type: none"> • Ethernet interface • Ethernet port-channel interface <p>When a port ACL is applied to a trunk port, the ACL filters traffic on all VLANs on the trunk port.</p>	<p>IPv4 ACLs</p> <p>IPv6 ACLs</p> <p>MAC ACLs</p>
VLAN ACL (VACL)	<p>An ACL is a VACL when you use an access map to associate the ACL with an action, and then apply the map to a VLAN.</p>	<p>IPv4 ACLs</p> <p>IPv6 ACLs</p> <p>MAC ACLs</p>

Application Order

When the switch processes a packet, it determines the forwarding path of the packet. The path determines which ACLs that the switch applies to the traffic. The switch applies the Port ACLs first.

Rules

You can create rules in access-list configuration mode by using the **permit** or **deny** command. The switch allows traffic that matches the criteria in a permit rule and blocks traffic that matches the criteria in a deny rule. You have many options for configuring the criteria that traffic must meet in order to match the rule.

Source and Destination

In each rule, you specify the source and the destination of the traffic that matches the rule. You can specify both the source and destination as a specific host, a network or group of hosts, or any host.

Protocols

ACLs allow you to identify traffic by protocol. For your convenience, you can specify some protocols by name. For example, in an IPv4 ACL, you can specify ICMP by name.

You can specify any protocol by number. In IPv4 ACLs, you can specify protocols by the integer that represents the Internet protocol number. For example, you can use 115 to specify Layer 2 Tunneling Protocol (L2TP) traffic.

Implicit Rules

IP ACLs have implicit rules, which means that although these rules do not appear in the running configuration, the switch applies them to traffic when no other rules in an ACL match.

All IPv4 ACLs include the following implicit rule:

```
deny ip any any
```

This implicit rule ensures that the switch denies unmatched IP traffic.

Additional Filtering Options

You can identify traffic by using additional options. IPv4 ACLs support the following additional filtering options:

- Layer 4 protocol
- TCP and UDP ports
- ICMP types and codes
- IGMP types
- Precedence level
- Differentiated Services Code Point (DSCP) value
- TCP packets with the ACK, FIN, PSH, RST, SYN, or URG bit set
- Established TCP connections

IPv6 ACLs support the following additional filtering options:

- Layer 4 protocol
- Authentication Header Protocol
- Encapsulating Security Payload
- Payload Compression Protocol
- Stream Control Transmission Protocol (SCTP)
- SCTP, TCP, and UDP ports
- ICMP types and codes
- IGMP types
- Flow label
- DSCP value
- TCP packets with the ACK, FIN, PSH, RST, SYN, or URG bit set
- Established TCP connections
- Packet length

Sequence Numbers

The switch supports sequence numbers for rules. Every rule that you enter receives a sequence number, either assigned by you or assigned automatically by the switch. Sequence numbers simplify the following ACL tasks:

- Adding new rules between existing rules—By specifying the sequence number, you specify where in the ACL a new rule should be positioned. For example, if you need to insert a rule between rules numbered 100 and 110, you could assign a sequence number of 105 to the new rule.
- Removing a rule—Without using a sequence number, removing a rule requires that you enter the whole rule, as follows:

```
switch(config-acl)# no permit tcp 10.0.0.0/8 any
```

However, if the same rule had a sequence number of 101, removing the rule requires only the following command:

```
switch(config-acl)# no 101
```

- Moving a rule—With sequence numbers, if you need to move a rule to a different position within an ACL, you can add a second instance of the rule using the sequence number that positions it correctly, and then you can remove the original instance of the rule. This action allows you to move the rule without disrupting traffic.

If you enter a rule without a sequence number, the switch adds the rule to the end of the ACL and assigns a sequence number that is 10 greater than the sequence number of the preceding rule to the rule. For example, if the last rule in an ACL has a sequence number of 225 and you add a rule without a sequence number, the switch assigns the sequence number 235 to the new rule.

In addition, the Cisco Nexus 5000 Series switch allows you to reassign sequence numbers to rules in an ACL. Resequencing is useful when an ACL has rules numbered contiguously, such as 100 and 101, and you need to insert one or more rules between those rules.

Logical Operators and Logical Operation Units

IP ACL rules for TCP and UDP traffic can use logical operators to filter traffic based on port numbers.

The switch stores operator-operand couples in registers called logical operator units (LOUs).

LOU usage for the "eq" operator is never stored in an LOU. The range operation is inclusive of boundary values.

The following guidelines determine when the switch stores operator-operand couples in LOUs:

- If the operator or operand differs from other operator-operand couples that are used in other rules, the couple is stored in an LOU.
For example, the operator-operand couples "gt 10" and "gt 11" would be stored separately in half an LOU each. The couples "gt 10" and "lt 10" would also be stored separately.
- Whether the operator-operand couple is applied to a source port or a destination port in the rule affects LOU usage. Identical couples are stored separately when one of the identical couples is applied to a source port and the other couple is applied to a destination port.
For example, if a rule applies the operator-operand couple "gt 10" to a source port and another rule applies a "gt 10" couple to a destination port, both couples would also be stored in half an LOU, resulting

in the use of one whole LOU. Any additional rules using a "gt 10" couple would not result in further LOU usage.

Configuring IP ACLs

Creating an IP ACL

You can create an IPv4 or IPv6 ACL on the switch and add rules to it.

SUMMARY STEPS

1. switch# **configure terminal**
2. switch(config)# **{ip | ipv6 } access-list name**
3. switch(config-acl)# *[sequence-number] {permit|deny} protocol source destination*
4. (Optional) switch(config-acl)# **statistics**
5. (Optional) switch# **show {ip | ipv6} access-lists name**
6. (Optional) switch# **copy running-config startup-config**

DETAILED STEPS

	Command or Action	Purpose
Step 1	switch# configure terminal	Enters configuration mode.
Step 2	switch(config)# {ip ipv6 } access-list name	Creates the IP ACL and enters IP ACL configuration mode. The <i>name</i> argument can be up to 64 characters.
Step 3	switch(config-acl)# <i>[sequence-number] {permit deny} protocol source destination</i>	Creates a rule in the IP ACL. You can create many rules. The <i>sequence-number</i> argument can be a whole number between 1 and 4294967295. The permit and deny commands support many ways of identifying traffic. For more information, see the <i>Cisco Nexus 5000 Series Command Reference</i> .
Step 4	switch(config-acl)# statistics	(Optional) Specifies that the switch maintains global statistics for packets matching the rules in the ACL.
Step 5	switch# show {ip ipv6} access-lists name	(Optional) Displays the IP ACL configuration.
Step 6	switch# copy running-config startup-config	(Optional) Copies the running configuration to the startup configuration.

The following example shows how to create an IPv4 ACL:

```
switch# configure terminal
switch(config)# ip access-list acl-01
switch(config-acl)# permit ip 192.168.2.0/24 any
switch(config-acl)# statistics
```

The following example shows how to create an IPv6 ACL:

```
switch# configure terminal
switch(config)# ipv6 access-list acl-01-ipv6
switch(config-ipv6-acl)# permit tcp 2001:0db8:85a3::/48 2001:0db8:be03:2112::/64
```

Changing an IP ACL

You can add and remove rules in an existing IPv4 or IPv6 ACL. You cannot change existing rules. Instead, to change a rule, you can remove it and recreate it with the desired changes.

If you need to add more rules between existing rules than the current sequence numbering allows, you can use the **resequence** command to reassign sequence numbers.

SUMMARY STEPS

1. switch# **configure terminal**
2. switch(config)# **{ip | ipv6} access-list name**
3. switch(config-acl)# **[sequence-number] {permit | deny} protocol source destination**
4. (Optional) switch(config-acl)# **no {sequence-number | {permit | deny} protocol source destination}**
5. (Optional) switch(config-acl)# **[no] statistics**
6. (Optional) switch# **show {ip | ipv6} access-lists name**
7. (Optional) switch# **copy running-config startup-config**

DETAILED STEPS

	Command or Action	Purpose
Step 1	switch# configure terminal	Enters configuration mode.
Step 2	switch(config)# {ip ipv6} access-list name	Enters IP ACL configuration mode for the ACL that you specify by name.
Step 3	switch(config-acl)# [sequence-number] {permit deny} protocol source destination	Creates a rule in the IP ACL. Using a sequence number allows you to specify a position for the rule in the ACL. Without a sequence number, the rule is added to the end of the rules. The <i>sequence-number</i> argument can be a whole number between 1 and 4294967295. The permit and deny commands support many ways of identifying traffic. For more information, see the <i>Cisco Nexus 5000 Series Command Reference</i> .
Step 4	switch(config-acl)# no {sequence-number {permit deny} protocol source destination}	(Optional) Removes the rule that you specified from the IP ACL. The permit and deny commands support many ways of identifying traffic. For more information, see the <i>Cisco Nexus 5000 Series Command Reference</i> .

	Command or Action	Purpose
Step 5	switch(config-acl)# [no] statistics	(Optional) Specifies that the switch maintains global statistics for packets matching the rules in the ACL. The no option stops the switch from maintaining global statistics for the ACL.
Step 6	switch# show {ip ipv6} access-lists name	(Optional) Displays the IP ACL configuration.
Step 7	switch# copy running-config startup-config	(Optional) Copies the running configuration to the startup configuration.

Related Topics

[Changing Sequence Numbers in an IP ACL, on page 314](#)

Removing an IP ACL

You can remove an IP ACL from the switch.

Before you remove an IP ACL from the switch, be sure that you know whether the ACL is applied to an interface. The switch allows you to remove ACLs that are currently applied. Removing an ACL does not affect the configuration of interfaces where you have applied the ACL. Instead, the switch considers the removed ACL to be empty.

To remove an IP ACL from the switch, perform this task:

SUMMARY STEPS

1. switch# **configure terminal**
2. switch(config)# no {ip | ipv6} **access-list name**
3. (Optional) switch# **show running-config**
4. (Optional) switch# **copy running-config startup-config**

DETAILED STEPS

	Command or Action	Purpose
Step 1	switch# configure terminal	Enters configuration mode.
Step 2	switch(config)# no {ip ipv6} access-list name	Removes the IP ACL that you specified by name from the running configuration.

	Command or Action	Purpose
Step 3	switch# show running-config	(Optional) Displays ACL configuration. The removed IP ACL should not appear.
Step 4	switch# copy running-config startup-config	(Optional) Copies the running configuration to the startup configuration.

Changing Sequence Numbers in an IP ACL

You can change all the sequence numbers assigned to the rules in an IP ACL. To change sequence numbers, perform this task:

SUMMARY STEPS

1. switch# **configure terminal**
2. switch(config)# **resequence {ip | ipv6} access-list name starting-sequence-number increment**
3. (Optional) switch# **show {ip | ipv6} access-lists name**
4. (Optional) switch# **copy running-config startup-config**

DETAILED STEPS

	Command or Action	Purpose
Step 1	switch# configure terminal	Enters configuration mode.
Step 2	switch(config)# resequence {ip ipv6} access-list name starting-sequence-number increment	Assigns sequence numbers to the rules contained in the ACL, where the first rule receives the starting sequence number that you specify. Each subsequent rule receives a number larger than the preceding rule. The difference in numbers is determined by the increment that you specify. The <i>starting-sequence-number</i> argument and the <i>increment</i> argument can be a whole number between 1 and 4294967295.
Step 3	switch# show {ip ipv6} access-lists name	(Optional) Displays the IP ACL configuration.
Step 4	switch# copy running-config startup-config	(Optional) Copies the running configuration to the startup configuration.

Applying an IP ACL as a Port ACL

You can apply an IPv4 or IPv6 ACL to a physical Ethernet interface or a EtherChannel. ACLs applied to these interface types are considered port ACLs.



Note Some configuration parameters when applied to an EtherChannel are not reflected on the configuration of the member ports.

SUMMARY STEPS

1. switch# **configure terminal**
2. switch(config)# **interface** {**ethernet** [*chassis/*]*slot/port* | **port-channel** *channel-number*}
3. switch(config-if)# {**ip port access-group** | **ipv6 port traffic-filter**} *access-list in*
4. (Optional) switch# **show running-config**
5. (Optional) switch# **copy running-config startup-config**

DETAILED STEPS

	Command or Action	Purpose
Step 1	switch# configure terminal	Enters configuration mode.
Step 2	switch(config)# interface { ethernet [<i>chassis/</i>] <i>slot/port</i> port-channel <i>channel-number</i> }	Enters interface configuration mode for the specified interface.
Step 3	switch(config-if)# { ip port access-group ipv6 port traffic-filter } <i>access-list in</i>	Applies an IPv4 or IPv6 ACL to the interface or EtherChannel. Only inbound filtering is supported with port ACLs. You can apply one port ACL to an interface.
Step 4	switch# show running-config	(Optional) Displays ACL configuration.
Step 5	switch# copy running-config startup-config	(Optional) Copies the running configuration to the startup configuration.

Verifying IP ACL Configurations

To display IP ACL configuration information, perform one of the following tasks:

- switch# **show running-config**
Displays ACL configuration, including IP ACL configuration and interfaces that IP ACLs are applied to.
- switch# **show running-config interface**
Displays the configuration of an interface to which you have applied an ACL.

For detailed information about the fields in the output from these commands, refer to the *Cisco Nexus 5000 Series Command Reference*.

Displaying and Clearing IP ACL Statistics

Use the **show ip access-lists** and **show ipv6 access-list** commands to display statistics about an IP ACL, including the number of packets that have matched each rule. For detailed information about the fields in the output from this command, refer to the *Cisco Nexus 5000 Series Command Reference*.



Note The mac access-list is applicable to non-IPv4 and non-IPv6 traffic only.

- switch# **show {ip | ipv6} access-lists name**
Displays IP ACL configuration. If the IP ACL includes the **statistics** command, then the **show ip access-lists** and **show ipv6 access-list** command output includes the number of packets that have matched each rule.
- switch# **clear {ip | ipv6} access-list counters [access-list-name]**
Clears statistics for all IP ACLs or for a specific IP ACL.

Configuring MAC ACLs

Creating a MAC ACL

To create a MAC ACL and add rules to it, perform this task:

SUMMARY STEPS

1. switch# **configure terminal**
2. switch# **mac access-list name**
3. switch(config-mac-acl)# **[sequence-number] {permit | deny} source destination protocol**
4. (Optional) switch(config-mac-acl)# **statistics**
5. (Optional) switch# **show mac access-lists name**
6. (Optional) switch# **copy running-config startup-config**

DETAILED STEPS

	Command or Action	Purpose
Step 1	switch# configure terminal	Enters configuration mode.
Step 2	switch# mac access-list name	Creates the MAC ACL and enters ACL configuration mode.
Step 3	switch(config-mac-acl)# [sequence-number] {permit deny} source destination protocol	Creates a rule in the MAC ACL. The permit and deny options support many ways of identifying traffic. For more information, see the <i>Cisco Nexus 5000 Series Command Reference</i> .

	Command or Action	Purpose
Step 4	switch(config-mac-acl)# statistics	(Optional) Specifies that the switch maintains global statistics for packets matching the rules in the ACL.
Step 5	switch# show mac access-lists name	(Optional) Displays the MAC ACL configuration.
Step 6	switch# copy running-config startup-config	(Optional) Copies the running configuration to the startup configuration.

The following example shows how to create a MAC ACL and add rules to it:

```
switch# configure terminal
switch(config)# mac access-list acl-mac-01
switch(config-mac-acl)# permit 00c0.4f00.0000 0000.00ff.ffff any
switch(config-mac-acl)# statistics
```

Changing a MAC ACL

In an existing MAC ACL, you can add and remove rules. You cannot change existing rules. Instead, to change a rule, you can remove it and recreate it with the desired changes.

If you need to add more rules between existing rules than the current sequence numbering allows, you can use the **resequence** command to reassign sequence numbers.

To change a MAC ACL, perform this task:

SUMMARY STEPS

1. switch# **configure terminal**
2. switch(config)# **mac access-list name**
3. switch(config-mac-acl)# [*sequence-number*] {**permit** | **deny**} *source destination protocol*
4. (Optional) switch(config-mac-acl)# **no** {*sequence-number* | {**permit**|**deny**} *source destination protocol*}
5. (Optional) switch(config-mac-acl)# [**no**] **statistics**
6. (Optional) switch# **show mac access-lists name**
7. (Optional) switch# **copy running-config startup-config**

DETAILED STEPS

	Command or Action	Purpose
Step 1	switch# configure terminal	Enters configuration mode.
Step 2	switch(config)# mac access-list name	Enters ACL configuration mode for the ACL that you specify by name.

	Command or Action	Purpose
Step 3	switch(config-mac-acl)# [<i>sequence-number</i>] { permit deny } <i>source destination protocol</i>	Creates a rule in the MAC ACL. Using a sequence number allows you to specify a position for the rule in the ACL. Without a sequence number, the rule is added to the end of the rules. The permit and deny commands support many ways of identifying traffic.
Step 4	switch(config-mac-acl)# no { <i>sequence-number</i> { permit deny } <i>source destination protocol</i> }	(Optional) Removes the rule that you specify from the MAC ACL. The permit and deny commands support many ways of identifying traffic.
Step 5	switch(config-mac-acl)# [no] statistics	(Optional) Specifies that the switch maintains global statistics for packets matching the rules in the ACL. The no option stops the switch from maintaining global statistics for the ACL.
Step 6	switch# show mac access-lists <i>name</i>	(Optional) Displays the MAC ACL configuration.
Step 7	switch# copy running-config startup-config	(Optional) Copies the running configuration to the startup configuration.

The following example shows how to change a MAC ACL:

```
switch# configure terminal
switch(config)# mac access-list acl-mac-01
switch(config-mac-acl)# 100 permit mac 00c0.4f00.00 0000.00ff.ffff any
switch(config-mac-acl)# statistics
```

Removing a MAC ACL

You can remove a MAC ACL from the switch.

Be sure that you know whether the ACL is applied to an interface. The switch allows you to remove ACLs that are current applied. Removing an ACL does not affect the configuration of interfaces where you have applied the ACL. Instead, the switch considers the removed ACL to be empty.

SUMMARY STEPS

1. switch# **configure terminal**
2. switch(config)# **no mac access-list** *name*
3. (Optional) switch# **show mac access-lists**
4. (Optional) switch# **copy running-config startup-config**

DETAILED STEPS

	Command or Action	Purpose
Step 1	switch# configure terminal	Enters configuration mode.
Step 2	switch(config)# no mac access-list <i>name</i>	Removes the MAC ACL that you specify by name from the running configuration.
Step 3	switch# show mac access-lists	(Optional) Displays the MAC ACL configuration.
Step 4	switch# copy running-config startup-config	(Optional) Copies the running configuration to the startup configuration.

Changing Sequence Numbers in a MAC ACL

You can change all the sequence numbers assigned to rules in a MAC ACL. Resequencing is useful when you need to insert rules into an ACL and there are not enough available sequence numbers.

To change all the sequence numbers assigned to rules in a MAC ACL, perform this task:

SUMMARY STEPS

1. switch# **configure terminal**
2. switch(config)# **resequence mac access-list** *name starting-sequence-number increment*
3. (Optional) switch# **show mac access-lists** *name*
4. (Optional) switch# **copy running-config startup-config**

DETAILED STEPS

	Command or Action	Purpose
Step 1	switch# configure terminal	Enters configuration mode.
Step 2	switch(config)# resequence mac access-list <i>name starting-sequence-number increment</i>	Assigns sequence numbers to the rules contained in the ACL, where the first rule receives the number specified by the starting-sequence number that you specify. Each subsequent rule receives a number larger than the preceding rule. The difference in numbers is determined by the increment number that you specify.
Step 3	switch# show mac access-lists <i>name</i>	(Optional) Displays the MAC ACL configuration.
Step 4	switch# copy running-config startup-config	(Optional) Copies the running configuration to the startup configuration.

Related Topics[Rules, on page 308](#)**Applying a MAC ACL as a Port ACL**

You can apply a MAC ACL as a port ACL to any of the following interface types:

- Ethernet interfaces
- EtherChannel interfaces

Be sure that the ACL that you want to apply exists and is configured to filter traffic as necessary for this application.

**Note**

Some configuration parameters when applied to an EtherChannel are not reflected on the configuration of the member ports.

SUMMARY STEPS

1. switch# **configure terminal**
2. switch(config)# **interface** {**ethernet** [*chassis*]/*slot/port* | **port-channel** *channel-number*}
3. switch(config-if)# **mac port access-group** *access-list*
4. (Optional) switch# **show running-config**
5. (Optional) switch# **copy running-config startup-config**

DETAILED STEPS

	Command or Action	Purpose
Step 1	switch# configure terminal	Enters configuration mode.
Step 2	switch(config)# interface { ethernet [<i>chassis</i>]/ <i>slot/port</i> port-channel <i>channel-number</i> }	Enters interface configuration mode for the Ethernet specified interface.
Step 3	switch(config-if)# mac port access-group <i>access-list</i>	Applies a MAC ACL to the interface.
Step 4	switch# show running-config	(Optional) Displays ACL configuration.
Step 5	switch# copy running-config startup-config	(Optional) Copies the running configuration to the startup configuration.

Related Topics[Creating an IP ACL, on page 311](#)

Verifying MAC ACL Configurations

To display MAC ACL configuration information, perform one of the following tasks:

- **switch# show mac access-lists**
Displays the MAC ACL configuration
- **switch# show running-config**
Displays ACL configuration, including MAC ACLs and the interfaces that ACLs are applied to.
- **switch# show running-config interface**
Displays the configuration of the interface to which you applied the ACL.

Displaying and Clearing MAC ACL Statistics

Use the **show mac access-lists** command to display statistics about a MAC ACL, including the number of packets that have matched each rule.

- **switch# show mac access-lists**
Displays MAC ACL configuration. If the MAC ACL includes the **statistics** command, the **show mac access-lists** command output includes the number of packets that have matched each rule.
- **switch# clear mac access-list counters**
Clears statistics for all MAC ACLs or for a specific MAC ACL.

Example Configuration for MAC ACLs

This example shows how to create a MAC ACL named `acl-mac-01` and apply it to Ethernet interface 1/1:

```
switch# configure terminal
switch(config)# mac access-list acl-mac-01
switch(config-mac-acl)# permit 00c0.4f00.0000 0000.00ff.ffff any
switch(config-mac-acl)# exit
switch(config)# interface ethernet 1/1
switch(config-if)# mac access-group acl-mac-01
```

Information About VLAN ACLs

A VLAN ACL (VACL) is one application of a MAC ACL or IP ACL. You can configure VACLs to apply to all packets that are bridged within a VLAN. VACLs are used strictly for security packet filtering. VACLs are not defined by direction (ingress or egress).

VACLs and Access Maps

VACLs use access maps to link an IP ACL or a MAC ACL to an action. The switch takes the configured action on packets permitted by the VACL.

VACLs and Actions

In access map configuration mode, you use the **action** command to specify one of the following actions:

- Forward—Sends the traffic to the destination determined by normal operation of the switch.
- Drop—Drops the traffic.

Statistics

The switch can maintain global statistics for each rule in a VACL. If a VACL is applied to multiple VLANs, the maintained rule statistics are the sum of packet matches (hits) on all the interfaces on which that VACL is applied.



Note

The Cisco Nexus 5000 Series switch does not support interface-level VACL statistics.

For each VLAN access map that you configure, you can specify whether the switch maintains statistics for that VACL. This allows you to turn VACL statistics on or off as needed to monitor traffic filtered by a VACL or to help troubleshoot VLAN access-map configuration.

Configuring VACLs

Creating or Changing a VACL

You can create or change a VACL. Creating a VACL includes creating an access map that associates an IP ACL or MAC ACL with an action to be applied to the matching traffic.

To create or change a VACL, perform this task:

SUMMARY STEPS

1. switch# **configure terminal**
2. switch(config)# **vlan access-map** *map-name*
3. switch(config-access-map)# **match ip address** *ip-access-list*
4. switch(config-access-map)# **match mac address** *mac-access-list*
5. switch(config-access-map)# **action** {drop | forward}
6. (Optional) switch(config-access-map)# [**no**] **statistics**
7. (Optional) switch(config-access-map)# **show running-config**
8. (Optional) switch(config-access-map)# **copy running-config startup-config**

DETAILED STEPS

	Command or Action	Purpose
Step 1	switch# configure terminal	Enters configuration mode.
Step 2	switch(config)# vlan access-map <i>map-name</i>	Enters access map configuration mode for the access map specified.

	Command or Action	Purpose
Step 3	switch(config-access-map)# match ip address <i>ip-access-list</i>	Specifies an IPv4 and IPV6 ACL for the map.
Step 4	switch(config-access-map)# match mac address <i>mac-access-list</i>	Specifies a MAC ACL for the map.
Step 5	switch(config-access-map)# action {drop forward}	Specifies the action that the switch applies to traffic that matches the ACL.
Step 6	switch(config-access-map)# [no] statistics	(Optional) Specifies that the switch maintains global statistics for packets matching the rules in the VACL. The no option stops the switch from maintaining global statistics for the VACL.
Step 7	switch(config-access-map)# show running-config	(Optional) Displays ACL configuration.
Step 8	switch(config-access-map)# copy running-config startup-config	(Optional) Copies the running configuration to the startup configuration.

Removing a VACL

You can remove a VACL, which means that you will delete the VLAN access map.

Be sure that you know whether the VACL is applied to a VLAN. The switch allows you to remove VACLs that are current applied. Removing a VACL does not affect the configuration of VLANs where you have applied the VACL. Instead, the switch considers the removed VACL to be empty.

SUMMARY STEPS

1. switch# **configure terminal**
2. switch(config)# **no vlan access-map** *map-name*
3. (Optional) switch(config)# **show running-config**
4. (Optional) switch(config)# **copy running-config startup-config**

DETAILED STEPS

	Command or Action	Purpose
Step 1	switch# configure terminal	Enters configuration mode.
Step 2	switch(config)# no vlan access-map <i>map-name</i>	Removes the VLAN access map configuration for the specified access map.

	Command or Action	Purpose
Step 3	switch(config)# show running-config	(Optional) Displays ACL configuration.
Step 4	switch(config)# copy running-config startup-config	(Optional) Copies the running configuration to the startup configuration.

Applying a VACL to a VLAN

You can apply a VACL to a VLAN.

SUMMARY STEPS

1. switch# **configure terminal**
2. switch(config)# [**no**] **vlan filter** *map-name* **vlan-list** *list*
3. (Optional) switch(config)# **show running-config**
4. (Optional) switch(config)# **copy running-config startup-config**

DETAILED STEPS

	Command or Action	Purpose
Step 1	switch# configure terminal	Enters configuration mode.
Step 2	switch(config)# [no] vlan filter <i>map-name</i> vlan-list <i>list</i>	Applies the VACL to the VLANs by the list that you specified. The no option unapplies the VACL. The vlan-list command can specify a list of up to 32 VLANs, but multiple vlan-list commands can be configured to cover more than 32 VLANs.
Step 3	switch(config)# show running-config	(Optional) Displays ACL configuration.
Step 4	switch(config)# copy running-config startup-config	(Optional) Copies the running configuration to the startup configuration.

Verifying VACL Configuration

To display VACL configuration information, perform one of the following tasks:

- switch# **show running-config aclmgr**
Displays ACL configuration, including VACL-related configuration.

- **switch# show vlan filter**
Displays information about VACLs that are applied to a VLAN.
- **switch# show vlan access-map**
Displays information about VLAN access maps.

Displaying and Clearing VACL Statistics

To display or clear VACL statistics, perform one of the following tasks:

- **switch# show vlan access-list**
Displays VACL configuration. If the VLAN access-map includes the **statistics** command, then the **show vlan access-list** command output includes the number of packets that have matched each rule.
- **switch# clear vlan access-list counters**
Clears statistics for all VACLs or for a specific VACL.

Example Configuration for VACL

This example shows how to configure a VACL to forward traffic permitted by an IP ACL named `acl-ip-01` and how to apply the VACL to VLANs 50 through 82:

```
switch# configure terminal
switch(config)# vlan access-map acl-ip-map
switch(config-access-map)# match ip address acl-ip-01
switch(config-access-map)# action forward
switch(config-access-map)# exit
switch(config)# vlan filter acl-ip-map vlan-list 50-82
```

Default ACL Settings

The following table lists the default settings for IP ACLs parameters.

Table 29: Default IP ACLs Parameters

Parameters	Default
IP ACLs	No IP ACLs exist by default.
ACL rules	Implicit rules apply to all ACLs .

The following table lists the default settings for MAC ACLs parameters.

Table 30: Default MAC ACLs Parameters

Parameters	Default
MAC ACLs	No MAC ACLs exist by default.

Parameters	Default
ACL rules	Implicit rules apply to all ACLs .

The following table lists the default settings for VACL parameters.

Table 31: Default VACL Parameters

Parameters	Default
VACLs	No IP ACLs exist by default.
ACL rules	Implicit rules apply to all ACLs.



PART **IV**

System Management

- [Using Cisco Fabric Services, page 329](#)
- [Configuring User Accounts and RBAC, page 347](#)
- [Configuring Session Manager, page 359](#)
- [Configuring Online Diagnostics, page 363](#)
- [Configuring System Message Logging, page 369](#)
- [Configuring Smart Call Home, page 385](#)
- [Configuring SNMP, page 415](#)
- [Configuring RMON, page 429](#)



Using Cisco Fabric Services

This chapter contains the following sections:

- [Using Cisco Fabric Services, page 329](#)

Using Cisco Fabric Services

Cisco Nexus 5000 Series switches provide Cisco Fabric Services (CFS) capability, which simplifies provisioning by automatically distributing configuration information to all switches in the network. Switch features can use the CFS infrastructure to distribute feature data or configuration data required by the feature.

Information About CFS

Some features in the Cisco Nexus 5000 Series switch require configuration synchronization with other switches in the network to function correctly. Synchronization through manual configuration at each switch in the network can be a tedious and error-prone process.

Cisco Fabric Services (CFS) provides a common infrastructure for automatic configuration synchronization in the network. It provides the transport function and a set of common services to the features. CFS has the ability to discover CFS capable switches in the network and discovering feature capabilities in all CFS capable switches.

Cisco Nexus 5000 Series switches support CFS message distribution over Fibre Channel, IPv4 or IPv6 networks. If the switch is provisioned with Fibre Channel ports, CFS over Fibre Channel is enabled by default. CFS over IP must be explicitly enabled.

CFS provides the following features:

- Peer-to-peer protocol with no client-server relationship at the CFS layer.
- CFS message distribution over Fibre Channel, IPv4 or IPv6 networks.
- Three modes of distribution.
 - Coordinated distributions: Only one distribution is allowed in the network at any given time.
 - Uncoordinated distributions: Multiple parallel distributions are allowed in the network except when a coordinated distribution is in progress.

- Unrestricted uncoordinated distributions: Multiple parallel distributions are allowed in the network in the presence of an existing coordinated distribution. Unrestricted uncoordinated distributions are allowed to run in parallel with all other types of distributions.

The following features are supported for CFS distribution over IP:

- One scope of distribution over an IP network:
 - Physical scope: The distribution spans the entire IP network.

The following features are supported for CFS distribution over Fibre Channel SANs:

- Three scopes of distribution over SAN fabrics.
 - Logical scope: The distribution occurs within the scope of a VSAN.
 - Physical scope: The distribution spans the entire physical topology.
 - Over a selected set of VSANs: Some features require configuration distribution over some specific VSANs. These features can specify to CFS the set of VSANs over which to restrict the distribution.
- Supports a merge protocol that facilitates the merge of feature configuration during a fabric merge event (when two independent SAN fabrics merge).

CFS Distribution

The CFS distribution functionality is independent of the lower layer transport. Cisco Nexus 5000 Series switches support CFS distribution over IP and CFS distribution over Fibre Channel. Features that use CFS are unaware of the lower layer transport.

CFS Distribution Modes

CFS supports three distribution modes to accommodate different feature requirements:

- Uncoordinated Distribution
- Coordinated Distribution
- Unrestricted Uncoordinated Distributions

Only one mode is allowed at any given time.

Uncoordinated Distribution

Uncoordinated distributions are used to distribute information that is not expected to conflict with that from a peer. Parallel uncoordinated distributions are allowed for a feature.

Coordinated Distribution

Coordinated distributions allow only one feature distribution at a given time. CFS uses locks to enforce this. A coordinated distribution is not allowed to start if locks are taken for the feature anywhere in the network. A coordinated distribution consists of three stages:

- A network lock is acquired.
- The configuration is distributed and committed.

- The network lock is released.

Coordinated distribution has two variants:

- CFS driven —The stages are executed by CFS in response to a feature request without intervention from the feature.
- Feature driven—The stages are under the complete control of the feature.

Coordinated distributions are used to distribute information that can be manipulated and distributed from multiple switches, for example, the port security configuration.

Unrestricted Uncoordinated Distributions

Unrestricted uncoordinated distributions allow multiple parallel distributions in the network in the presence of an existing coordinated distribution. Unrestricted uncoordinated distributions are allowed to run in parallel with all other types of distributions.

Disabling or Enabling CFS Distribution on a Switch

If the switch is provisioned with Fibre Channel ports, CFS over Fibre Channel is enabled by default. CFS over IP must be explicitly enabled.

You can globally disable CFS on a switch to isolate the features using CFS from network-wide distributions while maintaining physical connectivity. When CFS is globally disabled on a switch, CFS operations are restricted to the switch.

SUMMARY STEPS

1. switch# **configure terminal**
2. switch(config)# **no cfs distribute**
3. (Optional) switch(config)# **cfs distribute**

DETAILED STEPS

	Command or Action	Purpose
Step 1	switch# configure terminal	Enters configuration mode.
Step 2	switch(config)# no cfs distribute	Globally disables CFS distribution (CFS over Fibre Channel or IP) for all applications on the switch.
Step 3	switch(config)# cfs distribute	(Optional) Enables CFS distribution on the switch. This is the default.

Verifying CFS Distribution Status

The `show cfs status` command displays the status of CFS distribution on the switch.

```
switch# show cfs status
Distribution : Enabled
Distribution over IP : Enabled - mode IPv4
IPv4 multicast address : 239.255.70.83
IPv6 multicast address : ff15::efff:4653
Distribution over Ethernet : Enabled
```

CFS Distribution over IP

CFS distribution over IP supports the following features:

- Physical distribution over an entirely IP network.
- Physical distribution over a hybrid Fibre Channel and IP network with the distribution reaching all switches that are reachable over either Fibre Channel or IP.



Note

The switch attempts to distribute information over Fibre Channel first and then over the IP network if the first attempt over Fibre Channel fails. CFS does not send duplicate messages if distribution over both IP and Fibre Channel is enabled.

- Distribution over IP version 4 (IPv4) or IP version 6 (IPv6).



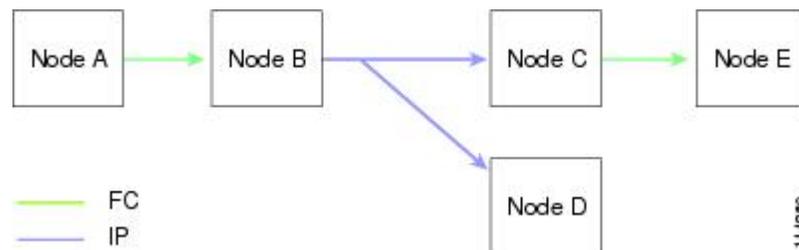
Note

CFS cannot distribute over both IPv4 and IPv6 from the same switch.

- Keepalive mechanism to detect network topology changes using a configurable multicast address.
- Compatibility with Cisco MDS 9000 Family switches running release 2.x or later.

The following figure (*Network Example 1*) shows a network with both Fibre Channel and IP connections. Node A forwards an event to node B over Fibre Channel. Node B forwards the event node C and node D using unicast IP. Node C forwards the event to node E using Fibre Channel.

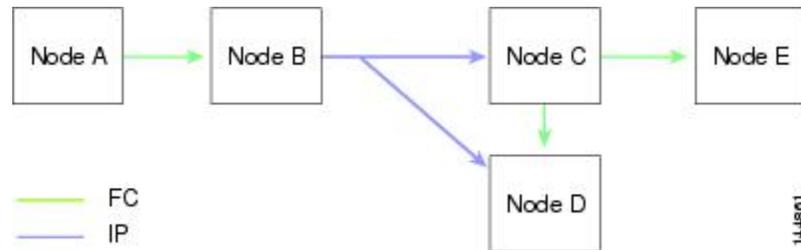
Figure 34: Network Example 1 with Fibre Channel and IP Connections



The following figure (*Network Example 2*) is the same as the previous figure except that node C and node D are connected using Fibre Channel. All processes is the same in this example because node B has node C and

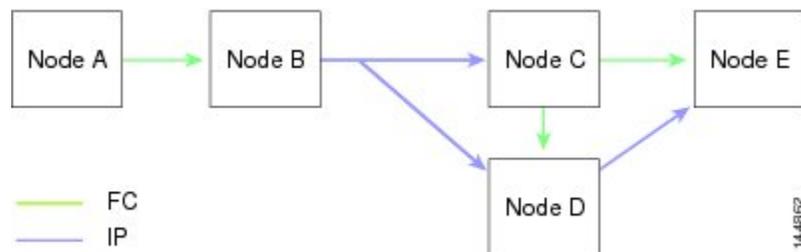
node D the distribution list for IP. Node C does not forward to node D because node D is already in the distribution list from node B.

Figure 35: Network Example 2 with Fibre Channel and IP Connections



The following figure (*Network Example 3*) is the same as the previous figure except that node D and node E are connected using IP. Both node C and node D forward the event to E because the node E is not in the distribution list from node B.

Figure 36: Network Example 3 with Fibre Channel and IP Connections



CFS Distribution over Fibre Channel

For FCS distribution over Fibre Channel, the CFS protocol layer resides on top of the FC2 layer. CFS uses the FC2 transport services to send information to other switches. CFS uses a proprietary SW_ILS (0x77434653) protocol for all CFS packets. CFS packets are sent to or from the switch domain controller addresses.

CFS Distribution Scopes

Different applications on the Cisco Nexus 5000 Series switches need to distribute the configuration at various levels. The following levels are available when using CFS distribution over Fibre Channel:

- VSAN level (logical scope)

Applications that operate within the scope of a VSAN have the configuration distribution restricted to the VSAN. An example application is port security where the configuration database is applicable only within a VSAN.



Note Logical scope is not supported for FCS distribution over IP.

- Physical topology level (physical scope)

Some applications (such as NTP) need to distribute the configuration to the entire physical topology.

- Between two selected switches

Some applications operate only between selected switches in the network.

CFS Merge Support

CFS Merge is supported for CFS distribution over Fibre Channel.

An application keeps the configuration synchronized in a SAN fabric through CFS. Two such fabrics might merge as a result of an ISL coming up between them. These two fabrics could have two different sets of configuration information that need to be reconciled in the event of a merge. CFS provides notification each time an application peer comes online. If a fabric with M application peers merges with another fabric with N application peers, and if an application triggers a merge action on every notification, a link-up event results in M×N merges in the fabric.

CFS supports a protocol that reduces the number of merges required to one by handling the complexity of the merge at the CFS layer. This protocol runs per application per scope. The protocol involves selecting one switch in a fabric as the merge manager for that fabric. The other switches do not have a role in the merge process.

During a merge, the merge manager in the two fabrics exchange their configuration databases with each other. The application on one of them merges the information, decides if the merge is successful, and informs all switches in the combined fabric of the status of the merge.

In case of a successful merge, the merged database is distributed to all switches in the combined fabric and the entire new fabric remains in a consistent state. You can recover from a merge failure by starting a distribution from any of the switches in the new fabric. This distribution restores all peers in the fabric to the same configuration database.

CFS Support for Applications

CFS Application Requirements

All switches in the network must be CFS capable. Switches that are not CFS capable do not receive distributions and result in part of the network not receiving the intended distribution. CFS has the following requirements:

- Implicit CFS usage—The first time you issue a CFS task for a CFS-enabled application, the configuration modification process begins and the application locks the network.
- Pending database—The pending database is a temporary buffer to hold uncommitted information. The uncommitted changes are not applied immediately to ensure that the database is synchronized with the database in the other switches in the network. When you commit the changes, the pending database overwrites the configuration database (also known as the active database or the effective database).
- CFS distribution enabled or disabled on a per-application basis—The default (enable or disable) for CFS distribution state differs between applications. If CFS distribution is disabled for an application, then that application does not distribute any configuration nor does it accept a distribution from other switches in the network.
- Explicit CFS commit—Most applications require an explicit commit operation to copy the changes in the temporary buffer to the application database, to distribute the new database to the network, and to release the network lock. The changes in the temporary buffer are not applied if you do not perform the commit operation.

Enabling CFS for an Application

All CFS-based applications provide an option to enable or disable the distribution capabilities.

Applications have the distribution enabled by default.

The application configuration is not distributed by CFS unless distribution is explicitly enabled for that application.

Verifying Application Registration Status

The **show cfs application** command displays the applications that are currently registered with CFS. The first column displays the application name. The second column indicates whether the application is enabled or disabled for distribution (enabled or disabled). The last column indicates the scope of distribution for the application (logical, physical, or both).



Note

The **show cfs application** command only displays applications registered with CFS. Conditional services that use CFS do not appear in the output unless these services are running.

```
switch# show cfs application
```

```
-----
Application      Enabled      Scope
-----
ntp              No          Physical-all
fscm             Yes         Physical-fc
rscn             No          Logical
fctimer         No          Physical-fc
syslogd         No          Physical-all
callhome        No          Physical-all
fcdomain        Yes         Logical
device-alias    Yes         Physical-fc
Total number of entries = 8
```

The **show cfs application name** command displays the details for a particular application. It displays the enabled/disabled state, timeout as registered with CFS, merge capability (if it has registered with CFS for merge support), and lastly the distribution scope.

```
switch# show cfs application name fscm
```

```
Enabled          : Yes
Timeout          : 100s
Merge Capable    : No
Scope           : Physical-fc
```

Locking the Network

When you configure (first time configuration) a feature (or application) that uses the CFS infrastructure, that feature starts a CFS session and locks the network. When a network is locked, the switch software allows configuration changes to this feature only from the switch holding the lock. If you make configuration changes to the feature from another switch, the switch issues a message to inform the user about the locked status. The configuration changes are held in a pending database by that application.

If you start a CFS session that requires a network lock but forget to end the session, an administrator can clear the session. If you lock a network at any time, your user name is remembered across restarts and switchovers. If another user (on the same machine) tries to perform configuration tasks, that user's attempts are rejected.

Verifying CFS Lock Status

The **show cfs lock** command displays all the locks that are currently acquired by any application. For each application the command displays the application name and scope of the lock taken. If the application lock is taken in the physical scope, then this command displays the switch WWN, IP address, user name, and user type of the lock holder. If the application is taken in the logical scope, then this command displays the VSAN in which the lock is taken, the domain, IP address, user name, and user type of the lock holder.

```
switch# show cfs lock

Application: ntp
Scope      : Physical
-----
Switch WWN          IP Address      User Name      User Type
-----
20:00:00:05:30:00:6b:9e  10.76.100.167  admin         CLI/SNMP v3
Total number of entries = 1

Application: port-security
Scope      : Logical
-----
VSAN   Domain   IP Address      User Name      User Type
-----
1      238      10.76.100.167  admin         CLI/SNMP v3
2      211      10.76.100.167  admin         CLI/SNMP v3
Total number of entries = 2
```

The **show cfs lock name** command displays the lock details for the specified application:

```
switch# show cfs lock name ntp
Scope      : Physical
-----
Switch WWN          IP Address      User Name      User Type
-----
20:00:00:05:30:00:6b:9e  10.76.100.167  admin         CLI/SNMP v3
Total number of entries = 1
```

Committing Changes

A commit operation saves the pending database for all application peers and releases the lock for all switches.

In general, the commit function does not start a session, only a lock function starts a session. However, an empty commit is allowed if configuration changes are not previously made. In this case, a commit operation results in a session that acquires locks and distributes the current database.

When you commit configuration changes to a feature using the CFS infrastructure, you receive a notification about one of the following responses:

- One or more external switches report a successful status—The application applies the changes locally and releases the network lock.
- None of the external switches report a successful state—The application considers this state a failure and does not apply the changes to any switch in the network. The network lock is not released.

You can commit changes for a specified feature by entering the **commit** command for that feature.

Discarding Changes

If you discard configuration changes, the application flushes the pending database and releases locks in the network. Both the abort and commit functions are only supported from the switch from which the network lock is acquired.

You can discard changes for a specified feature by using the **abort** command for that feature.

Saving the Configuration

Configuration changes that have not been applied yet (still in the pending database) are not shown in the running configuration. The configuration changes in the pending database overwrite the configuration in the effective database when you commit the changes.



Caution

If you do not commit the changes, they are not saved to the running configuration.

Clearing a Locked Session

You can clear locks held by an application from any switch in the network to recover from situations where locks are acquired and not released. This function requires Admin permissions.



Caution

Exercise caution when using this function to clear locks in the network. Any pending configurations in any switch in the network is flushed and lost.

CFS Regions

About CFS Regions

A CFS region is a user-defined subset of switches for a given feature or application in its physical distribution scope. When a network spans a vast geography, you may need to localize or restrict the distribution of certain profiles among a set of switches based on their physical proximity. CFS regions allow you to create multiple islands of distribution within the network for a given CFS feature or application. CFS regions are designed to restrict the distribution of a feature's configuration to a specific set or grouping of switches in a network.



Note

You can only configure a CFS region based on physical switches. You cannot configure a CFS region in a VSAN.

Example Scenario

The Call Home application triggers alerts to network administrators when a situation arises or something abnormal occurs. When the network covers many geographies, and there are multiple network administrators who are each responsible for a subset of switches in the network, the Call Home application sends alerts to all network administrators regardless of their location. For the Call Home application to send message alerts selectively to network administrators, the physical scope of the application has to be fine tuned or narrowed down. This is achieved by implementing CFS regions.

CFS regions are identified by numbers ranging from 0 through 200. Region 0 is reserved as the default region, and contains every switch in the network. You can configure regions from 1 through 200. The default region maintains backward compatibility.

If the feature is moved, that is, assigned to a new region, its scope is restricted to that region; it ignores all other regions for distribution or merging purposes. The assignment of the region to a feature has precedence in distribution over its initial physical scope.

You can configure a CFS region to distribute configurations for multiple features. However, on a given switch, you can configure only one CFS region at a time to distribute the configuration for a given feature. Once you assign a feature to a CFS region, its configuration cannot be distributed within another CFS region.

Managing CFS Regions

Creating CFS Regions

You can create a CFS region.

SUMMARY STEPS

1. switch# **configure terminal**
2. switch(config)# **cfs region** *region-id*

DETAILED STEPS

	Command or Action	Purpose
Step 1	switch# configure terminal	Enters configuration mode.
Step 2	switch(config)# cfs region <i>region-id</i>	Creates a region.

Assigning Applications to CFS Regions

You can assign an application on a switch to a region.

SUMMARY STEPS

1. switch# **configure terminal**
2. switch(config)# **cfs region** *region-id*
3. switch(config-cfs-region)# *application*

DETAILED STEPS

	Command or Action	Purpose
Step 1	switch# configure terminal	Enters configuration mode.
Step 2	switch(config)# cfs region <i>region-id</i>	Creates a region.
Step 3	switch(config-cfs-region)# <i>application</i>	Adds application(s) to the region. Note You can add any number of applications on the switch to a region. If you try adding an application to the same region more than once, you see the error message, "Application already present in the same region."

	Command or Action	Purpose
--	-------------------	---------

The following example shows how to assign applications to a region:

```
switch# configure terminal
switch(config)# cfs region 1
switch(config-cfs-region)# ntp
switch(config-cfs-region)# callhome
```

Moving an Application to a Different CFS Region

You can move an application from one region to another region.

SUMMARY STEPS

1. switch# **configure**
2. switch(config)# **cfs region region-id**
3. switch(config-cfs-region)# *application*

DETAILED STEPS

	Command or Action	Purpose
Step 1	switch# configure	Enters configuration mode.
Step 2	switch(config)# cfs region region-id	Enters CFS region configuration submenu.
Step 3	switch(config-cfs-region)# <i>application</i>	Indicates application(s) to be moved from one region into another. Note If you try moving an application to the same region more than once, you see the error message, "Application already present in the same region."

The following example shows how to move an application into Region 2 that was originally assigned to Region 1:

```
switch# configure terminal
switch(config)# cfs region 2
switch(config-cfs-region)# ntp
```

Removing an Application from a Region

Removing an application from a region is the same as moving the application back to the default region (Region 0). This brings the entire network into the scope of distribution for the application.

SUMMARY STEPS

1. switch# **configure**
2. switch(config)# **cfs region** *region-id*
3. switch(config-cfs-region)# **no application**

DETAILED STEPS

	Command or Action	Purpose
Step 1	switch# configure	Enters configuration mode.
Step 2	switch(config)# cfs region <i>region-id</i>	Enters CFS region configuration submode.
Step 3	switch(config-cfs-region)# no application	Removes application(s) that belong to the region.

Deleting CFS Regions

Deleting a region nullifies the region definition. All the applications bound by the region are released back to the default region.

SUMMARY STEPS

1. switch# **configure**
2. switch(config)# **no cfs region** *region-id*

DETAILED STEPS

	Command or Action	Purpose
Step 1	switch# configure	Enters configuration mode.
Step 2	switch(config)# no cfs region <i>region-id</i>	Deletes the region. Note You see the warning, "All the applications in the region will be moved to the default region."

Configuring CFS over IP**Enabling CFS over IPv4**

You can enable or disable CFS over IPv4.



Note CFS cannot distribute over both IPv4 and IPv6 from the same switch.

SUMMARY STEPS

1. switch# **configure**
2. switch(config)# **cfs ipv4 distribute**
3. (Optional) switch(config)# **no cfs ipv4 distribute**

DETAILED STEPS

	Command or Action	Purpose
Step 1	switch# configure	Enters configuration mode.
Step 2	switch(config)# cfs ipv4 distribute	Globally enables CFS over IPv6 for all applications on the switch.
Step 3	switch(config)# no cfs ipv4 distribute	(Optional) Disables (default) CFS over IPv6 on the switch.

Enabling CFS over IPv6

You can enable or disable CFS over IPv6.



Note CFS cannot distribute over both IPv4 and IPv6 from the same switch.

SUMMARY STEPS

1. switch# **configure**
2. switch(config)# **cfs ipv6 distribute**
3. (Optional) switch(config)# **no cfs ipv6 distribute**

DETAILED STEPS

	Command or Action	Purpose
Step 1	switch# configure	Enters configuration mode.
Step 2	switch(config)# cfs ipv6 distribute	Globally enables CFS over IPv6 for all applications on the switch.

	Command or Action	Purpose
Step 3	switch(config)# no cfs ipv6 distribute	(Optional) Disables (default) CFS over IPv6 on the switch.

Verifying the CFS Over IP Configuration

To verify the CFS over IP configuration, use the **show cfs status** command.

```
switch# show cfs status
Distribution : Enabled
Distribution over IP : Enabled - mode IPv4
IPv4 multicast address : 239.255.70.83
IPv6 multicast address : ff15::efff:4653
```

Configuring IP Multicast Address for CFS over IP

All CFS over IP enabled switches with similar multicast addresses form one CFS over IP network. CFS protocol-specific distributions, such as the keepalive mechanism for detecting network topology changes, use the IP multicast address to send and receive information.



Note CFS distributions for application data use directed unicast.

Configuring IPv4 Multicast Address for CFS

You can configure a CFS over IP multicast address value for IPv4. The default IPv4 multicast address is 239.255.70.83.

SUMMARY STEPS

1. switch# **configure**
2. switch(config)# **cfs ipv4 mcast-address ipv4-address**
3. (Optional) switch(config)# **no cfs ipv4 mcast-address ipv4-address**

DETAILED STEPS

	Command or Action	Purpose
Step 1	switch# configure	Enters configuration mode.
Step 2	switch(config)# cfs ipv4 mcast-address ipv4-address	Configures the IPv4 multicast address for CFS distribution over IPv4. The ranges of valid IPv4 addresses are 239.255.0.0 through 239.255.255.255 and 239.192/16 through 239.251/16.

	Command or Action	Purpose
Step 3	switch(config)# no cfs ipv4 mcast-address <i>ipv4-address</i>	(Optional) Reverts to the default IPv4 multicast address for CFS distribution over IPv4. The default IPv4 multicast address for CFS is 239.255.70.83.

Configuring IPv6 Multicast Address for CFS

You can configure a CFS over IP multicast address value for IPv6. The default IPv6 multicast address is ff13:7743:4653.

SUMMARY STEPS

1. switch# **configure**
2. switch(config)# **cfs ipv6 mcast-address** *ipv4-address*
3. (Optional) switch(config)# **no cfs ipv6 mcast-address** *ipv4-address*

DETAILED STEPS

	Command or Action	Purpose
Step 1	switch# configure	Enters configuration mode.
Step 2	switch(config)# cfs ipv6 mcast-address <i>ipv4-address</i>	Configures the IPv6 multicast address for CFS distribution over IPv6. The range of valid IPv6 addresses is ff15::/16 (ff15::0000:0000 through ff15::ffff:ffff) and ff18::/16 (ff18::0000:0000 through ff18::ffff:ffff).
Step 3	switch(config)# no cfs ipv6 mcast-address <i>ipv4-address</i>	(Optional) Reverts to the default IPv6 multicast address for CFS distribution over IPv6. The default IPv6 multicast address for CFS over IP is ff15::eff:4653.

Verifying IP Multicast Address Configuration for CFS over IP

To verify the IP multicast address configuration for CFS over IP, use the **show cfs status** command:

```
switch# show cfs status
Fabric distribution Enabled
IP distribution Enabled mode ipv4
IPv4 multicast address : 10.1.10.100
IPv6 multicast address : ff13::e244:4754
```

Displaying CFS Distribution Information

The **show cfs merge status name** command displays the merge status for a given application. The following example displays the output for an application distributing in logical scope. It shows the merge status in all

valid VSANs on the switch. The command output shows the merge status as one of the following: Success, Waiting, or Failure or In Progress. In case of a successful merge, all the switches in the network are shown under the local network. In case of a merge failure or a merge in progress, the local network and the remote network involved in the merge are indicated separately. The application server in each network that is mainly responsible for the merge is indicated by the term Merge Master.

```
switch# show cfs merge status name port-security

Logical [VSAN 1] Merge Status: Failed
Local Fabric
-----
Domain Switch WWN          IP Address
-----
238    20:00:00:05:30:00:6b:9e  10.76.100.167    [Merge Master]

Remote Fabric
-----
Domain Switch WWN          IP Address
-----
236    20:00:00:0e:d7:00:3c:9e    10.76.100.169    [Merge Master]

Logical [VSAN 2] Merge Status: Success
Local Fabric
-----
Domain Switch WWN          IP Address
-----
211    20:00:00:05:30:00:6b:9e    10.76.100.167    [Merge Master]
1      20:00:00:0e:d7:00:3c:9e    10.76.100.169

Logical [VSAN 3] Merge Status: Success
Local Fabric
-----
Domain Switch WWN          IP Address
-----
221    20:00:00:05:30:00:6b:9e    10.76.100.167    [Merge Master]
103    20:00:00:0e:d7:00:3c:9e    10.76.100.169
```

The following example of the **show cfs merge status name** command output displays an application using the physical scope with a merge failure. The command uses the specified application name to display the merge status based on the application scope.

```
switch# show cfs merge status name ntp

Physical Merge Status: Failed
Local Fabric
-----
Switch WWN          IP Address
-----
20:00:00:05:30:00:6b:9e  10.76.100.167    [Merge Master]

Remote Fabric
-----
Switch WWN          IP Address
-----
20:00:00:0e:d7:00:3c:9e  10.76.100.169    [Merge Master]
```

The **show cfs peers** command output displays all the switches in the physical network in terms of the switch WWN and the IP address. The local switch is indicated as Local.

```
switch# show cfs peers

Physical Fabric
-----
Switch WWN          IP Address
-----
20:00:00:05:30:00:6b:9e  10.76.100.167    [Local]
20:00:00:0e:d7:00:3c:9e  10.76.100.169

Total number of entries = 2
```

The **show cfs peers name** command displays all the peers for which a particular application is registered with CFS. The command output shows all the peers for the physical scope or for each of the valid VSANs on the switch, depending on the application scope. For physical scope, the switch WWNs for all the peers are indicated. The local switch is indicated as Local.

```
switch# show cfs peers name ntp
Scope      : Physical
-----
Switch WWN          IP Address
-----
20:00:00:44:22:00:4a:9e  172.22.92.27   [Local]
20:00:00:05:30:01:1b:c2  172.22.92.215
```

The following example **show cfs peers name** command output displays all the application peers (all switches in which that application is registered). The local switch is indicated as Local.

```
switch# show cfs peers name port-security
Scope      : Logical [VSAN 1]
-----
Domain     Switch WWN          IP Address
-----
124       20:00:00:44:22:00:4a:9e  172.22.92.27   [Local]
98        20:00:00:05:30:01:1b:c2  172.22.92.215

Total number of entries = 2

Scope      : Logical [VSAN 3]
-----
Domain     Switch WWN          IP Address
-----
224       20:00:00:44:22:00:4a:9e  172.22.92.27   [Local]
151       20:00:00:05:30:01:1b:c2  172.22.92.215

Total number of entries = 2
```

Default CFS Settings

The following table lists the default settings for CFS configurations.

Table 32: Default CFS Parameters

Parameters	Default
CFS distribution on the switch	Enabled.
Database changes	Implicitly enabled with the first configuration change.
Application distribution	Differs based on application.
Commit	Explicit configuration is required.
CFS over IP	Disabled.
IPv4 multicast address	239.255.70.83.
IPv6 multicast address	ff15::eff:4653.

The CISCO-CFS-MIB contains SNMP configuration information for any CFS-related functions. Refer to the *Cisco Nexus 5000 Series MIB Quick Reference* for more information on this MIB.



Configuring User Accounts and RBAC

This chapter contains the following sections:

- [Configuring User Accounts and RBAC, page 347](#)

Configuring User Accounts and RBAC

This section describes how to configure user accounts and role-based access control (RBAC) on the Cisco Nexus 5000 Series switch.

Information About User Accounts and RBAC

You can create and manage users accounts and assign roles that limit access to operations on the Cisco Nexus 5000 Series switch. RBAC allows you to define the rules for an assign role that restrict the authorization that the user has to access management operations.

About User Accounts



Tip

The following words are reserved and cannot be used to configure users: bin, daemon, adm, lp, sync, shutdown, halt, mail, news, uucp, operator, games, gopher, ftp, nobody, nscd, mailnull, rpc, rpcuser, xfs, gdm, mtsuser, ftpuser, man, and sys.



Note

User passwords are not displayed in the configuration files.

**Caution**

The Cisco Nexus 5000 Series switch does not support all numeric usernames, whether created with TACACS+ or RADIUS, or created locally. Local users with all numeric names cannot be created. If an all numeric user name exists on an AAA server and is entered during login, the user is not logged in. Usernames must begin with an alphanumeric character and can contain only these special characters: (+ = . _ \ -). The # and ! symbols are not supported. If the username contains characters that are not allowed, the specified user is unable to log in.

Characteristics of Strong Passwords

A strong password has the following characteristics:

- At least eight characters long
- Does not contain many consecutive characters (such as "abcd")
- Does not contain many repeating characters (such as "aaabbb")
- Does not contain dictionary words
- Does not contain proper names
- Contains both uppercase and lowercase characters
- Contains numbers

The following are examples of strong passwords:

- If2CoM18
- 2009AsdfLkj30
- Cb1955S21

**Note**

Clear text passwords can contain alphanumeric characters only. Special characters, such as the dollar sign (\$) or the percent sign (%), are not allowed.

**Tip**

If a password is trivial (such as a short, easy-to-decipher password), the Cisco Nexus 5000 Series switch will reject your password configuration. Be sure to configure a strong password as shown in the sample configuration. Passwords are case sensitive.

About User Roles

User roles contain rules that define the operations allowed for the user who is assigned the role. Each user role can contain multiple rules and each user can have multiple roles. For example, if role1 allows access only to configuration operations, and role2 allows access only to debug operations, then users who belong to both role1 and role2 can access configuration and debug operations. You can also limit access to specific VSANs, VLANs and interfaces.

The Cisco Nexus 5000 Series switch provides the following default user roles:

- network-admin (superuser)—Complete read and write access to the entire Cisco Nexus 5000 Series switch.
- network-operator—Complete read access to the Cisco Nexus 5000 Series switch.

**Note**

If you belong to multiple roles, you can execute a combination of all the commands permitted by these roles. Access to a command takes priority over being denied access to a command. For example, suppose a user has RoleA, which denied access to the configuration commands. However, the users also has RoleB, which has access to the configuration commands. In this case, the users has access to the configuration commands.

About Rules

The rule is the basic element of a role. A rule defines what operations the role allows the user to perform. You can apply rules for the following parameters:

- Command—A command or group of commands defined in a regular expression.
- Feature—Commands that apply to a function provided by the switch.
 - Enter the **show role feature** command to display the feature names available for this parameter.
- Feature group—Default or user-defined group of features.
 - Enter the **show role feature-group** command to display the default feature groups available for this parameter.

These parameters create a hierarchical relationship. The most basic control parameter is the command. The next control parameter is the feature, which represents all commands associated with the feature. The last control parameter is the feature group. The feature group combines related features and allows you to easily manage of the rules.

You can configure up to 256 rules for each role. The user-specified rule number determines the order in which the rules are applied. Rules are applied in descending order. For example, if a role has three rules, rule 3 is applied before rule 2, which is applied before rule 1.

About User Role Policies

You can define user role policies to limit the switch resources that the user can access. You can define user role policies to limit access to interfaces, VLANs and VSANs.

User role policies are constrained by the rules defined for the role. For example, if you define an interface policy to permit access to specific interfaces, the user will not have access to the interfaces unless you configure a command rule for the role to permit the interface command.

If a command rule permits access to specific resources (interfaces, VLANs or VSANs), the user is permitted to access these resources, even if they are not listed in the user role policies associated with that user.

Related Topics

[Changing User Role Interface Policies, on page 353](#)

Guidelines and Limitations for User Accounts

User account and RBAC have the following configuration guidelines and limitations:

- You can add up to 256 rules to a user role.
- You can assign a maximum of 64 user roles to a user account.



Note A user account must have at least one user role.

Configuring User Accounts

You can create a maximum of 256 user accounts on a Cisco Nexus 5000 Series switch. User accounts have the following attributes:

- Username
- Password
- Expiry date
- User roles

User accounts can have a maximum of 64 user roles.



Note Changes to user account attributes do not take effect until the user logs in and creates a new session.

To configure a user account, perform this task:

SUMMARY STEPS

1. (Optional) `switch(config)# show role`
2. `switch# configure terminal`
3. `switch(config)# username user-id [password password] [expire date] [role role-name]`
4. (Optional) `switch# show user-account`
5. (Optional) `switch# copy running-config startup-config`

DETAILED STEPS

	Command or Action	Purpose
Step 1	<code>switch(config)# show role</code>	(Optional) Displays the user roles available. You can configure other user roles, if necessary.

	Command or Action	Purpose
Step 2	switch# configure terminal	Enters configuration mode.
Step 3	switch(config)# username <i>user-id</i> [password <i>password</i>] [expire <i>date</i>] [role <i>role-name</i>]	Configures a user account. The <i>user-id</i> argument is a case-sensitive, alphanumeric character string with a maximum length of 28 characters. The default password is undefined. Note If you do not specify a password, the user might not be able to log in to the Cisco Nexus 5000 Series switch. The expire <i>date</i> option format is YYYY-MM-DD. The default is no expiry date.
Step 4	switch# show user-account	(Optional) Displays the role configuration.
Step 5	switch# copy running-config startup-config	(Optional) Copies the running configuration to the startup configuration.

The following example shows how to configure a user account:

```
switch# configure terminal
switch(config)# username NewUser password 4Ty18Rnt

switch(config)# exit
switch# show user-account
```

Configuring RBAC

Creating User Roles and Rules

Each user role can have up to 256 rules. You can assign a user role to more than one user account.

The rule number you specify determines the order in which the rules are applied. Rules are applied in descending order. For example, if a role has three rules, rule 3 is applied before rule 2, which is applied before rule 1.

SUMMARY STEPS

1. switch# **configure terminal**
2. switch(config)# **role name** *role-name*
3. switch(config-role)# **rule number** {deny | permit} **command** *command-string*
4. switch(config-role)# **rule number** {deny | permit} {read | read-write}
5. switch(config-role)# **rule number** {deny | permit} {read | read-write} **feature** *feature-name*
6. switch(config-role)# **rule number** {deny | permit} {read | read-write} **feature-group** *group-name*
7. (Optional) switch(config-role)# **description** *text*
8. (Optional) switch# **show role**
9. (Optional) switch# **copy running-config startup-config**

DETAILED STEPS

	Command or Action	Purpose
Step 1	switch# configure terminal	Enters configuration mode.
Step 2	switch(config)# role name <i>role-name</i>	Specifies a user role and enters role configuration mode. The <i>role-name</i> argument is a case-sensitive, alphanumeric character string with a maximum length of 16 characters.
Step 3	switch(config-role)# rule <i>number</i> { deny permit } command <i>command-string</i>	Configures a command rule. The <i>command-string</i> argument can contain spaces and regular expressions. For example, "interface ethernet *" includes all Ethernet interfaces. Repeat this command for as many rules as needed.
Step 4	switch(config-role)# rule <i>number</i> { deny permit } { read read-write }	Configures a read only or read and write rule for all operations.
Step 5	switch(config-role)# rule <i>number</i> { deny permit } { read read-write } feature <i>feature-name</i>	Configures a read-only or read-and-write rule for a feature. Use the show role feature command to display a list of features. Repeat this command for as many rules as needed.
Step 6	switch(config-role)# rule <i>number</i> { deny permit } { read read-write } feature-group <i>group-name</i>	Configures a read-only or read-and-write rule for a feature group. Use the show role feature-group command to display a list of feature groups. Repeat this command for as many rules as needed.
Step 7	switch(config-role)# description <i>text</i>	(Optional) Configures the role description. You can include spaces in the description.
Step 8	switch# show role	(Optional) Displays the user role configuration.
Step 9	switch# copy running-config startup-config	(Optional) Copies the running configuration to the startup configuration.

The following example shows how to create user roles and specify rules:

```
switch# configure terminal
switch(config)# role name UserA
switch(config-role)# rule deny command clear users
switch(config-role)# rule deny read-write
switch(config-role)# description This role does not allow users to use clear commands
switch(config-role)# end
switch(config)# show role
```

Creating Feature Groups

You can create feature groups.

SUMMARY STEPS

1. switch# **configure terminal**
2. switch(config)# **role feature-group** *group-name*
3. (Optional) switch# **show role feature-group**
4. (Optional) switch# **copy running-config startup-config**

DETAILED STEPS

	Command or Action	Purpose
Step 1	switch# configure terminal	Enters configuration mode.
Step 2	switch(config)# role feature-group <i>group-name</i>	Specifies a user role feature group and enters role feature group configuration mode. The <i>group-name</i> argument is a case-sensitive, alphanumeric character string with a maximum length of 32 characters.
Step 3	switch# show role feature-group	(Optional) Displays the role feature group configuration.
Step 4	switch# copy running-config startup-config	(Optional) Copies the running configuration to the startup configuration.

Changing User Role Interface Policies

You can change a user role interface policy to limit the interfaces that the user can access.

SUMMARY STEPS

1. switch# **configure terminal**
2. switch(config)# **role name** *role-name*
3. switch(config-role)# **interface policy deny**
4. switch(config-role-interface)# **permit interface** *interface-list*
5. switch(config-role-interface)# **exit**
6. (Optional) switch(config-role)# **show role**
7. (Optional) switch(config-role)# **copy running-config startup-config**

DETAILED STEPS

	Command or Action	Purpose
Step 1	switch# configure terminal	Enters configuration mode.
Step 2	switch(config)# role name <i>role-name</i>	Specifies a user role and enters role configuration mode.
Step 3	switch(config-role)# interface policy deny	Enters role interface policy configuration mode.
Step 4	switch(config-role-interface)# permit interface <i>interface-list</i>	Specifies a list of interfaces that the role can access. Repeat this command for as many interfaces as needed. For this command, you can specify Ethernet interfaces, Fibre Channel interfaces, and virtual Fibre Channel interfaces.
Step 5	switch(config-role-interface)# exit	Exits role interface policy configuration mode.
Step 6	switch(config-role)# show role	(Optional) Displays the role configuration.
Step 7	switch(config-role)# copy running-config startup-config	(Optional) Copies the running configuration to the startup configuration.

The following example shows how to change a user role interface policy to limit the interfaces that the user can access:

```
switch# configure terminal
switch(config)# role name UserB
switch(config-role)# interface policy deny
switch(config-role-interface)# permit interface ethernet 2/1
switch(config-role-interface)# permit interface fc 3/1
switch(config-role-interface)# permit interface vfc 30/1
```

You can specify a list of interfaces that the role can access. You can specify it for as many interfaces as needed.

Changing User Role VLAN Policies

You can change a user role VLAN policy to limit the VLANs that the user can access.

SUMMARY STEPS

1. switch# **configure terminal**
2. switch(config)# **role name** *role-name*
3. switch(config-role)# **vlan policy deny**
4. switch(config-role-vlan)# **permit vlan** *vlan-list*
5. (Optional) switch# **show role**
6. (Optional) switch# **copy running-config startup-config**

DETAILED STEPS

	Command or Action	Purpose
Step 1	switch# configure terminal	Enters configuration mode.
Step 2	switch(config)# role name <i>role-name</i>	Specifies a user role and enters role configuration mode.
Step 3	switch(config-role)# vlan policy deny	Enters role VLAN policy configuration mode.
Step 4	switch(config-role-vlan)# permit vlan <i>vlan-list</i>	Specifies a range of VLANs that the role can access. Repeat this command for as many VLANs as needed.
Step 5	switch# show role	(Optional) Displays the role configuration.
Step 6	switch# copy running-config startup-config	(Optional) Copies the running configuration to the startup configuration.

Changing User Role VSAN Policies

You can change a user role VSAN policy to limit the VSANs that the user can access.

SUMMARY STEPS

1. switch# **configure terminal**
2. switch(config-role)# **role name** *role-name*
3. switch(config-role)# **vsan policy deny**
4. switch(config-role-vsan)# **permit vsan** *vsan-list*
5. (Optional) switch# **show role**
6. (Optional) switch# **copy running-config startup-config**

DETAILED STEPS

	Command or Action	Purpose
Step 1	switch# configure terminal	Enters configuration mode.
Step 2	switch(config-role)# role name <i>role-name</i>	Specifies a user role and enters role configuration mode.
Step 3	switch(config-role)# vsan policy deny	Enters role VSAN policy configuration mode.
Step 4	switch(config-role-vsan)# permit vsan <i>vsan-list</i>	Specifies a range of VSANs that the role can access. Repeat this command for as many VSANs as needed.

	Command or Action	Purpose
Step 5	switch# show role	(Optional) Displays the role configuration.
Step 6	switch# copy running-config startup-config	(Optional) Copies the running configuration to the startup configuration.

Verifying User Accounts and RBAC Configuration

To display user account and RBAC configuration information, perform one of the following tasks:

Command	Purpose
switch# show role	Displays the user role configuration
switch# show role feature	Displays the feature list.
switch# show role feature-group	Displays the feature group configuration.
switch# show startup-config security	Displays the user account configuration in the startup configuration.
switch# show running-config security [all]	Displays the user account configuration in the running configuration. The all keyword displays the default values for the user accounts.
switch# show user-account	Displays user account information.

Default User Account and RBAC Settings

The following table lists the default settings for user accounts and RBAC parameters.

Table 33: Default User Accounts and RBAC Parameters

Parameters	Default
User account password	Undefined.
User account expiry date.	None.
Interface policy	All interfaces are accessible.
VLAN policy	All VLANs are accessible.
VFC policy	All VFCs are accessible.

Parameters	Default
VETH policy	All VETHs are accessible.



Configuring Session Manager

This chapter contains the following sections:

- [Configuring Session Manager, page 359](#)

Configuring Session Manager

This section describes how to configure the Session Manager features in Cisco NX-OS.

Information About Session Manager

Session Manager allows you to implement your configuration changes in batch mode. Session Manager works in the following phases:

- **Configuration session**—Creates a list of commands that you want to implement in session manager mode.
- **Validation**—Provides a basic semantic check on your configuration. Cisco NX-OS returns an error if the semantic check fails on any part of the configuration.
- **Verification**—Verifies the configuration as a whole, based on the existing hardware and software configuration and resources. Cisco NX-OS returns an error if the configuration does not pass this verification phase.
- **Commit**— Cisco NX-OS verifies the complete configuration and implements the changes atomically to the device. If a failure occurs, Cisco NX-OS reverts to the original configuration.
- **Abort**—Discards the configuration changes before implementation.

You can optionally end a configuration session without committing the changes. You can also save a configuration session.

Configuration Guidelines and Limitations

Session Manager has the following configuration guidelines and limitations:

- Session Manager supports only the ACL feature.

- You can create up to 32 configuration sessions.
- You can configure a maximum of 20,000 commands across all sessions.

Configuring Session Manager

Creating a Session

You can create up to 32 configuration sessions. To create a configuration session, perform this task:

SUMMARY STEPS

1. switch# **configure session** *name*
2. (Optional) switch(config-s)# **show configuration session** [*name*]
3. (Optional) switch(config-s)# **save** *location*

DETAILED STEPS

	Command or Action	Purpose
Step 1	switch# configure session <i>name</i>	Creates a configuration session and enters session configuration mode. The name can be any alphanumeric string.
Step 2	switch(config-s)# show configuration session [<i>name</i>]	(Optional) Displays the contents of the session.
Step 3	switch(config-s)# save <i>location</i>	(Optional) Saves the session to a file. The location can be in bootflash or volatile.

Configuring ACLs in a Session

You can configure ACLs within a configuration session. To configure ACLs within a configuration session, perform this task:

SUMMARY STEPS

1. switch# **configure session** *name*
2. switch(config-s)# **ip access-list** *name*
3. (Optional) switch(config-s-acl)# **permit** *protocol source destination*
4. switch(config-s-acl)# **interface** *interface-type number*
5. switch(config-s-if)# **ip port access-group** *name in*
6. (Optional) switch# **show configuration session** [*name*]

DETAILED STEPS

	Command or Action	Purpose
Step 1	switch# configure session <i>name</i>	Creates a configuration session and enters session configuration mode. The name can be any alphanumeric string.
Step 2	switch(config-s)# ip access-list <i>name</i>	Creates an ACL.
Step 3	switch(config-s-acl)# permit <i>protocol source destination</i>	(Optional) Adds a permit statement to the ACL.
Step 4	switch(config-s-acl)# interface <i>interface-type number</i>	Enters interface configuration mode.
Step 5	switch(config-s-if)# ip port access-group <i>name in</i>	Adds a port access group to the interface.
Step 6	switch# show configuration session [<i>name</i>]	(Optional) Displays the contents of the session.

Verifying a Session

To verify a session, use the following command in session mode:

Command	Purpose
switch(config-s)# verify [verbose]	Verifies the commands in the configuration session.

Committing a Session

To commit a session, use the following command in session mode:

Command	Purpose
switch(config-s)# commit [verbose]	Commits the commands in the configuration session.

Saving a Session

To save a session, use the following command in session mode:

Command	Purpose
switch(config-s)# save <i>location</i>	(Optional) Saves the session to a file. The location can be in bootflash or volatile.

Discarding a Session

To discard a session, use the following command in session mode:

Command	Purpose
switch(config-s)# abort	Discards the configuration session without applying the commands.

Session Manager Example Configuration

This example shows how to create a configuration session for ACLs:

```
switch# configure session name test2
switch(config-s)# ip access-list acl2
switch(config-s-acl)# permit tcp any any
switch(config-s-acl)# exit
switch(config-s)# interface Ethernet 1/4
switch(config-s-ip)# ip port access-group acl2 in
switch(config-s-ip)# exit
switch(config-s)# verify
switch(config-s)# exit
switch# show configuration session test2
```

Verifying Session Manager Configuration

To verify Session Manager configuration information, use the following commands:

Command	Purpose
switch# show configuration session [<i>name</i>]	Displays the contents of the configuration session.
switch# show configuration session status [<i>name</i>]	Displays the status of the configuration session.
switch# show configuration session summary	Displays a summary of all the configuration sessions.



Configuring Online Diagnostics

This chapter describes how to configure the generic online diagnostics (GOLD) feature. It contains the following sections:

- [Information About Online Diagnostics, page 363](#)
- [Configuring Online Diagnostics, page 366](#)
- [Verifying Online Diagnostics Configuration, page 367](#)
- [Default GOLD Settings, page 367](#)

Information About Online Diagnostics

Online diagnostics provide verification of hardware components during switch bootup or reset, and they monitor the health of the hardware during normal switch operation.

Online Diagnostics Overview

Cisco Nexus 5000 Series switches support bootup diagnostics and runtime diagnostics. Bootup diagnostics include disruptive tests and nondisruptive tests that run during system bootup and system reset.

Runtime diagnostics (also known as health monitoring diagnostics) include nondisruptive tests that run in the background during normal operation of the switch.

Bootup Diagnostics

Bootup diagnostics detect faulty hardware before bringing the switch online. Bootup diagnostics also check the data path and control path connectivity between the supervisor and the ASICs. The following table describes the diagnostics that are run only during switch bootup or reset.

Table 34: Bootup Diagnostics

Diagnostic	Description
PCIe	Tests PCI express (PCIe) access.

Diagnostic	Description
NVRAM	Verifies the integrity of the NVRAM.
In band port	Tests connectivity of the inband port to the supervisor.
Management port	Tests the management port.
Memory	Verifies the integrity of the DRAM.

Bootup diagnostics also include a set of tests that are common with health monitoring diagnostics.

Bootup diagnostics log any failures to the onboard failure logging (OBFL) system. Failures also trigger an LED display to indicate diagnostic test states (on, off, pass, or fail).

You can configure Cisco Nexus 5000 Series switches to either bypass the bootup diagnostics, or run the complete set of bootup diagnostics.

Health Monitoring Diagnostics

Health monitoring diagnostics provide information about the health of the switch. They detect runtime hardware errors, memory errors, software faults, and resource exhaustion.

Health monitoring diagnostics are nondisruptive and run in the background to ensure the health of a switch that is processing live network traffic.

The following table describes the health monitoring diagnostics for the switch.

Table 35: Health Monitoring Diagnostics Tests

Diagnostic	Description
LED	Monitors port and system status LEDs.
Power Supply	Monitors the power supply health state.
Temperature Sensor	Monitors temperature sensor readings.
Test Fan	Monitors fan speed and fan control.

The following table describes the health monitoring diagnostics that also run during system boot or system reset.

Table 36: Health Monitoring and Bootup Diagnostics Tests

Diagnostic	Description
SPROM	Verifies the integrity of backplane and supervisor SPROMs.

Diagnostic	Description
Fabric engine	Tests the switch fabric ASICs.
Fabric port	Tests the ports on the switch fabric ASIC.
Forwarding engine	Tests the forwarding engine ASICs.
Forwarding engine port	Tests the ports on the forwarding engine ASICs.
Front port	Tests the components (such as PHY and MAC) on the front ports.

Expansion Module Diagnostics

During switch bootup or reset, the bootup diagnostics include tests for the in-service expansion modules in the switch.

When you insert an expansion module into a running switch, a set of diagnostics tests are run. The following table describes the bootup diagnostics for an expansion module. These tests are common with the bootup diagnostics. If the bootup diagnostics fail, the expansion module is not placed into service.

Table 37: Expansion Module Bootup and Health Monitoring Diagnostics

Diagnostic	Description
SPROM	Verifies the integrity of backplane and supervisor SPROMs.
Fabric engine	Tests the switch fabric ASICs.
Fabric port	Tests the ports on the switch fabric ASIC.
Forwarding engine	Tests the forwarding engine ASICs.
Forwarding engine port	Tests the ports on the forwarding engine ASICs.
Front port	Tests the components (such as PHY and MAC) on the front ports.

Health monitoring diagnostics are run on in-service expansion modules. The following table describes the additional tests that are specific to health monitoring diagnostics for expansion modules.

Table 38: Expansion Module Health Monitoring Diagnostics

Diagnostic	Description
LED	Monitors port and system status LEDs.

Diagnostic	Description
Temperature Sensor	Monitors temperature sensor readings.

Configuring Online Diagnostics

You can configure the bootup diagnostics to run the complete set of tests, or you can bypass all bootup diagnostic tests for a faster module boot up time.



Note

We recommend that you set the bootup online diagnostics level to complete. We do not recommend bypassing the bootup online diagnostics.

SUMMARY STEPS

1. switch# **configure terminal**
2. switch(config)# **diagnostic bootup level [complete | bypass]**
3. (Optional) switch# **show diagnostic bootup level**

DETAILED STEPS

	Command or Action	Purpose
Step 1	switch# configure terminal	Enters configuration mode.
Step 2	switch(config)# diagnostic bootup level [complete bypass]	Configures the bootup diagnostic level to trigger diagnostics when the device boots, as follows: <ul style="list-style-type: none"> • complete—Performs all bootup diagnostics. This is the default value. • bypass—Does not perform any bootup diagnostics.
Step 3	switch# show diagnostic bootup level	(Optional) Displays the bootup diagnostic level (bypass or complete) that is currently in place on the switch.

The following example shows how to configure the bootup diagnostics level to trigger the complete diagnostics:

```
switch# configure terminal
switch(config)# diagnostic bootup level complete
```

Verifying Online Diagnostics Configuration

To display online diagnostics configuration information, perform one of the following tasks:

Command	Purpose
show diagnostic bootup level	Displays the bootup diagnostics level.
show diagnostic result module <i>slot</i>	Displays the results of the diagnostics tests.

Default GOLD Settings

The following table lists the default settings for online diagnostics parameters.

Table 39: Default Online Diagnostics Parameters

Parameters	Default
Bootup diagnostics level	complete



CHAPTER 27

Configuring System Message Logging

This chapter describes how to configure system message logging on the Cisco Nexus 5000 Series switch and contains the following sections:

- [Information About System Message Logging, page 369](#)
- [Configuring System Message Logging, page 370](#)
- [Verifying System Message Logging Configuration, page 382](#)
- [Default System Message Logging Settings, page 382](#)

Information About System Message Logging

You can use system message logging to control the destination and to filter the severity level of messages that system processes generate. You can configure logging to terminal sessions, a log file, and syslog servers on remote systems.

By default, the Cisco Nexus 5000 Series switch outputs messages to terminal sessions.

By default, the switch logs system messages to a log file.

The following table describes the severity levels used in system messages. When you configure the severity level, the system outputs messages at that level and lower.

Table 40: System Message Severity Levels

Level	Description
0 – emergency	System unusable
1 – alert	Immediate action needed
2 – critical	Critical condition
3 – error	Error condition
4 – warning	Warning condition

Level	Description
5 – notification	Normal but significant condition
6 – informational	Informational message only
7 – debugging	Appears during debugging only

The switch logs the most recent 100 messages of severity 0, 1, or 2 to the NVRAM log. You cannot configure logging to the NVRAM.

You can configure which system messages should be logged based on the facility that generated the message and its severity level.

Related Topics

[Configuring Module and Facility Messages Logging, on page 374](#)

[Configuring System Message Logging to a File, on page 373](#)

[Configuring System Message Logging to Terminal Sessions, on page 370](#)

syslog Servers

syslog servers run on remote systems that are configured to log system messages based on the syslog protocol. You can configure the Cisco Nexus 5000 Series to send its logs to up to three syslog servers.

To support the same configuration of syslog servers on all switches in a fabric, you can use the Cisco Fabric Services (CFS) to distribute the syslog server configuration.



Note

When the switch first initializes, messages are sent to syslog servers only after the network is initialized.

Configuring System Message Logging

Configuring System Message Logging to Terminal Sessions

You can configure the switch to log messages by their severity level to console, Telnet, and SSH sessions.

By default, logging is enabled for terminal sessions.

SUMMARY STEPS

1. switch# **terminal monitor**
2. switch# **configure terminal**
3. switch(config)# **logging console** [*severity-level*]
4. (Optional) switch(config)# **no logging console** [*severity-level*]
5. switch(config)# **logging monitor** [*severity-level*]
6. (Optional) switch(config)# **no logging monitor** [*severity-level*]
7. (Optional) switch# **show logging console**
8. (Optional) switch# **show logging monitor**
9. (Optional) switch# **copy running-config startup-config**

DETAILED STEPS

	Command or Action	Purpose
Step 1	switch# terminal monitor	Copies syslog messages from the console to the current terminal session.
Step 2	switch# configure terminal	Enters configuration mode.
Step 3	switch(config)# logging console [<i>severity-level</i>]	<p>Enables the switch to log messages to the console session based on a specified severity level or higher (a lower number value indicates a higher severity level). Severity levels range from 0 to 7:</p> <ul style="list-style-type: none"> • 0 – emergency • 1 – alert • 2 – critical • 3 – error • 4 – warning • 5 – notification • 6 – informational • 7 – debugging <p>If the severity level is not specified, the default of 2 is used.</p>
Step 4	switch(config)# no logging console [<i>severity-level</i>]	(Optional) Disables logging messages to the console.
Step 5	switch(config)# logging monitor [<i>severity-level</i>]	<p>Enables the switch to log messages to the monitor based on a specified severity level or higher (a lower number value indicates a higher severity level). Severity levels range from 0 to 7:</p> <ul style="list-style-type: none"> • 0 – emergency • 1 – alert • 2 – critical

	Command or Action	Purpose
		<ul style="list-style-type: none"> • 3 – error • 4 – warning • 5 – notification • 6 – informational • 7 – debugging <p>If the severity level is not specified, the default of 2 is used. The configuration applies to Telnet and SSH sessions.</p>
Step 6	switch(config)# no logging monitor [severity-level]	(Optional) Disables logging messages to telnet and SSH sessions.
Step 7	switch# show logging console	(Optional) Displays the console logging configuration.
Step 8	switch# show logging monitor	(Optional) Displays the monitor logging configuration.
Step 9	switch# copy running-config startup-config	(Optional) Copies the running configuration to the startup configuration.

The following example shows how to configure a logging level of 3 for the console:

```
switch# configure terminal
switch(config)# logging console 3
```

The following example shows how to display the console logging configuration:

```
switch# show logging console
Logging console:                enabled (Severity: error)
```

The following example shows how to disable logging for the console:

```
switch# configure terminal
switch(config)# no logging console
```

The following example shows how to configure a logging level of 4 for the terminal session:

```
switch# terminal monitor
switch# configure terminal
switch(config)# logging monitor 4
```

The following example shows how to display the terminal session logging configuration:

```
switch# show logging monitor
Logging monitor:                enabled (Severity: warning)
```

The following example shows how to disable logging for the terminal session:

```
switch# configure terminal
switch(config)# no logging monitor
```

Configuring System Message Logging to a File

You can configure the switch to log system messages to a file. By default, system messages are logged to the file `log:messages`.

SUMMARY STEPS

1. `switch# configure terminal`
2. `switch(config)# logging logfile logfile-name severity-level [size bytes]`
3. (Optional) `switch(config)# no logging logfile [logfile-name severity-level [size bytes]]`
4. (Optional) `switch# show logging info`
5. (Optional) `switch# copy running-config startup-config`

DETAILED STEPS

	Command or Action	Purpose
Step 1	<code>switch# configure terminal</code>	Enters configuration mode.
Step 2	<code>switch(config)# logging logfile logfile-name severity-level [size bytes]</code>	Configures the name of the log file used to store system messages and the minimum severity level to log. You can optionally specify a maximum file size. The default severity level is 5 and the file size is 4194304. Severity levels range from 0 to 7: <ul style="list-style-type: none"> • 0 – emergency • 1 – alert • 2 – critical • 3 – error • 4 – warning • 5 – notification • 6 – informational • 7 – debugging The file size is from 4096 to 10485760 bytes.
Step 3	<code>switch(config)# no logging logfile [logfile-name severity-level [size bytes]]</code>	(Optional) Disables logging to the log file.
Step 4	<code>switch# show logging info</code>	(Optional) Displays the logging configuration.
Step 5	<code>switch# copy running-config startup-config</code>	(Optional) Copies the running configuration to the startup configuration.

The following example shows how to configure a switch to log system messages to a file:

```
switch# configure terminal
switch(config)# logging logfile my_log 6 size 4194304
```

The following example shows how to display the logging configuration (some of the output has been removed for brevity):

```
switch# show logging info
Logging console:          enabled (Severity: debugging)
Logging monitor:         enabled (Severity: debugging)
Logging linecard:        enabled (Severity: notifications)
Logging fex:             enabled (Severity: notifications)
Logging timestamp:       Seconds
Logging server:          disabled
Logging logfile:         enabled
                          Name - my_log: Severity - informational Size - 4194304
Facility                 Default Severity          Current Session Severity
-----
aaa                      3                          3
aclmgr                   3                          3
afm                      3                          3
altos                    3                          3
auth                     0                          0
authpriv                 3                          3
bootvar                  5                          5
callhome                 2                          2
capability                2                          2
cdp                      2                          2
cert_enroll              2                          2
...
```

Related Topics

[Displaying and Clearing Log Files, on page 381](#)

Configuring Module and Facility Messages Logging

You can configure the severity level and time-stamp units of messages logged by modules and facilities.

SUMMARY STEPS

1. switch# **configure terminal**
2. switch(config)# **logging module** [*severity-level*]
3. switch(config)# **logging level** *facility severity-level*
4. (Optional) switch(config)# **no logging module** [*severity-level*]
5. (Optional) switch(config)# **no logging level** [*facility severity-level*]
6. (Optional) switch# **show logging module**
7. (Optional) switch# **show logging level** [*facility*]
8. (Optional) switch# **copy running-config startup-config**

DETAILED STEPS

	Command or Action	Purpose
Step 1	switch# configure terminal	Enters configuration mode.
Step 2	switch(config)# logging module [<i>severity-level</i>]	<p>Enables module log messages that have the specified severity level or higher. Severity levels range from 0 to 7:</p> <ul style="list-style-type: none"> • 0 – emergency • 1 – alert • 2 – critical • 3 – error • 4 – warning • 5 – notification • 6 – informational • 7 – debugging <p>If the severity level is not specified, the default of 5 is used.</p>
Step 3	switch(config)# logging level <i>facility</i> <i>severity-level</i>	<p>Enables logging messages from the specified facility that have the specified severity level or higher. Severity levels from 0 to 7:</p> <ul style="list-style-type: none"> • 0 – emergency • 1 – alert • 2 – critical • 3 – error • 4 – warning • 5 – notification • 6 – informational • 7 – debugging <p>To apply the same severity level to all facilities, use the all facility. For defaults, see the show logging level command.</p>
Step 4	switch(config)# no logging module [<i>severity-level</i>]	(Optional) Disables module log messages.
Step 5	switch(config)# no logging level [<i>facility</i> <i>severity-level</i>]	(Optional) Resets the logging severity level for the specified facility to its default level. If you do not specify a facility and severity level, the switch resets all facilities to their default levels.

	Command or Action	Purpose
Step 6	switch# show logging module	(Optional) Displays the module logging configuration.
Step 7	switch# show logging level [<i>facility</i>]	(Optional) Displays the logging level configuration and the system default level by facility. If you do not specify a facility, the switch displays levels for all facilities.
Step 8	switch# copy running-config startup-config	(Optional) Copies the running configuration to the startup configuration.

The following example shows how to configure the severity level of module and specific facility messages:

```
switch# configure terminal
switch(config)# logging module 3
switch(config)# logging level aaa 2
```

Configuring Logging Timestamps

You can configure the time-stamp units of messages logged by the Cisco Nexus 5000 Series switch.

SUMMARY STEPS

1. switch# **configure terminal**
2. switch(config)# **logging timestamp** {microseconds | milliseconds | seconds}
3. (Optional) switch(config)# **no logging timestamp** {microseconds | milliseconds | seconds}
4. (Optional) switch# **show logging timestamp**
5. (Optional) switch# **copy running-config startup-config**

DETAILED STEPS

	Command or Action	Purpose
Step 1	switch# configure terminal	Enters configuration mode.
Step 2	switch(config)# logging timestamp {microseconds milliseconds seconds}	Sets the logging time-stamp units. By default, the units are seconds.
Step 3	switch(config)# no logging timestamp {microseconds milliseconds seconds}	(Optional) Resets the logging time-stamp units to the default of seconds.
Step 4	switch# show logging timestamp	(Optional) Displays the logging time-stamp units configured.

	Command or Action	Purpose
Step 5	switch# copy running-config startup-config	(Optional) Copies the running configuration to the startup configuration.

The following example shows how to configure the time-stamp units of messages:

```
switch# configure terminal
switch(config)# logging timestamp milliseconds
switch(config)# exit
switch# show logging timestamp
Logging timestamp:                               Milliseconds
```

Configuring syslog Servers

You can configure up to three syslog servers that reference remote systems where you want to log system messages.

SUMMARY STEPS

1. switch# **configure terminal**
2. switch(config)# **logging server host [severity-level [usr-vrf vrf-name [facility facility]]]**
3. (Optional) switch(config)# **no logging server host**
- 4.
5. (Optional) switch# **show logging server**
6. (Optional) switch# **copy running-config startup-config**

DETAILED STEPS

	Command or Action	Purpose
Step 1	switch# configure terminal	Enters configuration mode.
Step 2	switch(config)# logging server host [severity-level [usr-vrf vrf-name [facility facility]]]	Configures a syslog server at the specified host name or IPv4 or IPv6 address. You can limit logging of messages with a minimum severity level and for a specific facility. Severity levels range from 0 to 7: <ul style="list-style-type: none"> • 0 – emergency • 1 – alert • 2 – critical • 3 – error • 4 – warning • 5 – notification

	Command or Action	Purpose
		<ul style="list-style-type: none"> • 6 – informational • 7 – debugging <p>The default outgoing facility is local7.</p>
Step 3	switch(config)# no logging server <i>host</i>	(Optional) Removes the logging server for the specified host.
Step 4		
Step 5	switch# show logging server	(Optional) Displays the syslog server configuration.
Step 6	switch# copy running-config startup-config	(Optional) Copies the running configuration to the startup configuration.

The following example shows how to configure a syslog server:

```
switch# configure terminal
switch(config)# logging server 172.28.254.254 5 use-vrf VRfname facility local3
```

Configuring syslog on a UNIX or Linux System

You can configure a syslog server on a UNIX or Linux system by adding the following line to the /etc/syslog.conf file:

```
facility.level <five tab characters> action
```

The following table describes the syslog fields that you can configure.

Table 41: syslog Fields in syslog.conf

Field	Description
Facility	<p>Creator of the message, which can be auth, authpriv, cron, daemon, kern, lpr, mail, mark, news, syslog, user, local0 through local7, or an asterisk (*) for all. These facility designators allow you to control the destination of messages based on their origin.</p> <p>Note Check your configuration before using a local facility.</p>
Level	<p>Minimum severity level at which messages are logged, which can be debug, info, notice, warning, err, crit, alert, emerg, or an asterisk (*) for all. You can use none to disable a facility.</p>

Field	Description
Action	Destination for messages, which can be a filename, a host name preceded by the at sign (@), or a comma-separated list of users or an asterisk (*) for all logged-in users.

SUMMARY STEPS

1. Log debug messages with the local7 facility in the file /var/log/myfile.log by adding the following line to the /etc/syslog.conf file:
2. Create the log file by entering these commands at the shell prompt:
3. Make sure the system message logging daemon reads the new changes by checking myfile.log after entering this command:

DETAILED STEPS

-
- Step 1** Log debug messages with the local7 facility in the file /var/log/myfile.log by adding the following line to the /etc/syslog.conf file:
- ```
debug.local7 /var/log/myfile.log
```
- Step 2** Create the log file by entering these commands at the shell prompt:
- ```
$ touch /var/log/myfile.log
$ chmod 666 /var/log/myfile.log
```
- Step 3** Make sure the system message logging daemon reads the new changes by checking myfile.log after entering this command:
- ```
$ kill -HUP ~cat /etc/syslog.pid~
```
- 

## Configuring syslog Server Configuration Distribution

You can distribute the syslog server configuration to other switches in the network by using the Cisco Fabric Services (CFS) infrastructure.

After you enable syslog server configuration distribution, you can modify the syslog server configuration and view the pending changes before committing the configuration for distribution. As long as distribution is enabled, the switch maintains pending changes to the syslog server configuration.



**Note** If the switch is restarted, the syslog server configuration changes that are kept in volatile memory may be lost.

### Before You Begin

You must have configured one or more syslog servers.

## SUMMARY STEPS

1. switch# **configure terminal**
2. switch(config)# **logging distribute**
3. switch(config)# **logging commit**
4. switch(config)# **logging abort**
5. (Optional) switch(config)# **no logging distribute**
6. (Optional) switch# **show logging pending**
7. (Optional) switch# **show logging pending-diff**
8. (Optional) switch# **show logging internal info**
9. (Optional) switch# **copy running-config startup-config**

## DETAILED STEPS

|               | Command or Action                                 | Purpose                                                                                                                                                                                                                                                                                             |
|---------------|---------------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Step 1</b> | switch# <b>configure terminal</b>                 | Enters configuration mode.                                                                                                                                                                                                                                                                          |
| <b>Step 2</b> | switch(config)# <b>logging distribute</b>         | Enables distribution of syslog server configuration to network switches using the CFS infrastructure. By default, distribution is disabled.                                                                                                                                                         |
| <b>Step 3</b> | switch(config)# <b>logging commit</b>             | Commits the pending changes to the syslog server configuration for distribution to the switches in the fabric.                                                                                                                                                                                      |
| <b>Step 4</b> | switch(config)# <b>logging abort</b>              | Cancels the pending changes to the syslog server configuration.                                                                                                                                                                                                                                     |
| <b>Step 5</b> | switch(config)# <b>no logging distribute</b>      | (Optional)<br>Disables distribution of syslog server configuration to network switches using the CFS infrastructure. You cannot disable distribution when configuration changes are pending. See the <b>logging commit</b> and <b>logging abort</b> commands. By default, distribution is disabled. |
| <b>Step 6</b> | switch# <b>show logging pending</b>               | (Optional)<br>Displays the pending changes to the syslog server configuration.                                                                                                                                                                                                                      |
| <b>Step 7</b> | switch# <b>show logging pending-diff</b>          | (Optional)<br>Displays the differences from the current syslog server configuration to the pending changes of the syslog server configuration.                                                                                                                                                      |
| <b>Step 8</b> | switch# <b>show logging internal info</b>         | (Optional)<br>Displays information about the current state of syslog server distribution and the last action taken.                                                                                                                                                                                 |
| <b>Step 9</b> | switch# <b>copy running-config startup-config</b> | (Optional)<br>Copies the running configuration to the startup configuration.                                                                                                                                                                                                                        |

## Related Topics

[Information About CFS, on page 329](#)

## Displaying and Clearing Log Files

You can display or clear messages in the log file and the NVRAM.

### SUMMARY STEPS

1. switch# **show logging last** *number-lines*
2. switch# **show logging logfile** [**start-time** *yyyy mmm dd hh:mm:ss*] [**end-time** *yyyy mmm dd hh:mm:ss*]
3. switch# **show logging nvram** [**last** *number-lines*]
4. switch# **clear logging logfile**
5. switch# **clear logging nvram**

### DETAILED STEPS

|               | Command or Action                                                                                                                     | Purpose                                                                                                                                                                                                                                            |
|---------------|---------------------------------------------------------------------------------------------------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Step 1</b> | switch# <b>show logging last</b> <i>number-lines</i>                                                                                  | Displays the last number of lines in the logging file. You can specify from 1 to 9999 for the last number of lines.                                                                                                                                |
| <b>Step 2</b> | switch# <b>show logging logfile</b> [ <b>start-time</b> <i>yyyy mmm dd hh:mm:ss</i> ] [ <b>end-time</b> <i>yyyy mmm dd hh:mm:ss</i> ] | Displays the messages in the log file that have a time stamp within the span entered. If you do not enter an end time, the current time is used. You enter three characters for the month time field, and digits for the year and day time fields. |
| <b>Step 3</b> | switch# <b>show logging nvram</b> [ <b>last</b> <i>number-lines</i> ]                                                                 | Displays the messages in the NVRAM. To limit the number of lines displayed, you can enter the last number of lines to display. You can specify from 1 to 100 for the last number of lines.                                                         |
| <b>Step 4</b> | switch# <b>clear logging logfile</b>                                                                                                  | Clears the contents of the log file.                                                                                                                                                                                                               |
| <b>Step 5</b> | switch# <b>clear logging nvram</b>                                                                                                    | Clears the logged messages in NVRAM.                                                                                                                                                                                                               |

The following example shows how to display messages in a log file:

```
switch# show logging last 40
switch# show logging logfile start-time 2007 nov 1 15:10:0
switch# show logging nvram last 10
```

The following example shows how to clear messages in a log file:

```
switch# clear logging logfile
switch# clear logging nvram
```

## Verifying System Message Logging Configuration

To display system message logging configuration information, perform one of the following tasks:

| Command                                                                                                                               | Purpose                                                                    |
|---------------------------------------------------------------------------------------------------------------------------------------|----------------------------------------------------------------------------|
| switch# <b>show logging console</b>                                                                                                   | Displays the console logging configuration.                                |
| switch# <b>show logging info</b>                                                                                                      | Displays the logging configuration.                                        |
| switch# <b>show logging internal info</b>                                                                                             | Displays the syslog distribution information.                              |
| switch# <b>show logging last</b> <i>number-lines</i>                                                                                  | Displays the last number of lines of the log file.                         |
| switch# <b>show logging level</b> [ <i>facility</i> ]                                                                                 | Displays the facility logging severity level configuration.                |
| switch# <b>show logging logfile</b> [ <b>start-time</b> <i>yyyy mmm dd hh:mm:ss</i> ] [ <b>end-time</b> <i>yyyy mmm dd hh:mm:ss</i> ] | Displays the messages in the log file.                                     |
| switch# <b>show logging module</b>                                                                                                    | Displays the module logging configuration.                                 |
| switch# <b>show logging monitor</b>                                                                                                   | Displays the monitor logging configuration.                                |
| switch# <b>show logging nvram</b> [ <b>last</b> <i>number-lines</i> ]                                                                 | Displays the messages in the NVRAM log.                                    |
| switch# <b>show logging pending</b>                                                                                                   | Displays the syslog server pending distribution configuration.             |
| switch# <b>show logging pending-diff</b>                                                                                              | Displays the syslog server pending distribution configuration differences. |
| switch# <b>show logging server</b>                                                                                                    | Displays the syslog server configuration.                                  |
| switch# <b>show logging session</b>                                                                                                   | Displays the logging session status.                                       |
| switch# <b>show logging status</b>                                                                                                    | Displays the logging status.                                               |
| switch# <b>show logging timestamp</b>                                                                                                 | Displays the logging time-stamp units configuration.                       |

## Default System Message Logging Settings

The following table lists the default settings for system message logging parameters.

**Table 42: Default System Message Logging Parameters**

| <b>Parameters</b>                        | <b>Default</b>                              |
|------------------------------------------|---------------------------------------------|
| Console logging                          | Enabled at severity level 2                 |
| Monitor logging                          | Enabled at severity level 2                 |
| Log file logging                         | Enabled to log:messages at severity level 5 |
| Module logging                           | Enabled at severity level 5                 |
| Facility logging                         | Enabled;                                    |
| Time-stamp units                         | Seconds                                     |
| syslog server logging                    | Disabled                                    |
| syslog server configuration distribution | Disabled                                    |





## CHAPTER 28

# Configuring Smart Call Home

This chapter contains the following sections:

- [Configuring Smart Call Home, page 385](#)

## Configuring Smart Call Home

### Information About Call Home

Call Home provides e-mail-based notification of critical system events. Cisco Nexus 5000 Series switches provide a range of message formats for optimal compatibility with pager services, standard e-mail, or XML-based automated parsing applications. You can use this feature to page a network support engineer, e-mail a Network Operations Center, or use Cisco Smart Call Home services to automatically generate a case with the Technical Assistance Center.

### Call Home Overview

You can use Call Home to notify an external entity when an important event occurs on your device. Call Home delivers alerts to multiple recipients that you configure in *destination profiles*.

Call Home includes a fixed set of predefined alerts on your switch. These alerts are grouped into alert groups and CLI commands to be assigned to execute when an alert in an alert group occurs. The switch includes the command output in the transmitted Call Home message.

The Call Home feature offers the following advantages:

- Automatic execution and attachment of relevant CLI command output.
- Multiple message format options such as the following:
  - Short Text—Suitable for pagers or printed reports.
  - Full Text—Fully formatted message information suitable for human reading.
  - XML—Matching readable format that uses the Extensible Markup Language (XML) and the Adaptive Messaging Language (AML) XML schema definition (XSD). The XML format enables communication with the Cisco Systems Technical Assistance Center (Cisco-TAC).

- Multiple concurrent message destinations. You can configure up to 50 e-mail destination addresses for each destination profile.

## Destination Profiles

A destination profile includes the following information:

- One or more alert groups—The group of alerts that trigger a specific Call Home message if the alert occurs.
- One or more e-mail destinations—The list of recipients for the Call Home messages generated by alert groups assigned to this destination profile.
- Message format—The format for the Call Home message (short text, full text, or XML).
- Message severity level—The Call Home severity level that the alert must meet before the switch generates a Call Home message to all e-mail addresses in the destination profile. The Cisco Nexus 5000 Series switch does not generate an alert if the Call Home severity level of the alert is lower than the message severity level set for the destination profile.

You can also configure a destination profile to allow periodic inventory update messages by using the inventory alert group that will send out periodic messages daily, weekly, or monthly.

Cisco Nexus 5000 Series switches support the following predefined destination profiles:

- CiscoTAC-1—Supports the Cisco-TAC alert group in XML message format.
- full-text-destination—Supports the full text message format.
- short-text-destination—Supports the short text message format.

## Call Home Alert Groups

An alert group is a predefined subset of Call Home alerts that are supported in all Cisco Nexus 5000 Series switches. Alert groups allow you to select the set of Call Home alerts that you want to send to a predefined or custom destination profile. The switch sends Call Home alerts to e-mail destinations in a destination profile only if that Call Home alert belongs to one of the alert groups associated with that destination profile and if the alert has a Call Home message severity at or above the message severity set in the destination profile.

The following table lists supported alert groups and the default CLI command output included in Call Home messages generated for the alert group.

**Table 43: Alert Groups and Executed Commands**

| Alert Group | Description                                                                   | Executed Commands                                                    |
|-------------|-------------------------------------------------------------------------------|----------------------------------------------------------------------|
| Cisco-TAC   | All critical alerts from the other alert groups destined for Smart Call Home. | Execute commands based on the alert group that originates the alert. |

| <b>Alert Group</b>  | <b>Description</b>                                                                                                                                                                                            | <b>Executed Commands</b>                                                                                                                              |
|---------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------|
| Diagnostic          | Events generated by diagnostics.                                                                                                                                                                              | <b>show diagnostic result module all detail</b><br><b>show moduleshow version</b><br><b>show tech-support platform callhome</b>                       |
| Supervisor hardware | Events related to supervisor modules.                                                                                                                                                                         | <b>show diagnostic result module all detail</b><br><b>show moduleshow version</b><br><b>show tech-support platform callhome</b>                       |
| Linecard hardware   | Events related to standard or intelligent switching modules.                                                                                                                                                  | <b>show diagnostic result module all detail</b><br><b>show moduleshow version</b><br><b>show tech-support platform callhome</b>                       |
| Configuration       | Periodic events related to configuration.                                                                                                                                                                     | <b>show version</b><br><b>show module</b><br><b>show running-config all</b><br><b>show startup-config</b>                                             |
| System              | Events generated by failure of a software system that is critical to unit operation.                                                                                                                          | <b>show system redundancy status</b><br><b>show tech-support</b>                                                                                      |
| Environmental       | Events related to power, fan, and environment-sensing elements such as temperature alarms.                                                                                                                    | <b>show environment</b><br><b>show logging last 1000</b><br><b>show module show version</b><br><b>show tech-support platform callhome</b>             |
| Inventory           | Inventory status that is provided whenever a unit is cold booted, or when FRUs are inserted or removed. This alert is considered a noncritical event, and the information is used for status and entitlement. | <b>show module</b><br><b>show version</b><br><b>show license usage</b><br><b>show inventory</b><br><b>show sprom all</b><br><b>show system uptime</b> |

Call Home maps the syslog severity level to the corresponding Call Home severity level for syslog port group messages

You can customize predefined alert groups to execute additional CLI **show** commands when specific events occur and send that **show** output with the Call Home message.

You can add **show** commands only to full text and XML destination profiles. Short text destination profiles do not support additional **show** commands because they only allow 128 bytes of text.

### Related Topics

[Call Home Message Levels](#) , on page 388

## Call Home Message Levels

Call Home allows you to filter messages based on their level of urgency. You can associate each destination profile (predefined and user defined) with a Call Home message level threshold. The switch does not generate any Call Home messages with a value lower than this threshold for the destination profile. The Call Home message level ranges from 0 (lowest level of urgency) to 9 (highest level of urgency), and the default is 0 (Cisco Nexus 5000 Series sends all messages).

Call Home messages that are sent for syslog alert groups have the syslog severity level mapped to the Call Home message level.



#### Note

Call Home does not change the syslog message level in the message text.

The following table lists each Call Home message level keyword and the corresponding syslog level for the syslog port alert group.

**Table 44: Severity and syslog Level Mapping**

| Call Home Level | Keyword      | syslog Level  | Description                                                           |
|-----------------|--------------|---------------|-----------------------------------------------------------------------|
| 9               | Catastrophic | N/A           | Network-wide catastrophic failure.                                    |
| 8               | Disaster     | N/A           | Significant network impact.                                           |
| 7               | Fatal        | Emergency (0) | System is unusable.                                                   |
| 6               | Critical     | Alert (1)     | Critical conditions that indicate that immediate attention is needed. |
| 5               | Major        | Critical (2)  | Major conditions.                                                     |
| 4               | Minor        | Error (3)     | Minor conditions.                                                     |
| 3               | Warning      | Warning (4)   | Warning conditions.                                                   |

| Call Home Level | Keyword      | syslog Level    | Description                                                                          |
|-----------------|--------------|-----------------|--------------------------------------------------------------------------------------|
| 2               | Notification | Notice (5)      | Basic notification and informational messages. Possibly independently insignificant. |
| 1               | Normal       | Information (6) | Normal event signifying return to normal state.                                      |
| 0               | Debugging    | Debug (7)       | Debugging messages.                                                                  |

### Obtaining Smart Call Home

If you have a service contract directly with Cisco Systems, you can register your devices for the Smart Call Home service. Smart Call Home provides fast resolution of system problems by analyzing Call Home messages sent from your devices and providing background information and recommendations. For issues that can be identified as known, particularly GOLD diagnostics failures, Automatic Service Requests will be generated with the Cisco-TAC.

Smart Call Home offers the following features:

- Continuous device health monitoring and real-time diagnostic alerts.
- Analysis of Call Home messages from your device and, where appropriate, Automatic Service Request generation, routed to the appropriate TAC team, including detailed diagnostic information to speed problem resolution.
- Secure message transport directly from your device or through a downloadable Transport Gateway (TG) aggregation point. You can use a TG aggregation point in cases that require support for multiple devices or in cases where security requirements mandate that your devices may not be connected directly to the Internet.
- Web-based access to Call Home messages and recommendations, inventory and configuration information for all Call Home devices. Provides access to associated field notices, security advisories and end-of-life information.

You need the following items to register:

- The SMARTnet contract number for your switch.
- Your e-mail address
- Your Cisco.com ID

For more information about Smart Call Home, see the Smart Call Home page at this URL: <http://www.cisco.com/go/smartcall/>

### Prerequisites for Call Home

Call Home has the following prerequisites:

- You must configure an e-mail server.

- You must configure the contact name (SNMP server contact), phone, and street address information before you enable Call Home. This step is required to determine the origin of messages received.
- Your switch must have IP connectivity to an e-mail server.
- If you use Smart Call Home, you need an active service contract for the device that you are configuring.

## Configuration Guidelines and Limitations

Call Home has the following configuration guidelines and limitations:

- If there is no IP connectivity or if the interface in the VRF to the profile destination is down, the switch cannot send the Call Home message.
- Operates with any SMTP server.

## Configuring Call Home

### Procedures for Configuring Call Home

#### SUMMARY STEPS

1. Assign contact information.
2. Configure destination profiles.
3. Associate one or more alert groups to each profile.
4. (Optional) Add additional **show** commands to the alert groups.
5. Configure transport options.
6. Enable Call Home.
7. (Optional) Test Call Home messages.

#### DETAILED STEPS

- 
- |               |                                                                     |
|---------------|---------------------------------------------------------------------|
| <b>Step 1</b> | Assign contact information.                                         |
| <b>Step 2</b> | Configure destination profiles.                                     |
| <b>Step 3</b> | Associate one or more alert groups to each profile.                 |
| <b>Step 4</b> | (Optional) Add additional <b>show</b> commands to the alert groups. |
| <b>Step 5</b> | Configure transport options.                                        |
| <b>Step 6</b> | Enable Call Home.                                                   |
| <b>Step 7</b> | (Optional) Test Call Home messages.                                 |
-

## Configuring Contact Information

You must configure the e-mail, phone, and street address information for Call Home. You can optionally configure the contract ID, customer ID, site ID, and switch priority information.

### SUMMARY STEPS

1. switch# **configuration terminal**
2. switch(config)# **snmp-server contact** *sys-contact*
3. switch(config)# **callhome**
4. switch(config-callhome)# **email-contact** *email-address*
5. switch(config-callhome)# **phone-contact** *international-phone-number*
6. switch(config-callhome)# **streetaddress** *address*
7. (Optional) switch(config-callhome)# **contract-id** *contract-number*
8. (Optional) switch(config-callhome)# **customer-id** *customer-number*
9. (Optional) switch(config-callhome)# **site-id** *site-number*
10. (Optional) switch(config-callhome)# **switch-priority** *number*
11. (Optional) switch# **show callhome**
12. (Optional) switch# **copy running-config startup-config**

### DETAILED STEPS

|               | Command or Action                                                                  | Purpose                                                                                                                                                                                                                                                                                              |
|---------------|------------------------------------------------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Step 1</b> | switch# <b>configuration terminal</b>                                              | Enters configuration mode.                                                                                                                                                                                                                                                                           |
| <b>Step 2</b> | switch(config)# <b>snmp-server contact</b><br><i>sys-contact</i>                   | Configures the SNMP sysContact.                                                                                                                                                                                                                                                                      |
| <b>Step 3</b> | switch(config)# <b>callhome</b>                                                    | Enters callhome configuration mode.                                                                                                                                                                                                                                                                  |
| <b>Step 4</b> | switch(config-callhome)# <b>email-contact</b><br><i>email-address</i>              | Configures the e-mail address for the primary person responsible for the switch. Up to 255 alphanumeric characters are accepted in e-mail address format.<br><br><b>Note</b> You can use any valid e-mail address. The address cannot contain spaces.                                                |
| <b>Step 5</b> | switch(config-callhome)# <b>phone-contact</b><br><i>international-phone-number</i> | Configures the phone number in international phone number format for the primary person responsible for the device. Up to 17 alphanumeric characters are accepted in international format.<br><br><b>Note</b> The phone number cannot contain spaces. Be sure to use the + prefix before the number. |
| <b>Step 6</b> | switch(config-callhome)# <b>streetaddress</b><br><i>address</i>                    | Configures the street address as an alphanumeric string with white paces for the primary person responsible for the switch. Up to 255 alphanumeric characters are accepted, including spaces.                                                                                                        |

|                | Command or Action                                                     | Purpose                                                                                                                                                               |
|----------------|-----------------------------------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Step 7</b>  | switch(config-callhome)# <b>contract-id</b><br><i>contract-number</i> | (Optional)<br>Configures the contract number for this switch from the service agreement. The contract number can be up to 255 alphanumeric characters in free format. |
| <b>Step 8</b>  | switch(config-callhome)# <b>customer-id</b><br><i>customer-number</i> | (Optional)<br>Configures the customer number for this switch from the service agreement. The customer number can be up to 255 alphanumeric characters in free format. |
| <b>Step 9</b>  | switch(config-callhome)# <b>site-id</b><br><i>site-number</i>         | (Optional)<br>Configures the site number for this switch. The site number can be up to 255 alphanumeric characters in free format.                                    |
| <b>Step 10</b> | switch(config-callhome)# <b>switch-priority</b><br><i>number</i>      | (Optional)<br>Configures the switch priority for this switch. The range is from 0 to 7, with 0 being the highest priority and 7 the lowest. The default is 7.         |
| <b>Step 11</b> | switch# <b>show callhome</b>                                          | (Optional)<br>Displays a summary of the Call Home configuration.                                                                                                      |
| <b>Step 12</b> | switch# <b>copy running-config</b><br><b>startup-config</b>           | (Optional)<br>Saves this configuration change.                                                                                                                        |

This example shows how to configure the contact information for Call Home:

```
switch# configuration terminal
switch(config)# snmp-server contact personname@companyname.com
switch(config)# callhome
switch(config-callhome)# email-contact personname@companyname.com
switch(config-callhome)# phone-contact +1-800-123-4567
switch(config-callhome)# street-address 123 Anystreet St., Anycity, Anywhere
```

## Creating a Destination Profile

You must create a user-defined destination profile and configure the message format for that new destination profile.

## SUMMARY STEPS

1. switch# **configuration terminal**
2. switch(config)# **callhome**
3. switch(config-callhome)# **destination-profile** {ciscoTAC-1 {**alert-group** *group* | **email-addr** *address* | **http** *URL* | **transport-method** {**email** | **http**}} | *profile-name* {**alert-group** *group* | **email-addr** *address* | **format** {**XML** | **full-txt** | **short-txt**} | **http** *URL* | **message-level** *level* | **message-size** *size* | **transport-method** {**email** | **http**}} | **full-txt-destination** {**alert-group** *group* | **email-addr** *address* | **http** *URL* | **message-level** *level* | **message-size** *size* | **transport-method** {**email** | **http**}} | **short-txt-destination** {**alert-group** *group* | **email-addr** *address* | **http** *URL* | **message-level** *level* | **message-size** *size* | **transport-method** {**email** | **http**}}}
4. (Optional) switch# **show callhome destination-profile** [*profile name*]
5. (Optional) switch# **copy running-config startup-config**

## DETAILED STEPS

|               | Command or Action                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         | Purpose                                                                                                                                                                                                                                                       |
|---------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Step 1</b> | switch# <b>configuration terminal</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     | Enters configuration mode.                                                                                                                                                                                                                                    |
| <b>Step 2</b> | switch(config)# <b>callhome</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           | Enters callhome configuration mode.                                                                                                                                                                                                                           |
| <b>Step 3</b> | switch(config-callhome)# <b>destination-profile</b> {ciscoTAC-1 { <b>alert-group</b> <i>group</i>   <b>email-addr</b> <i>address</i>   <b>http</b> <i>URL</i>   <b>transport-method</b> { <b>email</b>   <b>http</b> }}   <i>profile-name</i> { <b>alert-group</b> <i>group</i>   <b>email-addr</b> <i>address</i>   <b>format</b> { <b>XML</b>   <b>full-txt</b>   <b>short-txt</b> }   <b>http</b> <i>URL</i>   <b>message-level</b> <i>level</i>   <b>message-size</b> <i>size</i>   <b>transport-method</b> { <b>email</b>   <b>http</b> }}   <b>full-txt-destination</b> { <b>alert-group</b> <i>group</i>   <b>email-addr</b> <i>address</i>   <b>http</b> <i>URL</i>   <b>message-level</b> <i>level</i>   <b>message-size</b> <i>size</i>   <b>transport-method</b> { <b>email</b>   <b>http</b> }}   <b>short-txt-destination</b> { <b>alert-group</b> <i>group</i>   <b>email-addr</b> <i>address</i>   <b>http</b> <i>URL</i>   <b>message-level</b> <i>level</i>   <b>message-size</b> <i>size</i>   <b>transport-method</b> { <b>email</b>   <b>http</b> }}} | Creates a new destination profile and sets the message format for the profile. The profile-name can be any alphanumeric string up to 31 characters.<br><br>For further details about this command, see the <i>Cisco Nexus 5000 Series Command Reference</i> . |
| <b>Step 4</b> | switch# <b>show callhome destination-profile</b> [ <i>profile name</i> ]                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  | (Optional)<br>Displays information about one or more destination profiles.                                                                                                                                                                                    |
| <b>Step 5</b> | switch# <b>copy running-config startup-config</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         | (Optional)<br>Saves this configuration change.                                                                                                                                                                                                                |

This example shows how to create a destination profile for Call Home:

```
switch# configuration terminal
switch(config)# callhome
switch(config-callhome)# destination-profile Noc101 format full-text
```

## Modifying a Destination Profile

You can modify the following attributes for a predefined or user-defined destination profile:

- Destination address—The actual address, pertinent to the transport mechanism, to which the alert should be sent.
- Message formatting—The message format used for sending the alert (full text, short text, or XML).
- Message level—The Call Home message severity level for this destination profile.
- Message size—The allowed length of a Call Home message sent to the e-mail addresses in this destination profile.



### Note

You cannot modify or delete the CiscoTAC-1 destination profile.

## SUMMARY STEPS

1. switch# **configuration terminal**
2. switch(config)# **callhome**
3. switch(config-callhome)# **destination-profile** {*name* | **full-txt-destination** | **short-txt-destination**} **email-addr** *address*
4. **destination-profile** {*name* | **full-txt-destination** | **short-txt-destination**} **message-level** *number*
5. switch(config-callhome)# **destination-profile** {*name* | **full-txt-destination** | **short-txt-destination**} **message-size** *number*
6. (Optional) switch# **show callhome destination-profile** [**profile name**]
7. (Optional) switch# **copy running-config startup-config**

## DETAILED STEPS

|               | Command or Action                                                                                                                                                  | Purpose                                                                                                                                                                                                                                                              |
|---------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Step 1</b> | switch# <b>configuration terminal</b>                                                                                                                              | Enters configuration mode.                                                                                                                                                                                                                                           |
| <b>Step 2</b> | switch(config)# <b>callhome</b>                                                                                                                                    | Enters callhome configuration mode.                                                                                                                                                                                                                                  |
| <b>Step 3</b> | switch(config-callhome)# <b>destination-profile</b> { <i>name</i>   <b>full-txt-destination</b>   <b>short-txt-destination</b> } <b>email-addr</b> <i>address</i>  | Configures an e-mail address for a user-defined or predefined destination profile. You can configure up to 50 e-mail addresses in a destination profile.                                                                                                             |
| <b>Step 4</b> | <b>destination-profile</b> { <i>name</i>   <b>full-txt-destination</b>   <b>short-txt-destination</b> } <b>message-level</b> <i>number</i>                         | Configures the Call Home message severity level for this destination profile. The switch sends only alerts that have a matching or higher Call Home severity level to destinations in this profile. The range is from 0 to 9, where 9 is the highest severity level. |
| <b>Step 5</b> | switch(config-callhome)# <b>destination-profile</b> { <i>name</i>   <b>full-txt-destination</b>   <b>short-txt-destination</b> } <b>message-size</b> <i>number</i> | Configures the maximum message size for this destination profile. The range is from 0 to 5000000 for full-txt-destination and the default is 2500000; from 0 to 100000 for short-txt-destination and                                                                 |

|               | Command or Action                                                           | Purpose                                                                    |
|---------------|-----------------------------------------------------------------------------|----------------------------------------------------------------------------|
|               |                                                                             | the default is 4000; 5000000 for CiscoTAC-1, which is not changeable.      |
| <b>Step 6</b> | switch# <b>show callhome destination-profile</b><br>[ <b>profile name</b> ] | (Optional)<br>Displays information about one or more destination profiles. |
| <b>Step 7</b> | switch# <b>copy running-config startup-config</b>                           | (Optional)<br>Saves this configuration change.                             |

This example shows how to modify a destination profile for Call Home:

```
switch# configuration terminal
switch(config)# callhome
switch(config-callhome)# destination-profile full-text-destination email-addr
person@example.com
switch(config-callhome)# destination-profile full-text-destination message-level 5
switch(config-callhome)# destination-profile full-text-destination message-size 10000
```

### Related Topics

[Associating an Alert Group with a Destination Profile, on page 395](#)

## Associating an Alert Group with a Destination Profile

To associate one or more alert groups with a destination profile, perform this task:

### SUMMARY STEPS

1. switch# **configuration terminal**
2. switch(config)# **callhome**
3. switch(config-callhome)# **destination-profile** *name* **alert-group** {**All** | **Cisco-TAC** | **Configuration** | **Diagnostic** | **Environmental** | **Inventory** | **License** | **Linecard-Hardware** | **Supervisor-Hardware** | **Syslog-group-port** | **System** | **Test**}
4. (Optional) switch# **show callhome destination-profile** [**profile name**]
5. (Optional) switch# **copy running-config startup-config**

### DETAILED STEPS

|               | Command or Action                     | Purpose                             |
|---------------|---------------------------------------|-------------------------------------|
| <b>Step 1</b> | switch# <b>configuration terminal</b> | Enters configuration mode.          |
| <b>Step 2</b> | switch(config)# <b>callhome</b>       | Enters callhome configuration mode. |

|               | Command or Action                                                                                                                                                                                                                                                                                                                              | Purpose                                                                                                                                         |
|---------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Step 3</b> | switch(config-callhome)# <b>destination-profile</b> <i>name</i><br><b>alert-group</b> { <b>All</b>   <b>Cisco-TAC</b>   <b>Configuration</b>   <b>Diagnostic</b>   <b>Environmental</b>   <b>Inventory</b>   <b>License</b>   <b>Linecard-Hardware</b>   <b>Supervisor-Hardware</b>   <b>Syslog-group-port</b>   <b>System</b>   <b>Test</b> } | Associates an alert group with this destination profile. Use the <b>All</b> keyword to associate all alert groups with the destination profile. |
| <b>Step 4</b> | switch# <b>show callhome destination-profile</b> [ <b>profile name</b> ]                                                                                                                                                                                                                                                                       | (Optional)<br>Displays information about one or more destination profiles.                                                                      |
| <b>Step 5</b> | switch# <b>copy running-config startup-config</b>                                                                                                                                                                                                                                                                                              | (Optional)<br>Saves this configuration change.                                                                                                  |

This example shows how to associate all alert groups with the destination profile Noc101:

```
switch# configuration terminal
switch(config)# callhome
switch(config-callhome)# destination-profile Noc101 alert-group All
```

### Adding show Commands to an Alert Group

You can assign a maximum of five user-defined CLI **show** commands to an alert group.

### SUMMARY STEPS

1. switch# **configuration terminal**
2. switch(config)# **callhome**
3. switch(config-callhome)# **alert-group** {**Configuration** | **Diagnostic** | **Environmental** | **Inventory** | **License** | **Linecard-Hardware** | **Supervisor-Hardware** | **Syslog-group-port** | **System** | **Test**} **user-def-cmd** *show-cmd*
4. (Optional) switch# **show callhome user-def-cmds**
5. (Optional) switch# **copy running-config startup-config**

### DETAILED STEPS

|               | Command or Action                                                                                                                                                                                                                                                                                        | Purpose                                                                                                                                                                                                                                                                                                      |
|---------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Step 1</b> | switch# <b>configuration terminal</b>                                                                                                                                                                                                                                                                    | Enters configuration mode.                                                                                                                                                                                                                                                                                   |
| <b>Step 2</b> | switch(config)# <b>callhome</b>                                                                                                                                                                                                                                                                          | Enters callhome configuration mode.                                                                                                                                                                                                                                                                          |
| <b>Step 3</b> | switch(config-callhome)# <b>alert-group</b> { <b>Configuration</b>   <b>Diagnostic</b>   <b>Environmental</b>   <b>Inventory</b>   <b>License</b>   <b>Linecard-Hardware</b>   <b>Supervisor-Hardware</b>   <b>Syslog-group-port</b>   <b>System</b>   <b>Test</b> } <b>user-def-cmd</b> <i>show-cmd</i> | Adds the <b>show</b> command output to any Call Home messages sent for this alert group. You must enclose the <b>show</b> command in double quotes. Only valid <b>show</b> commands are accepted.<br><b>Note</b> You cannot add user-defined CLI <b>show</b> commands to the CiscoTAC-1 destination profile. |

|               | Command or Action                                 | Purpose                                                                                               |
|---------------|---------------------------------------------------|-------------------------------------------------------------------------------------------------------|
| <b>Step 4</b> | switch# <b>show callhome user-def-cmds</b>        | (Optional)<br>Displays information about all user-defined <b>show</b> commands added to alert groups. |
| <b>Step 5</b> | switch# <b>copy running-config startup-config</b> | (Optional)<br>Saves this configuration change.                                                        |

This example shows how to add the **show ip routing** command to the Cisco-TAC alert group:

```
switch# configuration terminal
switch(config)# callhome
switch(config-callhome)# alert-group Configuration user-def-cmd "show ip routing"
```

### Configuring E-Mail

You must configure the SMTP server address for the Call Home functionality to work. You can also configure the from and reply-to e-mail addresses.

### SUMMARY STEPS

1. switch# **configuration terminal**
2. switch(config)# **callhome**
3. switch(config-callhome)# **transport email smtp-server ip-address [port number] [use-vrf vrf-name]**
4. (Optional) switch(config-callhome)# **transport email from email-address**
5. (Optional) switch(config-callhome)# **transport email reply-to email-address**
6. (Optional) switch# **show callhome transport-email**
7. (Optional) switch# **copy running-config startup-config**

### DETAILED STEPS

|               | Command or Action                                                                                       | Purpose                                                                                                                                                                                                                                                                                                                    |
|---------------|---------------------------------------------------------------------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Step 1</b> | switch# <b>configuration terminal</b>                                                                   | Enters configuration mode.                                                                                                                                                                                                                                                                                                 |
| <b>Step 2</b> | switch(config)# <b>callhome</b>                                                                         | Enters callhome configuration mode.                                                                                                                                                                                                                                                                                        |
| <b>Step 3</b> | switch(config-callhome)# <b>transport email smtp-server ip-address [port number] [use-vrf vrf-name]</b> | Configures the SMTP server as either the domain name server (DNS) name, IPv4 address, or IPv6 address). Optionally you can configure the port number. The port ranges is from 1 to 65535. The default port number is 25.<br><br>Also optionally you can configure the VRF to use when communicating with this SMTP server. |

|               | Command or Action                                                             | Purpose                                                                          |
|---------------|-------------------------------------------------------------------------------|----------------------------------------------------------------------------------|
| <b>Step 4</b> | switch(config-callhome)# <b>transport email from</b> <i>email-address</i>     | (Optional)<br>Configures the e-mail from field for Call Home messages.           |
| <b>Step 5</b> | switch(config-callhome)# <b>transport email reply-to</b> <i>email-address</i> | (Optional)<br>Configures the e-mail reply-to field for Call Home messages.       |
| <b>Step 6</b> | switch# <b>show callhome transport-email</b>                                  | (Optional)<br>Displays information about the e-mail configuration for Call Home. |
| <b>Step 7</b> | switch# <b>copy running-config startup-config</b>                             | (Optional)<br>Saves this configuration change.                                   |

This example shows how to configure the e-mail options for Call Home messages:

```
switch# configuration terminal
switch(config)# callhome
switch(config-callhome)# transport email smtp-server 192.0.2.10 use-vrf Red
switch(config-callhome)# transport email from person@example.com
switch(config-callhome)# transport email reply-to person@example.com
```

## Configuring Periodic Inventory Notification

You can configure the switch to periodically send a message with an inventory of all software services currently enabled and running on the device along with hardware inventory information. The switch generates two Call Home notifications; periodic configuration messages and periodic inventory messages.

### SUMMARY STEPS

1. switch# **configuration terminal**
2. switch(config)# **callhome**
3. switch(config-callhome)# **periodic-inventory notification** [interval *days*] [timeofday *time*]
4. (Optional) switch# **show callhome**
5. (Optional) switch# **copy running-config startup-config**

### DETAILED STEPS

|               | Command or Action                     | Purpose                             |
|---------------|---------------------------------------|-------------------------------------|
| <b>Step 1</b> | switch# <b>configuration terminal</b> | Enters configuration mode.          |
| <b>Step 2</b> | switch(config)# <b>callhome</b>       | Enters callhome configuration mode. |

|               | Command or Action                                                                                                | Purpose                                                                                                                                             |
|---------------|------------------------------------------------------------------------------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Step 3</b> | switch(config-callhome)# <b>periodic-inventory notification</b> [interval <i>days</i> ] [timeofday <i>time</i> ] | Configures the periodic inventory messages. The interval range is from 1 to 30 days. The default is 7 days. The timeofday value is in HH:MM format. |
| <b>Step 4</b> | switch# <b>show callhome</b>                                                                                     | (Optional)<br>Displays information about Call Home.                                                                                                 |
| <b>Step 5</b> | switch# <b>copy running-config startup-config</b>                                                                | (Optional)<br>Saves this configuration change.                                                                                                      |

This example shows how to configure the periodic inventory messages to generate every 20 days:

```
switch# configuration terminal
switch(config)# callhome
switch(config-callhome)# periodic-inventory notification interval 20
```

### Disabling Duplicate Message Throttle

You can limit the number of duplicate messages received for the same event. By default, the switch limits the number of duplicate messages received for the same event. If the number of duplicate messages sent exceeds 30 messages within a 2-hour time frame, then the switch discards further messages for that alert type.

| Command                                                       | Purpose                                                                  |
|---------------------------------------------------------------|--------------------------------------------------------------------------|
| switch(config-callhome)# <b>no duplicate-message throttle</b> | Disables duplicate message throttling for Call Home. Enabled by default. |

### Enabling or Disabling Call Home

Once you have configured the contact information, you can enable the Call Home function in callhome configuration mode.

| Command                                | Purpose                                 |
|----------------------------------------|-----------------------------------------|
| switch(config-callhome)# <b>enable</b> | Enables Call Home. Disabled by default. |

You can disable Call Home in the callhome configuration mode.

| Command                                   | Purpose                                 |
|-------------------------------------------|-----------------------------------------|
| switch(config-callhome)# <b>no enable</b> | Disables Call Home. Disabled by default |

You can enable Call Home distribution using CFS in the callhome configuration mode.

| Command                                    | Purpose                                                        |
|--------------------------------------------|----------------------------------------------------------------|
| switch(config-callhome)# <b>distribute</b> | Enables Call Home distribution using CFS. Disabled by default. |

You can commit Call Home configuration changes and distribute using CFS in the callhome configuration mode.

| Command                                | Purpose                                                                                          |
|----------------------------------------|--------------------------------------------------------------------------------------------------|
| switch(config-callhome)# <b>commit</b> | Commits Call Home configuration changes and distributes the changes to call CFS-enabled devices. |

You can discard Call Home configuration changes and release the CFS lock in callhome configuration mode.

| Command                               | Purpose                                                                                                                                                                         |
|---------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| switch(config-callhome)# <b>abort</b> | Discards Call Home configuration changes and releases the CFS lock. Use this command if you are the CFS lock owner or if you are logged into the device that holds the CFS lock |

### Testing Call Home Communications

You can generate a test message to test your Call Home communications.

| Command                                                  | Purpose                                                                                                                              |
|----------------------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------|
| switch(config-callhome)# <b>callhome send diagnostic</b> | Sends the specified Call Home test message to all configured destinations.                                                           |
| switch(config-callhome)# <b>callhome test</b>            | Sends a test message to all configured destinations. <b>callhome test</b> and <b>callhome test inventory</b> commands are supported. |

### Verifying Call Home Configuration

To display Call Home configuration information, perform one of the following tasks:

| Command                                                      | Purpose                                              |
|--------------------------------------------------------------|------------------------------------------------------|
| switch# <b>show callhome</b>                                 | Displays the status for Call Home.                   |
| switch# <b>show callhome destination-profile</b> <i>name</i> | Displays one or more Call Home destination profiles. |

| Command                                                      | Purpose                                                                           |
|--------------------------------------------------------------|-----------------------------------------------------------------------------------|
| switch# <b>show callhome merge</b>                           | Displays the status of the last CFS merge for Call Home.                          |
| switch# <b>show callhome pending</b>                         | Displays the Call Home configuration changes in the pending CFS database.         |
| switch# <b>show callhome pending-diff</b>                    | Displays the differences between the pending and running Call Home configuration. |
| switch# <b>show callhome session</b>                         | Displays the status of the last Call Home CFS command.                            |
| switch# <b>show callhome status</b>                          | Displays the Call Home status.                                                    |
| switch# <b>show callhome transport-email</b>                 | Displays the e-mail configuration for Call Home.                                  |
| switch# <b>show callhome user-def-cmds</b>                   | Displays CLI commands added to any alert groups.                                  |
| switch# <b>show running-config [callhome   callhome-all]</b> | Displays the running configuration for Call Home.                                 |
| switch# <b>show startup-config callhome</b>                  | Displays the startup configuration for Call Home.                                 |
| switch# <b>show tech-support callhome</b>                    | Displays the technical support output for Call Home.                              |

## Default Call Home Settings

The following table lists the default settings for Call Home parameters.

**Table 45: Default Call Home Parameters**

| Parameters                                                        | Default                                                                                                                              |
|-------------------------------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------|
| Destination message size for a message sent in full text format.  | 4000000                                                                                                                              |
| Destination message size for a message sent in XML format.        | 4000000                                                                                                                              |
| Destination message size for a message sent in short text format. | 4000                                                                                                                                 |
| SMTP server port number if no port is specified.                  | 25                                                                                                                                   |
| Alert group association with profile.                             | All for full-text-destination and short-text-destination profiles. The cisco-tac alert group for the CiscoTAC-1 destination profile. |

| Parameters               | Default  |
|--------------------------|----------|
| Format type.             | XML      |
| Call Home message level. | 0 (zero) |

## Additional References

### Call Home Message Formats

Call Home supports the following message formats:

- Short Text Message Format
- Common Fields for All Full Text and XML Messages
- Inserted Fields for a Reactive or Proactive Event Message
- Inserted Fields for an Inventory Event Message
- Inserted Fields for a User-Generated Test Message

The following table describes the short text formatting option for all message types.

**Table 46: Short Text Message Format**

| Data Item               | Description                                        |
|-------------------------|----------------------------------------------------|
| Device identification   | Configured device name                             |
| Date/time stamp         | Time stamp of the triggering event                 |
| Error isolation message | Plain English description of triggering event      |
| Alarm urgency level     | Error level such as that applied to system message |

The following table describes the common event message format for full text or XML.

**Table 47: Common Fields for All Full Text and XML Messages**

| Data Item(Plain Text and XML) | Description(Plain Text and XML)                                                                                | XML Tag (XML Only) |
|-------------------------------|----------------------------------------------------------------------------------------------------------------|--------------------|
| Time stamp                    | Date and time stamp of event in ISO time notation:<br><i>YYYY-MM-DD</i><br><i>HH:MM:SS</i><br><i>GMT+HH:MM</i> | /aml/header/time   |

| <b>Data Item(Plain Text and XML)</b> | <b>Description(Plain Text and XML)</b>                                   | <b>XML Tag (XML Only)</b> |
|--------------------------------------|--------------------------------------------------------------------------|---------------------------|
| Message name                         | Name of message. Specific event names are listed in the preceding table. | /aml/header/name          |
| Message type                         | Name of message type, such as reactive or proactive.                     | /aml/header/type          |
| Message group                        | Name of alert group, such as syslog.                                     | /aml/header/group         |
| Severity level                       | Severity level of message.                                               | /aml/header/level         |
| Source ID                            | Product type for routing. Specifically Catalyst 6500.                    | /aml/header/source        |

| Data Item(Plain Text and XML) | Description(Plain Text and XML)                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            | XML Tag (XML Only)          |
|-------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------------------------|
| Device ID                     | <p>Unique device identifier (UDI) for end device that generated the message. This field should be empty if the message is nonspecific to a device. The format is <i>type@Sid@serial</i>:</p> <ul style="list-style-type: none"> <li>• <i>type</i> is the product model number from backplane IDPROM.</li> <li>• <i>@</i> is a separator character.</li> <li>• <i>Sid</i> is C, identifying the serial ID as a chassis serial number.</li> <li>• <i>serial</i> is the number identified by the Sid field.</li> </ul> <p>An example is WSC609@C@12345678</p> | /aml/<br>header/deviceID    |
| Customer ID                   | Optional user-configurable field used for contract information or other ID by any support service.                                                                                                                                                                                                                                                                                                                                                                                                                                                         | /aml/<br>header/customerID  |
| Contract ID                   | Optional user-configurable field used for contract information or other ID by any support service.                                                                                                                                                                                                                                                                                                                                                                                                                                                         | /aml/ header<br>/contractID |

| Data Item(Plain Text and XML) | Description(Plain Text and XML)                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          | XML Tag (XML Only)   |
|-------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|----------------------|
| Site ID                       | Optional user-configurable field used for Cisco-supplied site ID or other data meaningful to alternate support service.                                                                                                                                                                                                                                                                                                                                                                                                  | /aml/ header/siteID  |
| Server ID                     | <p>If the message is generated from the device, this is the unique device identifier (UDI) of the device.</p> <p>The format is <i>type@Sid@serial</i>:</p> <ul style="list-style-type: none"> <li>• <i>type</i> is the product model number from backplane IDPROM.</li> <li>• <i>@</i> is a separator character.</li> <li>• <i>Sid</i> is C, identifying the serial ID as a chassis serial number.</li> <li>• <i>serial</i> is the number identified by the Sid field.</li> </ul> <p>An example is WSC600@C@12345678</p> | /aml/header/serverID |
| Message description           | Short text that describes the error.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     | /aml/body/msgDesc    |
| Device name                   | Node that experienced the event (host name of the device).                                                                                                                                                                                                                                                                                                                                                                                                                                                               | /aml/body/sysName    |

| Data Item(Plain Text and XML)                                                                    | Description(Plain Text and XML)                                                                   | XML Tag (XML Only)                 |
|--------------------------------------------------------------------------------------------------|---------------------------------------------------------------------------------------------------|------------------------------------|
| Contact name                                                                                     | Name of person to contact for issues associated with the node that experienced the event.         | /aml/body/sysContact               |
| Contact e-mail                                                                                   | E-mail address of person identified as the contact for this unit.                                 | /aml/body/sysContactEmail          |
| Contact phone number                                                                             | Phone number of the person identified as the contact for this unit.                               | /aml/body/sysContactPhoneNo        |
| Street address                                                                                   | Optional field that contains the street address for RMA part shipments associated with this unit. | /aml/body/sysStreetAddress         |
| Model name                                                                                       | Model name of the device (the specific model as part of a product family name).                   | /aml/body/chassis/name             |
| Serial number                                                                                    | Chassis serial number of the unit.                                                                | /aml/body/chassis/serialNo         |
| Chassis part number                                                                              | Top assembly number of the chassis.                                                               | /aml/body/chassis/partNo           |
| Fields specific to a particular alert group message are inserted here.                           |                                                                                                   |                                    |
| The following fields may be repeated if multiple CLI commands are executed for this alert group. |                                                                                                   |                                    |
| Command output name                                                                              | Exact name of the issued CLI command.                                                             | /aml/body/alertGroup/cmdOutputName |

| Data Item(Plain Text and XML) | Description(Plain Text and XML)           | XML Tag (XML Only) |
|-------------------------------|-------------------------------------------|--------------------|
| Attachment type               | Specific command output.                  | /aml/msg/cmdType   |
| MIME type                     | Either plain text or encoding type.       | /aml/msg/cmdMime   |
| Command output text           | Output of command automatically executed. | /aml/msg/cmdData   |

The following table describes the reactive event message format for full text or XML.

**Table 48: Inserted Fields for a Reactive or Proactive Event Message**

| Data Item(Plain Text and XML)      | Description(Plain Text and XML)                                | XML Tag (XML Only)          |
|------------------------------------|----------------------------------------------------------------|-----------------------------|
| Chassis hardware version           | Hardware version of chassis.                                   | /aml/body/chassis/hwVersion |
| Supervisor module software version | Top-level software version.                                    | /aml/body/chassis/swVersion |
| Affected FRU name                  | Name of the affected FRU that is generating the event message. | /aml/body/fru/name          |
| Affected FRU serial number         | Serial number of the affected FRU.                             | /aml/body/fru/serialNo      |
| Affected FRU part number           | Part number of the affected FRU.                               | /aml/body/fru/partNo        |
| FRU slot                           | Slot number of the FRU that is generating the event message.   | /aml/body/fru/slot          |
| FRU hardware version               | Hardware version of the affected FRU.                          | /aml/body/fru/hwVersion     |
| FRU software version               | Software version(s) that is running on the affected FRU.       | /aml/body/fru/swVersion     |

The following table describes the inventory event message format for full text or XML.

**Table 49: Inserted Fields for an Inventory Event Message**

| Data Item(Plain Text and XML) | Description(Plain Text and XML)  | XML Tag(XML Only)           |
|-------------------------------|----------------------------------|-----------------------------|
| Chassis hardware version      | Hardware version of the chassis. | /aml/body/chassis/hwVersion |

| Data Item(Plain Text and XML)      | Description(Plain Text and XML)                                | XML Tag(XML Only)           |
|------------------------------------|----------------------------------------------------------------|-----------------------------|
| Supervisor module software version | Top-level software version.                                    | /aml/body/chassis/swVersion |
| FRU name                           | Name of the affected FRU that is generating the event message. | /aml/body/fru/name          |
| FRU s/n                            | Serial number of the FRU.                                      | /aml/body/fru/serialNo      |
| FRU part number                    | Part number of the FRU.                                        | /aml/body/fru/partNo        |
| FRU slot                           | Slot number of the FRU.                                        | /aml/body/fru/slot          |
| FRU hardware version               | Hardware version of the FRU.                                   | /aml/body/fru/hwVersion     |
| FRU software version               | Software version(s) that is running on the FRU.                | /aml/body/fru/swVersion     |

The following table describes the user-generated test message format for full text or XML.

**Table 50: Inserted Fields for a User-Generated Test Message**

| Data Item(Plain Text and XML) | Description(Plain Text and XML)                    | XML Tag(XML Only)              |
|-------------------------------|----------------------------------------------------|--------------------------------|
| Process ID                    | Unique process ID.                                 | /aml/body/process/id           |
| Process state                 | State of process (for example, running or halted). | /aml/body/process/processState |
| Process exception             | Exception or reason code.                          | /aml/body/process/exception    |

### Sample syslog Alert Notification in Full-Text Format

This sample shows the full-text format for a syslog port alert-group notification:

```
source:MDS9000
Switch Priority:7
Device Id:WS-C6509@C@FG@07120011
Customer Id:Example.com
Contract Id:123
Site Id:San Jose
Server Id:WS-C6509@C@FG@07120011
Time of Event:2004-10-08T11:10:44
Message Name:SYSLOG_ALERT
Message Type:Syslog
Severity Level:2
System Name:10.76.100.177
Contact Name:User Name
Contact Email:person@example.com
Contact Phone:+1-408-555-1212
Street Address:#1234 Any Street, Any City, Any State, 12345
Event Description:2006 Oct 8 11:10:44 10.76.100.177 %PORT-5-IF_TRUNK_UP: %$VLAN 1%$ Interface
e2/5, vlan 1 is up

syslog_facility:PORT
start Chassis information:
Affected Chassis:WS-C6509
Affected Chassis Serial Number:FG@07120011
Affected Chassis Hardware Version:0.104
Affected Chassis Software Version:3.1(1)
Affected Chassis Part No:73-8607-01
end chassis information:
```

### Sample syslog Alert Notification in XML Format

This sample shows the XML format for a syslog port alert-group notification:

```
From: example
Sent: Wednesday, April 25, 2007 7:20 AM
To: User (user)
Subject: System Notification From Router - syslog - 2007-04-25 14:19:55
GMT+00:00
```

```

<?xml version="1.0" encoding="UTF-8"?>
<soap-env:Envelope xmlns:soap-env="http://www.w3.org/2003/05/soap-envelope">
<soap-env:Header>
<aml-session:Session xmlns:aml-session="http://www.example.com/2004/01/aml-session"
soap-env:mustUnderstand="true"
soap-env:role="http://www.w3.org/2003/05/soap-envelope/role/next">
<aml-session:To>http://tools.example.com/services/DDCEService</aml-session:To>
<aml-session:Path>
<aml-session:Via>http://www.example.com/appliance/uri</aml-session:Via>
</aml-session:Path>
<aml-session:From>http://www.example.com/appliance/uri</aml-session:From>
<aml-session:MessageId>M2:69000101:C9D9E20B</aml-session:MessageId>
</aml-session:Session>
</soap-env:Header>
<soap-env:Body>
<aml-block:Block xmlns:aml-block="http://www.example.com/2004/01/aml-block">
<aml-block:Header>
<aml-block:Type>http://www.example.com/2005/05/callhome/syslog</aml-block:Type>
<aml-block:CreationDate>2007-04-25 14:19:55 GMT+00:00</aml-block:CreationDate>
<aml-block:Builder>
<aml-block:Name>Cat6500</aml-block:Name>
<aml-block:Version>2.0</aml-block:Version>
</aml-block:Builder>
<aml-block:BlockGroup>
<aml-block:GroupId>G3:69000101:C9F9E20C</aml-block:GroupId>
<aml-block:Number>0</aml-block:Number>
<aml-block:IsLast>true</aml-block:IsLast>
<aml-block:IsPrimary>true</aml-block:IsPrimary>
<aml-block:WaitForPrimary>false</aml-block:WaitForPrimary>
</aml-block:BlockGroup>
<aml-block:Severity>2</aml-block:Severity>
</aml-block:Header>
<aml-block:Content>
<ch:Call Home xmlns:ch="http://www.example.com/2005/05/callhome" version="1.0">
<ch:EventTime>2007-04-25 14:19:55 GMT+00:00</ch:EventTime>
<ch:MessageDescription>03:29:29: %CLEAR-5-COUNTERS: Clear counter on all interfaces by
console</ch:MessageDescription>
<ch:Event>
<ch:Type>syslog</ch:Type>
<ch:SubType></ch:SubType>
<ch:Brand>Cisco Systems</ch:Brand>
<ch:Series>Catalyst 6500 Series Switches</ch:Series>
</ch:Event>
<ch:CustomerData>
<ch:UserData>
<ch:Email>person@example.com</ch:Email>
</ch:UserData>
<ch:ContractData>
<ch:CustomerId>12345</ch:CustomerId>
<ch:SiteId>building 1</ch:SiteId>
<ch:ContractId>abcdefg12345</ch:ContractId>
<ch:DeviceId>WS-C6509@C@69000101</ch:DeviceId>
</ch:ContractData>
<ch:SystemInfo>
<ch:Name>Router</ch:Name>
<ch:Contact></ch:Contact>
<ch:ContactEmail>user@example.com</ch:ContactEmail>
<ch:ContactPhoneNumber>+1-408-555-1212</ch:ContactPhoneNumber>
<ch:StreetAddress>#1234 Any Street, Any City, Any State, 12345</ch:StreetAddress>
</ch:SystemInfo>
</ch:CustomerData>
<ch:Device>
<rme:Chassis xmlns:rme="http://www.example.com/rme/4.0">
<rme:Model>WS-C6509</rme:Model>
<rme:HardwareVersion>1.0</rme:HardwareVersion>
<rme:SerialNumber>69000101</rme:SerialNumber>
<rme:AdditionalInformation>
<rme:AD name="PartNumber" value="73-3438-03 01" />
<rme:AD name="SoftwareVersion" value="4.0(20080421:012711)" />
</rme:AdditionalInformation>
</rme:Chassis>
</ch:Device>
</ch:Call Home>

```

```

</aml-block:Content>
<aml-block:Attachments>
<aml-block:Attachment type="inline">
<aml-block:Name>show logging</aml-block:Name>
<aml-block:Data encoding="plain">
<![CDATA[
Syslog logging: enabled (0 messages dropped, 0 messages rate-limited, 0 flushes, 0 overruns,
xml disabled, filtering disabled)
 Console logging: level debugging, 53 messages logged, xml disabled,
 filtering disabled
 Monitor logging: level debugging, 0 messages logged, xml disabled,
 filtering disabled
 Buffer logging: level debugging, 53 messages logged, xml disabled,
 filtering disabled
 Exception Logging: size (4096 bytes)
 Count and timestamp logging messages: disabled
 Trap logging: level informational, 72 message lines logged

Log Buffer (8192 bytes):

00:00:54: curr is 0x20000

00:00:54: RP: Currently running ROMMON from F2 region
00:01:05: %SYS-5-CONFIG I: Configured from memory by console
00:01:09: %SYS-5-RESTART: System restarted --
Cisco IOS Software, s72033_rp Software (s72033_rp-ADVENTERPRISEK9_DBG-VM), Experimental
Version 12.2(20070421:012711)

Copyright (c) 1986-2007 by Cisco Systems, Inc.
Compiled Thu 26-Apr-07 15:54 by xxx

Firmware compiled 11-Apr-07 03:34 by integ Build [100]

00:01:01: %PFREDUN-6-ACTIVE: Initializing as ACTIVE processor for this switch
00:01:01: %SYS-3-LOGGER_FLUSHED: System was paused for 00:00:00 to ensure console debugging
output.
00:03:00: SP: SP: Currently running ROMMON from F1 region
00:03:07: %C6K_PLATFORM-SP-4-CONFREG_BREAK_ENABLED: The default factory setting for config
register is 0x2102.It is advisable to retain 1 in 0x2102 as it prevents returning to ROMMON
when break is issued.

00:03:18: %SYS-SP-5-RESTART: System restarted --
Cisco IOS Software, s72033_sp Software (s72033_sp-ADVENTERPRISEK9_DBG-VM), Experimental
Version 12.2(20070421:012711)

```

```

Copyright (c) 1986-2007 by Cisco Systems, Inc.
Compiled Thu 26-Apr-07 18:00 by xxx
00:03:18: %SYS-SP-6-BOOTTIME: Time taken to reboot after reload = 339 seconds
00:03:18: %OIR-SP-6-INSPS: Power supply inserted in slot 1
00:03:18: %C6KPWR-SP-4-PSOK: power supply 1 turned on.
00:03:18: %OIR-SP-6-INSPS: Power supply inserted in slot 2
00:01:09: %SSH-5-ENABLED: SSH 1.99 has been enabled
00:03:18: %C6KPWR-SP-4-PSOK: power supply 2 turned on.
00:03:18: %C6KPWR-SP-4-PSREDUNDANTMISMATCH: power supplies rated outputs do not match.
00:03:18: %C6KPWR-SP-4-PSREDUNDANTBOTHSUPPLY: in power-redundancy mode, system is operating
on both power supplies.
00:01:10: %CRYPTO-6-ISAKMP_ON_OFF: ISAKMP is OFF
00:01:10: %CRYPTO-6-ISAKMP_ON_OFF: ISAKMP is OFF
00:03:20: %C6KENV-SP-4-FANHIOUTPUT: Version 2 high-output fan-tray is in effect
00:03:22: %C6KPWR-SP-4-PSNOREDUNDANCY: Power supplies are not in full redundancy, power
usage exceeds lower capacity supply
00:03:26: %FABRIC-SP-5-FABRIC_MODULE_ACTIVE: The Switch Fabric Module in slot 6 became
active.
00:03:28: %DIAG-SP-6-RUN_MINIMUM: Module 6: Running Minimal Diagnostics...
00:03:50: %DIAG-SP-6-DIAG_OK: Module 6: Passed Online Diagnostics
00:03:50: %OIR-SP-6-INSCARD: Card inserted in slot 6, interfaces are now online
00:03:51: %DIAG-SP-6-RUN_MINIMUM: Module 3: Running Minimal Diagnostics...
00:03:51: %DIAG-SP-6-RUN_MINIMUM: Module 7: Running Minimal Diagnostics...
00:03:51: %DIAG-SP-6-RUN_MINIMUM: Module 9: Running Minimal Diagnostics...
00:01:51: %MFIB_CONST_REP-6-REPLICATION_MODE_CHANGE: Replication Mode Change Detected. Current
system replication mode is Ingress
00:04:01: %DIAG-SP-6-DIAG_OK: Module 3: Passed Online Diagnostics
00:04:01: %OIR-SP-6-DOWNGRADE: Fabric capable module 3 not at an appropriate hardware
revision level, and can only run in flowthrough mode
00:04:02: %OIR-SP-6-INSCARD: Card inserted in slot 3, interfaces are now online
00:04:11: %DIAG-SP-6-DIAG_OK: Module 7: Passed Online Diagnostics
00:04:14: %OIR-SP-6-INSCARD: Card inserted in slot 7, interfaces are now online
00:04:35: %DIAG-SP-6-DIAG_OK: Module 9: Passed Online Diagnostics
00:04:37: %OIR-SP-6-INSCARD: Card inserted in slot 9, interfaces are now online
00:00:09: DaughterBoard (Distributed Forwarding Card 3)

```

Firmware compiled 11-Apr-07 03:34 by integ Build [100]

```

00:00:22: %SYS-DFC4-5-RESTART: System restarted --
Cisco DCOS Software, c6lc2 Software (c6lc2-SPDBG-VM), Experimental Version
4.0(20080421:012711)

```

```

Copyright (c) 1986-2008 by Cisco Systems, Inc.
Compiled Thu 26-Apr-08 17:20 by xxx
00:00:23: DFC4: Currently running ROMMON from F2 region
00:00:25: %SYS-DFC2-5-RESTART: System restarted --
Cisco IOS Software, c6slc Software (c6slc-SPDBG-VM), Experimental Version
12.2(20070421:012711)

```

```

Copyright (c) 1986-2007 by Cisco Systems, Inc.
Compiled Thu 26-Apr-08 16:40 by username1
00:00:26: DFC2: Currently running ROMMON from F2 region
00:04:56: %DIAG-SP-6-RUN_MINIMUM: Module 4: Running Minimal Diagnostics...
00:00:09: DaughterBoard (Distributed Forwarding Card 3)

```

Firmware compiled 11-Apr-08 03:34 by integ Build [100]

slot\_id is 8

```

00:00:31: %FLASHFS_HES-DFC8-3-BADCARD: /bootflash:: The flash card seems to be corrupted
00:00:31: %SYS-DFC8-5-RESTART: System restarted --
Cisco DCOS Software, c6lc2 Software (c6lc2-SPDBG-VM), Experimental Version
4.0(20080421:012711)

```

```

Copyright (c) 1986-2008 by Cisco Systems, Inc.
Compiled Thu 26-Apr-08 17:20 by username1
00:00:31: DFC8: Currently running ROMMON from S (Gold) region
00:04:59: %DIAG-SP-6-RUN_MINIMUM: Module 2: Running Minimal Diagnostics...
00:05:12: %DIAG-SP-6-RUN_MINIMUM: Module 8: Running Minimal Diagnostics...
00:05:13: %DIAG-SP-6-RUN_MINIMUM: Module 1: Running Minimal Diagnostics...
00:00:24: %SYS-DFC1-5-RESTART: System restarted --
Cisco DCOS Software, c6slc Software (c6slc-SPDBG-VM), Experimental Version
4.0(20080421:012711)

```

```

Copyright (c) 1986-2008 by Cisco Systems, Inc.
Compiled Thu 26-Apr-08 16:40 by username1
00:00:25: DFC1: Currently running ROMMON from F2 region
00:05:30: %DIAG-SP-6-DIAG_OK: Module 4: Passed Online Diagnostics
00:05:31: %SPAN-SP-6-SPAN_EGRESS_REPLICATION_MODE_CHANGE: Span Egress HW Replication Mode
Change Detected. Current replication mode for unused asic session 0 is Centralized
00:05:31: %SPAN-SP-6-SPAN_EGRESS_REPLICATION_MODE_CHANGE: Span Egress HW Replication Mode
Change Detected. Current replication mode for unused asic session 1 is Centralized
00:05:31: %OIR-SP-6-INSCARD: Card inserted in slot 4, interfaces are now online
00:06:02: %DIAG-SP-6-DIAG_OK: Module 1: Passed Online Diagnostics
00:06:03: %OIR-SP-6-INSCARD: Card inserted in slot 1, interfaces are now online
00:06:31: %DIAG-SP-6-DIAG_OK: Module 2: Passed Online Diagnostics
00:06:33: %OIR-SP-6-INSCARD: Card inserted in slot 2, interfaces are now online
00:04:30: %XDR-6-XDRIPCNOTIFY: Message not sent to slot 4/0 (4) because of IPC error timeout.
Disabling linecard. (Expected during linecard OIR)
00:06:59: %DIAG-SP-6-DIAG_OK: Module 8: Passed Online Diagnostics
00:06:59: %OIR-SP-6-DOWNGRADE_EARL: Module 8 DFC installed is not identical to system PFC
and will perform at current system operating mode.
00:07:06: %OIR-SP-6-INSCARD: Card inserted in slot 8, interfaces are now online

```

```

Router#]]></aml-block:Data>
</aml-block:Attachment>
</aml-block:Attachments>
</aml-block:Block>
</soap-env:Body>
</soap-env:Envelope>

```





## Configuring SNMP

This chapter describes the configuration of the Simple Network Management Protocol (SNMP) on Cisco Nexus 5000 Series switches and contains the following sections:

- [Information About SNMP, page 415](#)
- [Configuration Guidelines and Limitations, page 419](#)
- [Configuring SNMP, page 419](#)
- [Verifying SNMP Configuration, page 427](#)
- [Default SNMP Settings, page 427](#)

### 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 the managed device that maintains the data for the device and reports these data, as needed, to managing systems. The Cisco Nexus 5000 Series switch supports the agent and MIB. To enable the SNMP agent, you must define the relationship between the manager and the agent.
- A managed information base (MIB)—The collection of managed objects on the SNMP agent



**Note**

Cisco NX-OS does not support SNMP sets for Ethernet MIBs.

The Cisco Nexus 5000 Series switch supports SNMPv1, SNMPv2c and SNMPv3. Both SNMPv1 and SNMPv2c use a community-based form of security. SNMP is defined in 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>), and 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 NX-OS 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. The switch 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 Nexus 5000 Series switch never receives a response, it can send the inform request again.

You can configure Cisco NX-OS to send notifications to multiple host receivers.

### Related Topics

[Configuring SNMP Notification Receivers, on page 421](#)

## SNMPv3

SNMPv3 provides secure access to devices by a combination of authenticating and encrypting frames over the network. The security features provided in SNMPv3 are the following:

- Message integrity—Ensures that a packet has not been tampered with in-transit.
- Authentication—Determines the message is from a valid source.
- Encryption—Scrambles the packet contents to prevent it from being seen by unauthorized sources.

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.

### Security Models and Levels for SNMPv1, v2, v3

The security level determines if an SNMP message needs to be protected from disclosure and if the message needs to be authenticated. The various security levels that exist within a security model are as follows:

- noAuthNoPriv—Security level that does not provide authentication or encryption.
- authNoPriv—Security level that provides authentication but does not provide encryption.
- authPriv—Security level that provides both authentication and encryption.

Three security models are available: SNMPv1, SNMPv2c, and SNMPv3. The security model combined with the security level determine the security mechanism applied when the SNMP message is processed.

### User-Based Security Model

The following table identifies what the combinations of security models and levels mean.

**Table 51: 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 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—Ensures that information is not made available or disclosed to unauthorized individuals, entities, or processes.

SNMPv3 authorizes management operations only by configured users and encrypts SNMP messages.

Cisco NX-OS uses two authentication protocols for SNMPv3:

- HMAC-MD5-96 authentication protocol
- HMAC-SHA-96 authentication protocol

Cisco NX-OS uses Advanced Encryption Standard (AES) as one of the privacy protocols for SNMPv3 message encryption and conforms with RFC 3826.

The `priv` option offers a choice of DES or 128-bit AES encryption for SNMP security encryption. The `priv` option along with the `aes-128` token indicates that this privacy password is for generating a 128-bit AES key. The AES `priv` password can have a minimum of eight characters. If the passphrases are specified in clear text, you can specify a maximum of 64 characters. If you use the localized key, you can specify a maximum of 130 characters.



**Note**

---

For an SNMPv3 operation using the external AAA server, you must use AES for the privacy protocol in user configuration on the external AAA server.

---

## CLI and SNMP User Synchronization

SNMPv3 user management can be centralized at the Access Authentication and Accounting (AAA) server level. This centralized user management allows the SNMP agent in Cisco NX-OS to leverage the user authentication service of the AAA server. Once user authentication is verified, the SNMP PDUs are processed further. Additionally, the AAA server is also used to store user group names. SNMP uses the group names to apply the access/role policy that is locally available in the switch.

Any configuration changes made to the user group, role, or password results in database synchronization for both SNMP and AAA.

Cisco NX-OS synchronizes user configuration in the following ways:

- The auth passphrase specified in the `snmp-server user` command becomes the password for the CLI user.
- The password specified in the `username` command becomes as the auth and `priv` passphrases for the SNMP user.
- Deleting a user using either SNMP or the CLI results in the user being deleted for both SNMP and the CLI.
- User-role mapping changes are synchronized in SNMP and the CLI.



**Note** When you configure passphrase/password in localized key/encrypted format, Cisco NX-OS does not synchronize the password.

### Group-Based SNMP Access



**Note** Because group is a standard SNMP term used industry-wide, roles are referred to as groups in this SNMP section.

SNMP access rights are organized by groups. Each group in SNMP is similar to a role through the CLI. Each group is defined with three accesses: read access, write access, and notification access. Each access can be enabled or disabled within each group.

You can begin communicating with the agent once your user name is created, your roles are set up by your administrator, and you are added to the roles.

## Configuration Guidelines and Limitations

SNMP has the following configuration guidelines and limitations:

- Cisco NX-OS supports read-only access to Ethernet MIBs.

## Configuring SNMP

### Configuring SNMP Users

To configure a user for SNMP, perform this task:

#### SUMMARY STEPS

1. switch# **configuration terminal**
2. switch(config)# **snmp-server user** *name* [**auth** {**md5** | **sha**} *passphrase* [**auto**] [**priv** [**aes-128**] *passphrase*] [**engineID** *id*] [**localizedkey**]]
3. (Optional) switch# **show snmp user**
4. (Optional) switch# **copy running-config startup-config**

#### DETAILED STEPS

|        | Command or Action                     | Purpose                    |
|--------|---------------------------------------|----------------------------|
| Step 1 | switch# <b>configuration terminal</b> | Enters configuration mode. |

|               | Command or Action                                                                                                                                                                                                                           | Purpose                                                             |
|---------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|---------------------------------------------------------------------|
| <b>Step 2</b> | switch(config)# <b>snmp-server user</b> <i>name</i> [ <b>auth</b> { <b>md5</b>   <b>sha</b> } <i>passphrase</i> [ <b>auto</b> ] [ <b>priv</b> [ <b>aes-128</b> ] <i>passphrase</i> ] [ <b>engineID</b> <i>id</i> ] [ <b>localizedkey</b> ]] | Configures an SNMP user with authentication and privacy parameters. |
| <b>Step 3</b> | switch# <b>show snmp user</b>                                                                                                                                                                                                               | (Optional)<br>Displays information about one or more SNMP users.    |
| <b>Step 4</b> | switch# <b>copy running-config startup-config</b>                                                                                                                                                                                           | (Optional)<br>Saves this configuration change.                      |

## Enforcing SNMP Message Encryption

You can configure SNMP to require authentication or encryption for incoming requests. By default the SNMP agent accepts SNMPv3 messages without authentication and encryption. When you enforce privacy, Cisco NX-OS responds with an authorization Error for any SNMPv3 PDU request using securityLevel parameter of either noAuthNoPriv or authNoPriv.

You can enforce SNMP message encryption for a specific user.

| Command                                                                | Purpose                                         |
|------------------------------------------------------------------------|-------------------------------------------------|
| switch(config)# <b>snmp-server user</b> <i>name</i> <b>enforcePriv</b> | Enforces SNMP message encryption for this user. |

You can enforce SNMP message encryption for all users.

| Command                                              | Purpose                                         |
|------------------------------------------------------|-------------------------------------------------|
| switch(config)# <b>snmp-server globalEnforcePriv</b> | Enforces SNMP message encryption for all users. |

## Assigning SNMPv3 Users to Multiple Roles

After you configure an SNMP user, you can assign multiple roles for the user.



### Note

Only users belonging to a network-admin role can assign roles to other users.

| Command                                                          | Purpose                                                  |
|------------------------------------------------------------------|----------------------------------------------------------|
| switch(config)# <b>snmp-server user</b> <i>name</i> <i>group</i> | Associates this SNMP user with the configured user role. |

## Creating SNMP Communities

You can create SNMP communities for SNMPv1 or SNMPv2c.

To create an SNMP community string in a global configuration mode, perform this task:

| Command                                                                                            | Purpose                           |
|----------------------------------------------------------------------------------------------------|-----------------------------------|
| switch(config)# <b>snmp-server community</b> <i>name</i><br><i>group</i> { <b>ro</b>   <b>rw</b> } | Creates an SNMP community string. |

## Configuring SNMP Notification Receivers

You can configure Cisco NX-OS to generate SNMP notifications to multiple host receivers.

You can configure a host receiver for SNMPv1 traps in a global configuration mode.

| Command                                                                                                                                       | Purpose                                                                                                                                                       |
|-----------------------------------------------------------------------------------------------------------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------|
| switch(config)# <b>snmp-server host</b> <i>ip-address</i> <b>traps</b><br><b>version 1</b> <i>community</i> [ <b>udp_port</b> <i>number</i> ] | Configures a host receiver for SNMPv1 traps. The community can be any alphanumeric string up to 255 characters. The UDP port number range is from 0 to 65535. |

You can configure a host receiver for SNMPv2c traps or informs in a global configuration mode.

| Command                                                                                                                                                             | Purpose                                                                                                                                                                   |
|---------------------------------------------------------------------------------------------------------------------------------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| switch(config)# <b>snmp-server host</b> <i>ip-address</i> { <b>traps</b><br>  <b>informs</b> } <b>version 2c</b> <i>community</i> [ <b>udp_port</b> <i>number</i> ] | Configures a host receiver for SNMPv2c traps or informs. The community can be any alphanumeric string up to 255 characters. The UDP port number range is from 0 to 65535. |

You can configure a host receiver for SNMPv3 traps or informs in a global configuration mode.

| Command                                                                                                                                                                                                            | Purpose                                                                                                                                                                  |
|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| switch(config)# <b>snmp-server host</b> <i>ip-address</i> { <b>traps</b><br>  <b>informs</b> } <b>version 3</b> { <b>auth</b>   <b>noauth</b>   <b>priv</b> } <i>username</i><br>[ <b>udp_port</b> <i>number</i> ] | Configures a host receiver for SNMPv2c traps or informs. The username can be any alphanumeric string up to 255 characters. The UDP port number range is from 0 to 65535. |



### Note

The SNMP manager must know the user credentials (authKey/PrivKey) based on the SNMP engineID of the Cisco Nexus 5000 Series switch to authenticate and decrypt the SNMPv3 messages.

The following example shows how to configure a host receiver for an SNMPv1 trap:

```
switch(config)# snmp-server host 192.0.2.1 traps version 1 public
```

The following example shows how to configure a host receiver for an SNMPv2 inform:

```
switch(config)# snmp-server host 192.0.2.1 informs version 2c public
```

The following example shows how to configure a host receiver for an SNMPv3 inform:

```
switch(config)# snmp-server host 192.0.2.1 informs version 3 auth NMS
```

## Configuring the Notification Target User

You must configure a notification target user on the device to send SNMPv3 inform notifications to a notification host receiver.

The Cisco Nexus 5000 Series switch uses the credentials of the notification target user to encrypt the SNMPv3 inform notification messages to the configured notification host receiver.



**Note**

For authenticating and decrypting the received INFORM PDU, The notification host receiver should have the same user credentials as configured in the Cisco Nexus 5000 Series switch to authenticate and decrypt the informs.

| Command                                                                                                                                                                                                            | Purpose                                                                                                                                                                    |
|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| switch(config)# <b>snmp-server user</b> <i>name</i> [ <b>auth</b> { <b>md5</b>   <b>sha</b> } <i>passphrase</i> [ <b>auto</b> ] [ <b>priv</b> [ <b>aes-128</b> ] <i>passphrase</i> ] [ <b>engineID</b> <i>id</i> ] | Configures the notification target user with the specified engine ID for notification host receiver. The engineID format is a 12-digit colon-separated hexadecimal number. |

The following example shows how to configure a notification target user:

```
switch(config)# snmp-server user NMS auth sha abcd1234 priv abcdefgh engineID 00:00:00:63:00:01:00:a1:ac:15:10:03
```

## Enabling SNMP Notifications

You can enable or disable notifications. If you do not specify a notification name, Cisco NX-OS enables all notifications.



**Note**

The **snmp-server enable traps** CLI command enables both traps and informs, depending on the configured notification host receivers.

The following table lists the CLI commands that enable the notifications for Cisco NX-OS MIBs.

**Table 52: Enabling SNMP Notifications**

| <b>MIB</b>                                                              | <b>Related Commands</b>                                                                                                                                                                                                                                                                                                         |
|-------------------------------------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| All notifications                                                       | <b>snmp-server enable traps</b>                                                                                                                                                                                                                                                                                                 |
| CISCO-AAA-SERVER-MIB                                                    | <b>snmp-server enable traps aaa</b>                                                                                                                                                                                                                                                                                             |
| ENTITY-MIB,<br>CISCO-ENTITY-FRU-CONTROL-MIB,<br>CISCO-ENTITY-SENSOR-MIB | <b>snmp-server enable traps entity</b><br><b>snmp-server enable traps entity fru</b>                                                                                                                                                                                                                                            |
| CISCO-LICENSE-MGR-MIB                                                   | <b>snmp-server enable traps license</b>                                                                                                                                                                                                                                                                                         |
| IF-MIB                                                                  | <b>snmp-server enable traps link</b>                                                                                                                                                                                                                                                                                            |
| CISCO-PSM-MIB                                                           | <b>snmp-server enable traps port-security</b>                                                                                                                                                                                                                                                                                   |
| SNMPv2-MIB                                                              | <b>snmp-server enable traps snmp</b><br><b>snmp-server enable traps snmp authentication</b>                                                                                                                                                                                                                                     |
| CISCO-FCC-MIB                                                           | <b>snmp-server enable traps fcc</b>                                                                                                                                                                                                                                                                                             |
| CISCO-DM-MIB                                                            | <b>snmp-server enable traps fcdomain</b>                                                                                                                                                                                                                                                                                        |
| CISCO-NS-MIB                                                            | <b>snmp-server enable traps fcns</b>                                                                                                                                                                                                                                                                                            |
| CISCO-FCS-MIB                                                           | <b>snmp-server enable traps fcs discovery-complete</b><br><b>snmp-server enable traps fcs request-reject</b>                                                                                                                                                                                                                    |
| CISCO-FDMI-MIB                                                          | <b>snmp-server enable traps fdmi</b>                                                                                                                                                                                                                                                                                            |
| CISCO-FSPF-MIB                                                          | <b>snmp-server enable traps fspf</b>                                                                                                                                                                                                                                                                                            |
| CISCO-PSM-MIB                                                           | <b>snmp-server enable traps port-security</b>                                                                                                                                                                                                                                                                                   |
| CISCO-RSCN-MIB                                                          | <b>snmp-server enable traps rscn</b><br><b>snmp-server enable traps rscn els</b><br><b>snmp-server enable traps rscn ils</b>                                                                                                                                                                                                    |
| CISCO-ZS-MIB                                                            | <b>snmp-server enable traps zone</b><br><b>snmp-server enable traps zone default-zone-behavior-change</b><br><b>snmp-server enable traps zone merge-failure</b><br><b>snmp-server enable traps zone merge-success</b><br><b>snmp-server enable traps zone request-reject</b><br><b>snmp-server enable traps zone unsupp-mem</b> |

**Note**

The license notifications are enabled by default. All other notifications are disabled by default.

To enable the specified notification in the global configuration mode, perform one of the following tasks:

| Command                                                                      | Purpose                                       |
|------------------------------------------------------------------------------|-----------------------------------------------|
| switch(config)# <b>snmp-server enable traps</b>                              | Enables all SNMP notifications.               |
| switch(config)# <b>snmp-server enable traps aaa</b><br>[server-state-change] | Enables the AAA SNMP notifications.           |
| switch(config)# <b>snmp-server enable traps entity</b><br>[fru]              | Enables the ENTITY-MIB SNMP notifications.    |
| switch(config)# <b>snmp-server enable traps license</b>                      | Enables the license SNMP notification.        |
| switch(config)# <b>snmp-server enable traps port-security</b>                | Enables the port security SNMP notifications. |
| switch(config)# <b>snmp-server enable traps snmp</b><br>[authentication]     | Enables the SNMP agent notifications.         |

## Configuring Link Notifications

You can configure which linkUp/linkDown notifications to enable on a device. You can enable the following types of linkUp/linkDown notifications:

- Cisco—Cisco NX-OS sends only the Cisco-defined notifications (cieLinkUp, cieLinkDown in CISCO-IF-EXTENSION-MIB.my), if ifLinkUpDownTrapEnable (defined in IF-MIB) is enabled for that interface.
- IETF—Cisco NX-OS sends only the IETF-defined notifications (linkUp, linkDown in IF-MIB) with only the defined varbinds, if ifLinkUpDownTrapEnable (defined in IF-MIB) is enabled for that interface.
- IETF extended—Cisco NX-OS sends only the IETF-defined notifications (linkUp, linkDown defined in IF-MIB), if ifLinkUpDownTrapEnable (defined in IF-MIB) is enabled for that interface. Cisco NX-OS adds additional varbinds specific to Cisco Systems in addition to the varbinds defined in the IF-MIB. This is the default setting.
- IETF Cisco—Cisco NX-OS sends the notifications (linkUp, linkDown) defined in IF-MIB and notifications (cieLinkUp, cieLinkDown) defined in CISCO-IF-EXTENSION-MIB.my, if ifLinkUpDownTrapEnable (defined in IF-MIB) is enabled for that interface. Cisco NX-OS sends only the varbinds defined in the linkUp and linkDown notifications.
- IETF extended Cisco—Cisco NX-OS sends the notifications (linkUp, linkDown) defined in IF-MIB and notifications (cieLinkUp, cieLinkDown) defined in CISCO-IF-EXTENSION-MIB.my, if ifLinkUpDownTrapEnable (defined in IF-MIB) is enabled for that interface. Cisco NX-OS adds additional varbinds specific to Cisco Systems in addition to the varbinds defined in the IF-MIB for the linkUp and linkDown notifications.

**SUMMARY STEPS**

1. switch# **configure terminal**
2. switch(config)# **snmp-server enable traps link [cisco] [ietf | ietf-extended]**

**DETAILED STEPS**

|               | Command or Action                                                                   | Purpose                              |
|---------------|-------------------------------------------------------------------------------------|--------------------------------------|
| <b>Step 1</b> | switch# <b>configure terminal</b>                                                   | Enters configuration mode.           |
| <b>Step 2</b> | switch(config)# <b>snmp-server enable traps link [cisco] [ietf   ietf-extended]</b> | Enables the link SNMP notifications. |

**Disabling Link Notifications on an Interface**

You can disable linkUp and linkDown notifications on an individual interface. You can use this limit notifications on flapping interface (an interface that transitions between up and down repeatedly).

**SUMMARY STEPS**

1. switch# **configure terminal**
2. switch(config)# **interface type slot/port**
3. switch(config-if)# **no snmp trap link-status**

**DETAILED STEPS**

|               | Command or Action                                  | Purpose                                                               |
|---------------|----------------------------------------------------|-----------------------------------------------------------------------|
| <b>Step 1</b> | switch# <b>configure terminal</b>                  | Enters configuration mode.                                            |
| <b>Step 2</b> | switch(config)# <b>interface type slot/port</b>    | Specifies the interface to be changed.                                |
| <b>Step 3</b> | switch(config-if)# <b>no snmp trap link-status</b> | Disables SNMP link-state traps for the interface. Enabled by default. |

**Enabling One-Time Authentication for SNMP over TCP**

You can enable a one-time authentication for SNMP over a TCP session.

| Command                                               | Purpose                                                                             |
|-------------------------------------------------------|-------------------------------------------------------------------------------------|
| switch(config)# <b>snmp-server tcp-session [auth]</b> | Enables a one-time authentication for SNMP over a TCP session. Default is disabled. |

## Assigning SNMP Switch Contact and Location Information

You can assign the switch contact information, which is limited to 32 characters (without spaces), and the switch location.

### SUMMARY STEPS

1. switch# **configuration terminal**
2. switch(config)# **snmp-server contact** *name*
3. switch(config)# **snmp-server location** *name*
4. (Optional) switch# **show snmp**
5. (Optional) switch# **copy running-config startup-config**

### DETAILED STEPS

|               | Command or Action                                       | Purpose                                                                    |
|---------------|---------------------------------------------------------|----------------------------------------------------------------------------|
| <b>Step 1</b> | switch# <b>configuration terminal</b>                   | Enters configuration mode.                                                 |
| <b>Step 2</b> | switch(config)# <b>snmp-server contact</b> <i>name</i>  | Configures sysContact, the SNMP contact name.                              |
| <b>Step 3</b> | switch(config)# <b>snmp-server location</b> <i>name</i> | Configures sysLocation, the SNMP location.                                 |
| <b>Step 4</b> | switch# <b>show snmp</b>                                | (Optional)<br>Displays information about one or more destination profiles. |
| <b>Step 5</b> | switch# <b>copy running-config startup-config</b>       | (Optional)<br>Saves this configuration change.                             |

## Configuring the Context to Network Entity Mapping

You can configure an SNMP context to map to a logical network entity, such as a protocol instance or VRF.

### SUMMARY STEPS

1. switch# **configuration terminal**
2. switch(config)# **snmp-server context** *context-name* [**instance** *instance-name*] [**vrf** *vrf-name*] [**topology** *topology-name*]
3. switch(config)# **snmp-server mib community-map** *community-name* **context** *context-name*
4. (Optional) switch(config)# **no snmp-server context** *context-name* [**instance** *instance-name*] [**vrf** *vrf-name*] [**topology** *topology-name*]

## DETAILED STEPS

|               | Command or Action                                                                                                                                                                  | Purpose                                                                                                                                                                                                                                                                                                                                                                                                              |
|---------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Step 1</b> | switch# <b>configuration terminal</b>                                                                                                                                              | Enters configuration mode.                                                                                                                                                                                                                                                                                                                                                                                           |
| <b>Step 2</b> | switch(config)# <b>snmp-server context</b> <i>context-name</i> [ <b>instance</b> <i>instance-name</i> ] [ <b>vrf</b> <i>vrf-name</i> ] [ <b>topology</b> <i>topology-name</i> ]    | Maps an SNMP context to a protocol instance, VRF, or topology. The names can be any alphanumeric string up to 32 characters.                                                                                                                                                                                                                                                                                         |
| <b>Step 3</b> | switch(config)# <b>snmp-server mib community-map</b> <i>community-name</i> <b>context</b> <i>context-name</i>                                                                      | Maps an SNMPv2c community to an SNMP context. The names can be any alphanumeric string up to 32 characters.                                                                                                                                                                                                                                                                                                          |
| <b>Step 4</b> | switch(config)# <b>no snmp-server context</b> <i>context-name</i> [ <b>instance</b> <i>instance-name</i> ] [ <b>vrf</b> <i>vrf-name</i> ] [ <b>topology</b> <i>topology-name</i> ] | <p>(Optional)<br/>Deletes the mapping between an SNMP context and a protocol instance, VRF, or topology. The names can be any alphanumeric string up to 32 characters.</p> <p><b>Note</b> Do not enter an instance, VRF, or topology to delete a context mapping. If you use the <b>instance</b>, <b>vrf</b>, or <b>topology</b> keywords, you configure a mapping between the context and a zero-length string.</p> |

## Verifying SNMP Configuration

To display SNMP configuration information, perform one of the following tasks:

| Command                            | Purpose                                              |
|------------------------------------|------------------------------------------------------|
| switch# <b>show snmp</b>           | Displays the SNMP status.                            |
| switch# <b>show snmp community</b> | Displays the SNMP community strings.                 |
| switch# <b>show snmp engineID</b>  | Displays the SNMP engineID.                          |
| switch# <b>show snmp group</b>     | Displays SNMP roles.                                 |
| switch# <b>show snmp sessions</b>  | Displays SNMP sessions.                              |
| switch# <b>show snmp trap</b>      | Displays the SNMP notifications enabled or disabled. |
| switch# <b>show snmp user</b>      | Displays SNMPv3 users.                               |

## Default SNMP Settings

The following table lists the default settings for SNMP parameters.

**Table 53: Default SNMP Parameters**

| <b>Parameters</b>             | <b>Default</b> |
|-------------------------------|----------------|
| license notifications         | enabled        |
| linkUp/Down notification type | ietf-extended  |



## Configuring RMON

This chapter contains the following sections:

- [Configuring RMON, page 429](#)

### Configuring RMON

#### Information About RMON

RMON is an Internet Engineering Task Force (IETF) standard monitoring specification that allows various network agents and console systems to exchange network monitoring data. The Cisco NX-OS supports RMON alarms, events and logs to monitor Cisco Nexus 5000 Series switches

An RMON alarm monitors a specific management information base (MIB) object for a specified interval, triggers an alarm at a specified threshold value (threshold), and resets the alarm at another threshold value. You can use alarms with RMON events to generate a log entry or an SNMP notification when the RMON alarm triggers.

RMON is disabled by default and no events or alarms are configured in Cisco Nexus 5000 Series. You can configure your RMON alarms and events by using the CLI or an SNMP-compatible network management station

#### RMON Alarms

You can set an alarm on any MIB object that resolves into an SNMP INTEGER type. The specified object must be an existing SNMP MIB object in standard dot notation (for example, 1.3.6.1.2.1.2.2.1.17 represents ifOutOctets.17).

When you create an alarm, you specify the following parameters:

- MIB object to monitor
- Sampling interval—The interval that the Cisco Nexus 5000 Series switch uses to collect a sample value of the MIB object.
- The sample type—Absolute samples take the current snapshot of the MIB object value. Delta samples take two consecutive samples and calculate the difference between them.

- Rising threshold—The value at which the Cisco Nexus 5000 Series switch triggers a rising alarm or resets a falling alarm.
- Falling threshold—The value at which the Cisco Nexus 5000 Series switch triggers a falling alarm or resets a rising alarm.
- Events—The action that the Cisco Nexus 5000 Series switch takes when an alarm (rising or falling) triggers.




---

**Note** Use the `hcalarms` option to set an alarm on a 64-bit integer MIB object.

---

For example, you can set a delta type rising alarm on an error counter MIB object. If the error counter delta exceeds this value, you can trigger an event that sends an SNMP notification and logs the rising alarm event. This rising alarm will not occur again until the delta sample for the error counter drops below the falling threshold.




---

**Note** The falling threshold must be less than the rising threshold.

---

## RMON Events

You can associate a particular event to each RMON alarm. RMON supports the following event types:

- SNMP notification—Sends an SNMP `risingAlarm` or `fallingAlarm` notification when the associated alarm triggers.
- Log—Adds an entry in the RMON log table when the associated alarm triggers.
- Both—Sends an SNMP notification and adds an entry in the RMON log table when the associated alarm triggers.

You can specify a different event for a falling alarm and a rising alarm.

## Configuration Guidelines and Limitations

RMON has the following configuration guidelines and limitations:

- You must configure an SNMP user as a notification receiver to use the SNMP notification event type.
- You can only configure an RMON alarm on a MIB object that resolves to an integer.

## Configuring RMON

### Configuring RMON Alarms

You can configure RMON alarms on any integer-based SNMP MIB object.

You can optionally specify the following parameters:

- The event-number to trigger if the rising or falling threshold exceeds the specified limit.

- The owner of the alarm.

Ensure you have configured an SNMP user and enabled SNMP notifications.

## SUMMARY STEPS

1. switch# **configure terminal**
2. switch(config)# **rmon alarm** *index mib-object sample-interval* {absolute | delta} **rising-threshold** *value* [*event-index*] **falling-threshold** *value* [*event-index*] [**owner name**]
3. switch(config)# **rmon hcalarm** *index mib-object sample-interval* {absolute | delta} **rising-threshold-high** *value* **rising-threshold-low** *value* [*event-index*] **falling-threshold-high** *value* **falling-threshold-low** *value* [*event-index*] [**owner name**] [**storagetype** *type*]
4. (Optional) switch# **show rmon** {alarms | hcalarms}
5. (Optional) switch# **copy running-config startup-config**

## DETAILED STEPS

|               | Command or Action                                                                                                                                                                                                                                                                                                                                                             | Purpose                                                                                                                                                                              |
|---------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Step 1</b> | switch# <b>configure terminal</b>                                                                                                                                                                                                                                                                                                                                             | Enters configuration mode.                                                                                                                                                           |
| <b>Step 2</b> | switch(config)# <b>rmon alarm</b> <i>index mib-object sample-interval</i> {absolute   delta} <b>rising-threshold</b> <i>value</i> [ <i>event-index</i> ] <b>falling-threshold</b> <i>value</i> [ <i>event-index</i> ] [ <b>owner name</b> ]                                                                                                                                   | Creates an RMON alarm. The value range is from -2147483647 to 2147483647. The owner name can be any alphanumeric string.                                                             |
| <b>Step 3</b> | switch(config)# <b>rmon hcalarm</b> <i>index mib-object sample-interval</i> {absolute   delta} <b>rising-threshold-high</b> <i>value</i> <b>rising-threshold-low</b> <i>value</i> [ <i>event-index</i> ] <b>falling-threshold-high</b> <i>value</i> <b>falling-threshold-low</b> <i>value</i> [ <i>event-index</i> ] [ <b>owner name</b> ] [ <b>storagetype</b> <i>type</i> ] | Creates an RMON high-capacity alarm. The value range is from -2147483647 to 2147483647. The owner name can be any alphanumeric string.<br><br>The storage type range is from 1 to 5. |
| <b>Step 4</b> | switch# <b>show rmon</b> {alarms   hcalarms}                                                                                                                                                                                                                                                                                                                                  | (Optional)<br>Displays information about RMON alarms or high-capacity alarms.                                                                                                        |
| <b>Step 5</b> | switch# <b>copy running-config startup-config</b>                                                                                                                                                                                                                                                                                                                             | (Optional)<br>Saves this configuration change.                                                                                                                                       |

The following example shows how to configure RMON alarms:

```
switch# configure terminal
switch(config)# rmon alarm 1 1.3.6.1.2.1.2.2.1.17.83886080 5 delta rising-threshold 5 1
falling-threshold 0 owner test
switch(config)# exit
switch# show rmon alarms
Alarm 1 is active, owned by test
Monitors 1.3.6.1.2.1.2.2.1.17.83886080 every 5 second(s)
Taking delta samples, last value was 0
Rising threshold is 5, assigned to event 1
Falling threshold is 0, assigned to event 0
On startup enable rising or falling alarm
```

## Configuring RMON Events

You can configure RMON events to associate with RMON alarms. You can reuse the same event with multiple RMON alarms.

Ensure you have configured an SNMP user and enabled SNMP notifications.

### SUMMARY STEPS

1. switch# **configure terminal**
2. switch(config)# **rmon event** *index* [**description** *string*] [**log**] [**trap**] [**owner** *name*]
3. (Optional) switch(config)# **show rmon** {**alarms** | **hcalarms**}
4. (Optional) switch# **copy running-config startup-config**

### DETAILED STEPS

|               | Command or Action                                                                                                                               | Purpose                                                                                         |
|---------------|-------------------------------------------------------------------------------------------------------------------------------------------------|-------------------------------------------------------------------------------------------------|
| <b>Step 1</b> | switch# <b>configure terminal</b>                                                                                                               | Enters configuration mode.                                                                      |
| <b>Step 2</b> | switch(config)# <b>rmon event</b> <i>index</i> [ <b>description</b> <i>string</i> ] [ <b>log</b> ] [ <b>trap</b> ] [ <b>owner</b> <i>name</i> ] | Configures an RMON event. The description string and owner name can be any alphanumeric string. |
| <b>Step 3</b> | switch(config)# <b>show rmon</b> { <b>alarms</b>   <b>hcalarms</b> }                                                                            | (Optional)<br>Displays information about RMON alarms or high-capacity alarms.                   |
| <b>Step 4</b> | switch# <b>copy running-config startup-config</b>                                                                                               | (Optional)<br>Saves this configuration change.                                                  |

## Verifying RMON Configuration

To display RMON configuration information, perform one of the following tasks:

| Command                           | Purpose                                   |
|-----------------------------------|-------------------------------------------|
| switch# <b>show rmon alarms</b>   | Displays information about RMON alarms.   |
| switch# <b>show rmon events</b>   | Displays information about RMON events.   |
| switch# <b>show rmon hcalarms</b> | Displays information about RMON hcalarms. |
| switch# <b>show rmon logs</b>     | Displays information about RMON logs.     |

## Default RMON Settings

The following table lists the default settings for RMON parameters.

**Table 54: Default RMON Parameters**

| Parameters | Default          |
|------------|------------------|
| Alarms     | None configured. |
| Events     | None configured. |





# PART **V**

## **Fibre Channel over Ethernet**

- [Configuring FCoE, page 437](#)
- [Configuring FCoE VLANs and Virtual Interfaces, page 457](#)





# Configuring FCoE

This chapter describes how to configure Fibre Channel over Ethernet (FCoE) on Cisco Nexus 5000 Series switches. It contains the following sections:

- [Information About FCoE, page 437](#)
- [FCoE Topologies, page 442](#)
- [FCoE Best Practices, page 444](#)
- [Licensing Requirements for FCoE, page 447](#)
- [Configuring FCoE, page 447](#)
- [Configuring LLDP, page 452](#)
- [Verifying FCoE Configuration, page 453](#)

## Information About FCoE

Fibre Channel over Ethernet (FCoE) provides a method of transporting Fibre Channel traffic over a physical Ethernet connection. FCoE requires the underlying Ethernet to be full duplex and to provide lossless behavior for Fibre Channel traffic.



**Note**

Lossless behavior on Ethernet is provided by using a priority flow control (PFC) mechanism that prevents packet loss during congestion conditions.

Cisco Nexus 5000 Series switches support T11-compliant FCoE on all 10-Gigabit Ethernet interfaces.

## Information About FCoE and FIP

### FCoE Initiation Protocol

The FCoE Initialization Protocol (FIP) allows the switch to discover and initialize FCoE-capable entities that are connected to an Ethernet LAN. Two versions of FIP are supported by the Cisco Nexus 5000 Series switch:

- FIP—The Converged Enhanced Ethernet Data Center Bridging Exchange (CEE-DCBX) protocol supports T11-compliant Gen-2 CNAs.
- Pre-FIP—The Cisco, Intel, Nuova Data Center Bridging Exchange (CIN-DCBX) protocol supports Gen-1 converged network adapters (CNAs).

The Cisco Nexus 5000 Series switch detects the capabilities of the attached CNA and switches to the correct FIP mode.

### FIP Virtual Link Instantiation

Cisco NX-OS Release 4.1(3)N1(1) adds support for the T11-compliant FIP on the Cisco Nexus 5000 Series switch.

FIP is used to perform device discovery, initialization, and link maintenance. FIP performs the following protocols:

- FIP Discovery—When a FCoE device is connected to the fabric, it sends out a Discovery Solicitation message. A Fibre Channel Forwarder (FCF) or a switch responds to the message with a Solicited Advertisement that provides an FCF MAC address to use for subsequent logins.
- FCoE Virtual Link instantiation—FIP defines the encapsulation of fabric login (FLOGI), fabric discovery (FDISC), logout (LOGO), and exchange link parameters (ELP) frames along with the corresponding reply frames. The FCoE devices use these messages to perform a fabric login.
- FCoE Virtual Link maintenance—FIP periodically sends maintenance messages between the switch and the CNA to ensure the connection is still valid.

### FCoE Frame Format

FCoE is encapsulated in an Ethernet packet with a dedicated EtherType, 0x8906. That packet has a 4-bit version field. The other header fields in the frame (the source and destination MAC addresses, VLAN tags, and frame markers) are all standard Ethernet fields. Reserved bits pad the FCoE frame to the IEEE 802.3 minimum packet length of 64 bytes.

A Fibre Channel frame consists of 36 bytes of headers and up to 2112 bytes of data for a total maximum size of 2148 bytes. The encapsulated Fibre Channel frame has all the standard headers, which allow it to be passed to the storage network without further modification. To accommodate the maximum Fibre Channel frame in an FCoE frame, the class-foe is defined with a default MTU of 2240 bytes.

### VLAN Tagging for FCoE Frames

The Ethernet frames that are sent by the switch to the adapter may include the IEEE 802.1Q tag. This tag includes a field for the class of service (CoS) value used by the priority flow control (PFC). The IEEE 802.1Q tag also includes a VLAN field.

The Cisco Nexus 5000 Series switch expects frames from a FIP T11-compliant CNA to be tagged with the VLAN tag for the FCoE VLAN. Frames that are not correctly tagged are discarded.

The switch expects frames from a pre-FIP CNA to be priority tagged with the FCoE CoS value. The switch will still accept untagged frames from the CNA.

## FIP Ethernet Frame Format

FIP is encapsulated in an Ethernet packet with a dedicated EtherType, 0x8914. The packet has a 4-bit version field. Along with the source and destination MAC addresses, the FIP packet also contains a FIP operation code and a FIP operation subcode. The following table describes the FIP operation codes.

**Table 55: FIP Operation Codes**

| FIP Operation Code | FIP Subcode | FIP Operation                      |
|--------------------|-------------|------------------------------------|
| 0x0001             | 0x01        | Discovery Solicitation             |
|                    | 0x02        | Discovery Advertisement            |
| 0x0002             | 0x01        | Virtual Link Instantiation Request |
|                    | 0x02        | Virtual Link Instantiation Reply   |
| 0x0003             | 0x01        | FIP Keep Alive                     |
|                    | 0x02        | FIP Clear Virtual Links            |
| 0x0004             | 0x01        | FIP VLAN Request                   |
|                    | 0x02        | FIP VLAN Notification              |

## Pre-FIP Virtual Link Instantiation

Pre-FIP virtual link instantiation consists of two phases; link discovery using the Data Center Bridging Exchange protocol (DCBX), which is followed by Fabric Login.

The Cisco Nexus 5000 Series switch is backward compatible with Gen-1 CNAs that operate in pre-FIP mode.



### Note

Pre-FIP is also known as the Cisco, Intel, Nuova Data Center Bridging Exchange (CIN-DCBX) protocol.

## Information About DCBX

### Data Center Bridging Exchange Protocol

The Data Center Bridging Exchange (DCBX) protocol is an extension of the Link Layer Discovery Protocol (LLDP). DCBX end points exchange request and acknowledgment messages. For flexibility, parameters are coded in a type-length-value (TLV) format.

The Cisco Nexus 5000 Series switch supports two versions of DCBX:

- CEE-DCBX—The Converged Enhanced Ethernet DCBX is supported on all T11-compliant Gen-2 CNAs

- CIN-DCBX—The Cisco, Intel, Nuova DCBX is supported on Gen-1 converged network adapters (CNAs). CIN-DCBX is used to perform link detection in addition to other functions.

DCBX runs on the physical Ethernet link between the Cisco Nexus 5000 Series switch and the CNA. By default, DCBX is enabled on Ethernet interfaces. When an Ethernet interface is brought up, the switch automatically starts to communicate with the CNA.

During the normal operation of FCoE between the switch and the CNA, DCBX provides link-error detection. DCBX is also used to negotiate capabilities between the switch and the CNA and to send configuration values to the CNA.

The CNAs that are connected to a Cisco Nexus 5000 Series switch are programmed to accept the configuration values sent by the switch, allowing the switch to distribute configuration values to all attached CNAs, which reduces the possibility of configuration errors and simplifies CNA administration.

### DCBX Feature Negotiation

The switch and CNA exchange capability information and configuration values. The Cisco Nexus 5000 Series switches support the following capabilities:

- FCoE—If the CNA supports FCoE capability, the switch sends the IEEE 802.1p CoS value to be used with FCoE packets.
- Priority Flow Control (PFC)—If the adapter supports PFC, the switch sends the IEEE 802.1p CoS values to be enabled with PFC.
- Priority group type-length-value (TLV)
- Ethernet logical link up and down signal
- FCoE logical link up and down signal for pre-FIP CNAs

The following rules determine whether the negotiation results in a capability being enabled:

- If a capability and its configuration values match between the switch and the CNA, the feature is enabled.
- If a capability matches, but the configuration values do not match, the following occurs:
  - If the CNA is configured to accept the switch configuration value, the capability is enabled using the switch value.
  - If the CNA is not configured to accept the switch configuration value, the capability remains disabled.
- If the CNA does not support a DCBX capability, that capability remains disabled.
- If the CNA does not implement DCBX, all capabilities remain disabled.



#### Note

The Cisco Nexus 5000 Series switch provides CLI commands to manually override the results of the PFC negotiation with the adapter. On a per-interface basis, you can force capabilities to be enabled or disabled.

## Lossless Ethernet

Standard Ethernet is a best-effort medium which means that it lacks any form of flow control. In the event of congestion or collisions, Ethernet will drop packets. The higher level protocols detect the missing data and retransmit the dropped packets.

To properly support Fibre Channel, Ethernet has been enhanced with a priority flow control (PFC) mechanism.

## Logical Link Up/Down

The optional N5K-M1404 or N5K-M1008 expansion modules provide native 1/2/4-Gigabit Fibre Channel ports to connect the Cisco Nexus 5000 Series switch to other Fibre Channel devices. On a native Fibre Channel link, some configuration actions (such as changing the VSAN) require that you reset the interface status. When you reset the interface status, the switch disables the interface and then immediately reenables the interface.

If an Ethernet link provides FCoE service, do not reset the physical link because this action is disruptive to all traffic on the link.

The logical link up/down feature allows the switch to reset an individual virtual link. The logical link down is signaled with a FIP Clear Virtual Link message.

For pre-FIP CNAs, the switch sends a DCBX message to request the CNA to reset only the virtual Fibre Channel interface.



### Note

If the CNA does not support the logical link level up/down feature, the CNA resets the physical link. In this case, all traffic on the Ethernet interface is disrupted.

DCBX-based FC Logical Link Status signaling only applies to FCoE sessions to pre-FIP CNAs.

## Converged Network Adapters

The following types of CNAs are available:

- Hardware adapter
  - Works with the existing Fibre Channel host bus adapter (HBA) driver and Ethernet Network Interface Card (NIC) driver in the server.
  - Server operating system view of the network is unchanged; the CNA presents a SAN interface and a LAN interface to the operating system.
- FCoE software stack
  - Runs on existing 10-Gigabit Ethernet adapters.

Two generations of CNAs are supported by the Cisco Nexus 5000 Series switch:

- A FIP adapter uses the FIP to exchange information about its available capabilities and to negotiate the configurable values with the switch.

- A pre-FIP adapter uses DCBX to exchange information about its available capabilities and to negotiate the configurable values with the switch.

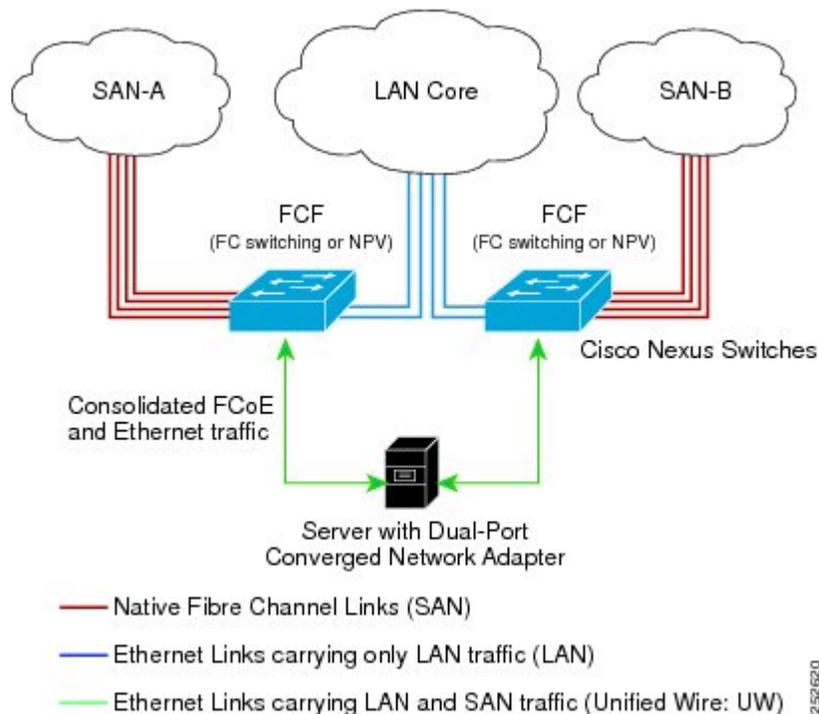
To reduce configuration errors and simplify administration, the switch distributes the configuration data to all the connected adapters.

## FCoE Topologies

### Directly Connected CNA Topology

The Cisco Nexus 5000 Series switch can be deployed as a Fibre Channel Forwarder (FCF) as shown in the following figure.

**Figure 37: Directly Connected Fibre Channel Forwarder**



The following rules are used to process FIP frames to avoid the FCF being used as a transit between an FCoE node (ENode) and another FCF. These rules also prevent login sessions between ENodes and FCFs in different fabrics.

- FIP solicitation and login frames received from the CNAs are processed by the FCF and are not forwarded.
- If an FCF receives solicitations and advertisements from other FCFs over an interface, the following occurs:
  - The frames are ignored and discarded if the FC-MAP value in the frame matches the value of the FCF (the FCF is in the same fabric)

- The interface is placed in the "FCoE Isolated" state if the FC-MAP value in the FIP frame does not match that of the FCF (the FCF is in a different fabric)

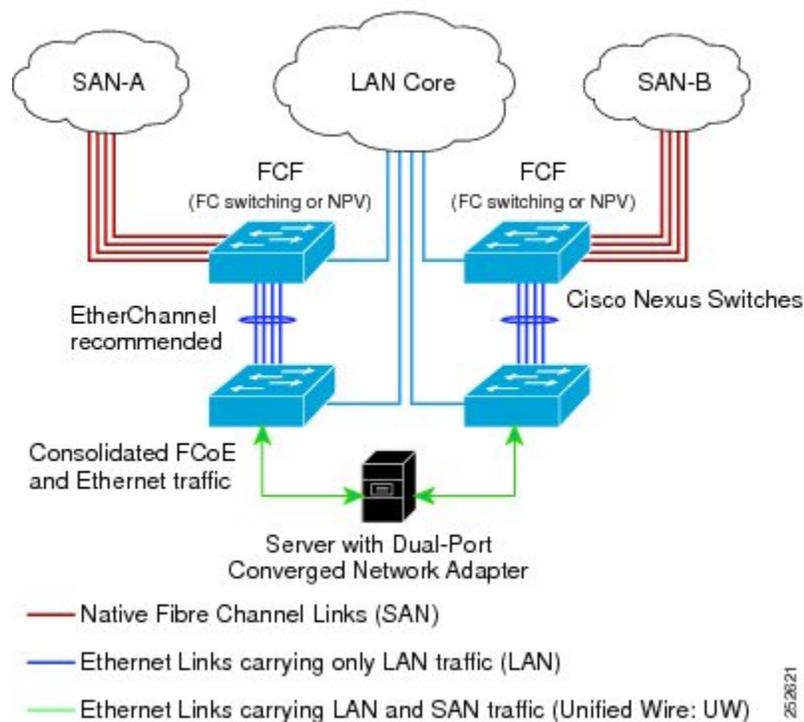
CNAs cannot discover or login to FCFs that are reachable only through a transit Cisco Nexus 5000 Series FCF. The Cisco Nexus 5000 Series switch cannot perform the FCoE transit function between a CNA and another FCF due to hardware limitations.

Because the Cisco Nexus 5000 Series FCF cannot perform the transit FCoE function, you must design your network topology so that the active STP path of FCoE VLANs is always over the directly connected links between the CNA and the FCF. Make sure that you configure the FCoE VLAN on the directly connected links only.

## Remotely Connected CNA Topology

The Cisco Nexus 5000 Series switch can be deployed as a Fibre Channel Forwarder (FCF) for remotely connected CNAs, but not as a FIP Snooping Bridge, as shown in the following figure.

**Figure 38: Remotely Connected Fibre Channel Forwarder**



The following rules are used to process FIP frames to avoid the FCF being used as a transit between an ENode and another FCF. These rules also prevent login sessions between ENodes and FCFs in different fabrics.

- FIP solicitation and login frames received from the CNAs are processed by the FCF and are not forwarded.
- If an FCF receives solicitations and advertisements from other FCFs over an interface, the following occurs:

- The frames are ignored and discarded if the FC-MAP value in the frame matches the value of the FCF (the FCF is in the same fabric)
- The interface is placed in the "FCoE Isolated" state if the FC-MAP value in the FIP frame does not match that of the FCF (the FCF is in a different fabric)

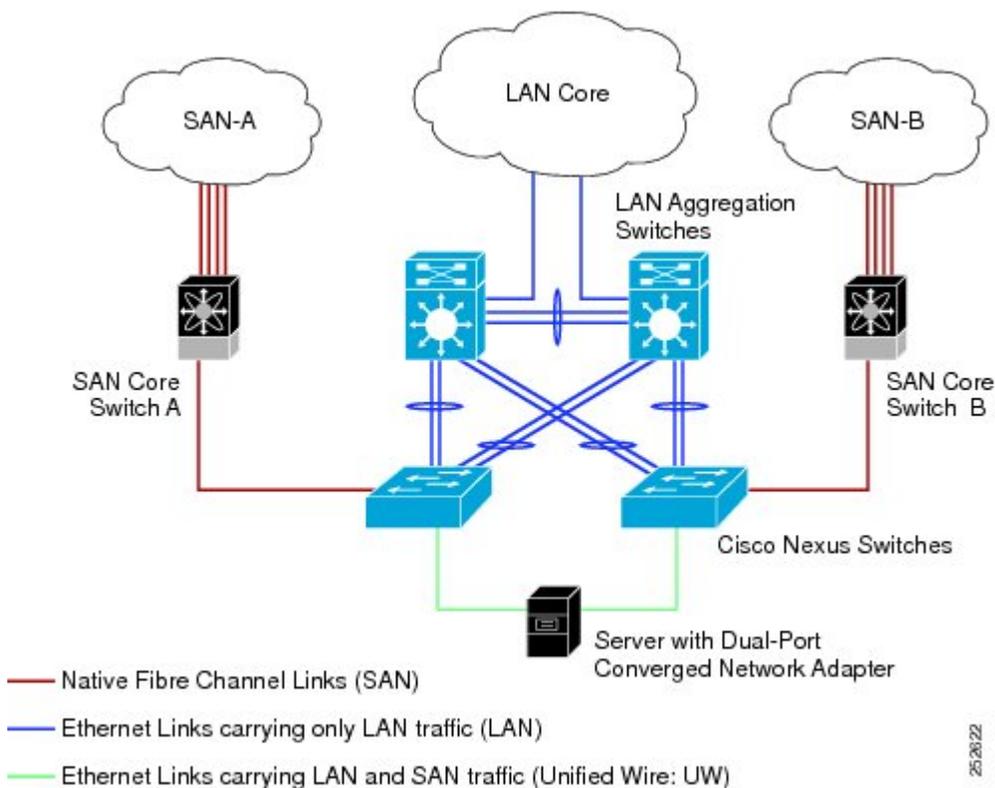
Because the Cisco Nexus 5000 Series FCF cannot perform the transit FCoE function, you must design your network topology so that the active STP path of FCoE VLANs is always over the directly connected links between the CNA and the FCF. Make sure that you configure the FCoE VLAN on the directly connected links only.

## FCoE Best Practices

### Directly Connected CNA Best Practice

The following figure shows a best practices topology for an access network using directly connected CNAs with Cisco Nexus 5000 Series switches.

**Figure 39: Directly Connected CNA**



Follow these configuration best practices for the deployment topology in the preceding figure:

- 1 You must configure a unique dedicated VLAN at every converged access switch to carry traffic for each Virtual Fabric (VSAN) in the SAN (for example, VLAN 1002 for VSAN 1, VLAN 1003 for VSAN 2, and so on). If you enable MST, you must use a separate MST instance for FCoE VLANs.
- 2 You must configure the unified fabric (UF) links as trunk ports. Do not configure the FCoE VLAN as a native VLAN. You must configure all FCoE VLANs as members of the UF links to allow extensions for VF\_Port trunking and VSAN management for the virtual Fibre Channel interfaces.

**Note**

---

A unified fabric link carries both Ethernet and FCoE traffic.

---

- 3 You must configure the UF links as spanning-tree edge ports.
- 4 You must not configure the FCoE VLANs as members of Ethernet links that are not designated to carry FCoE traffic because you want to ensure the scope of the STP for the FCoE VLANs is limited to UF links only.
- 5 If the converged access switches (in the same SAN fabric or in another) need to be connected to each other over Ethernet links for a LAN alternate path, then such links must explicitly be configured to exclude all FCoE VLANs from membership. This action ensures that the scope of the STP for the FCoE VLANs is limited to UF links only.
- 6 You must use separate FCoE VLANs for FCoE in SAN-A and SAN-B.

**Note**

---

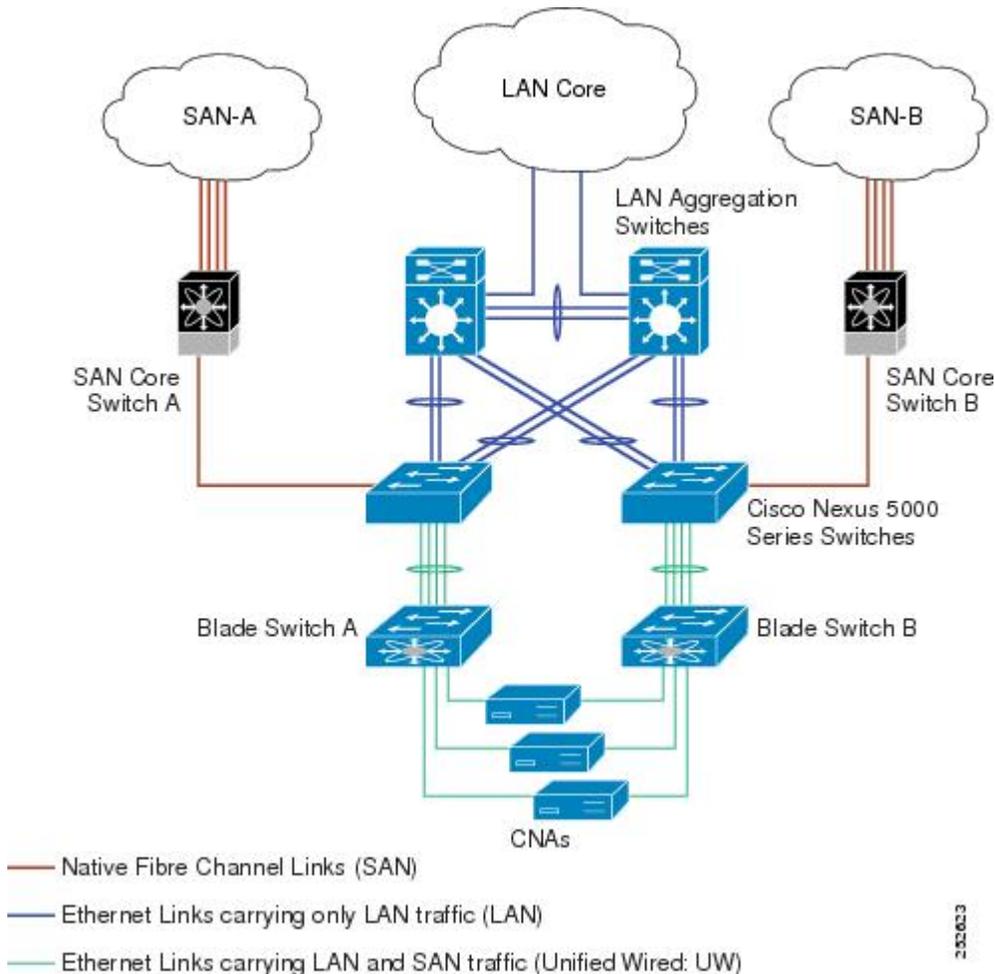
All Gen-1 (pre-FIP) and Gen-2 (FIP) CNAs are supported in a directly connected topology.

---

## Remotely Connected CNA Best Practice

The following figure shows a best practices topology for an access network using remotely connected CNAs with Cisco Nexus 5000 Series switches.

**Figure 40: Remotely Connected CNAs**



Follow these configuration best practices for the deployment topology in the preceding figure:

- 1 You must configure a unique dedicated VLAN at every converged access switch to carry traffic for each Virtual Fabric (VSAN) in the SAN (for example, VLAN 1002 for VSAN 1, VLAN 1003 for VSAN 2, and so on). If you enable MST, you must use a separate MST instance for FCoE VLANs.
- 2 You must configure the unified fabric (UF) links as trunk ports. Do not configure the FCoE VLAN as a native VLAN. You must configure all FCoE VLANs as members of the UF links to allow extensions for VF\_Port trunking and VSAN management for the virtual Fibre Channel interfaces.

**Note**


---

A unified fabric link carries both Ethernet and FCoE traffic.

---

- 3 You must configure the CNAs and the blade switches as spanning-tree edge ports.
- 4 A blade switch must connect to exactly one Cisco Nexus 5000 Series converged access switch, preferably over an EtherChannel, to avoid disruption due to STP reconvergence on events such as provisioning new links or blade switches.
- 5 You must configure the Cisco Nexus 5000 Series converged access switch with a better STP priority than the blade switches that are connected to it. This requirement allows you to create an island of FCoE VLANs where the converged access switch is the spanning-tree root and all the blade switches connected to it become downstream nodes.
- 6 Do not configure the FCoE VLANs as members of Ethernet links that are not designated to carry FCoE traffic because you want to ensure that the scope of the STP for the FCoE VLANs is limited to UF links only.
- 7 If the converged access switches and/or the blade switches need to be connected to each over Ethernet links for the purposes of LAN alternate pathing, then such links must explicitly be configured to exclude all FCoE VLANs from membership. This will ensure the scope of the spanning-tree protocol for FCoE VLANs is limited to UF links only.
- 8 You must use separate FCoE VLANs for FCoE in SAN-A and SAN-B.

**Note**


---

A remotelyconnected topology is supported only with Gen-2 (FIP) CNAs.

---

## Licensing Requirements for FCoE

On Cisco Nexus 5000 Series switches, FCoE capability is included in the Storage Protocol Services License.

Before using FCoE capabilities, you must ensure the following:

- The correct license is installed (N5010SS or N5020SS).
- You activated FCoE by entering the **feature fcoe** command in configuration mode.

## Configuring FCoE

### Enabling FCoE

You can enable FCoE on the switch.

**Note**

All the Fibre Channel features of the Cisco Nexus 5000 Series switch are packaged in the FC Plugin. When you enable FCoE, the switch software checks for the FC\_FEATURES\_PKG license. If it finds the license, the software loads the plugin. If the license is not found, the software loads the plugin with a grace period of 180 days.

After the FC Plugin is loaded, the following occurs:

- All Fibre Channel and FCoE related CLI are available
- The Fibre Channel interfaces of any installed Expansion Modules are available

If after 180 days, a valid license is not found, the FC Plugin is disabled. At the next switch reboot, all FCoE commands are removed from the CLI and the FCoE configuration is deleted.

**Before You Begin**

You need to have the FC\_FEATURES\_PKG (N5010SS or N5020SS) license installed.

**SUMMARY STEPS**

1. switch# **configure terminal**
2. switch(config)# **feature fcoe**

**DETAILED STEPS**

|               | <b>Command or Action</b>            | <b>Purpose</b>               |
|---------------|-------------------------------------|------------------------------|
| <b>Step 1</b> | switch# <b>configure terminal</b>   | Enters configuration mode.   |
| <b>Step 2</b> | switch(config)# <b>feature fcoe</b> | Enables the FCoE capability. |

This example shows how to enable FCoE on the switch:

```
switch# configure terminal
switch(config)# feature fcoe
```

**Disabling FCoE**

After you disable FCoE, all FCoE commands are removed from the CLI and the FCoE configuration is deleted.

**SUMMARY STEPS**

1. switch# **configure terminal**
2. switch(config)# **no feature fcoe**

## DETAILED STEPS

|               | Command or Action                      | Purpose                       |
|---------------|----------------------------------------|-------------------------------|
| <b>Step 1</b> | switch# <b>configure terminal</b>      | Enters configuration mode.    |
| <b>Step 2</b> | switch(config)# <b>no feature fcoe</b> | Disables the FCoE capability. |

This example shows how to disable FCoE on the switch:

```
switch# configure terminal
switch(config)# no feature fcoe
```

## Disabling LAN Traffic on an FCoE Link

You can disable LAN traffic on an FCoE link.

DCBX allows the switch to send a LAN Logical Link Status (LLS) message to a directly-connected CNA. Enter the **shutdown lan** command to send an LLS-Down message to the CNA. This command causes all VLANs on the interface that are not enabled for FCoE to be brought down. If a VLAN on the interface is enabled for FCoE, it continues to carry SAN traffic without any interruption.

## SUMMARY STEPS

1. switch# **configure terminal**
2. switch(config)# **interface ethernet slot/port**
3. switch(config-if)# **shutdown lan**
4. (Optional) switch(config-if)# **no shutdown lan**

## DETAILED STEPS

|               | Command or Action                                   | Purpose                                                                                                                                 |
|---------------|-----------------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------|
| <b>Step 1</b> | switch# <b>configure terminal</b>                   | Enters configuration mode.                                                                                                              |
| <b>Step 2</b> | switch(config)# <b>interface ethernet slot/port</b> | Specifies an interface to configure, and enters interface configuration mode.                                                           |
| <b>Step 3</b> | switch(config-if)# <b>shutdown lan</b>              | Shuts down Ethernet traffic on the interface. If the interface is part of an FCoE VLAN, the shutdown has no impact on the FCoE traffic. |
| <b>Step 4</b> | switch(config-if)# <b>no shutdown lan</b>           | (Optional)<br>Reenables Ethernet traffic on the interface.                                                                              |

## Configuring the FC-Map

You can prevent data corruption due to cross-fabric talk by configuring an FC-Map which identifies the Fibre Channel fabric for this Cisco Nexus 5000 Series switch. When the FC-Map is configured, the switch discards the MAC addresses that are not part of the current fabric.

### SUMMARY STEPS

1. switch# **configure terminal**
2. switch(config)# **fcoe fcmmap fabric-map**
3. (Optional) switch(config)# **no fcoe fcmmap fabric-map**

### DETAILED STEPS

|               | Command or Action                                | Purpose                                                                                              |
|---------------|--------------------------------------------------|------------------------------------------------------------------------------------------------------|
| <b>Step 1</b> | switch# <b>configure terminal</b>                | Enters configuration mode.                                                                           |
| <b>Step 2</b> | switch(config)# <b>fcoe fcmmap fabric-map</b>    | Configures the global FC-Map. The default value is 0E.FC.00. The range is from 0E.FC.00 to 0E.FC.FF. |
| <b>Step 3</b> | switch(config)# <b>no fcoe fcmmap fabric-map</b> | (Optional)<br>Resets the global FC-Map to the default value of 0E.FC.00.                             |

This example shows how to configure the global FC-Map:

```
switch# configure terminal
switch(config)# fcoe fcmmap 0e.fc.2a
```

## Configuring the Fabric Priority

The Cisco Nexus 5000 Series switch advertises its priority. The priority is used by the CNAs in the fabric to determine the best switch to connect to.

### SUMMARY STEPS

1. switch# **configure terminal**
2. switch(config)# **fcoe fcf-priority fabric-priority**
3. (Optional) switch(config)# **no fcoe fcf-priority fabric-priority**

### DETAILED STEPS

|               | Command or Action                 | Purpose                    |
|---------------|-----------------------------------|----------------------------|
| <b>Step 1</b> | switch# <b>configure terminal</b> | Enters configuration mode. |

|               | Command or Action                                                  | Purpose                                                                                                       |
|---------------|--------------------------------------------------------------------|---------------------------------------------------------------------------------------------------------------|
| <b>Step 2</b> | switch(config)# <b>fcoe fcf-priority</b> <i>fabric-priority</i>    | Configures the global fabric priority. The default value is 128. The range is from 0 (higher) to 255 (lower). |
| <b>Step 3</b> | switch(config)# <b>no fcoe fcf-priority</b> <i>fabric-priority</i> | (Optional)<br>Resets the global fabric priority to the default value of 128.                                  |

This example shows how to configure the global fabric priority:

```
switch# configure terminal
switch(config)# fcoe fcf-priority 42
```

## Setting the Advertisement Interval

You can configure the interval for Fibre Channel fabric advertisement on the switch.

### SUMMARY STEPS

1. switch# **configure terminal**
2. switch(config)# **fcoe fka-adv-period** *interval*
3. (Optional) switch(config)# **no fcoe fka-adv-period** *interval*

### DETAILED STEPS

|               | Command or Action                                             | Purpose                                                                                                                  |
|---------------|---------------------------------------------------------------|--------------------------------------------------------------------------------------------------------------------------|
| <b>Step 1</b> | switch# <b>configure terminal</b>                             | Enters configuration mode.                                                                                               |
| <b>Step 2</b> | switch(config)# <b>fcoe fka-adv-period</b> <i>interval</i>    | Configures the advertisement interval for the fabric. The default value is 8 seconds. The range is from 4 to 60 seconds. |
| <b>Step 3</b> | switch(config)# <b>no fcoe fka-adv-period</b> <i>interval</i> | (Optional)<br>Resets the advertisement interval for the fabric to its default value of 8 seconds.                        |

This example shows how to configure the advertisement interval for the fabric:

```
switch# configure terminal
switch(config)# fcoe fka-adv-period 42
```

## Configuring LLDP

### Configuring Global LLDP Commands

You can set global LLDP settings. These settings include the length of time before discarding LLDP information received from peers, the length of time to wait before performing LLDP initialization on any interface, and the rate at which LLDP packets are sent.

To configure LLDP settings, perform this task:

#### SUMMARY STEPS

1. switch# **configure terminal**
2. switch(config)# **lldp** {holdtime seconds | reinit seconds | timer seconds}
3. switch(config)# **no lldp** {holdtime | reinit | timer}

#### DETAILED STEPS

|               | Command or Action                                                               | Purpose                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
|---------------|---------------------------------------------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Step 1</b> | switch# <b>configure terminal</b>                                               | Enters configuration mode.                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>Step 2</b> | switch(config)# <b>lldp</b> {holdtime seconds   reinit seconds   timer seconds} | Configures LLDP options.<br><br>Use the <b>holdtime</b> option to set the length of time (10 to 255 seconds, default 120 seconds) that a device should save LLDP information received before discarding it.<br><br>Use the <b>reinit</b> option to set the length of time (1 to 10 seconds, default 2 seconds) to wait before performing LLDP initialization on any interface.<br><br>Use the <b>timer</b> option to set the rate (5 to 254 seconds, default 30 seconds) at which LLDP packets are sent. |
| <b>Step 3</b> | switch(config)# <b>no lldp</b> {holdtime   reinit   timer}                      | Reset the LLDP values to their defaults.                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |

This example shows how to set LLDP timer option to 15 seconds:

```
switch# configure terminal
switch(config)# lldp timer 15
```

### Configuring Interface LLDP Commands

To configure the LLDP feature for a physical Ethernet interface, perform this task:

## SUMMARY STEPS

1. switch# **configure terminal**
2. switch(config)# **interface** *type slot/port*
3. switch(config-if)# [**no**] **lldp** {**receive** | **transmit**}

## DETAILED STEPS

|               | Command or Action                                                                 | Purpose                                                                                                                                |
|---------------|-----------------------------------------------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------|
| <b>Step 1</b> | switch# <b>configure terminal</b>                                                 | Enters configuration mode.                                                                                                             |
| <b>Step 2</b> | switch(config)# <b>interface</b> <i>type slot/port</i>                            | Selects the interface to change.                                                                                                       |
| <b>Step 3</b> | switch(config-if)# [ <b>no</b> ] <b>lldp</b> { <b>receive</b>   <b>transmit</b> } | Sets the selected interface to either receive or transmit.<br>The <b>no</b> form of the command disables the LLDP transmit or receive. |

This example shows how to set an interface to transmit LLDP packets:

```
switch# configure terminal
switch(config)# interface ethernet 1/2
switch(config-if)# lldp transmit
```

This example shows how to configure an interface to disable LLDP:

```
switch# configure terminal
switch(config)# interface ethernet 1/2
switch(config-if)# no lldp transmit
switch(config-if)# no lldp receive
```

## Verifying FCoE Configuration

To verify FCoE configuration information, perform one of these tasks:

| Command                                                               | Purpose                                                        |
|-----------------------------------------------------------------------|----------------------------------------------------------------|
| switch# <b>show fcoe</b>                                              | Displays whether FCoE is enabled on the switch.                |
| switch# <b>show fcoe database</b>                                     | Displays the contents of the FCoE database.                    |
| switch# <b>show interface</b> [ <i>interface number</i> ] <b>fcoe</b> | Displays the FCoE settings for an interface or all interfaces. |
| switch# <b>show lldp</b>                                              | Displays LLDP configuration.                                   |

This example shows how to verify that the FCoE capability is enabled:

```
switch# show fcoe
Global FCF details
 FCF-MAC is 00:0d:ec:6d:95:00
 FCF-MAP is 0e:fc:00
 FCF Priority is 128
 FKA Advertisement period for FCF is 8 seconds
```

This example shows how to display the FCoE database:

```
switch# show fcoe database

INTERFACE FCID PORT NAME MAC ADDRESS

vfc3 0x490100 21:00:00:1b:32:0a:e7:b8 00:c0:dd:0e:5f:76
```

This example shows how to display the FCoE settings for an interface.

```
switch# show interface ethernet 1/37 fcoe
Ethernet1/37 is FCoE UP
 vfc3 is Up
 FCID is 0x490100
 PWWN is 21:00:00:1b:32:0a:e7:b8
 MAC addr is 00:c0:dd:0e:5f:76
```

This example shows how to display LLDP interface information:

```
switch# show lldp interface ethernet 1/2
tx_enabled: TRUE
rx_enabled: TRUE
dcbx_enabled: TRUE
Port MAC address: 00:0d:ec:a3:5f:48
Remote Peers Information
No remote peers exist
```

This example shows how to display LLDP neighbor information:

```
switch# show lldp neighbors
LLDP Neighbors

Remote Peers Information on interface Eth1/40
Remote peer's MSAP: length 12 Bytes:
00 c0 dd 0e 5f 3a 00 c0 dd 0e 5f 3a

LLDP TLV's
LLDP TLV type:Chassis ID LLDP TLV Length: 7
LLDP TLV type:Port ID LLDP TLV Length: 7
LLDP TLV type:Time to Live LLDP TLV Length: 2
LLDP TLV type:LLDP Organizationally Specific LLDP TLV Length: 55
LLDP TLV type:LLDP Organizationally Specific LLDP TLV Length: 5
LLDP TLV type:END of LLDPDU LLDP TLV Length: 0

Remote Peers Information on interface Eth1/34
Remote peer's MSAP: length 12 Bytes:
00 0d ec a3 27 40 00 0d ec a3 27 69

LLDP TLV's
LLDP TLV type:Chassis ID LLDP TLV Length: 7
LLDP TLV type:Port ID LLDP TLV Length: 7
LLDP TLV type:Time to Live LLDP TLV Length: 2
LLDP TLV type:LLDP Organizationally Specific LLDP TLV Length: 55
LLDP TLV type:LLDP Organizationally Specific LLDP TLV Length: 5
LLDP TLV type:END of LLDPDU LLDP TLV Length: 0

Remote Peers Information on interface Eth1/33
Remote peer's MSAP: length 12 Bytes:
00 0d ec a3 27 40 00 0d ec a3 27 68
```

```
LLDP TLV's
LLDP TLV type:Chassis ID LLDP TLV Length: 7
LLDP TLV type:Port ID LLDP TLV Length: 7
LLDP TLV type:Time to Live LLDP TLV Length: 2
LLDP TLV type:LLDP Organizationally Specific LLDP TLV Length: 55
LLDP TLV type:LLDP Organizationally Specific LLDP TLV Length: 5
LLDP TLV type:END of LLDPDU LLDP TLV Length: 0
```

This example shows how to display LLDP timer information:

```
switch# show lldp timers
LLDP Timers
holdtime 120 seconds
reinit 2 seconds
msg_tx_interval 30 seconds
```

This example shows how to display LLDP counters:

```
switch# show lldp traffic
LLDP traffic statistics:

Total frames out: 8464
Total Entries aged: 6
Total frames in: 6342
Total frames received in error: 2
Total frames discarded: 2
Total TLVs unrecognized: 0
```





## Configuring FCoE VLANs and Virtual Interfaces

This chapter describes how to configure Fibre Channel over Ethernet (FCoE) VLANs and virtual interfaces on Cisco Nexus 5000 Series switches. It contains the following sections:

- [Information About Virtual Interfaces, page 457](#)
- [Guidelines and Limitations for FCoE VLANs and Virtual Interfaces, page 457](#)
- [Configuring Virtual Interfaces, page 459](#)
- [Verifying the Virtual Interface, page 461](#)
- [Mapping VSANs to VLANs Example Configuration, page 463](#)

### Information About Virtual Interfaces

Cisco Nexus 5000 Series switches support Fibre Channel over Ethernet (FCoE), which allows Fibre Channel and Ethernet traffic to be carried on the same physical Ethernet connection between the switch and the servers.

The Fibre Channel portion of FCoE is configured as a virtual Fibre Channel interface. Logical Fibre Channel features (such as interface mode) can be configured on virtual Fibre Channel interfaces.

A virtual Fibre Channel interface must be bound to an interface before it can be used. The binding is to a physical Ethernet interface (when the converged network adapter (CNA) is directly connected to the Cisco Nexus 5000 Series switch), a MAC address (when the CNA is remotely connected over a Layer 2 bridge), or an EtherChannel when the CNA connects to the Fibre Channel Forwarder (FCF) over a virtual port channel (vPC).

### Guidelines and Limitations for FCoE VLANs and Virtual Interfaces

Follow these guidelines and limitations when configuring FCoE VLANs and Virtual Interfaces:

- Each virtual Fibre Channel interface must be bound to an FCoE-enabled Ethernet or EtherChannel interface or to the MAC address of a remotely connected adapter. FCoE is supported on 10-Gigabit Ethernet interfaces.

The Ethernet or EtherChannel interface that you bind the virtual Fibre Channel interface to must be configured as follows:

- The Ethernet or EtherChannel interface must be a trunk port (use the **switchport mode trunk** command).
- The FCoE VLAN that corresponds to a virtual Fibre Channel's VSAN must be in the allowed VLAN list.
- You must not configure an FCoE VLAN as the native VLAN of the trunk port.




---

**Note** The native VLAN is the default VLAN on a trunk. Any untagged frames transit the trunk as native VLAN traffic.

---

- You should use an FCoE VLAN only for FCoE.
- Do not use the default VLAN, VLAN1, as an FCoE VLAN.
- You must configure the Ethernet interface as PortFast (use the **spanning-tree port type edge trunk** command).




---

**Note** You are not required to configure trunking on the server interface even if the switch interface is configured with trunking enabled. All non-FCoE traffic from the server will be passed on the native VLAN.

---

- Each virtual Fibre Channel interface is associated with only one VSAN.
- You must map any VSAN with associated virtual Fibre Channel interfaces to a dedicated FCOE-enabled VLAN.
- FCoE is not supported on private VLANs.
- If the converged access switches (in the same SAN fabric or in another) need to be connected to each other over Ethernet links for a LAN alternate path, then you must explicitly configure such links to exclude all FCoE VLANs from membership.
- You must use separate FCoE VLANs for FCoE in SAN-A and SAN-B fabrics.
- FCoE connectivity to pre-FIP CNAs over virtual port channels (vPCs) is not supported.




---

**Note** Virtual interfaces are created with the administrative state set to down. You must explicitly configure the administrative state to bring the virtual interface into operation.

---

## Configuring Virtual Interfaces

### Mapping a VSAN to a VLAN

A unique, dedicated VLAN must be configured at every converged access switch to carry traffic for each Virtual Fabric (VSAN) in the SAN (for example, VLAN 1002 for VSAN 1, VLAN 1003 for VSAN 2, and so on). If MST is enabled, a separate MST instance must be used for FCoE VLANs.

#### SUMMARY STEPS

1. switch# **configure terminal**
2. switch(config)# **vlan *vlan-id***
3. switch(config-vlan)# **fcoe [*vsan vsan-id*]**
4. switch(config-vlan)# **exit**

#### DETAILED STEPS

|               | Command or Action                                      | Purpose                                                                                                                                                                                                             |
|---------------|--------------------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Step 1</b> | switch# <b>configure terminal</b>                      | Enters configuration mode.                                                                                                                                                                                          |
| <b>Step 2</b> | switch(config)# <b>vlan <i>vlan-id</i></b>             | Enters VLAN configuration mode. The VLAN number range is from 1 to 4096.                                                                                                                                            |
| <b>Step 3</b> | switch(config-vlan)# <b>fcoe [<i>vsan vsan-id</i>]</b> | Enables FCoE for the specified VLAN. If you do not specify a VSAN number, a mapping is created from this VLAN to the VSAN with the same number.<br><br>Configures the mapping from this VLAN to the specified VSAN. |
| <b>Step 4</b> | switch(config-vlan)# <b>exit</b>                       | Exits VLAN configuration mode.                                                                                                                                                                                      |

This example shows how to map VLAN 200 to VSAN 2:

```
switch(config)# vlan 200
switch(config-vlan)# fcoe vsan 2
```

### Creating a Virtual Fibre Channel Interface

You can create a virtual Fibre Channel interface. You must bind the virtual Fibre Channel interface to a physical interface before it can be used.

## SUMMARY STEPS

1. switch# **configure terminal**
2. switch(config)# **interface vfc vfc-id**
3. switch(config-if)# **bind {interface {ethernet slot/port | port-channel channel-number} | mac-address MAC-address}**
4. (Optional) switch(config-if)# **no bind {interface {ethernet slot/port | port-channel channel-number} | mac-address MAC-address}**
5. (Optional) switch(config)# **no interface vfc vfc-id**

## DETAILED STEPS

|               | Command or Action                                                                                                          | Purpose                                                                                                                                                                                  |
|---------------|----------------------------------------------------------------------------------------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Step 1</b> | switch# <b>configure terminal</b>                                                                                          | Enters configuration mode.                                                                                                                                                               |
| <b>Step 2</b> | switch(config)# <b>interface vfc vfc-id</b>                                                                                | Creates a virtual Fibre Channel interface (if it does not already exist) and enters interface configuration mode.<br><br>The virtual Fibre Channel interface ID range is from 1 to 8192. |
| <b>Step 3</b> | switch(config-if)# <b>bind {interface {ethernet slot/port   port-channel channel-number}   mac-address MAC-address}</b>    | Binds the virtual Fibre Channel interface to the specified interface.                                                                                                                    |
| <b>Step 4</b> | switch(config-if)# <b>no bind {interface {ethernet slot/port   port-channel channel-number}   mac-address MAC-address}</b> | (Optional)<br>Unbinds the virtual Fibre Channel interface from the specified interface.                                                                                                  |
| <b>Step 5</b> | switch(config)# <b>no interface vfc vfc-id</b>                                                                             | (Optional)<br>Deletes a virtual Fibre Channel interface.                                                                                                                                 |

This example shows how to bind a virtual Fibre Channel interface to an Ethernet interface:

```
switch# configure terminal
switch(config)# interface vfc 4
switch(config-if)# bind interface ethernet 1/4
```

This example shows how to bind a virtual Fibre Channel interface to create a vPC:

```
switch# configure terminal
switch(config)# interface vfc 3
switch(config-if)# bind interface port-channel 1
```

This example shows how to bind a virtual Fibre Channel interface to a MAC address:

```
switch# configure terminal
switch(config)# interface vfc 2
switch(config-if)# bind mac-address 00:0a:00:00:00:36
```

This example shows how to delete a virtual Fibre Channel interface:

```
switch# configure terminal
switch(config)# no interface vfc 4
```

## Associating a Virtual Fibre Channel Interface to a VSAN

A unique, dedicated VLAN must be configured at every converged access switch to carry traffic for each Virtual Fabric (VSAN) in the SAN (for example, VLAN 1002 for VSAN 1, VLAN 1003 for VSAN 2, and so on). If MST is enabled, a separate MST instance must be used for FCoE VLANs.

### SUMMARY STEPS

1. switch# **configure terminal**
2. switch(config)# **vsan database**
3. switch(config-vsan)# **vsan vsan-id interface vfc vfc-id**
4. (Optional) switch(config-vsan)# **no vsan vsan-id interface vfc vfc-id**

### DETAILED STEPS

|               | Command or Action                                                | Purpose                                                                                                                                                                                                             |
|---------------|------------------------------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Step 1</b> | switch# <b>configure terminal</b>                                | Enters configuration mode.                                                                                                                                                                                          |
| <b>Step 2</b> | switch(config)# <b>vsan database</b>                             | Enters VSAN configuration mode.                                                                                                                                                                                     |
| <b>Step 3</b> | switch(config-vsan)# <b>vsan vsan-id interface vfc vfc-id</b>    | Configures the association between the VSAN and virtual Fibre Channel interface.<br><br>The VSAN number must map to a VLAN on the physical Ethernet interface that is bound to the virtual Fibre Channel interface. |
| <b>Step 4</b> | switch(config-vsan)# <b>no vsan vsan-id interface vfc vfc-id</b> | (Optional)<br>Disassociates the connection between the VSAN and virtual Fibre Channel interface.                                                                                                                    |

This example shows how to associate a virtual Fibre Channel interface to a VSAN:

```
switch# configure terminal
switch(config)# vsan database
switch(config-vsan)# vsan 2 interface vfc 4
```

## Verifying the Virtual Interface

To display configuration information about virtual interfaces, perform one of the following tasks:

| Command                                         | Purpose                                                                       |
|-------------------------------------------------|-------------------------------------------------------------------------------|
| switch# <b>show interface vfc</b> <i>vfc-id</i> | Displays the detailed configuration of the specified Fibre Channel interface. |
| switch# <b>show interface brief</b>             | Displays the status of all interfaces.                                        |
| switch# <b>show vlan fcoe</b>                   | Displays the mapping of FCoE VLANs to VSANs.                                  |

This example shows how to display a virtual Fibre Channel interface bound to an Ethernet interface:

```
switch# show interface vfc 3
vfc3 is up
 Bound interface is Ethernet1/37
 Hardware is Virtual Fibre Channel
 Port WWN is 20:02:00:0d:ec:6d:95:3f
 Admin port mode is F, trunk mode is on
 snmp link state traps are enabled
 Port mode is F, FCID is 0x490100
 Port vsan is 931
 1 minute input rate 0 bits/sec, 0 bytes/sec, 0 frames/sec
 1 minute output rate 0 bits/sec, 0 bytes/sec, 0 frames/sec
 0 frames input, 0 bytes
 0 discards, 0 errors
 0 frames output, 0 bytes
 0 discards, 0 errors
 Interface last changed at Thu May 21 04:44:42 2009
```

This example shows how to display a virtual Fibre Channel interface bound to a MAC address:

```
switch# show interface vfc 1001
vfc1001 is down
 Bound MAC is 00:0a:00:00:00:01
 Hardware is Virtual Fibre Channel
 Port WWN is 23:e8:00:0d:ec:6d:95:3f
 Admin port mode is F, trunk mode is on
 snmp link state traps are enabled
 Port vsan is 901
 1 minute input rate 0 bits/sec, 0 bytes/sec, 0 frames/sec
 1 minute output rate 0 bits/sec, 0 bytes/sec, 0 frames/sec
 0 frames input, 0 bytes
 0 discards, 0 errors
 0 frames output, 0 bytes
 0 discards, 0 errors
```

This example shows how to display the status of all the interfaces on the switch (some output has been removed for brevity):

```
switch# show interface brief

Interface Vsan Admin Admin Status SFP Oper Oper Port
 Mode Mode Mode Mode Mode Speed Channel
 (Gbps)

fc3/1 1 auto on trunking swl TE 2 --
fc3/2 1 auto on sfpAbsent -- -- -- --
...
fc3/8 1 auto on sfpAbsent -- -- -- --

Interface Status IP Address Speed MTU Port
 Channel

Ethernet1/1 hwFailure -- -- 1500 --
Ethernet1/2 hwFailure -- -- 1500 --
Ethernet1/3 up -- 10000 1500 --
...
Ethernet1/39 sfpIsAbsen -- -- 1500 --
Ethernet1/40 sfpIsAbsen -- -- 1500 --

Interface Status IP Address Speed MTU

mgmt0 up 172.16.24.41 100 1500

Interface Vsan Admin Admin Status SFP Oper Oper Port
 Mode Mode Mode Mode Mode Speed Channel
 (Gbps)

vfc 1 1 F -- down -- -- -- --
...

```

This example shows how to display the mapping between the VLANs and VSANs on the switch:

```
switch# show vlan fcoe
VLAN VSAN Status

15 15 Operational
20 20 Operational
25 25 Operational
30 30 Non-operational

```

## Mapping VSANs to VLANs Example Configuration

The following example shows how to configure the FCoE VLAN and a virtual Fibre Channel interface:

### SUMMARY STEPS

1. Configure the VLAN on a physical Ethernet address.
2. Create a virtual Fibre Channel interface and bind it to a physical Ethernet interface.
3. Enable the associated VLAN and map the VLAN to a VSAN.
4. Associate the virtual Fibre Channel interface to the VSAN.

## DETAILED STEPS

---

- Step 1** Configure the VLAN on a physical Ethernet address.
- ```
switch# configure terminal
switch(config)# interface ethernet 1/4
switch(config-if)# spanning-tree port type edge trunk
switch(config-if)# switchport mode trunk
switch(config-if)# switchport trunk allowed vlan 1,200
switch(config-if)# exit
```
- Step 2** Create a virtual Fibre Channel interface and bind it to a physical Ethernet interface.
- ```
switch(config)# interface vfc 4
switch(config-if)# bind interface ethernet 1/4
switch(config-if)# exit
```
- Step 3** Enable the associated VLAN and map the VLAN to a VSAN.
- ```
switch(config)# vlan 200
switch(config-vlan)# fcoe vsan 2
switch(config-vlan)# exit
```
- Step 4** Associate the virtual Fibre Channel interface to the VSAN.
- ```
switch(config)# vsan database
switch(config-vsan)# vsan 2 interface vfc 4
switch(config-vsan)# exit
```
-



# PART VI

## Quality of Service

- [Configuring QoS, page 467](#)





## Configuring QoS

This chapter describes how to configure quality of service (QoS) on Cisco Nexus 5000 Series switches. It contains the following sections:

- [Information About QoS, page 467](#)
- [QoS Configuration Guidelines and Limitations, page 475](#)
- [Configuring System Classes, page 476](#)
- [Configuring QoS on Interfaces, page 494](#)
- [Configuring Priority Flow Control and Link-Level Flow Control, page 496](#)
- [Verifying QoS Configuration, page 498](#)
- [Example QoS Configurations, page 504](#)

### Information About QoS

The configurable Cisco NX-OS QoS features on the Cisco Nexus 5000 Series switch allow you to classify the network traffic, police and prioritize the traffic flow, and provide congestion avoidance.

The default QoS configuration on the switch provides lossless service for Fibre Channel and Fibre Channel Over Ethernet (FCoE) traffic and best-effort service for Ethernet traffic. QoS can be configured to provide additional classes of service for Ethernet traffic. Cisco Nexus 5000 Series QoS features are configured using Cisco Modular QoS CLI (MQC).



#### Note

Standard Ethernet is a best-effort medium which means that it lacks any form of flow control. In the event of congestion or collisions, Ethernet will drop packets. The higher level protocols detect the missing data and retransmit the dropped packets.

Fibre Channel requires a reliable transport system that guarantees the delivery of every packet. To properly support FCoE, Ethernet has been enhanced with a priority flow control (PFC) mechanism to prevent congestion.

## MQC

The Cisco Modular QoS CLI (MQC) provides a standard set of commands for configuring QoS.

You can use MQC to define additional traffic classes and to configure QoS policies for the whole system and for individual interfaces. Configuring a QoS policy with MQC consists of the following steps:

- 1 Define traffic classes.
- 2 Associate policies and actions with each traffic class.
- 3 Attach policies to logical or physical interfaces as well as at the global system level.

MQC provides two command types to define traffic classes and policies:

- **class-map**—Defines a class map that represents a class of traffic based on packet-matching criteria. Class maps are referenced in policy maps.  
The class map classifies incoming packets based on matching criteria, such as the IEEE 802.1p CoS value. Unicast and multicast packets are classified.
- **policy-map**—Defines a policy map that represents a set of policies to be applied on a class-by-class basis to class maps.  
The policy map defines a set of actions to take on the associated traffic class, such as limiting the bandwidth or dropping packets.

You define the following class-map and policy-map object types when you create them:

- **network-qos**—Defines MQC objects that you can use for system level related actions.
- **qos**—Defines MQC objects that you can use for classification.
- **queuing**—Defines MQC objects that you can use for queuing and scheduling.



### Note

The **qos** type is the default for the **class-map** and **policy-map** commands, but not for the **service-policy** which requires that you specify an explicit type.

You can attach policies to interfaces or EtherChannels as well as at the global system level by using the **service-policy** command.

You can view all or individual values for MQC objects by using the **show class-map** and **show policy-map** commands.

An MQC target is an entity (such as an Ethernet interface) that represents a flow of packets. A service policy associates a policy map with an MQC target, and specifies whether to apply the policy on incoming or outgoing packets. This mapping enables the configuration of QoS policies such as marking, bandwidth allocation, buffer allocation, and so on.

## System Classes

The system qos is a type of MQC target. You use a service-policy to associate a policy map with the system qos target. A system qos policy applies to all interfaces on the switch unless a specific interface has an overriding service-policy configuration. The system qos policies are used to define system classes, the classes of traffic across the entire switch, and their attributes. To ensure QoS consistency (and for ease of configuration),

the switch distributes the system class parameter values to all its attached network adapters using the Data Center Bridging Exchange (DCBX) protocol.

If service policies are configured at the interface level, the interface-level policy always takes precedence over system class configuration or defaults.

On the Cisco Nexus 5000 Series switch, a system class is uniquely identified by a qos-group value. A total of six system classes are supported. Two of the six system classes are defaults and are always present on the switch. Up to four additional system classes can be created by the administrator.

## Default System Classes

The Cisco Nexus 5000 Series switch provides the following default system classes:

- Drop system class

By default, the software classifies all unicast and multicast Ethernet traffic into the default drop system class. This class is identified by qos-group 0.

This class is created automatically when the system starts up (the class is named class-default in the CLI). You cannot delete this class and you cannot change the match criteria associated with the default class.

- FCoE system class

All Fibre Channel and FCoE control and data traffic is automatically classified into the FCoE system class, which provides no-drop service.

This class is created automatically when the system starts up (the class is named class-fcoe in the CLI). You cannot delete this class, and you can only modify the IEEE 802.1p CoS value to associate with this class. This class is identified by qos-group 1.

The switch classifies packets into the FCoE system class as follows:

- FCoE traffic is classified based on EtherType.
- Native Fibre Channel traffic is classified based on the physical interface type.



**Note**

The optional N5K-M1404 or N5K-M1008 expansion modules provide native 1/2/4-Gigabit Fibre Channel ports.

## Policy Types

The Cisco Nexus 5000 Series switch supports a number of policy types. You create class maps in the policy types.

There are three policy types. The following QoS parameters can be specified for each type of class:

- Type network-qos—A network-qos policy is used to instantiate system classes and associate parameters with those classes that are of system-wide scope.
  - Classification—The traffic that matches this class are as follows:

- QoS Group—A class-map of type network-qos identifies a system-class and is matched by its associated qos-group.
- Policy—The actions that are performed on the matching traffic are as follows:




---

**Note** A network-qos policy can only be attached to the system qos target.

---

- MTU—The MTU that needs to be enforced for the traffic that is mapped to a system class. Each system class has a default MTU and the system class MTU is configurable.
  - Multicast optimization—This configuration specifies if the performance of multicast traffic mapped to this class will be optimized.
  - Pause no-drop—No drop specifies lossless service for the system class. Drop specifies that tail drop is used (arriving packets are dropped when the queue reaches its allocated size) when a queue for this system class is full.  
An additional parameter pfc-cos can be configured. This parameter identifies the class of service (CoS) values to assert priority flow control (PFC) when traffic for a no-drop system class is not mapped based purely on CoS experiences congestion.
  - Queue Limit—This configuration specifies the number of buffers that need to be reserved to the queues of this system class. This option is not configurable for no-drop system classes.
  - Set CoS value—This configuration is used to mark 802.1p values for all traffic mapped to this system class. The marking value for a system class needs to be unique and cannot be used as a marking value for any other system class.
- Type queuing—A type queuing policy is used to define the scheduling characteristics of the queues associated with system classes.




---

**Note** Some configuration parameters when applied to an EtherChannel are not reflected on the configuration of the member ports.

---

- Classification—The traffic that matches this class are as follows:
  - QoS Group—A class-map of type queuing identifies a system-class and is matched by its associated qos-group.
- Policy—The actions that are performed on the matching traffic are as follows:




---

**Note** These policies can be attached to the system qos target or to any interface. The output queuing policy is used to configure output queues on the switch associated with system classes. The input queuing policy is used to configure scheduling for queues in the CNA. The input queuing policy parameters are signalled to the CNA over the DCBX protocol.

---

- Bandwidth—Sets the guaranteed scheduling deficit weighted round robin (DWRR) percentage for the system class.
- Priority—Sets a system class for strict-priority scheduling. Only one system class can be configured for priority in a given queuing policy.
- Type qos—A type qos policy is used to classify traffic that is based on various Layer 2, Layer 3, and Layer 4 fields in the frame and to map it to system classes.




---

**Note** Some configuration parameters when applied to an EtherChannel are not reflected on the configuration of the member ports.

---

- Classification—The traffic that matches this class are as follows:
  - Access Control Lists—Classifies traffic based on the criteria in existing ACLs.
  - Class of Service—Matches traffic based on the CoS field in the frame header.
  - DSCP—Classifies traffic based on the Differentiated Services Code Point (DSCP) value in the DiffServ field of the IP header.
  - IP Real Time Protocol—Classifies traffic on the port numbers used by real-time applications.
  - Precedence—Classifies traffic based on the precedence value in the type of service (ToS) field of the IP header.
  - Protocol—Classifies traffic based on the protocol field of the IP header.
- Policy—The actions that are performed on the matching traffic are as follows:




---

**Note** This policy can be attached to the system or to any interface. It applies to input traffic only.

---

- QoS Group—Sets the qos-group corresponding to the system class this traffic flow is mapped to.

## Link-Level Flow Control

IEEE 802.3x link-level flow control allows a congested receiver to communicate a transmitter at the other end of the link to pause its data transmission for a short period of time. The link-level flow control feature applies to all the traffic on the link.

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

On the Cisco Nexus 5000 Series switch, Ethernet interfaces do not automatically detect the link-level flow control capability. You must configure the capability explicitly on the Ethernet interfaces.

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

## Priority Flow Control

Priority flow control (PFC) allows you to apply pause functionality to specific classes of traffic on a link instead of all the traffic on the link. PFC applies pause functionality based on the IEEE 802.1p CoS value. When the switch enables PFC, it communicates to the adapter which CoS values to apply the pause.

Ethernet interfaces use PFC to provide lossless service to no-drop system classes. PFC implements pause frames on a per-class basis and uses the IEEE 802.1p CoS value to identify the classes that require lossless service.

In the switch, each system class has an associated IEEE 802.1p CoS value that is assigned by default or configured on the system class. If you enable PFC, the switch sends the no-drop CoS values to the adapter, which then applies PFC to these CoS values.

The default CoS value for the FCoE system class is 3. This value is configurable.

By default, the switch 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.

If you do not enable PFC on an interface, you can enable IEEE 802.3X link-level pause. By default, link-level pause is disabled.

## MTU

The Cisco Nexus 5000 Series switch is a Layer 2 switch, and it does not support packet fragmentation. A maximum transmission unit (MTU) configuration mismatch between ingress and egress interfaces may result in packets being truncated.

When configuring MTU, follow these guidelines:

- MTU is specified per system class. The system class allows a different MTU for each class of traffic but they must be consistent on all ports across the entire switch. You cannot configure MTU on the interfaces.
- Fibre Channel and FCoE payload MTU is 2112 bytes across the switch. As a result, the rxbufsize for Fibre Channel interfaces is fixed at 2112 bytes. If the Cisco Nexus 5000 Series switch receives an rxbufsize from a peer that is different than 2112 bytes, it will fail the exchange of link parameters (ELP) negotiation and not bring the link up.
- Enter the **system jumbomtu** command to define the upper bound of any MTU in the system. The system jumbo MTU has a default value of 9216 bytes. The minimum MTU is 2240 bytes and the maximum MTU is 9216 bytes.
- The system class MTU sets the MTU for all packets in the class. The system class MTU cannot be configured larger than the global jumbo MTU.
- The FCoE system class (for Fibre Channel and FCoE traffic) has a default MTU of 2240 bytes. This value cannot be modified.
- The default drop system class has a default MTU of 1538 bytes. You can configure this value.
- The switch sends the MTU configuration to network adapters that support DCBX.



---

**Note** MTU is not supported in Converged Enhanced Ethernet (CEE) mode for DCBX.

---

## Trust Boundaries

The trust boundary is enforced by the incoming interface as follows:

- All Fibre Channel and virtual Fibre Channel interfaces are automatically classified into the FCoE system class.
- By default, all Ethernet interfaces are trusted interfaces. A packet tagged with an 802.1p CoS value is classified into a system class using the value in the packet.
- Any packet that is not tagged with an 802.1p CoS value is classified into the default drop system class. If the untagged packet is sent over a trunk, it is tagged with the default untagged CoS value, which is zero.
- You can override the default untagged CoS value for an Ethernet interface or port channel.

After the system applies the untagged CoS value, QoS functions the same as for a packet that entered the system tagged with the CoS value.

## Ingress Queuing Policies

You can associate an ingress policy map with an Ethernet interface to guarantee bandwidth for the specified traffic class or to specify a priority queue.

The ingress policy is applied in the adapter to all outgoing traffic that matches the specified CoS value.

When you configure an ingress policy for an interface, the switch sends the configuration data to the adapter. If the adapter does not support the DCBX protocol or the ingress policy type-length-value (TLV), the ingress policy configuration is ignored.

## Ingress Classification Policies

You use classification to partition traffic into classes. You classify the traffic based on the port characteristics (CoS field) or the packet header fields that include IP precedence, Differentiated Services Code Point (DSCP), and Layer 2 to Layer 4 parameters. The values used to classify traffic are called match criteria. When you define a traffic class, you can specify multiple match criteria, you can choose to not match on a particular criterion, or you can determine traffic class by matching any or all criteria.

Traffic that fails to match any class is assigned to a default class of traffic called class-default.

## Egress Queuing Policies

You can associate an egress policy map with an Ethernet interface to guarantee the bandwidth for the specified traffic class or to configure the egress queues.

The bandwidth allocation limit applies to all traffic on the interface including any FCoE traffic.

Each Ethernet interface supports up to six queues, one for each system class. The queues have the following default configuration:

- In addition to the six queues, control traffic that is destined for the CPU uses strict priority queues. These queues are not accessible for user configuration.
- FCoE traffic (traffic that maps to the FCoE system class) is assigned a queue. This queue uses weighted round-robin (WRR) scheduling with 50 percent of the bandwidth.
- Standard Ethernet traffic in the default drop system class is assigned a queue. This queue uses WRR scheduling with 50 percent of the bandwidth.

If you add a system class, a queue is assigned to the class. You must reconfigure the bandwidth allocation on all affected interfaces. Bandwidth is not dedicated automatically to user-defined system classes.

You can configure a strict priority queue. This queue is serviced before all other queues except the control traffic queue (which carries control rather than data traffic).

## QoS for Multicast Traffic

The system provides six multicast queues per interface and allocates one queue for each system class. By default, all multicast Ethernet traffic is classified into the default drop system class. This traffic is serviced by one multicast queue.

Optimized multicasting allows use of the unused multicast queues to achieve better throughput for multicast frames. If optimized multicast is enabled for the default drop system class, the system will use all six queues to service the multicast traffic (all six queues are given equal priority).

If you define a new system class, a dedicated multicast queue is assigned for this class. This queue is removed from the set of queues available for the optimized multicast class.

Optimized multicasting achieves better throughput and improves performance for multicast frames.

The system provides two predefined class maps for matching broadcast or multicast traffic. These class maps are convenient for creating separate policy maps for unicast and multicast traffic. The predefined class maps are as follows:

- class-all-flood

The class-all-flood class map matches all broadcast, multicast, and unknown unicast traffic (across all CoS values). If you configure a policy map with the class-all-flood class map, the system automatically uses all available multicast queues for this traffic.

- class-ip-multicast

The class-ip-multicast class map matches all IP multicast traffic. Policy options configured in this class map apply to traffic across all Ethernet CoS values. For example, if you enable optimized multicast for this class, the IP multicast traffic for all CoS values is optimized.



### Note

If you configure either of these predefined class maps as a no-drop class, the priority flow control capability is applied across all Ethernet CoS values. In this configuration, pause will be applied to unicast and multicast traffic.

## Policy for Fibre Channel Interfaces

The egress queues are not configurable for native Fibre Channel interfaces. Two queues are available as follows:

- A strict priority queue to serve high-priority control traffic.
- A queue to serve all data traffic and low-priority control traffic.

## QoS for Traffic Directed to the CPU

The switch automatically applies QoS policies to traffic that is directed to the CPU to ensure that the CPU is not flooded with packets. Control traffic, such as BPDU frames, is given higher priority to ensure delivery.

## QoS Configuration Guidelines and Limitations

Switch resources (such as buffers, virtual output queues, and egress queues) are partitioned based on the default and user-defined system classes. Cisco NX-OS automatically adjusts the resource allocation to accommodate the configured system classes.

To maintain optimal switch performance, follow these guidelines when configuring system classes and policies:

- If less than four Ethernet classes are defined, up to two of these classes can be configured as no-drop classes. If more than three Ethernet classes are defined, only one of these classes can be configured as a no-drop class. The default drop class is counted as an Ethernet class.
- If priority flow control is enabled on an Ethernet interface, pause will never be applied to traffic with a drop system class. PFC does not apply pause to drop classes and the link-level pause feature is never enabled on an interface with PFC.
- All FCoE traffic on an Ethernet interface is mapped to one no-drop system class. By default, this class is associated with CoS value 3, although you can configure a different value. If you configure standard Ethernet traffic to use the same CoS value as FCoE, this traffic is still mapped to the FCoE system class and the switch will apply priority flow control on the FCoE CoS value.
- When a Cisco Nexus 2148T Fabric Extender is connected and in use, data traffic should never be marked with a CoS value of 7. CoS 7 is reserved for control traffic transiting the Fabric Extender.



### Note

Type qos policies can be activated only on Cisco Nexus 5000 Series interfaces and Cisco Nexus 2000 Series Fabric Extender interfaces. Type qos policies on Fabric Extender fabric interfaces or Fabric Extender fabric EtherChannel interfaces are ineffective, though the Cisco NX-OS CLI does not reject the configuration.

We recommend that you do not configure type qos policy-maps on Fabric Extender fabric interfaces or Fabric Extender fabric EtherChannel interfaces to avoid wasting hardware resources.

When configuring EtherChannels, note the following guidelines:

- The service policy configured on an EtherChannel applies to all member interfaces.
- The priority flow control configured on an EtherChannel applies to all member interfaces.

# Configuring System Classes

## Configuring Class Maps

You can create or modify a class map with the **class-map** command. The class map is a named object that represents a class of traffic. In the class map, you specify a set of match criteria for classifying the packets. You can then reference class maps in policy maps.

### SUMMARY STEPS

1. switch# **configure terminal**
2. switch(config)# **class-map** [type {network-qos | qos | queuing}] *class-name*
3. (Optional) switch(config)# **no class-map** [type {network-qos | queuing | qos | queuing}] *class-name*

### DETAILED STEPS

|               | Command or Action                                                                                    | Purpose                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
|---------------|------------------------------------------------------------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Step 1</b> | switch# <b>configure terminal</b>                                                                    | Enters configuration mode.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Step 2</b> | switch(config)# <b>class-map</b> [type {network-qos   qos   queuing}] <i>class-name</i>              | <p>Creates or accesses a named object that represents the specified class of traffic. Class-map names can contain alphabetic, hyphen, or underscore characters, are case sensitive, and can be up to 40 characters.</p> <p>The three class-map configuration modes are as follows:</p> <ul style="list-style-type: none"> <li>• network-qos—Network-wide (global) mode. CLI prompt: switch(config-cmap-nq)#</li> <li>• qos—Classification mode; this is the default mode. CLI prompt: switch(config-cmap-qos)#</li> <li>• queuing—Queuing mode. CLI prompt: switch(config-cmap-que)#</li> </ul> |
| <b>Step 3</b> | switch(config)# <b>no class-map</b> [type {network-qos   queuing   qos   queuing}] <i>class-name</i> | <p>(Optional)<br/>Deletes the specified class map.</p> <p><b>Note</b> You cannot delete the two system-defined class maps: class-fcoe and class-default.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                    |

## Configuring ACL Classification

You can classify traffic by matching packets based on an existing access control list (ACL). Traffic is classified by the criteria defined in the ACL. The **permit** and **deny** ACL keywords are ignored in the matching; even if a match criteria in the access-list has a **deny** action, it is still used for matching for this class.

## SUMMARY STEPS

1. switch# **configure terminal**
2. switch(config)# **class-map type qos class-name**
3. switch(config-cmap-qos)# **match access-group name acl-name**
4. (Optional) switch(config-cmap-qos)# **no match access-group name acl-name**

## DETAILED STEPS

|               | Command or Action                                                   | Purpose                                                                                                                                                                                                                                                                                                                           |
|---------------|---------------------------------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Step 1</b> | switch# <b>configure terminal</b>                                   | Enters configuration mode.                                                                                                                                                                                                                                                                                                        |
| <b>Step 2</b> | switch(config)# <b>class-map type qos class-name</b>                | Creates a named object that represents a class of traffic. Class-map names can contain alphabetic, hyphen, or underscore characters, are case sensitive, and can be up to 40 characters.                                                                                                                                          |
| <b>Step 3</b> | switch(config-cmap-qos)# <b>match access-group name acl-name</b>    | Configures a traffic class by matching packets based on the <i>acl-name</i> . The <b>permit</b> and <b>deny</b> ACL keywords are ignored in the matching.<br><br><b>Note</b> You can only define a single ACL in a class map.<br><br>You cannot add any other match criteria to a class with a <b>match access-group</b> defined. |
| <b>Step 4</b> | switch(config-cmap-qos)# <b>no match access-group name acl-name</b> | (Optional)<br>Removes the match from the traffic class.                                                                                                                                                                                                                                                                           |

This example shows how to classify traffic by matching packets based on existing ACLs:

```
switch# configure terminal
switch(config)# class-map type qos class_acl
switch(config-cmap-qos)# match access-group name acl-01
```

Use the **show class-map** command to display the ACL class-map configuration:

```
switch# show class-map class_acl
```

## Configuring CoS Classification

You can classify traffic based on the class of service (CoS) in the IEEE 802.1Q header. This 3-bit field is defined in IEEE 802.1p to support QoS traffic classes. CoS is encoded in the high order 3 bits of the VLAN ID Tag field and is referred to as *user\_priority*.

If a system class is configured with a no-drop function, the **match cos** command serves an additional purpose. The switch sends the CoS value to the adapter so that the adapter will apply a PFC pause for this CoS value.

The FCoE system class has a default CoS value of 3. You can add a **match cos** configuration to the FCoE system class to set a different CoS value. A PFC pause will be applied to traffic that matches the new value.

## SUMMARY STEPS

1. switch# **configure terminal**
2. switch(config)# **class-map type qos class-name**
3. switch(config-cmap-qos)# **match cos cos-value**
4. (Optional) switch(config-cmap-qos)# **no match cos cos-value**

## DETAILED STEPS

|               | Command or Action                                      | Purpose                                                                                                                                                                                                                                                                                                                                               |
|---------------|--------------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Step 1</b> | switch# <b>configure terminal</b>                      | Enters configuration mode.                                                                                                                                                                                                                                                                                                                            |
| <b>Step 2</b> | switch(config)# <b>class-map type qos class-name</b>   | Creates a named object that represents a class of traffic. Class-map names can contain alphabetic, hyphen, or underscore characters, are case sensitive, and can be up to 40 characters.                                                                                                                                                              |
| <b>Step 3</b> | switch(config-cmap-qos)# <b>match cos cos-value</b>    | Specifies the CoS value to match for classifying packets into this class. You can configure a CoS value in the range of 0 to 7.<br><br><b>Note</b> When a Cisco Nexus 2148T Fabric Extender is connected and in use, data traffic should never be marked with a CoS value of 7. CoS 7 is reserved for control traffic transiting the Fabric Extender. |
| <b>Step 4</b> | switch(config-cmap-qos)# <b>no match cos cos-value</b> | (Optional)<br>Removes the match from the traffic class.                                                                                                                                                                                                                                                                                               |

This example shows how to classify traffic by matching packets based on a defined CoS value:

```
switch# configure terminal
switch(config)# class-map type qos class_cos
switch(config-cmap-qos)# match cos 4, 5-6
```

Use the **show class-map** command to display the CoS value class-map configuration:

```
switch# show class-map class_cos
```

## Configuring DSCP Classification

You can classify traffic based on the Differentiated Services Code Point (DSCP) value in the DiffServ field of the IP header (either IPv4 or IPv6). The standard DSCP values are listed in the following table:

**Table 56: Standard DSCP Values**

| Value | List of DSCP Values                 |
|-------|-------------------------------------|
| af11  | AF11 dscp (001010)—decimal value 10 |
| af12  | AF12 dscp (001100)—decimal value 12 |

| <b>Value</b> | <b>List of DSCP Values</b>                        |
|--------------|---------------------------------------------------|
| af13         | AF13 dscp (001110)—decimal value 14               |
| af21         | AF21 dscp (010010)—decimal value 18               |
| af22         | AF22 dscp (010100)—decimal value 20               |
| af23         | AF23 dscp (010110)—decimal value 22               |
| af31         | AF31 dscp (011010)—decimal value 26               |
| af32         | AF40 dscp (011100)—decimal value 28               |
| af33         | AF33 dscp (011110)—decimal value 30               |
| af41         | AF41 dscp (100010)—decimal value 34               |
| af42         | AF42 dscp (100100)—decimal value 36               |
| af43         | AF43 dscp (100110)—decimal value 38               |
| cs1          | CS1 (precedence 1) dscp (001000)—decimal value 8  |
| cs2          | CS2 (precedence 2) dscp (010000)—decimal value 16 |
| cs3          | CS3 (precedence 3) dscp (011000)—decimal value 24 |
| cs4          | CS4 (precedence 4) dscp (100000)—decimal value 32 |
| cs5          | CS5 (precedence 5) dscp (101000)—decimal value 40 |
| cs6          | CS6 (precedence 6) dscp (110000)—decimal value 48 |
| cs7          | CS7 (precedence 7) dscp (111000)—decimal value 56 |
| default      | Default dscp (000000)—decimal value 0             |
| ef           | EF dscp (101110)—decimal value 46                 |

**SUMMARY STEPS**

1. switch# **configure terminal**
2. switch(config)# **class-map type qos class-name**
3. switch(config-cmap-qos)# **match dscp dscp-list**
4. (Optional) switch(config-cmap-qos)# **no match dscp dscp-list**

**DETAILED STEPS**

|               | <b>Command or Action</b>                                | <b>Purpose</b>                                                                                                                                                                           |
|---------------|---------------------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Step 1</b> | switch# <b>configure terminal</b>                       | Enters configuration mode.                                                                                                                                                               |
| <b>Step 2</b> | switch(config)# <b>class-map type qos class-name</b>    | Creates a named object that represents a class of traffic. Class-map names can contain alphabetic, hyphen, or underscore characters, are case sensitive, and can be up to 40 characters. |
| <b>Step 3</b> | switch(config-cmap-qos)# <b>match dscp dscp-list</b>    | Configures the traffic class by matching packets based on the values in the <i>dscp-list</i> . The standard DSCP values are shown in the preceding table.                                |
| <b>Step 4</b> | switch(config-cmap-qos)# <b>no match dscp dscp-list</b> | (Optional)<br>Removes the match from the traffic class.                                                                                                                                  |

This example shows how to classify traffic by matching packets based on the DSCP value in the DiffServ field of the IP header:

```
switch# configure terminal
switch(config)# class-map type qos class dscp
switch(config-cmap-qos)# match dscp af21, af32
```

Use the **show class-map** command to display the DSCP class-map configuration:

```
switch# show class-map class_dscp
```

**Configuring IP RTP Classification**

The IP Real-time Transport Protocol (RTP) is a transport protocol for real-time applications that transmits data such as audio or video and is defined by RFC 3550. Although RTP does not use a common TCP or UDP port, you typically configure RTP to use ports 16384 to 32767. UDP communications use an even port and the next higher odd port is used for RTP Control Protocol (RTCP) communications.

You can classify based on UDP port ranges, which are likely to target applications using RTP.

## SUMMARY STEPS

1. switch# **configure terminal**
2. switch(config)# **class-map type qos class-name**
3. switch(config-cmap-qos)# **match ip rtp port-number**
4. (Optional) switch(config-cmap-qos)# **no match ip rtp port-number**

## DETAILED STEPS

|               | Command or Action                                           | Purpose                                                                                                                                                                                       |
|---------------|-------------------------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Step 1</b> | switch# <b>configure terminal</b>                           | Enters configuration mode.                                                                                                                                                                    |
| <b>Step 2</b> | switch(config)# <b>class-map type qos class-name</b>        | Creates a named object that represents a class of traffic. Class-map names can contain alphabetic, hyphen, or underscore characters, are case sensitive, and can be up to 40 characters.      |
| <b>Step 3</b> | switch(config-cmap-qos)# <b>match ip rtp port-number</b>    | Configures the traffic class by matching packets based on a range of lower and upper UDP port numbers, which is likely to target applications using RTP. Values can range from 2000 to 65535. |
| <b>Step 4</b> | switch(config-cmap-qos)# <b>no match ip rtp port-number</b> | (Optional)<br>Removes the match from the traffic class.                                                                                                                                       |

This example shows how to classify traffic by matching packets based on UDP port ranges that are typically used by RTP applications:

```
switch# configure terminal
switch(config)# class-map type qos class_rtp
switch(config-cmap-qos)# match ip rtp 2000-2100, 4000-4100
```

Use the **show class-map** command to display the RTP class-map configuration:

```
switch# show class-map class_rtp
```

## Configuring Precedence Classification

You can classify traffic based on the precedence value in the type of service (ToS) byte field of the IP header (either IPv4 or IPv6). The following table shows the precedence values:

**Table 57: Precedence Values**

| Value    | List of Precedence Values |
|----------|---------------------------|
| <0-7>    | IP precedence value       |
| critical | Critical precedence (5)   |
| flash    | Flash precedence (3)      |

| Value          | List of Precedence Values           |
|----------------|-------------------------------------|
| flash-override | Flash override precedence (4)       |
| immediate      | Immediate precedence (2)            |
| internet       | Internetwork control precedence (6) |
| network        | Network control precedence (7)      |
| priority       | Priority precedence (1)             |
| routine        | Routine precedence (0)              |

## SUMMARY STEPS

1. switch# **configure terminal**
2. switch(config)# **class-map type qos class-name**
3. switch(config-cmap-qos)# **match precedence precedence-values**
4. (Optional) switch(config-cmap-qos)# **no match precedence precedence-values**

## DETAILED STEPS

|               | Command or Action                                                     | Purpose                                                                                                                                                                                  |
|---------------|-----------------------------------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Step 1</b> | switch# <b>configure terminal</b>                                     | Enters configuration mode.                                                                                                                                                               |
| <b>Step 2</b> | switch(config)# <b>class-map type qos class-name</b>                  | Creates a named object that represents a class of traffic. Class-map names can contain alphabetic, hyphen, or underscore characters, are case sensitive, and can be up to 40 characters. |
| <b>Step 3</b> | switch(config-cmap-qos)# <b>match precedence precedence-values</b>    | Configures the traffic class by matching packets based on <i>precedence-values</i> . Values are shown in the preceding table.                                                            |
| <b>Step 4</b> | switch(config-cmap-qos)# <b>no match precedence precedence-values</b> | (Optional)<br>Removes the match from the traffic class.                                                                                                                                  |

This example shows how to classify traffic by matching packets based on the precedence value in the ToS byte field of the IP header:

```
switch# configure terminal
switch(config)# class-map type qos class_precedence
switch(config-cmap-qos)# match precedence 1-2, critical
```

Use the **show class-map** command to display the IP precedence value class-map configuration:

```
switch# show class-map class_precedence
```

## Configuring Protocol Classification

You can classify traffic based on the protocol field in the IP header. The following table shows the protocol arguments:

**Table 58: Protocol Arguments**

| Argument | Description                               |
|----------|-------------------------------------------|
| arp      | Address Resolution Protocol (ARP)         |
| clns_es  | CLNS End Systems                          |
| clns_is  | CLNS Intermediate System                  |
| dhcp     | Dynamic Host Configuration (DHCP)         |
| ldp      | Label Distribution Protocol (LDP)         |
| netbios  | NetBIOS Extended User Interface (NetBEUI) |

### SUMMARY STEPS

1. switch# **configure terminal**
2. switch(config)# **class-map type qos class-name**
3. switch(config-cmap-qos)# **match protocol {arp | clns\_es | clns\_is | dhcp | ldp | netbios}**
4. (Optional) switch(config-cmap-qos)# **no match protocol {arp | clns\_es | clns\_is | dhcp | ldp | netbios}**

### DETAILED STEPS

|               | Command or Action                                                                                  | Purpose                                                                                                                                                                                  |
|---------------|----------------------------------------------------------------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Step 1</b> | switch# <b>configure terminal</b>                                                                  | Enters configuration mode.                                                                                                                                                               |
| <b>Step 2</b> | switch(config)# <b>class-map type qos class-name</b>                                               | Creates a named object that represents a class of traffic. Class-map names can contain alphabetic, hyphen, or underscore characters, are case sensitive, and can be up to 40 characters. |
| <b>Step 3</b> | switch(config-cmap-qos)# <b>match protocol {arp   clns_es   clns_is   dhcp   ldp   netbios}</b>    | Configures the traffic class by matching packets based on the specified protocol.                                                                                                        |
| <b>Step 4</b> | switch(config-cmap-qos)# <b>no match protocol {arp   clns_es   clns_is   dhcp   ldp   netbios}</b> | (Optional)<br>Removes the match from the traffic class.                                                                                                                                  |

This example shows how to classify traffic by matching packets based on the protocol field:

```
switch# configure terminal
switch(config)# class-map type qos class_protocol
switch(config-cmap-qos)# match protocol arp
```

Use the **show class-map** command to display the protocol class-map configuration:

```
switch# show class-map class_protocol
```

## Configuring QoS Group Classification

You can classify traffic based on the value of the QoS group internal label, that represents a system class. You can set the value of the QoS group within a policy map using the **set qos-group** command.

### SUMMARY STEPS

1. switch# **configure terminal**
2. switch(config)# **class-map type {network-qos | queuing} class-name**
3. switch(config-cmap-que)# **match qos-group qos-group-value**
4. (Optional) switch(config-cmap-que)# **no match qos-group qos-group-value**

### DETAILED STEPS

|               | Command or Action                                                        | Purpose                                                                                                                                                                                                                                                                                                      |
|---------------|--------------------------------------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Step 1</b> | switch# <b>configure terminal</b>                                        | Enters configuration mode.                                                                                                                                                                                                                                                                                   |
| <b>Step 2</b> | switch(config)# <b>class-map type {network-qos   queuing} class-name</b> | Creates a named object that represents a class of traffic. Class-map names can contain alphabetic, hyphen, or underscore characters, are case sensitive, and can be up to 40 characters.                                                                                                                     |
| <b>Step 3</b> | switch(config-cmap-que)# <b>match qos-group qos-group-value</b>          | Configures the traffic class by matching packets based on a list of QoS group values. Values can range from 0 to 5. QoS group 0 is equivalent to class-default and QoS group 1 is equivalent to class-fcoe.<br><br><b>Note</b> qos-groups 0 and 1 are reserved for default classes and cannot be configured. |
| <b>Step 4</b> | switch(config-cmap-que)# <b>no match qos-group qos-group-value</b>       | (Optional)<br>Removes the match from the traffic class.                                                                                                                                                                                                                                                      |

This example shows how to classify traffic based on the value of the QoS group:

```
switch# configure terminal
switch(config)# class-map type queuing class_qos_group
switch(config-cmap-que)# match qos-group 4
```

Use the **show class-map** command to display the QoS group class-map configuration:

```
switch# show class-map class_qos_group
```

## Configuring Policy Maps

The **policy-map** command is used to create a named object that represents a set of policies that are to be applied to a set of traffic classes.

The switch provides two default system classes: a no-drop class for lossless service (**class-foe**) and a drop class for best-effort service (**class-default**). You can define up to four additional system classes for Ethernet traffic.

The following predefined policy maps are used as default service policies:

- network-qos: default-nq-policy
- Input qos: default-in-policy
- Input queuing: default-in-policy
- Output queuing: default-out-policy

You need to create a policy map to specify the policies for any user-defined class. In the policy map, you can configure the QoS parameters for each class. You can use the same policy map to modify the configuration of the default classes.

The switch distributes all the policy-map configuration values to the attached network adapters.

### Before You Begin

Before creating the policy map, define a class map for each new system class.

## SUMMARY STEPS

1. switch# **configure terminal**
2. switch(config)# **policy-map** [**type** {**network-qos** | **qos** | **queuing**}] *policy-name*
3. (Optional) switch(config)# **no policy-map** [**type** {**network-qos** | **qos** | **queuing**}] *policy-name*
4. switch(config-pmap)# **class** [**type** {**network-qos** | **qos** | **queuing**}] *class-name*
5. (Optional) switch(config-pmap)# **no class** [**type** {**network-qos** | **qos** | **queuing**}] *class-name*

## DETAILED STEPS

|               | Command or Action                                                                                                           | Purpose                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
|---------------|-----------------------------------------------------------------------------------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Step 1</b> | switch# <b>configure terminal</b>                                                                                           | Enters configuration mode.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| <b>Step 2</b> | switch(config)# <b>policy-map</b> [ <b>type</b> { <b>network-qos</b>   <b>qos</b>   <b>queuing</b> }]<br><i>policy-name</i> | Creates a named object representing a set of policies that are to be applied to a set of traffic classes. Policy-map names can contain alphabetic, hyphen, or underscore characters, are case sensitive, and can be up to 40 characters.<br><br>The three policy-map configuration modes are as follows: <ul style="list-style-type: none"> <li>• <b>network-qos</b>—Network-wide (global) mode. CLI prompt:<br/>switch(config-pmap-nq)#</li> <li>• <b>qos</b>—Classification mode; this is the default mode. CLI prompt:<br/>switch(config-pmap-qos)#</li> </ul> |

|               | Command or Action                                                                                                      | Purpose                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
|---------------|------------------------------------------------------------------------------------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
|               |                                                                                                                        | <ul style="list-style-type: none"> <li>queuing—Queuing mode. CLI prompt: switch(config-pmap-que)#</li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| <b>Step 3</b> | switch(config)# <b>no policy-map</b> [type { <b>network-qos</b>   <b>qos</b>   <b>queuing</b> }]<br><i>policy-name</i> | (Optional)<br>Deletes the specified policy map.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>Step 4</b> | switch(config-pmap)# <b>class</b> [type { <b>network-qos</b>   <b>qos</b>   <b>queuing</b> }]<br><i>class-name</i>     | <p>Associates a class map with the policy map, and enters configuration mode for the specified system class. The three class-map configuration modes are as follows:</p> <ul style="list-style-type: none"> <li>network-qos—Network-wide (global) mode. CLI prompt: switch(config-pmap-c-nq)#</li> <li>qos—Classification mode; this is the default mode. CLI prompt: switch(config-pmap-c-qos)#</li> <li>queuing—Queuing mode. CLI prompt: switch(config-pmap-c-que)#</li> </ul> <p><b>Note</b> The associated class map must be the same type as the policy-map type.</p> |
| <b>Step 5</b> | switch(config-pmap)# <b>no class</b> [type { <b>network-qos</b>   <b>qos</b>   <b>queuing</b> }]<br><i>class-name</i>  | (Optional)<br>Deletes the class map association.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |

### Configuring Type Network QoS Policies

Type network qos policies can only be configured on the system qos attachment point. They are applied to the entire switch for a particular class.

#### SUMMARY STEPS

1. switch# **configure terminal**
2. switch(config)# **policy-map type network-qos** *policy-name*
3. switch(config-pmap-nq)# **class type network-qos** *class-name*
4. switch(config-pmap-c-nq)# **mtu** *mtu-value*
5. (Optional) switch(config-pmap-c-nq)# **no mtu**
6. switch(config-pmap-c-nq)# **multicast-optimize**
7. (Optional) switch(config-pmap-c-nq)# **no multicast-optimize**
8. switch(config-pmap-c)# **pause no-drop** [**pfc-cos** *pfc-cos-value*]
9. (Optional) switch(config-pmap-c-nq)# **no pause no-drop**
10. switch(config-pmap-c-nq)# **queue-limit** *number-bytes* **bytes**
11. (Optional) switch(config-pmap-c-nq)# **no queue-limit** *number-bytes* **bytes**
12. switch(config-pmap-c-nq)# **set cos** *cos-value*
13. (Optional) switch(config-pmap-c-nq)# **no set cos** *cos-value*

## DETAILED STEPS

|         | Command or Action                                                             | Purpose                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
|---------|-------------------------------------------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Step 1  | switch# <b>configure terminal</b>                                             | Enters configuration mode.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| Step 2  | switch(config)# <b>policy-map type network-qos</b> <i>policy-name</i>         | Creates a named object that represents a set of policies that are to be applied to a set of traffic classes. Policy-map names can contain alphabetic, hyphen, or underscore characters, are case sensitive, and can be up to 40 characters.                                                                                                                                                                                                                                                                                                                                                                                                  |
| Step 3  | switch(config-pmap-nq)# <b>class type network-qos</b> <i>class-name</i>       | Associates a class map with the policy map, and enters configuration mode for the specified system class.<br><b>Note</b> The associated class map must be the same type as the policy map type.                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| Step 4  | switch(config-pmap-c-nq)# <b>mtu</b> <i>mtu-value</i>                         | Specifies the MTU value in bytes.<br><b>Note</b> The <i>mtu-value</i> that you configure must be less than the value set by the <b>system jumbomtu</b> command.                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| Step 5  | switch(config-pmap-c-nq)# <b>no mtu</b>                                       | (Optional)<br>Resets the MTU value in this class.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| Step 6  | switch(config-pmap-c-nq)# <b>multicast-optimize</b>                           | Enables multicast optimization. Multicast traffic in this class will be served by all available multicast queues.<br><b>Note</b> Only one class in a policy map can be configured for multicast optimization.                                                                                                                                                                                                                                                                                                                                                                                                                                |
| Step 7  | switch(config-pmap-c-nq)# <b>no multicast-optimize</b>                        | (Optional)<br>Disables multicast optimization.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| Step 8  | switch(config-pmap-c)# <b>pause no-drop</b> [ <i>pfc-cos pfc-cos-value</i> ]  | Configures a no-drop class. If you do not specify this command, the default policy is drop.<br><b>Note</b> The operation for the drop policy is a simple tail drop, where arriving packets will be dropped if the queue increases to its allocated size.<br>The <i>pfc-cos-value</i> range is from 0 to 7. This option is supported only for for a ACL-based system class (which filters traffic using criteria other than cos-based matches).<br><b>Caution</b> The list of CoS values can potentially include the CoS value that is used for FCoE traffic in class-fcoe. You must determine if this is desired behavior for your topology. |
| Step 9  | switch(config-pmap-c-nq)# <b>no pause no-drop</b>                             | (Optional)<br>Removes the no-drop option from this class.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| Step 10 | switch(config-pmap-c-nq)# <b>queue-limit</b> <i>number-bytes</i> <b>bytes</b> | Specifies the tail drop threshold on this interface. The threshold range is from 20480 to 204800 bytes.<br><b>Note</b> The queue limit can only be configured on drop classes. If you try to configure a queue limit on a no-drop class, or try to configure no-drop on a class where a queue limit is already defined, the CLI will return an error.                                                                                                                                                                                                                                                                                        |

|                | Command or Action                                                        | Purpose                                                                                                       |
|----------------|--------------------------------------------------------------------------|---------------------------------------------------------------------------------------------------------------|
| <b>Step 11</b> | <code>switch(config-pmap-c-nq)# no queue-limit number-bytes bytes</code> | (Optional)<br>Disables the queue limit specification in this class.                                           |
| <b>Step 12</b> | <code>switch(config-pmap-c-nq)# set cos cos-value</code>                 | Specifies a 802.1Q CoS value which is used to mark packets on this interface. The value range is from 0 to 7. |
| <b>Step 13</b> | <code>switch(config-pmap-c-nq)# no set cos cos-value</code>              | (Optional)<br>Disables the marking operation in this class.                                                   |

This example shows how to define a type network-qos policy map:

```
switch# configure terminal
switch(config)# policy-map type network-qos policy-que1
switch(config-pmap-nq)# class type network-qos class-que1
switch(config-pmap-c-nq)# mtu 5000
switch(config-pmap-c-nq)# pause no-drop pfc-cos 1-3,5
switch(config-pmap-c-nq)# set cos 4
```

## Configuring Type Queuing Policies

Type queuing policies are used for scheduling and buffering the traffic of a specific system class. A type queuing policy is identified by its qos-group and can be attached to the system or to individual interfaces (except for Fabric Extender host interfaces) for input or output traffic.

### SUMMARY STEPS

1. `switch# configure terminal`
2. `switch(config)# policy-map type queuing policy-name`
3. `switch(config-pmap-que)# class type queuing class-name`
4. `switch(config-pmap-c-que)# bandwidth percent percentage`
5. (Optional) `switch(config-pmap-c-que)# no bandwidth percent percentage`
6. `switch(config-pmap-c-que)# priority`
7. (Optional) `switch(config-pmap-c-que)# no priority`

### DETAILED STEPS

|               | Command or Action                                                | Purpose                                                                                                                                                                                                                                     |
|---------------|------------------------------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Step 1</b> | <code>switch# configure terminal</code>                          | Enters configuration mode.                                                                                                                                                                                                                  |
| <b>Step 2</b> | <code>switch(config)# policy-map type queuing policy-name</code> | Creates a named object that represents a set of policies that are to be applied to a set of traffic classes. Policy-map names can contain alphabetic, hyphen, or underscore characters, are case sensitive, and can be up to 40 characters. |

|               | Command or Action                                                        | Purpose                                                                                                                                                                                                                                                                                                       |
|---------------|--------------------------------------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Step 3</b> | switch(config-pmap-que)# <b>class type queuing</b> <i>class-name</i>     | Associates a class map with the policy map, and enters configuration mode for the specified system class.                                                                                                                                                                                                     |
| <b>Step 4</b> | switch(config-pmap-c-que)# <b>bandwidth percent</b> <i>percentage</i>    | Specifies the guaranteed percentage of interface bandwidth allocated to this class. By default, no bandwidth is specified for a class.<br><b>Note</b> Before you can successfully allocate bandwidth to the class, you must first reduce the default bandwidth configuration on class-default and class-fcoe. |
| <b>Step 5</b> | switch(config-pmap-c-que)# <b>no bandwidth percent</b> <i>percentage</i> | (Optional)<br>Removes the bandwidth specification from this class.                                                                                                                                                                                                                                            |
| <b>Step 6</b> | switch(config-pmap-c-que)# <b>priority</b>                               | Specifies that traffic in this class is mapped to a strict priority queue.<br><b>Note</b> Only one class in each policy map can have strict priority set on it.                                                                                                                                               |
| <b>Step 7</b> | switch(config-pmap-c-que)# <b>no priority</b>                            | (Optional)<br>Removes the strict priority queuing from the traffic in this class.                                                                                                                                                                                                                             |

This example shows how to define a type queuing policy map:

```
switch# configure terminal
switch(config)# policy-map type queuing policy-queue1
switch(config-pmap-que)# class type queuing class-queue1
switch(config-pmap-c-que)# bandwidth 20
```

## Configuring Type QoS Policies

Type qos policies are used for classifying the traffic of a specific system class identified by a unique qos-group value. A type qos policy can be attached to the system or to individual interfaces (including Fabric Extender host interfaces) for input traffic only.

### SUMMARY STEPS

1. switch# **configure terminal**
2. switch(config)# **policy-map type qos** *policy-name*
3. switch(config-pmap-qos)# **class type qos** *class-name*
4. switch(config-pmap-c-qos)# **set qos-group** *qos-group-value*
5. (Optional) switch(config-pmap-c-qos)# **no set qos-group** *qos-group-value*

## DETAILED STEPS

|               | Command or Action                                                         | Purpose                                                                                                                                                                                                                                                                                                             |
|---------------|---------------------------------------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Step 1</b> | switch# <b>configure terminal</b>                                         | Enters configuration mode.                                                                                                                                                                                                                                                                                          |
| <b>Step 2</b> | switch(config)# <b>policy-map type qos</b><br><i>policy-name</i>          | Creates a named object that represents a set of policies that are to be applied to a set of traffic classes. Policy-map names can contain alphabetic, hyphen, or underscore characters, are case sensitive, and can be up to 40 characters.                                                                         |
| <b>Step 3</b> | switch(config-pmap-qos)# <b>class type qos</b><br><i>class-name</i>       | Associates a class map with the policy map, and enters configuration mode for the specified system class.<br><br><b>Note</b> The associated class map must be the same type as the policy map type.                                                                                                                 |
| <b>Step 4</b> | switch(config-pmap-c-qos)# <b>set qos-group</b><br><i>qos-group-value</i> | Configures one or more <b>qos-group</b> values to match on for classification of traffic into this class map. The range of <i>qos-group-value</i> is from 2 to 5. There is no default value.<br><br><b>Note</b> The Cisco Nexus 5000 Series switch can only support a maximum of five qos-groups within this range. |
| <b>Step 5</b> | switch(config-pmap-c-qos)# <b>no set qos-group</b> <i>qos-group-value</i> | (Optional)<br>Removes the <b>qos-group</b> values from this class.                                                                                                                                                                                                                                                  |

This example shows how to define a type qos policy map:

```
switch# configure terminal
switch(config)# policy-map type qos policy-s1
switch(config-pmap-qos)# class type qos class-s1
switch(config-pmap-c-qos)# set qos-group 2
```

## Attaching the System Service Policy

You can use the **service-policy** command to associate the system class policy map as the service policy for the system.

## SUMMARY STEPS

1. switch# **configure terminal**
2. switch(config)# **system qos**
3. switch(config-sys-qos)# **service-policy type {network-qos | qos | queuing} [input | output] policy-name**

## DETAILED STEPS

|               | Command or Action                                                                                             | Purpose                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
|---------------|---------------------------------------------------------------------------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Step 1</b> | switch# <b>configure terminal</b>                                                                             | Enters configuration mode.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Step 2</b> | switch(config)# <b>system qos</b>                                                                             | Enters system class configuration mode.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Step 3</b> | switch(config-sys-qos)# <b>service-policy type {network-qos   qos   queuing} [input   output] policy-name</b> | <p>Specifies the policy map to use as the service policy for the system. There are three policy-map configuration modes:</p> <ul style="list-style-type: none"> <li>• network-qos—Network-wide (system qos) mode.</li> <li>• qos—Classification mode (system qos input or interface input only).</li> <li>• queuing—Queuing mode (input and output at system qos and interface).</li> </ul> <p><b>Note</b> There is no default policy-map configuration mode; you must specify the <b>type</b>. The <b>input</b> keyword specifies that this policy map should be applied to traffic received on an interface. The <b>output</b> keyword specifies that this policy-map should be applied to traffic transmitted from an interface. You can only apply <b>input</b> to a qos policy; you can apply both <b>input</b> and <b>output</b> to a queuing policy.</p> |

This example shows how to set a no-drop Ethernet policy map as the system class:

```
switch(config)# class-map type network-qos ethCoS4
switch(config-cmap-nq)# match cos 4
switch(config-cmap-nq)# exit
switch(config)# policy-map type network-qos ethNoDrop
switch(config-pmap-nq)# class type network-qos ethCoS4
switch(config-pmap-c-nq)# pause no-drop
switch(config-pmap-c-nq)# exit
switch(config-pmap-nq)# exit
switch(config)# system qos
switch(config-sys-qos)# service-policy type network-qos ethNoDrop
```

## Restoring the Default System Service Policies

If you have created and attached new policies to the system qos configuration, you must reapply the default policies to restore the system.

## SUMMARY STEPS

1. switch# **configure terminal**
2. switch(config)# **system qos**
3. switch(config-sys-qos)# **service-policy type qos input default-in-policy**
4. switch(config-sys-qos)# **service-policy type network-qos default-nq-policy**
5. switch(config-sys-qos)# **service-policy type queuing output default-out-policy**
6. switch(config-sys-qos)# **service-policy type queuing input default-in-policy**

## DETAILED STEPS

|        | Command or Action                                                                    | Purpose                                                                                                                   |
|--------|--------------------------------------------------------------------------------------|---------------------------------------------------------------------------------------------------------------------------|
| Step 1 | switch# <b>configure terminal</b>                                                    | Enters configuration mode.                                                                                                |
| Step 2 | switch(config)# <b>system qos</b>                                                    | Enters system class configuration mode.                                                                                   |
| Step 3 | switch(config-sys-qos)# <b>service-policy type qos input default-in-policy</b>       | Resets the classification mode policy map. This policy-map configuration is for system qos input or interface input only: |
| Step 4 | switch(config-sys-qos)# <b>service-policy type network-qos default-nq-policy</b>     | Resets the network-wide policy map.                                                                                       |
| Step 5 | switch(config-sys-qos)# <b>service-policy type queuing output default-out-policy</b> | Resets the output queuing mode policy map.                                                                                |
| Step 6 | switch(config-sys-qos)# <b>service-policy type queuing input default-in-policy</b>   | Resets the input queuing mode policy map.                                                                                 |

This example shows how to reset the system qos configuration:

```
switch# configure terminal
switch(config)# system qos
switch(config-sys-qos)# service-policy type qos input default-in-policy
switch(config-sys-qos)# service-policy type network-qos default-nq-policy
switch(config-sys-qos)# service-policy type queuing output default-out-policy
switch(config-sys-qos)# service-policy type queuing input default-in-policy
```

The default service policies are shown in this example:

```
switch# show policy-map

Type qos policy-maps
=====

policy-map type qos default-in-policy
 class type qos class-fcoe
 set qos-group 1
 class type qos class-default
 set qos-group 0

Type queuing policy-maps
=====

policy-map type queuing default-in-policy
 class type queuing class-fcoe
 bandwidth percent 50
 class type queuing class-default
 bandwidth percent 50
policy-map type queuing default-out-policy
 class type queuing class-fcoe
 bandwidth percent 50
 class type queuing class-default
 bandwidth percent 50

Type network-qos policy-maps
=====
```

```

policy-map type network-qos default-nq-policy
 class type network-qos class-fcoe
 pause no-drop
 mtu 2240
 class type network-qos class-default
 mtu 1538

```

## Enabling the Jumbo MTU

You can enable the jumbo MTU for the whole switch by setting the MTU to its maximum size (9216 bytes) in the policy map for the default Ethernet system class (class-default).

This example shows how to configure the default Ethernet system class to support the jumbo MTU:

```

switch(config)# policy-map type network-qos jumbo
switch(config-pmap-nq)# class type network-qos class-default
switch(config-pmap-c-nq)# mtu 9216
switch(config-pmap-c-nq)# exit
switch(config-pmap-nq)# exit
switch(config)# system qos
switch(config-sys-qos)# service-policy type network-qos jumbo

```



### Note

The **system jumbomtu** command defines the maximum MTU size for the switch. However, jumbo MTU is only supported for system classes that have MTU configured.

## Verifying the Jumbo MTU

To verify that the jumbo MTU is enabled, enter the **show interface ethernet slot/port** command for an Ethernet interface that carries traffic with jumbo MTU.

This example shows how to display summary jumbo MTU information for Ethernet 1/2 (the relevant part of the output is shown in bold font):

```

switch# show interface ethernet 1/2
Ethernet1/2 is up
...
Rx
1547805598 Input Packets 1547805596 Unicast Packets 0 Multicast Packets
0 Broadcast Packets 1301767362 Jumbo Packets 33690 Storm Suppression Packets
7181776513802 Bytes
Tx
1186564478 Output Packets 7060 Multicast Packets
0 Broadcast Packets 997813205 Jumbo Packets
4813632103603 Bytes
...

```

This example shows how to display detailed jumbo MTU information for Ethernet 1/2 (the relevant part of the output is shown in bold font):

```
switch# show interface ethernet 1/2 counters detailed
Rx Packets: 1547805598
Rx Unicast Packets: 1547805596
Rx Jumbo Packets: 1301767362
Rx Bytes: 7181776513802
Rx Storm Suppression: 33690
Rx Packets from 0 to 64 bytes: 169219
Rx Packets from 65 to 127 bytes: 10657133
Rx Packets from 128 to 255 bytes: 21644488
Rx Packets from 256 to 511 bytes: 43290596
Rx Packets from 512 to 1023 bytes: 86583071
Rx Packets from 1024 to 1518 bytes: 83693729
Rx Trunk Packets: 1547805596
Tx Packets: 1186564481
Tx Unicast Packets: 1005445334
Tx Multicast Packets: 7063
Tx Jumbo Packets: 997813205
Tx Bytes: 4813632103819
Tx Packets from 0 to 64 bytes: 137912
Tx Packets from 65 to 127 bytes: 8288443
Tx Packets from 128 to 255 bytes: 16596457
Tx Packets from 256 to 511 bytes: 33177999
Tx Packets from 512 to 1023 bytes: 66363944
Tx Packets from 1024 to 1518 bytes: 64186521
Tx Trunk Packets: 1005451729
```

## Configuring QoS on Interfaces

### Configuring Untagged CoS

Any incoming packet not tagged with an 802.1p CoS value is assigned the default untagged CoS value of zero (which maps to the default Ethernet drop system class). You can override the default untagged CoS value for an Ethernet or EtherChannel interface.



#### Note

Untagged CoS and type qos input policies are mutually exclusive on an Ethernet or EtherChannel interface. If a type qos policy is configured at the interface, untagged frames received over that interface will not match any **match cos 0** commands in the policy.

### SUMMARY STEPS

1. switch# **configure terminal**
2. switch(config)# **interface** {ethernet [chassis]/slot/port | port-channel channel-number}
3. switch(config-if)# **untagged cos** cos-value

### DETAILED STEPS

|        | Command or Action                 | Purpose                    |
|--------|-----------------------------------|----------------------------|
| Step 1 | switch# <b>configure terminal</b> | Enters configuration mode. |

|               | Command or Action                                                                                                                              | Purpose                                                                |
|---------------|------------------------------------------------------------------------------------------------------------------------------------------------|------------------------------------------------------------------------|
| <b>Step 2</b> | switch(config)# <b>interface</b> { <b>ethernet</b> [ <i>chassis</i> ]/ <i>slot</i> / <i>port</i>   <b>port-channel</b> <i>channel-number</i> } | Enters configuration mode for the specified interface or port channel. |
| <b>Step 3</b> | switch(config-if)# <b>untagged cos</b> <i>cos-value</i>                                                                                        | Configures the untagged CoS value.                                     |

## Configuring Interface Service Policy

An input qos policy is a service policy applied to incoming traffic on an Ethernet interface for classification. For type queuing, the output policy is applied to all outgoing traffic that matches the specified class. When you configure an input queuing policy on an interface or EtherChannel, the switch sends the configuration data to the adapter using the DCBX protocol.



### Note

Type qos policies can be activated only on Cisco Nexus 5000 Series interfaces and Cisco Nexus 2000 Series Fabric Extender interfaces. Type qos policies on Fabric Extender fabric interfaces or Fabric Extender fabric EtherChannel interfaces are ineffective, though the Cisco NX-OS CLI does not reject the configuration.

We recommend that you do not configure type qos policy-maps on Fabric Extender fabric interfaces or Fabric Extender fabric EtherChannel interfaces to avoid wasting hardware resources.

## SUMMARY STEPS

1. switch# **configure terminal**
2. switch(config)# **interface** {**ethernet** [*chassis*]/*slot*/*port* | **port-channel** *channel-number*}
3. switch(config-if)# **service-policy** [**type** {**qos** | **queuing**}] [**input** | **output**]*policy-name*
4. switch(config-if)# **service-policy input** *policy-name*

## DETAILED STEPS

|               | Command or Action                                                                                                                              | Purpose                                                                                                                                                                                                                                                     |
|---------------|------------------------------------------------------------------------------------------------------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Step 1</b> | switch# <b>configure terminal</b>                                                                                                              | Enters configuration mode.                                                                                                                                                                                                                                  |
| <b>Step 2</b> | switch(config)# <b>interface</b> { <b>ethernet</b> [ <i>chassis</i> ]/ <i>slot</i> / <i>port</i>   <b>port-channel</b> <i>channel-number</i> } | Enters configuration mode for the specified interface.<br><b>Note</b> The service policy on a port channel applies to all member interfaces.                                                                                                                |
| <b>Step 3</b> | switch(config-if)# <b>service-policy</b> [ <b>type</b> { <b>qos</b>   <b>queuing</b> }] [ <b>input</b>   <b>output</b> ] <i>policy-name</i>    | Specifies the policy map to use as the service policy for the system. There are two policy-map configuration modes: <ul style="list-style-type: none"> <li>• qos—Classification mode; this is the default mode.</li> <li>• queuing—Queuing mode.</li> </ul> |

|               | Command or Action                                              | Purpose                                                                                                                                                                                                                                                                                                                                                                        |
|---------------|----------------------------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
|               |                                                                | <p><b>Note</b> The <b>input</b> keyword specifies that this policy map should be applied to traffic received on an interface. The <b>output</b> keyword specifies that this policy map should be applied to traffic transmitted from an interface. You can only apply <b>input</b> to a qos policy; you can apply both <b>input</b> and <b>output</b> to a queuing policy.</p> |
| <b>Step 4</b> | <pre>switch(config-if)# service-policy input policy-name</pre> | <p>Applies the policy map to the interface.</p> <p><b>Note</b> There is a restriction that system type qos policy cannot be the same as any the type qos policy applied to an interface or EtherChannel.</p>                                                                                                                                                                   |

This example shows how to apply a policy to an Ethernet interface:

```
switch# configure terminal
switch(config)# interface ethernet 1/1
switch(config-if)# service-policy type qos input policy1
```

## Configuring Priority Flow Control and Link-Level Flow Control

Cisco Nexus 5000 Series switches support priority flow control (PFC) and Link-Level Flow Control (LLC) on Ethernet interfaces. The Ethernet interface can operate in two different modes: FCoE mode or standard Ethernet mode.

If the interface is operating in FCoE mode, the Ethernet link is connected at the server port using a converged network adapter (CNA).

If the interface is operating in a standard Ethernet mode, the Ethernet link is connected at the server port with a standard Ethernet network adapter (NIC). The network adapter must support the Data Center Bridging Exchange protocol (DCBX) for PFC or ingress policing to be supported on the interface.



**Note** You must configure a system class with the pause no-drop parameter for PFC to operate on Ethernet traffic (PFC will be applied to traffic that matches the CoS value configured for this class).

### Configuring Priority Flow Control

By default, Ethernet interfaces negotiate PFC with the network adapter using the DCBX protocol. When PFC is enabled, PFC is applied to traffic that matches the CoS value configured for the no-drop class.

You can override the negotiation result by forcing the interface to enable PFC.

**SUMMARY STEPS**

1. switch# **configure terminal**
2. switch(config)# **interface** *type slot/port*
3. switch(config-if)# **priority-flow-control mode** {auto | on}
4. (Optional) switch(config-if)# **no priority-flow-control mode on**

**DETAILED STEPS**

|               | Command or Action                                                | Purpose                                                                                                                                            |
|---------------|------------------------------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Step 1</b> | switch# <b>configure terminal</b>                                | Enters configuration mode.                                                                                                                         |
| <b>Step 2</b> | switch(config)# <b>interface</b> <i>type slot/port</i>           | Specifies the interface to be changed.                                                                                                             |
| <b>Step 3</b> | switch(config-if)# <b>priority-flow-control mode</b> {auto   on} | Sets PFC mode for the selected interface.<br>Specifies auto to negotiate PFC capability. This is the default.<br>Specifies on to force-enable PFC. |
| <b>Step 4</b> | switch(config-if)# <b>no priority-flow-control mode on</b>       | (Optional)<br>Disables the PFC setting for the selected interface.                                                                                 |

This example shows how to force-enable PFC on an interface:

```
switch# configure terminal
switch(config)# interface ethernet 1/2
switch(config-if)# priority-flow-control mode on
```

**Configuring Link-Level Flow Control**

By default, LLC on Ethernet interfaces is disabled. You can enable LLC for the transmit and receive directions.

**SUMMARY STEPS**

1. switch# **configure terminal**
2. switch(config)# **interface** *type slot/port*
3. switch(config-if)# **flowcontrol** [receive {on | off}] [transmit {on | off}]
4. (Optional) switch(config-if)# **no flowcontrol** [receive {on | off}] [transmit {on | off}]

**DETAILED STEPS**

|               | Command or Action                 | Purpose                    |
|---------------|-----------------------------------|----------------------------|
| <b>Step 1</b> | switch# <b>configure terminal</b> | Enters configuration mode. |

|               | Command or Action                                                                      | Purpose                                                                                              |
|---------------|----------------------------------------------------------------------------------------|------------------------------------------------------------------------------------------------------|
| <b>Step 2</b> | switch(config)# <b>interface</b> <i>type slot/port</i>                                 | Specifies the interface to be changed.                                                               |
| <b>Step 3</b> | switch(config-if)# <b>flowcontrol</b> [receive {on   off}]<br>[transmit {on   off}]    | Enables LLC for the selected interface. Set <b>receive</b> and/or <b>transmit on</b> or <b>off</b> . |
| <b>Step 4</b> | switch(config-if)# <b>no flowcontrol</b> [receive {on   off}]<br>[transmit {on   off}] | (Optional)<br>Disables LLC for the selected interface.                                               |

This example shows how to enable LLC on an interface:

```
switch# configure terminal
switch(config)# interface ethernet 1/2
switch(config-if)# flowcontrol receive on transmit on
```

## Verifying QoS Configuration

To verify QoS configuration information, perform one of these tasks:

| Command                                                                           | Purpose                                                                                                         |
|-----------------------------------------------------------------------------------|-----------------------------------------------------------------------------------------------------------------|
| switch# <b>show class-map</b>                                                     | Displays the class maps defined on the switch.                                                                  |
| switch# <b>show policy-map</b> [ <i>name</i> ]                                    | Displays the policy maps defined on the switch. Optionally, you can display the named policy only.              |
| switch# <b>show policy-map interface</b> [ <i>interface number</i> ]              | Displays the policy map settings for an interface or all interfaces.                                            |
| switch# <b>show policy-map system</b>                                             | Displays the policy map settings attached to the system qos.                                                    |
| switch# <b>show policy-map type</b> {network-qos   qos   queuing} [ <i>name</i> ] | Displays the policy map settings for a specific policy type. Optionally, you can display the named policy only. |
| switch# <b>show queuing interface</b> [ <i>interface number</i> ]                 | Displays the queue configuration and statistics.                                                                |

You can clear the QoS policy statistics.

| Command                             | Purpose                       |
|-------------------------------------|-------------------------------|
| switch# <b>clear qos statistics</b> | Clears the policy statistics. |

This example shows how to display the class maps defined on the switch:

```
switch# show class-map
```

```
Type qos class-maps
=====

class-map type qos c1
 match cos 0,7

class-map type qos c2
 match protocol ldp
 match ip rtp 2000-65535
 match dscp 10,12
 match precedence 6-7
 match protocol dhcp
 match protocol arp

class-map type qos c3
 match cos 2,4-6

class-map type qos c4
 match access-group name ipv4

class-map type qos class-fcoe
 match cos 3

class-map type qos class-default
 match any

class-map type qos class-ip-multicast
 match ip multicast

Type queuing class-maps
=====

class-map type queuing c1
 match qos-group 2

class-map type queuing c2
 match qos-group 3

class-map type queuing c3
 match qos-group 4

class-map type queuing class-fcoe
 match qos-group 1

class-map type queuing class-default
 match qos-group 0

Type network-qos class-maps
=====

class-map type network-qos c1
 match qos-group 2

class-map type network-qos c2
 match qos-group 3

class-map type network-qos c3
 match qos-group 4

class-map type network-qos c4
 match qos-group 5

class-map type network-qos class-fcoe
 match qos-group 1
```

```
class-map type network-qos class-default
 match qos-group 0
```

This example shows how to display the policy maps defined on the switch:

```
switch# show policy-map
```

```
Type qos policy-maps
=====

policy-map type qos p1
 class type qos c1
 set qos-group 2
 class type qos c3
 set qos-group 4
 class type qos c4
 set qos-group 5
 class type qos c2
 set qos-group 3
 class type qos c22
 set qos-group 3
 class type qos class-fcoe
 set qos-group 1
 class type qos class-default
 set qos-group 0

Type queuing policy-maps
=====

policy-map type queuing p1
 class type queuing c2
 bandwidth percent 10
 class type queuing c4
 bandwidth percent 25
 class type queuing c1
 bandwidth percent 20
 class type queuing c3
 bandwidth percent 5
 class type queuing class-fcoe
 bandwidth percent 30
 class type queuing class-default
 bandwidth percent 10

Type network-qos policy-maps
=====

policy-map type network-qos p1
 class type network-qos c1
 mtu 5000
 class type network-qos c2
 mtu 9216
 queue-limit 30000 bytes
 class type network-qos c3
 mtu 8000
 class type network-qos c4
 pause no-drop
 class type network-qos class-fcoe
 pause no-drop
 mtu 2240
 class type network-qos class-default
 mtu 1538
```

This example shows how to display the policy maps attached on the system qos:

```
switch# show policy-map system
```

```
Type network-qos policy-maps
=====
```

```

policy-map type network-qos p1
 class type network-qos c1 match qos-group 2

 mtu 5000
 class type network-qos c2 match qos-group 3

 mtu 9216
 queue-limit 30000 bytes
 class type network-qos c3 match qos-group 4

 mtu 8000
 class type network-qos c4 match qos-group 5

 pause no-drop
 class type network-qos class-fcoe match qos-group 1

 pause no-drop
 mtu 2240
 class type network-qos class-default match qos-group 0

 mtu 1538

Service-policy (queuing) input: p1
policy statistics status: disabled

Class-map (queuing): c2 (match-any)
 Match: qos-group 3
 bandwidth percent 10

Class-map (queuing): c4 (match-any)
 Match: qos-group 5
 bandwidth percent 25

Class-map (queuing): c1 (match-any)
 Match: qos-group 2
 bandwidth percent 20

Class-map (queuing): c3 (match-any)
 Match: qos-group 4
 bandwidth percent 5

Class-map (queuing): class-fcoe (match-any)
 Match: qos-group 1
 bandwidth percent 30

Class-map (queuing): class-default (match-any)
 Match: qos-group 0
 bandwidth percent 10

Service-policy (queuing) output: default-out-policy
policy statistics status: disabled

Class-map (queuing): class-fcoe (match-any)
 Match: qos-group 1
 bandwidth percent 50

Class-map (queuing): class-default (match-any)
 Match: qos-group 0
 bandwidth percent 50

```

This example shows how to display the policy maps attached to an interface:

```

switch# show policy-map interface ethernet 1/1

Global statistics status : disabled

Ethernet1/1

Service-policy (qos) input: p2
policy statistics status: disabled

```

```
Class-map (qos): c1 (match-any)
 Match: cos 0,7
 set qos-group 2

Class-map (qos): c2 (match-any)
 Match: protocol ldp
 Match: ip rtp 2000-65535
 Match: dscp 10,12
 Match: precedence 6-7
 Match: protocol dhcp
 Match: protocol arp
 set qos-group 3

Class-map (qos): c3 (match-any)
 Match: cos 2,4-6
 set qos-group 4

Class-map (qos): class-ip-multicast (match-any)
 Match: ip multicast
 set qos-group 5
Class-map (qos): class-fcoe (match-any)
 Match: cos 3
 set qos-group 1

Class-map (qos): class-default (match-any)
 Match: any
 set qos-group 0

Service-policy (queuing) input: p1
 policy statistics status: disabled

Class-map (queuing): c2 (match-any)
 Match: qos-group 3
 bandwidth percent 10

Class-map (queuing): c4 (match-any)
 Match: qos-group 5
 bandwidth percent 25

Class-map (queuing): c1 (match-any)
 Match: qos-group 2
 bandwidth percent 20

Class-map (queuing): c3 (match-any)
 Match: qos-group 4
 bandwidth percent 5

Class-map (queuing): class-fcoe (match-any)
 Match: qos-group 1
 bandwidth percent 30

Class-map (queuing): class-default (match-any)
 Match: qos-group 0
 bandwidth percent 10

Service-policy (queuing) output: p2
 policy statistics status: disabled

Class-map (queuing): c1 (match-any)
 Match: qos-group 2
 bandwidth percent 5
 priority

Class-map (queuing): c2 (match-any)
 Match: qos-group 3
 bandwidth percent 20

Class-map (queuing): c3 (match-any)
 Match: qos-group 4
 bandwidth percent 20
```

```

Class-map (queuing): c4 (match-any)
 Match: qos-group 5
 bandwidth percent 40

Class-map (queuing): class-fcoe (match-any)
 Match: qos-group 1
 bandwidth percent 10

Class-map (queuing): class-default (match-any)
 Match: qos-group 0
 bandwidth percent 5

```

This example shows how to display the queue configuration and statistics:

```

switch# show queuing interface ethernet 1/1

Interface Ethernet1/1 TX Queuing
qos-group sched-type oper-bandwidth
0 WRR 5
1 WRR 10
2 priority 5
3 WRR 20
4 WRR 20
5 WRR 40

Interface Ethernet1/1 RX Queuing
qos-group 0:
q-size: 21120, MTU: 1538
drop-type: drop, xon: 0, xoff: 132
Statistics:
 Pkts received over the port : 1265258330
 Ucast pkts sent to the cross-bar : 0
 Mcast pkts sent to the cross-bar : 2883444539
 Ucast pkts received from the cross-bar : 0
 Pkts sent to the port : 367529517
 Pkts discarded on ingress : 781087 (36419)
 Per-priority-pause status : Rx (Inactive), Tx (Inactive)

qos-group 1:
q-size: 76800, MTU: 2240
drop-type: no-drop, xon: 128, xoff: 240
Statistics:
 Pkts received over the port : 0
 Ucast pkts sent to the cross-bar : 0
 Mcast pkts sent to the cross-bar : 0
 Ucast pkts received from the cross-bar : 0
 Pkts sent to the port : 0
 Pkts discarded on ingress : 0 (0)
 Per-priority-pause status : Rx (Inactive), Tx (Inactive)

qos-group 2:
q-size: 20480, MTU: 5000
drop-type: drop, xon: 0, xoff: 128
Statistics:
 Pkts received over the port : 0
 Ucast pkts sent to the cross-bar : 0
 Mcast pkts sent to the cross-bar : 0
 Ucast pkts received from the cross-bar : 0
 Pkts sent to the port : 0
 Pkts discarded on ingress : 0 (0)
 Per-priority-pause status : Rx (Inactive), Tx (Inactive)

```

```

qos-group 3:
 q-size: 30080, MTU: 9216
 drop-type: drop, xon: 0, xoff: 188
 Statistics:
 Pkts received over the port : 0
 Ucast pkts sent to the cross-bar : 0
 Mcast pkts sent to the cross-bar : 0
 Ucast pkts received from the cross-bar : 0
 Pkts sent to the port : 0
 Pkts discarded on ingress : 0 (0)
 Per-priority-pause status : Rx (Inactive), Tx (Inactive)

qos-group 4:
 q-size: 20480, MTU: 8000
 drop-type: drop, xon: 0, xoff: 128
 Statistics:
 Pkts received over the port : 0
 Ucast pkts sent to the cross-bar : 0
 Mcast pkts sent to the cross-bar : 0
 Ucast pkts received from the cross-bar : 0
 Pkts sent to the port : 0
 Pkts discarded on ingress : 0(0)
 Per-priority-pause status : Rx (Inactive), Tx (Inactive)

qos-group 5:
 q-size: 76800, MTU: 1538
 drop-type: no-drop, xon: 128, xoff: 242
 Statistics:
 Pkts received over the port : 0
 Ucast pkts sent to the cross-bar : 0
 Mcast pkts sent to the cross-bar : 0
 Ucast pkts received from the cross-bar : 0
 Pkts sent to the port : 0
 Pkts discarded on ingress : 0(0)
 Per-priority-pause status : Rx (Inactive), Tx (Inactive)

Total Multicast crossbar statistics:
 Mcast pkts received from the cross-bar : 557155515

```

## Example QoS Configurations

### QoS Example 1

This example shows how to configure traffic in the entire system matching an access control list to have the frame CoS fields rewritten to the value 5.

#### SUMMARY STEPS

1. Set up the ingress classification policy (the access control list was defined previously).
2. Attach the classification policy to the system.
3. Set up the system class allocation and rewrite policy. Allocate the system class for qos-group 4 and define the rewrite action.
4. Attach the allocation and rewrite policy to the system.

## DETAILED STEPS

|               | Command or Action                                                                                                               | Purpose                                                                                                                                                                                                                                                                                                     |
|---------------|---------------------------------------------------------------------------------------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Step 1</b> | Set up the ingress classification policy (the access control list was defined previously).                                      | <pre>(config)# class-map type qos cmap-qos-acl (config-cmap-qos)# match access-group ACL-CoS (config-cmap-qos)# exit (config)# policy-map type qos pmap-qos-acl (config-pmap-qos)# class cmap-qos-acl (config-pmap-c-qos)# set qos-group 4 (config-pmap-c-qos)# exit (config-pmap-qos)# exit</pre>          |
| <b>Step 2</b> | Attach the classification policy to the system.                                                                                 | <pre>(config)# system qos (config-sys-qos)# service-policy type qos input pmap-qos-acl (config-sys-qos)# exit</pre>                                                                                                                                                                                         |
| <b>Step 3</b> | Set up the system class allocation and rewrite policy. Allocate the system class for qos-group 4 and define the rewrite action. | <pre>(config)# class-map type network-qos cmap-nq-acl (config-cmap-nq)# match qos-group 4 (config-cmap-nq)# exit (config)# policy-map type network-qos pmap-nq-acl (config-pmap-nq)# class type network-qos cmap-nq-acl (config-pmap-c-nq)# set cos 5 (config-pmap-c-nq)# exit (config-pmap-nq)# exit</pre> |
| <b>Step 4</b> | Attach the allocation and rewrite policy to the system.                                                                         | <pre>(config)# system qos (config-sys-qos)# service-policy type network-qos pmap-nq-acl (config-sys-qos)# exit</pre>                                                                                                                                                                                        |

## QoS Example 2

This example shows how to use an access control list to apply 50% bandwidth to traffic on Ethernet interface 1/3 that matches traffic on Ethernet interface 1/1.

## SUMMARY STEPS

1. Set up the ingress classification policy.
2. Attach the classification policy to the interface Ethernet 1/1.
3. Set up the system-wide definition of the qos-group first.
4. Set up the egress bandwidth policy.
5. Attach the bandwidth policy to the egress interface.
6. Allocate the system class for qos-group 2.
7. Set up the network-qos policy.
8. Attach the network-qos policy to the system.

## DETAILED STEPS

|               | Command or Action                                               | Purpose                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
|---------------|-----------------------------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Step 1</b> | Set up the ingress classification policy.                       | <pre>(config)# class-map type qos cmap-qos-bandwidth (config-cmap-qos)# match access-group ACL-bandwidth (config-cmap-qos)# exit (config)# policy-map type qos pmap-qos-eth1-1 (config-pmap-qos)# class cmap-qos-bandwidth (config-pmap-c-qos)# set qos-group 2 (config-pmap-c-qos)# exit (config-pmap-qos)# exit</pre>                                                                                                                                                                                                                                                                                                                                                   |
| <b>Step 2</b> | Attach the classification policy to the interface Ethernet 1/1. | <pre>(config)# interface ethernet 1/1 (config-if)# service-policy type qos input pmap-qos-eth1-1 (config-if)# exit</pre>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Step 3</b> | Set up the system-wide definition of the qos-group first.       | <pre>(config)# class-map type queuing cmap-que-bandwidth (config-cmap-que)# match qos-group 2 (config-cmap-que)# exit</pre>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>Step 4</b> | Set up the egress bandwidth policy.                             | <p><b>Note</b> Before you can successfully allocate bandwidth to the user-defined class cmap-que-bandwidth, you must first reduce the default bandwidth configuration on class-default and class-fcoe.</p> <pre>(config)# policy-map type queuing pmap-que-eth1-2 (config-pmap-que)# class type queuing class-default (config-pmap-c-que)# bandwidth percent 10 (config-pmap-c-que)# exit (config-pmap-que)# class type queuing class-fcoe (config-pmap-c-que)# bandwidth percent 40 (config-pmap-c-que)# exit (config-pmap-que)# class type queuing cmap-que-bandwidth (config-pmap-c-que)# bandwidth percent 50 (config-pmap-c-que)# exit (config-pmap-que)# exit</pre> |

|               | Command or Action                                    | Purpose                                                                                                                                                                                                 |
|---------------|------------------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Step 5</b> | Attach the bandwidth policy to the egress interface. | (config)# <b>interface ethernet 1/3</b><br>(config-if)# <b>service-policy type queuing output pmap-que-eth1-2</b><br>(config-if)# <b>exit</b>                                                           |
| <b>Step 6</b> | Allocate the system class for qos-group 2.           | (config)# <b>class-map type network-qos cmap-nq-bandwidth</b><br>(config-cmap-nq)# <b>match qos-group 2</b><br>(config-cmap-nq)# <b>exit</b>                                                            |
| <b>Step 7</b> | Set up the network-qos policy.                       | (config)# <b>policy-map type network-qos pmap-nq-bandwidth</b><br>(config-pmap-nq)# <b>class type network-qos cmap-nq-bandwidth</b><br>(config-pmap-c-nq)# <b>exit</b><br>(config-pmap-nq)# <b>exit</b> |
| <b>Step 8</b> | Attach the network-qos policy to the system.         | (config)# <b>system qos</b><br>(config-sys-qos)# <b>service-policy type network-qos pmap-nq-bandwidth</b><br>(config-sys-qos)# <b>exit</b>                                                              |

### QoS Example 3

This example shows how to attach a 802.1p tag with a CoS value of 3 to incoming untagged packets, and force priority-flow-control negotiation on Ethernet interface 1/15.

#### SUMMARY STEPS

1. Set up the ingress classification policy (the access control list was defined previously).

#### DETAILED STEPS

|               | Command or Action                                                                          | Purpose                                                                                                                                                         |
|---------------|--------------------------------------------------------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Step 1</b> | Set up the ingress classification policy (the access control list was defined previously). | (config)# <b>interface Ethernet 1/15</b><br>(config-if)# <b>untagged cos 3</b><br>(config-if)# <b>priority-flow-control mode on</b><br>(config-if)# <b>exit</b> |





# PART VII

## SAN Switching

- [Configuring Fibre Channel Interfaces, page 511](#)
- [Configuring Domain Parameters, page 531](#)
- [Configuring N Port Virtualization, page 553](#)
- [Configuring VSAN Trunking, page 563](#)
- [Configuring SAN Port Channel, page 571](#)
- [Configuring and Managing VSANs, page 589](#)
- [Configuring and Managing Zones, page 603](#)
- [Distributing Device Alias Services, page 631](#)
- [Configuring Fibre Channel Routing Services and Protocols, page 643](#)
- [Managing FLOGI, Name Server, FDMI, and RSCN Databases, page 663](#)
- [Discovering SCSI Targets, page 675](#)
- [Advanced Fibre Channel Features and Concepts, page 679](#)
- [Configuring FC-SP and DHCHAP, page 699](#)
- [Configuring Port Security, page 713](#)
- [Configuring Fabric Binding, page 735](#)
- [Configuring Fabric Configuration Servers, page 743](#)
- [Configuring Port Tracking, page 747](#)





# Configuring Fibre Channel Interfaces

This chapter contains the following sections:

- [Configuring Fibre Channel Interfaces](#), page 511

## Configuring Fibre Channel Interfaces

### Information About Fibre Channel Interfaces

#### Licensing Requirements for Fibre Channel

On Cisco Nexus 5000 Series switches, Fibre Channel capability is included in the Storage Protocol Services license.

Ensure that you have the correct license installed (N5010SS or N5020SS) before using Fibre Channel interfaces and capabilities.



**Note**

You can configure virtual Fibre Channel interfaces without a Storage Protocol Services license, but these interfaces will not become operational until the license is activated.

#### Physical Fibre Channel Interfaces

Cisco Nexus 5000 Series switches provide up to sixteen physical Fibre Channel uplinks. The Fibre Channel interfaces are supported on optional expansion modules. The Fibre Channel expansion module contains eight Fibre Channel interfaces. The Fibre Channel plus Ethernet expansion module contains four Fibre Channel interfaces.

Each Fibre Channel port can be used as a downlink (connected to a server) or as an uplink (connected to the data center SAN network). The Fibre Channel interfaces support the following modes: F, NP, E, TE, and SD.

## Virtual Fibre Channel Interfaces

Fibre Channel over Ethernet (FCoE) encapsulation allows a physical Ethernet cable to simultaneously carry Fibre Channel and Ethernet traffic. In Cisco Nexus 5000 Series switches, an FCoE-capable physical Ethernet interface can carry traffic for one virtual Fibre Channel interface.

Native Fibre Channel and virtual Fibre Channel interfaces are configured using the same CLI commands. Virtual Fibre Channel interfaces support only F mode, and offer a subset of the features that are supported on native Fibre Channel interfaces.

The following capabilities are not supported for virtual Fibre Channel interfaces:

- SAN port channels.
- VSAN trunking. The virtual Fibre Channel is associated with one VSAN.
- The SPAN destination cannot be a virtual Fibre Channel interface.
- Buffer-to-buffer credits.
- Exchange link parameters (ELP), or Fabric Shortest Path First (FSPF) protocol.
- Configuration of physical attributes (speed, rate, mode, transmitter information, MTU size).
- Port tracking.

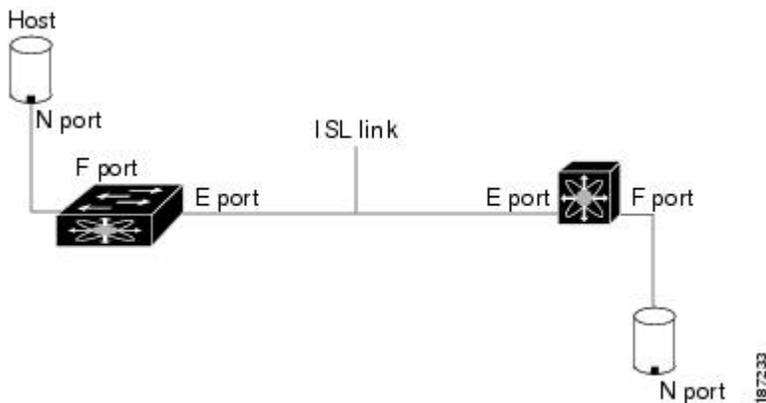
## Interface Modes

Each physical Fibre Channel interface in a switch may operate in one of several port modes: E mode, TE mode, F mode, and SD mode (see the figure below). A physical Fibre Channel interface can be configured as an E port, an F port, or an SD port. Interfaces may also be configured in Auto mode; the port type is determined during interface initialization.

In NPV mode, Fibre Channel interfaces may operate in NP mode, F mode or SD mode.

Virtual Fibre Channel interfaces can only be configured in F mode.

**Figure 41: Switch Port Modes**



**Note**


---

Interfaces are automatically assigned VSAN 1 by default.

---

Each interface has an associated administrative configuration and an operational status:

- The administrative configuration does not change unless you modify it. This configuration has various attributes that you can configure in administrative mode.
- The operational status represents the current status of a specified attribute such as the interface speed. This status cannot be changed and is read-only. Some values may not be valid when the interface is down (for example, the operational speed).

**Related Topics**

[Configuring and Managing VSANs, on page 589](#)

[Configuring N Port Virtualization, on page 553](#)

*E Port*

In expansion port (E port) mode, an interface functions as a fabric expansion port. This port may be connected to another E port to create an Inter-Switch Link (ISL) between two switches. E ports carry frames between switches for configuration and fabric management. They serve as a conduit between switches for frames destined to remote N ports. E ports support class 3 and class F service.

An E port connected to another switch may also be configured to form a SAN port channel.

**Related Topics**

[Configuring SAN Port Channel, on page 571](#)

*F Port*

In fabric port (F port) mode, an interface functions as a fabric port. This port may be connected to a peripheral device (host or disk) operating as a node port (N port). An F port can be attached to only one N port. F ports support class 3 service.

*NP Port*

When the switch is operating in NPV mode, the interfaces that connect the switch to the core network switch are configured as NP ports. NP ports operate like N ports that function as proxies for multiple physical N ports.

**Related Topics**

[Configuring N Port Virtualization, on page 553](#)

*TE Port*

In trunking E port (TE port) mode, an interface functions as a trunking expansion port. It may be connected to another TE port to create an extended ISL (EISL) between two switches. TE ports connect to another Cisco Nexus 5000 Series switch or a Cisco MDS 9000 Family switch. They expand the functionality of E ports to support the following:

- VSAN trunking
- Fibre Channel trace (fctrace) feature

In TE port mode, all frames are transmitted in EISL frame format, which contains VSAN information. Interconnected switches use the VSAN ID to multiplex traffic from one or more VSANs across the same

physical link. This feature is referred to as VSAN trunking in the Cisco Nexus 5000 Series switch. TE ports support class 3 and class F service.

### Related Topics

[Configuring VSAN Trunking, on page 563](#)

### SD Port

In SPAN destination port (SD port) mode, an interface functions as a switched port analyzer (SPAN). The SPAN feature monitors network traffic that passes through a Fibre Channel interface. This monitoring is done using a standard Fibre Channel analyzer (or a similar switch probe) that is attached to an SD port. SD ports do not receive frames, instead they transmit a copy of the source traffic. The SPAN feature is nonintrusive and does not affect switching of network traffic for any SPAN source ports.

### Auto Mode

Interfaces configured in auto mode can operate in one of the following modes: F port, E port, or TE port. The port mode is determined during interface initialization. For example, if the interface is connected to a node (host or disk), it operates in F port mode. If the interface is attached to a third-party switch, it operates in E port mode. If the interface is attached to another switch in the Cisco Nexus 5000 Series or Cisco MDS 9000 Family, it may become operational in TE port mode.

SD ports are not determined during initialization and are administratively configured.

### Related Topics

[Configuring VSAN Trunking, on page 563](#)

## Interface States

The interface state depends on the administrative configuration of the interface and the dynamic state of the physical link.

### Administrative States

The administrative state refers to the administrative configuration of the interface. The table below describes the administrative states.

**Table 59: Administrative States**

| Administrative State | Description                                                                                                                                           |
|----------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------|
| Up                   | Interface is enabled.                                                                                                                                 |
| Down                 | Interface is disabled. If you administratively disable an interface by shutting down that interface, the physical link layer state change is ignored. |

### Operational States

The operational state indicates the current operational state of the interface. The table below describes the operational states.

**Table 60: Operational States**

| Operational State | Description                                                                                                                                                                                                              |
|-------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Up                | Interface is transmitting or receiving traffic as desired. To be in this state, an interface must be administratively up, the interface link layer state must be up, and the interface initialization must be completed. |
| Down              | Interface cannot transmit or receive (data) traffic.                                                                                                                                                                     |
| Trunking          | Interface is operational in TE mode.                                                                                                                                                                                     |

**Reason Codes**

Reason codes are dependent on the operational state of the interface. The following table describes the reason codes for operational states.

**Table 61: Reason Codes for Interface States**

| Administrative Configuration | Operational Status | Reason Code                                                                                                                                      |
|------------------------------|--------------------|--------------------------------------------------------------------------------------------------------------------------------------------------|
| Up                           | Up                 | None.                                                                                                                                            |
| Down                         | Down               | Administratively down. If you administratively configure an interface as down, you disable the interface. No traffic is received or transmitted. |
| Up                           | Down               | See the table below.                                                                                                                             |

If the administrative state is up and the operational state is down, the reason code differs based on the nonoperational reason code. The table below describes the reason codes for nonoperational states.

**Note**

Only some of the reason codes are listed in the table.

**Table 62: Reason Codes for Nonoperational States**

| Reason Code (long version)    | Description                                                       | Applicable Modes |
|-------------------------------|-------------------------------------------------------------------|------------------|
| Link failure or not connected | The physical layer link is not operational.                       | All              |
| SFP not present               | The small form-factor pluggable (SFP) hardware is not plugged in. | All              |

| Reason Code (long version)                                   | Description                                                                                                                                                                                                                                                                                                                                                                                                                   | Applicable Modes          |
|--------------------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|---------------------------|
| Initializing                                                 | The physical layer link is operational and the protocol initialization is in progress.                                                                                                                                                                                                                                                                                                                                        | All                       |
| Reconfigure fabric in progress                               | The fabric is currently being reconfigured.                                                                                                                                                                                                                                                                                                                                                                                   |                           |
| Offline                                                      | The switch software waits for the specified R_A_TOV time before retrying initialization.                                                                                                                                                                                                                                                                                                                                      |                           |
| Inactive                                                     | The interface VSAN is deleted or is in a suspended state.<br><br>To make the interface operational, assign that port to a configured and active VSAN.                                                                                                                                                                                                                                                                         |                           |
| Hardware failure                                             | A hardware failure is detected.                                                                                                                                                                                                                                                                                                                                                                                               |                           |
| Error disabled                                               | Error conditions require administrative attention. Interfaces may be error-disabled for various reasons. For example: <ul style="list-style-type: none"> <li>• Configuration failure.</li> <li>• Incompatible buffer-to-buffer credit configuration.</li> </ul><br>To make the interface operational, you must first fix the error conditions causing this state and then administratively shut down or enable the interface. |                           |
| Isolation because limit of active port channels is exceeded. | The interface is isolated because the switch is already configured with the maximum number of active SAN port channels.                                                                                                                                                                                                                                                                                                       |                           |
| Isolation due to ELP failure                                 | The port negotiation failed.                                                                                                                                                                                                                                                                                                                                                                                                  | Only E ports and TE ports |
| Isolation due to ESC failure                                 | The port negotiation failed.                                                                                                                                                                                                                                                                                                                                                                                                  |                           |
| Isolation due to domain overlap                              | The Fibre Channel domains (fcdomain) overlap.                                                                                                                                                                                                                                                                                                                                                                                 |                           |
| Isolation due to domain ID assignment failure                | The assigned domain ID is not valid.                                                                                                                                                                                                                                                                                                                                                                                          |                           |

| Reason Code (long version)                                  | Description                                                                                                                                                          | Applicable Modes                      |
|-------------------------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------|---------------------------------------|
| Isolation due to the other side of the link E port isolated | The E port at the other end of the link is isolated.                                                                                                                 |                                       |
| Isolation due to invalid fabric reconfiguration             | The port is isolated due to fabric reconfiguration.                                                                                                                  |                                       |
| Isolation due to domain manager disabled                    | The fcdomain feature is disabled.                                                                                                                                    |                                       |
| Isolation due to zone merge failure                         | The zone merge operation failed.                                                                                                                                     |                                       |
| Isolation due to VSAN mismatch                              | The VSANs at both ends of an ISL are different.                                                                                                                      |                                       |
| port channel administratively down                          | The interfaces belonging to the SAN port channel are down.                                                                                                           | Only SAN port channel interfaces      |
| Suspended due to incompatible speed                         | The interfaces belonging to the SAN port channel have incompatible speeds.                                                                                           |                                       |
| Suspended due to incompatible mode                          | The interfaces belonging to the SAN port channel have incompatible modes.                                                                                            |                                       |
| Suspended due to incompatible remote switch WWN             | An improper connection is detected. All interfaces in a SAN port channel must be connected to the same pair of switches.                                             |                                       |
| Bound physical interface down                               | The Ethernet interface bound to a virtual Fibre Channel interface is not operational.                                                                                | Only virtual Fibre Channel interfaces |
| STP not forwarding in FCoE mapped VLAN                      | The Ethernet interface bound to a virtual Fibre Channel interface is not in an STP forwarding state for the VLAN associated with the virtual Fibre Channel interface | Only virtual Fibre Channel interfaces |

### Buffer-to-Buffer Credits

Buffer-to-buffer credits (BB\_credits) are a flow-control mechanism to ensure that Fibre Channel interfaces do not drop frames. BB\_credits are negotiated on a per-hop basis.

In Cisco Nexus 5000 Series switches, the BB\_credit mechanism is used on Fibre Channel interfaces but not on virtual Fibre Channel interfaces. Virtual Fibre Channel interfaces provide flow control based on capabilities of the underlying physical Ethernet interface.

The receive BB\_credit value (fcrxbbcredit) may be configured for each Fibre Channel interface. In most cases, you do not need to modify the default configuration.



**Note** The receive BB\_credit values depend on the port mode. For physical Fibre Channel interfaces, the default value is 16 for F mode and E mode interfaces. This value can be changed as required. The maximum value is 64.

For virtual Fibre Channel interfaces, BB\_credits are not used.

## Configuring Fibre Channel Interfaces

### Configuring a Fibre Channel Interface

To configure a Fibre Channel interface, perform this task:

#### SUMMARY STEPS

1. switch# **configuration terminal**
2. switch(config)# **interface** { **fc slot/port** } | { **vfc vfc-id** }

#### DETAILED STEPS

|               | Command or Action                                                                | Purpose                                                                                                                                                                                                                                                                                              |
|---------------|----------------------------------------------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Step 1</b> | switch# <b>configuration terminal</b>                                            | Enters configuration mode.                                                                                                                                                                                                                                                                           |
| <b>Step 2</b> | switch(config)# <b>interface</b> { <b>fc slot/port</b> }   { <b>vfc vfc-id</b> } | Selects a Fibre Channel interface and enters interface configuration mode.<br><br><b>Note</b> When a Fibre Channel interface is configured, it is automatically assigned a unique world wide name (WWN). If the interface's operational state is up, it is also assigned a Fibre Channel ID (FC ID). |

### Configuring a Range of Fibre Channel Interfaces

To configure a range of Fibre Channel interfaces, perform this task:

#### SUMMARY STEPS

1. switch# **configuration terminal**
2. switch(config)# **interface** { **fc slot/port - port** [ , **fc slot/port - port** ] | **vfc vfc-id - vfc-id** [ , **vfc vfc-id - vfc-id** ] }

**DETAILED STEPS**

|               | <b>Command or Action</b>                                                                                                                                       | <b>Purpose</b>                                                                         |
|---------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------|----------------------------------------------------------------------------------------|
| <b>Step 1</b> | switch# <b>configuration terminal</b>                                                                                                                          | Enters configuration mode.                                                             |
| <b>Step 2</b> | switch(config)# <b>interface</b> { <b>fc slot/port - port</b> [ , <b>fc slot/port - port</b> ]   <b>vfc vfc-id - vfc-id</b> [ , <b>vfc vfc-id - vfc-id</b> ] } | Selects the range of Fibre Channel interfaces and enters interface configuration mode. |

**Setting the Interface Administrative State**

To gracefully shut down an interface, perform this task:

To enable traffic flow, perform this task:

**SUMMARY STEPS**

1. switch# **configuration terminal**
2. switch(config)# **interface** { **fc slot/port** } | { **vfc vfc-id** }
3. switch(config-if)# **shutdown**

**DETAILED STEPS**

|               | <b>Command or Action</b>                                                         | <b>Purpose</b>                                                                            |
|---------------|----------------------------------------------------------------------------------|-------------------------------------------------------------------------------------------|
| <b>Step 1</b> | switch# <b>configuration terminal</b>                                            | Enters configuration mode.                                                                |
| <b>Step 2</b> | switch(config)# <b>interface</b> { <b>fc slot/port</b> }   { <b>vfc vfc-id</b> } | Selects a Fibre Channel interface and enters interface configuration mode.                |
| <b>Step 3</b> | switch(config-if)# <b>shutdown</b>                                               | Gracefully shuts down the interface and administratively disables traffic flow (default). |

**Configuring Interface Modes**

To configure the interface mode, perform this task:

**SUMMARY STEPS**

1. switch# **configuration terminal**
2. switch(config)# **interface** { **fc slot/port** } | { **vfc vfc-id** }
3. switch(config-if)# **switchport mode** **E** | **F** | **SD** | **auto**

## DETAILED STEPS

|               | Command or Action                                            | Purpose                                                                                                                                                                                                                                                                                                                                                  |
|---------------|--------------------------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Step 1</b> | switch# <b>configuration terminal</b>                        | Enters configuration mode.                                                                                                                                                                                                                                                                                                                               |
| <b>Step 2</b> | switch(config)# <b>interface</b> {fc slot/port} {vfc vfc-id} | Selects a Fibre Channel interface and enters interface configuration mode.                                                                                                                                                                                                                                                                               |
| <b>Step 3</b> | switch(config-if)# <b>switchport mode E   F   SD   auto</b>  | For a Fibre Channel interface, you can set the mode to E, F, or SD port mode. Set the mode to auto to auto-negotiate an E, F, TE port mode (not SD port mode) of operation.<br><br><b>Note</b> SD ports cannot be configured automatically. They must be administratively configured.<br>For a virtual Fibre Channel, only the F port mode is supported. |

## Configuring the Interface Description

Interface descriptions should help you identify the traffic or use for that interface. The interface description can be any alphanumeric string.

To configure a description for an interface, perform this task:

## SUMMARY STEPS

1. switch# **configuration terminal**
2. switch(config)# **interface** {fc slot/port}|{vfc vfc-id}
3. switch(config-if)# **switchport description cisco-HBA2**
4. switch(config-if)# **no switchport description**

## DETAILED STEPS

|               | Command or Action                                            | Purpose                                                                                  |
|---------------|--------------------------------------------------------------|------------------------------------------------------------------------------------------|
| <b>Step 1</b> | switch# <b>configuration terminal</b>                        | Enters configuration mode.                                                               |
| <b>Step 2</b> | switch(config)# <b>interface</b> {fc slot/port} {vfc vfc-id} | Selects a Fibre Channel interface and enters interface configuration mode.               |
| <b>Step 3</b> | switch(config-if)# <b>switchport description cisco-HBA2</b>  | Configures the description of the interface. The string can be up to 80 characters long. |
| <b>Step 4</b> | switch(config-if)# <b>no switchport description</b>          | Clears the description of the interface.                                                 |

## Configuring Port Speeds

Port speed can be configured on a physical Fibre Channel interface but not on a virtual Fibre Channel interface. By default, the port speed for an interface is automatically calculated by the switch.



### Caution

Changing the interface speed is a disruptive operation.

To configure the port speed of the interface, perform this task:

## SUMMARY STEPS

1. switch# **configuration terminal**
2. switch(config)# **interface fc slot/port**
3. switch(config-if)# **switchport speed 1000**
4. switch(config-if)# **no switchport speed**

## DETAILED STEPS

|               | Command or Action                               | Purpose                                                                                                                                                                                                                                                         |
|---------------|-------------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Step 1</b> | switch# <b>configuration terminal</b>           | Enters configuration mode.                                                                                                                                                                                                                                      |
| <b>Step 2</b> | switch(config)# <b>interface fc slot/port</b>   | Selects the specified interface and enters interface configuration mode.<br><b>Note</b> You cannot configure the port speed of a virtual Fibre Channel interface.                                                                                               |
| <b>Step 3</b> | switch(config-if)# <b>switchport speed 1000</b> | Configures the port speed of the interface to 1000 Mbps.<br>The number indicates the speed in megabits per second (Mbps). You can set the speed to 1000 (for 1-Gbps interfaces), 2000 (for 2-Gbps interfaces), 4000 (for 4-Gbps interfaces), or auto (default). |
| <b>Step 4</b> | switch(config-if)# <b>no switchport speed</b>   | Reverts to the factory default (auto) administrative speed of the interface.                                                                                                                                                                                    |

### Autosensing

Autosensing speed is enabled on all 4-Gbps interfaces by default. This configuration enables the interfaces to operate at speeds of 1 Gbps, 2 Gbps, or 4 Gbps on the 4-Gbps ports. When autosensing is enabled for an interface operating in dedicated rate mode, 4-Gbps of bandwidth is reserved, even if the port negotiates at an operating speed of 1-Gbps or 2-Gbps.

## Configuring SD Port Frame Encapsulation

The **switchport encap eisl** command only applies to SD port interfaces. This command determines the frame format for all frames transmitted by the interface in SD port mode. If the encapsulation is set to EISL, all outgoing frames are transmitted in the EISL frame format, for all SPAN sources.

The **switchport encap eisl** command is disabled by default. If you enable encapsulation, all outgoing frames are encapsulated, and you will see a new line (Encapsulation is eisl) in the **show interface** *SD\_port\_interface* command output.

### Configuring Receive Data Field Size

You can configure the receive data field size for native Fibre Channel interfaces (but not for virtual Fibre Channel interfaces). If the default data field size is 2112 bytes, the frame length will be 2148 bytes.

To configure the receive data field size, perform this task:

#### SUMMARY STEPS

1. switch# **configuration terminal**
2. switch(config)# **interface fc** *slot/port*
3. switch(config-if)# **switchport fcrxbufsize 2000**

#### DETAILED STEPS

|               | Command or Action                                     | Purpose                                                                                                                                  |
|---------------|-------------------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Step 1</b> | switch# <b>configuration terminal</b>                 | Enters configuration mode.                                                                                                               |
| <b>Step 2</b> | switch(config)# <b>interface fc</b> <i>slot/port</i>  | Selects a Fibre Channel interface and enters interface configuration mode.                                                               |
| <b>Step 3</b> | switch(config-if)# <b>switchport fcrxbufsize 2000</b> | Reduces the data field size for the selected interface to 2000 bytes. The default is 2112 bytes and the range is from 256 to 2112 bytes. |

### Understanding Bit Error Thresholds

The bit error rate threshold is used by the switch to detect an increased error rate before performance degradation seriously affects traffic.

The bit errors can occur for the following reasons:

- Faulty or bad cable.
- Faulty or bad GBIC or SFP.
- GBIC or SFP is specified to operate at 1 Gbps but is used at 2 Gbps.
- GBIC or SFP is specified to operate at 2 Gbps but is used at 4 Gbps.
- Short haul cable is used for long haul or long haul cable is used for short haul.
- Momentary synchronization loss.
- Loose cable connection at one or both ends.
- Improper GBIC or SFP connection at one or both ends.

A bit error rate threshold is detected when 15 error bursts occur in a 5-minute period. By default, the switch disables the interface when the threshold is reached.

You can enter the **shutdown/no shutdown** command sequence to reenable the interface.

You can configure the switch to not disable an interface when the threshold is crossed.

**Note**

The switch generates a syslog message when bit error threshold events are detected, even if the interface is configured not to be disabled by bit-error threshold events.

To disable the bit error threshold for an interface, perform this task:

**SUMMARY STEPS**

1. switch# **configuration terminal**
2. switch(config)# **interface fc slot/port**
3. switch(config-if)# **switchport ignore bit-errors**
4. switch(config-if)# **no switchport ignore bit-errors**

**DETAILED STEPS**

|               | <b>Command or Action</b>                                  | <b>Purpose</b>                                                                     |
|---------------|-----------------------------------------------------------|------------------------------------------------------------------------------------|
| <b>Step 1</b> | switch# <b>configuration terminal</b>                     | Enters configuration mode.                                                         |
| <b>Step 2</b> | switch(config)# <b>interface fc slot/port</b>             | Selects a Fibre Channel interface and enters interface configuration mode.         |
| <b>Step 3</b> | switch(config-if)# <b>switchport ignore bit-errors</b>    | Prevents the detection of bit error threshold events from disabling the interface. |
| <b>Step 4</b> | switch(config-if)# <b>no switchport ignore bit-errors</b> | Prevents the detection of bit error threshold events from enabling the interface.  |

**Configuring Buffer-to-Buffer Credits**

To configure BB\_credits for a Fibre Channel interface, perform this task:

**SUMMARY STEPS**

1. switch# **configuration terminal**
2. switch(config)# **interface fc slot/port**
3. switch(config-if)# **switchport fcrxbbcredit default**
4. switch(config-if)# **switchport fcrxbbcredit 5**
5. switch(config-if)# **switchport fcrxbbcredit 5 mode E**
6. switch(config-if)# **switchport fcrxbbcredit 5 mode Fx**
7. switch(config-if)# **do show int fc slot/port**

**DETAILED STEPS**

|               | <b>Command or Action</b>                                    | <b>Purpose</b>                                                                                                                                                                                                                                                               |
|---------------|-------------------------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Step 1</b> | switch# <b>configuration terminal</b>                       | Enters configuration mode.                                                                                                                                                                                                                                                   |
| <b>Step 2</b> | switch(config)# <b>interface fc slot/port</b>               | Selects a Fibre Channel interface and enters interface configuration mode.                                                                                                                                                                                                   |
| <b>Step 3</b> | switch(config-if)# <b>switchport fcrxbbcredit default</b>   | Applies the default operational value to the selected interface. The operational value depends on the port mode. The default values are assigned based on the port capabilities.                                                                                             |
| <b>Step 4</b> | switch(config-if)# <b>switchport fcrxbbcredit 5</b>         | Assigns a BB_credit of 5 to the selected interface. The range to assign BB_credits is between 1 and 64.                                                                                                                                                                      |
| <b>Step 5</b> | switch(config-if)# <b>switchport fcrxbbcredit 5 mode E</b>  | Assigns this value if the port is operating in E or TE mode. The range to assign BB_credits is between 1 and 64.                                                                                                                                                             |
| <b>Step 6</b> | switch(config-if)# <b>switchport fcrxbbcredit 5 mode Fx</b> | Assigns this value if the port is operating in F mode. The range to assign BB_credits is between 1 and 64.                                                                                                                                                                   |
| <b>Step 7</b> | switch(config-if)# <b>do show int fc slot/port</b>          | Displays the receive and transmit BB_credit along with other pertinent interface information for this interface.<br><br><b>Note</b> The BB_credit values are correct at the time the registers are read. They are useful to verify situations when the data traffic is slow. |

**Configuring Global Attributes for Fibre Channel Interfaces****Configuring Switch Port Attribute Default Values**

You can configure attribute default values for various switch port attributes. These attributes will be applied globally to all future switch port configurations, even if you do not individually specify them at that time.

To configure switch port attributes, perform this task:

## SUMMARY STEPS

1. switch# **configuration terminal**
2. switch(config)# **no system default switchport shutdown san**
3. switch(config)# **system default switchport shutdown san**
4. switch(config)# **system default switchport trunk mode auto**

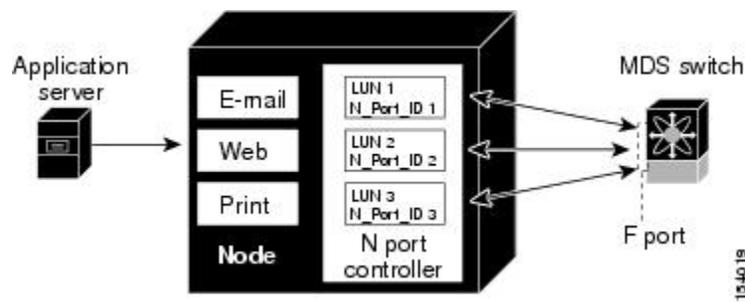
## DETAILED STEPS

|               | Command or Action                                                | Purpose                                                                                                                                                                                                                                                    |
|---------------|------------------------------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Step 1</b> | switch# <b>configuration terminal</b>                            | Enters configuration mode.                                                                                                                                                                                                                                 |
| <b>Step 2</b> | switch(config)# <b>no system default switchport shutdown san</b> | Configures the default setting for administrative state of an interface as Up. (The factory default setting is Down).<br><br><b>Tip</b> This command is applicable only to interfaces for which no user configuration exists for the administrative state. |
| <b>Step 3</b> | switch(config)# <b>system default switchport shutdown san</b>    | Configures the default setting for administrative state of an interface as Down. This is the factory default setting.<br><br><b>Tip</b> This command is applicable only to interfaces for which no user configuration exists for the administrative state. |
| <b>Step 4</b> | switch(config)# <b>system default switchport trunk mode auto</b> | Configures the default setting for administrative trunk mode state of an interface as Auto.<br><br><b>Note</b> The default setting is trunk mode on.                                                                                                       |

### About N Port Identifier Virtualization

N port identifier virtualization (NPIV) provides a means to assign multiple FC IDs to a single N port. This feature allows multiple applications on the N port to use different identifiers and allows access control, zoning, and port security to be implemented at the application level. The following figure shows an example application using NPIV.

**Figure 42: NPIV Example**



## Enabling N Port Identifier Virtualization

To enable or disable NPIV on the switch, perform this task:

### Before You Begin

You must globally enable NPIV for all VSANs on the switch to allow the NPIV-enabled applications to use multiple N port identifiers.



#### Note

All of the N port identifiers are allocated in the same VSAN.

## SUMMARY STEPS

1. switch# **configuration terminal**
2. switch(config)# **feature npiv**
3. switch(config)# **no npiv enable**

## DETAILED STEPS

|               | Command or Action                     | Purpose                                   |
|---------------|---------------------------------------|-------------------------------------------|
| <b>Step 1</b> | switch# <b>configuration terminal</b> | Enters configuration mode.                |
| <b>Step 2</b> | switch(config)# <b>feature npiv</b>   | Enables NPIV for all VSANs on the switch. |
| <b>Step 3</b> | switch(config)# <b>no npiv enable</b> | Disables (default) NPIV on the switch.    |

## Verifying Fibre Channel Interfaces

### Verifying SFP Transmitter Types

The SFP transmitter type can be displayed for a physical Fibre Channel interface (but not for a virtual Fibre Channel).

The small form-factor pluggable (SFP) hardware transmitters are identified by their acronyms when displayed in the **show interface brief** command. If the related SFP has a Cisco-assigned extended ID, then the **show interface** and **show interface brief** commands display the ID instead of the transmitter type. The **show interface transceiver** command and the **show interface fc slot/port** transceiver command display both values for Cisco supported SFPs.

### Verifying Interface Information

The **show interface** command displays interface configurations. If no arguments are provided, this command displays the information for all the configured interfaces in the switch.

You can also specify arguments (a range of interfaces or multiple, specified interfaces) to display interface information. You can specify a range of interfaces by entering a command with the following example format:

```
interface fc2/1 - 4 , fc3/2 - 3
```

The following example shows how to display all interfaces:

```
switch# show interface
fc3/1 is up
...
fc3/3 is up
...
Ethernet1/3 is up
...
mgmt0 is up
...
vethernet1/1 is up
...
vfc 1 is up
...
```

The following example shows how to display multiple specified interfaces:

```
switch# show interface fc3/1 , fc3/3
fc3/1 is up
...
fc3/3 is up
...
```

The following example shows how to display a specific interface:

```
switch# show interface vfc 1
vfc 1 is up
...
```

The following example shows how to display interface descriptions:

```
switch# show interface description
```

```

Interface Description

fc3/1 test intest
Ethernet1/1 --
vfc 1 --
...
```

The following example shows how to display all interfaces in brief:

```
switch# show interface brief
```

The following example shows how to display interface counters:

```
switch# show interface counters
```

The following example shows how to display transceiver information for a specific interface:

```
switch# show interface fc3/1 transceiver
```



#### Note

The **show interface transceiver** command is only valid if the SFP is present.

The **show running-configuration** command displays the entire running configuration with information for all interfaces. The interfaces have multiple entries in the configuration files to ensure that the interface configuration commands execute in the correct order when the switch reloads. If you display the running configuration for a specific interface, all the configuration commands for that interface are grouped together.

The following example shows the interface display when showing the running configuration for all interfaces:

```
switch# show running configuration
...
interface fc3/5
 switchport speed 2000
...
interface fc3/5
 switchport mode E

...
interface fc3/5
 channel-group 11 force
 no shutdown
```

The following example shows the interface display when showing the running configuration for a specific interface:

```
switch# show running configuration fc3/5
interface fc3/5
 switchport speed 2000
 switchport mode E
 channel-group 11 force
 no shutdown
```

### Verifying BB\_Credit Information

The following example shows how to display the BB\_credit information for all Fibre Channel interfaces:

```
switch# show interface bccredit
...
fc2/3 is trunking
 Transmit B2B Credit is 255
 Receive B2B Credit is 12
 Receive B2B Credit performance buffers is 375
 12 receive B2B credit remaining
 255 transmit B2B credit remaining
```

## Default Fibre Channel Interface Settings

The following table lists the default settings for native Fibre Channel interface parameters.

**Table 63: Default Native Fibre Channel Interface Parameters**

| Parameters           | Default                                        |
|----------------------|------------------------------------------------|
| Interface mode       | Auto                                           |
| Interface speed      | Auto                                           |
| Administrative state | Shutdown (unless changed during initial setup) |
| Trunk mode           | On (unless changed during initial setup)       |
| Trunk-allowed VSANs  | 1 to 4093                                      |
| Interface VSAN       | Default VSAN (1)                               |
| Beacon mode          | Off (disabled)                                 |
| EISL encapsulation   | Disabled                                       |

| Parameters      | Default    |
|-----------------|------------|
| Data field size | 2112 bytes |

The following table lists the default settings for virtual Fibre Channel interface parameters.

**Table 64: Default Virtual Fibre Channel Interface Parameters**

| Parameters           | Default                                        |
|----------------------|------------------------------------------------|
| Interface mode       | Auto                                           |
| Interface speed      | n/a                                            |
| Administrative state | Shutdown (unless changed during initial setup) |
| Trunk mode           | n/a                                            |
| Trunk-allowed VSANs  | n/a                                            |
| Interface VSAN       | Default VSAN (1)                               |
| EISL encapsulation   | n/a                                            |
| Data field size      | n/a                                            |





## Configuring Domain Parameters

This chapter contains the following sections:

- [Configuring Domain Parameters, page 531](#)

### Configuring Domain Parameters

The Fibre Channel domain (fcdomain) feature performs principal switch selection, domain ID distribution, FC ID allocation, and fabric reconfiguration functions as described in the FC-SW-2 standards. The domains are configured on a per-VSAN basis. If you do not configure a domain ID, the local switch uses a random ID.



#### Caution

Changes to fcdomain parameters should not be performed on a daily basis. These changes should be made by an administrator or individual who is completely familiar with switch operations.

When you change the configuration, be sure to save the running configuration. The next time you reboot the switch, the saved configuration is used. If you do not save the configuration, the previously saved startup configuration is used.

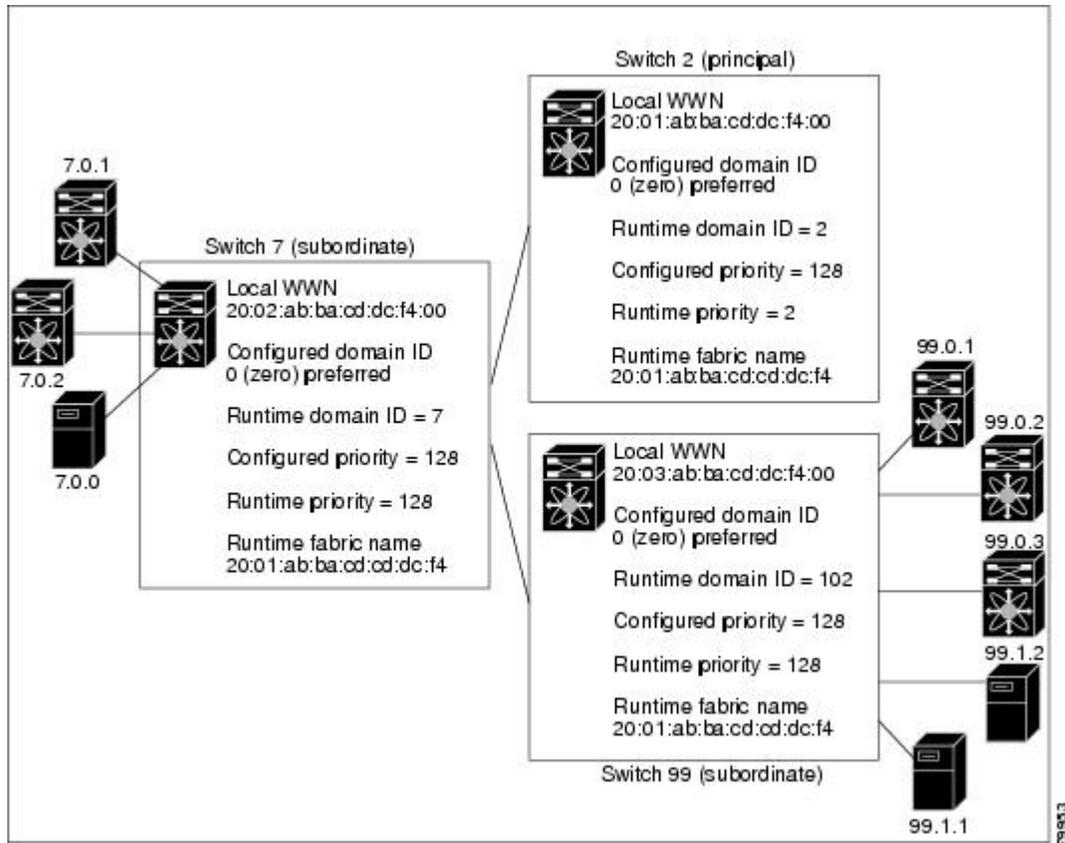
### Information About Fibre Channel Domains

This section describes each fcdomain phase:

- Principal switch selection—This phase guarantees the selection of a unique principal switch across the fabric.
- Domain ID distribution—This phase guarantees each switch in the fabric obtains a unique domain ID.
- FC ID allocation—This phase guarantees a unique FC ID assignment to each device attached to the corresponding switch in the fabric.
- Fabric reconfiguration—This phase guarantees a resynchronization of all switches in the fabric to ensure they simultaneously restart a new principal switch selection phase.

The following figure illustrates an example fcdomain configuration.

**Figure 43: Sample fcdomain Configuration**



**About Domain Restart**

Fibre Channel domains can be started disruptively or nondisruptively. If you perform a disruptive restart, reconfigure fabric (RCF) frames are sent to other switches in the fabric and data traffic is disrupted on all the switches in the VSAN (including remotely segmented ISLs). If you perform a nondisruptive restart, build fabric (BF) frames are sent to other switches in the fabric and data traffic is disrupted only on the switch.

If you are attempting to resolve a domain ID conflict, you must manually assign domain IDs. A disruptive restart is required to apply most configuration changes, including manually assigned domain IDs. Nondisruptive domain restarts are acceptable only when changing a preferred domain ID into a static one (and the actual domain ID remains the same).



**Note**

A static domain is specifically configured by the user and may be different from the runtime domain. If the domain IDs are different, the runtime domain ID changes to take on the static domain ID after the next restart, either disruptively or nondisruptively.

If a VSAN is in interop mode, you cannot disruptively restart the fcdomain for that VSAN.

You can apply most of the configurations to their corresponding runtime values. Each of the following sections provide further details on how the fcdomain parameters are applied to the runtime values.

The **fcdomain restart** command applies your changes to the runtime settings. Use the disruptive option to apply most of the configurations to their corresponding runtime values, including preferred domain IDs.

### Restarting a Domain

To restart the fabric disruptively or nondisruptively, perform this task:

#### SUMMARY STEPS

1. switch# **configuration terminal**
2. switch(config)# **fcdomain restart vsan vsan-id**
3. switch(config)# **fcdomain restart disruptive vsan vsan-id**

#### DETAILED STEPS

|               | Command or Action                                               | Purpose                                                      |
|---------------|-----------------------------------------------------------------|--------------------------------------------------------------|
| <b>Step 1</b> | switch# <b>configuration terminal</b>                           | Enters configuration mode.                                   |
| <b>Step 2</b> | switch(config)# <b>fcdomain restart vsan vsan-id</b>            | Forces the VSAN to reconfigure without traffic disruption.   |
| <b>Step 3</b> | switch(config)# <b>fcdomain restart disruptive vsan vsan-id</b> | Forces the VSAN to reconfigure with data traffic disruption. |

### About Domain Manager Fast Restart

When a principal link fails, the domain manager must select a new principal link. By default, the domain manager starts a build fabric (BF) phase, followed by a principal switch selection phase. Both of these phases involve all the switches in the VSAN, and together take at least 15 seconds to complete. To reduce the time required for the domain manager to select a new principal link, you can enable the domain manager fast restart feature.

When fast restart is enabled and a backup link is available, the domain manager needs only a few milliseconds to select a new principal link to replace the one that failed. Also, the reconfiguration required to select the new principal link only affects the two switches that are directly attached to the failed link, not the entire VSAN. When a backup link is not available, the domain manager reverts to the default behavior and starts a BF phase, followed by a principal switch selection phase. The fast restart feature can be used in any interoperability mode.

### Enabling Domain Manager Fast Restart

To enable the domain manager fast restart feature, perform this task:

**SUMMARY STEPS**

1. switch# **configuration terminal**
2. switch(config)# **fcdomain optimize fast-restart vsan** *vsan-id*
3. switch(config)# **fcdomain optimize fast-restart vsan** *vsan-id - vsan-id*
4. switch(config)# **no fcdomain optimize fast-restart vsan** *vsan-id*

**DETAILED STEPS**

|               | <b>Command or Action</b>                                                            | <b>Purpose</b>                                                        |
|---------------|-------------------------------------------------------------------------------------|-----------------------------------------------------------------------|
| <b>Step 1</b> | switch# <b>configuration terminal</b>                                               | Enters configuration mode.                                            |
| <b>Step 2</b> | switch(config)# <b>fcdomain optimize fast-restart vsan</b> <i>vsan-id</i>           | Enables domain manager fast restart in the specified VSAN.            |
| <b>Step 3</b> | switch(config)# <b>fcdomain optimize fast-restart vsan</b> <i>vsan-id - vsan-id</i> | Enables domain manager fast restart in the specified range of VSANs.  |
| <b>Step 4</b> | switch(config)# <b>no fcdomain optimize fast-restart vsan</b> <i>vsan-id</i>        | Disables (default) domain manager fast restart in the specified VSAN. |

**About Switch Priority**

By default, the configured priority is 128. The valid range to set the priority is between 1 and 254. Priority 1 has the highest priority. Value 255 is accepted from other switches, but cannot be locally configured.

Any new switch cannot become the principal switch when it joins a stable fabric. During the principal switch selection phase, the switch with the highest priority becomes the principal switch. If two switches have the same configured priority, the switch with the lower world-wide name (WWN) becomes the principal switch.

The priority configuration is applied to runtime when the fcdomain is restarted. This configuration is applicable to both disruptive and nondisruptive restarts.

**Configuring Switch Priority**

To configure the priority for the principal switch, perform this task:

**SUMMARY STEPS**

1. switch# **configuration terminal**
2. switch(config)# **fcdomain priority** *number* **VSAN** *vsan-id*
3. switch(config)# **no fcdomain priority** *number* **VSAN** *vsan-id*

**DETAILED STEPS**

|               | <b>Command or Action</b>                                                             | <b>Purpose</b>                                                                |
|---------------|--------------------------------------------------------------------------------------|-------------------------------------------------------------------------------|
| <b>Step 1</b> | switch# <b>configuration terminal</b>                                                | Enters configuration mode.                                                    |
| <b>Step 2</b> | switch(config)# <b>fcdomain priority</b> <i>number</i> <b>VSAN</b> <i>vsan-id</i>    | Configures the specified priority for the local switch in the specified VSAN. |
| <b>Step 3</b> | switch(config)# <b>no fcdomain priority</b> <i>number</i> <b>VSAN</b> <i>vsan-id</i> | Reverts the priority to the factory default (128) in the specified VSAN.      |

**About fcdomain Initiation**

By default, the fcdomain feature is enabled on each switch. If you disable the fcdomain feature in a switch, that switch can no longer participate with other switches in the fabric. The fcdomain configuration is applied to runtime through a disruptive restart.

**Disabling or Reenabling fcdomains**

To disable or reenable fcdomains in a single VSAN or a range of VSANs, perform this task:

**SUMMARY STEPS**

1. switch# **configuration terminal**
2. switch(config)# **no fcdomain vsan** *vsan-id* - *vsan-id*
3. switch(config)# **fcdomain vsan** *vsan-id*

**DETAILED STEPS**

|               | <b>Command or Action</b>                                                | <b>Purpose</b>                                                   |
|---------------|-------------------------------------------------------------------------|------------------------------------------------------------------|
| <b>Step 1</b> | switch# <b>configuration terminal</b>                                   | Enters configuration mode.                                       |
| <b>Step 2</b> | switch(config)# <b>no fcdomain vsan</b> <i>vsan-id</i> - <i>vsan-id</i> | Disables the fcdomain configuration in the specified VSAN range. |
| <b>Step 3</b> | switch(config)# <b>fcdomain vsan</b> <i>vsan-id</i>                     | Enables the fcdomain configuration in the specified VSAN.        |

**Configuring Fabric Names**

To set the fabric name value for a disabled fcdomain, perform this task:

**SUMMARY STEPS**

1. switch# **configuration terminal**
2. switch(config)# **fcdomain fabric-name 20:1:ac:16:5e:0:21:01 vsan vsan-id**
3. switch(config)# **no fcdomain fabric-name 20:1:ac:16:5e:0:21:01 vsan vsan-id**

**DETAILED STEPS**

|               | <b>Command or Action</b>                                                          | <b>Purpose</b>                                                                               |
|---------------|-----------------------------------------------------------------------------------|----------------------------------------------------------------------------------------------|
| <b>Step 1</b> | switch# <b>configuration terminal</b>                                             | Enters configuration mode.                                                                   |
| <b>Step 2</b> | switch(config)# <b>fcdomain fabric-name 20:1:ac:16:5e:0:21:01 vsan vsan-id</b>    | Assigns the configured fabric name value in the specified VSAN.                              |
| <b>Step 3</b> | switch(config)# <b>no fcdomain fabric-name 20:1:ac:16:5e:0:21:01 vsan vsan-id</b> | Changes the fabric name value to the factory default (20:01:00:05:30:00:28:df) in VSAN 3010. |

**About Incoming RCFs**

You can configure the rcf-reject option on a per-interface, per-VSAN basis. By default, the rcf-reject option is disabled (that is, RCF request frames are not automatically rejected).

The rcf-reject option takes effect immediately.

No fcdomain restart is required.

**Note**

You do not need to configure the RFC reject option on virtual Fibre Channel interfaces, because these interfaces operate only in F port mode.

**Rejecting Incoming RCFs**

To reject incoming RCF request frames, perform this task:

**SUMMARY STEPS**

1. switch# **configuration terminal**
2. switch(config)# **interface fc slot/port**
3. switch(config-if)# **fcdomain rcf-reject vsan vsan-id**
4. switch(config-if)# **no fcdomain rcf-reject vsan vsan-id**

## DETAILED STEPS

|               | Command or Action                                             | Purpose                                                                             |
|---------------|---------------------------------------------------------------|-------------------------------------------------------------------------------------|
| <b>Step 1</b> | switch# <b>configuration terminal</b>                         | Enters configuration mode.                                                          |
| <b>Step 2</b> | switch(config)# <b>interface fc slot/port</b>                 | Configures the specified interface.                                                 |
| <b>Step 3</b> | switch(config-if)# <b>fcdomain rcf-reject vsan vsan-id</b>    | Enables the RCF filter on the specified interface in the specified VSAN.            |
| <b>Step 4</b> | switch(config-if)# <b>no fcdomain rcf-reject vsan vsan-id</b> | Disables (default) the RCF filter on the specified interface in the specified VSAN. |

### About Autoreconfiguring Merged Fabrics

By default, the autoreconfigure option is disabled. When you join two switches belonging to two different stable fabrics that have overlapping domains, the following situations can occur:

- If the autoreconfigure option is enabled on both switches, a disruptive reconfiguration phase is started.
- If the autoreconfigure option is disabled on either or both switches, the links between the two switches become isolated.

The autoreconfigure option takes immediate effect at runtime. You do not need to restart the fcdomain. If a domain is currently isolated due to domain overlap, and you later enable the autoreconfigure option on both switches, the fabric continues to be isolated. If you enabled the autoreconfigure option on both switches before connecting the fabric, a disruptive reconfiguration (RCF) will occur. A disruptive reconfiguration may affect data traffic. You can nondisruptively reconfigure the fcdomain by changing the configured domains on the overlapping links and eliminating the domain overlap.

### Enabling Autoreconfiguration

To enable automatic reconfiguration in a specific VSAN (or range of VSANs), perform this task:

## SUMMARY STEPS

1. switch# **configuration terminal**
2. switch(config)# **fcdomain auto-reconfigure vsan vsan-id**
3. switch(config)# **no fcdomain auto-reconfigure vsan vsan-id**

## DETAILED STEPS

|               | Command or Action                     | Purpose                    |
|---------------|---------------------------------------|----------------------------|
| <b>Step 1</b> | switch# <b>configuration terminal</b> | Enters configuration mode. |

|               | Command or Action                                                         | Purpose                                                                                                    |
|---------------|---------------------------------------------------------------------------|------------------------------------------------------------------------------------------------------------|
| <b>Step 2</b> | switch(config)# <b>fdomain auto-reconfigure vsan</b><br><i>vsan-id</i>    | Enables the automatic reconfiguration option in the specified VSAN.                                        |
| <b>Step 3</b> | switch(config)# <b>no fdomain auto-reconfigure vsan</b><br><i>vsan-id</i> | Disables the automatic reconfiguration option and reverts it to the factory default in the specified VSAN. |

## Domain IDs

Domain IDs uniquely identify a switch in a VSAN. A switch may have different domain IDs in different VSANs. The domain ID is part of the overall FC ID.

### About Domain IDs

The configured domain ID can be preferred or static. By default, the configured domain ID is 0 (zero) and the configured type is preferred.



#### Note

The 0 (zero) value can be configured only if you use the preferred option.

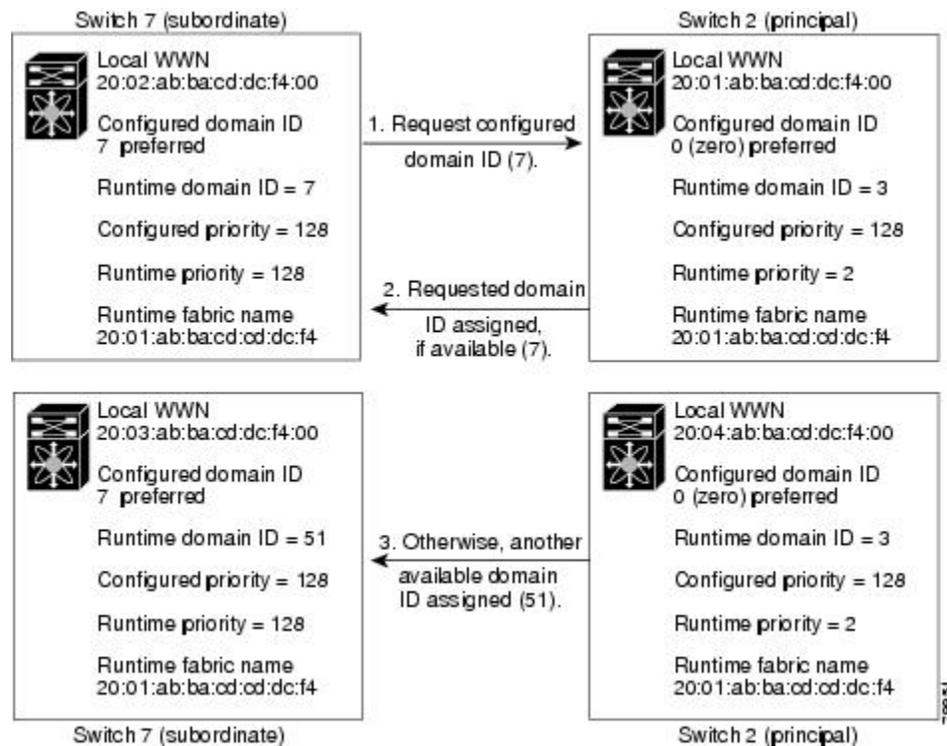
If you do not configure a domain ID, the local switch sends a random ID in its request. We recommend that you use static domain IDs.

When a subordinate switch requests a domain, the following process takes place (see the figure below):

- The local switch sends a configured domain ID request to the principal switch.

- The principal switch assigns the requested domain ID if available. Otherwise, it assigns another available domain ID.

**Figure 44: Configuration Process Using the Preferred Option**



The operation of a subordinate switch changes based on three factors:

- The allowed domain ID lists.
- The configured domain ID.
- The domain ID that the principal switch has assigned to the requesting switch.

In specific situations, the changes are as follows:

- When the received domain ID is not within the allowed list, the requested domain ID becomes the runtime domain ID and all interfaces on that VSAN are isolated.
- When the assigned and requested domain IDs are the same, the preferred and static options are not relevant, and the assigned domain ID becomes the runtime domain ID.
- When the assigned and requested domain IDs are different, the following cases apply:
  - If the configured type is static, the assigned domain ID is discarded, all local interfaces are isolated, and the local switch assigns itself the configured domain ID, which becomes the runtime domain ID.
  - If the configured type is preferred, the local switch accepts the domain ID assigned by the principal switch and the assigned domain ID becomes the runtime domain ID.

If you change the configured domain ID, the change is only accepted if the new domain ID is included in all the allowed domain ID lists currently configured in the VSAN. Alternatively, you can also configure zero-preferred domain ID.



**Caution**

You must enter the `fcdomain restart` command if you want to apply the configured domain changes to the runtime domain.



**Note**

If you have configured an allow domain ID list, the domain IDs that you add must be in that range for the VSAN.

**Related Topics**

[About Allowed Domain ID Lists, on page 541](#)

**Specifying Static or Preferred Domain IDs**

When you assign a static domain ID type, you are requesting a particular domain ID. If the switch does not obtain the requested address, it will isolate itself from the fabric. When you specify a preferred domain ID, you are also requesting a particular domain ID; however, if the requested domain ID is unavailable, then the switch will accept another domain ID.

While the static option can be applied at runtime after a disruptive or nondisruptive restart, the preferred option is applied at runtime only after a disruptive restart.



**Note**

Within a VSAN all switches should have the same domain ID type (either static or preferred). If a configuration is mixed (some switches with static domain types and others with preferred), you may experience link isolation.

To specify a static or preferred domain ID, perform this task:

**SUMMARY STEPS**

1. `switch# configuration terminal`
2. `switch(config)# fcdomain domain domain-id static vsan vsan-id`
3. `switch(config)# no fcdomain domain domain-id static vsan vsan-id`
4. `switch(config)# fcdomain domain domain-id preferred vsan vsan-id`
5. `switch(config)# no fcdomain domain domain-id preferred vsan vsan-id`

**DETAILED STEPS**

|        | Command or Action                           | Purpose                    |
|--------|---------------------------------------------|----------------------------|
| Step 1 | <code>switch# configuration terminal</code> | Enters configuration mode. |

|               | Command or Action                                                                                  | Purpose                                                                                                                                                                                          |
|---------------|----------------------------------------------------------------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Step 2</b> | switch(config)# <b>fcdomain domain</b> <i>domain-id</i><br><b>static vsan</b> <i>vsan-id</i>       | Configures the switch in the specified VSAN to accept only a specific value and moves the local interfaces in the specified VSAN to an isolated state if the requested domain ID is not granted. |
| <b>Step 3</b> | switch(config)# <b>no fcdomain domain</b> <i>domain-id</i><br><b>static vsan</b> <i>vsan-id</i>    | Resets the configured domain ID to factory defaults in the specified VSAN. The configured domain ID becomes 0 preferred.                                                                         |
| <b>Step 4</b> | switch(config)# <b>fcdomain domain</b> <i>domain-id</i><br><b>preferred vsan</b> <i>vsan-id</i>    | Configures the switch in the specified VSAN to request a preferred domain ID 3 and accepts any value assigned by the principal switch. The domain is range is 1 to 239.                          |
| <b>Step 5</b> | switch(config)# <b>no fcdomain domain</b> <i>domain-id</i><br><b>preferred vsan</b> <i>vsan-id</i> | Resets the configured domain ID to 0 (default) in the specified VSAN. The configured domain ID becomes 0 preferred.                                                                              |

### About Allowed Domain ID Lists

By default, the valid range for an assigned domain ID list is from 1 to 239. You can specify a list of ranges to be in the allowed domain ID list and separate each range with a comma. The principal switch assigns domain IDs that are available in the locally configured allowed domain list.

Use allowed domain ID lists to design your VSANs with nonoverlapping domain IDs. This helps you in the future if you need to implement IVR without the NAT feature.

If you configure an allowed list on one switch in the fabric, we recommend that you configure the same list in all other switches in the fabric to ensure consistency or use CFS to distribute the configuration.

An allowed domain ID list must satisfy the following conditions:

- If this switch is a principal switch, all the currently assigned domain IDs must be in the allowed list.
- If this switch is a subordinate switch, the local runtime domain ID must be in the allowed list.
- The locally configured domain ID of the switch must be in the allowed list.
- The intersection of the assigned domain IDs with other already configured domain ID lists must not be empty.

### Configuring Allowed Domain ID Lists

To configure the allowed domain ID list, perform this task:

#### SUMMARY STEPS

1. switch# **configuration terminal**
2. switch(config)# **fcdomain allowed** *domain-id range* **vsan** *vsan-id*
3. switch(config)# **no fcdomain allowed** *domain-id range* **vsan** *vsan-id*

## DETAILED STEPS

|               | Command or Action                                                                               | Purpose                                                                                         |
|---------------|-------------------------------------------------------------------------------------------------|-------------------------------------------------------------------------------------------------|
| <b>Step 1</b> | switch# <b>configuration terminal</b>                                                           | Enters configuration mode.                                                                      |
| <b>Step 2</b> | switch(config)# <b>fcdomain allowed</b> <i>domain-id range</i><br><b>vsan</b> <i>vsan-id</i>    | Configures the list to allow switches with the domain ID range in the specified VSAN.           |
| <b>Step 3</b> | switch(config)# <b>no fcdomain allowed</b> <i>domain-id range</i><br><b>vsan</b> <i>vsan-id</i> | Reverts to the factory default of allowing domain IDs from 1 through 239 in the specified VSAN. |

### About CFS Distribution of Allowed Domain ID Lists

You can enable the distribution of the allowed domain ID list configuration information to all Cisco SAN switches in the fabric using the Cisco Fabric Services (CFS) infrastructure. This feature allows you to synchronize the configuration across the fabric from the console of a single switch. Because the same configuration is distributed to the entire VSAN, you can avoid possible misconfiguration and the possibility that two switches in the same VSAN have configured incompatible allowed domains.

Use CFS to distribute the allowed domain ID list to ensure consistency in the allowed domain ID lists on all switches in the VSAN.



#### Note

We recommend configuring the allowed domain ID list and committing it on the principal switch.

### Related Topics

[Using Cisco Fabric Services, on page 329](#)

### Enabling Distribution

CFS distribution of allowed domain ID lists is disabled by default. You must enable distribution on all switches to which you want to distribute the allowed domain ID lists.

To enable (or disable) allowed domain ID list configuration distribution, perform this task:

## SUMMARY STEPS

1. switch# **configuration terminal**
2. switch(config)# **fcdomain distribute**
3. switch(config)# **no fcdomain distribute**

**DETAILED STEPS**

|               | <b>Command or Action</b>                      | <b>Purpose</b>                                        |
|---------------|-----------------------------------------------|-------------------------------------------------------|
| <b>Step 1</b> | switch# <b>configuration terminal</b>         | Enters configuration mode.                            |
| <b>Step 2</b> | switch(config)# <b>fcdomain distribute</b>    | Enables domain configuration distribution.            |
| <b>Step 3</b> | switch(config)# <b>no fcdomain distribute</b> | Disables (default) domain configuration distribution. |

**Locking the Fabric**

The first action that modifies the existing configuration creates the pending configuration and locks the feature in the fabric. After you lock the fabric, the following conditions apply:

- No other user can make any configuration changes to this feature.
- A pending configuration is created by copying the active configuration. Subsequent modifications are made to the pending configuration and remain there until you commit the changes to the active configuration (and other switches in the fabric) or discard them.

**Committing Changes**

To apply the pending domain configuration changes to other SAN switches in the VSAN, you must commit the changes. The pending configuration changes are distributed and, on a successful commit, the configuration changes are applied to the active configuration in the SAN switches throughout the VSAN and the fabric lock is released.

To commit pending domain configuration changes and release the lock, perform this task:

**SUMMARY STEPS**

1. switch# **configuration terminal**
2. switch(config)# **fcdomain commit vsan** *vsan-id*

**DETAILED STEPS**

|               | <b>Command or Action</b>                                   | <b>Purpose</b>                                    |
|---------------|------------------------------------------------------------|---------------------------------------------------|
| <b>Step 1</b> | switch# <b>configuration terminal</b>                      | Enters configuration mode.                        |
| <b>Step 2</b> | switch(config)# <b>fcdomain commit vsan</b> <i>vsan-id</i> | Commits the pending domain configuration changes. |

## Discarding Changes

At any time, you can discard the pending changes to the domain configuration and release the fabric lock. If you discard (abort) the pending changes, the configuration remains unaffected and the lock is released.

To discard pending domain configuration changes and release the lock, perform this task:

### SUMMARY STEPS

1. `switch# configuration terminal`
2. `switch(config)# fcdomain abort vsan vsan-id`

### DETAILED STEPS

|               | Command or Action                                        | Purpose                                            |
|---------------|----------------------------------------------------------|----------------------------------------------------|
| <b>Step 1</b> | <code>switch# configuration terminal</code>              | Enters configuration mode.                         |
| <b>Step 2</b> | <code>switch(config)# fcdomain abort vsan vsan-id</code> | Discards the pending domain configuration changes. |

## Clearing a Fabric Lock

If you have performed a domain configuration task and have not released the lock by either committing or discarding the changes, an administrator can release the lock from any switch in the fabric. If the administrator performs this task, your pending changes are discarded and the fabric lock is released.

The pending changes are only available in the volatile directory and are discarded if the switch is restarted.

To release a fabric lock, enter the `clear fcdomain session vsan` command in EXEC mode using a login ID that has administrative privileges.

```
switch# clear fcdomain session vsan 10
```

## Displaying CFS Distribution Status

You can display the status of CFS distribution for allowed domain ID lists using the `show fcdomain status` command.

```
switch# show fcdomain status
CFS distribution is enabled
```

## Displaying Pending Changes

You can display the pending configuration changes using the **show fcdomain pending** command.

```
switch# show fcdomain pending vsan 10
Pending Configured Allowed Domains

VSAN 10
Assigned or unallowed domain IDs: 1-9,24,100,231-239.
[User] configured allowed domain IDs: 10-230.
```

You can display the differences between the pending configuration and the current configuration using the **show fcdomain pending-diff** command.

```
switch# show fcdomain pending-diff vsan 10
Current Configured Allowed Domains

VSAN 10
Assigned or unallowed domain IDs: 24,100.
[User] configured allowed domain IDs: 1-239.
Pending Configured Allowed Domains

VSAN 10
Assigned or unallowed domain IDs: 1-9,24,100,231-239.
[User] configured allowed domain IDs: 10-230.
```

## Displaying Session Status

You can display the status of the distribution session using the **show fcdomain session-status vsan** command.

```
switch# show fcdomain session-status vsan 1
Last Action: Distribution Enable
Result: Success
```

## About Contiguous Domain ID Assignments

By default, the contiguous domain assignment is disabled. When a subordinate switch requests the principal switch for two or more domains and the domains are not contiguous, the following situations can occur:

- If the contiguous domain assignment is enabled in the principal switch, the principal switch locates contiguous domains and assigns them to the subordinate switches. If contiguous domains are not available, the switch software rejects this request.
- If the contiguous domain assignment is disabled in the principal switch, the principal switch assigns the available domains to the subordinate switch.

## Enabling Contiguous Domain ID Assignments

To enable contiguous domains in a specific VSAN (or a range of VSANs), perform this task:

### SUMMARY STEPS

1. switch# **configuration terminal**
2. switch(config)# **fcdomain contiguous-allocation vsan vsan-id - vsan-id**
3. switch(config)# **no fcdomain contiguous-allocation vsan vsan-id**

## DETAILED STEPS

|               | Command or Action                                                            | Purpose                                                                                                                                                                                                      |
|---------------|------------------------------------------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Step 1</b> | switch# <b>configuration terminal</b>                                        | Enters configuration mode.                                                                                                                                                                                   |
| <b>Step 2</b> | switch(config)# <b>fcdomain contiguous-allocation vsan vsan-id - vsan-id</b> | Enables the contiguous allocation option in the specified VSAN range.<br><br><b>Note</b> The <b>contiguous-allocation</b> option takes immediate effect at runtime. You do not need to restart the fcdomain. |
| <b>Step 3</b> | switch(config)# <b>no fcdomain contiguous-allocation vsan vsan-id</b>        | Disables the contiguous allocation option and reverts it to the factory default in the specified VSAN.                                                                                                       |

## FC IDs

When an N port logs into a Cisco Nexus 5000 Series switch, it is assigned an FC ID. By default, the persistent FC ID feature is enabled. If this feature is disabled, the following situations can occur:

- An N port logs into a Cisco Nexus 5000 Series switch. The WWN of the requesting N port and the assigned FC ID are retained and stored in a volatile cache. The contents of this volatile cache are not saved across reboots.
- The switch is designed to preserve the binding FC ID to the WWN on a best-effort basis. For example, if one N port disconnects from the switch and its FC ID is requested by another device, this request is granted and the WWN with the initial FC ID association is released.
- The volatile cache stores up to 4000 entries of WWN to FC ID binding. If this cache is full, a new (more recent) entry overwrites the oldest entry in the cache. In this case, the corresponding WWN to FC ID association for the oldest entry is lost.
- N ports receive the same FC IDs if disconnected and reconnected to any port within the same switch (as long as it belongs to the same VSAN).

### About Persistent FC IDs

When persistent FC IDs are enabled, the following occurs:

- The current FC IDs in use in the fcdomain are saved across reboots.
- The fcdomain automatically populates the database with dynamic entries that the switch has learned about after a device (host or disk) is plugged into a port interface.



#### Note

If you connect to the switch from an AIX or HP-UX host, be sure to enable the persistent FC ID feature in the VSAN that connects these hosts.

**Note**

When persistent FC IDs are enabled, FC IDs cannot be changed after a reboot. FC IDs are enabled by default, but can be disabled for each VSAN.

A persistent FC ID assigned to an F port can be moved across interfaces and can continue to maintain the same persistent FC ID.

**Enabling the Persistent FC ID Feature**

To enable the persistent FC ID feature, perform this task:

**SUMMARY STEPS**

1. switch# **configuration terminal**
2. switch(config)# **fcdomain fcid persistent vsan** *vsan-id*
3. switch(config)# **no fcdomain fcid persistent vsan** *vsan-id*

**DETAILED STEPS**

|               | <b>Command or Action</b>                                               | <b>Purpose</b>                                                   |
|---------------|------------------------------------------------------------------------|------------------------------------------------------------------|
| <b>Step 1</b> | switch# <b>configuration terminal</b>                                  | Enters configuration mode.                                       |
| <b>Step 2</b> | switch(config)# <b>fcdomain fcid persistent vsan</b> <i>vsan-id</i>    | Activates (default) persistency of FC IDs in the specified VSAN. |
| <b>Step 3</b> | switch(config)# <b>no fcdomain fcid persistent vsan</b> <i>vsan-id</i> | Disables the FC ID persistency feature in the specified VSAN.    |

**Persistent FC ID Configuration Guidelines**

When the persistent FC ID feature is enabled, you can enter the persistent FC ID submode and add static or dynamic entries in the FC ID database. By default, all added entries are static. Persistent FC IDs are configured on a per-VSAN basis.

When manually configuring a persistent FC ID, follow these requirements:

- Ensure that the persistent FC ID feature is enabled in the required VSAN.
- Ensure that the required VSAN is an active VSAN. Persistent FC IDs can only be configured on active VSANs.
- Verify that the domain part of the FC ID is the same as the runtime domain ID in the required VSAN. If the software detects a domain mismatch, the command is rejected.
- Verify that the port field of the FC ID is 0 (zero) when configuring an area.

## Configuring Persistent FC IDs

To configure persistent FC IDs, perform this task:

### SUMMARY STEPS

1. switch# **configuration terminal**
2. switch(config)# **fcdomain fcid database**
3. switch(config-fcid-db)# **vsan vsan-id wwn 33:e8:00:05:30:00:16:df fcid fcid**
4. switch(config-fcid-db)# **vsan vsan-id wwn 11:22:11:22:33:44:33:44 fcid fcid dynamic**
5. switch(config-fcid-db)# **vsan vsan-id wwn 11:22:11:22:33:44:33:44 fcid fcid area**

### DETAILED STEPS

|               | Command or Action                                                                         | Purpose                                                                                                                                                                                                                                              |
|---------------|-------------------------------------------------------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Step 1</b> | switch# <b>configuration terminal</b>                                                     | Enters configuration mode.                                                                                                                                                                                                                           |
| <b>Step 2</b> | switch(config)# <b>fcdomain fcid database</b>                                             | Enters FC ID database configuration submode.                                                                                                                                                                                                         |
| <b>Step 3</b> | switch(config-fcid-db)# <b>vsan vsan-id wwn 33:e8:00:05:30:00:16:df fcid fcid</b>         | Configures a device WWN (33:e8:00:05:30:00:16:df) with the FC ID 0x070128 in the specified VSAN.<br><br><b>Note</b> To avoid assigning a duplicate FC ID, use the <b>show fcdomain address-allocation vsan</b> command to display the FC IDs in use. |
| <b>Step 4</b> | switch(config-fcid-db)# <b>vsan vsan-id wwn 11:22:11:22:33:44:33:44 fcid fcid dynamic</b> | Configures a device WWN (11:22:11:22:33:44:33:44) with the FC ID 0x070123 in the specified VSAN in dynamic mode.                                                                                                                                     |
| <b>Step 5</b> | switch(config-fcid-db)# <b>vsan vsan-id wwn 11:22:11:22:33:44:33:44 fcid fcid area</b>    | Configures a device WWN (11:22:11:22:33:44:33:44) with the FC IDs 0x070100 through 0x701FF in the specified VSAN.<br><br><b>Note</b> To secure the entire area for this fcdomain, assign 00 as the last two characters of the FC ID.                 |

### About Unique Area FC IDs for HBAs



**Note** Only read this section if the Host Bus Adapter (HBA) port and the storage port are connected to the same switch.

Some HBA ports require a different area ID than for the storage ports when they are both connected to the same switch. For example, if the storage port FC ID is 0x6f7704, the area for this port is 77. In this case, the HBA port's area can be anything other than 77. The HBA port's FC ID must be manually configured to be different from the storage port's FC ID.

Cisco Nexus 5000 Series switches facilitate this requirement with the FC ID persistence feature. You can use this feature to preassign an FC ID with a different area to either the storage port or the HBA port.

## Configuring Unique Area FC IDs for an HBA

The following task uses an example configuration with a switch domain of 111(6f hex). The server connects to the switch over FCoE. The HBA port connects to interface vfc20/1 and the storage port connects to interface fc2/3 on the same switch.

To configure a different area ID for the HBA port, perform this task:

### SUMMARY STEPS

1. Obtain the port WWN (Port Name field) ID of the HBA using the **show flogi database** command.
2. Shut down the HBA interface in the Cisco Nexus 5000 Series switch.
3. Verify that the FC ID feature is enabled using the **show fcdomain vsan** command.
4. Enable the persistent FC ID feature in the Cisco Nexus 5000 Series switch.
5. Assign a new FC ID with a different area allocation. In this example, replace 77 with *ee*.
6. Enable the HBA interface in the Cisco Nexus 5000 Series switch.
7. Verify the pWWN ID of the HBA by using the **show flogi database** command.

### DETAILED STEPS

**Step 1** Obtain the port WWN (Port Name field) ID of the HBA using the **show flogi database** command.

```
switch# show flogi database
```

```

INTERFACE VSAN FCID PORT NAME NODE NAME

vfc10/1 3 0x6f7703 50:05:08:b2:00:71:c8:c2 50:05:08:b2:00:71:c8:c0
fc2/3 3 0x6f7704 50:06:0e:80:03:29:61:0f 50:06:0e:80:03:29:61:0f
```

**Note** Both FC IDs in this setup have the same area 77 assignment.

**Step 2** Shut down the HBA interface in the Cisco Nexus 5000 Series switch.

```
switch# configuration terminal
switch(config)# interface vfc20/1
```

```
switch(config-if)# shutdown
```

```
switch(config-if)# end
```

**Step 3** Verify that the FC ID feature is enabled using the **show fcdomain vsan** command.

```
switch# show fcdomain vsan 1
```

```
...
Local switch configuration information:
 State: Enabled
 FCID persistence: Disabled
```

If this feature is disabled, continue to the next step to enable the persistent FC ID.

If this feature is already enabled, skip to the following step.

**Step 4** Enable the persistent FC ID feature in the Cisco Nexus 5000 Series switch.

```
switch# configuration terminal
switch(config)# fcdomain fcid persistent vsan 1
switch(config)# end
```

**Step 5** Assign a new FC ID with a different area allocation. In this example, replace 77 with ee.

```
switch# configuration terminal
switch(config)# fcdomain fcid database
switch(config-fcid-db)# vsan 3 wwn 50:05:08:b2:00:71:c8:c2 fcid 0x6fee00 area
```

**Step 6** Enable the HBA interface in the Cisco Nexus 5000 Series switch.

```
switch# configuration terminal
switch(config)# interface vfc20/1
switch(config-if)# no shutdown
```

```
switch(config-if)# end
```

**Step 7** Verify the pWWN ID of the HBA by using the **show flogi database** command.

```
switch# show flogi database
```

```

INTERFACE VSAN FCID PORT NAME NODE NAME

vfc20/1 3 0x6fee00 50:05:08:b2:00:71:c8:c2 50:05:08:b2:00:71:c8:c0
fc2/3 3 0x6f7704 50:06:0e:80:03:29:61:0f 50:06:0e:80:03:29:61:0f
```

**Note** Both FC IDs now have different area assignments.

### About Persistent FC ID Selective Purging

Persistent FC IDs can be purged selectively. Static entries and FC IDs currently in use cannot be deleted. The table below identifies the FC ID entries that are deleted or retained when persistent FC IDs are purged.

**Table 65: Purged FC IDs**

| Persistent FC ID state | Persistent Usage State | Action      |
|------------------------|------------------------|-------------|
| Static                 | In use                 | Not deleted |
| Static                 | Not in use             | Not deleted |
| Dynamic                | In use                 | Not deleted |
| Dynamic                | Not in use             | Deleted     |

### Purging Persistent FC IDs

To purge persistent FC IDs, perform this task:

## SUMMARY STEPS

1. switch# **purge fcdomain fcid vsan** *vsan-id*
2. switch# **purge fcdomain fcid vsan** *vsan-id - vsan-id*

## DETAILED STEPS

|        | Command or Action                                                | Purpose                                                       |
|--------|------------------------------------------------------------------|---------------------------------------------------------------|
| Step 1 | switch# <b>purge fcdomain fcid vsan</b> <i>vsan-id</i>           | Purges all dynamic and unused FC IDs in the specified VSAN.   |
| Step 2 | switch# <b>purge fcdomain fcid vsan</b> <i>vsan-id - vsan-id</i> | Purges dynamic and unused FC IDs in the specified VSAN range. |

## Verifying fcdomain Information



### Note

If the fcdomain feature is disabled, the runtime fabric name in the display is the same as the configured fabric name.

This example shows how to display information about fcdomain configurations:

```
switch# show fcdomain vsan 2
```

Use the **show fcdomain domain-list** command to display the list of domain IDs of all switches belonging to a specified VSAN. This list provides the WWN of the switches owning each domain ID. The next example uses the following values:

- A switch with WWN of 20:01:00:05:30:00:47:df is the principal switch and has domain 200.
- A switch with WWN of 20:01:00:0d:ec:08:60:c1 is the local switch (the one where you typed the CLI command to show the domain-list) and has domain 99.
- The IVR manager obtained virtual domain 97 using 20:01:00:05:30:00:47:df as the WWN for a virtual switch.

```
switch# show fcdomain domain-list vsan 76
Number of domains: 3
Domain ID WWN

0xc8 (200) 20:01:00:05:30:00:47:df [Principal]
0x63 (99) 20:01:00:0d:ec:08:60:c1 [Local]
0x61 (97) 50:00:53:0f:ff:f0:10:06 [Virtual (IVR)]
```

Use the **show fcdomain allowed vsan** command to display the list of allowed domain IDs configured on this switch..

```
switch# show fcdomain allowed vsan 1
Assigned or unallowed domain IDs: 1-96,100,111-239.
[Interoperability Mode 1] allowed domain IDs: 97-127.
[User] configured allowed domain IDs: 50-110.
```

Ensure that the requested domain ID passes the switch software checks, if interop 1 mode is required in this switch.

The following example shows how to display all existing, persistent FC IDs for a specified VSAN. You can also specify the unused option to view only persistent FC IDs that are still not in use.

```
switch# show fcdomain fcid persistent vsan 1000
```

The following example shows how to display frame and other fcdomain statistics for a specified VSAN or SAN port channel:

```
switch# show fcdomain statistics vsan 1
VSAN Statistics
 Number of Principal Switch Selections: 5
 Number of times Local Switch was Principal: 0
 Number of 'Build Fabric's: 3
 Number of 'Fabric Reconfigurations': 0
```

The following example shows how to display FC ID allocation statistics including a list of assigned and free FC IDs:

```
switch# show fcdomain address-allocation vsan 1
```

The following example shows how to display the valid address allocation cache. The cache is used by the principal switch to reassign the FC IDs for a device (disk or host) that exited and reentered the fabric. In the cache content, VSAN refers to the VSAN that contains the device, WWN refers to the device that owned the FC IDs, and mask refers to a single or entire area of FC IDs.

```
switch# show fcdomain address-allocation cache
```

## Default Fibre Channel Domain Settings

The table below lists the default settings for all fcdomain parameters.

**Table 66: Default fcdomain Parameters**

| Parameters                                        | Default                 |
|---------------------------------------------------|-------------------------|
| fcdomain feature                                  | Enabled                 |
| Configured domain ID                              | 0 (zero)                |
| Configured domain                                 | Preferred               |
| auto-reconfigure option                           | Disabled                |
| contiguous-allocation option                      | Disabled                |
| Priority                                          | 128                     |
| Allowed list                                      | 1 to 239                |
| Fabric name                                       | 20:01:00:05:30:00:28:df |
| ref-reject                                        | Disabled                |
| Persistent FC ID                                  | Enabled                 |
| Allowed domain ID list configuration distribution | Disabled                |



# Configuring N Port Virtualization

---

This chapter contains the following sections:

- [Configuring N Port Virtualization, page 553](#)

## Configuring N Port Virtualization

### Information About NPV

#### NPV Overview

By default, Cisco Nexus 5000 Series switches operate in fabric mode. In this mode, the switch provides standard Fibre Channel switching capability and features.

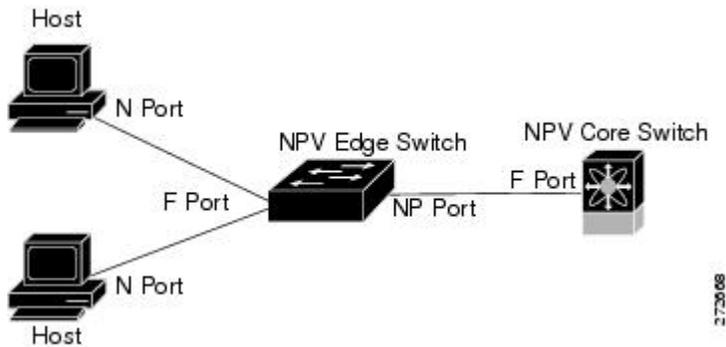
In fabric mode, each switch that joins a SAN is assigned a domain ID. Each SAN (or VSAN) supports a maximum of 239 domain IDs, so the SAN has a limit of 239 switches. In a SAN topology with a large number of edge switches, the SAN may need to grow beyond this limit. NPV alleviates the domain ID limit by sharing the domain ID of the core switch among multiple edge switches.

In NPV mode, the edge switch relays all traffic from server-side ports to the core switch. The core switch provides F port functionality (such as login and port security) and all the Fibre Channel switching capabilities.

The edge switch appears as a Fibre Channel host to the core switch and as a regular Fibre Channel switch to its connected devices.

The figure below shows an interface-level view of an NPV configuration.

**Figure 45: NPV Interface Configuration**



## NPV Mode

In NPV mode, the edge switch relays all traffic to the core switch, which provides the Fibre Channel switching capabilities. The edge switch shares the domain ID of the core switch.

To convert a switch into NPV mode, you set the NPV feature to enabled. This configuration command automatically triggers a switch reboot. You cannot configure NPV mode on a per-interface basis. NPV mode applies to the entire switch.

In NPV mode, a subset of fabric mode CLI commands and functionality is supported. For example, commands related to fabric login and name server registration are not required on the edge switch, because these functions are provided in the core switch. To display the fabric login and name server registration databases, you must enter the **show flogi database** and **show fcns database** commands on the core switch.

## Server Interfaces

Server interfaces are F ports on the edge switch that connect to the servers. A server interface may support multiple end devices by enabling the N port identifier virtualization (NPIV) feature. NPIV provides a means to assign multiple FC IDs to a single N port, which allows the server to assign unique FC IDs to different applications.



**Note** To use NPIV, enable the NPIV feature and reinitialize the server interfaces that will support multiple devices.

Server interfaces are automatically distributed among the NP uplinks to the core switch. All of the end devices connected to a server interface are mapped to the same NP uplink.

In Cisco Nexus 5000 Series switches, server interfaces can be physical or virtual Fibre Channel interfaces.

## Related Topics

[About N Port Identifier Virtualization, on page 525](#)

## NP Uplinks

All interfaces from the edge switch to the core switch are configured as proxy N ports (NP ports).

An NP uplink is a connection from an NP port on the edge switch to an F port on the core switch. When an NP uplink is established, the edge switch sends a fabric login message (FLOGI) to the core switch, and then (if the FLOGI is successful) it registers itself with the name server on the core switch. Subsequent FLOGIs from end devices connected to this NP uplink are converted to fabric discovery messages (FDISCs).



### Note

In the switch CLI configuration commands and output displays, NP uplinks are called External Interfaces.

In Cisco Nexus 5000 Series switches, NP uplink interfaces must be native Fibre Channel interfaces.

### Related Topics

[Information About Fabric Login, on page 663](#)

## FLOGI Operation

When an NP port becomes operational, the switch first logs itself in to the core switch by sending a FLOGI request (using the port WWN of the NP port).

After completing the FLOGI request, the switch registers itself with the fabric name server on the core switch (using the symbolic port name of the NP port and the IP address of the edge switch).

The following table identifies port and node names in the edge switch used in NPV mode.

**Table 67: Edge Switch FLOGI Parameters**

| Parameter          | Derived From                                                                                                                                                  |
|--------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------|
| pWWN               | The fWWN of the NP port on the edge switch.                                                                                                                   |
| nWWN               | The VSAN-based sWWN of the edge switch.                                                                                                                       |
| symbolic port name | The edge switch name and NP port interface string.<br><b>Note</b> If no switch name is available, the output will read "switch." For example, switch: fc 1/5. |
| IP address         | The IP address of the edge switch.                                                                                                                            |
| symbolic node name | The edge switch name.                                                                                                                                         |



### Note

The buffer-to-buffer state change number (BB\_SCN) of internal FLOGIs on an NP port is always set to zero. The BB\_SCN is supported by the F port on the edge switch.

We do not recommend using fWWN-based zoning on the edge switch for the following reasons:

- Zoning is not enforced at the edge switch (rather, it is enforced on the core switch).
- Multiple devices attached to an edge switch log in through the same F port on the core, so they cannot be separated into different zones.
- The same device might log in using different fWWNs on the core switch (depending on the NPV link it uses) and may need to be zoned using different fWWNs.

### Related Topics

[Configuring and Managing Zones, on page 603](#)

## NPV Traffic Management

### *Automatic Uplink Selection*

NPV supports automatic selection of NP uplinks. When a server interface is brought up, the NP uplink interface with the minimum load is selected from the available NP uplinks in the same VSAN as the server interface.

When a new NP uplink interface becomes operational, the existing load is not redistributed automatically to include the newly available uplink. Server interfaces that become operational after the NP uplink can select the new NP uplink.

### *Traffic Maps*

In Release 4.0(1a)N2(1) and later software releases, NPV supports traffic maps. A traffic map allows you to specify the NP uplinks that a server interface can use to connect to the core switches.



#### Note

When an NPV traffic map is configured for a server interface, the server interface must select only from the NP uplinks in its traffic map. If none of the specified NP uplinks are operational, the server remains in a non-operational state.

The NPV traffic map feature provides the following benefits:

- Facilitates traffic engineering by allowing configuration of a fixed set of NP uplinks for a specific server interface (or range of server interfaces).
- Ensures correct operation of the persistent FC ID feature, because a server interface will always connect to the same NP uplink (or one of a specified set of NP uplinks) after an interface reinitialization or switch reboot.

### *Disruptive Load Balancing*

In Release 4.0(0)N1(2a) and later software releases, NPV supports disruptive load balancing. When disruptive load balancing is enabled, NPV redistributes the server interfaces across all available NP uplinks when a new NP uplink becomes operational. To move a server interface from one NP uplink to another NP uplink, NPV forces reinitialization of the server interface so that the server performs a new login to the core switch.

Only server interfaces that are moved to a different uplink are reinitialized. A system message is generated for each server interface that is moved.



#### Note

Redistributing a server interface causes traffic disruption to the attached end devices.

To avoid disruption of server traffic, you should enable this feature only after adding a new NP uplink, and then disable it again after the server interfaces have been redistributed.

If disruptive load balancing is not enabled, you can manually reinitialize some or all of the server interfaces to distribute server traffic to new NP uplink interfaces.

### NPV Traffic Management Guidelines

When deploying NPV traffic management, follow these guidelines:

- Use NPV traffic management only when automatic traffic engineering does not meet your network requirements.
- You do not need to configure traffic maps for all server interfaces. By default, NPV will use automatic traffic management.
- Server interfaces configured to use a set of NP uplink interfaces cannot use any other available NP uplink interfaces, even if none of the configured interfaces are available.
- When disruptive load balancing is enabled, a server interface may be moved from one NP uplink to another NP uplink. Moving between NP uplink interfaces requires NPV to relogin to the core switch, causing traffic disruption.
- To link a set of servers to a specific core switch, associate the server interfaces with a set of NP uplink interfaces that all connect to that core switch.
- Configure Persistent FC IDs on the core switch and use the Traffic Map feature to direct server interface traffic onto NP uplinks that all connect to the associated core switch.

### NPV Guidelines and Limitations

When configuring NPV, note the following guidelines and limitations:

- In-order data delivery is not required in NPV mode because the exchange between two end devices always takes the same uplink from the edge switch to the core. Upstream of the edge switch, core switches will enforce in-order delivery if configured.
- You can configure zoning for end devices that are connected to edge switches using all available member types on the core switch. For fWWN, sWWN, domain, or port-based zoning, use the fWWN, sWWN, domain, or port of the core switch in the configuration commands.
- Port tracking is not supported in NPV mode.
- Port security is supported on the core switch for devices logged in through the NPV switch. Port security is enabled on the core switch on a per-interface basis. To enable port security on the core switch for devices that log in through an NPV switch, you must adhere to the following requirements:
  - The internal FLOGI must be in the port security database; in this way, the port on the core switch will allow communications and links.
  - All the end device pWWNs must also be in the port security database.
- Edge switches can connect to multiple core switches. In other words, different NP ports can be connected to different core switches.

- NPV uses a load-balancing algorithm to automatically assign end devices in a VSAN to one of the NP uplinks (in the same VSAN) upon initial login. If there are multiple NP uplinks in the same VSAN, you cannot assign an end device to a specific NP uplink.
- If a server interface goes down and then returns to service, the interface is not guaranteed to be assigned to the same NP uplink.
- The server interface is only operational when its assigned NP uplink is operational.
- Both servers and targets can be connected to the switch when in NPV mode.
- Fibre Channel switching is not performed in the edge switch; all traffic is switched in the core switch.
- NPV supports NPIV-capable module servers. This capability is called nested NPIV.
- Only F, NP, and SD ports are supported in NPV mode.

## Configuring NPV

### Enabling NPV

When you enable NPV, the system configuration is erased and the switch reboots.



#### Note

We recommend that you save your current configuration either in boot flash memory or to a TFTP server before you enable NPV.

To enable NPV, perform this task:

### SUMMARY STEPS

1. switch# **configure terminal**
2. switch(config)# **npv enable**
3. switch(config-npv)# **no npv enable**

### DETAILED STEPS

|               | Command or Action                        | Purpose                                                                                                                                          |
|---------------|------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Step 1</b> | switch# <b>configure terminal</b>        | Enters configuration mode.                                                                                                                       |
| <b>Step 2</b> | switch(config)# <b>npv enable</b>        | Enables NPV mode. The switch reboots, and it comes back up in NPV mode.<br><br><b>Note</b> A write-erase is performed during the initialization. |
| <b>Step 3</b> | switch(config-npv)# <b>no npv enable</b> | Disables NPV mode, which results in a reload of the switch.                                                                                      |

## Configuring NPV Interfaces

After you enable NPV, you should configure the NP uplink interfaces and the server interfaces.

### Configuring an NP Interface

After you enable NPV, you should configure the NP uplink interfaces and the server interfaces. To configure an NP uplink interface, perform this task:

To configure a server interface, perform this task:

## SUMMARY STEPS

1. switch# **configure terminal**
2. switch(config)# **interface fc slot/port**
3. switch(config-if)# **switchport mode NP**
4. switch(config-if)# **no shutdown**

## DETAILED STEPS

|               | Command or Action                             | Purpose                                                             |
|---------------|-----------------------------------------------|---------------------------------------------------------------------|
| <b>Step 1</b> | switch# <b>configure terminal</b>             | Enters configuration mode.                                          |
| <b>Step 2</b> | switch(config)# <b>interface fc slot/port</b> | Selects an interface that will be connected to the core NPV switch. |
| <b>Step 3</b> | switch(config-if)# <b>switchport mode NP</b>  | Configures the interface as an NP port.                             |
| <b>Step 4</b> | switch(config-if)# <b>no shutdown</b>         | Brings up the interface.                                            |

### Configuring a Server Interface

To configure a server interface, perform this task:

## SUMMARY STEPS

1. switch# **configure terminal**
2. switch(config)# **interface {fc slot/port | vfc vfc-id}**
3. switch(config-if)# **switchport mode F**
4. switch(config-if)# **no shutdown**

## DETAILED STEPS

|               | Command or Action                 | Purpose                    |
|---------------|-----------------------------------|----------------------------|
| <b>Step 1</b> | switch# <b>configure terminal</b> | Enters configuration mode. |

|               | Command or Action                                            | Purpose                                |
|---------------|--------------------------------------------------------------|----------------------------------------|
| <b>Step 2</b> | switch(config)# <b>interface</b> {fc slot/port   vfc vfc-id} | Selects a server interface.            |
| <b>Step 3</b> | switch(config-if)# <b>switchport mode F</b>                  | Configures the interface as an F port. |
| <b>Step 4</b> | switch(config-if)# <b>no shutdown</b>                        | Brings up the interface.               |

### Configuring NPV Traffic Management

#### Configuring NPV Traffic Maps

An NPV traffic map associates one or more NP uplink interfaces with a server interface. The switch associates the server interface with one of these NP uplinks.



**Note** If a server interface is already mapped to an NP uplink, you should include this mapping in the traffic map configuration.

To configure a traffic map, perform this task:

#### SUMMARY STEPS

1. switch# **config t**
2. switch(config)# **npv traffic-map server-interface** {fc slot/port | vfc vfc-id} **external-interface fc slot/port**
3. switch(config)# **no npv traffic-map server-interface** {fc slot/port | vfc vfc-id} **external-interface fc slot/port**

#### DETAILED STEPS

|               | Command or Action                                                                                                             | Purpose                                                                                                                                        |
|---------------|-------------------------------------------------------------------------------------------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Step 1</b> | switch# <b>config t</b>                                                                                                       | Enters configuration mode on the NPV.                                                                                                          |
| <b>Step 2</b> | switch(config)# <b>npv traffic-map server-interface</b> {fc slot/port   vfc vfc-id} <b>external-interface fc slot/port</b>    | Configures a mapping between a server interface (or range of server interfaces) and an NP uplink interface (or range of NP uplink interfaces). |
| <b>Step 3</b> | switch(config)# <b>no npv traffic-map server-interface</b> {fc slot/port   vfc vfc-id} <b>external-interface fc slot/port</b> | Removes the mapping between the specified server interfaces and NP uplink interfaces.                                                          |

#### Enabling Disruptive Load Balancing

If you configure additional NP uplinks, you can enable the disruptive load-balancing feature to distribute the server traffic load evenly among all the NP uplinks.

To enable disruptive load balancing, perform this task:

## SUMMARY STEPS

1. switch# **configure terminal**
2. switch(config)# **npv auto-load-balance disruptive**
3. switch (config)# **no npv auto-load-balance disruptive**

## DETAILED STEPS

|               | Command or Action                                           | Purpose                                           |
|---------------|-------------------------------------------------------------|---------------------------------------------------|
| <b>Step 1</b> | switch# <b>configure terminal</b>                           | Enters configuration mode on the NPV.             |
| <b>Step 2</b> | switch(config)# <b>npv auto-load-balance disruptive</b>     | Enables disruptive load balancing on the switch.  |
| <b>Step 3</b> | switch (config)# <b>no npv auto-load-balance disruptive</b> | Disables disruptive load balancing on the switch. |

## Verifying NPV

To display information about NPV, perform the following task:

## SUMMARY STEPS

1. switch# **show npv flogi-table [all]**

## DETAILED STEPS

|               | Command or Action                         | Purpose                         |
|---------------|-------------------------------------------|---------------------------------|
| <b>Step 1</b> | switch# <b>show npv flogi-table [all]</b> | Displays the NPV configuration. |

## Verifying NPV Examples

To display a list of devices on a server interface and their assigned NP uplinks, enter the **show npv flogi-table** command on the Cisco Nexus 5000 Series switch:

```
switch# show npv flogi-table

SERVER
INTERFACE VSAN FCID PORT NAME NODE NAME EXTERNAL
INTERFACE

vfc3/1 1 0xee0008 10:00:00:00:c9:60:e4:9a 20:00:00:00:c9:60:e4:9a fc2/1
vfc3/1 1 0xee0009 20:00:00:00:0a:00:00:01 20:00:00:00:c9:60:e4:9a fc2/2
vfc3/1 1 0xee000a 20:00:00:00:0a:00:00:02 20:00:00:00:c9:60:e4:9a fc2/3
```

```
vfc3/1 1 0xee000b 33:33:33:33:33:33:33:33 20:00:00:00:c9:60:e4:9a fc2/4
Total number of flogi = 4
```



**Note** For each server interface, the External Interface value displays the assigned NP uplink.

To display the status of the server interfaces and the NP uplink interfaces, enter the **show npv status** command:

```
switch# show npv status
npiv is enabled

External Interfaces:
=====
Interface: fc2/1, VSAN: 1, FCID: 0x1c0000, State: Up
Interface: fc2/2, VSAN: 1, FCID: 0x040000, State: Up
Interface: fc2/3, VSAN: 1, FCID: 0x260000, State: Up
Interface: fc2/4, VSAN: 1, FCID: 0x1a0000, State: Up
Number of External Interfaces: 4

Server Interfaces:
=====
Interface: vfc3/1, VSAN: 1, NPIV: No, State: Up
Number of Server Interfaces: 1
```



**Note** To view fcns database entries for NPV edge switches, you must enter the **show fcns database** command on the core switch.

To view all the NPV edge switches, enter the **show fcns database** command on the core switch:

```
core-switch# show fcns database
```

For additional details (such as IP addresses, switch names, interface names) about the NPV edge switches that you see in the **show fcns database** output, enter the **show fcns database detail** command on the core switch:

```
core-switch# show fcns database detail
```

## Verifying NPV Traffic Management

To display the NPV traffic map, enter the **show npv traffic-map** command.

```
switch# show npv traffic-map
NPV Traffic Map Information:

Server-If External-If(s)

fc1/3 fc1/10,fc1/11
fc1/5 fc1/1,fc1/2

```

To display the NPV internal traffic details, enter the **show npv internal info traffic-map** command.

To display the disruptive load-balancing status, enter the **show npv status** command:

```
switch# show npv status
npiv is enabled
disruptive load balancing is enabled
External Interfaces:
=====
Interface: fc2/1, VSAN: 2, FCID: 0x1c0000, State: Up
...
```

## Configuring VSAN Trunking

This chapter contains the following sections:

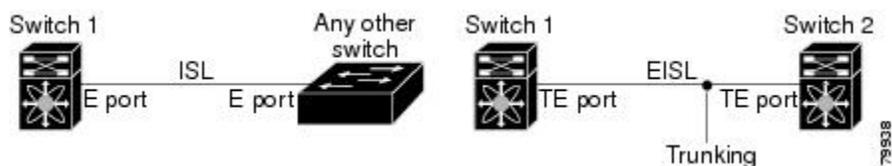
- [Configuring VSAN Trunking, page 563](#)

### Configuring VSAN Trunking

#### Information About VSAN Trunking

VSAN trunking enables interconnect ports to transmit and receive frames in more than one VSAN, over the same physical link, using enhanced ISL (EISL) frame format (see the following figure).

**Figure 46: VSAN Trunking**



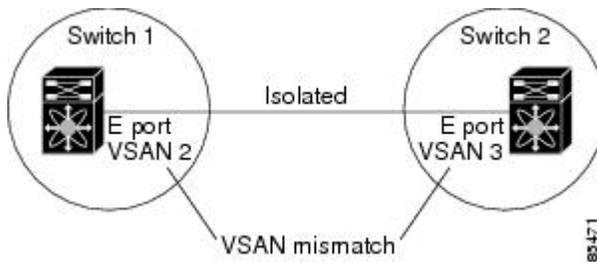
VSAN trunking is supported on native Fibre Channel interfaces, but not on virtual Fibre Channel interfaces. The VSAN trunking feature includes the following restrictions:

- Trunking configurations are only applicable to E ports. If trunk mode is enabled in an E port and that port becomes operational as a trunking E port, it is referred to as a TE port.
- The trunk-allowed VSANs configured for TE ports are used by the trunking protocol to determine the allowed-active VSANs in which frames can be received or transmitted.
- If a trunking-enabled E port is connected to a third-party switch, the trunking protocol ensures seamless operation as an E port.

## VSAN Trunking Mismatches

If you misconfigure VSAN configurations across E ports, issues can occur such as the merging of traffic in two VSANs (causing both VSANs to mismatch). The VSAN trunking protocol validates the VSAN interfaces at both ends of an ISL to avoid merging VSANs (see the following figure).

**Figure 47: VSAN Mismatch**



In this example, the trunking protocol detects potential VSAN merging and isolates the ports involved.

The trunking protocol cannot detect merging of VSANs when a third-party switch is placed in between two Cisco Nexus 5000 Series switches (see the following figure).

**Figure 48: Third-Party Switch VSAN Mismatch**



VSAN 2 and VSAN 3 are effectively merged with overlapping entries in the name server and the zone applications. The Cisco MDS 9000 Fabric Manager helps detect such topologies.

## VSAN Trunking Protocol

The trunking protocol is important for E-port and TE-port operations. It supports the following capabilities:

- Dynamic negotiation of operational trunk mode.
- Selection of a common set of trunk-allowed VSANs.
- Detection of a VSAN mismatch across an ISL.

By default, the VSAN trunking protocol is enabled. If the trunking protocol is disabled on a switch, no port on that switch can apply new trunk configurations. Existing trunk configurations are not affected: the TE port continues to function in trunk mode, but only supports traffic in VSANs that it negotiated with previously (when the trunking protocol was enabled). Other switches that are directly connected to this switch are similarly affected on the connected interfaces. If you need to merge traffic from different port VSANs across a nontrunking ISL, disable the trunking protocol.

## Configuring VSAN Trunking

### Guidelines and Restrictions

When configuring VSAN trunking, note the following guidelines:

- We recommend that both ends of a VSAN trunking ISL belong to the same port VSAN. On platforms or fabric switches where the port VSANs are different, one end returns an error, and the other is not connected.
- To avoid inconsistent configurations, disable all E ports with a **shutdown** command before enabling or disabling the VSAN trunking protocol.

### Enabling or Disabling the VSAN Trunking Protocol

To enable or disable the VSAN trunking protocol, perform this task:

#### SUMMARY STEPS

1. switch# **configuration terminal**
2. switch(config)# **no trunk protocol enable**
3. switch(config)# **trunk protocol enable**

#### DETAILED STEPS

|               | Command or Action                               | Purpose                              |
|---------------|-------------------------------------------------|--------------------------------------|
| <b>Step 1</b> | switch# <b>configuration terminal</b>           | Enters configuration mode.           |
| <b>Step 2</b> | switch(config)# <b>no trunk protocol enable</b> | Disables the trunking protocol.      |
| <b>Step 3</b> | switch(config)# <b>trunk protocol enable</b>    | Enables trunking protocol (default). |

### About Trunk Mode

By default, trunk mode is enabled in all Fibre Channel interfaces. However, trunk mode configuration takes effect only in E-port mode. You can configure trunk mode as on (enabled), off (disabled), or auto (automatic). The default trunk mode is on. The trunk mode configurations at the two ends of the link determine the trunking state of the link and the port modes at both ends (see the following table).

**Table 68: Trunk Mode Status Between Switches**

| Your Trunk Mode Configuration | Resulting State and Port Mode |                   |                |
|-------------------------------|-------------------------------|-------------------|----------------|
|                               | Switch 1                      | Switch 2          | Trunking State |
| On                            | Auto or on                    | Trunking (EISL)   | TE port        |
| Off                           | Auto, on, or off              | No trunking (ISL) | E port         |
| Auto                          | Auto                          | No trunking (ISL) | E port         |

The preferred configuration on the Cisco Nexus 5000 Series switches is that one side of the trunk is set to auto and the other is set to on.



**Note**

When connected to a third-party switch, the trunk mode configuration has no effect. The ISL is always in a trunking disabled state.

**Configuring Trunk Mode**

To configure trunk mode, perform this task:

**SUMMARY STEPS**

1. switch# **configuration terminal**
2. switch(config)# **interface fc slot/port**
3. switch(config-if)# **switchport trunk mode on**
4. switch(config-if)# **switchport trunk mode off**
5. switch(config-if)# **switchport trunk mode auto**

**DETAILED STEPS**

|               | Command or Action                                   | Purpose                                                       |
|---------------|-----------------------------------------------------|---------------------------------------------------------------|
| <b>Step 1</b> | switch# <b>configuration terminal</b>               | Enters configuration mode.                                    |
| <b>Step 2</b> | switch(config)# <b>interface fc slot/port</b>       | Configures the specified interface.                           |
| <b>Step 3</b> | switch(config-if)# <b>switchport trunk mode on</b>  | Enables (default) the trunk mode for the specified interface. |
| <b>Step 4</b> | switch(config-if)# <b>switchport trunk mode off</b> | Disables the trunk mode for the specified interface.          |

|        | Command or Action                                    | Purpose                                                                                            |
|--------|------------------------------------------------------|----------------------------------------------------------------------------------------------------|
| Step 5 | switch(config-if)# <b>switchport trunk mode auto</b> | Configures the trunk mode to <b>auto</b> mode, which provides automatic sensing for the interface. |

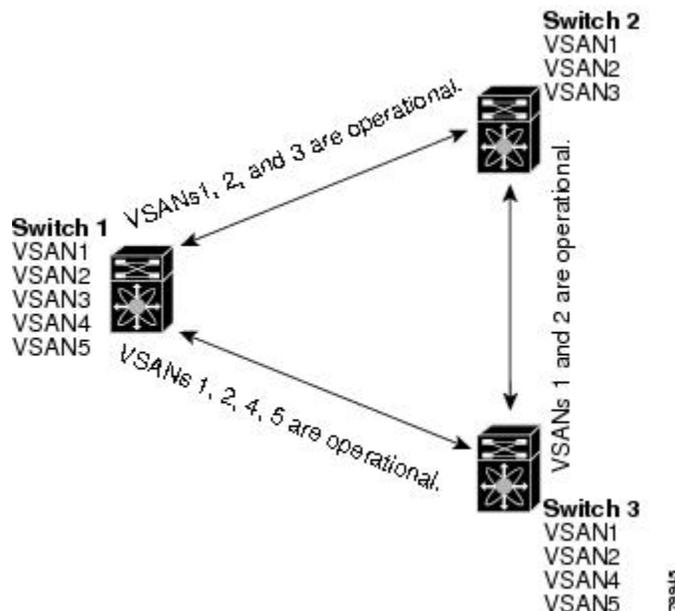
### About Trunk-Allowed VSAN Lists

Each Fibre Channel interface has an associated trunk-allowed VSAN list. In TE-port mode, frames are transmitted and received in one or more VSANs specified in this list. By default, the complete VSAN range (1 through 4093) is included in the trunk-allowed list.

The common set of VSANs that are configured and active in the switch are included in the trunk-allowed VSAN list for an interface, and they are called *allowed-active VSANs*. The trunking protocol uses the list of allowed-active VSANs at the two ends of an ISL to determine the list of operational VSANs in which traffic is allowed.

In the following figure, switch 1 has VSANs 1 through 5, switch 2 has VSANs 1 through 3, and switch 3 has VSANs 1, 2, 4, and 5 with a default configuration of trunk-allowed VSANs. All VSANs configured in all three switches are allowed-active. However, only the common set of allowed-active VSANs at the ends of the ISL become operational as shown in below.

**Figure 49: Default Allowed-Active VSAN Configuration**



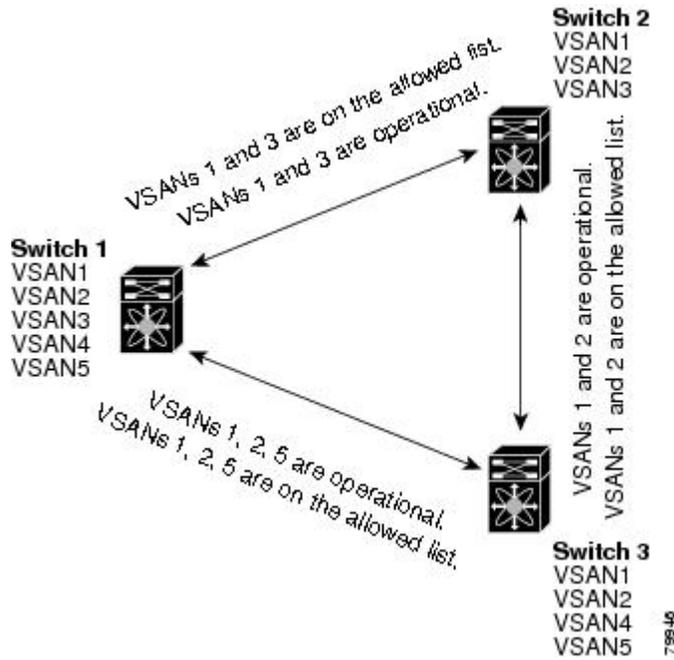
You can configure a selected set of VSANs (from the allowed-active list) to control access to the VSANs specified in a trunking ISL.

Using the figure above as an example, you can configure the list of allowed VSANs on a per-interface basis (see the following figure). For example, if VSANs 2 and 4 are removed from the allowed VSAN list of ISLs connecting to switch 1, the operational allowed list of VSANs for each ISL would be as follows:

- The ISL between switch 1 and switch 2 includes VSAN 1 and VSAN 3.
- The ISL between switch 2 and switch 3 includes VSAN 1 and VSAN 2.
- The ISL between switch 3 and switch 1 includes VSAN 1, 2, and 5.

Consequently, VSAN 2 can only be routed from switch 1 through switch 3 to switch 2.

**Figure 50: Operational and Allowed VSAN Configuration**



### Configuring an Allowed-Active List of VSANs

To configure an allowed-active list of VSANs for an interface, perform this task:

#### SUMMARY STEPS

1. switch# **configuration terminal**
2. switch(config)# **interface fc slot/port**
3. switch(config-if)# **switchport trunk allowed vsan vsan-id - vsan-id**
4. switch(config-if)# **switchport trunk allowed vsan add vsan-id**
5. switch(config-if)# **no switchport trunk allowed vsan vsan-id - vsan-id**
6. switch(config-if)# **no switchport trunk allowed vsan add vsan-id**

## DETAILED STEPS

|        | Command or Action                                                            | Purpose                                                |
|--------|------------------------------------------------------------------------------|--------------------------------------------------------|
| Step 1 | switch# <b>configuration terminal</b>                                        | Enters configuration mode.                             |
| Step 2 | switch(config)# <b>interface fc slot/port</b>                                | Configures the specified interface.                    |
| Step 3 | switch(config-if)# <b>switchport trunk allowed vsan vsan-id - vsan-id</b>    | Changes the allowed list for the specified VSAN range. |
| Step 4 | switch(config-if)# <b>switchport trunk allowed vsan add vsan-id</b>          | Expands the specified VSAN to the new allowed list.    |
| Step 5 | switch(config-if)# <b>no switchport trunk allowed vsan vsan-id - vsan-id</b> | Deletes the specified VSAN range.                      |
| Step 6 | switch(config-if)# <b>no switchport trunk allowed vsan add vsan-id</b>       | Deletes the expanded allowed list.                     |

## Displaying VSAN Trunking Information

The **show interface** command is invoked from the EXEC mode and displays VSAN trunking configurations for a TE port. Without any arguments, this command displays the information for all of the configured interfaces in the switch.

The following example shows how to display the trunk mode of a Fibre Channel interface:

```
switch# show interface fc3/3
fc3/3 is up
 Hardware is Fibre Channel, SFP is short wave laser w/o OFC (SN)
 Port WWN is 20:83:00:0d:ec:6d:78:40
 Peer port WWN is 20:0c:00:0d:ec:0d:d0:00
 Admin port mode is auto, trunk mode is on
...
```

The following example shows how to display the trunk protocol of a Fibre Channel interface:

```
switch# show trunk protocol
Trunk protocol is enabled
```

The following example shows how to display the VSAN information for all trunk interfaces:

```
switch# show interface trunk vsan 1-1000
fc3/1 is not trunking
...
fc3/11 is trunking
 Belongs to san-port-channel 6
 Vsan 1 is up, FCID is 0xef0000
 Vsan 2 is up, FCID is 0xef0000
...
san-port-channel 6 is trunking
 Vsan 1 is up, FCID is 0xef0000
 Vsan 2 is up, FCID is 0xef0000
```

## Default Trunk Configuration Settings

The following table lists the default settings for trunking parameters.

**Table 69: Default Trunk Configuration Parameters**

| <b>Parameters</b>      | <b>Default</b>                  |
|------------------------|---------------------------------|
| Switch port trunk mode | On                              |
| Allowed VSAN list      | 1 to 4093 user-defined VSAN IDs |
| Trunking protocol      | Enabled                         |



## Configuring SAN Port Channel

---

This chapter contains the following sections:

- [Configuring SAN Port Channels, page 571](#)

### Configuring SAN Port Channels

SAN port channels refer to the aggregation of multiple physical interfaces into one logical interface to provide higher aggregated bandwidth, load balancing, and link redundancy.

On Cisco Nexus 5000 Series switches, SAN port channels can include physical Fibre Channel interfaces, but not virtual Fibre Channel interfaces. A SAN port channel can include up to eight Fibre Channel interfaces.

#### Information About SAN Port Channels

A SAN port channel has the following functionality:

- Provides a point-to-point connection over ISL (E ports) or EISL (TE ports). Multiple links can be combined into a SAN port channel.
- Increases the aggregate bandwidth on an ISL by distributing traffic among all functional links in the channel.
- Load balances across multiple links and maintains optimum bandwidth utilization. Load balancing is based on the source ID, destination ID, and exchange ID (OX ID).
- Provides high availability on an ISL. If one link fails, traffic previously carried on this link is switched to the remaining links. If a link goes down in a SAN port channel, the upper layer protocol is not aware of it. To the upper layer protocol, the link is still there, although the bandwidth is diminished. The routing tables are not affected by link failure.

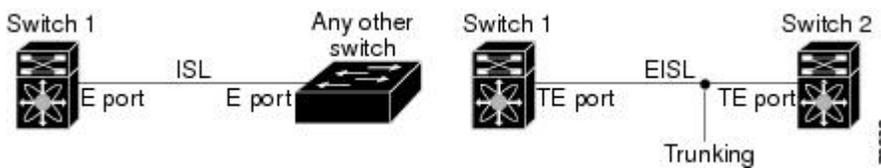
Cisco Nexus 5000 Series switches support a maximum of four SAN port channels (with eight interfaces per port channel). A port channel number refers to the unique (within each switch) identifier associated with each channel group. This number ranges from 1 to 256.

## Understanding Port Channels and VSAN Trunking

Switches in the Cisco Nexus 5000 Series implement VSAN trunking and port channels as follows:

- A SAN port channel enables several physical links to be combined into one aggregated logical link.
- An industry standard E port can link to other vendor switches and is referred to as inter-switch link (ISL), as shown on the left side of the figure below.
- VSAN trunking enables a link transmitting frames in the EISL format to carry traffic for multiple VSAN . When trunking is operational on an E port, that E port becomes a TE port. EISLs connects only between Cisco switches, as shown on the right side of the figure below.

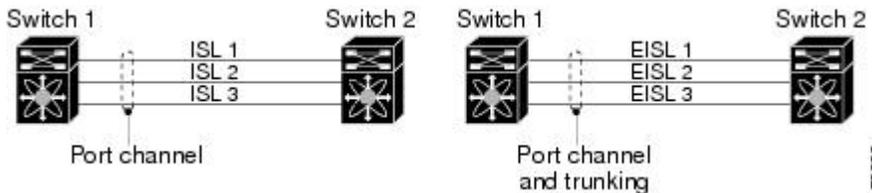
**Figure 51: VSAN Trunking Only**



You can create a SAN port channel with members that are E ports, as shown on the left side of the figure below. In this configuration, the port channel implements a logical ISL (carrying traffic for one VSAN).

You can create a SAN port channel with members that are TE-ports, as shown on the right side of the figure below. In this configuration, the port channel implements a logical EISL (carrying traffic for multiple VSANs).

**Figure 52: Port Channels and VSAN Trunking**



### Related Topics

[Configuring VSAN Trunking, on page 563](#)

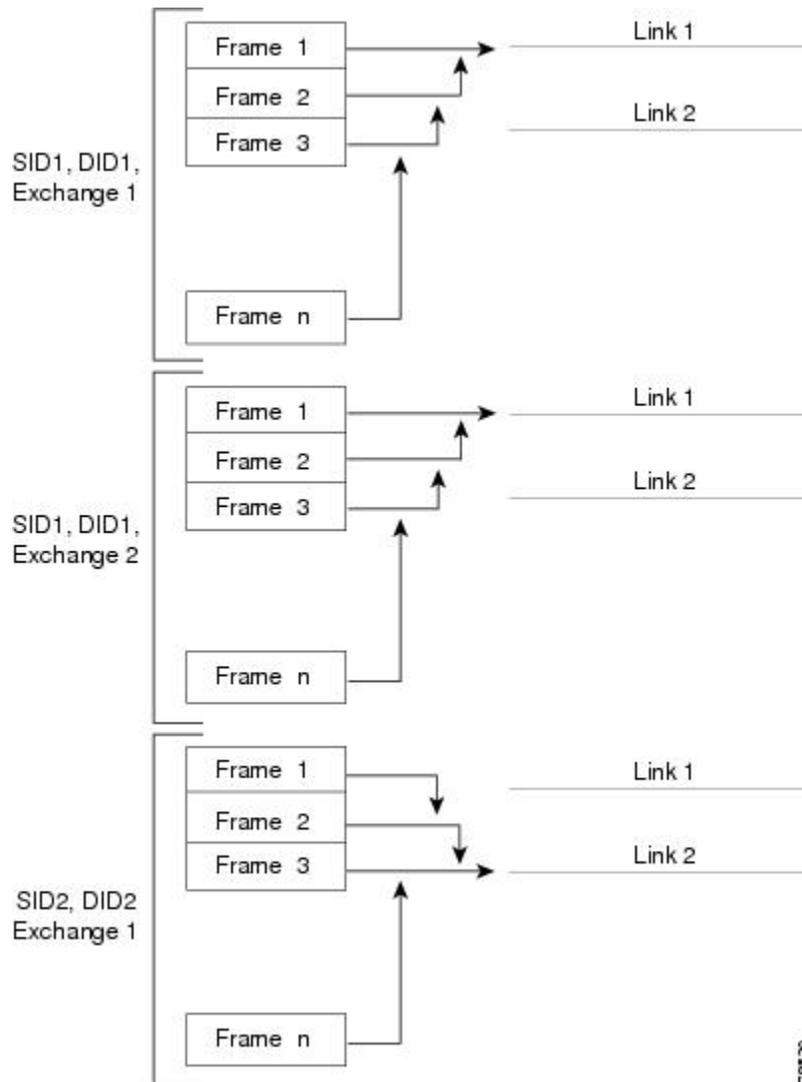
## Understanding Load Balancing

Load-balancing functionality can be provided using the following methods:

- Flow based—All frames between source and destination follow the same links for a given flow. That is, whichever link is selected for the first exchange of the flow is used for all subsequent exchanges.
- Exchange based—The first frame in an exchange is assigned to a link, and then subsequent frames in the exchange follow the same link. However, subsequent exchanges can use a different link. This method provides finer granularity for load balancing while preserving the order of frames for each exchange.

The following figure illustrates how flow-based load balancing works. When the first frame in a flow is received on an interface for forwarding, link 1 is selected. Each subsequent frame in that flow is sent over the same link. No frame in SID1 and DID1 utilizes link 2.

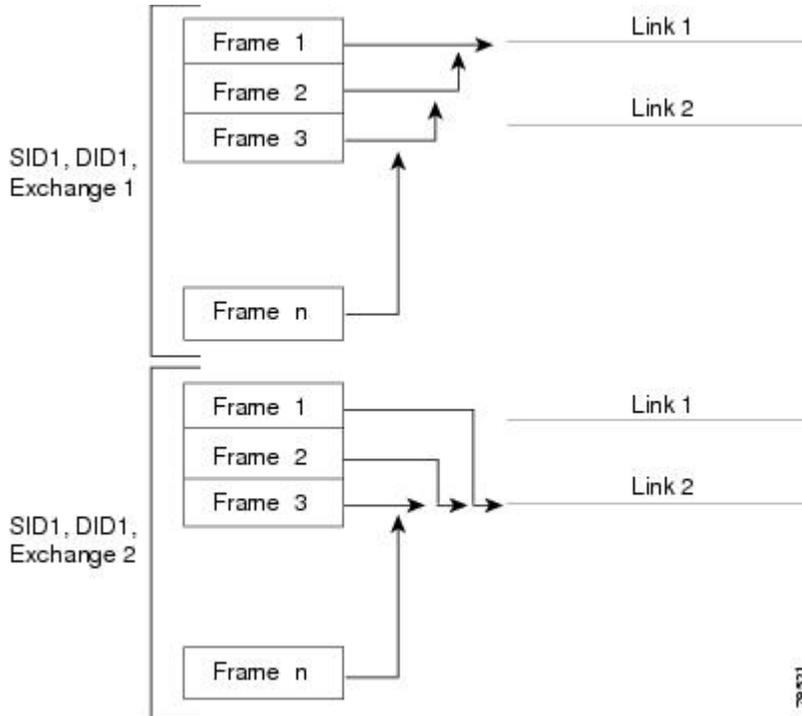
**Figure 53: SID1, DID1, and Flow-Based Load Balancing**



The following figure illustrates how exchange-based load balancing works. When the first frame in an exchange is received for forwarding on an interface, link 1 is chosen by a hash algorithm. All remaining frames in that

particular exchange are sent on the same link. For exchange 1, no frame uses link 2. For the next exchange, link 2 is chosen by the hash algorithm. Now all frames in exchange 2 use link 2.

**Figure 54: SID1, DID1, and Exchange-Based Load Balancing**

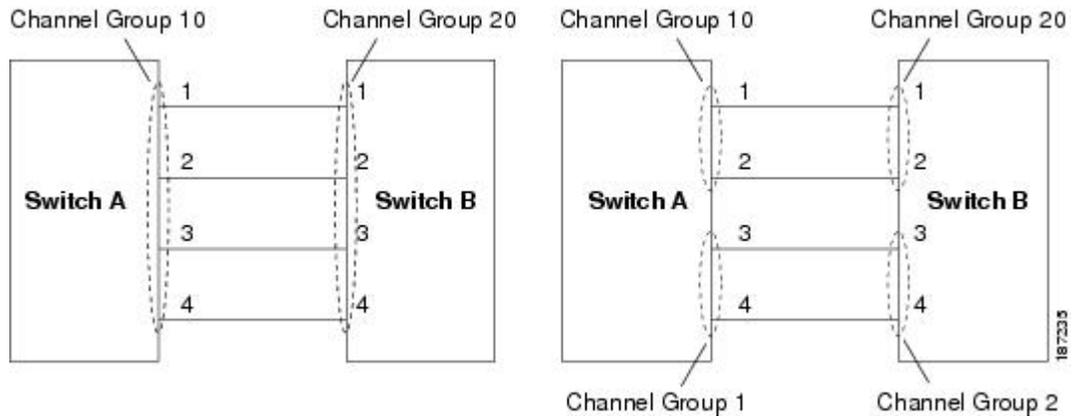


## Configuring SAN Port Channels

SAN port channels are created with default values. You can change the default configuration just as any other physical interface.

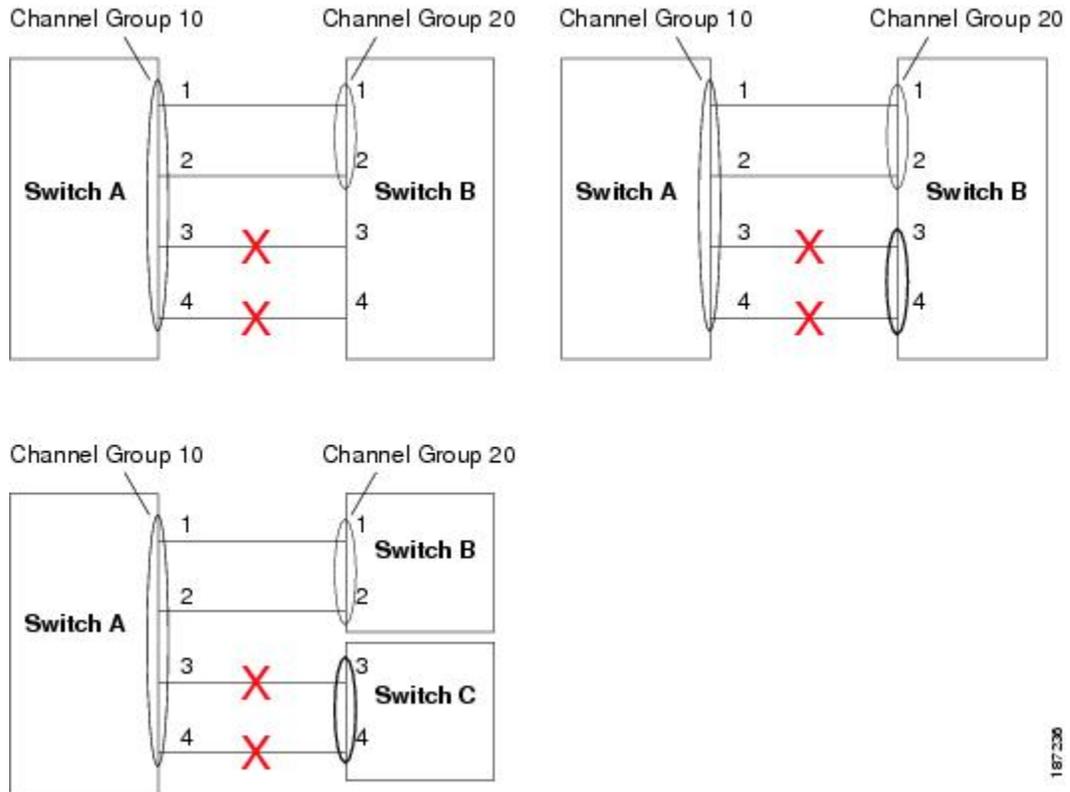
The following figure provides examples of valid SAN port channel configurations.

**Figure 55: Valid SAN Port Channel Configurations**



The following figure shows examples of invalid configurations. Assuming that the links are brought up in the 1, 2, 3, 4 sequence, links 3 and 4 will be operationally down as the fabric is misconfigured.

**Figure 56: Misconfigured Configurations**



1872306

### SAN Port Channel Configuration Guidelines

Before configuring a SAN port channel, consider the following guidelines:

- Configure the SAN port channel using Fibre Channel ports from both expansion modules to provide increased availability (if one of the expansion modules failed).
- Ensure that one SAN port channel is not connected to different sets of switches. SAN port channels require point-to-point connections between the same set of switches.

If you misconfigure SAN port channels, you may receive a misconfiguration message. If you receive this message, the port channel's physical links are disabled because an error has been detected.

If the following requirements are not met, a SAN port channel error is detected:

- Each switch on either side of a SAN port channel must be connected to the same number of interfaces.
- Each interface must be connected to a corresponding interface on the other side.
- Links in a SAN port channel cannot be changed after the port channel is configured. If you change the links after the port channel is configured, be sure to reconnect the links to interfaces within the port channel and reenable the links.

If all three conditions are not met, the faulty link is disabled.

Enter the **show interface** command for that interface to verify that the SAN port channel is functioning as required.

### Creating a SAN Port Channel

To create a SAN port channel, perform this task:

#### SUMMARY STEPS

1. switch# **configuration terminal**
2. switch(config)# **interface san-port-channel** *channel-number*

#### DETAILED STEPS

|               | Command or Action                                                          | Purpose                                                                                                                      |
|---------------|----------------------------------------------------------------------------|------------------------------------------------------------------------------------------------------------------------------|
| <b>Step 1</b> | switch# <b>configuration terminal</b>                                      | Enters configuration mode.                                                                                                   |
| <b>Step 2</b> | switch(config)# <b>interface san-port-channel</b><br><i>channel-number</i> | Creates the specified SAN port channel using the default mode (on). The SAN port channel number is in the range of 1 to 256. |

### SAN Port Channel Configuration Guidelines

You can configure each SAN port channel with a channel group mode parameter to determine the port channel protocol behavior for all member ports in this channel group. The possible values for a channel group mode are as follows:

- On (default)—The member ports only operate as part of a SAN port channel or remain inactive. In this mode, the port channel protocol is not initiated. However, if a port channel protocol frame is received from a peer port, the software indicates its nonnegotiable status. Port channels configured in the On mode require you to explicitly enable and disable the port channel member ports at either end if you add or remove ports from the port channel configuration. You must physically verify that the local and remote ports are connected to each other.
- Active—The member ports initiate port channel protocol negotiation with the peer port(s) regardless of the channel group mode of the peer port. If the peer port, while configured in a channel group, does not support the port channel protocol, or responds with a nonnegotiable status, it will default to the On mode behavior. The Active port channel mode allows automatic recovery without explicitly enabling and disabling the port channel member ports at either end.

The table below compares On and Active modes.

**Table 70: Channel Group Configuration Differences**

| <b>On Mode</b>                                                                                                                                                        | <b>Active Mode</b>                                                                                                                                                     |
|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| No protocol is exchanged.                                                                                                                                             | A port channel protocol negotiation is performed with the peer ports.                                                                                                  |
| Moves interfaces to the suspended state if its operational values are incompatible with the SAN port channel.                                                         | Moves interfaces to the isolated state if its operational values are incompatible with the SAN port channel.                                                           |
| When you add or modify a port channel member port configuration, you must explicitly disable (shut) and enable (no shut) the port channel member ports at either end. | When you add or modify a port channel interface, the SAN port channel automatically recovers.                                                                          |
| Port initialization is not synchronized.                                                                                                                              | There is synchronized startup of all ports in a channel across peer switches.                                                                                          |
| All misconfigurations are not detected as no protocol is exchanged.                                                                                                   | Consistently detect misconfigurations using a port channel protocol.                                                                                                   |
| Transitions misconfigured ports to the suspended state. You must explicitly disable (shut) and enable (no shut) the member ports at either end.                       | Transitions misconfigured ports to the isolated state to correct the misconfiguration. Once you correct the misconfiguration, the protocol ensures automatic recovery. |
| This is the default mode.                                                                                                                                             | You must explicitly configure this mode.                                                                                                                               |

### About SAN Port Channel Deletion

When you delete the SAN port channel, the corresponding channel membership is also deleted. All interfaces in the deleted SAN port channel convert to individual physical links. After the SAN port channel is removed, regardless of the mode (active and on) used, the ports at either end are gracefully brought down, indicating that no frames are lost when the interface is going down.

If you delete the SAN port channel for one port, then the individual ports within the deleted SAN port channel retain the compatibility parameter settings (speed, mode, port VSAN, allowed VSAN, and port security). You can explicitly change those settings as required.

- If you use the default On mode to avoid inconsistent states across switches and to maintain consistency across switches, then the ports shut down. You must explicitly enable those ports again.
- If you use the Active mode, then the port channel ports automatically recover from the deletion.

### Related Topics

[Setting the Interface Administrative State, on page 519](#)

#### *Configuring Active Mode SAN Port Channel*

To configure active mode, perform this task:

**SUMMARY STEPS**

1. switch# **configuration terminal**
2. switch(config)# **interface san-port-channel** *channel-number*
3. switch(config-if)# **channel mode active**
4. switch(config-if)# **no channel mode active**

**DETAILED STEPS**

|               | <b>Command or Action</b>                                                | <b>Purpose</b>                                                                                                            |
|---------------|-------------------------------------------------------------------------|---------------------------------------------------------------------------------------------------------------------------|
| <b>Step 1</b> | switch# <b>configuration terminal</b>                                   | Enters configuration mode.                                                                                                |
| <b>Step 2</b> | switch(config)# <b>interface san-port-channel</b> <i>channel-number</i> | Configures the specified port channel using the default On mode. The SAN port channel number is in the range of 1 to 256. |
| <b>Step 3</b> | switch(config-if)# <b>channel mode active</b>                           | Configures the Active mode.                                                                                               |
| <b>Step 4</b> | switch(config-if)# <b>no channel mode active</b>                        | Reverts to the default On mode.                                                                                           |

**Example of Configuring Active Modes**

The following example shows how to configure active mode:

```
switch(config)# interface san-port-channel 1
switch(config-if)# channel mode active
```

**Deleting SAN Port Channels**

To delete a SAN port channel, perform this task:

**SUMMARY STEPS**

1. switch# **configuration terminal**
2. switch(config)# **no interface san-port-channel** *channel-number*

**DETAILED STEPS**

|               | <b>Command or Action</b>                                                   | <b>Purpose</b>                                                                                                                  |
|---------------|----------------------------------------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------|
| <b>Step 1</b> | switch# <b>configuration terminal</b>                                      | Enters configuration mode.                                                                                                      |
| <b>Step 2</b> | switch(config)# <b>no interface san-port-channel</b> <i>channel-number</i> | Deletes the specified port channel, its associated interface mappings, and the hardware associations for this SAN port channel. |

## Interfaces in a SAN Port Channel

You can add or remove a physical Fibre Channel interface (or a range of interfaces) to an existing SAN port channel. The compatible parameters on the configuration are mapped to the SAN port channel. Adding an interface to a SAN port channel increases the channel size and bandwidth of the SAN port channel. Removing an interface from a SAN port channel decreases the channel size and bandwidth of the SAN port channel.


**Note**


---

Virtual Fibre Channel interfaces cannot be added to SAN port channels.

---

### About Interface Addition to a SAN Port Channel

You can add a physical interface (or a range of interfaces) to an existing SAN port channel. The compatible parameters on the configuration are mapped to the SAN port channel. Adding an interface to a SAN port channel increases the channel size and bandwidth of the SAN port channel.

After the members are added, regardless of the mode (Active and On) used, the ports at either end are gracefully brought down, indicating that no frames are lost when the interface is going down.

#### Compatibility Check

A compatibility check ensures that the same parameter settings are used in all physical ports in the channel. Otherwise, they cannot become part of a SAN port channel. The compatibility check is performed before a port is added to the SAN port channel.

The check ensures that the following parameters and settings match at both ends of a SAN port channel:

- Capability parameters (type of interface, Fibre Channel at both ends).
- Administrative compatibility parameters (speed, mode, port VSAN, allowed VSAN, and port security).
- Operational parameters (speed and remote switch's WWN).

A port addition procedure fails if the capability and administrative parameters in the remote switch are incompatible with the capability and administrative parameters in the local switch. If the compatibility check is successful, the interfaces are operational and the corresponding compatibility parameter settings apply to these interfaces.

#### Suspended and Isolated States

If the operational parameters are incompatible, the compatibility check fails and the interface is placed in a suspended or isolated state based on the configured mode:

- An interface enters the suspended state if the interface is configured in the On mode.
- An interface enters the isolated state if the interface is configured in the Active mode.

#### Related Topics

[Reason Codes](#), on page 515

### Adding an Interface to a SAN Port Channel

To add an interface to a SAN port channel, perform this task:

## SUMMARY STEPS

1. switch# **configuration terminal**
2. switch(config)# **interface** *type slot/port*
3. switch(config-if)# **channel-group** *channel-number*

## DETAILED STEPS

|               | Command or Action                                                | Purpose                                                                                                                                     |
|---------------|------------------------------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Step 1</b> | switch# <b>configuration terminal</b>                            | Enters configuration mode.                                                                                                                  |
| <b>Step 2</b> | switch(config)# <b>interface</b> <i>type slot/port</i>           | Enters configuration mode for the specified interface.                                                                                      |
| <b>Step 3</b> | switch(config-if)# <b>channel-group</b><br><i>channel-number</i> | Adds the Fibre Channel interface to the specified channel group. If the channel group does not exist, it is created. The port is shut down. |

### Forcing an Interface Addition

You can force the port configuration to be overwritten by the SAN port channel. In this case, the interface is added to a SAN port channel.

- If you use the default On mode to avoid inconsistent states across switches and to maintain consistency across switches, then the ports shut down. You must explicitly enable those ports again.
- If you use the Active mode, then the port channel ports automatically recover from the addition.



#### Note

When SAN port channels are created from within an interface, the **force** option cannot be used.

After the members are forcefully added, regardless of the mode (Active and On) used, the ports at either end are gracefully brought down, indicating that no frames are lost when the interface is going down.

To force the addition of a port to a SAN port channel, perform this task:

## SUMMARY STEPS

1. switch# **configuration terminal**
2. switch(config)# **interface** *type slot/port*
3. switch(config-if)# **channel-group** *channel-number* **force**

**DETAILED STEPS**

|               | <b>Command or Action</b>                                                      | <b>Purpose</b>                                                                                  |
|---------------|-------------------------------------------------------------------------------|-------------------------------------------------------------------------------------------------|
| <b>Step 1</b> | switch# <b>configuration terminal</b>                                         | Enters configuration mode.                                                                      |
| <b>Step 2</b> | switch(config)# <b>interface</b> <i>type slot/port</i>                        | Enters configuration mode for the specified interface.                                          |
| <b>Step 3</b> | switch(config-if)# <b>channel-group</b> <i>channel-number</i><br><b>force</b> | Forces the addition of the interface into the specified channel group. The E port is shut down. |

**About Interface Deletion from a SAN Port Channel**

When a physical interface is deleted from the SAN port channel, the channel membership is automatically updated. If the deleted interface is the last operational interface, then the port channel status is changed to a down state. Deleting an interface from a SAN port channel decreases the channel size and bandwidth of the SAN port channel.

- If you use the default On mode to avoid inconsistent states across switches and to maintain consistency across switches, then the ports shut down. You must explicitly enable those ports again.
- If you use the Active mode, then the port channel ports automatically recover from the deletion.

After the members are deleted, regardless of the mode (Active and On) used, the ports at either end are gracefully brought down, indicating that no frames are lost when the interface is going down.

**Deleting an Interface from a SAN Port Channel**

To delete a physical interface (or a range of physical interfaces) from a SAN port channel, perform this task:

**SUMMARY STEPS**

1. switch(config)# **interface** *type slot/port*
2. switch(config-if)# **no channel-group** *channel-number*

**DETAILED STEPS**

|               | <b>Command or Action</b>                                         | <b>Purpose</b>                                                                 |
|---------------|------------------------------------------------------------------|--------------------------------------------------------------------------------|
| <b>Step 1</b> | switch(config)# <b>interface</b> <i>type slot/port</i>           | Enters configuration mode for the specified interface.                         |
| <b>Step 2</b> | switch(config-if)# <b>no channel-group</b> <i>channel-number</i> | Deletes the physical Fibre Channel interface from the specified channel group. |

## SAN Port Channel Protocol

The switch software provides robust error detection and synchronization capabilities. You can manually configure channel groups, or they can be automatically created. In both cases, the channel groups have the same capability and configurational parameters. Any change in configuration applied to the associated SAN port channel interface is propagated to all members of the channel group.

Cisco SAN switches support a protocol to exchange SAN port channel configurations, which simplifies port channel management with incompatible ISLs. An additional autocreation mode enables ISLs with compatible parameters to automatically form channel groups without manual intervention.

The port channel protocol is enabled by default.

The port channel protocol expands the port channel functional model in Cisco SAN switches. It uses the exchange peer parameters (EPP) services to communicate across peer ports in an ISL. Each switch uses the information received from the peer ports along with its local configuration and operational values to decide if it should be part of a SAN port channel. The protocol ensures that a set of ports are eligible to be part of the same SAN port channel. They are only eligible to be part of the same port channel if all the ports have a compatible partner.

The port channel protocol uses two subprotocols:

- Bringup protocol—Automatically detects misconfigurations so you can correct them. This protocol synchronizes the SAN port channel at both ends so that all frames for a given flow (as identified by the source FC ID, destination FC ID and OX\_ID) are carried over the same physical link in both directions. This helps make applications such as write acceleration work for SAN port channels over FCIP links.
- Autocreation protocol—Automatically aggregates compatible ports into a SAN port channel.

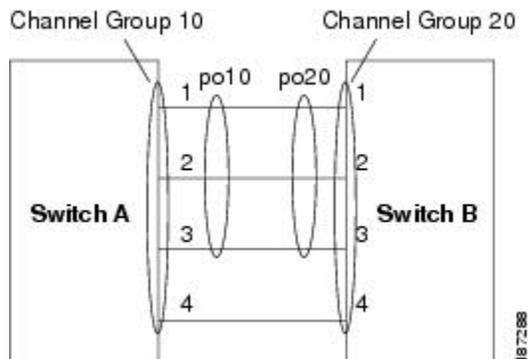
### About Channel Group Creation

If channel group autocreation is enabled, ISLs can be configured automatically into channel groups without manual intervention. The following figure shows an example of channel group autocreation.

The first ISL comes up as an individual link. In the example shown in the following figure, this is link A1-B1. When the next link comes up (A2-B2 in the example), the port channel protocol determines if this link is compatible with link A1-B1 and automatically creates channel groups 10 and 20 in the respective switches. Link A3-B3 can join the channel groups (and the port channels) if the respective ports have compatible

configurations. Link A4-B4 operates as an individual link, because it is not compatible with the existing member ports in the channel group.

**Figure 57: Autocreating Channel Groups**



The channel group numbers are assigned dynamically (when the channel group is formed).

The channel group number may change across reboots for the same set of port channels depending on the initialization order of the ports.

The following table identifies the differences between user-configured and auto-configured channel groups.

**Table 71: Channel Group Configuration Differences**

| User-Configured Channel Group                                                                                                                                                                   | Autocreating Channel Group                                                                                                                                                                                                                                                               |
|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Manually configured by the user.                                                                                                                                                                | Created automatically when compatible links come up between two compatible switches, if channel group autocreation is enabled in all ports at both ends.                                                                                                                                 |
| Member ports cannot participate in autocreation of channel groups. The autocreation feature cannot be configured.                                                                               | None of these ports are members of a user-configured channel group.                                                                                                                                                                                                                      |
| You can form the SAN port channel with a subset of the ports in the channel group. Incompatible ports remain in a suspended or isolated state depending on the On or Active mode configuration. | All ports included in the channel group participate in the SAN port channel. No member port becomes isolated or suspended; instead, the member port is removed from the channel group when the link is found to be incompatible.                                                         |
| Any administrative configuration made to the SAN port channel is applied to all ports in the channel group, and you can save the configuration for the port channel interface.                  | Any administrative configuration made to the SAN port channel is applied to all ports in the channel group, but the configurations are saved for the member ports; no configuration is saved for the port channel interface. You can explicitly convert this channel group, if required. |

| User-Configured Channel Group                                        | Autocreated Channel Group                                                                                                                                  |
|----------------------------------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------|
| You can remove any channel group and add members to a channel group. | You cannot remove a channel group. You cannot add members to the channel group or remove members. The channel group is removed when no member ports exist. |

## Autocreation Guidelines

When using the autocreation protocol, follow these guidelines:

- A port is not allowed to be configured as part of a SAN port channel when the autocreation feature is enabled. These two configurations are mutually exclusive.
- Autocreation must be enabled in both the local and peer ports to negotiate a SAN port channel.
- Aggregation occurs in one of two ways:
  - A port is aggregated into a compatible autocreated SAN port channel.
  - A port is aggregated with another compatible port to form a new SAN port channel.
- Newly created SAN port channels are allocated from the maximum possible port channel number in a decreasing order based on availability. If all port channel numbers are used up, aggregation is not allowed.
- You cannot change the membership or delete an autocreated SAN port channel.
- When you disable autocreation, all member ports are removed from the autocreated SAN port channel.
- Once the last member is removed from an autocreated SAN port channel, the channel is automatically deleted and the number is released for reuse.
- An autocreated SAN port channel is not persistent through a reboot. An autocreated SAN port channel can be manually configured to appear the same as a persistent SAN port channel. Once the SAN port channel is made persistent, the autocreation feature is disabled in all member ports.
- You can enable or disable the autocreation feature on a per-port basis or for all ports in the switch. When this configuration is enabled, the channel group mode is assumed to be active. The default for this task is disabled.
- If autocreation of channel groups is enabled for an interface, you must first disable autocreation before downgrading to earlier software versions or before configuring the interface in a manually configured channel group.



### Tip

When enabling autocreation in any switch in the Cisco Nexus 5000 Series, we recommend that you retain at least one interconnected port between the switches without any autocreation configuration. If all ports between two switches are configured with the autocreation feature at the same time, a possible traffic disruption may occur between these two switches as ports are automatically disabled and reenabled when they are added to an autocreated SAN port channel.

## Enabling and Configuring Autocreation

To configure automatic channel groups, perform this task:

### SUMMARY STEPS

1. switch# **configuration terminal**
2. switch(config)# **interface** *type slot/port*
3. switch(config-if)# **channel-group auto**
4. switch(config-if)# **no channel-group auto**

### DETAILED STEPS

|               | Command or Action                                      | Purpose                                                                                                                                 |
|---------------|--------------------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------|
| <b>Step 1</b> | switch# <b>configuration terminal</b>                  | Enters configuration mode.                                                                                                              |
| <b>Step 2</b> | switch(config)# <b>interface</b> <i>type slot/port</i> | Enters configuration mode for the specified interface.                                                                                  |
| <b>Step 3</b> | switch(config-if)# <b>channel-group auto</b>           | Automatically creates the channel group for the selected interface(s).                                                                  |
| <b>Step 4</b> | switch(config-if)# <b>no channel-group auto</b>        | Disables the autocreation of channel groups for this interface, even if the system default configuration may have autocreation enabled. |

### Example of Configuring Autocreation

The following example configures an automatic channel group:

```
switch(config)# interface fc2/3
switch(config-if)# channel-group auto
```

### About Manually Configured Channel Groups

A user-configured channel group cannot be converted to an autocreated channel group. However, you can convert an autocreated channel group to a manual channel group. This task is irreversible. The channel group number does not change, but the member ports operate according to the properties of the manually configured channel group, and channel group autocreation is implicitly disabled for all the member ports.

If you enable persistence, be sure to enable it at both ends of the SAN port channel.

### Converting to Manually Configured Channel Groups

You can convert autocreated channel group to a user-configured channel group using the **san-port-channel channel-group-number** persistent EXEC command. If the SAN port channel does not exist, this command is not executed.

## Verifying SAN Port Channel Configuration

You can view specific information about existing SAN port channels at any time from EXEC mode. The following **show** commands provide further details on existing SAN port channels.

The **show san-port-channel summary** command displays a summary of SAN port channels within the switch. A one-line summary of each SAN port channel provides the administrative state, the operational state, the number of attached and active interfaces (up), and the first operational port (FOP), which is the primary operational interface selected in the SAN port channel to carry control-plane traffic (no load-balancing). The FOP is the first port that comes up in a SAN port channel and can change if the port goes down. The FOP is also identified by an asterisk ( \* ).

To display VSAN configuration information, perform one of the following tasks:

### SUMMARY STEPS

1. switch# **show san-port-channel summary | database | consistency [ details ] | usage | compatibility-parameters**
2. switch# **show san-port-channel database interface san-port-channel *channel-number***
3. switch# switch# **show interface fc *slot/port***

### DETAILED STEPS

|               | Command or Action                                                                                                    | Purpose                                                                            |
|---------------|----------------------------------------------------------------------------------------------------------------------|------------------------------------------------------------------------------------|
| <b>Step 1</b> | switch# <b>show san-port-channel summary   database   consistency [ details ]   usage   compatibility-parameters</b> | Displays SAN port channel information.                                             |
| <b>Step 2</b> | switch# <b>show san-port-channel database interface san-port-channel <i>channel-number</i></b>                       | Displays information for the specified SAN port channel.                           |
| <b>Step 3</b> | switch# switch# <b>show interface fc <i>slot/port</i></b>                                                            | Displays VSAN configuration information for the specified Fibre Channel interface. |

### Example of Verification Commands

The following example shows how to display a summary of SAN port channel information:

```
switch# show san-port-channel summary

Interface Total Ports Oper Ports First Oper Port

san-port-channel 7 2 0 --
san-port-channel 8 2 0 --
san-port-channel 9 2 2
```

The following example shows how to display SAN port channel consistency:

```
switch# show san-port-channel consistency
Database is consistent
```

The following example shows how to display details of the used and unused port channel numbers:

```
switch# show san-port-channel usage
Totally 3 port-channel numbers used
=====
Used : 77 - 79
Unused: 1 - 76 , 80 - 256
```

Autocreated SAN port channels are indicated explicitly to help differentiate them from the manually created SAN port channels. The following example shows how to display an autocreated port channel:

```
switch# show interface fc2/1
fc2/1 is trunking
 Hardware is Fibre Channel, FCOT is short wave laser
 Port WWN is 20:0a:00:0b:5f:3b:fe:80
 ...
 Receive data field Size is 2112
 Beacon is turned off
 Port-channel auto creation is enabled
Belongs to port-channel 123
...
```

## Default Settings for SAN Port Channels

The table below lists the default settings for SAN port channels.

**Table 72: Default SAN Port Channel Parameters**

| Parameters                | Default                     |
|---------------------------|-----------------------------|
| Port channels             | FSPF is enabled by default. |
| Create port channel       | Administratively up.        |
| Default port channel mode | On.                         |
| Autocreation              | Disabled.                   |





## Configuring and Managing VSANs

This chapter contains the following sections:

- [Configuring and Managing VSANs](#), page 589

### Configuring and Managing VSANs

You can achieve higher security and greater stability in Fibre Channel fabrics by using virtual SANs (VSANs). VSANs provide isolation among devices that are physically connected to the same fabric. With VSANs you can create multiple logical SANs over a common physical infrastructure. Each VSAN can contain up to 239 switches and has an independent address space that allows identical Fibre Channel IDs (FC IDs) to be used simultaneously in different VSANs.

#### Information About VSANs

A VSAN is a virtual storage area network (SAN). A SAN is a dedicated network that interconnects hosts and storage devices primarily to exchange SCSI traffic. In SANs you use the physical links to make these interconnections. A set of protocols run over the SAN to handle routing, naming, and zoning. You can design multiple SANs with different topologies.

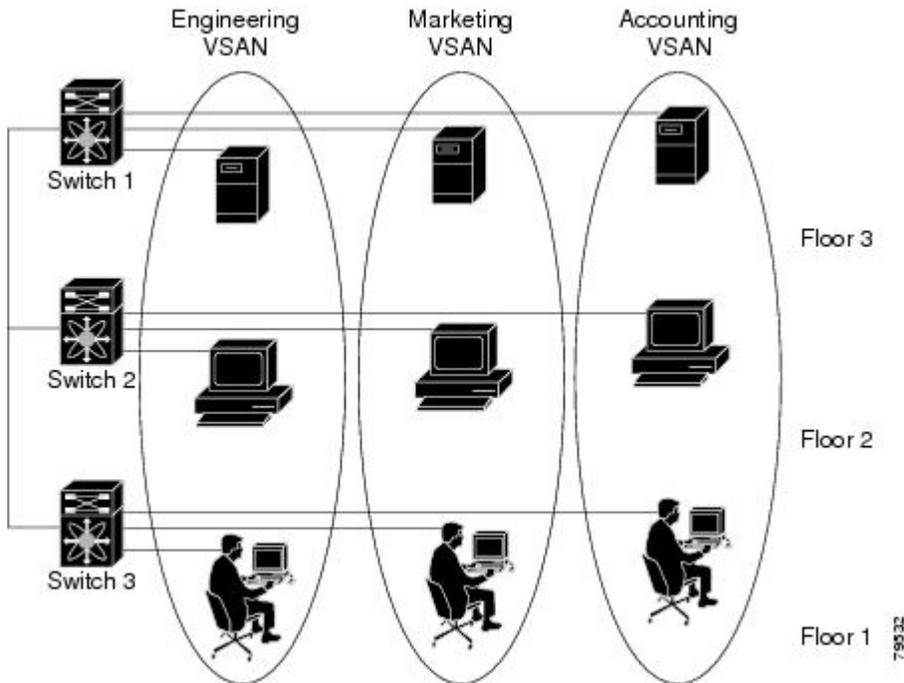
#### VSAN Topologies

With the introduction of VSANs, the network administrator can build a single topology containing switches, links, and one or more VSANs. Each VSAN in this topology has the same operation and property of a SAN. A VSAN has the following additional features:

- Multiple VSANs can share the same physical topology.
- The same Fibre Channel IDs (FC IDs) can be assigned to a host in another VSAN, which increases VSAN scalability.
- Every instance of a VSAN runs all required protocols such as FSPF, domain manager, and zoning.
- Fabric-related configurations in one VSAN do not affect the associated traffic in another VSAN.
- Events causing traffic disruptions in one VSAN are contained within that VSAN and are not propagated to other VSANs.

The following figure shows a fabric with three switches, one on each floor. The geographic location of the switches and the attached devices is independent of their segmentation into logical VSANs. No communication between VSANs is possible. Within each VSAN, all members can talk to one another.

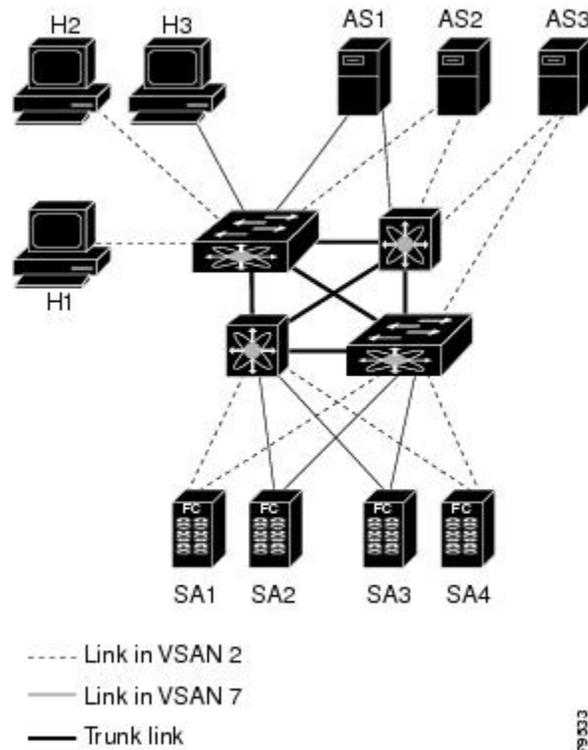
**Figure 58: Logical VSAN Segmentation**



The application servers or storage arrays can be connected to the switch using Fibre Channel or virtual Fibre Channel interfaces. A VSAN can include a mixture of Fibre Channel and virtual Fibre Channel interfaces.

The following figure shows a physical Fibre Channel switching infrastructure with two defined VSANs: VSAN 2 (dashed) and VSAN 7 (solid). VSAN 2 includes hosts H1 and H2, application servers AS2 and AS3, and storage arrays SA1 and SA4. VSAN 7 connects H3, AS1, SA2, and SA3.

**Figure 59: Example of Two VSANs**



The four switches in this network are interconnected by VSAN trunk links that carry both VSAN 2 and VSAN 7 traffic. You can configure a different inter-switch topology for each VSAN. In the preceding figure, the inter-switch topology is identical for VSAN 2 and VSAN 7.

Without VSANs, a network administrator would need separate switches and links for separate SANs. By enabling VSANs, the same switches and links may be shared by multiple VSANs. VSANs allow SANs to be built on port granularity instead of switch granularity. The preceding figure illustrates that a VSAN is a group of hosts or storage devices that communicate with each other using a virtual topology defined on the physical SAN.

The criteria for creating such groups differ based on the VSAN topology:

- VSANs can separate traffic based on the following requirements:
  - Different customers in storage provider data centers
  - Production or test in an enterprise network
  - Low and high security requirements
  - Backup traffic on separate VSANs
  - Replicating data from user traffic

- VSANs can meet the needs of a particular department or application.

### VSAN Advantages

VSANs offer the following advantages:

- Traffic isolation—Traffic is contained within VSAN boundaries and devices reside only in one VSAN ensuring absolute separation between user groups, if desired.
- Scalability—VSANs are overlaid on top of a single physical fabric. The ability to create several logical VSAN layers increases the scalability of the SAN.
- Per VSAN fabric services—Replication of fabric services on a per VSAN basis provides increased scalability and availability.
- Redundancy—Several VSANs created on the same physical SAN ensure redundancy. If one VSAN fails, redundant protection (to another VSAN in the same physical SAN) is configured using a backup path between the host and the device.
- Ease of configuration—Users can be added, moved, or changed between VSANs without changing the physical structure of a SAN. Moving a device from one VSAN to another only requires configuration at the port level, not at a physical level.

Up to 256 VSANs can be configured in a switch. Of these, one is a default VSAN (VSAN 1), and another is an isolated VSAN (VSAN 4094). User-specified VSAN IDs range from 2 to 4093.

### VSANs Versus Zones

Zones are always contained within a VSAN. You can define multiple zones in a VSAN.

Because two VSANs are equivalent to two unconnected SANs, zone A on VSAN 1 is different and separate from zone A in VSAN 2. The following table lists the differences between VSANs and zones.

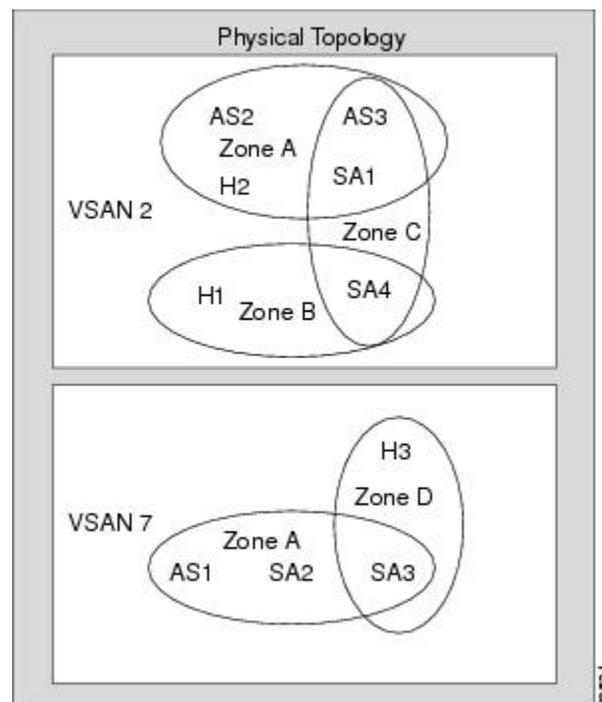
**Table 73: VSAN and Zone Comparison**

| VSAN Characteristic                                                                                | Zone Characteristic                                                                 |
|----------------------------------------------------------------------------------------------------|-------------------------------------------------------------------------------------|
| VSANs equal SANs with routing, naming, and zoning protocols.                                       | Routing, naming, and zoning protocols are not available on a per-zone basis.        |
| VSANs limit unicast, multicast, and broadcast traffic.                                             | Zones limit unicast traffic.                                                        |
| Membership is typically defined using the VSAN ID to F ports.                                      | Membership is typically defined by the pWWN.                                        |
| An HBA or a storage device can belong only to a single VSAN (the VSAN associated with the F port). | An HBA or storage device can belong to multiple zones.                              |
| VSANs enforce membership at each E port, source port, and destination port.                        | Zones enforce membership only at the source and destination ports.                  |
| VSANs are defined for larger environments (storage service providers).                             | Zones are defined for a set of initiators and targets not visible outside the zone. |

| VSAN Characteristic                | Zone Characteristic                      |
|------------------------------------|------------------------------------------|
| VSANs encompass the entire fabric. | Zones are configured at the fabric edge. |

The following figure shows the possible relationships between VSANs and zones. In VSAN 2, three zones are defined: zone A, zone B, and zone C. Zone C overlaps both zone A and zone B as permitted by Fibre Channel standards. In VSAN 7, two zones are defined: zone A and zone D. No zone crosses the VSAN boundary. Zone A defined in VSAN 2 is different and separate from zone A defined in VSAN 7.

**Figure 60: VSANS with Zoning**



## Configuring VSANs

VSANs have the following attributes:

- **VSAN ID**—The VSAN ID identifies the VSAN as the default VSAN (VSAN 1), user-defined VSANs (VSAN 2 to 4093), and the isolated VSAN (VSAN 4094).
- **State**—The administrative state of a VSAN can be configured to an active (default) or suspended state. Once VSANs are created, they may exist in various conditions or states.
  - The active state of a VSAN indicates that the VSAN is configured and enabled. By enabling a VSAN, you activate the services for that VSAN.
  - The suspended state of a VSAN indicates that the VSAN is configured but not enabled. If a port is configured in this VSAN, it is disabled. Use this state to deactivate a VSAN without losing the

VSAN's configuration. All ports in a suspended VSAN are disabled. By suspending a VSAN, you can preconfigure all the VSAN parameters for the whole fabric and activate the VSAN immediately.

- VSAN name—This text string identifies the VSAN for management purposes. The name can be from 1 to 32 characters long and it must be unique across all VSANs. By default, the VSAN name is a concatenation of VSAN and a four-digit string representing the VSAN ID. For example, the default name for VSAN 3 is VSAN0003.



**Note** A VSAN name must be unique.

- Load-balancing attributes—These attributes indicate the use of the source-destination ID (src-dst-id) or the originator exchange OX ID (src-dst-ox-id, the default) for load-balancing path selection.

### About VSAN Creation

A VSAN is in the operational state if the VSAN is active and at least one port is up. This state indicates that traffic can pass through this VSAN. This state cannot be configured.

### Creating VSANs Statically

You cannot configure any application-specific parameters for a VSAN before creating the VSAN.

To create VSANs, perform this task:

### SUMMARY STEPS

1. switch# **configuration terminal**
2. switch(config)# **vsan database**
3. switch(config-vsan-db)# **vsan vsan-id**
4. switch(config-vsan-db)# **vsan vsan-id name name**
5. switch(config-vsan-db)# **vsan vsan-id suspend**
6. switch(config-vsan-db)# **no vsan vsan-id suspend**
7. switch(config-vsan-db)# **end**

### DETAILED STEPS

|               | Command or Action                           | Purpose                                                                                                         |
|---------------|---------------------------------------------|-----------------------------------------------------------------------------------------------------------------|
| <b>Step 1</b> | switch# <b>configuration terminal</b>       | Enters configuration mode.                                                                                      |
| <b>Step 2</b> | switch(config)# <b>vsan database</b>        | Configures the database for a VSAN. Application specific VSAN parameters cannot be configured from this prompt. |
| <b>Step 3</b> | switch(config-vsan-db)# <b>vsan vsan-id</b> | Creates a VSAN with the specified ID if that VSAN does not exist already.                                       |

|               | Command or Action                                                          | Purpose                                                         |
|---------------|----------------------------------------------------------------------------|-----------------------------------------------------------------|
| <b>Step 4</b> | switch(config-vsan-db)# <b>vsan</b> <i>vsan-id</i> <b>name</b> <i>name</i> | Updates the VSAN with the assigned name.                        |
| <b>Step 5</b> | switch(config-vsan-db)# <b>vsan</b> <i>vsan-id</i> <b>suspend</b>          | Suspends the selected VSAN.                                     |
| <b>Step 6</b> | switch(config-vsan-db)# <b>no vsan</b> <i>vsan-id</i> <b>suspend</b>       | Negates the <b>suspend</b> command issued in the previous step. |
| <b>Step 7</b> | switch(config-vsan-db)# <b>end</b>                                         | Returns you to EXEC mode.                                       |

### About Port VSAN Membership

Port VSAN membership on the switch is assigned on a port-by-port basis. By default each port belongs to the default VSAN. You can assign VSAN membership to ports using one of two methods:

- Statically—Assigning VSANs to ports.
- Dynamically—Assigning VSANs based on the device WWN. This method is referred to as dynamic port VSAN membership (DPVM). Cisco Nexus 5000 Series switches do not support DPVM.

VSAN trunking ports have an associated list of VSANs that are part of an allowed list.

### Related Topics

[Assigning Static Port VSAN Membership, on page 595](#)

[Configuring VSAN Trunking, on page 563](#)

### Assigning Static Port VSAN Membership

To statically assign VSAN membership for an interface port, perform this task:

#### SUMMARY STEPS

1. switch# **configuration terminal**
2. switch(config)# **vsan database**
3. switch(config-vsan-db)# **vsan** *vsan-id*
4. switch(config-vsan-db)# **vsan** *vsan-id* **interface** {*fc slot/port* | **vfc** *vfc-id*}
5. switch(config-vsan-db)# **vsan** *vsan-id* {*fc slot/port* | **vfc** *vfc-id*}
6. switch(config-vsan-db)# **no vsan** *vsan-id* {*fc slot/port* | **vfc** *vfc-id*}

#### DETAILED STEPS

|               | Command or Action                     | Purpose                    |
|---------------|---------------------------------------|----------------------------|
| <b>Step 1</b> | switch# <b>configuration terminal</b> | Enters configuration mode. |

|               | Command or Action                                                                 | Purpose                                                                          |
|---------------|-----------------------------------------------------------------------------------|----------------------------------------------------------------------------------|
| <b>Step 2</b> | switch(config)# <b>vsan database</b>                                              | Configures the database for a VSAN.                                              |
| <b>Step 3</b> | switch(config-vsan-db)# <b>vsan vsan-id</b>                                       | Creates a VSAN with the specified ID if that VSAN does not exist already.        |
| <b>Step 4</b> | switch(config-vsan-db)# <b>vsan vsan-id interface</b> {fc slot/port   vfc vfc-id} | Assigns the membership of the specified interface to the VSAN.                   |
| <b>Step 5</b> | switch(config-vsan-db)# <b>vsan vsan-id</b> {fc slot/port   vfc vfc-id}           | Updates the membership information of the interface to reflect the changed VSAN. |
| <b>Step 6</b> | switch(config-vsan-db)# <b>no vsan vsan-id</b> {fc slot/port   vfc vfc-id}        | Removes the interface from the VSAN.                                             |

### Displaying VSAN Static Membership

To display the VSAN static membership information, use the **show vsan membership** command.

The following example displays membership information for the specified VSAN:

```
switch # show vsan 1 membership
vsan 1 interfaces:
 fc2/1 fc2/2 fc2/3 fc2/4
 san-port-channel 3 vfc1/1
```



**Note** Interface information is not displayed if interfaces are not configured on this VSAN.

The following example displays membership information for all VSANs:

```
switch # show vsan membership
vsan 1 interfaces:
 fc2/1 fc2/2 fc2/3 fc2/4
 san-port-channel 3 vfc3/1
vsan 2 interfaces:
 fc2/3 vfc4/1
vsan 7 interfaces:
vsan 100 interfaces:
vsan 4094(isolated vsan) interfaces:
```

The following example displays static membership information for the specified interface:

```
switch # show vsan membership interface fc2/1
fc2/1
 vsan:1
 allowed list:1-4093
```

### About the Default VSAN

The factory settings for switches in the Cisco Nexus 5000 Series have only the default VSAN 1 enabled. We recommend that you do not use VSAN 1 as your production environment VSAN. If no VSANs are configured, all devices in the fabric are considered part of the default VSAN. By default, all ports are assigned to the default VSAN.




---

**Note** VSAN 1 cannot be deleted, but it can be suspended.

Up to 256 VSANs can be configured in a switch. Of these, one is a default VSAN (VSAN 1), and another is an isolated VSAN (VSAN 4094). User-specified VSAN IDs range from 2 to 4093.

---

### About the Isolated VSAN

VSAN 4094 is an isolated VSAN. When a VSAN is deleted, all nontrunking ports are transferred to the isolated VSAN to avoid an implicit transfer of ports to the default VSAN or to another configured VSAN. This action ensures that all ports in the deleted VSAN become isolated (disabled).




---

**Note** When you configure a port in VSAN 4094 or move a port to VSAN 4094, that port is immediately isolated.

---




---

**Caution** Do not use an isolated VSAN to configure ports.

---




---

**Note** Up to 256 VSANs can be configured in a switch. Of these, one is a default VSAN (VSAN 1), and another is an isolated VSAN (VSAN 4094). User-specified VSAN IDs range from 2 to 4093.

---

### Displaying Isolated VSAN Membership

The **show vsan 4094 membership** command displays all ports associated with the isolated VSAN.

### Operational State of a VSAN

A VSAN is in the operational state if the VSAN is active and at least one port is up. This state indicates that traffic can pass through this VSAN. This state cannot be configured.

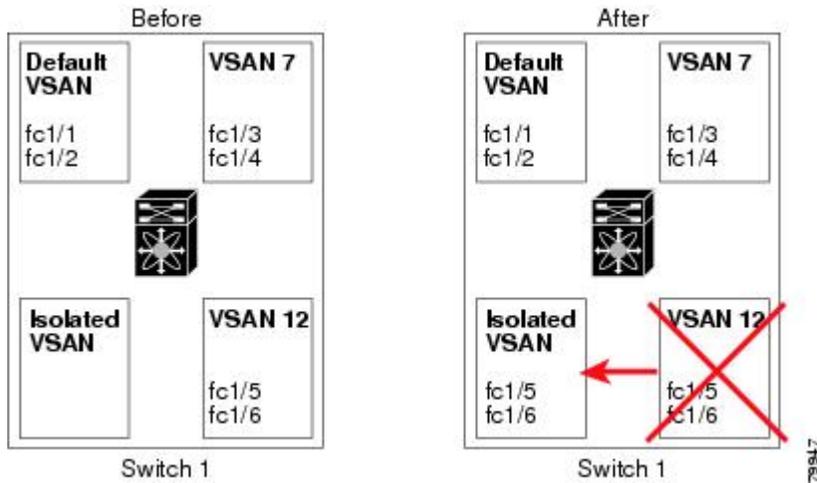
### About Static VSAN Deletion

When an active VSAN is deleted, all of its attributes are removed from the running configuration. VSAN-related information is maintained by the system software as follows:

- VSAN attributes and port membership details are maintained by the VSAN manager. This feature is affected when you delete a VSAN from the configuration. When a VSAN is deleted, all the ports in that VSAN are made inactive and the ports are moved to the isolated VSAN. If the same VSAN is recreated,

the ports do not automatically get assigned to that VSAN. You must explicitly reconfigure the port VSAN membership (see the figure below).

**Figure 61: VSAN Port Membership Details**



- VSAN-based runtime (name server), zoning, and configuration (static routes) information is removed when the VSAN is deleted.
- Configured VSAN interface information is removed when the VSAN is deleted.



**Note** The allowed VSAN list is not affected when a VSAN is deleted.

Any commands for a nonconfigured VSAN are rejected. For example, if VSAN 10 is not configured in the system, then a command request to move a port to VSAN 10 is rejected.

**Related Topics**

[Configuring VSAN Trunking, on page 563](#)

**Deleting Static VSANs**

To delete a VSAN and its various attributes, perform this task:

**SUMMARY STEPS**

1. switch# **configuration terminal**
2. switch(config)# **vsan database**
3. switch-config-db# **vsan 2**
4. switch(config-vsan-db)# **no vsan 5**
5. switch(config-vsan-db)# **end**

**DETAILED STEPS**

|               | <b>Command or Action</b>                 | <b>Purpose</b>                               |
|---------------|------------------------------------------|----------------------------------------------|
| <b>Step 1</b> | switch# <b>configuration terminal</b>    | Enters configuration mode.                   |
| <b>Step 2</b> | switch(config)# <b>vsan database</b>     | Configures the VSAN database.                |
| <b>Step 3</b> | switch-config-db# <b>vsan 2</b>          | Places you in VSAN configuration mode.       |
| <b>Step 4</b> | switch(config-vsan-db)# <b>no vsan 5</b> | Deletes VSAN 5 from the database and switch. |
| <b>Step 5</b> | switch(config-vsan-db)# <b>end</b>       | Places you in EXEC mode.                     |

**About Load Balancing**

Load-balancing attributes indicate the use of the source-destination ID (src-dst-id) or the originator exchange OX ID (src-dst-ox-id, the default) for load-balancing path selection.

**Configuring Load Balancing**

To configure load balancing on an existing VSAN, perform this task:

**SUMMARY STEPS**

1. switch# **configuration terminal**
2. switch(config)# **vsan database**
3. switch(config-vsan-db)# **vsan vsan-id**
4. switch(config-vsan-db)# **vsan vsan-id loadbalancing src-dst-id**
5. switch(config-vsan-db)# **no vsan vsan-id loadbalancing src-dst-id**
6. switch(config-vsan-db)# **vsan vsan-id loadbalancing src-dst-ox-id**
7. switch(config-vsan-db)# **vsan vsan-id suspend**
8. switch(config-vsan-db)# **no vsan vsan-id suspend**
9. switch(config-vsan-db)# **end**

**DETAILED STEPS**

|               | <b>Command or Action</b>                    | <b>Purpose</b>                             |
|---------------|---------------------------------------------|--------------------------------------------|
| <b>Step 1</b> | switch# <b>configuration terminal</b>       | Enters configuration mode.                 |
| <b>Step 2</b> | switch(config)# <b>vsan database</b>        | Enters VSAN database configuration submode |
| <b>Step 3</b> | switch(config-vsan-db)# <b>vsan vsan-id</b> | Specifies an existing VSAN.                |

|               | Command or Action                                                                        | Purpose                                                                                                                                                |
|---------------|------------------------------------------------------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Step 4</b> | switch(config-vsan-db)# <b>vsan</b> <i>vsan-id</i><br><b>loadbalancing src-dst-id</b>    | Enables the load-balancing guarantee for the selected VSAN and directs the switch to use the source and destination ID for its path selection process. |
| <b>Step 5</b> | switch(config-vsan-db)# <b>no vsan</b> <i>vsan-id</i><br><b>loadbalancing src-dst-id</b> | Negates the command entered in the previous step and reverts to the default values of the load-balancing parameters.                                   |
| <b>Step 6</b> | switch(config-vsan-db)# <b>vsan</b> <i>vsan-id</i><br><b>loadbalancing src-dst-ox-id</b> | Changes the path selection setting to use the source ID, the destination ID, and the OX ID (default).                                                  |
| <b>Step 7</b> | switch(config-vsan-db)# <b>vsan</b> <i>vsan-id</i> <b>suspend</b>                        | Suspends the selected VSAN.                                                                                                                            |
| <b>Step 8</b> | switch(config-vsan-db)# <b>no vsan</b> <i>vsan-id</i> <b>suspend</b>                     | Negates the <b>suspend</b> command entered in the previous step.                                                                                       |
| <b>Step 9</b> | switch(config-vsan-db)# <b>end</b>                                                       | Returns you to EXEC mode.                                                                                                                              |

### About Interop Mode

Interoperability enables the products of multiple vendors to connect with each other. Fibre Channel standards guide vendors to create common external Fibre Channel interfaces.

#### Related Topics

[Switch Interoperability](#), on page 688

## Displaying Static VSAN Configuration

The following example shows how to display information about a specific VSAN:

```
switch# show vsan 100
```

The following example shows how to display VSAN usage:

```
switch# show vsan usage
4 vsan configured
configured vsans:1-4
vsans available for configuration:5-4093
```

The following example shows how to display all VSANs:

```
switch# show vsan
```

## Default VSAN Settings

The following table lists the default settings for all configured VSANs.

**Table 74: Default VSAN Parameters**

| Parameters   | Default |
|--------------|---------|
| Default VSAN | VSAN 1. |

| <b>Parameters</b>        | <b>Default</b>                                                                                           |
|--------------------------|----------------------------------------------------------------------------------------------------------|
| State                    | Active state.                                                                                            |
| Name                     | Concatenation of VSAN and a four-digit string representing the VSAN ID. For example, VSAN 3 is VSAN0003. |
| Load-balancing attribute | OX ID (src-dst-ox-id).                                                                                   |





## Configuring and Managing Zones

---

This chapter contains the following sections:

- [Configuring and Managing Zones](#), page 603

### Configuring and Managing Zones

Zoning enables you to set up access control between storage devices or user groups. If you have administrator privileges in your fabric, you can create zones to increase network security and to prevent data loss or corruption. Zoning is enforced by examining the source-destination ID field.

Advanced zoning capabilities specified in the FC-GS-4 and FC-SW-3 standards are supported. You can use either the existing basic zoning capabilities or the advanced, standards-compliant zoning capabilities.

#### Information About Zoning

##### Zoning Features

Zoning includes the following features:

- A zone consists of multiple zone members.
  - Members in a zone can access each other; members in different zones cannot access each other.
  - If zoning is not activated, all devices are members of the default zone.
  - If zoning is activated, any device that is not in an active zone (a zone that is part of an active zone set) is a member of the default zone.
  - Zones can vary in size.
  - Devices can belong to more than one zone.
  - A physical fabric can have a maximum of 16,000 members. This includes all VSANs in the fabric.
- A zone set consists of one or more zones.
  - A zone set can be activated or deactivated as a single entity across all switches in the fabric.

- Only one zone set can be activated at any time.
- A zone can be a member of more than one zone set.
- A zone switch can have a maximum of 500 zone sets.
- Zoning can be administered from any switch in the fabric.
  - When you activate a zone (from any switch), all switches in the fabric receive the active zone set. Additionally, full zone sets are distributed to all switches in the fabric, if this feature is enabled in the source switch.
  - If a new switch is added to an existing fabric, zone sets are acquired by the new switch.
- Zone changes can be configured nondisruptively.
  - New zones and zone sets can be activated without interrupting traffic on unaffected ports or devices.
- Zone membership can be specified using the following identifiers:
  - Port world wide name (pWWN)—Specifies the pWWN of an N port attached to the switch as a member of the zone.
  - Fabric pWWN—Specifies the WWN of the fabric port (switch port's WWN). This membership is also referred to as port-based zoning.
  - FC ID—Specifies the FC ID of an N port attached to the switch as a member of the zone.
  - Interface and switch WWN (sWWN)—Specifies the interface of a switch identified by the sWWN. This membership is also referred to as interface-based zoning.
  - Interface and domain ID—Specifies the interface of a switch identified by the domain ID.
  - Domain ID and port number—Specifies the domain ID of a Cisco switch domain and additionally specifies a port belonging to a non-Cisco switch.

**Note**


---

For N ports attached to the switch over a virtual Fibre Channel interface, you can specify zone membership using the pWWN of the N port, the FC ID of the N port, or the fabric pWWN of the virtual Fibre Channel interface.

---

- Default zone membership includes all ports or WWNs that do not have a specific membership association. Access between default zone members is controlled by the default zone policy.
- You can configure up to 8000 zones per VSAN and a maximum of 8000 zones for all VSANs on the switch.

**Note**


---

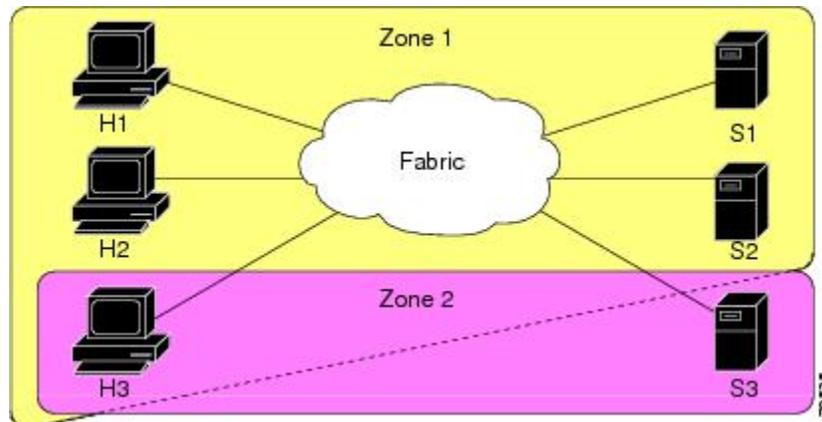
Interface-based zoning only works with Cisco SAN switches. Interface-based zoning does not work for VSANs configured in interop mode.

---

### Zoning Example

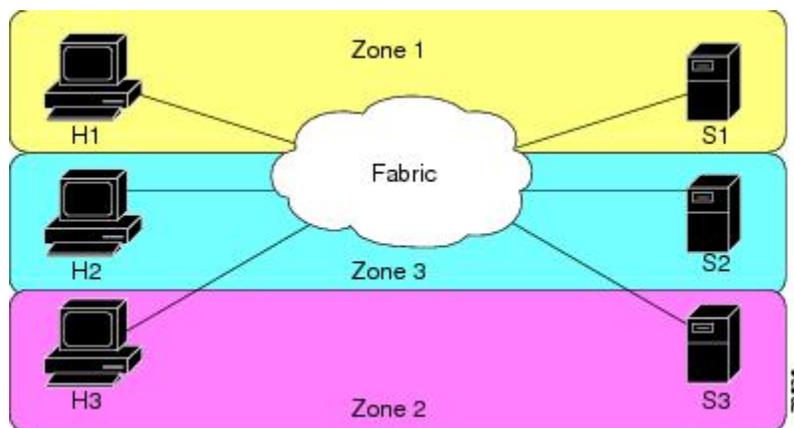
The following figure shows a zone set with two zones, zone 1 and zone 2, in a fabric. Zone 1 provides access from all three hosts (H1, H2, H3) to the data residing on storage systems S1 and S2. Zone 2 restricts the data on S3 to access only by H3. H3 resides in both zones.

**Figure 62: Fabric with Two Zones**



You can use other ways to partition this fabric into zones. The following figure shows another possibility. Assume that there is a need to isolate storage system S2 for the purpose of testing new software. To achieve this, zone 3 is configured, which contains only host H2 and storage S2. You can restrict access to only H2 and S2 in zone 3, and to H1 and S1 in zone 1.

**Figure 63: Fabric with Three Zones**



### Zone Implementation

Cisco Nexus 5000 Series switches automatically support the following basic zone features (no additional configuration is required):

- Zones are contained in a VSAN.

- Hard zoning cannot be disabled.
- Name server queries are soft-zoned.
- Only active zone sets are distributed.
- Unzoned devices cannot access each other.
- A zone or zone set with the same name can exist in each VSAN.
- Each VSAN has a full database and an active database.
- Active zone sets cannot be changed, without activating a full zone database.
- Active zone sets are preserved across switch reboots.
- Changes to the full database must be explicitly saved.
- Zone reactivation (a zone set is active and you activate another zone set) does not disrupt existing traffic.

If required, you can additionally configure the following zone features:

- Propagate full zone sets to all switches on a per VSAN basis.
- Change the default policy for unzoned members.
- Interoperate with other vendors by configuring a VSAN in the interop mode. You can also configure one VSAN in the interop mode and another VSAN in the basic mode in the same switch without disrupting each other.
- Bring E ports out of isolation.

### Active and Full Zone Set Configuration Guidelines

Before configuring a zone set, consider the following guidelines:

- Each VSAN can have multiple zone sets but only one zone set can be active at any given time.
- When you create a zone set, that zone set becomes a part of the full zone set.
- When you activate a zone set, a copy of the zone set from the full zone set is used to enforce zoning, and is called the active zone set. An active zone set cannot be modified. A zone that is part of an active zone set is called an active zone.
- The administrator can modify the full zone set even if a zone set with the same name is active. However, the modification will be enforced only upon reactivation.
- When the activation is done, the active zone set is automatically stored in persistent configuration. This enables the switch to preserve the active zone set information across switch resets.
- All other switches in the fabric receive the active zone set so they can enforce zoning in their respective switches.
- Hard and soft zoning are implemented using the active zone set. Modifications take effect during zone set activation.
- An FC ID or Nx port that is not part of the active zone set belongs to the default zone and the default zone information is not distributed to other switches.

**Note**

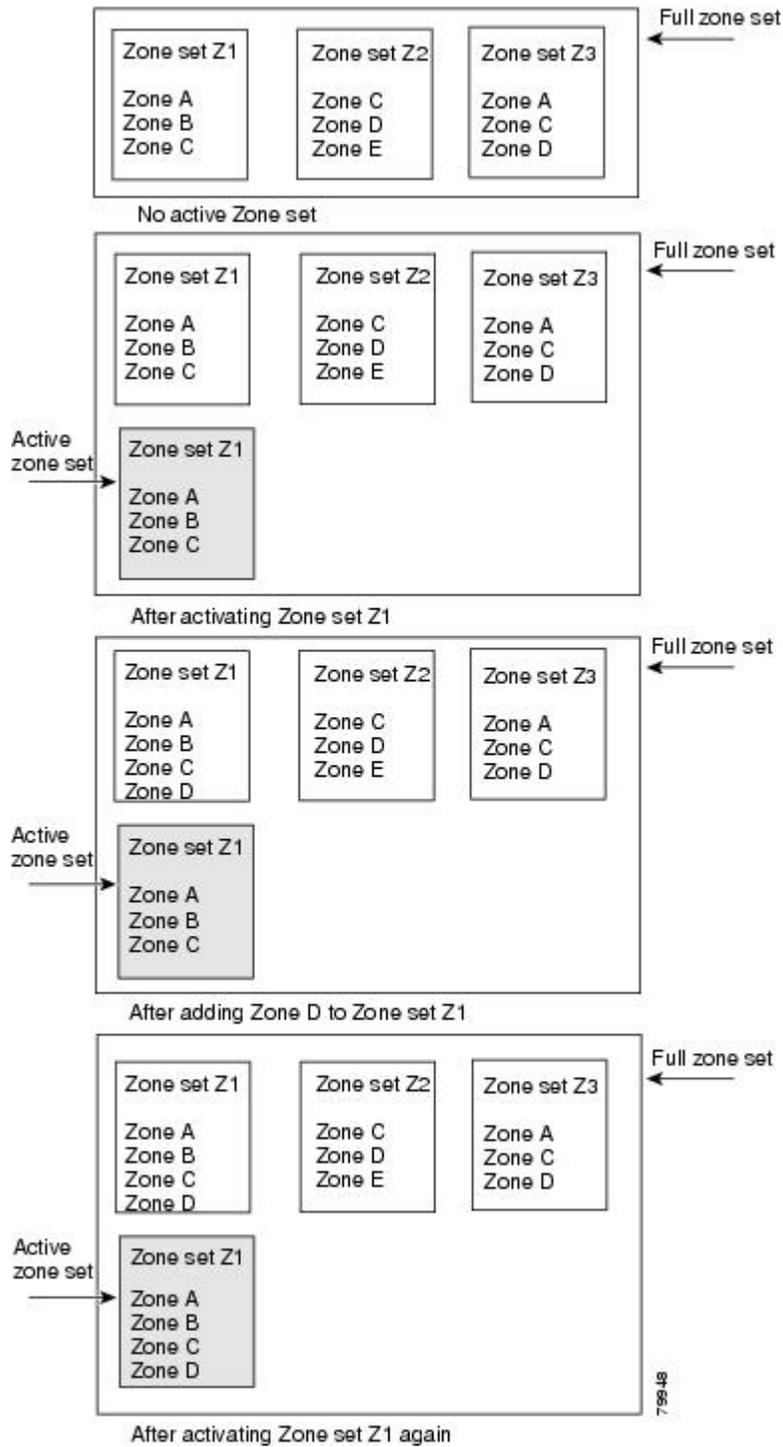
---

If one zone set is active and you activate another zone set, the currently active zone set is automatically deactivated. You do not need to explicitly deactivate the currently active zone set before activating a new zone set.

---

The following figure shows a zone being added to an activated zone set.

**Figure 64: Active and Full Zone Sets**



## Configuring Zones

To configure a zone and assign a zone name, perform this task:

### SUMMARY STEPS

1. switch# **configuration terminal**
2. switch(config)# **zone name** *zone-name* **vsan** *vsan-id*
3. switch(config-zone)# **member** *type value*

### DETAILED STEPS

|               | Command or Action                                                               | Purpose                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
|---------------|---------------------------------------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Step 1</b> | switch# <b>configuration terminal</b>                                           | Enters configuration mode.                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| <b>Step 2</b> | switch(config)# <b>zone name</b> <i>zone-name</i><br><b>vsan</b> <i>vsan-id</i> | Configures a zone in the specified VSAN.<br><b>Note</b> All alphanumeric characters or one of the following symbols (\$, -, ^, _) are supported.                                                                                                                                                                                                                                                                                                                                          |
| <b>Step 3</b> | switch(config-zone)# <b>member</b> <i>type value</i>                            | Configures a member for the specified zone based on the type (pWWN, fabric pWWN, FC ID, fcalias, domain ID, or interface) and value specified.<br><b>Caution</b> You must only configure pWWN-type zoning on all SAN switches running Cisco NX-OS if there is a Cisco MDS 9020 switch running FabricWare in the same fabric.<br><b>Tip</b> Use a relevant display command (for example, <b>show interface</b> or <b>show flogi database</b> ) to obtain the required value in hex format. |

### Configuring Zones Example

**Table 75: Type and Value Syntax for the member Command**

|                       |                                                                                 |
|-----------------------|---------------------------------------------------------------------------------|
| Domain ID             | <b>member domain-id</b> <i>domain-id</i> <b>portnumber</b> <i>number</i>        |
| FC alias              | <b>member fcalias</b> <i>fc-alias-name</i>                                      |
| FC ID                 | <b>member fcid</b> <i>fcid</i>                                                  |
| Fabric pWWN           | <b>member fwwn</b> <i>fwwn-id</i>                                               |
| Local sWWN interface  | <b>member interface</b> <i>type slot/port</i>                                   |
| Domain ID interface   | <b>member interface</b> <i>type slot/port</i> <b>domain-id</b> <i>domain-id</i> |
| Remote sWWN interface | <b>member interface</b> <i>type slot/port</i> <b>swwn</b> <i>swwn-id</i>        |

pWWN

member pwwn *pwwn-id***Tip**

Use the **show wwn switch** command to retrieve the sWWN. If you do not provide a sWWN, the software automatically uses the local sWWN.

The following examples show how to configure zone members:

```
switch(config)# zone name MyZone vsan 2
```

pWWN example:

```
switch(config-zone)# member pwwn 10:00:00:23:45:67:89:ab
```

Fabric pWWN example:

```
switch(config-zone)# member fwwn 10:01:10:01:10:ab:cd:ef
```

FC ID example:

```
switch(config-zone)# member fcid 0xce00d1
```

FC alias example:

```
switch(config-zone)# member fcalias Payroll
```

Domain ID example:

```
switch(config-zone)# member domain-id 2 portnumber 23
```

Local sWWN interface example:

```
switch(config-zone)# member interface fc 2/1
```

Remote sWWN interface example:

```
switch(config-zone)# member interface fc 2/1 swwn 20:00:00:05:30:00:4a:de
```

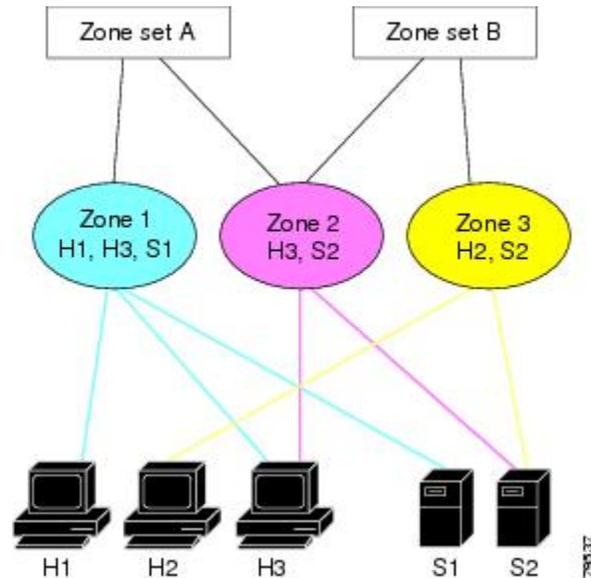
Domain ID interface example:

```
switch(config-zone)# member interface fc 2/1 domain-id 25
```

## Zone Sets

In the following figure, two separate sets are created, each with its own membership hierarchy and zone members.

**Figure 65: Hierarchy of Zone Sets, Zones, and Zone Members**



Zones provide a method for specifying access control, while zone sets are a grouping of zones to enforce access control in the fabric. Either zone set A or zone set B can be activated (but not together).



**Tip** Zone sets are configured with the names of the member zones and the VSAN (if the zone set is in a configured VSAN).

### Activating a Zone Set

Changes to a zone set do not take effect in a full zone set until you activate it.

To activate or deactivate an existing zone set, perform this task:

### SUMMARY STEPS

1. switch# **configuration terminal**
2. switch(config)# **zoneset activate name zoneset-name vsan vsan-id**
3. switch(config)# **no zoneset activate name zoneset-name vsan vsan-id**

## DETAILED STEPS

|               | Command or Action                                                                                 | Purpose                             |
|---------------|---------------------------------------------------------------------------------------------------|-------------------------------------|
| <b>Step 1</b> | switch# <b>configuration terminal</b>                                                             | Enters configuration mode.          |
| <b>Step 2</b> | switch(config)# <b>zoneset activate name</b> <i>zoneset-name</i> <b>vsan</b><br><i>vsan-id</i>    | Activates the specified zone set.   |
| <b>Step 3</b> | switch(config)# <b>no zoneset activate name</b> <i>zoneset-name</i> <b>vsan</b><br><i>vsan-id</i> | Deactivates the specified zone set. |

## About the Default Zone

Each member of a fabric (in effect a device attached to an Nx port) can belong to any zone. If a member is not part of any active zone, it is considered to be part of the default zone. Therefore, if no zone set is active in the fabric, all devices are considered to be in the default zone. Even though a member can belong to multiple zones, a member that is part of the default zone cannot be part of any other zone. The switch determines whether a port is a member of the default zone when the attached port comes up.

**Note**

Unlike configured zones, default zone information is not distributed to the other switches in the fabric.

Traffic can either be permitted or denied among members of the default zone. This information is not distributed to all switches; it must be configured in each switch.

**Note**

When the switch is initialized for the first time, no zones are configured and all members are considered to be part of the default zone. Members are not permitted to communicate with each other.

Configure the default zone policy on each switch in the fabric. If you change the default zone policy on one switch in a fabric, be sure to change it on all the other switches in the fabric.

**Note**

The default settings for default zone configurations can be changed.

The default zone members are explicitly listed when the default policy is configured as permit or when a zone set is active. When the default policy is configured as deny, the members of this zone are not explicitly enumerated when you view the active zone set.

## Configuring the Default Zone Access Permission

To permit or deny traffic to members in the default zone, perform this task:

## SUMMARY STEPS

1. switch# **configuration terminal**
2. switch(config)# **zone default-zone permit vsan** *vsan-id*
3. switch(config)# **no zone default-zone permit vsan** *vsan-id*

## DETAILED STEPS

|               | Command or Action                                                      | Purpose                                                |
|---------------|------------------------------------------------------------------------|--------------------------------------------------------|
| <b>Step 1</b> | switch# <b>configuration terminal</b>                                  | Enters configuration mode.                             |
| <b>Step 2</b> | switch(config)# <b>zone default-zone permit vsan</b> <i>vsan-id</i>    | Permits traffic flow to default zone members.          |
| <b>Step 3</b> | switch(config)# <b>no zone default-zone permit vsan</b> <i>vsan-id</i> | Denies (default) traffic flow to default zone members. |

### About FC Alias Creation

You can assign an alias name and configure an alias member using the following values:

- pWWN—The WWN of the N port is in hex format (for example, 10:00:00:23:45:67:89:ab).
- fWWN—The WWN of the fabric port name is in hex format (for example, 10:00:00:23:45:67:89:ab).
- FC ID—The N port ID is in 0xhhhhh format (for example, 0xce00d1).
- Domain ID—The domain ID is an integer from 1 to 239. A mandatory port number of a non-Cisco switch is required to complete this membership configuration.
- Interface—Interface-based zoning is similar to port-based zoning because the switch interface is used to configure the zone. You can specify a switch interface as a zone member for both local and remote switches. To specify a remote switch, enter the remote switch WWN (sWWN) or the domain ID in the particular VSAN.



**Tip** The switch supports a maximum of 2048 aliases per VSAN.

### Creating FC Aliases

To create an alias, perform this task:

## SUMMARY STEPS

1. switch# **configuration terminal**
2. switch(config)# **falias name AliasSample vsan** *vsan-id*
3. switch(config-falias)# **member** *type value*

## DETAILED STEPS

|               | Command or Action                                            | Purpose                                                                                                                                                                                                                        |
|---------------|--------------------------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Step 1</b> | switch# <b>configuration terminal</b>                        | Enters configuration mode.                                                                                                                                                                                                     |
| <b>Step 2</b> | switch(config)# <b>fcalias name AliasSample vsan vsan-id</b> | Configures an alias name (AliasSample).                                                                                                                                                                                        |
| <b>Step 3</b> | switch(config-fcalias)# <b>member type value</b>             | Configures a member for the specified fcalias (AliasSample) based on the type (pWWN, fabric pWWN, FC ID, domain ID, or interface) and value specified.<br><br><b>Note</b> Multiple members can be specified on multiple lines. |

## Creating FC Aliases Example

Table 76: Type and Value Syntax for the member Command

|                       |                                                            |
|-----------------------|------------------------------------------------------------|
| Device alias          | <b>member device-alias device-alias</b>                    |
| Domain ID             | <b>member domain-id domain-id portnumber number</b>        |
| FC ID                 | <b>member fcid fcid</b>                                    |
| Fabric pWWN           | <b>member fwwn fwwn-id</b>                                 |
| Local sWWN interface  | <b>member interface type slot/port</b>                     |
| Domain ID interface   | <b>member interface type slot/port domain-id domain-id</b> |
| Remote sWWN interface | <b>member interface type slot/port swwn swwn-id</b>        |
| pWWN                  | <b>member pwwn pwwn-id</b>                                 |

The following example shows how to configure different types of member alias:

```
switch(config)# fcalias name AliasSample vsan 3
pWWN example:
switch(config-fcalias)# member pwwn 10:00:00:23:45:67:89:ab
fWWN example:
switch(config-fcalias)# member fwwn 10:01:10:01:10:ab:cd:ef
FC ID example:
switch(config-fcalias)# member fcid 0x222222
Domain ID example:
switch(config-fcalias)# member domain-id 2 portnumber 23
```

Local sWWN interface example:

```
switch(config-fcalias)# member interface fc 2/1
```

Remote sWWN interface example:

```
switch(config-fcalias)# member interface fc 2/1 swwn 20:00:00:05:30:00:4a:de
```

Domain ID interface example:

```
switch(config-fcalias)# member interface fc2/1 domain-id 25
```

Device alias example:

```
switch(config-fcalias)# member device-alias devName
```

## Creating Zone Sets and Adding Member Zones

To create a zone set to include several zones, perform this task:

### SUMMARY STEPS

1. switch# **configuration terminal**
2. switch(config)# **zone set name zoneset-name vsan vsan-id**
3. switch(config-zoneset)# **member name**
4. switch(config-zoneset)# **zone name zone-name**
5. switch(config-zoneset-zone)# **member fcid fcid**

### DETAILED STEPS

|               | Command or Action                                              | Purpose                                                                                                                                                                                            |
|---------------|----------------------------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Step 1</b> | switch# <b>configuration terminal</b>                          | Enters configuration mode.                                                                                                                                                                         |
| <b>Step 2</b> | switch(config)# <b>zone set name zoneset-name vsan vsan-id</b> | Configures a zone set with the configured zoneset-name.<br><b>Tip</b> To activate a zone set, you must first create the zone and a zone set.                                                       |
| <b>Step 3</b> | switch(config-zoneset)# <b>member name</b>                     | Adds a zone as a member of the previously specified zone set.<br><b>Tip</b> If the specified zone name was not previously configured, this command will return a "zone not present" error message: |
| <b>Step 4</b> | switch(config-zoneset)# <b>zone name zone-name</b>             | Adds a zone to the specified zone set.<br><b>Tip</b> Execute this step only if you need to create a zone from a zone set prompt.                                                                   |
| <b>Step 5</b> | switch(config-zoneset-zone)# <b>member fcid fcid</b>           | Adds a new member to the new zone.<br><b>Tip</b> Execute this step only if you need to add a member to a zone from a zone set prompt.                                                              |



**Tip**

You do not have to **copy** the running configuration to the startup configuration to store the active zone set. However, you need to copy the running configuration to the startup configuration to explicitly store full zone sets.

**Zone Enforcement**

Zoning can be enforced in two ways: soft and hard. Each end device (N port) discovers other devices in the fabric by querying the name server. When a device logs in to the name server, the name server returns the list of other devices that can be accessed by the querying device. If an N port does not know about the FC IDs of other devices outside its zone, it cannot access those devices.

In soft zoning, zoning restrictions are applied only during interaction between the name server and the end device. If an end device somehow knows the FC ID of a device outside its zone, it can access that device.

Hard zoning is enforced by the hardware on each frame sent by an N port. As frames enter the switch, source-destination IDs are compared with permitted combinations to allow the frame at wire speed. Hard zoning is applied to all forms of zoning.



**Note**

Hard zoning enforces zoning restrictions on every frame, and prevents unauthorized access.

Cisco Nexus 5000 Series switches support both hard and soft zoning.

**Zone Set Distribution**

You can distribute full zone sets using one of two methods: one-time distribution using the **zoneset distribute vsan** command at the EXEC mode level or full zone set distribution using the **zoneset distribute full vsan** command at the configuration mode level. The following table lists the differences between the methods.

*Table 77: Zone Set Distribution Differences*

| <b>One-Time Distribution<br/>zoneset distribute vsan Command (EXEC Mode)</b>                                                        | <b>Full Zone Set Distribution<br/>zoneset distribute full vsan Command(Configuration Mode)</b>                                             |
|-------------------------------------------------------------------------------------------------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------|
| Distributes the full zone set immediately.                                                                                          | Does not distribute the full zone set immediately.                                                                                         |
| Does not distribute the full zone set information along with the active zone set during activation, deactivation, or merge process. | Remembers to distribute the full zone set information along with the active zone set during activation, deactivation, and merge processes. |

**Enabling Full Zone Set Distribution**

All switches in the Cisco Nexus 5000 Series distribute active zone sets when new E port links come up or when a new zone set is activated in a VSAN. The zone set distribution takes effect while sending merge requests to the adjacent switch or while activating a zone set.

To enable full zone set and active zone set distribution to all switches on a per VSAN basis, perform this task:

## SUMMARY STEPS

1. switch# **configuration terminal**
2. switch(config)# **zoneset distribute full vsan** *vsan-id*

## DETAILED STEPS

|               | Command or Action                                                  | Purpose                                                        |
|---------------|--------------------------------------------------------------------|----------------------------------------------------------------|
| <b>Step 1</b> | switch# <b>configuration terminal</b>                              | Enters configuration mode.                                     |
| <b>Step 2</b> | switch(config)# <b>zoneset distribute full vsan</b> <i>vsan-id</i> | Enables sending a full zone set along with an active zone set. |

### Enabling a One-Time Distribution

You can perform a one-time distribution of inactive, unmodified zone sets throughout the fabric.

Use the **zoneset distribute vsan** *vsan-id* command in EXEC mode to perform this distribution.

```
switch# zoneset distribute vsan 2
Zoneset distribution initiated. check zone status
```

This command only distributes the full zone set information, as it does not save the information to the startup configuration. You must explicitly enter the **copy running-config startup-config** command to save the full zone set information to the startup configuration.



#### Note

The one-time distribution of the full zone set is supported in interop 2 and interop 3 modes, and not in interop 1 mode.

Use the **show zone status vsan** *vsan-id* command to check the status of the one-time zone set distribution request.

```
switch# show zone status vsan 3
VSAN: 3 default-zone: permit distribute: active only Interop: 100
mode:basic merge-control:allow

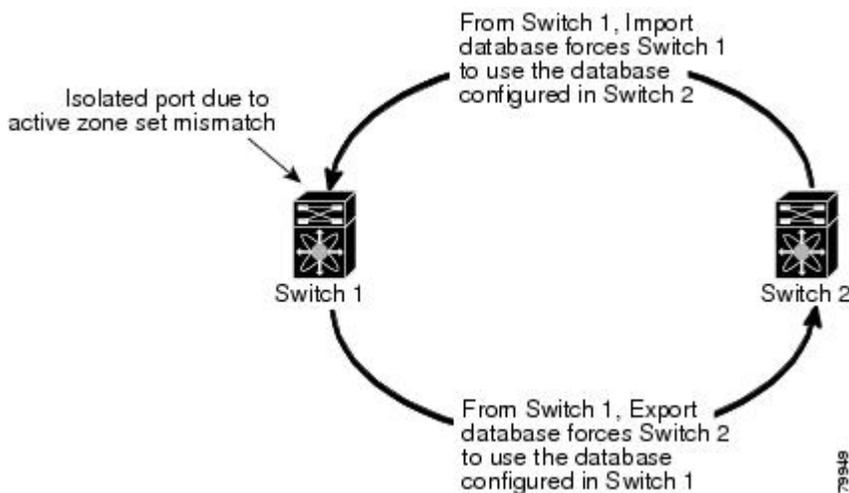
session:none
hard-zoning:enabled
Default zone:
qos:none broadcast:disabled ronly:disabled
Full Zoning Database :
Zonesets:0 Zones:0 Aliases: 0
Active Zoning Database :
Name: nozoneset Zonesets:1 Zones:2
Status: Zoneset distribution completed at 04:01:06 Aug 28 2004
```

## About Recovering from Link Isolation

When two switches in a fabric are merged using a TE or E port, these TE and E ports may become isolated when the active zone set databases are different between the two switches or fabrics. When a TE port or an E port become isolated, you can recover that port from its isolated state using one of three options:

- Import the neighboring switch's active zone set database and replace the current active zone set (see the figure below).
- Export the current database to the neighboring switch.
- Manually resolve the conflict by editing the full zone set, activating the corrected zone set, and then bringing up the link.

**Figure 66: Importing and Exporting the Database**



## Importing and Exporting Zone Sets

To import or export the zone set information from or to an adjacent switch, perform this task:

### SUMMARY STEPS

1. switch# **zoneset import interface fc slot/port vsan vsan-id**
2. switch# **zoneset import interface fc slot/port vsan vsan-id**
3. switch# **zoneset export vsan vsan-id**
4. switch# **zoneset export vsan vsan-id**

## DETAILED STEPS

|               | Command or Action                                                 | Purpose                                                                                                     |
|---------------|-------------------------------------------------------------------|-------------------------------------------------------------------------------------------------------------|
| <b>Step 1</b> | switch# <b>zoneset import interface fc slot/port vsan vsan-id</b> | Imports the zone set from the adjacent switch connected through the specified interface for the VSAN .      |
| <b>Step 2</b> | switch# <b>zoneset import interface fc slot/port vsan vsan-id</b> | Imports the zone set from the adjacent switch connected through the specified interface for the VSAN range. |
| <b>Step 3</b> | switch# <b>zoneset export vsan vsan-id</b>                        | Exports the zone set to the adjacent switch connected through the specified VSAN.                           |
| <b>Step 4</b> | switch# <b>zoneset export vsan vsan-id</b>                        | Exports the zone set to the adjacent switch connected through the specified range of VSANs.                 |

## Zone Set Duplication

You can make a copy and then edit it without altering the existing active zone set. You can copy an active zone set from the bootflash: directory, volatile: directory, or slot0 to one of the following areas:

- To the full zone set
- To a remote location (using FTP, SCP, SFTP, or TFTP)

The active zone set is not part of the full zone set. You cannot make changes to an existing zone set and activate it if the full zone set is lost or is not propagated.



### Caution

Copying an active zone set to a full zone set may overwrite a zone with the same name if it already exists in the full zone set database.

## Copying Zone Sets

On Cisco Nexus 5000 Series switches, you cannot edit an active zone set. However, you can copy an active zone set to create a new zone set that you can edit.

To make a copy of a zone set, perform this task:

## SUMMARY STEPS

1. switch# **zone copy active-zoneset full-zoneset vsan vsan-id**
2. switch# **zone copy vsan vsan-id active-zoneset scp://guest@myserver/tmp/active\_zoneset.txt**

## DETAILED STEPS

|               | Command or Action                                                                                                        | Purpose                                                                         |
|---------------|--------------------------------------------------------------------------------------------------------------------------|---------------------------------------------------------------------------------|
| <b>Step 1</b> | switch# <b>zone copy active-zoneset full-zoneset vsan</b><br><i>vsan-id</i>                                              | Makes a copy of the active zone set in the specified VSAN to the full zone set. |
| <b>Step 2</b> | switch# <b>zone copy vsan</b> <i>vsan-id</i> <b>active-zoneset</b><br><b>scp://guest@myserver/tmp/active_zoneset.txt</b> | Copies the active zone in the specified VSAN to a remote location using SCP.    |

## Renaming Zones, Zone Sets, and Aliases

To rename a zone, zone set, fcalias, or zone-attribute-group, perform this task:

## SUMMARY STEPS

1. switch# **configuration terminal**
2. switch(config)# **zoneset rename** oldname newname **vsan** *vsan-id*
3. switch(config)# **zone rename** oldname newname **vsan** *vsan-id*
4. switch(config)# **fcalias rename** oldname newname **vsan** *vsan-id*
5. switch(config)# **zone-attribute-group rename** oldname newname **vsan** *vsan-id*
6. switch(config)# **zoneset activate name** newname **vsan** *vsan-id*

## DETAILED STEPS

|               | Command or Action                                                                                | Purpose                                                                      |
|---------------|--------------------------------------------------------------------------------------------------|------------------------------------------------------------------------------|
| <b>Step 1</b> | switch# <b>configuration terminal</b>                                                            | Enters configuration mode.                                                   |
| <b>Step 2</b> | switch(config)# <b>zoneset rename</b> oldname newname <b>vsan</b><br><i>vsan-id</i>              | Renames a zone set in the specified VSAN.                                    |
| <b>Step 3</b> | switch(config)# <b>zone rename</b> oldname newname <b>vsan</b><br><i>vsan-id</i>                 | Renames a zone in the specified VSAN.                                        |
| <b>Step 4</b> | switch(config)# <b>fcalias rename</b> oldname newname <b>vsan</b><br><i>vsan-id</i>              | Renames a fcalias in the specified VSAN.                                     |
| <b>Step 5</b> | switch(config)# <b>zone-attribute-group rename</b> oldname<br>newname <b>vsan</b> <i>vsan-id</i> | Renames a zone attribute group in the specified VSAN.                        |
| <b>Step 6</b> | switch(config)# <b>zoneset activate name</b> newname <b>vsan</b><br><i>vsan-id</i>               | Activates the zone set and updates the new zone name in the active zone set. |

## Cloning Zones, Zone Sets, FC Aliases, and Zone Attribute Groups

To clone a zone, zone set, fcalias, or zone-attribute-group, perform this task:

### SUMMARY STEPS

1. switch# **configuration terminal**
2. switch(config)# **zoneset clone** *oldname newname vsan vsan-id*
3. switch(config)# **zone clone** *oldname newname vsan number*
4. switch(config)# **fcalias clone** *oldname newname vsan vsan-id*
5. switch(config)# **zone-attribute-group clone** *oldname newname vsan vsan-id*
6. switch(config)# **zoneset activate name** *newname vsan vsan-id*

### DETAILED STEPS

|               | Command or Action                                                                     | Purpose                                                                      |
|---------------|---------------------------------------------------------------------------------------|------------------------------------------------------------------------------|
| <b>Step 1</b> | switch# <b>configuration terminal</b>                                                 | Enters configuration mode.                                                   |
| <b>Step 2</b> | switch(config)# <b>zoneset clone</b> <i>oldname newname vsan vsan-id</i>              | Clones a zone set in the specified VSAN.                                     |
| <b>Step 3</b> | switch(config)# <b>zone clone</b> <i>oldname newname vsan number</i>                  | Clones a zone in the specified VSAN.                                         |
| <b>Step 4</b> | switch(config)# <b>fcalias clone</b> <i>oldname newname vsan vsan-id</i>              | Clones a fcalias in the specified VSAN.                                      |
| <b>Step 5</b> | switch(config)# <b>zone-attribute-group clone</b> <i>oldname newname vsan vsan-id</i> | Clones a zone attribute group in the specified VSAN.                         |
| <b>Step 6</b> | switch(config)# <b>zoneset activate name</b> <i>newname vsan vsan-id</i>              | Activates the zone set and updates the new zone name in the active zone set. |

## Clearing the Zone Server Database

You can clear all configured information in the zone server database for the specified VSAN.

To clear the zone server database, use the following command:

```
switch# clear zone database vsan 2
```



**Note** After entering a **clear zone database** command, you must explicitly enter the **copy running-config startup-config** to ensure that the running configuration is used when the switch reboots.



**Note** Clearing a zone set only erases the full zone database, not the active zone database.

## Verifying Zone Information

You can view any zone information by using the **show** command. If you request information for a specific object (for example, a specific zone, zone set, VSAN, or alias, or keywords such as brief or active), only information for the specified object is displayed.

The following example shows how to display zone information for all VSANs:

```
switch# show zone
```

The following example shows how to display zone information for a specific VSAN:

```
switch# show zone vsan 1
```

The following example shows how to display the configured zone sets for a range of VSANs:

```
switch# show zoneset vsan 2-3
```

The following example shows how to display the members of a specific zone:

```
switch# show zone name Zone1
```

The following example shows how to display fcalias configuration:

```
switch# show fcalias vsan 1
```

The following example shows how to display all zones to which a member belongs:

```
switch# show zone member pwnn 21:00:00:20:37:9c:48:e5
```

The following example shows how to display the number of control frames exchanged with other switches:

```
switch# show zone statistics
```

The following example shows how to display the active zone set:

```
switch# show zoneset active
```

The following example shows how to display the active zones:

```
switch# show zone active
```

The following example shows how to display the zone status:

```
switch# show zone status
```

## Enhanced Zoning

The zoning feature complies with the FC-GS-4 and FC-SW-3 standards. Both standards support the basic zoning functionalities explained in the previous section and the enhanced zoning functionalities described in this section.

### About Enhanced Zoning

The following table lists the advantages of the enhanced zoning feature in all switches in the Cisco Nexus 5000 Series.

**Table 78: Advantages of Enhanced Zoning**

| Basic Zoning                                                                                                                                  | Enhanced Zoning                                                                                                                                          | Enhanced Zoning Advantages                                                               |
|-----------------------------------------------------------------------------------------------------------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------|------------------------------------------------------------------------------------------|
| Administrators can make simultaneous configuration changes. Upon activation, one administrator can overwrite another administrator's changes. | Performs all configurations within a single configuration session. When you begin a session, the switch locks the entire fabric to implement the change. | One configuration session for the entire fabric to ensure consistency within the fabric. |

| Basic Zoning                                                                                                                                                                                                 | Enhanced Zoning                                                                                  | Enhanced Zoning Advantages                                                                            |
|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|--------------------------------------------------------------------------------------------------|-------------------------------------------------------------------------------------------------------|
| If a zone is part of multiple zone sets, you create an instance of this zone in each zone set                                                                                                                | References to the zone are used by the zone sets as required once you define the zone.           | Reduced payload size as the zone is referenced. The size is more pronounced with bigger databases.    |
| The default zone policy is defined per switch. To ensure smooth fabric operation, all switches in the fabric must have the same default zone setting.                                                        | Enforces and exchanges the default zone setting throughout the fabric.                           | Fabric-wide policy enforcement reduces troubleshooting time.                                          |
| To retrieve the results of the activation on a per switch basis, the managing switch provides a combined status about the activation. It does not identify the failure switch.                               | Retrieves the activation results and the nature of the problem from each remote switch.          | Enhanced error reporting eases the troubleshooting process                                            |
| To distribute the zoning database, you must reactivate the same zone set. The reactivation may affect hardware changes for hard zoning on the local switch and on remote switches.                           | Implements changes to the zoning database and distributes it without reactivation.               | Distribution of zone sets without activation avoids hardware changes for hard zoning in the switches. |
| The Cisco-specific zone member types (symbolic node name, and other types) may be used by other non-Cisco switches. During a merge, the Cisco-specific types can be misunderstood by the non-Cisco switches. | Provides a vendor ID along with a vendor-specific type value to uniquely identify a member type. | Unique vendor type.                                                                                   |
| The fWWN-based zone membership is only supported in Cisco interop mode.                                                                                                                                      | Supports fWWN-based membership in the standard interop mode (interop mode 1).                    | The fWWN-based member type is standardized.                                                           |

### Changing from Basic Zoning to Enhanced Zoning

To change to the enhanced zoning mode from the basic mode, perform this task:

#### SUMMARY STEPS

1. Verify that all switches in the fabric are capable of working in the enhanced mode.
2. If one or more switches are not capable of working in enhanced mode, then your request to move to enhanced mode is rejected.
3. Set the operation mode to enhanced zoning mode.

## DETAILED STEPS

- 
- Step 1** Verify that all switches in the fabric are capable of working in the enhanced mode.
- Step 2** If one or more switches are not capable of working in enhanced mode, then your request to move to enhanced mode is rejected.
- Step 3** Set the operation mode to enhanced zoning mode.
- 

### Changing from Enhanced Zoning to Basic Zoning

Cisco SAN switches allow you to change from enhanced zoning to basic zoning to enable you to downgrade and upgrade to other Cisco NX-OS releases.

To change to the basic zoning mode from the enhanced mode, perform this task:

## SUMMARY STEPS

1. Verify that the active and full zone set do not contain any configuration that is specific to the enhanced zoning mode.
2. If such configurations exist, delete them before proceeding with this procedure. If you do not delete the existing configuration, the switch software automatically removes them.
3. Set the operation mode to basic zoning mode.

## DETAILED STEPS

- 
- Step 1** Verify that the active and full zone set do not contain any configuration that is specific to the enhanced zoning mode.
- Step 2** If such configurations exist, delete them before proceeding with this procedure. If you do not delete the existing configuration, the switch software automatically removes them.
- Step 3** Set the operation mode to basic zoning mode.
- 

### Enabling Enhanced Zoning

By default, the enhanced zoning feature is disabled in all switches in the Cisco Nexus 5000 Series.

To enable enhanced zoning in a VSAN, perform this task:

## SUMMARY STEPS

1. switch# **configuration terminal**
2. switch(config)# **zone mode enhanced vsan** *vsan-id*
3. switch(config)# **no zone mode enhanced vsan** *vsan-id*

**DETAILED STEPS**

|               | <b>Command or Action</b>                                         | <b>Purpose</b>                                  |
|---------------|------------------------------------------------------------------|-------------------------------------------------|
| <b>Step 1</b> | switch# <b>configuration terminal</b>                            | Enters configuration mode.                      |
| <b>Step 2</b> | switch(config)# <b>zone mode enhanced vsan</b> <i>vsan-id</i>    | Enables enhanced zoning in the specified VSAN.  |
| <b>Step 3</b> | switch(config)# <b>no zone mode enhanced vsan</b> <i>vsan-id</i> | Disables enhanced zoning in the specified VSAN. |

**Modifying the Zone Database**

Modifications to the zone database is done within a session. A session is created at the time of the first successful configuration command. On creation of a session, a copy of the zone database is created. Any changes done within the session are performed on this copy of the zoning database. These changes in the copy zoning database are not applied to the effective zoning database until you commit the changes. Once you apply the changes, the session is closed.

If the fabric is locked by another user and for some reason the lock is not cleared, you can force the operation and close the session. You must have permission (role) to clear the lock in this switch and perform the operation on the switch from where the session was originally created.

To commit or discard changes to the zoning database in a VSAN, perform this task:

**SUMMARY STEPS**

1. switch# **configuration terminal**
2. switch(config)# **zone commit vsan** *vsan-id*
3. switch(config)# **zone commit vsan** *vsan-id* **force**
4. switch(config)# **no zone commit vsan** *vsan-id*
5. switch(config)# **no zone commit vsan** *vsan-id* **force**

**DETAILED STEPS**

|               | <b>Command or Action</b>                                            | <b>Purpose</b>                                                                                               |
|---------------|---------------------------------------------------------------------|--------------------------------------------------------------------------------------------------------------|
| <b>Step 1</b> | switch# <b>configuration terminal</b>                               | Enters configuration mode.                                                                                   |
| <b>Step 2</b> | switch(config)# <b>zone commit vsan</b> <i>vsan-id</i>              | Applies the changes to the enhanced zone database and closes the session.                                    |
| <b>Step 3</b> | switch(config)# <b>zone commit vsan</b> <i>vsan-id</i> <b>force</b> | Forcefully applies the changes to the enhanced zone database and closes the session created by another user. |
| <b>Step 4</b> | switch(config)# <b>no zone commit vsan</b> <i>vsan-id</i>           | Discards the changes to the enhanced zone database and closes the session.                                   |

|        | Command or Action                                                         | Purpose                                                                                                       |
|--------|---------------------------------------------------------------------------|---------------------------------------------------------------------------------------------------------------|
| Step 5 | switch(config)# <b>no zone commit vsan</b> <i>vsan-id</i><br><b>force</b> | Forcefully discards the changes to the enhanced zone database and closes the session created by another user. |

### Releasing Zone Database Locks

To release the session lock on the zoning database on the switches in a VSAN, use the **no zone commit vsan** command from the switch where the database was initially locked.

```
switch# configuration terminal
switch(config)# no zone commit vsan 2
```

If session locks remain on remote switches after using the **no zone commit vsan** command, you can use the **clear zone lock vsan** command on the remote switches.

```
switch# clear zone lock vsan 2
```



**Note**

We recommend using the **no zone commit vsan** command first to release the session lock in the fabric. If that fails, use the **clear zone lock vsan** command on the remote switches where the session is still locked.

### Merging the Database

The merge method depends on the fabric-wide merge control setting:

- Restrict—If the two databases are not identical, the ISLs between the switches are isolated.
- Allow—The two databases are merged using the merge rules specified in the following table.

**Table 79: Database Zone Merge Status**

| Local Database                                                                                                                                                                                                                                            | Adjacent Database | Merge Status | Results of the Merge                                            |
|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-------------------|--------------|-----------------------------------------------------------------|
| The databases contain zone sets with the same name. In the enhanced zoning mode, the active zone set does not have a name in interop mode 1. The zone set names are only present for full zone sets, but different zones, aliases, and attributes groups. |                   | Successful.  | ISLs are isolated.                                              |
| The databases contains a zone, zone alias, or zone attribute group object with same name but different members.                                                                                                                                           |                   | Failed.      | The adjacent database information populates the local database. |
| Empty.                                                                                                                                                                                                                                                    | Contains data.    | Successful.  | The union of the local and adjacent databases.                  |

| Local Database | Adjacent Database | Merge Status | Results of the Merge                                            |
|----------------|-------------------|--------------|-----------------------------------------------------------------|
| Contains data. | Empty.            | Successful.  | The local database information populates the adjacent database. |

The merge process operates as follows:

- The software compares the protocol versions. If the protocol versions differ, then the ISL is isolated.
- If the protocol versions are the same, then the zone policies are compared. If the zone policies differ, then the ISL is isolated.
- If the zone merge options are the same, then the comparison is implemented based on the merge control setting.
  - If the setting is restrict, the active zone set and the full zone set should be identical. Otherwise, the link is isolated.
  - If the setting is allow, then the merge rules are used to perform the merge.

### Configuring Zone Merge Control Policies

To configure merge control policies, perform this task:

#### SUMMARY STEPS

1. switch# **configuration terminal**
2. switch(config)# **zone merge-control restrict vsan** *vsan-id*
3. switch(config)# **no zone merge-control restrict vsan** *vsan-id*
4. switch(config)# **zone commit vsan** *vsan-id*

#### DETAILED STEPS

|               | Command or Action                                                         | Purpose                                                          |
|---------------|---------------------------------------------------------------------------|------------------------------------------------------------------|
| <b>Step 1</b> | switch# <b>configuration terminal</b>                                     | Enters configuration mode.                                       |
| <b>Step 2</b> | switch(config)# <b>zone merge-control restrict vsan</b> <i>vsan-id</i>    | Configures a restricted merge control setting for this VSAN.     |
| <b>Step 3</b> | switch(config)# <b>no zone merge-control restrict vsan</b> <i>vsan-id</i> | Defaults to using the allow merge control setting for this VSAN. |
| <b>Step 4</b> | switch(config)# <b>zone commit vsan</b> <i>vsan-id</i>                    | Commits the changes made to the specified VSAN.                  |

## Default Zone Policies

To permit or deny traffic in the default zone, perform this task:

### SUMMARY STEPS

1. switch# **configuration terminal**
2. switch(config)# **zone default-zone permit vsan vsan-id**
3. switch(config)# **no zone default-zone permit vsan vsan-id**
4. switch(config)# **zone commit vsan vsan-id**

### DETAILED STEPS

|               | Command or Action                                               | Purpose                                                                     |
|---------------|-----------------------------------------------------------------|-----------------------------------------------------------------------------|
| <b>Step 1</b> | switch# <b>configuration terminal</b>                           | Enters configuration mode.                                                  |
| <b>Step 2</b> | switch(config)# <b>zone default-zone permit vsan vsan-id</b>    | Permits traffic flow to default zone members.                               |
| <b>Step 3</b> | switch(config)# <b>no zone default-zone permit vsan vsan-id</b> | Denies traffic flow to default zone members and reverts to factory default. |
| <b>Step 4</b> | switch(config)# <b>zone commit vsan vsan-id</b>                 | Commits the changes made to the specified VSAN.                             |

## Configuring System Default Zoning Settings

You can configure default settings for default zone policies and full zone distribution for new VSANs on the switch. To configure switch-wide default settings, perform this task:

### SUMMARY STEPS

1. switch# **configuration terminal**
2. switch(config)# **system default zone default-zone permit**
3. switch(config)# **no system default zone default-zone permit**
4. switch(config)# **system default zone distribute full**
5. switch(config)# **no system default zone distribute full**

### DETAILED STEPS

|               | Command or Action                                              | Purpose                                                                     |
|---------------|----------------------------------------------------------------|-----------------------------------------------------------------------------|
| <b>Step 1</b> | switch# <b>configuration terminal</b>                          | Enters configuration mode.                                                  |
| <b>Step 2</b> | switch(config)# <b>system default zone default-zone permit</b> | Configures permit as the default zoning policy for new VSANs on the switch. |

|        | Command or Action                                                 | Purpose                                                                                                                                      |
|--------|-------------------------------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------|
| Step 3 | switch(config)# <b>no system default zone default-zone permit</b> | Configures deny (default) as the default zoning policy for new VSANs on the switch.                                                          |
| Step 4 | switch(config)# <b>system default zone distribute full</b>        | Enables full zone database distribution as the default for new VSANs on the switch.                                                          |
| Step 5 | switch(config)# <b>no system default zone distribute full</b>     | Disables (default) full zone database distribution as the default for new VSANs on the switch. Only the active zone database is distributed. |

### Verifying Enhanced Zone Information

The following example shows how to display the zone status for a specified VSAN:

```
switch# show zone status vsan 2
```

### Compacting the Zone Database

You can delete excess zones and compact the zone database for the VSAN.



#### Note

A merge failure occurs when a switch supports more than 2000 zones per VSAN but its neighbor does not. Also, zone set activation can fail if the switch has more than 2000 zones per VSAN and not all switches in the fabric support more than 2000 zones per VSAN.

To delete zones and compact the zone database for a VSAN, perform this task:

### SUMMARY STEPS

1. switch# **configuration terminal**
2. switch(config)# **no zone name zone-name vsan vsan-id**
3. switch(config)# **zone compact vsan vsan-id**

### DETAILED STEPS

|        | Command or Action                                          | Purpose                                                                                                    |
|--------|------------------------------------------------------------|------------------------------------------------------------------------------------------------------------|
| Step 1 | switch# <b>configuration terminal</b>                      | Enters configuration mode.                                                                                 |
| Step 2 | switch(config)# <b>no zone name zone-name vsan vsan-id</b> | Deletes a zone to reduce the number of zones to 2000 or fewer.                                             |
| Step 3 | switch(config)# <b>zone compact vsan vsan-id</b>           | Compacts the zone database for the specified VSAN to recover the zone ID released when a zone was deleted. |

## Zone and Zone Set Analysis

To better manage the zones and zone sets on your switch, you can display zone and zone set information using the **show zone analysis** command.

The following example shows how to display full zoning analysis:

```
switch# show zone analysis vsan 1
```

The following example shows how to display active zoning analysis:

```
switch# show zone analysis active vsan 1
```

See the Cisco Nexus 5000 Series *Switch Command Reference* for the description of the information displayed in the command output.

## Default Basic Zone Settings

The following table lists the default settings for basic zone parameters.

**Table 80: Default Basic Zone Parameters**

| Parameters               | Default                                  |
|--------------------------|------------------------------------------|
| Default zone policy      | Denied to all members.                   |
| Full zone set distribute | The full zone set(s) is not distributed. |
| Enhanced zoning          | Disabled.                                |



# Distributing Device Alias Services

This chapter contains the following sections:

- [Distributing Device Alias Services, page 631](#)

## Distributing Device Alias Services

Switches in the Cisco Nexus 5000 Series support Distributed Device Alias Services (device aliases) on a fabric-wide basis.

### Information About Device Aliases

When the port WWN (pWWN) of a device must be specified to configure features (for example, zoning, DPVM, or port security) in a Cisco Nexus 5000 Series switch, you must assign the correct device name each time you configure these features. An inaccurate device name may cause unexpected results. You can circumvent this problem if you define a user-friendly name for a pWWN and use this name in all the configuration commands as required. These user-friendly names are referred to as *device aliases*.

### Device Alias Features

Device aliases have the following features:

- The device alias information is independent of the VSAN configuration.
- The device alias configuration and distribution is independent of the zone server and the zone server database.
- You can import legacy zone alias configurations without losing data.
- The device alias application uses the Cisco Fabric Services (CFS) infrastructure to enable efficient database management and distribution. Device aliases use the coordinated distribution mode and the fabric-wide distribution scope.
- Basic and enhanced modes.
- Device aliases used to configure zones, IVR zones, or port security features are displayed automatically with their respective pWWNs in the **show** command output.

## Related Topics

[Device Alias Modes, on page 634](#)

[Using Cisco Fabric Services, on page 329](#)

## Device Alias Requirements

Device aliases have the following requirements:

- You can only assign device aliases to pWWNs.
- There must be a one-to-one relationship between the pWWN and the device alias that maps to it.
- A device alias name is restricted to 64 alphanumeric characters and may include one or more of the following characters:
  - a to z and A to Z
  - Device alias names must begin with an alphabetic character (a to z or A to Z).
  - 1 to 9
  - - (hyphen) and \_ (underscore)
  - \$ (dollar sign) and ^ (up caret)

## Zone Aliases Versus Device Aliases

The following table compares the configuration differences between zone-based alias configuration and device alias configuration.

**Table 81: Comparison Between Zone Aliases and Device Aliases**

| Zone-Based Aliases                                                                                               | Device Aliases                                                                                                                                        |
|------------------------------------------------------------------------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------|
| Aliases are limited to the specified VSAN.                                                                       | You can define device aliases without specifying the VSAN number. You can also use the same definition in one or more VSANs without any restrictions. |
| Zone aliases are part of the zoning configuration. The alias mapping cannot be used to configure other features. | Device aliases can be used with any feature that uses the pWWN.                                                                                       |
| You can use any zone member type to specify the end devices.                                                     | Only pWWNs are supported.                                                                                                                             |
| Configuration is contained within the zone server database and is not available to other features.               | Device aliases are not restricted to zoning. Device alias configuration is available to the FCNS, zone, fcping, and traceroute applications.          |

## Device Alias Databases

The device alias feature uses two databases to accept and implement device alias configurations.

- Effective database—The database currently used by the fabric.
- Pending database—Your subsequent device alias configuration changes are stored in the pending database.

If you modify the device alias configuration, you need to commit or discard the changes as the fabric remains locked during this period.

Device alias database changes are validated with the applications. If any of the applications cannot accept the device alias database changes, then those changes are rejected; this applies to device alias database changes resulting from either a commit or merge operation.

### Creating Device Aliases

To create a device alias in the pending database, perform this task:

#### SUMMARY STEPS

1. switch# **configuration terminal**
2. switch(config)# **device-alias database**
3. switch(config-device-alias-db)# **device-alias name** *device-name* **pwwn** *pwwn-id*
4. switch(config-device-alias-db)# **no device-alias name** *device-name*
5. switch(config-device-alias-db)# **device-alias rename** *old-device-name* *new-device-name*

#### DETAILED STEPS

|               | Command or Action                                                                                        | Purpose                                                                                                                                                                                                           |
|---------------|----------------------------------------------------------------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Step 1</b> | switch# <b>configuration terminal</b>                                                                    | Enters configuration mode.                                                                                                                                                                                        |
| <b>Step 2</b> | switch(config)# <b>device-alias database</b>                                                             | Enters the pending database configuration submode.                                                                                                                                                                |
| <b>Step 3</b> | switch(config-device-alias-db)# <b>device-alias name</b> <i>device-name</i> <b>pwwn</b> <i>pwwn-id</i>   | Specifies a device name for the device that is identified by its pWWN. Starts writing to the pending database and simultaneously locks the fabric as this is the first-issued device alias configuration command. |
| <b>Step 4</b> | switch(config-device-alias-db)# <b>no device-alias name</b> <i>device-name</i>                           | Removes the device name for the device that is identified by its pWWN.                                                                                                                                            |
| <b>Step 5</b> | switch(config-device-alias-db)# <b>device-alias rename</b> <i>old-device-name</i> <i>new-device-name</i> | Renames an existing device alias with a new name.                                                                                                                                                                 |

### Example of Creating a Device Alias

To display the device alias configuration, use the **show device-alias name** command:

```
switch# show device-alias name x
device-alias name x pwnn 21:01:00:e0:8b:2e:80:93
```

### Device Alias Modes

You can specify that aliases operate in basic or enhanced modes.

When operating in basic mode, which is the default mode, the device alias is immediately expanded to a pWWN. In basic mode, when device aliases are changed to point to a new HBA, for example, that change is not reflected in the zone server. Users must remove the previous HBA's pWWN, add the new HBA's pWWN, and then reactivate the zoneset.

When operating in enhanced mode, applications accept a device alias name in its "native" format. Instead of expanding the device alias to a pWWN, the device alias name is stored in the configuration and distributed in its native device alias format. So applications such as zone server, PSM or DPVM can automatically keep track of the device alias membership changes and enforce them accordingly. The primary benefit of operating in enhanced mode is that you have a single point of change.

Whenever you change device alias modes, the change is distributed to other switches in the network only if device alias distribution is enabled or on. Otherwise, the mode change only takes place on the local switch.



#### Note

---

Enhanced mode, or native device alias-based configurations are not accepted in interop mode VSANs. IVR zoneset activation will fail in interop mode VSANs if the corresponding zones have native device alias-based members.

---

### Changing Device Alias Mode Guidelines

When changing device alias modes, follow these guidelines:

- If two fabrics running in different device alias modes are joined together, the device alias merge will fail. There is no automatic conversion to one mode or the other during the merge process. In this situation, you must to select one mode over the other.
- Before changing from enhanced to basic mode, you must first explicitly remove all native device alias-based configurations from both local and remote switches, or, replace all device alias-based configuration members with the corresponding pWWN.
- If you remove a device alias from the device alias database, all applications will automatically stop enforcing the corresponding device alias. If that corresponding device alias is part of an active zoneset, all the traffic to and from that pWWN is disrupted.
- Renaming the device alias not only changes the device alias name in the device alias database, but also replaces the corresponding device alias configuration in all the applications.
- When a new device alias is added to the device alias database, and the application configuration is present on that device alias, it automatically takes effect. For example, if the corresponding device alias is part of the active zoneset and the device is online, then zoning is enforced automatically. You do not have to reactivate the zoneset.

- If a device alias name is mapped to a new HBA's pWWN, then the application's enforcement changes accordingly. In this case, the zone server automatically enforces zoning based on the new HBA's pWWN.

## Configuring Device Alias Modes

To configure device aliases to operate in enhanced mode, perform this task:

### SUMMARY STEPS

1. switch# **configuration terminal**
2. switch(config)# **device-alias mode enhanced**
3. switch(config)# **no device-alias mode enhance**

### DETAILED STEPS

|               | Command or Action                                   | Purpose                                               |
|---------------|-----------------------------------------------------|-------------------------------------------------------|
| <b>Step 1</b> | switch# <b>configuration terminal</b>               | Enters configuration mode.                            |
| <b>Step 2</b> | switch(config)# <b>device-alias mode enhanced</b>   | Assigns the device alias to operate in enhanced mode. |
| <b>Step 3</b> | switch(config)# <b>no device-alias mode enhance</b> | Assigns the device alias to operate in basic mode.    |

### Viewing the Device Alias Mode Setting

To view the current device alias mode setting, enter the **show device-alias status** command.

```
switch# show device-alias status
Fabric Distribution: Enabled
Database:- Device Aliases 0 Mode: Basic
Locked By:- User "admin" SWWN 20:00:00:0d:ec:30:90:40
Pending Database:- Device Aliases 0 Mode: Basic
```

### About Device Alias Distribution

By default, device alias distribution is enabled. The device alias feature uses CFS to distribute the modifications to all switches in a fabric.

If device alias distribution is disabled, database changes are not distributed to the switches in the fabric. The same changes would have to be performed manually on all switches in the fabric to keep the device alias database up-to-date. Database changes immediately take effect, so there would not be any pending database and commit or abort operations either. If you have not committed the changes and you disable distribution, then a commit task will fail.

The following example displays a failed device alias status:

```
switch# show device-alias status
Fabric Distribution: Disabled
Database:- Device Aliases 25
Status of the last CFS operation issued from this switch:
=====
Operation: Commit
Status: Failed (Reason: Operation is not permitted as the fabric distribution is
currently disabled.)
```

## Locking the Fabric

When you perform any device alias configuration task (regardless of which device alias task), the fabric is automatically locked for the device alias feature. Once you lock the fabric, the following situations apply:

- No other user can make any configuration changes to this feature.
- A copy of the effective database is obtained and used as the pending database. Subsequent modifications are made to the pending database. The pending database remains in use until you commit the modifications to the pending database or discard (**abort**) the changes to the pending database.

## Committing Changes

If you commit the changes made to the pending database, the following events occur:

- The pending database content overwrites the effective database content.
- The pending database is distributed to the switches in the fabric and the effective database on those switches is overwritten with the new changes.
- The pending database is emptied of its contents.
- The fabric lock is released for this feature.

To commit the changes, perform this task:

## SUMMARY STEPS

1. switch# **configuration terminal**
2. switch(config)# **device-alias commit**

## DETAILED STEPS

|               | Command or Action                          | Purpose                                                   |
|---------------|--------------------------------------------|-----------------------------------------------------------|
| <b>Step 1</b> | switch# <b>configuration terminal</b>      | Enters configuration mode.                                |
| <b>Step 2</b> | switch(config)# <b>device-alias commit</b> | Commits the changes made to the currently active session. |

## Discarding Changes

If you discard the changes made to the pending database, the following events occur:

- The effective database contents remain unaffected.
- The pending database is emptied of its contents.
- The fabric lock is released for this feature.

To discard the device alias session, perform this task:

## SUMMARY STEPS

1. switch# **configuration terminal**
2. switch(config)# **device-alias abort**

## DETAILED STEPS

|        | Command or Action                         | Purpose                                |
|--------|-------------------------------------------|----------------------------------------|
| Step 1 | switch# <b>configuration terminal</b>     | Enters configuration mode.             |
| Step 2 | switch(config)# <b>device-alias abort</b> | Discards the currently active session. |

### Displaying the Status of a Discard Operation

To display the status of the discard operation, use the show **device alias status** command.

```
switch# show device-alias status
Fabric Distribution: Enabled
Database:- Device Aliases 24
Status of the last CFS operation issued from this switch:
=====
Operation: Abort
Status: Success
```

### Fabric Lock Override

You can use locking operations (clear, commit, abort) only when device alias distribution is enabled. If you have performed a device alias task and have forgotten to release the lock by either committing or discarding the changes, an administrator can release the lock from any switch in the fabric. If the administrator performs this task, your changes to the pending database are discarded and the fabric lock is released.

The changes are only available in the volatile directory and may be discarded if the switch is restarted.

To use administrative privileges and release a locked device alias session, use the **clear device-alias session** command in EXEC mode.

```
switch# clear device-alias session
To display the status of the clear operation, use the show device-alias status command.

switch# show device-alias status
Fabric Distribution: Enabled
Database:- Device Aliases 24
Status of the last CFS operation issued from this switch:
=====
Operation: Clear Session<-----Lock released by administrator
Status: Success<-----Successful status of the operation
```

### Disabling and Enabling Device Alias Distribution

To disable or enable the device alias distribution, perform this task:

## SUMMARY STEPS

1. switch# **configuration terminal**
2. switch(config)# **no device-alias distribute**
3. switch(config)# **device-alias distribute**

## DETAILED STEPS

|               | Command or Action                                 | Purpose                             |
|---------------|---------------------------------------------------|-------------------------------------|
| <b>Step 1</b> | switch# <b>configuration terminal</b>             | Enters configuration mode.          |
| <b>Step 2</b> | switch(config)# <b>no device-alias distribute</b> | Disables the distribution.          |
| <b>Step 3</b> | switch(config)# <b>device-alias distribute</b>    | Enables the distribution (default). |

### Viewing the Status of Device Alias Distribution

To display the status of device alias distribution, use the **show device-alias status** command. The following example shows the device alias display when distribution is enabled:

```
switch# show device-alias status
Fabric Distribution: Enabled <-----Distribution is enabled

Database:-Device Aliases 24

Locked By:-User "Test" SWWN 20:00:00:0c:cf:f4:02:83<-Lock holder's user name and switch ID

Pending Database:- Device Aliases 24

Status of the last CFS operation issued from this switch:
=====

Operation: Enable Fabric Distribution

Status: Success
```

The following example shows the device alias display when distribution is disabled:

```
switch# show device-alias status
Fabric Distribution: Disabled

Database:- Device Aliases 24

Status of the last CFS operation issued from this switch:
=====

Operation: Disable Fabric Distribution

Status: Success
```

## About Legacy Zone Alias Configuration

You can import legacy zone alias configurations to use this feature without losing data if they satisfy the following restrictions:

- Each zone alias has only one member.
- The member type is pWWN.

If any name or definition conflict exists, the zone aliases are not imported.

Ensure that you copy any required zone aliases to the device alias database as required by your configuration.

When an import operation is complete, the modified alias database is distributed to all other switches in the physical fabric when you perform the **commit** operation. If you do not want to distribute the configuration to other switches in the fabric, you can perform the **abort** operation and the merge changes are completely discarded.

### Importing a Zone Alias

To import the zone alias for a specific VSAN, perform this task:

#### SUMMARY STEPS

1. switch# **configuration terminal**
2. switch(config)# **device-alias import fcalias vsan** *vlan-id*

#### DETAILED STEPS

|               | Command or Action                                                      | Purpose                                                 |
|---------------|------------------------------------------------------------------------|---------------------------------------------------------|
| <b>Step 1</b> | switch# <b>configuration terminal</b>                                  | Enters configuration mode.                              |
| <b>Step 2</b> | switch(config)# <b>device-alias import fcalias vsan</b> <i>vlan-id</i> | Imports the fcalias information for the specified VSAN. |

## Device Alias Database Merge Guidelines

When merging two device alias databases, follow these guidelines:

- Verify that two device aliases with different names are not mapped to the same pWWN.
- Verify that two identical pWWNs are not mapped to two different device aliases.
- Verify that the combined number of device aliases in both databases does not exceed 8K (8191 device aliases) in fabrics running Cisco MDS SAN-OS Release 3.0 (x) and earlier, and 20K in fabrics running Cisco MDS SAN-OS Release 3.1(x) and later.

If the combined number of device entries in both databases exceeds the supported configuration limit, then the merge will fail. For example, if database *N* has 6000 device aliases and database *M* has 2192 device aliases, and you are running SAN-OS Release 3.0(x) or earlier, then this merge operation will fail. Merge operations will also fail if there is a device alias mode mismatch.

#### Related Topics

[CFS Merge Support](#), on page 334

## Verifying Device Alias Configuration

To display device alias information, perform one of the following tasks:

### SUMMARY STEPS

1. switch# **show zoneset** [active]
2. switch# **show device-alias database** [pending | pending-diffs]
3. switch# **show device-alias** {pwwn *pwwn-id* | name *device-name* } [pending]
4. switch# **show flogi database** [pending]
5. switch# **show fcns database** [pending]

### DETAILED STEPS

|               | Command or Action                                                                           | Purpose                                                                |
|---------------|---------------------------------------------------------------------------------------------|------------------------------------------------------------------------|
| <b>Step 1</b> | switch# <b>show zoneset</b> [active]                                                        | Displays the device aliases in the zone set information.               |
| <b>Step 2</b> | switch# <b>show device-alias database</b> [pending   pending-diffs]                         | Displays the device alias database.                                    |
| <b>Step 3</b> | switch# <b>show device-alias</b> {pwwn <i>pwwn-id</i>   name <i>device-name</i> } [pending] | Displays the device alias information for the specified pwwn or alias. |
| <b>Step 4</b> | switch# <b>show flogi database</b> [pending]                                                | Displays device alias information the the flogi database.              |
| <b>Step 5</b> | switch# <b>show fcns database</b> [pending]                                                 | Displays device alias information the the fcns database.               |

### Examples of Verifying Device Alias Configuration

The following example shows how to display device alias information in the zone set:

```
switch# show zoneset
zoneset name s1 vsan 1
 zone name z1 vsan 1
 pwwn 21:01:00:e0:8b:2e:80:93 [x] <-----Device alias displayed for each pWWN.
 pwwn 21:00:00:20:37:39:ab:5f [y]
 zone name z2 vsan 1
 pwwn 21:00:00:e0:8b:0b:66:56 [SampleName]
 pwwn 21:00:00:20:37:39:ac:0d [z]
```

The following example shows how to display pending changes in the device alias database:

```
switch# show device-alias database pending
```

The following example shows how to display a specific pWWN in the device alias database:

```
switch# show device-alias pwwn 21:01:00:e0:8b:2e:80:93 pending
```

The following example shows how to display the difference between the pending and effective device alias databases:

```
switch# show device-alias database pending-diff
- device-alias name Doc pwwn 21:01:02:03:00:01:01:01
+ device-alias name SampleName pwwn 21:00:00:e0:8b:0b:66:56
```

Where available, device aliases are displayed regardless of a member being configured using a **device-alias** command or a zone-specific **member pwwn** command.

## Default Device Alias Settings

The following table lists the default settings for device alias parameters.

**Table 82: Default Device Alias Parameters**

| <b>Parameters</b>              | <b>Default</b>                           |
|--------------------------------|------------------------------------------|
| Device alias distribution      | Enabled.                                 |
| Device alias mode              | Basic.                                   |
| Database in use                | Effective database.                      |
| Database to accept changes     | Pending database.                        |
| Device alias fabric lock state | Locked with the first device alias task. |





# Configuring Fibre Channel Routing Services and Protocols

---

This chapter contains the following sections:

- [Configuring Fibre Channel Routing Services and Protocols](#), page 643

## Configuring Fibre Channel Routing Services and Protocols

Fabric Shortest Path First (FSPF) is the standard path selection protocol used by Fibre Channel fabrics. The FSPF feature is enabled by default on the E mode and TE mode Fibre Channel interfaces on Cisco Nexus 5000 Series switches. Except in configurations that require special consideration, you do not need to configure any FSPF services. FSPF automatically calculates the best path between any two switches in a fabric. FSPF provides the following capabilities:

- Dynamically computes routes throughout a fabric by establishing the shortest and quickest path between any two switches.
- Selects an alternative path in the event of the failure of a given path. FSPF supports multiple paths and automatically computes an alternative path around a failed link. It provides a preferred route when two equal paths are available.

### Information About FSPF

FSPF is the protocol currently standardized by the T11 committee for routing in Fibre Channel networks. The FSPF protocol has the following characteristics and features:

- Supports multipath routing.
- Bases path status on a link state protocol.
- Routes hop by hop, based only on the domain ID.
- Runs only on E ports or TE ports and provides a loop free topology.
- Runs on a per VSAN basis. Connectivity in a given VSAN in a fabric is guaranteed only for the switches configured in that VSAN.

- Uses a topology database to keep track of the state of the links on all switches in the fabric and associates a cost with each link.
- Guarantees a fast reconvergence time in case of a topology change. Uses the standard Dijkstra algorithm, but there is a static dynamic option for a more robust, efficient, and incremental Dijkstra algorithm. The reconvergence time is fast and efficient as the route computation is done on a per VSAN basis.

**Note**

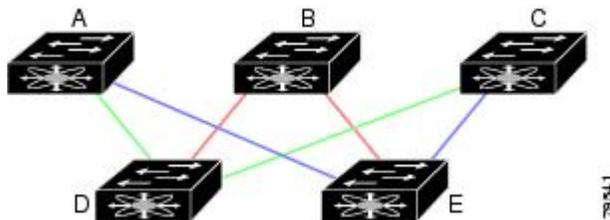
The FSPF feature can be used on any topology.

## FSPF Examples

### *Fault Tolerant Fabric Example*

The following figure depicts a fault tolerant fabric using a partial mesh topology. If a link goes down anywhere in the fabric, any switch can still communicate with all others in the fabric. In the same way, if any switch goes down, the connectivity of the rest of the fabric is preserved.

**Figure 67: Fault Tolerant Fabric**



For example, if all links are of equal speed, the FSPF calculates two equal paths from A to C: A-D-C (green) and A-E-C (blue).

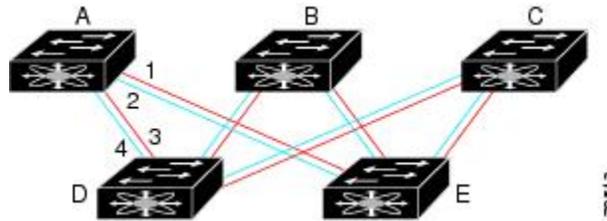
### *Redundant Link Example*

To improve on the topology, each connection between any pair of switches can be replicated; two or more links can be present between a pair of switches. The following figure shows this arrangement. Because switches in the Cisco Nexus 5000 Series support SAN port channels, each pair of physical links can appear to the FSPF protocol as one single logical link.

By bundling pairs of physical links, FSPF efficiency is considerably improved by the reduced database size and the frequency of link updates. Once physical links are aggregated, failures are not attached to a single link but to the entire SAN port channel. This configuration also improves the resiliency of the network. The

failure of a link in a SAN port channel does not trigger a route change, which reduces the risks of routing loops, traffic loss, or fabric downtime for route reconfiguration.

**Figure 68: Fault Tolerant Fabric with Redundant Links**



For example, if all links are of equal speed and no SAN port channels exist, the FSPF calculates four equal paths from A to C: A1-E-C, A2-E-C, A3-D-C, and A4-D-C. If SAN port channels exist, these paths are reduced to two.

## FSPF Global Configuration

By default, FSPF is enabled on switches in the Cisco Nexus 5000 Series .

Some FSPF features can be globally configured in each VSAN. By configuring a feature for the entire VSAN, you do not have to specify the VSAN number for every command. This global configuration feature also reduces the chance of typing errors or other minor configuration errors.



### Note

FSPF is enabled by default. Generally, you do not need to configure these advanced features.



### Caution

The default for the backbone region is 0 (zero). You do not need to change this setting unless your region is different from the default. If you are operating with other vendors using the backbone region, you can change this default to be compatible with those settings.

## About SPF Computational Hold Times

The SPF computational hold time sets the minimum time between two consecutive SPF computations on the VSAN. Setting this to a small value means that FSPF reacts faster to any fabric changes by recomputing paths on the VSAN. A small SPF computational hold time uses more switch CPU time.

## About Link State Records

Each time a new switch enters the fabric, a link state record (LSR) is sent to the neighboring switches, and then flooded throughout the fabric.

The following table displays the default settings for switch responses.

**Table 83: LSR Default Settings**

| LSR Option                             | Default    | Description                                                                       |
|----------------------------------------|------------|-----------------------------------------------------------------------------------|
| Acknowledgment interval (RxmtInterval) | 5 seconds  | The time a switch waits for an acknowledgment from the LSR before retransmission. |
| Refresh time (LSRefreshTime)           | 30 minutes | The time a switch waits before sending an LSR refresh transmission.               |
| Maximum age (MaxAge)                   | 60 minutes | The time a switch waits before dropping the LSR from the database.                |

The LSR minimum arrival time is the period between receiving LSR updates on this VSAN. Any LSR updates that arrive before the LSR minimum arrival time are discarded.

The LSR minimum interval time is the frequency at which this switch sends LSR updates on a VSAN.

### Configuring FSPF on a VSAN

To configure an FSPF feature for the entire VSAN, perform this task:

#### SUMMARY STEPS

1. switch# **configuration terminal**
2. switch(config)# **fspf config vsan** *vsan-id*
3. switch-config-(fspf-config)# **spf static**
4. switch-config-(fspf-config)# **spf hold-time** *value*
5. switch-config-(fspf-config)# **region** *region-id*

#### DETAILED STEPS

|               | Command or Action                                              | Purpose                                                                                                                                                                                                                                                           |
|---------------|----------------------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Step 1</b> | switch# <b>configuration terminal</b>                          | Enters configuration mode.                                                                                                                                                                                                                                        |
| <b>Step 2</b> | switch(config)# <b>fspf config vsan</b> <i>vsan-id</i>         | Enters FSPF global configuration mode for the specified VSAN.                                                                                                                                                                                                     |
| <b>Step 3</b> | switch-config-(fspf-config)# <b>spf static</b>                 | Forces static SPF computation for the dynamic (default) incremental VSAN.                                                                                                                                                                                         |
| <b>Step 4</b> | switch-config-(fspf-config)# <b>spf hold-time</b> <i>value</i> | Configures the hold time between two route computations in milliseconds (msec) for the entire VSAN. The default value is 0.<br><br><b>Note</b> If the specified time is shorter, the routing is faster. However, the processor consumption increases accordingly. |

|               | Command or Action                                           | Purpose                                                                     |
|---------------|-------------------------------------------------------------|-----------------------------------------------------------------------------|
| <b>Step 5</b> | switch-config-(fspf-config)# <b>region</b> <i>region-id</i> | Configures the autonomous region for this VSAN and specifies the region ID. |

### Resetting FSPF to the Default Configuration

To return the FSPF VSAN global configuration to its factory default, perform this task:

#### SUMMARY STEPS

1. switch# **configuration terminal**
2. switch(config)# **no fspf config vsan** *vsan-id*

#### DETAILED STEPS

|               | Command or Action                                         | Purpose                                                |
|---------------|-----------------------------------------------------------|--------------------------------------------------------|
| <b>Step 1</b> | switch# <b>configuration terminal</b>                     | Enters configuration mode.                             |
| <b>Step 2</b> | switch(config)# <b>no fspf config vsan</b> <i>vsan-id</i> | Deletes the FSPF configuration for the specified VSAN. |

### Enabling or Disabling FSPF

To enable or disable FSPF routing protocols, perform this task:

#### SUMMARY STEPS

1. switch# **configuration terminal**
2. switch(config)# **fspf enable vsan** *vsan-id*
3. switch(config)# **no fspf enable vsan** *vsan-id*

#### DETAILED STEPS

|               | Command or Action                                         | Purpose                                                   |
|---------------|-----------------------------------------------------------|-----------------------------------------------------------|
| <b>Step 1</b> | switch# <b>configuration terminal</b>                     | Enters configuration mode.                                |
| <b>Step 2</b> | switch(config)# <b>fspf enable vsan</b> <i>vsan-id</i>    | Enables the FSPF routing protocol in the specified VSAN.  |
| <b>Step 3</b> | switch(config)# <b>no fspf enable vsan</b> <i>vsan-id</i> | Disables the FSPF routing protocol in the specified VSAN. |

## Clearing FSPF Counters for the VSAN

To clear the FSPF statistics counters for the entire VSAN, perform this task:

### SUMMARY STEPS

1. `switch# clear fspf counters vsan vsan-id`

### DETAILED STEPS

|        | Command or Action                                     | Purpose                                                                                                                           |
|--------|-------------------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------|
| Step 1 | <code>switch# clear fspf counters vsan vsan-id</code> | Clears the FSPF statistics counters for the specified VSAN. If an interface reference is not specified, all counters are cleared. |

## FSPF Interface Configuration

Several FSPF commands are available on a per-interface basis. These configuration procedures apply to an interface in a specific VSAN.

### About FSPF Link Cost

FSPF tracks the state of links on all switches in the fabric, associates a cost with each link in its database, and then chooses the path with a minimal cost. The cost associated with an interface can be administratively changed to implement the FSPF route selection. The integer value to specify cost can range from 1 to 65,535. The default cost for 1 Gbps is 1000 and for 2 Gbps is 500.

### Configuring FSPF Link Cost

To configure FSPF link cost, perform this task:

### SUMMARY STEPS

1. `switch# configuration terminal`
2. `switch(config)# interface fc slot/port`
3. `switch(config-if)# fspf cost value vsan vsan-id`

**DETAILED STEPS**

|               | <b>Command or Action</b>                               | <b>Purpose</b>                                                                                                       |
|---------------|--------------------------------------------------------|----------------------------------------------------------------------------------------------------------------------|
| <b>Step 1</b> | switch# <b>configuration terminal</b>                  | Enters configuration mode.                                                                                           |
| <b>Step 2</b> | switch(config)# <b>interface fc slot/port</b>          | Configures the specified interface, or if already configured, enters configuration mode for the specified interface. |
| <b>Step 3</b> | switch(config-if)# <b>fspf cost value vsan vsan-id</b> | Configures the cost for the selected interface in the specified VSAN.                                                |

**About Hello Time Intervals**

You can set the FSPF Hello time interval to specify the interval between the periodic hello messages sent to verify the health of the link. The integer value can range from 1 to 65,535 seconds.



**Note** This value must be the same in the ports at both ends of the ISL.

**Configuring Hello Time Intervals**

To configure the FSPF Hello time interval, perform this task:

**SUMMARY STEPS**

1. switch# **configuration terminal**
2. switch(config)# **interface fc slot/port**
3. switch(config-if)# **fspf hello-interval value vsan vsan-id**

**DETAILED STEPS**

|               | <b>Command or Action</b>                                         | <b>Purpose</b>                                                                                                       |
|---------------|------------------------------------------------------------------|----------------------------------------------------------------------------------------------------------------------|
| <b>Step 1</b> | switch# <b>configuration terminal</b>                            | Enters configuration mode.                                                                                           |
| <b>Step 2</b> | switch(config)# <b>interface fc slot/port</b>                    | Configures the specified interface, or if already configured, enters configuration mode for the specified interface. |
| <b>Step 3</b> | switch(config-if)# <b>fspf hello-interval value vsan vsan-id</b> | Specifies the hello message interval to verify the health of the link in VSAN 175. The default is 20 seconds.        |

## About Dead Time Intervals

You can set the FSPF dead time interval to specify the maximum interval for which a hello message must be received before the neighbor is considered lost and removed from the database. The integer value can range from 1 to 65,535 seconds.



### Note

This value must be the same in the ports at both ends of the ISL.



### Caution

An error is reported at the command prompt if the configured dead time interval is less than the hello time interval.

## Configuring Dead Time Intervals

To configure the FSPF dead time interval, perform this task:

### SUMMARY STEPS

1. switch# **configuration terminal**
2. switch(config)# **interface fc slot/port**
3. switch(config-if)# **fspf dead-interval value vsan vsan-id**

### DETAILED STEPS

|               | Command or Action                                               | Purpose                                                                                                                                                                                          |
|---------------|-----------------------------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Step 1</b> | switch# <b>configuration terminal</b>                           | Enters configuration mode.                                                                                                                                                                       |
| <b>Step 2</b> | switch(config)# <b>interface fc slot/port</b>                   | Configures the specified interface, or if already configured, enters configuration mode for the specified interface.                                                                             |
| <b>Step 3</b> | switch(config-if)# <b>fspf dead-interval value vsan vsan-id</b> | Specifies the maximum interval for the specified VSAN before which a hello message must be received on the selected interface before the neighbor is considered lost. The default is 80 seconds. |

## About Retransmitting Intervals

You can specify the time after which an unacknowledged link state update should be transmitted on the interface. The integer value to specify retransmit intervals can range from 1 to 65,535 seconds.



### Note

This value must be the same on the switches on both ends of the interface.

## Configuring Retransmitting Intervals

To configure the FSPF retransmit time interval, perform this task:

### SUMMARY STEPS

1. switch# **configuration terminal**
2. switch(config)# **interface fc slot/port**
3. switch(config-if)# **fspf retransmit-interval value vsan vsan-id**

### DETAILED STEPS

|               | Command or Action                                                     | Purpose                                                                                                                       |
|---------------|-----------------------------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------|
| <b>Step 1</b> | switch# <b>configuration terminal</b>                                 | Enters configuration mode.                                                                                                    |
| <b>Step 2</b> | switch(config)# <b>interface fc slot/port</b>                         | Configures the specified interface, or if already configured, enters configuration mode for the specified interface.          |
| <b>Step 3</b> | switch(config-if)# <b>fspf retransmit-interval value vsan vsan-id</b> | Specifies the retransmit time interval for unacknowledged link state updates in the specified VSAN. The default is 5 seconds. |

### About Disabling FSPF for Specific Interfaces

You can disable the FSPF protocol for selected interfaces. By default, FSPF is enabled on all E ports and TE ports. This default can be disabled by setting the interface as passive.



#### Note

FSPF must be enabled at both ends of the interface for the protocol to work.

### Disabling FSPF for Specific Interfaces

You can disable the FSPF protocol for selected interfaces. By default, FSPF is enabled on all E ports and TE ports. This default can be disabled by setting the interface as passive.

To disable FSPF for a specific interface, perform this task:

### SUMMARY STEPS

1. switch# **configuration terminal**
2. switch(config)# **interface fc slot/port**
3. switch(config-if)# **fspf passive vsan vsan-id**
4. switch(config-if)# **no fspf passive vsan vsan-id**

**DETAILED STEPS**

|               | <b>Command or Action</b>                               | <b>Purpose</b>                                                                                                     |
|---------------|--------------------------------------------------------|--------------------------------------------------------------------------------------------------------------------|
| <b>Step 1</b> | switch# <b>configuration terminal</b>                  | Enters configuration mode.                                                                                         |
| <b>Step 2</b> | switch(config)# <b>interface fc slot/port</b>          | Configures a specified interface, or if already configured, enters configuration mode for the specified interface. |
| <b>Step 3</b> | switch(config-if)# <b>fspf passive vsan vsan-id</b>    | Disables the FSPF protocol for the specified interface in the specified VSAN.                                      |
| <b>Step 4</b> | switch(config-if)# <b>no fspf passive vsan vsan-id</b> | Reenables the FSPF protocol for the specified interface in the specified VSAN.                                     |

**Clearing FSPF Counters for an Interface**

To clear the FSPF statistics counters for an interface, perform this task:

**SUMMARY STEPS**

1. switch# **clear fspf counters vsan vsan-id interface fc slot/port**

**DETAILED STEPS**

|               | <b>Command or Action</b>                                               | <b>Purpose</b>                                                                         |
|---------------|------------------------------------------------------------------------|----------------------------------------------------------------------------------------|
| <b>Step 1</b> | switch# <b>clear fspf counters vsan vsan-id interface fc slot/port</b> | Clears the FSPF statistics counters for the specified interface in the specified VSAN. |

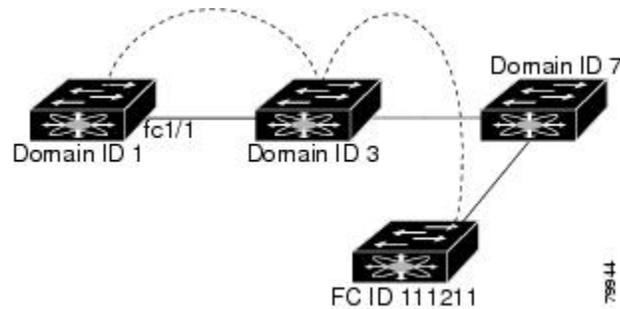
**FSPF Routes**

FSPF routes traffic across the fabric, based on entries in the FSPF database. These routes can be learned dynamically, or configured statically.

## About Fibre Channel Routes

Each port implements forwarding logic, which forwards frames based on its FC ID. Using the FC ID for the specified interface and domain, you can configure the specified route (for example, FC ID 111211 and domain ID 3) in the switch with domain ID 1 (see the following figure).

**Figure 69: Fibre Channel Routes**



## Configuring Fibre Channel Routes

To configure a Fibre Channel route, perform this task:

### SUMMARY STEPS

1. switch# **configuration terminal**
2. switch(config)#**fcroute** *fcid* **interface** *fc slot/port* **domain** *domain-id* **vsan** *vsan-id*
3. switch(config)#**fcroute** *fcid* **interface** *san-port-channel port* **domain** *domain-id* **vsan** *vsan-id*
4. switch(config)# **fcroute** *fcid* **interface** *fc slot/port* **domain** *domain-id* **metric** *value* **vsan** *vsan-id*
5. switch(config)#**fcroute** *fcid* **interface** *fc slot/port* **domain** *domain-id* **metric** *value* **remote** **vsan** *vsan-id*
6. switch(config)#**fcroute** *fcid* **netmask** **interface** *fc slot/port* **domain** *domain-id* **vsan** *vsan-id*

### DETAILED STEPS

|               | Command or Action                                                                                                                                                    | Purpose                                                                                                                                                                                             |
|---------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Step 1</b> | switch# <b>configuration terminal</b>                                                                                                                                | Enters configuration mode.                                                                                                                                                                          |
| <b>Step 2</b> | switch(config)# <b>fcroute</b> <i>fcid</i> <b>interface</b> <i>fc slot/port</i> <b>domain</b> <i>domain-id</i> <b>vsan</b> <i>vsan-id</i>                            | Configures the route for the specified Fibre Channel interface and domain. In this example, the specified interface is assigned an FC ID and a domain ID to the next hop switch.                    |
| <b>Step 3</b> | switch(config)# <b>fcroute</b> <i>fcid</i> <b>interface</b> <i>san-port-channel port</i> <b>domain</b> <i>domain-id</i> <b>vsan</b> <i>vsan-id</i>                   | Configures the route for the specified SAN port channel interface and domain. In this example, interface san-port-channel 1 is assigned an FC ID (0x111211) and a domain ID to the next hop switch. |
| <b>Step 4</b> | switch(config)# <b>fcroute</b> <i>fcid</i> <b>interface</b> <i>fc slot/port</i> <b>domain</b> <i>domain-id</i> <b>metric</b> <i>value</i> <b>vsan</b> <i>vsan-id</i> | Configures the static route for a specific FC ID and next hop domain ID and also assigns the cost of the route.<br><br>If the remote destination option is not specified, the default is direct.    |

|               | Command or Action                                                                                                                                                                  | Purpose                                                                                                                                                                                                                                            |
|---------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Step 5</b> | switch(config)# <b>fcroute</b> <i>fcid</i> <b>interface</b> <b>fc slot/port</b> <b>domain</b> <i>domain-id</i> <b>metric</b> <i>value</i> <b>remote</b> <b>vsan</b> <i>vsan-id</i> | Adds a static route to the RIB. If this is an active route and the FIB/FIB = Forwarding Information Base records are free, it is also added to the FIB.<br><br>If the cost (metric) of the route is not specified, the default is 10.              |
| <b>Step 6</b> | switch(config)# <b>fcroute</b> <i>fcid</i> <b>netmask</b> <b>interface</b> <b>fc slot/port</b> <b>domain</b> <i>domain-id</i> <b>vsan</b> <i>vsan-id</i>                           | Configures the netmask for the specified route the in interface (or SAN port channel). You can specify one of three routes: 0xff0000 matches only the domain, 0xffff00 matches the domain and the area, 0xffff matches the domain, area, and port. |

### In-Order Delivery

In-order delivery (IOD) of data frames guarantees frame delivery to a destination in the same order that they were sent by the originator.

Some Fibre Channel protocols or applications cannot handle out-of-order frame delivery. In these cases, switches in the Cisco Nexus 5000 Series preserve frame ordering in the frame flow. The source ID (SID), destination ID (DID), and optionally the originator exchange ID (OX ID) identify the flow of the frame.

On a switch with IOD enabled, all frames received by a specific ingress port and destined to a certain egress port are always delivered in the same order in which they were received.

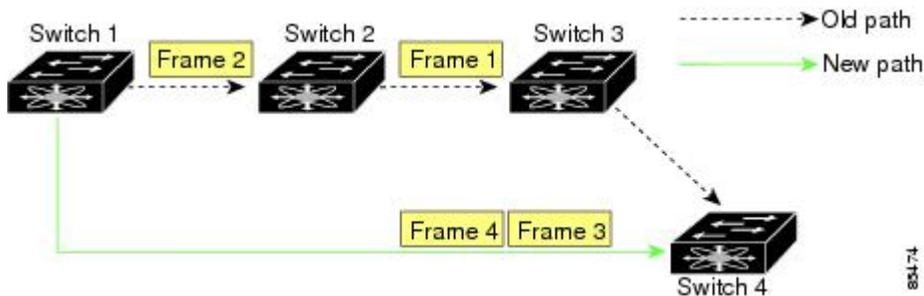
Use IOD only if your environment cannot support out-of-order frame delivery.

If you enable the in-order delivery feature, the graceful shutdown feature is not implemented.

### About Reordering Network Frames

When you experience a route change in the network, the new selected path may be faster or less congested than the old route.

**Figure 70: Route Change Delivery**



In the figure above, the new path from Switch 1 to Switch 4 is faster. In this scenario, Frame 3 and Frame 4 may be delivered before Frame 1 and Frame 2.

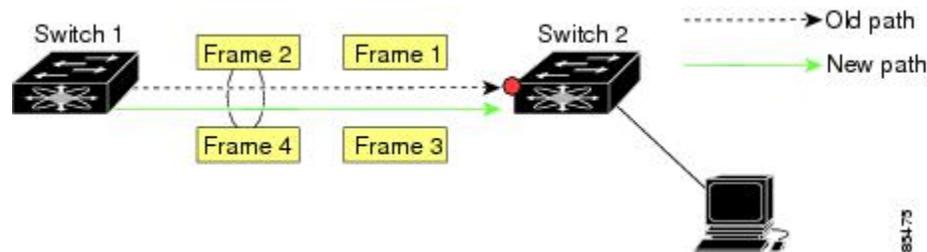
If the in-order guarantee feature is enabled, the frames within the network are delivered as follows:

- Frames in the network are delivered in the order in which they are transmitted.
- Frames that cannot be delivered in order within the network latency drop period are dropped inside the network.

### About Reordering SAN Port Channel Frames

When a link change occurs in a SAN port channel, the frames for the same exchange or the same flow can switch from one path to another faster path.

**Figure 71: Link Congestion Delivery**



In the figure above, the port of the old path (red dot) is congested. In this scenario, Frame 3 and Frame 4 can be delivered before Frame 1 and Frame 2.

When the in-order delivery feature is enabled and a port channel link change occurs, the frames crossing the SAN port channel are delivered as follows:

- Frames using the old path are delivered before new frames are accepted.
- The new frames are delivered through the new path after the network latency drop period has elapsed and all old frames are flushed.

Frames that cannot be delivered in order through the old path within the network latency drop period are dropped.

### Related Topics

[Configuring the Drop Latency Time, on page 657](#)

### About Enabling In-Order Delivery

You can enable the in-order delivery feature for a specific VSAN or for the entire switch. By default, in-order delivery is disabled on switches in the Cisco Nexus 5000 Series.

We recommend that you only enable this feature when devices that cannot handle any out-of-order frames are present in the switch. Load-balancing algorithms within the Cisco Nexus 5000 Series switch ensure that frames are delivered in order during normal fabric operation. The load-balancing algorithms based on source FC ID, destination FC ID, and exchange ID are enforced in hardware without any performance degradation. However, if the fabric encounters a failure and the in-order delivery feature is enabled, the recovery will be delayed because of an intentional pausing of fabric forwarding to purge the fabric of resident frames that could potentially be forwarded out-of-order.

## Enabling In-Order Delivery Globally

To ensure that the in-order delivery parameters are uniform across all VSANs on the switch, enable in-order delivery globally.

Only enable in-order delivery globally if this is a requirement across your entire fabric. Otherwise, enable IOD only for the VSANs that require this feature.

To enable in-order delivery for the switch, perform this task:

### SUMMARY STEPS

1. switch# **configuration terminal**
2. switch(config)# **in-order-guarantee**
3. switch(config)# **no in-order-guarantee**

### DETAILED STEPS

|               | Command or Action                            | Purpose                                                                                |
|---------------|----------------------------------------------|----------------------------------------------------------------------------------------|
| <b>Step 1</b> | switch# <b>configuration terminal</b>        | Enters configuration mode.                                                             |
| <b>Step 2</b> | switch(config)# <b>in-order-guarantee</b>    | Enables in-order delivery in the switch.                                               |
| <b>Step 3</b> | switch(config)# <b>no in-order-guarantee</b> | Reverts the switch to the factory defaults and disables the in-order delivery feature. |

## Enabling In-Order Delivery for a VSAN

When you create a VSAN, that VSAN automatically inherits the global in-order guarantee value. You can override this global value by enabling or disabling in-order guarantee for the new VSAN.

To use the lowest domain switch for the multicast tree computation, perform this task:

### SUMMARY STEPS

1. switch# **configuration terminal**
2. switch(config)# **in-order-guarantee vsan vsan-id**
3. switch(config)# **no in-order-guarantee vsan vsan-id**

### DETAILED STEPS

|               | Command or Action                                      | Purpose                                          |
|---------------|--------------------------------------------------------|--------------------------------------------------|
| <b>Step 1</b> | switch# <b>configuration terminal</b>                  | Enters configuration mode.                       |
| <b>Step 2</b> | switch(config)# <b>in-order-guarantee vsan vsan-id</b> | Enables in-order delivery in the specified VSAN. |

|        | Command or Action                                                | Purpose                                                                                                      |
|--------|------------------------------------------------------------------|--------------------------------------------------------------------------------------------------------------|
| Step 3 | switch(config)# <b>no in-order-guarantee vsan</b> <i>vsan-id</i> | Reverts the switch to the factory defaults and disables the in-order delivery feature in the specified VSAN. |

### Displaying the In-Order Delivery Status

Use the **show in-order-guarantee** command to display the present configuration status:

```
switch# show in-order-guarantee
global inorder delivery configuration:guaranteed
VSAN specific settings
vsan 1 inorder delivery:guaranteed
vsan 101 inorder delivery:not guaranteed
vsan 1000 inorder delivery:guaranteed
vsan 1001 inorder delivery:guaranteed
vsan 1682 inorder delivery:guaranteed
vsan 2001 inorder delivery:guaranteed
vsan 2009 inorder delivery:guaranteed
vsan 2456 inorder delivery:guaranteed
vsan 3277 inorder delivery:guaranteed
vsan 3451 inorder delivery:guaranteed
vsan 3452 inorder delivery:guaranteed
```

### Configuring the Drop Latency Time

You can change the default latency time for a network, a specified VSAN in a network, or for the entire switch.

To configure the network and the switch drop latency time, perform this task:

#### SUMMARY STEPS

1. switch# **configuration terminal**
2. switch(config)# **fcdroplateny network** *value*
3. switch(config)# **fcdroplateny network** *value vsan vsan-id*
4. switch(config)# **no fcdroplateny network** *value*

#### DETAILED STEPS

|        | Command or Action                                                     | Purpose                                                                                                                                                                                                                                                 |
|--------|-----------------------------------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Step 1 | switch# <b>configuration terminal</b>                                 | Enters configuration mode.                                                                                                                                                                                                                              |
| Step 2 | switch(config)# <b>fcdroplateny network</b> <i>value</i>              | Configures network drop latency time for the network. The valid range is 0 to 60000 msec. The default is 2000 msec.<br><br><b>Note</b> The network drop latency must be computed as the sum of all switch latencies of the longest path in the network. |
| Step 3 | switch(config)# <b>fcdroplateny network</b> <i>value vsan vsan-id</i> | Configures network drop latency time for the specified VSAN.                                                                                                                                                                                            |

|               | Command or Action                                              | Purpose                                                                                                |
|---------------|----------------------------------------------------------------|--------------------------------------------------------------------------------------------------------|
| <b>Step 4</b> | switch(config)# <b>no fcdroplateny network</b><br><i>value</i> | Removes the current fcdroplateny network configuration and reverts the switch to the factory defaults. |

### Displaying Latency Information

You can view the configured latency parameters using the **show fcdroplateny** command. The following example shows how to display network latency information:

```
switch# show fcdroplateny
switch latency value:500 milliseconds
global network latency value:2000 milliseconds
VSAN specific network latency settings
vsan 1 network latency:5000 milliseconds
vsan 2 network latency:2000 milliseconds
vsan 103 network latency:2000 milliseconds
vsan 460 network latency:500 milliseconds
```

### Flow Statistics Configuration

Flow statistics count the ingress traffic in the aggregated statistics table. You can collect two kinds of statistics:

- Aggregated flow statistics to count the traffic for a VSAN.
- Flow statistics to count the traffic for a source and destination ID pair in a VSAN.

#### About Flow Statistics

If you enable flow counters, you can enable a maximum of 1000 entries for aggregate flow and flow statistics. Be sure to assign an unused flow index for each new flow. The number space for flow index is shared between the aggregate flow statistics and the flow statistics.

#### Counting Aggregated Flow Statistics

To count the aggregated flow statistics for a VSAN, perform this task:

#### SUMMARY STEPS

1. switch# **configuration terminal**
2. switch(config)# **fcbflow stats aggregated index** *value vsan vsan-id*
3. switch(config)# **no fcbflow stats aggregated index** *value vsan vsan-id*

**DETAILED STEPS**

|               | <b>Command or Action</b>                                                   | <b>Purpose</b>                        |
|---------------|----------------------------------------------------------------------------|---------------------------------------|
| <b>Step 1</b> | switch# <b>configuration terminal</b>                                      | Enters configuration mode.            |
| <b>Step 2</b> | switch(config)# <b>fcflow stats aggregated index value vsan vsan-id</b>    | Enables the aggregated flow counter.  |
| <b>Step 3</b> | switch(config)# <b>no fcflow stats aggregated index value vsan vsan-id</b> | Disables the aggregated flow counter. |

**Counting Individual Flow Statistics**

To count the flow statistics for a source and destination FC ID in a VSAN, perform this task:

**SUMMARY STEPS**

1. switch# **configuration terminal**
2. switch(config)# **fcflow stats index value dest-fcid source-fcid netmask vsan vsan-id**
3. switch(config)# **no fcflow stats aggregated index value vsan vsan-id**

**DETAILED STEPS**

|               | <b>Command or Action</b>                                                                   | <b>Purpose</b>                                                                                                                                                                       |
|---------------|--------------------------------------------------------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Step 1</b> | switch# <b>configuration terminal</b>                                                      | Enters configuration mode.                                                                                                                                                           |
| <b>Step 2</b> | switch(config)# <b>fcflow stats index value dest-fcid source-fcid netmask vsan vsan-id</b> | Enables the flow counter.<br><b>Note</b> The source ID and the destination ID are specified in FC ID hex format (for example, 0x123aff). The mask can be one of 0xff0000 or 0xfffff. |
| <b>Step 3</b> | switch(config)# <b>no fcflow stats aggregated index value vsan vsan-id</b>                 | Disables the flow counter.                                                                                                                                                           |

**Clearing FIB Statistics**

Use the **clear fcflow stats** command to clear the aggregated flow counter. The following example clears the aggregated flow counters:

```
switch# clear fcflow stats aggregated index 1
```

The following example clears the flow counters for source and destination FC IDs:

```
switch# clear fcflow stats index 1
```

## Displaying Flow Statistics

Use the **show fcflow stats** commands to view flow statistics. The following example displays the aggregated flow summary:

```
switch# show fcflow stats aggregated
Idx VSAN frames

 6 1 42871
```

The following example displays the flow statistics:

```
switch# show fcflow stats
```

The following example displays flow index usage:

```
switch# show fcflow stats usage
2 flows configured
Configured flows : 3,7
```

The following example shows how to display global FSPF information for a specific VSAN:

```
switch# show fspf vsan 1
```

The following example shows how to display a summary of the FSPF database for a specified VSAN. If no additional parameters are specified, all LSRs in the database are displayed:

```
switch# show fspf database vsan 1
```

The following example shows how to display FSPF interface information:

```
switch# show fspf vsan 1 interface fc2/1
```

## Default FSPF Settings

The following table lists the default settings for FSPF features.

**Table 84: Default FSPF Settings**

| Parameters                             | Default                                        |
|----------------------------------------|------------------------------------------------|
| FSPF                                   | Enabled on all E ports and TE ports.           |
| SPF computation                        | Dynamic.                                       |
| SPF hold time                          | 0.                                             |
| Backbone region                        | 0.                                             |
| Acknowledgment interval (RxmtInterval) | 5 seconds.                                     |
| Refresh time (LSRefreshTime)           | 30 minutes.                                    |
| Maximum age (MaxAge)                   | 60 minutes.                                    |
| Hello interval                         | 20 seconds.                                    |
| Dead interval                          | 80 seconds.                                    |
| Distribution tree information          | Derived from the principal switch (root node). |

| <b>Parameters</b>         | <b>Default</b>                                                            |
|---------------------------|---------------------------------------------------------------------------|
| Routing table             | FSPF stores up to 16 equal cost paths to a given destination.             |
| Load balancing            | Based on destination ID and source ID on different, equal cost paths.     |
| In-order delivery         | Disabled.                                                                 |
| Drop latency              | Disabled.                                                                 |
| Static route cost         | If the cost (metric) of the route is not specified, the default is 10.    |
| Remote destination switch | If the remote destination switch is not specified, the default is direct. |
| Multicast routing         | Uses the principal switch to compute the multicast tree.                  |





# CHAPTER 43

## Managing FLOGI, Name Server, FDMI, and RSCN Databases

This chapter contains the following sections:

- [Managing FLOGI, Name Server, FDMI, and RSCN Databases, page 663](#)

### Managing FLOGI, Name Server, FDMI, and RSCN Databases

#### Information About Fabric Login

In a Fibre Channel fabric, each host or disk requires an FC ID. Use the **show flogi** command to verify if a storage device is displayed in the fabric login (FLOGI) table as in the following examples. If the required device is displayed in the FLOGI table, the fabric login is successful. Examine the FLOGI database on a switch that is directly connected to the host HBA and connected ports.

The following example shows how to verify the storage devices in the fabric login (FLOGI) table:

```
switch# show flogi database

INTERFACE VSAN FCID PORT NAME NODE NAME

fc2/3 1 0xb200e2 21:00:00:04:cf:27:25:2c 20:00:00:04:cf:27:25:2c
fc2/3 1 0xb200e1 21:00:00:04:cf:4c:18:61 20:00:00:04:cf:4c:18:61
fc2/3 1 0xb200d1 21:00:00:04:cf:4c:18:64 20:00:00:04:cf:4c:18:64
fc2/3 1 0xb200ce 21:00:00:04:cf:4c:16:fb 20:00:00:04:cf:4c:16:fb
fc2/3 1 0xb200cd 21:00:00:04:cf:4c:18:f7 20:00:00:04:cf:4c:18:f7
vfc3/1 2 0xb30100 10:00:00:05:30:00:49:63 20:00:00:05:30:00:49:5e
Total number of flogi = 6.
```

The following example shows how to verify the storage devices attached to a specific interface:

```
switch# show flogi database interface vfc1/1
INTERFACE VSAN FCID PORT NAME NODE NAME

vfc1/1 1 0x870000 20:00:00:1b:21:06:58:bc 10:00:00:1b:21:06:58:bc
Total number of flogi = 1.
```

The following example shows how to verify the storage devices associated with VSAN 1:

```
switch# show flogi database vsan 1
```

## Name Server Proxy

The name server functionality maintains a database containing the attributes for all hosts and storage devices in each VSAN. Name servers allow a database entry to be modified by a device that originally registered the information.

The proxy feature is useful when you need to modify (update or delete) the contents of a database entry that was previously registered by a different device.

### About Registering Name Server Proxies

All name server registration requests come from the same port whose parameter is registered or changed. If it does not, then the request is rejected.

This authorization enables WWNs to register specific parameters for another node.

### Registering Name Server Proxies

To register the name server proxy, perform this task:

#### SUMMARY STEPS

1. switch# **configuration terminal**
2. switch(config)#**fcns proxy-port *wwn-id* vsan *vsan-id***

#### DETAILED STEPS

|               | Command or Action                                                        | Purpose                                         |
|---------------|--------------------------------------------------------------------------|-------------------------------------------------|
| <b>Step 1</b> | switch# <b>configuration terminal</b>                                    | Enters configuration mode.                      |
| <b>Step 2</b> | switch(config)# <b>fcns proxy-port <i>wwn-id</i> vsan <i>vsan-id</i></b> | Configures a proxy port for the specified VSAN. |

### About Rejecting Duplicate pWWNs

You can prevent malicious or accidental log in using another device's pWWN by enabling the reject-duplicate-pwwn option. If you disable this option, these pWWNs are allowed to log in to the fabric and replace the first device in the name server database.

### Rejecting Duplicate pWWNs

To reject duplicate pWWNs, perform this task:

## SUMMARY STEPS

1. switch# **configuration terminal**
2. switch(config)# **fcns reject-duplicate-pwwn vsan vsan-id**
3. switch(config)# **no fcns reject-duplicate-pwwn vsan vsan-id**

## DETAILED STEPS

|               | Command or Action                                                 | Purpose                                                                                                             |
|---------------|-------------------------------------------------------------------|---------------------------------------------------------------------------------------------------------------------|
| <b>Step 1</b> | switch# <b>configuration terminal</b>                             | Enters configuration mode.                                                                                          |
| <b>Step 2</b> | switch(config)# <b>fcns reject-duplicate-pwwn vsan vsan-id</b>    | Logs out devices when they log into the fabric if the pWWNs already exist.                                          |
| <b>Step 3</b> | switch(config)# <b>no fcns reject-duplicate-pwwn vsan vsan-id</b> | Overwrites the first device's entry in the name server database with the new device having the same pWWN (default). |

### About Name Server Database Entries

The name server stores name entries for all hosts in the FCNS database. The name server permits an Nx port to register attributes during a PLOGI (to the name server) to obtain attributes of other hosts. These attributes are deregistered when the Nx port logs out either explicitly or implicitly.

In a multiswitch fabric configuration, the name server instances running on each switch shares information in a distributed database. One instance of the name server process runs on each switch.

### Displaying Name Server Database Entries

The following example shows how to display the name server database for all VSANs:

```
switch# show fcns database

FCID TYPE PWWN (VENDOR) FC4-TYPE:FEATURE

0x010000 N 50:06:0b:00:00:10:a7:80 (Cisco) scsi-fcp fc-gs
0x010001 N 10:00:00:05:30:00:24:63 (Cisco) ipfc
0x010002 N 50:06:04:82:c3:a0:98:52 (Company 1) scsi-fcp 250
0x010100 N 21:00:00:e0:8b:02:99:36 (Company A) scsi-fcp
0x020000 N 21:00:00:e0:8b:08:4b:20 (Company A)
0x020100 N 10:00:00:05:30:00:24:23 (Cisco) ipfc
0x020200 N 21:01:00:e0:8b:22:99:36 (Company A) scsi-fcp
```

The following example shows how to display the name server database and statistical information for a specified VSAN:

```
switch# show fcns database vsan 1
VSAN 1:

FCID TYPE PWWN (VENDOR) FC4-TYPE:FEATURE

0x030001 N 10:00:00:05:30:00:25:a3 (Cisco) ipfc
0x030101 NL 10:00:00:00:77:99:60:2c (Interphase)
0x030200 N 10:00:00:49:c9:28:c7:01
0xec0001 NL 21:00:00:20:37:a6:be:14 (Seagate) scsi-fcp
Total number of entries = 4
```

The following example shows how to display the name server database details for all VSANs:

```
switch# show fcns database detail
```

The following example shows how to display the name server database statistics for all VSANs:

```
switch# show fcns statistics
```

## FDMI

Cisco Nexus 5000 Series switches provide support for the Fabric-Device Management Interface (FDMI) functionality, as described in the FC-GS-4 standard. FDMI enables management of devices such as Fibre Channel host bus adapters (HBAs) through in-band communications. This addition complements the existing Fibre Channel name server and management server functions.

Using the FDMI functionality, the switch software can extract the following management information about attached HBAs and host operating systems without installing proprietary host agents:

- Manufacturer, model, and serial number
- Node name and node symbolic name
- Hardware, driver, and firmware versions
- Host operating system (OS) name and version number

All FDMI entries are stored in persistent storage and are retrieved when the FDMI process is started.

## Displaying FDMI

The following example shows how to display all HBA details for a specified VSAN:

```
switch# show fDMI database detail vsan 1
```

## RSCN

The Registered State Change Notification (RSCN) is a Fibre Channel service that informs hosts about changes in the fabric. Hosts can receive this information by registering with the fabric controller (through a State Change Registration (SCR) request). These notifications provide a timely indication of one or more of the following events:

- Disks joining or leaving the fabric
- A name server registration change
- A new zone enforcement
- IP address change

- Any other similar event that affects the operation of the host

### About RSCN Information

A switch RSCN (SW-RSCN) is sent to registered hosts and to all reachable switches in the fabric.



**Note** The switch sends an RSCN to notify registered nodes that a change has occurred. It is up to the nodes to query the name server again to obtain the new information. The details of the changed information are not delivered by the switch in the RSCN sent to the nodes.

### Displaying RSCN Information

The following example shows how to display registered device information:

```
switch# show rscn scr-table vsan 1
```



**Note** The SCR table is not configurable. It is populated when hosts send SCR frames with RSCN information. If hosts do not receive RSCN information, then the **show rscn scr-table** command will not return entries.

### About the multi-pid Option

If the RSCN multi-pid option is enabled, then RSCNs generated to the registered Nx ports may contain more than one affected port IDs. In this case, zoning rules are applied before putting the multiple affected port IDs together in a single RSCN. By enabling this option, you can reduce the number of RSCNs. For example, you have two disks (D1, D2) and a host (H) connected to switch 1. Host H is registered to receive RSCNs. D1, D2, and H belong to the same zone. If disks D1 and D2 are online at the same time, one of the following actions applies:

- The multi-pid option is disabled on switch 1— Two RSCNs are generated to host H: one for the disk D1 and another for disk D2.
- The multi-pid option is enabled on switch 1—A single RSCN is generated to host H, and the RSCN payload lists the affected port IDs (in this case, both D1 and D2).



**Note** Some Nx ports may not support multi-pid RSCN payloads. If so, disable the RSCN multi-pid option.

### Configuring the multi-pid Option

To configure the **multi-pid** option, perform this task:

#### SUMMARY STEPS

1. switch# **configuration terminal**
2. switch(config)# **rscn multi-pid vsan vsan-id**

**DETAILED STEPS**

|               | <b>Command or Action</b>                           | <b>Purpose</b>                                            |
|---------------|----------------------------------------------------|-----------------------------------------------------------|
| <b>Step 1</b> | switch# <b>configuration terminal</b>              | Enters configuration mode.                                |
| <b>Step 2</b> | switch(config)# <b>rscn multi-pid vsan vsan-id</b> | Sends RSCNs in a multi-pid format for the specified VSAN. |

**Suppressing Domain Format SW-RSCNs**

A domain format SW-RSCN is sent whenever the local switch name or the local switch management IP address changes. This SW-RSCN is sent to all other domains and switches over the ISLs. The remote switches can issue GMAL and GIELN commands to the switch that initiated the domain format SW-RSCN to determine what changed. Domain format SW-RSCNs can cause problems with some non-Cisco SAN switches.

To suppress the transmission of these SW-RSCNs over an ISL, perform this task:

**SUMMARY STEPS**

1. switch# **configuration terminal**
2. switch(config)# **rscn suppress domain-swrsn vsan vsan-id**

**DETAILED STEPS**

|               | <b>Command or Action</b>                                       | <b>Purpose</b>                                                            |
|---------------|----------------------------------------------------------------|---------------------------------------------------------------------------|
| <b>Step 1</b> | switch# <b>configuration terminal</b>                          | Enters configuration mode.                                                |
| <b>Step 2</b> | switch(config)# <b>rscn suppress domain-swrsn vsan vsan-id</b> | Suppresses transmission of domain format SW-RSCNs for the specified VSAN. |

**Clearing RSCN Statistics**

You can clear the counters and later view the counters for a different set of events. For example, you can keep track of how many RSCNs or SW-RSCNs are generated on a particular event (such as ONLINE or OFFLINE events). You can use these statistics to monitor responses for each event in the VSAN.

The following example shows how to clear the RSCN statistics for the specified VSAN:

```
switch# clear rscn statistics vsan 1
```

After clearing the RSCN statistics, you can view the cleared counters by entering the **show rscn statistics** command:

```
switch# show rscn statistics vsan 1
```

### Configuring the RSCN Timer

RSCN maintains a per VSAN event list queue, where the RSCN events are queued as they are generated. When the first RSCN event is queued, a per VSAN timer starts. Upon time-out, all the events are dequeued and coalesced RSCNs are sent to registered users. The default timer values minimize the number of coalesced RSCNs sent to registered users. Some deployments require smaller event timer values to track changes in the fabric.



**Note** The RSCN timer value must be the same on all switches in the VSAN.



**Note** Before performing a downgrade, make sure that you revert the RSCN timer value in your network to the default value. Failure to do so will disable the links across your VSANs and other devices.

To configure the RSCN timer, perform this task:

#### SUMMARY STEPS

1. switch# **configuration terminal**
2. switch(config)# **rscn distribute**
3. switch(config)# **rscn event-tov timeout vsan vsan-id**
4. switch(config)# **no rscn event-tov timeout vsan vsan-id**
5. switch(config)# **rscn commit vsan vsan-id**

#### DETAILED STEPS

|               | Command or Action                                             | Purpose                                                                                                                                                 |
|---------------|---------------------------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Step 1</b> | switch# <b>configuration terminal</b>                         | Enters configuration mode.                                                                                                                              |
| <b>Step 2</b> | switch(config)# <b>rscn distribute</b>                        | Enables RSCN timer configuration distribution.                                                                                                          |
| <b>Step 3</b> | switch(config)# <b>rscn event-tov timeout vsan vsan-id</b>    | Sets the event time-out value in milliseconds for the specified VSAN. The range is 0 to 2000 milliseconds. Setting a zero (0) value disables the timer. |
| <b>Step 4</b> | switch(config)# <b>no rscn event-tov timeout vsan vsan-id</b> | Reverts to the default value (2000 milliseconds for Fibre Channel VSANs).                                                                               |
| <b>Step 5</b> | switch(config)# <b>rscn commit vsan vsan-id</b>               | Commits the RSCN timer configuration to be distributed to the switches in the specified VSAN.                                                           |

## Verifying the RSCN Timer Configuration

You verify the RSCN timer configuration using the **show rscn event-tov vsan** command. The following example shows how to clear the RSCN statistics for VSAN 10:

```
switch# show rscn event-tov vsan 10
Event TOV : 1000 ms
```

## RSCN Timer Configuration Distribution

Because the timeout value for each switch is configured manually, a misconfiguration occurs when different switches time out at different times. This means different N-ports in a network can receive RSCNs at different times. Cisco Fabric Services (CFS) infrastructure alleviates this situation by automatically distributing the RSCN timer configuration information to all switches in a fabric. This also reduces the number of SW-RSCNs.

RSCN supports two modes, distributed and nondistributed. In distributed mode, RSCN uses CFS to distribute configuration to all switches in the fabric. In nondistributed mode, only the configuration commands on the local switch are affected.



### Note

All configuration commands are not distributed. Only the **rscn event-tov tov vsan vsan** command is distributed.



### Caution

Only the RSCN timer configuration is distributed.

The RSCN timer is registered with CFS during initialization and switchover. For high availability, if the RSCN timer distribution crashes and restarts or a switchover occurs, it resumes normal functionality from the state prior to the crash or switchover.

### Related Topics

[Using Cisco Fabric Services, on page 329](#)

### Enabling RSCN Timer Configuration Distribution

To enable RSCN timer configuration distribution, perform this task:

## SUMMARY STEPS

1. switch# **configuration terminal**
2. switch(config)# **rscn distribute**
3. switch(config)# **no rscn distribute**

## DETAILED STEPS

|               | Command or Action                      | Purpose                          |
|---------------|----------------------------------------|----------------------------------|
| <b>Step 1</b> | switch# <b>configuration terminal</b>  | Enters configuration mode.       |
| <b>Step 2</b> | switch(config)# <b>rscn distribute</b> | Enables RSCN timer distribution. |

|               | Command or Action                         | Purpose                                     |
|---------------|-------------------------------------------|---------------------------------------------|
| <b>Step 3</b> | switch(config)# <b>no rscn distribute</b> | Disables (default) RSCN timer distribution. |

#### *Locking the Fabric*

The first action that modifies the database creates the pending database and locks the feature in the VSAN. Once you lock the fabric, the following situations apply:

- No other user can make any configuration changes to this feature.
- A copy of the configuration database becomes the pending database along with the first active change.

#### *Committing the RSCN Timer Configuration Changes*

If you commit the changes made to the active database, the configuration is committed to all the switches in the fabric. On a successful commit, the configuration change is applied throughout the fabric and the lock is released.

To commit RSCN timer configuration changes, perform this task:

### SUMMARY STEPS

1. switch# **configuration terminal**
2. switch(config)# **rscn commit vsan timeout**

### DETAILED STEPS

|               | Command or Action                               | Purpose                         |
|---------------|-------------------------------------------------|---------------------------------|
| <b>Step 1</b> | switch# <b>configuration terminal</b>           | Enters configuration mode.      |
| <b>Step 2</b> | switch(config)# <b>rscn commit vsan timeout</b> | Commits the RSCN timer changes. |

#### *Discarding the RSCN Timer Configuration Changes*

If you discard (abort) the changes made to the pending database, the configuration database remains unaffected and the lock is released.

To discard RSCN timer configuration changes, perform this task:

### SUMMARY STEPS

1. switch# **configuration terminal**
2. switch(config)# **rscn abort vsan timeout**

**DETAILED STEPS**

|               | <b>Command or Action</b>                              | <b>Purpose</b>                                                                 |
|---------------|-------------------------------------------------------|--------------------------------------------------------------------------------|
| <b>Step 1</b> | switch# <b>configuration terminal</b>                 | Enters configuration mode.                                                     |
| <b>Step 2</b> | switch(config)# <b>rscn abort vsan <i>timeout</i></b> | Discards the RSCN timer changes and clears the pending configuration database. |

*Clearing a Locked Session*

If you have changed the RSCN timer configuration and have forgotten to release the lock by either committing or discarding the changes, an administrator can release the lock from any switch in the fabric. If the administrator performs this task, your changes to the pending database are discarded and the fabric lock is released.

The pending database is only available in the volatile directory and are subject to being discarded if the switch is restarted.

To use administrative privileges and release a locked DPVM session, use the **clear rscn session vsan** command in EXEC mode. The following example shows how to clear the RSCN session for VSAN 10:

```
switch# clear rscn session vsan 10
```

*Displaying RSCN Configuration Distribution Information*

The following example shows how to display the registration status for RSCN configuration distribution:

```
switch# show cfs application name rscn
Enabled : Yes
Timeout : 5s
Merge Capable : Yes
Scope : Logical
```



**Note** A merge failure results when the RSCN timer values are different on the merging fabrics.

The following example shows how to display the set of configuration commands that would take effect when you commit the configuration:



**Note** The pending database includes both existing and modified configuration.

```
switch# show rscn pending
rscn event-tov 2000 ms vsan 1
rscn event-tov 2000 ms vsan 2
rscn event-tov 300 ms vsan 10
```

The following example shows how to display the difference between pending and active configurations:

```
switch# show rscn pending-diff vsan 10
- rscn event-tov 2000 ms vsan 10
+ rscn event-tov 300 ms vsan 10
```

**Default RSCN Settings**

The following table lists the default settings for RSCN.

**Table 85: Default RSCN Settings**

| <b>Parameters</b>                     | <b>Default</b>                            |
|---------------------------------------|-------------------------------------------|
| RSCN timer value                      | 2000 milliseconds for Fibre Channel VSANs |
| RSCN timer configuration distribution | Disabled                                  |





# Discovering SCSI Targets

---

This chapter contains the following sections:

- [Discovering SCSI Targets, page 675](#)

## Discovering SCSI Targets

### Information About SCSI LUN Discovery

Small Computer System Interface (SCSI) targets include disks, tapes, and other storage devices. These targets do not register logical unit numbers (LUNs) with the name server.

The name server requires LUN information for the following reasons:

- To display LUN storage device information so that a Network Management System (NMS) can access this information.
- To report device capacity, serial number, and device ID information.
- To register the initiator and target features with the name server.

The SCSI LUN discovery feature uses the local domain controller Fibre Channel address. It uses the local domain controller as the source FC ID, and performs SCSI INQUIRY, REPORT LUNS, and READ CAPACITY commands on SCSI devices.

The SCSI LUN discovery feature is initiated on demand, through CLI or SNMP. This information is also synchronized with neighboring switches, if those switches belong to the Cisco Nexus 5000 Series.

### About Starting SCSI LUN Discovery

SCSI LUN discovery is done on demand.

Only Nx ports that are present in the name server database and that are registered as FC4 Type = SCSI\_FCP are discovered.

### Starting SCSI LUN Discovery

To start SCSI LUN discovery, perform this task:

**SUMMARY STEPS**

1. switch# **discover scsi-target** {**custom-list** | **local** | **remote** | **vsan** *vsan-id* **fcid** *fc-id*} **os** {**aix** | **hpux** | **linux** | **solaris** | **windows**} [**lun** | **target**]

**DETAILED STEPS**

|               | <b>Command or Action</b>                                                                                                                                                                                                                                               | <b>Purpose</b>                                                  |
|---------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------------------------------------------------------------|
| <b>Step 1</b> | switch# <b>discover scsi-target</b> { <b>custom-list</b>   <b>local</b>   <b>remote</b>   <b>vsan</b> <i>vsan-id</i> <b>fcid</b> <i>fc-id</i> } <b>os</b> { <b>aix</b>   <b>hpux</b>   <b>linux</b>   <b>solaris</b>   <b>windows</b> } [ <b>lun</b>   <b>target</b> ] | Discovers SCSI targets for the specified operating system (OS). |

**Examples of Starting SCSI LUN Discovery**

The following example discovers local SCSI targets for all operating systems (OSs):

```
switch# discover scsi-target local os all
discovery started
```

The following example discovers remote SCSI targets assigned to the AIX OS:

```
switch# discover scsi-target remote os aix
discovery started
```

The following example discovers SCSI targets for VSAN 1 and FC ID 0x9c03d6:

```
switch# discover scsi-target vsan 1 fcid 0x9c03d6
discover scsi-target vsan 1 fcid 0x9c03d6
VSAN: 1 FCID: 0x9c03d6 PWWN: 00:00:00:00:00:00:00:00
PRLI RSP: 0x01 SPARM: 0x0012
SCSI TYPE: 0 NLUNS: 1
Vendor: Company 4 Model: ST318203FC Rev: 0004
Other: 00:00:02:32:8b:00:50:0a
```

The following example discovers SCSI targets from the customized list assigned to the Linux OS:

```
switch# discover scsi-target custom-list os linux
discovery started
```

**About Initiating Customized Discovery**

Customized discovery consists of a list of VSAN and domain pairs that are selectively configured to initiate a discovery. Use the custom-list option to initiate this discovery. The domain ID is a number from 0 to 255 in decimal or a number from 0x0 to 0xFF in hex.

**Initiating Customized Discovery**

To initiate a customized discovery, perform this task:

**SUMMARY STEPS**

1. switch# **discover custom-list add vsan** *vsan-id* **domain** *domain-id*
2. switch# **discover custom-list delete vsan** *vsan-id* **domain** *domain-id*

## DETAILED STEPS

|               | Command or Action                                                                                | Purpose                                               |
|---------------|--------------------------------------------------------------------------------------------------|-------------------------------------------------------|
| <b>Step 1</b> | switch# <b>discover custom-list add vsan</b> <i>vsan-id</i> <b>domain</b><br><i>domain-id</i>    | Adds the specified entry to the custom list.          |
| <b>Step 2</b> | switch# <b>discover custom-list delete vsan</b> <i>vsan-id</i> <b>domain</b><br><i>domain-id</i> | Deletes the specified domain ID from the custom list. |

## Displaying SCSI LUN Information

Use the **show scsi-target** and **show fcns database** commands to display the results of the discovery.

The following example displays the discovered targets:

```
switch# show scsi-target status
discovery completed
```



### Note

This command takes several minutes to complete, especially if the fabric is large or if several devices are slow to respond.

The following example displays the FCNS database:

```
switch# show fcns database
```

The following example displays the SCSI target disks:

```
switch# show scsi-target disk
```

The following example displays the discovered LUNs on all operating systems:

```
switch# show scsi-target lun os all
```

The following example displays the port WWN that is assigned to each operating system (Windows, AIX, Solaris, Linux, or HPUX):

```
switch# show scsi-target pwn
```





# Advanced Fibre Channel Features and Concepts

This chapter contains the following sections:

- [Advanced Fibre Channel Features and Concepts, page 679](#)

## Advanced Fibre Channel Features and Concepts

### Fibre Channel Timeout Values

You can modify Fibre Channel protocol-related timer values for the switch by configuring the following timeout values (TOVs):

- Distributed services TOV (D\_S\_TOV)—The valid range is from 5,000 to 10,000 milliseconds. The default is 5,000 milliseconds.
- Error detect TOV (E\_D\_TOV)—The valid range is from 1,000 to 10,000 milliseconds. The default is 2,000 milliseconds. This value is matched with the other end during port initialization.
- Resource allocation TOV (R\_A\_TOV)—The valid range is from 5,000 to 10,000 milliseconds. The default is 10,000 milliseconds. This value is matched with the other end during port initialization.



**Note**

---

The fabric stability TOV (F\_S\_TOV) constant cannot be configured.

---

### Timer Configuration Across All VSANs

You can modify Fibre Channel protocol related timer values for the switch.



**Caution**

---

The D\_S\_TOV, E\_D\_TOV, and R\_A\_TOV values cannot be globally changed unless all VSANs in the switch are suspended.

---

**Note**

If a VSAN is not specified when you change the timer value, the changed value is applied to all VSANs in the switch.

To configure Fibre Channel timers across all VSANs, perform this task:

**SUMMARY STEPS**

1. switch# **configure terminal**
2. switch(config)# **fctimer R\_A\_TOV timeout**

**DETAILED STEPS**

|               | Command or Action                              | Purpose                                                                                                                                                           |
|---------------|------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Step 1</b> | switch# <b>configure terminal</b>              | Enters configuration mode.                                                                                                                                        |
| <b>Step 2</b> | switch(config)# <b>fctimer R_A_TOV timeout</b> | Configures the R_A_TOV timeout value for all VSANs. The units is milliseconds.<br><br>This type of configuration is not permitted unless all VSANs are suspended. |

**Timer Configuration Per-VSAN**

You can also issue the fctimer for a specified VSAN to configure different TOV values for VSANs with special links such as Fibre Channel. You can configure different E\_D\_TOV, R\_A\_TOV, and D\_S\_TOV values for individual VSANs. Active VSANs are suspended and activated when their timer values are changed.

**Note**

This configuration must be propagated to all switches in the fabric. Be sure to configure the same value in all switches in the fabric.

To configure per-VSAN Fibre Channel timers, perform this task:

**SUMMARY STEPS**

1. switch# **configuration terminal**
2. switch(config)# **fctimer D\_S\_TOV timeout vsan vsan-id**

**DETAILED STEPS**

|               | Command or Action                     | Purpose                    |
|---------------|---------------------------------------|----------------------------|
| <b>Step 1</b> | switch# <b>configuration terminal</b> | Enters configuration mode. |

|               | Command or Action                                           | Purpose                                                                                                                                                             |
|---------------|-------------------------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Step 2</b> | switch(config)# <b>fctimer D_S_TOV timeout vsan vsan-id</b> | Configures the D_S_TOV timeout value (in milliseconds) for the specified VSAN. Suspends the VSAN temporarily. You have the option to end this command, if required. |

The following example configures the timer value for VSAN 2:

```
switch(config)# fctimer D_S_TOV 6000 vsan 2
Warning: The vsan will be temporarily suspended when updating the timer value This
configuration would impact whole fabric. Do you want to continue? (y/n) y
Since this configuration is not propagated to other switches, please configure the same
value in all the switches
```

### About fctimer Distribution

You can enable per-VSAN fctimer fabric distribution for all Cisco SAN switches in the fabric. When you perform fctimer configurations, and distribution is enabled, that configuration is distributed to all the switches in the fabric.

You automatically acquire a fabric-wide lock when you enter the first configuration command after you enabled distribution in a switch. The fctimer application uses the effective and pending database model to store or commit the commands based on your configuration.

### Related Topics

[Using Cisco Fabric Services, on page 329](#)

### Enabling or Disabling fctimer Distribution

To enable or disable fctimer fabric distribution, perform this task:

#### SUMMARY STEPS

1. switch# **configuration terminal**
2. switch(config)# **fctimer distribute**
3. switch(config)# **no fctimer distribute**

#### DETAILED STEPS

|               | Command or Action                         | Purpose                                                                                                                                                               |
|---------------|-------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Step 1</b> | switch# <b>configuration terminal</b>     | Enters configuration mode.                                                                                                                                            |
| <b>Step 2</b> | switch(config)# <b>fctimer distribute</b> | Enables fctimer configuration distribution to all switches in the fabric. Acquires a fabric lock and stores all future configuration changes in the pending database. |

|        | Command or Action                            | Purpose                                                                              |
|--------|----------------------------------------------|--------------------------------------------------------------------------------------|
| Step 3 | switch(config)# <b>no fctimer distribute</b> | Disables (default) fctimer configuration distribution to all switches in the fabric. |

### Committing fctimer Changes

When you commit the fctimer configuration changes, the effective database is overwritten by the configuration changes in the pending database and all the switches in the fabric receive the same configuration. When you commit the fctimer configuration changes without implementing the session feature, the fctimer configurations are distributed to all the switches in the physical fabric.

To commit the fctimer configuration changes, perform this task:

#### SUMMARY STEPS

1. switch# **configuration terminal**
2. switch(config)# **fctimer commit**

#### DETAILED STEPS

|        | Command or Action                     | Purpose                                                                                                                                                                             |
|--------|---------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Step 1 | switch# <b>configuration terminal</b> | Enters configuration mode.                                                                                                                                                          |
| Step 2 | switch(config)# <b>fctimer commit</b> | Distributes the fctimer configuration changes to all switches in the fabric and releases the lock. Overwrites the effective database with the changes made to the pending database. |

### Discarding fctimer Changes

After making the configuration changes, you can choose to discard the changes by discarding the changes instead of committing them. In either case, the lock is released.

To discard the fctimer configuration changes, perform this task:

#### SUMMARY STEPS

1. switch# **configuration terminal**
2. switch(config)# **fctimer abort**

**DETAILED STEPS**

|               | <b>Command or Action</b>              | <b>Purpose</b>                                                                                   |
|---------------|---------------------------------------|--------------------------------------------------------------------------------------------------|
| <b>Step 1</b> | switch# <b>configuration terminal</b> | Enters configuration mode.                                                                       |
| <b>Step 2</b> | switch(config)# <b>fctimer abort</b>  | Discards the fctimer configuration changes in the pending database and releases the fabric lock. |

**Fabric Lock Override**

If you have performed a fctimer fabric task and have forgotten to release the lock by either committing or discarding the changes, an administrator can release the lock from any switch in the fabric. If the administrator performs this task, your changes to the pending database are discarded and the fabric lock is released.

The changes are only available in the volatile directory and are subject to being discarded if the switch is restarted.

To use administrative privileges and release a locked fctimer session, use the **clear fctimer session** command.

```
switch# clear fctimer session
```

**Fabric Database Merge Guidelines**

When merging two fabrics, follow these guidelines:

- Be aware of the following merge conditions:
  - The merge protocol is not implemented for distribution of the fctimer values. You must manually merge the fctimer values when a fabric is merged.
  - The per-VSAN fctimer configuration is distributed in the physical fabric.
  - The fctimer configuration is only applied to those switches containing the VSAN with a modified fctimer value.
  - The global fctimer values are not distributed.
- Do not configure global timer values when distribution is enabled.

**Note**

The number of pending fctimer configuration operations cannot be more than 15. After 15 operations, you must commit or abort the pending configurations before performing any more operations.

**Related Topics**

[CFS Merge Support, on page 334](#)

## Verifying Configured fctimer Values

Use the **show fctimer** command to display the configured fctimer values. The following example displays the configured global TOVs:

```
switch# show fctimer
F_S_TOV D_S_TOV E_D_TOV R_A_TOV

5000 ms 5000 ms 2000 ms 10000 ms
```



### Note

The F\_S\_TOV constant, though not configured, is displayed in the output of the **show fctimer** command.

The following example displays the configured TOV for VSAN 10:

```
switch# show fctimer vsan 10
vsan no. F_S_TOV D_S_TOV E_D_TOV R_A_TOV

10 5000 ms 5000 ms 3000 ms 10000 ms
```

## World Wide Names

The world wide name (WWN) in the switch is equivalent to the Ethernet MAC address. As with the MAC address, you must uniquely associate the WWN to a single device. The principal switch selection and the allocation of domain IDs rely on the WWN.

Cisco Nexus 5000 Series switches support three network address authority (NAA) address formats. (see the following table).

**Table 86: Standardized NAA WWN Formats**

| NAA Address         | NAA Type       | WWN Format               |                    |
|---------------------|----------------|--------------------------|--------------------|
| IEEE 48-bit address | Type 1 = 0001b | 000 0000 0000b           | 48-bit MAC address |
| IEEE extended       | Type 2 = 0010b | Locally assigned         | 48-bit MAC address |
| IEEE registered     | Type 5 = 0101b | IEEE company ID: 24 bits | VSID: 36 bits      |



### Caution

Changes to the world-wide names should be made by an administrator or individual who is completely familiar with switch operations.

## Verifying WWN Information

Use the **show wwn** commands to display the status of the WWN configuration. The following example displays the status of all WWNs:

```
switch# show wwn status
Type Configured Available Resvd. Alarm State

1 64 48 (75%) 16 NONE
2,5 524288 442368 (84%) 73728 NONE
```

The following example displays the information for block ID 51:

```
switch# show wwn status block-id 51
WWNs in this block: 21:00:ac:16:5e:52:00:03 to 21:ff:ac:16:5e:52:00:03
Num. of WWNs:: Configured: 256 Allocated: 0 Available: 256
```

Block Allocation Status: FREE

The following example displays the WWN for a specific switch:

```
switch# show wwn switch
Switch WWN is 20:00:ac:16:5e:52:00:00
```

## Link Initialization WWN Usage

Exchange Link Protocol (ELP) and Exchange Fabric Protocol (EFP) use WWNs during link initialization. ELPs and EFPs both use the VSAN WWN by default during link initialization. However, the ELP usage changes based on the peer switch's usage:

- If the peer switch ELP uses the switch WWN, then the local switch also uses the switch WWN.
- If the peer switch ELP uses the VSAN WWN, then the local switch also uses the VSAN WWN.

## Configuring a Secondary MAC Address

To allocate secondary MAC addresses, perform this task:

### SUMMARY STEPS

1. switch# **configuration terminal**
2. switch(config)# **wwn secondary-mac wwn-id range value**

### DETAILED STEPS

|        | Command or Action                                           | Purpose                                                              |
|--------|-------------------------------------------------------------|----------------------------------------------------------------------|
| Step 1 | switch# <b>configuration terminal</b>                       | Enters configuration mode.                                           |
| Step 2 | switch(config)# <b>wwn secondary-mac wwn-id range value</b> | Configures the secondary MAC address. This command cannot be undone. |

The following example shows how to configure the secondary MAC address:

```
switch(config)# wwn secondary-mac 00:99:55:77:55:55 range 64
This command CANNOT be undone.
Please enter the BASE MAC ADDRESS again: 00:99:55:77:55:55
Please enter the mac address RANGE again: 64
From now on WWN allocation would be based on new MACs. Are you sure? (yes/no) no
You entered: no. Secondary MAC NOT programmed
```

## FC ID Allocation for HBAs

Fibre Channel standards require a unique FC ID to be allocated to an N port attached to an F port in any switch. To conserve the number of FC IDs used, Cisco Nexus 5000 Series switches use a special allocation scheme.

Some HBAs do not discover targets that have FC IDs with the same domain and area. The switch software maintains a list of tested company IDs that do not exhibit this behavior. These HBAs are allocated with single FC IDs. If the HBA can discover targets within the same domain and area, a full area is allocated.

To allow further scalability for switches with numerous ports, the switch software maintains a list of HBAs that can discover targets within the same domain and area. Each HBA is identified by its company ID (also known as Organizational Unique Identifier, or OUI) used in the pWWN during a fabric log in. A full area is allocated to the N ports with company IDs that are listed and for the others, a single FC ID is allocated. Regardless of the type (whole area or single) of FC ID allocated, the FC ID entries remain persistent.

### Default Company ID List

All Cisco Nexus 5000 Series switches contain a default list of company IDs that require area allocation. Using the company ID reduces the number of configured persistent FC ID entries. You can configure or modify these entries using the CLI.



#### Caution

Persistent entries take precedence over company ID configuration. If the HBA fails to discover a target, verify that the HBA and the target are connected to the same switch and have the same area in their FC IDs, then perform the following procedure:

- 1 Shut down the port connected to the HBA.
- 2 Clear the persistent FC ID entry.
- 3 Get the company ID from the port WWN.
- 4 Add the company ID to the list that requires area allocation.
- 5 Bring up the port.

The list of company IDs have the following characteristics:

- A persistent FC ID configuration always takes precedence over the list of company IDs. Even if the company ID is configured to receive an area, the persistent FC ID configuration results in the allocation of a single FC ID.
- New company IDs added to subsequent releases are automatically added to existing company IDs.
- The list of company IDs is saved as part of the running and saved configuration.
- The list of company IDs is used only when the fcinterop FC ID allocation scheme is in auto mode. By default, the interop FC ID allocation is set to auto, unless changed.



**Tip** We recommend that you set the `fcinterop FC ID` allocation scheme to auto and use the company ID list and persistent FC ID configuration to manipulate the FC ID device allocation.

Use the **fcinterop FCID allocation auto** command to change the FC ID allocation and the **show running-config** command to view the currently allocated mode.

- When you enter a **write erase**, the list inherits the default list of company IDs shipped with a relevant release.

To allocate company IDs, perform this task:

## SUMMARY STEPS

1. `switch# configuration terminal`
2. `switch(config)# fcid-allocation area company-id value`
3. `switch(config)# no fcid-allocation area company-id value`

## DETAILED STEPS

|               | Command or Action                                                     | Purpose                                     |
|---------------|-----------------------------------------------------------------------|---------------------------------------------|
| <b>Step 1</b> | <code>switch# configuration terminal</code>                           | Enters configuration mode.                  |
| <b>Step 2</b> | <code>switch(config)# fcid-allocation area company-id value</code>    | Adds a new company ID to the default list.  |
| <b>Step 3</b> | <code>switch(config)# no fcid-allocation area company-id value</code> | Deletes a company ID from the default list. |

The following example adds a new company ID to the default list:

```
switch(config)# fcid-allocation area company-id 0x003223
```

## Verifying the Company ID Configuration

You can view the configured company IDs by entering the **show fcid-allocation area** command. Default entries are listed first and the user-added entries are listed next. Entries are listed even if they were part of the default list and you later removed them.

The following example displays the list of default and configured company IDs:

```
switch# show fcid-allocation area
FCID area allocation company id info:
00:50:2E <----- Default entry
00:50:8B
00:60:B0
00:A0:B8
00:E0:69
00:30:AE + <----- User-added entry
00:32:23 +
00:E0:8B * <----- Explicitly deleted entry (from the original default list)
Total company ids: 7
+ - Additional user configured company ids.
* - Explicitly deleted company ids from default list.
```

You can implicitly derive the default entries shipped with a specific release by combining the list of Company IDs displayed without any identification with the list of deleted entries.

You can also view or obtain the company IDs in a specific WWN by entering the **show fcid-allocation company-id-from-wwn** command. Some WWN formats do not support company IDs. In these cases, you may need to configure the FC ID persistent entry.

The following example displays the company ID for the specified WWN:

```
switch# show fcid-allocation company-id-from-wwn 20:00:00:05:30:00:21:60
Extracted Company ID: 0x000530
```

## Switch Interoperability

Interoperability enables the products of multiple vendors to interwork with each other. Fibre Channel standards guide vendors towards common external Fibre Channel interfaces.

Not all vendors follow the standards in the same way, which results in the need for interoperability modes. This section briefly explains the basic concepts of these modes.

Each vendor has a regular mode and an equivalent interoperability mode, which specifically turns off advanced or proprietary features and provides the product with a standards-compliant implementation.



### Note

For more information on configuring interoperability for Cisco Nexus 5000 Series switches, see the *Cisco MDS 9000 Family Switch-to-Switch Interoperability Configuration Guide*

## About Interop Mode

Cisco NX-OS software supports the following four interop modes:

- Mode 1— Standards-based interop mode that requires all other vendors in the fabric to be in interop mode.
- Mode 2—Brocade native mode (Core PID 0).
- Mode 3—Brocade native mode (Core PID 1).
- Mode 4—McData native mode.

For information about configuring interop modes 2, 3, and 4, see the *Cisco MDS 9000 Family Switch-to-Switch Interoperability Configuration Guide*, available at the following location: [http://www.cisco.com/en/US/docs/storage/san\\_switches/mds9000/interoperability/guide/intopgd.html](http://www.cisco.com/en/US/docs/storage/san_switches/mds9000/interoperability/guide/intopgd.html)

The following table lists the changes in switch operation when you enable interoperability mode. These changes are specific to Cisco Nexus 5000 Series switches while in interop mode.

**Table 87: Changes in Switch Operation When Interoperability Is Enabled**

| Switch Feature | Changes if Interoperability Is Enabled                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
|----------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Domain IDs     | <p>Some vendors cannot use the full range of 239 domains within a fabric.</p> <p>Domain IDs are restricted to the range 97 to 127, to accommodate McData's nominal restriction to this same range. Domain IDs can either be static or preferred, which operate as follows:</p> <ul style="list-style-type: none"> <li>• Static: Cisco switches accept only one domain ID; if a switch does not get that domain ID it isolates itself from the fabric.</li> <li>• Preferred: If the switch does not get its requested domain ID, it accepts any assigned domain ID.</li> </ul> |
| Timers         | All Fibre Channel timers must be the same on all switches as these values are exchanged by E ports when establishing an ISL. The timers are F_S_TOV, D_S_TOV, E_D_TOV, and R_A_TOV.                                                                                                                                                                                                                                                                                                                                                                                           |
| F_S_TOV        | Verify that the Fabric Stability Time Out Value timers match exactly.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| D_S_TOV        | Verify that the Distributed Services Time Out Value timers match exactly.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| E_D_TOV        | Verify that the Error Detect Time Out Value timers match exactly.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| R_A_TOV        | Verify that the Resource Allocation Time Out Value timers match exactly.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| Trunking       | Trunking is not supported between two different vendor's switches. This feature may be disabled on a per port or per switch basis.                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| Default zone   | The default zone operation of permit (all nodes can see all other nodes) or deny (all nodes are isolated when not explicitly placed in a zone) may change.                                                                                                                                                                                                                                                                                                                                                                                                                    |

| Switch Feature                       | Changes if Interoperability Is Enabled                                                                                                                                                                                                                                                                                                                                                                                                    |
|--------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Zoning attributes                    | <p>Zones may be limited to the pWWN and other proprietary zoning methods (physical port number) may be eliminated.</p> <p><b>Note</b> On a Brocade switch, use the <b>cfgsave</b> command to save fabric-wide zoning configuration. This command does not have any effect on Cisco Nexus 5000 Series switches if they are part of the same fabric. You must explicitly save the configuration on each Cisco Nexus 5000 Series switch.</p> |
| Zone propagation                     | <p>Some vendors do not pass the full zone configuration to other switches, only the active zone set gets passed.</p> <p>Verify that the active zone set or zone configuration has correctly propagated to the other switches in the fabric.</p>                                                                                                                                                                                           |
| VSAN                                 | <p>Interop mode only affects the specified VSAN.</p> <p><b>Note</b> Interop modes cannot be enabled on FICON-enabled VSANs.</p>                                                                                                                                                                                                                                                                                                           |
| TE ports and SAN port channels       | <p>TE ports and SAN port channels cannot be used to connect Cisco switches to non-Cisco SAN switches. Only E ports can be used to connect to non-Cisco SAN switches. TE ports and SAN port channels can still be used to connect a Cisco switch to other Cisco SAN switches even when in interop mode.</p>                                                                                                                                |
| FSPF                                 | <p>The routing of frames within the fabric is not changed by the introduction of interop mode. The switch continues to use src-id, dst-id, and ox-id to load balance across multiple ISL links.</p>                                                                                                                                                                                                                                       |
| Domain reconfiguration disruptive    | <p>This is a switch-wide impacting event. Brocade and McData require the entire switch to be placed in offline mode and/or rebooted when changing domain IDs.</p>                                                                                                                                                                                                                                                                         |
| Domain reconfiguration nondisruptive | <p>This event is limited to the affected VSAN. Cisco Nexus 5000 Series switches have the capability to restart only the domain manager process for the affected VSAN and not the entire switch.</p>                                                                                                                                                                                                                                       |
| Name server                          | <p>Verify that all vendors have the correct values in their respective name server database.</p>                                                                                                                                                                                                                                                                                                                                          |

## Configuring Interop Mode 1

The interop mode1 in Cisco Nexus 5000 Series switches can be enabled disruptively or nondisruptively.



**Note** Brocade's `msplmgmtdeactivate` command must explicitly be run prior to connecting from a Brocade switch to either Cisco Nexus 5000 Series switches or to McData switches. This command uses Brocade proprietary frames to exchange platform information, which Cisco Nexus 5000 Series switches or McData switches do not understand. Rejecting these frames causes the common E ports to become isolated.

To configure interop mode 1 in any switch in the Cisco Nexus 5000 Series, perform this task:

### SUMMARY STEPS

1. Place the VSAN of the E ports that connect to the OEM switch in interoperability mode.
2. Assign a domain ID in the range of 97 (0x61) through 127 (0x7F).
3. Change the Fibre Channel timers (if they have been changed from the system defaults).
4. When making changes to the domain, you may or may not need to restart the domain manager function for the altered VSAN.

### DETAILED STEPS

**Step 1** Place the VSAN of the E ports that connect to the OEM switch in interoperability mode.

```
switch# configuration terminal
switch(config)# vsan database
switch(config-vsan-db)# vsan 1 interop 1
switch(config-vsan-db)# exit
```

**Step 2** Assign a domain ID in the range of 97 (0x61) through 127 (0x7F).

**Note** This is an limitation imposed by the McData switches.

In Cisco Nexus 5000 Series switches, the default is to request an ID from the principal switch. If the preferred option is used, Cisco Nexus 5000 Series switches request a specific ID, but still join the fabric if the principal switch assigns a different ID. If the static option is used, the Cisco Nexus 5000 Series switches do not join the fabric unless the principal switch agrees and assigns the requested ID.

**Note** When changing the domain ID, the FC IDs assigned to N ports also change.

**Step 3** Change the Fibre Channel timers (if they have been changed from the system defaults).

**Note** The Cisco Nexus 5000 Series, Brocade, and McData FC Error Detect (ED\_TOV) and Resource Allocation (RA\_TOV) timers default to the same values. They can be changed if needed. The RA\_TOV default is 10 seconds, and the ED\_TOV default is 2 seconds. Per the FC-SW2 standard, these values must be the same on each switch within the fabric.

```
switch(config)# fctimer e_d_tov ?
<1000-100000> E_D_TOV in milliseconds(1000-100000)

switch(config)# fctimer r_a_tov ?
<5000-100000> R_A_TOV in milliseconds(5000-100000)
```

**Step 4** When making changes to the domain, you may or may not need to restart the domain manager function for the altered VSAN.

- Force a fabric reconfiguration with the **disruptive** option.  

```
switch(config)# fcdomain restart disruptive vsan 1
```

or

- Do not force a fabric reconfiguration.  

```
switch(config)# fcdomain restart vsan 1
```

### Verifying Interoperating Status

This section highlights the commands used to verify if the fabric is up and running in interoperability mode. To verify the resulting status of entering the interoperability command in any switch in the Cisco Nexus 5000 Series, perform this task:

### SUMMARY STEPS

1. Verify the software version.
2. Verify if the interface states are as required by your configuration.
3. Verify if you are running the desired configuration.
4. Verify if the interoperability mode is active.
5. Verify the domain ID.
6. Verify the local principal switch status.
7. Verify the next hop and destination for the switch.
8. Verify the name server information.

### DETAILED STEPS

**Step 1** Verify the software version.

**Example:**

```
switch# show version
Cisco Nexus Operating System (NX-OS) Software

TAC support: http://www.cisco.com/tac
Copyright (c) 2002-2008, Cisco Systems, Inc. All rights reserved.
The copyrights to certain works contained herein are owned by
other third parties and are used and distributed under license.

Some parts of this software are covered under the GNU Public
License. A copy of the license is available at
```

<http://www.gnu.org/licenses/gpl.html>.

```
Software
 BIOS: version 1.2.0
 loader: version N/A
 kickstart: version 4.0(1a)N1(1)
 system: version 4.0(1a)N1(1)
 BIOS compile time: 06/19/08
 kickstart image file is: bootflash:/n5000-uk9-kickstart.4.0.1a.N1.latest.bin
 kickstart compile time: 11/25/2008 6:00:00 [11/25/2008 14:17:12]
 system image file is: bootflash:/n5000-uk9.4.0.1a.N1.latest.bin
 system compile time: 11/25/2008 6:00:00 [11/25/2008 14:59:49]
Hardware
 cisco Nexus5020 Chassis ("40x10GE/Supervisor")
 Intel(R) Celeron(R) M CPU with 2074308 kB of memory.
 Processor Board ID JAB120900PJ
 Device name: switch
 bootflash: 1003520 kB

Kernel uptime is 0 day(s), 1 hour(s), 29 minute(s), 55 second(s)

Last reset at 510130 usecs after Wed Nov 26 18:12:23 2008
Reason: Reset Requested by CLI command reload
System version: 4.0(1a)N1(1)
Service:
```

```
plugin
 Core Plugin, Ethernet Plugin
```

**Step 2** Verify if the interface states are as required by your configuration.

**Example:**

```
switch# show interface brief
```

```

```

| Interface | Vsan | Admin Mode | Admin Trunk Mode | Status       | SFP | Oper Mode | Oper Speed (Gbps) | Port Channel |
|-----------|------|------------|------------------|--------------|-----|-----------|-------------------|--------------|
| fc3/1     | 1    | E          | on               | trunking     | sw1 | TE        | 2                 | --           |
| fc3/2     | 1    | auto       | on               | sfpAbsent    | --  | --        |                   | --           |
| fc3/3     | 1    | E          | on               | trunking     | sw1 | TE        | 2                 | --           |
| fc3/4     | 1    | auto       | on               | sfpAbsent    | --  | --        |                   | --           |
| fc3/5     | 1    | auto       | auto             | notConnected | sw1 | --        |                   | --           |
| fc3/6     | 1    | auto       | on               | sfpAbsent    | --  | --        |                   | --           |
| fc3/7     | 1    | auto       | auto             | sfpAbsent    | --  | --        |                   | --           |
| fc3/8     | 1    | auto       | auto             | sfpAbsent    | --  | --        |                   | --           |

```

```

**Step 3** Verify if you are running the desired configuration.

**Example:**

```
switch# show running-config
```

```
Building Configuration...
```

```

interface fc2/1
no shutdown
interface fc2/2
no shutdown
interface fc2/3
interface fc2/4
<snip>
interface mgmt0
ip address 6.1.1.96 255.255.255.0
switchport encap default
no shutdown
vsan database
vsan 1 interop
boot system bootflash:/nx5000-system-23e.bin
boot kickstart bootflash:/nx5000-kickstart-23e.bin
callhome
fcdomain domain 100 preferred vsan 1
ip route 6.1.1.0 255.255.255.0 6.1.1.1
ip routing
line console
 databits 5
 speed 110
logging linecard
ssh key rsa 512 force
ssh server enable
switchname switch
username admin password 5 1Li8/fBYX$SNc72.xt4nTXpSnR9OUFB/ role network-admin

```

**Step 4** Verify if the interoperability mode is active.

**Example:**

```

switch# show vsan 1
vsan 1 information
 name:VSAN0001 state:active
 interoperability mode:yes <----- verify mode
 loadbalancing:src-id/dst-id/oxid
 operational state:up

```

**Step 5** Verify the domain ID.**Example:**

```
switch# show fcdomain vsan 1
```

The local switch is a Subordinated Switch.

Local switch run time information:

```
State: Stable
Local switch WWN: 20:01:00:05:30:00:51:1f
Running fabric name: 10:00:00:60:69:22:32:91
Running priority: 128
Current domain ID: 0x64(100) <-----verify domain id
```

Local switch configuration information:

```
State: Enabled
Auto-reconfiguration: Disabled
Contiguous-allocation: Disabled
Configured fabric name: 41:6e:64:69:61:6d:6f:21
Configured priority: 128
Configured domain ID: 0x64(100) (preferred)
```

Principal switch run time information:

```
Running priority: 2
```

| Interface | Role       | RCF-reject |
|-----------|------------|------------|
| fc2/1     | Downstream | Disabled   |
| fc2/2     | Downstream | Disabled   |
| fc2/4     | Upstream   | Disabled   |

**Step 6** Verify the local principal switch status.**Example:**

```
switch# show fcdomain domain-list vsan 1
```

Number of domains: 5

| Domain ID | WWN                             |
|-----------|---------------------------------|
| 0x61(97)  | 10:00:00:60:69:50:0c:fe         |
| 0x62(98)  | 20:01:00:05:30:00:47:9f         |
| 0x63(99)  | 10:00:00:60:69:c0:0c:1d         |
| 0x64(100) | 20:01:00:05:30:00:51:1f [Local] |

```
0x65 (101) 10:00:00:60:69:22:32:91 [Principal]

```

**Step 7** Verify the next hop and destination for the switch.

**Example:**

```
switch# show fspf internal route vsan 1
```

```
FSPF Unicast Routes
```

```

```

| VSAN Number | Dest Domain | Route Cost | Next hops      |
|-------------|-------------|------------|----------------|
| 1           | 0x61 (97)   | 500        | fc2/2          |
| 1           | 0x62 (98)   | 1000       | fc2/1<br>fc2/2 |
| 1           | 0x63 (99)   | 500        | fc2/1          |
| 1           | 0x65 (101)  | 1000       | fc2/4          |

**Step 8** Verify the name server information.

**Example:**

```
switch# show fcns data vsan 1
```

```
VSAN 1:
```

```

```

| FCID     | TYPE | PWWN                    | (VENDOR)  | FC4-TYPE:FEATURE |
|----------|------|-------------------------|-----------|------------------|
| 0x610400 | N    | 10:00:00:00:c9:24:3d:90 | (Emulex)  | scsi-fcp         |
| 0x6105dc | NL   | 21:00:00:20:37:28:31:6d | (Seagate) | scsi-fcp         |
| 0x6105e0 | NL   | 21:00:00:20:37:28:24:7b | (Seagate) | scsi-fcp         |
| 0x6105e1 | NL   | 21:00:00:20:37:28:22:ea | (Seagate) | scsi-fcp         |
| 0x6105e2 | NL   | 21:00:00:20:37:28:2e:65 | (Seagate) | scsi-fcp         |
| 0x6105e4 | NL   | 21:00:00:20:37:28:26:0d | (Seagate) | scsi-fcp         |
| 0x630400 | N    | 10:00:00:00:c9:24:3f:75 | (Emulex)  | scsi-fcp         |
| 0x630500 | N    | 50:06:01:60:88:02:90:cb |           | scsi-fcp         |
| 0x6514e2 | NL   | 21:00:00:20:37:a7:ca:b7 | (Seagate) | scsi-fcp         |
| 0x6514e4 | NL   | 21:00:00:20:37:a7:c7:e0 | (Seagate) | scsi-fcp         |
| 0x6514e8 | NL   | 21:00:00:20:37:a7:c7:df | (Seagate) | scsi-fcp         |
| 0x651500 | N    | 10:00:00:e0:69:f0:43:9f | (JNI)     |                  |

```
Total number of entries = 12
```

**Note** The Cisco switch name server shows both local and remote entries, and does not time out the entries.

## Default Settings for Advanced Features

The following table lists the default settings for the features included in this chapter.

**Table 88: Default Settings for Advanced Features**

| Parameters                                 | Default             |
|--------------------------------------------|---------------------|
| CIM server                                 | Disabled            |
| CIM server security protocol               | HTTP                |
| D_S_TOV                                    | 5,000 milliseconds  |
| E_D_TOV                                    | 2,000 milliseconds  |
| R_A_TOV                                    | 10,000 milliseconds |
| Timeout period to invoke fctrace           | 5 seconds           |
| Number of frame sent by the fcping feature | 5 frames            |
| Remote capture connection protocol         | TCP                 |
| Remote capture connection mode             | Passive             |
| Local capture frame limits                 | 10 frames           |
| FC ID allocation mode                      | Auto mode           |
| Loop monitoring                            | Disabled            |
| Interop mode                               | Disabled            |





## Configuring FC-SP and DHCHAP

---

This chapter contains the following sections:

- [Configuring FC-SP and DHCHAP, page 699](#)

### Configuring FC-SP and DHCHAP

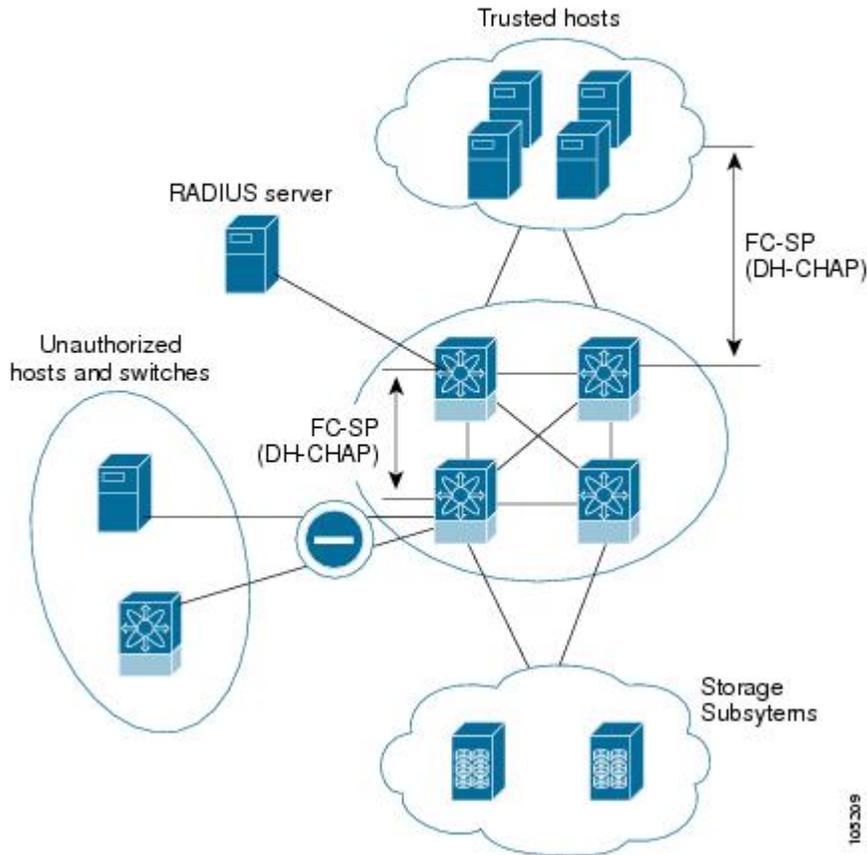
Fibre Channel Security Protocol (FC-SP) capabilities provide switch-to-switch and host-to-switch authentication to overcome security challenges for enterprise-wide fabrics. Diffie-Hellman Challenge Handshake Authentication Protocol (DHCHAP) is an FC-SP protocol that provides authentication between Cisco Nexus 5000 Series switches and other devices. DHCHAP consists of the CHAP protocol combined with the Diffie-Hellman exchange.

#### Information About Fabric Authentication

All Cisco Nexus 5000 Series switches enable fabric-wide authentication from one switch to another switch, or from a switch to a host. These switch and host authentications are performed locally or remotely in each fabric. As storage islands are consolidated and migrated to enterprise-wide fabrics, new security challenges arise. The approach of securing storage islands cannot always be guaranteed in enterprise-wide fabrics. For example, in a campus environment with geographically distributed switches, someone could maliciously interconnect incompatible switches or you could accidentally do so, resulting in Inter-Switch Link (ISL) isolation and link disruption.

Cisco Nexus 5000 Series switches support authentication features to address physical security (see the following figure).

**Figure 72: Switch and Host Authentication**



**Note** Fibre Channel Host Bus Adapters (HBAs) with appropriate firmware and drivers are required for host-switch authentication.

## DHCHAP

DHCHAP is an authentication protocol that authenticates the devices connecting to a switch. Fibre Channel authentication allows only trusted devices to be added to a fabric, which prevents unauthorized devices from accessing the switch.



**Note** The terms FC-SP and DHCHAP are used interchangeably in this chapter.

DHCHAP is a mandatory password-based, key-exchange authentication protocol that supports both switch-to-switch and host-to-switch authentication. DHCHAP negotiates hash algorithms and DH groups before performing authentication. It supports MD5 and SHA-1 algorithm-based authentication.

To configure DHCHAP authentication using the local password database, perform this task:

## SUMMARY STEPS

1. Enable DHCHAP.
2. Identify and configure the DHCHAP authentication modes.
3. Configure the hash algorithm and DH group.
4. Configure the DHCHAP password for the local switch and other switches in the fabric.
5. Configure the DHCHAP timeout value for reauthentication.
6. Verify the DHCHAP configuration.

## DETAILED STEPS

- 
- |               |                                                                                      |
|---------------|--------------------------------------------------------------------------------------|
| <b>Step 1</b> | Enable DHCHAP.                                                                       |
| <b>Step 2</b> | Identify and configure the DHCHAP authentication modes.                              |
| <b>Step 3</b> | Configure the hash algorithm and DH group.                                           |
| <b>Step 4</b> | Configure the DHCHAP password for the local switch and other switches in the fabric. |
| <b>Step 5</b> | Configure the DHCHAP timeout value for reauthentication.                             |
| <b>Step 6</b> | Verify the DHCHAP configuration.                                                     |
- 

### DHCHAP Compatibility with Fibre Channel Features

This section identifies the impact of configuring the DHCHAP feature along with existing Cisco NX-OS features:

- SAN port channel interfaces—If DHCHAP is enabled for ports belonging to a SAN port channel, DHCHAP authentication is performed at the physical interface level, not at the port channel level.
- Port security or fabric binding—Fabric-binding policies are enforced based on identities authenticated by DHCHAP.
- VSANs—DHCHAP authentication is not done on a per-VSAN basis.

### About Enabling DHCHAP

By default, the DHCHAP feature is disabled in all Cisco Nexus 5000 Series switches.

You must explicitly enable the DHCHAP feature to access the configuration and verification commands for fabric authentication. When you disable this feature, all related configurations are automatically discarded.

## Enabling DHCHAP

To enable DHCHAP for a Cisco Nexus 5000 Series switch, perform this task:

### SUMMARY STEPS

1. switch# **configuration terminal**
2. switch(config)# **fcsp enable**
3. switch(config)# **no fcsp enable**

### DETAILED STEPS

|               | Command or Action                     | Purpose                                       |
|---------------|---------------------------------------|-----------------------------------------------|
| <b>Step 1</b> | switch# <b>configuration terminal</b> | Enters configuration mode.                    |
| <b>Step 2</b> | switch(config)# <b>fcsp enable</b>    | Enables the DHCHAP in this switch.            |
| <b>Step 3</b> | switch(config)# <b>no fcsp enable</b> | Disables (default) the DHCHAP in this switch. |

## About DHCHAP Authentication Modes

The DHCHAP authentication status for each interface depends on the configured DHCHAP port mode.

When the DHCHAP feature is enabled in a switch, each Fibre Channel interface or FCIP interface may be configured to be in one of four DHCHAP port modes:

- On—During switch initialization, if the connecting device supports DHCHAP authentication, the software performs the authentication sequence. If the connecting device does not support DHCHAP authentication, the link is placed in an isolated state.
- Auto-Active—During switch initialization, if the connecting device supports DHCHAP authentication, the software performs the authentication sequence. If the connecting device does not support DHCHAP authentication, the software continues with the rest of the initialization sequence.
- Auto-Passive (default)—The switch does not initiate DHCHAP authentication, but participates in DHCHAP authentication if the connecting device initiates DHCHAP authentication.
- Off—The switch does not support DHCHAP authentication. Authentication messages sent to ports in this mode return error messages to the initiating switch.



#### Note

Whenever DHCHAP port mode is changed to a mode other than the Off mode, reauthentication is performed.

The following table identifies switch-to-switch authentication between two Cisco switches in various modes.

**Table 89: DHCHAP Authentication Status Between Two MDS Switches**

| Switch N<br>DHCHAP<br>Modes | Switch 1 DHCHAP Modes              |                                               |                                               |                       |
|-----------------------------|------------------------------------|-----------------------------------------------|-----------------------------------------------|-----------------------|
|                             | on                                 | auto-active                                   | auto-passive                                  | off                   |
| on                          | FC-SP authentication is performed. | FC-SP authentication is performed.            | FC-SP authentication is performed.            | Link is brought down. |
| auto-Active                 |                                    |                                               | FC-SP authentication is <i>not</i> performed. |                       |
| auto-Passive                |                                    |                                               | FC-SP authentication is <i>not</i> performed. |                       |
| off                         | Link is brought down.              | FC-SP authentication is <i>not</i> performed. |                                               |                       |

### Configuring the DHCHAP Mode

To configure the DHCHAP mode for a particular interface, perform this task:

#### SUMMARY STEPS

1. switch# **configuration terminal**
2. switch(config)# **interface fc slot/port - slot/port**
3. switch(config-if)# **fcsp on**
4. switch(config-if)# **no fcsp on**
5. switch(config-if)# **fcsp auto-active 0**
6. switch(config-if)# **fcsp auto-active timeout-period**
7. switch(config-if)# **fcsp auto-active**

#### DETAILED STEPS

|               | Command or Action                                         | Purpose                                                                                                                                                                                                                                                   |
|---------------|-----------------------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Step 1</b> | switch# <b>configuration terminal</b>                     | Enters configuration mode.                                                                                                                                                                                                                                |
| <b>Step 2</b> | switch(config)# <b>interface fc slot/port - slot/port</b> | Selects a range of interfaces and enters the interface configuration mode.                                                                                                                                                                                |
| <b>Step 3</b> | switch(config-if)# <b>fcsp on</b>                         | Sets the DHCHAP mode for the selected interfaces to be in the on state.                                                                                                                                                                                   |
| <b>Step 4</b> | switch(config-if)# <b>no fcsp on</b>                      | Reverts to the factory default of auto-passive for these three interfaces.                                                                                                                                                                                |
| <b>Step 5</b> | switch(config-if)# <b>fcsp auto-active 0</b>              | Changes the DHCHAP authentication mode for the selected interfaces to auto-active. Zero (0) indicates that the port does not perform reauthentication.<br><br><b>Note</b> The reauthorization interval configuration is the same as the default behavior. |

|               | Command or Action                                                   | Purpose                                                                                                                                                                                                                         |
|---------------|---------------------------------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Step 6</b> | switch(config-if)# <b>fcsp auto-active</b><br><i>timeout-period</i> | Changes the DHCHAP authentication mode to auto-active for the selected interfaces. The timeout period value (in minutes) sets how often reauthentication occurs after the initial authentication.                               |
| <b>Step 7</b> | switch(config-if)# <b>fcsp auto-active</b>                          | Changes the DHCHAP authentication mode to auto-active for the selected interfaces. Reauthentication is disabled (default).<br><br><b>Note</b> The reauthorization interval configuration is the same as setting it to zero (0). |

### About the DHCHAP Hash Algorithm

Cisco SAN switches support a default hash algorithm priority list of MD5 followed by SHA-1 for DHCHAP authentication.

If you change the hash algorithm configuration, then change it globally for all switches in the fabric.



#### Caution

RADIUS and TACACS+ protocols always use MD5 for CHAP authentication. Using SHA-1 as the hash algorithm may prevent RADIUS and TACACS+ usage, even if these AAA protocols are enabled for DHCHAP authentication.

### Configuring the DHCHAP Hash Algorithm

To configure the hash algorithm, perform this task:

#### SUMMARY STEPS

1. switch# **configuration terminal**
2. switch(config)# **fcsp dhchap hash [md5] [sha1]**
3. switch(config)# **no fcsp dhchap hash sha1**

#### DETAILED STEPS

|               | Command or Action                                    | Purpose                                                                                                      |
|---------------|------------------------------------------------------|--------------------------------------------------------------------------------------------------------------|
| <b>Step 1</b> | switch# <b>configuration terminal</b>                | Enters configuration mode.                                                                                   |
| <b>Step 2</b> | switch(config)# <b>fcsp dhchap hash [md5] [sha1]</b> | Configures the use of the the MD5 or SHA-1 hash algorithm.                                                   |
| <b>Step 3</b> | switch(config)# <b>no fcsp dhchap hash sha1</b>      | Reverts to the factory default priority list of the MD5 hash algorithm followed by the SHA-1 hash algorithm. |

### About the DHCHAP Group Settings

All Cisco Nexus 5000 Series switches support all DHCHAP groups specified in the standard: 0 (null DH group, which does not perform the Diffie-Hellman exchange), 1, 2, 3, or 4.

If you change the DH group configuration, change it globally for all switches in the fabric.

### Configuring the DHCHAP Group Settings

To change the DH group settings, perform this task:

#### SUMMARY STEPS

1. switch# **configuration terminal**
2. switch(config)# **fcsp dhchap dhgroup [0 | 1 | 2 | 3 | 4]**
3. switch(config)# **no fcsp dhchap dhgroup [0 | 1 | 2 | 3 | 4]**

#### DETAILED STEPS

|               | Command or Action                                                 | Purpose                                                           |
|---------------|-------------------------------------------------------------------|-------------------------------------------------------------------|
| <b>Step 1</b> | switch# <b>configuration terminal</b>                             | Enters configuration mode.                                        |
| <b>Step 2</b> | switch(config)# <b>fcsp dhchap dhgroup [0   1   2   3   4]</b>    | Prioritizes the use of DH groups in the configured order.         |
| <b>Step 3</b> | switch(config)# <b>no fcsp dhchap dhgroup [0   1   2   3   4]</b> | Reverts to the DHCHAP factory default order of 0, 4, 1, 2, and 3. |

### About the DHCHAP Password

DHCHAP authentication in each direction requires a shared secret password between the connected devices. To do this, you can use one of three configurations to manage passwords for all switches in the fabric that participate in DHCHAP:

- Configuration 1—Use the same password for all switches in the fabric. This is the simplest configuration. When you add a new switch, you use the same password to authenticate that switch in this fabric. It is also the most vulnerable configuration if someone from the outside maliciously attempts to access any one switch in the fabric.
- Configuration 2—Use a different password for each switch and maintain that password list in each switch in the fabric. When you add a new switch, you create a new password list and update all switches with the new list. Accessing one switch yields the password list for all switches in that fabric.
- Configuration 3—Use different passwords for different switches in the fabric. When you add a new switch, multiple new passwords corresponding to each switch in the fabric must be generated and configured in each switch. Even if one switch is compromised, the password of other switches are still protected. This configuration requires considerable password maintenance by the user.

**Note**

All passwords are restricted to 64 alphanumeric characters and can be changed, but not deleted.

We recommend using RADIUS or TACACS+ for fabrics with more than five switches. If you need to use a local password database, you can continue to do so using Configuration 3 and using the Cisco MDS 9000 Family Fabric Manager to manage the password database.

**Configuring DHCHAP Passwords for the Local Switch**

To configure the DHCHAP password for the local switch, perform this task:

**SUMMARY STEPS**

1. switch# **configuration terminal**
2. switch(config)# **fcsp dhchap password** [0 | 7] *password* [wwn *wwn-id*]

**DETAILED STEPS**

|               | <b>Command or Action</b>                                                                    | <b>Purpose</b>                                         |
|---------------|---------------------------------------------------------------------------------------------|--------------------------------------------------------|
| <b>Step 1</b> | switch# <b>configuration terminal</b>                                                       | Enters configuration mode.                             |
| <b>Step 2</b> | switch(config)# <b>fcsp dhchap password</b> [0   7] <i>password</i><br>[wwn <i>wwn-id</i> ] | Configures a clear text password for the local switch. |

**About Password Configuration for Remote Devices**

You can configure passwords in the local authentication database for other devices in a fabric. The other devices are identified by their device name, which is also known as the switch WWN or device WWN. The password is restricted to 64 characters and can be specified in clear text (0) or in encrypted text (7).

**Note**

The switch WWN identifies the physical switch. This WWN is used to authenticate the switch and is different from the VSAN node WWN.

**Configuring DHCHAP Passwords for Remote Devices**

To locally configure the remote DHCHAP password for another switch in the fabric, perform this task:

**SUMMARY STEPS**

1. switch# **configuration terminal**
2. switch(config)# **fcsp dhchap devicename** *switch-wwn* **password** *password*
3. switch(config)# **no fcsp dhchap devicename** *switch-wwn* **password** *password*

**DETAILED STEPS**

|               | <b>Command or Action</b>                                                                              | <b>Purpose</b>                                                                                           |
|---------------|-------------------------------------------------------------------------------------------------------|----------------------------------------------------------------------------------------------------------|
| <b>Step 1</b> | switch# <b>configuration terminal</b>                                                                 | Enters configuration mode.                                                                               |
| <b>Step 2</b> | switch(config)# <b>fcsp dhchap devicename</b> <i>switch-wwn</i><br><b>password</b> <i>password</i>    | Configures a password for another switch in the fabric that is identified by the switch WWN device name. |
| <b>Step 3</b> | switch(config)# <b>no fcsp dhchap devicename</b><br><i>switch-wwn</i> <b>password</b> <i>password</i> | Removes the password entry for this switch from the local authentication database.                       |

**About the DHCHAP Timeout Value**

During the DHCHAP protocol exchange, if the Cisco Nexus 5000 Series switch does not receive the expected DHCHAP message within a specified time interval, authentication failure is assumed. The time ranges from 20 (no authentication is performed) to 1000 seconds. The default is 30 seconds.

When changing the timeout value, consider the following factors:

- The existing RADIUS and TACACS+ timeout values.
- The same value must also be configured on all switches in the fabric.

**Configuring the DHCHAP Timeout Value**

To configure the DHCHAP timeout value, perform this task:

**SUMMARY STEPS**

1. switch# **configuration terminal**
2. switch(config)# **fcsp timeout** *timeout*
3. switch(config)# **no fcsp timeout** *timeout*

**DETAILED STEPS**

|               | <b>Command or Action</b>                              | <b>Purpose</b>                                                                       |
|---------------|-------------------------------------------------------|--------------------------------------------------------------------------------------|
| <b>Step 1</b> | switch# <b>configuration terminal</b>                 | Enters configuration mode.                                                           |
| <b>Step 2</b> | switch(config)# <b>fcsp timeout</b> <i>timeout</i>    | Configures the reauthentication timeout to the specified value. The unit is seconds. |
| <b>Step 3</b> | switch(config)# <b>no fcsp timeout</b> <i>timeout</i> | Reverts to the factory default of 30 seconds.                                        |

## Configuring DHCHAP AAA Authentication

You can configure AAA authentication to use a RADIUS or TACACS+ server group. If AAA authentication is not configured, local authentication is used by default.

## Displaying Protocol Security Information

Use the **show fcsp** commands to display configurations for the local database.

The following example shows how to display the DHCHAP configuration for the specified interface:

```
switch# show fcsp interface fc2/4
fc2/4:
 fcsp authentication mode:SEC_MODE_ON
 Status: Successfully authenticated
```

The following example shows how to display DHCHAP statistics for the specified interface:

```
switch# show fcsp interface fc2/4 statistics
```

The following example shows how to display the FC-SP WWN of the device connected to the specified interface:

```
switch# show fcsp interface fc2/1 wwn
```

The following example shows how to display the hash algorithm and DHCHAP groups configured in the switch:

```
switch# show fcsp dhchap
```

The following example shows how to display the DHCHAP local password database:

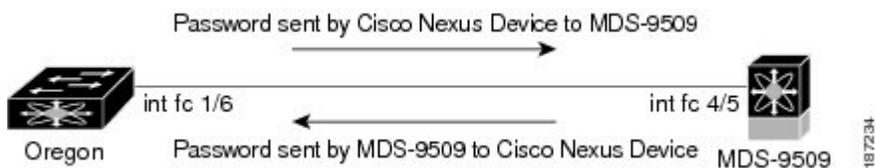
```
switch# show fcsp dhchap database
```

Use the ASCII representation of the device WWN to configure the switch information on RADIUS and TACACS+ servers.

## Sample Configuration

This section provides the steps to configure the example illustrated in the following figure.

**Figure 73: Sample DHCHAP Authentication**



To configure the authentication setup shown in the above figure, perform this task:

## SUMMARY STEPS

1. Obtain the device name of the Cisco Nexus 5000 Series switch in the fabric. The Cisco Nexus 5000 Series switch in the fabric is identified by the switch WWN.
2. Explicitly enable DHCHAP in this switch.
3. Configure a clear text password for this switch. This password will be used by the connecting device.
4. Configure a password for another switch in the fabric that is identified by the switch WWN device name.
5. Enable the DHCHAP mode for the required Fibre Channel interface.
6. Verify the protocol security information configured in this switch by displaying the DHCHAP local password database.
7. Display the DHCHAP configuration in the Fibre Channel interface.
8. Repeat these steps on the connecting MDS 9509 switch.

## DETAILED STEPS

**Step 1** Obtain the device name of the Cisco Nexus 5000 Series switch in the fabric. The Cisco Nexus 5000 Series switch in the fabric is identified by the switch WWN.

**Example:**

```
switch# show wwn switch
Switch WWN is 20:00:00:05:30:00:54:de
```

**Step 2** Explicitly enable DHCHAP in this switch.

**Note** When you disable DHCHAP, all related configurations are automatically discarded.

**Example:**

```
switch(config)# fcsp enable
```

**Step 3** Configure a clear text password for this switch. This password will be used by the connecting device.

**Example:**

```
switch(config)# fcsp dhchap password rtp9216
```

**Step 4** Configure a password for another switch in the fabric that is identified by the switch WWN device name.

**Example:**

```
switch(config)# fcsp dhchap devicename 20:00:00:05:30:00:38:5e password rtp9509
```

**Step 5** Enable the DHCHAP mode for the required Fibre Channel interface.

**Note** Whenever DHCHAP port mode is changed to a mode other than the Off mode, reauthentication is performed.

**Example:**

```
switch(config)# interface fc2/4
switch(config-if)# fcsp on
```

**Step 6** Verify the protocol security information configured in this switch by displaying the DHCHAP local password database.

**Example:**

```
switch# show fcsp dhchap database
DHCHAP Local Password:
 Non-device specific password:*****
Other Devices' Passwords:
 Password for device with WWN:20:00:00:05:30:00:38:5e is *****
```

**Step 7** Display the DHCHAP configuration in the Fibre Channel interface.

**Example:**

```
switch# show fcsp interface fc2/4
fc2/4
 fcsp authentication mode:SEC_MODE_ON
 Status:Successfully authenticated
```

**Step 8** Repeat these steps on the connecting MDS 9509 switch.

**Example:**

```
MDS-9509# show wwn switch
Switch WWN is 20:00:00:05:30:00:38:5e
MDS-9509(config)# fcsp enable
MDS-9509(config)# fcsp dhchap password rtp9509
MDS-9509(config)# fcsp dhchap devicename 20:00:00:05:30:00:54:de password rtp9216
MDS-9509(config)# interface fc 4/5
MDS-9509(config-if)# fcsp on
MDS-9509# show fcsp dhchap database
DHCHAP Local Password:
 Non-device specific password:*****
Other Devices' Passwords:
 Password for device with WWN:20:00:00:05:30:00:54:de is *****
MDS-9509# show fcsp interface fc2/4
Fc2/4
 fcsp authentication mode:SEC_MODE_ON
 Status:Successfully authenticated
```

You have now enabled and configured DHCHAP authentication for the sample setup in shown in the figure above.

## Default Fabric Security Settings

The following table lists the default settings for all fabric security features in any switch.

**Table 90: Default Fabric Security Settings**

| Parameters                                   | Default                                                            |
|----------------------------------------------|--------------------------------------------------------------------|
| DHCHAP feature                               | Disabled                                                           |
| DHCHAP hash algorithm                        | A priority list of MD5 followed by SHA-1 for DHCHAP authentication |
| DHCHAP authentication mode                   | Auto-passive                                                       |
| DHCHAP group default priority exchange order | 0, 4, 1, 2, and 3, respectively                                    |
| DHCHAP timeout value                         | 30 seconds                                                         |







# Configuring Port Security

This chapter contains the following sections:

- [Configuring Port Security, page 713](#)

## Configuring Port Security

Cisco Nexus 5000 Series switches provide port security features that reject intrusion attempts and report these intrusions to the administrator.



**Note**

---

Port security is supported on virtual Fibre Channel ports and physical Fibre Channel ports.

---

### Information About Port Security

Typically, any Fibre Channel device in a SAN can attach to any SAN switch port and access SAN services based on zone membership. Port security features prevent unauthorized access to a switch port in the Cisco Nexus 5000 Series switch, using the following methods:

- Login requests from unauthorized Fibre Channel devices (N ports) and switches (xE ports) are rejected.
- All intrusion attempts are reported to the SAN administrator through system messages.
- Configuration distribution uses the CFS infrastructure, and is limited to those switches that are CFS capable. Distribution is disabled by default.
- Configuring the port security policy requires the Storage Protocol Services license.

### Port Security Enforcement

To enforce port security, configure the devices and switch port interfaces through which each device or switch is connected, and activate the configuration.

- Use the port world wide name (pWWN) or the node world wide name (nWWN) to specify the N port connection for each device.

- Use the switch world wide name (sWWN) to specify the xE port connection for each switch.

Each N and xE port can be configured to restrict a single port or a range of ports.

Enforcement of port security policies are done on every activation and when the port tries to come up.

The port security feature uses two databases to accept and implement configuration changes.

- Configuration database—All configuration changes are stored in the configuration database.
- Active database—The database currently enforced by the fabric. The port security feature requires all devices connecting to a switch to be part of the port security active database. The software uses this active database to enforce authorization.

### About Auto-Learning

You can instruct the switch to automatically learn (auto-learn) the port security configurations over a specified period. This feature allows any Cisco Nexus 5000 Series switch to automatically learn about devices and switches that connect to it. Use this feature when you activate the port security feature for the first time as it saves tedious manual configuration for each port. You must configure auto-learning on a per-VSAN basis. If enabled, devices and switches that are allowed to connect to the switch are automatically learned, even if you have not configured any port access.

When auto-learning is enabled, learning happens only for the devices or interfaces that were not already logged into the switch. Learned entries on a port are cleaned up after you shut down that port if auto-learning is still enabled.

Learning does not override the existing configured port security policies. For example, if an interface is configured to allow a specific pWWN, then auto-learning will not add a new entry to allow any other pWWN on that interface. All other pWWNs will be blocked even in auto-learning mode.

No entries are learned for a port in the shutdown state.

When you activate the port security feature, auto-learning is also automatically enabled.



#### Note

---

If you enable auto-learning before activating port security, you cannot activate port security until auto-learning is disabled.

---

### Port Security Activation

By default, the port security feature is not activated in Cisco Nexus 5000 Series switches.

When you activate the port security feature, the following operations occur:

- Auto-learning is also automatically enabled, which means:
  - From this point, auto-learning happens only for the devices or interfaces that were not logged into the switch.
  - You cannot activate the database until you disable auto-learning.
- All the devices that are already logged in are learned and are added to the active database.
- All entries in the configured database are copied to the active database.

After the database is activated, subsequent device login is subject to the activated port bound WWN pairs, excluding the auto-learned entries. You must disable auto-learning before the auto-learned entries become activated.

When you activate the port security feature, auto-learning is also automatically enabled. You can choose to activate the port security feature and disable auto-learning.

If a port is shut down because of a denied login attempt, and you subsequently configure the database to allow that login, the port does not come up automatically. You must explicitly enter a **no shutdown** CLI command to bring that port back online.

## Configuring Port Security

### Configuring Port Security with Auto-Learning and CFS Distribution

To configure port security, using auto-learning and CFS distribution, perform this task:

#### SUMMARY STEPS

1. Enable port security.
2. Enable CFS distribution.
3. Activate port security on each VSAN.
4. Issue a CFS commit to copy this configuration to all switches in the fabric.
5. Wait until all switches and all hosts are automatically learned.
6. Disable auto-learn on each VSAN.
7. Issue a CFS commit to copy this configuration to all switches in the fabric.
8. Copy the active database to the configure database on each VSAN.
9. Issue a CFS commit to copy this configuration to all switches in the fabric.
10. Copy the running configuration to the startup configuration, using the fabric option.

#### DETAILED STEPS

- 
- |               |                                                                                                                                                                                                            |
|---------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Step 1</b> | Enable port security.                                                                                                                                                                                      |
| <b>Step 2</b> | Enable CFS distribution.                                                                                                                                                                                   |
| <b>Step 3</b> | Activate port security on each VSAN.<br>This action turns on auto-learning by default.                                                                                                                     |
| <b>Step 4</b> | Issue a CFS commit to copy this configuration to all switches in the fabric.<br>All switches have port security activated with auto-learning enabled.                                                      |
| <b>Step 5</b> | Wait until all switches and all hosts are automatically learned.                                                                                                                                           |
| <b>Step 6</b> | Disable auto-learn on each VSAN.                                                                                                                                                                           |
| <b>Step 7</b> | Issue a CFS commit to copy this configuration to all switches in the fabric.<br>The auto-learned entries from every switch are combined into a static active database that is distributed to all switches. |
| <b>Step 8</b> | Copy the active database to the configure database on each VSAN.                                                                                                                                           |
| <b>Step 9</b> | Issue a CFS commit to copy this configuration to all switches in the fabric.                                                                                                                               |

This ensures that the configure database is the same on all switches in the fabric.

**Step 10** Copy the running configuration to the startup configuration, using the fabric option.

---

### Related Topics

- [Activating Port Security, on page 718](#)
- [Committing the Changes, on page 727](#)
- [Copying the Port Security Database, on page 732](#)
- [Disabling Auto-Learning, on page 721](#)
- [Enabling Port Security, on page 717](#)
- [Enabling Port Security Distribution, on page 726](#)

## Configuring Port Security with Auto-Learning without CFS

To configure port security using auto-learning without CFS, perform this task:

### SUMMARY STEPS

1. Enable port security.
2. Activate port security on each VSAN, which turns on auto-learning by default.
3. Wait until all switches and all hosts are automatically learned.
4. Disable auto-learn on each VSAN.
5. Copy the active database to the configure database on each VSAN.
6. Copy the running configuration to the startup configuration, which saves the port security configuration database to the startup configuration.
7. Repeat the above steps for all switches in the fabric.

### DETAILED STEPS

- 
- |               |                                                                                                                                                 |
|---------------|-------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Step 1</b> | Enable port security.                                                                                                                           |
| <b>Step 2</b> | Activate port security on each VSAN, which turns on auto-learning by default.                                                                   |
| <b>Step 3</b> | Wait until all switches and all hosts are automatically learned.                                                                                |
| <b>Step 4</b> | Disable auto-learn on each VSAN.                                                                                                                |
| <b>Step 5</b> | Copy the active database to the configure database on each VSAN.                                                                                |
| <b>Step 6</b> | Copy the running configuration to the startup configuration, which saves the port security configuration database to the startup configuration. |
| <b>Step 7</b> | Repeat the above steps for all switches in the fabric.                                                                                          |
-

### Related Topics

- [Activating Port Security, on page 718](#)
- [Copying the Port Security Database, on page 732](#)
- [Disabling Auto-Learning, on page 721](#)
- [Enabling Port Security, on page 717](#)

## Configuring Port Security with Manual Database Configuration

To configure port security and manually configure the port security database, perform this task:

### SUMMARY STEPS

1. Enable port security.
2. Manually configure all port security entries into the configure database on each VSAN.
3. Activate port security on each VSAN. This turns on auto-learning by default.
4. Disable auto-learn on each VSAN.
5. Copy the running configuration to the startup configuration, which saves the port security configuration database to the startup configuration.
6. Repeat the above steps for all switches in the fabric.

### DETAILED STEPS

- 
- |               |                                                                                                                                                 |
|---------------|-------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Step 1</b> | Enable port security.                                                                                                                           |
| <b>Step 2</b> | Manually configure all port security entries into the configure database on each VSAN.                                                          |
| <b>Step 3</b> | Activate port security on each VSAN. This turns on auto-learning by default.                                                                    |
| <b>Step 4</b> | Disable auto-learn on each VSAN.                                                                                                                |
| <b>Step 5</b> | Copy the running configuration to the startup configuration, which saves the port security configuration database to the startup configuration. |
| <b>Step 6</b> | Repeat the above steps for all switches in the fabric.                                                                                          |
- 

## Enabling Port Security

By default, the port security feature is disabled in Cisco Nexus 5000 Series switches.

To enable port security, perform this task:

### SUMMARY STEPS

1. switch# **configuration terminal**
2. switch(config)# **port-security enable**
3. switch(config)# **no port-security enable**

**DETAILED STEPS**

|               | <b>Command or Action</b>                       | <b>Purpose</b>                                   |
|---------------|------------------------------------------------|--------------------------------------------------|
| <b>Step 1</b> | switch# <b>configuration terminal</b>          | Enters configuration mode.                       |
| <b>Step 2</b> | switch(config)# <b>port-security enable</b>    | Enables port security on that switch.            |
| <b>Step 3</b> | switch(config)# <b>no port-security enable</b> | Disables (default) port security on that switch. |

**Port Security Activation****Activating Port Security**

To activate port security, perform this task:

**SUMMARY STEPS**

1. switch# **configuration terminal**
2. switch(config)# **port-security activate vsan vsan-id**
3. switch(config)# **port-security activate vsan vsan-id no-auto-learn**
4. switch(config)# **no port-security activate vsan vsan-id**

**DETAILED STEPS**

|               | <b>Command or Action</b>                                                 | <b>Purpose</b>                                                                                           |
|---------------|--------------------------------------------------------------------------|----------------------------------------------------------------------------------------------------------|
| <b>Step 1</b> | switch# <b>configuration terminal</b>                                    | Enters configuration mode.                                                                               |
| <b>Step 2</b> | switch(config)# <b>port-security activate vsan vsan-id</b>               | Activates the port security database for the specified VSAN, and automatically enables auto-learning.    |
| <b>Step 3</b> | switch(config)# <b>port-security activate vsan vsan-id no-auto-learn</b> | Activates the port security database for the specified VSAN, and disables auto-learning.                 |
| <b>Step 4</b> | switch(config)# <b>no port-security activate vsan vsan-id</b>            | Deactivates the port security database for the specified VSAN, and automatically disables auto-learning. |

**Database Activation Rejection**

Database activation is rejected in the following cases:

- Missing or conflicting entries exist in the configuration database but not in the active database.

- The auto-learning feature was enabled before the activation. To reactivate a database in this state, disable auto-learning.
- The exact security is not configured for each port channel member.
- The configured database is empty but the active database is not.

If the database activation is rejected due to one or more conflicts listed in the previous section, you may decide to proceed by forcing the port security activation.

### Forcing Port Security Activation

If the port security activation request is rejected, you can force the activation.



#### Note

If you force the activation, existing devices are logged out if they violate the active database.

You can view missing or conflicting entries using the **port-security database diff active vsan** command in EXEC mode.

To forcefully activate the port security database, perform this task:

### SUMMARY STEPS

1. switch# **configuration terminal**
2. switch(config)# **port-security activate vsan vsan-id force**

### DETAILED STEPS

|               | Command or Action                                                | Purpose                                                                                       |
|---------------|------------------------------------------------------------------|-----------------------------------------------------------------------------------------------|
| <b>Step 1</b> | switch# <b>configuration terminal</b>                            | Enters configuration mode.                                                                    |
| <b>Step 2</b> | switch(config)# <b>port-security activate vsan vsan-id force</b> | Forces the port security database to activate for the specified VSAN even if conflicts occur. |

### Database Reactivation



#### Tip

If auto-learning is enabled, you cannot activate the database without the force option until you disable auto-learning.

To reactivate the port security database, perform this task:

## SUMMARY STEPS

1. switch# **configuration terminal**
2. switch(config)# **no port-security auto-learn vsan** *vsan-id*
3. switch(config)# **exit**
4. switch# **port-security database copy vsan** *vsan-id*
5. switch# **configuration terminal**
6. switch(config)# **port-security activate vsan** *vsan-id*

## DETAILED STEPS

|               | Command or Action                                                      | Purpose                                                                                                                                                                         |
|---------------|------------------------------------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Step 1</b> | switch# <b>configuration terminal</b>                                  | Enters configuration mode.                                                                                                                                                      |
| <b>Step 2</b> | switch(config)# <b>no port-security auto-learn vsan</b> <i>vsan-id</i> | Disables auto-learning and stops the switch from learning about new devices accessing the switch. Enforces the database contents based on the devices learned up to this point. |
| <b>Step 3</b> | switch(config)# <b>exit</b>                                            |                                                                                                                                                                                 |
| <b>Step 4</b> | switch# <b>port-security database copy vsan</b> <i>vsan-id</i>         | Copies from the active to the configured database.                                                                                                                              |
| <b>Step 5</b> | switch# <b>configuration terminal</b>                                  | Re-enters configuration mode.                                                                                                                                                   |
| <b>Step 6</b> | switch(config)# <b>port-security activate vsan</b> <i>vsan-id</i>      | Activates the port security database for the specified VSAN, and automatically enables auto-learning.                                                                           |

## Auto-Learning

### About Enabling Auto-Learning

The state of the auto-learning configuration depends on the state of the port security feature:

- If the port security feature is not activated, auto-learning is disabled by default.
- If the port security feature is activated, auto-learning is enabled by default (unless you explicitly disabled this option).



#### Tip

If auto-learning is enabled on a VSAN, you can only activate the database for that VSAN by using the force option.

### Enabling Auto-Learning

To enable auto-learning, perform this task:

**SUMMARY STEPS**

1. switch# **configuration terminal**
2. switch(config)# **port-security auto-learn vsan** *vsan-id*

**DETAILED STEPS**

|               | <b>Command or Action</b>                                            | <b>Purpose</b>                                                                                                                                                  |
|---------------|---------------------------------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Step 1</b> | switch# <b>configuration terminal</b>                               | Enters configuration mode.                                                                                                                                      |
| <b>Step 2</b> | switch(config)# <b>port-security auto-learn vsan</b> <i>vsan-id</i> | Enables auto-learning so the switch can learn about any device that is allowed to access VSAN 1. These devices are logged in the port security active database. |

**Disabling Auto-Learning**

To disable auto-learning, perform this task:

**SUMMARY STEPS**

1. switch# **configuration terminal**
2. switch(config)# **no port-security auto-learn vsan** *vsan-id*

**DETAILED STEPS**

|               | <b>Command or Action</b>                                               | <b>Purpose</b>                                                                                                                                                                  |
|---------------|------------------------------------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Step 1</b> | switch# <b>configuration terminal</b>                                  | Enters configuration mode.                                                                                                                                                      |
| <b>Step 2</b> | switch(config)# <b>no port-security auto-learn vsan</b> <i>vsan-id</i> | Disables auto-learning and stops the switch from learning about new devices accessing the switch. Enforces the database contents based on the devices learned up to this point. |

**Auto-Learning Device Authorization**

The following table summarizes the authorized connection conditions for device requests.

**Table 91: Authorized Auto-Learning Device Requests**

| Condition | Device (pWWN, nWWN, sWWN)                | Requests Connection to                   | Authorization                      |
|-----------|------------------------------------------|------------------------------------------|------------------------------------|
| 1         | Configured with one or more switch ports | A configured switch port                 | Permitted                          |
| 2         |                                          | Any other switch port                    | Denied                             |
| 3         | Not configured                           | A switch port that is not configured     | Permitted if auto-learning enabled |
| 4         |                                          |                                          | Denied if auto-learning disabled   |
| 5         | Configured or not configured             | A switch port that allows any device     | Permitted                          |
| 6         | Configured to log in to any switch port  | Any port on the switch                   | Permitted                          |
| 7         | Not configured                           | A port configured with some other device | Denied                             |

**Authorization Scenario**

Assume that the port security feature is activated and the following conditions are specified in the active database:

- A pWWN (P1) is allowed access through interface fc2/1 (F1).
- A pWWN (P2) is allowed access through interface fc2/2 (F1).
- A nWWN (N1) is allowed access through interface fc2/2 (F2).
- Any WWN is allowed access through interface vfc3/1 (F3).
- A nWWN (N3) is allowed access through any interface.
- A pWWN (P3) is allowed access through interface fc2/4 (F4).
- A sWWN (S1) is allowed access through interface fc3/1-3 (F10 to F13).
- A pWWN (P10) is allowed access through interface vfc4/1 (F11).

The following table summarizes the port security authorization results for this active database.

**Table 92: Authorization Results for Scenario**

| Device Connection Request | Authorization | Condition | Reason       |
|---------------------------|---------------|-----------|--------------|
| P1, N2, F1                | Permitted     | 1         | No conflict. |

| Device Connection Request      | Authorization | Condition | Reason                              |
|--------------------------------|---------------|-----------|-------------------------------------|
| P2, N2, F1                     | Permitted     | 1         | No conflict.                        |
| P3, N2, F1                     | Denied        | 2         | F1 is bound to P1/P2.               |
| P1, N3, F1                     | Permitted     | 6         | Wildcard match for N3.              |
| P1, N1, F3                     | Permitted     | 5         | Wildcard match for F3.              |
| P1, N4, F5                     | Denied        | 2         | P1 is bound to F1.                  |
| P5, N1, F5                     | Denied        | 2         | N1 is only allowed on F2.           |
| P3, N3, F4                     | Permitted     | 1         | No conflict.                        |
| S1, F10                        | Permitted     | 1         | No conflict.                        |
| S2, F11                        | Denied        | 7         | P10 is bound to F11.                |
| P4, N4, F5 (auto-learning on)  | Permitted     | 3         | No conflict.                        |
| P4, N4, F5 (auto-learning off) | Denied        | 4         | No match.                           |
| S3, F5 (auto-learning on)      | Permitted     | 3         | No conflict.                        |
| S3, F5 (auto-learning off)     | Denied        | 4         | No match.                           |
| P1, N1, F6 (auto-learning on)  | Denied        | 2         | P1 is bound to F1.                  |
| P5, N5, F1 (auto-learning on)  | Denied        | 7         | Only P1 and P2 bound to F1.         |
| S3, F4 (auto-learning on)      | Denied        | 7         | P3 paired with F4.                  |
| S1, F3 (auto-learning on)      | Permitted     | 5         | No conflict.                        |
| P5, N3, F3                     | Permitted     | 6         | Wildcard ( * ) match for F3 and N3. |
| P7, N3, F9                     | Permitted     | 6         | Wildcard ( * ) match for N3.        |

### Related Topics

[Auto-Learning Device Authorization, on page 721](#)

## Port Security Manual Configuration

To configure port security on a Cisco Nexus 5000 Series switch, perform this task:

### SUMMARY STEPS

1. Identify the WWN of the ports that need to be secured.
2. Secure the fWWN to an authorized nWWN or pWWN.
3. Activate the port security database.
4. Verify your configuration.

### DETAILED STEPS

- 
- |               |                                                        |
|---------------|--------------------------------------------------------|
| <b>Step 1</b> | Identify the WWN of the ports that need to be secured. |
| <b>Step 2</b> | Secure the fWWN to an authorized nWWN or pWWN.         |
| <b>Step 3</b> | Activate the port security database.                   |
| <b>Step 4</b> | Verify your configuration.                             |
- 

### WWN Identification Guidelines

If you decide to manually configure port security, note the following guidelines:

- Identify switch ports by the interface or by the fWWN.
- Identify devices by the pWWN or by the nWWN.
- If an N port is allowed to log in to SAN switch port F, then that N port can only log in through the specified F port.
- If an N port's nWWN is bound to an F port WWN, then all pWWNs in the N port are implicitly paired with the F port.
- TE port checking is done on each VSAN in the allowed VSAN list of the VSAN trunk port.
- All port channel xE ports must be configured with the same set of WWNs in the same SAN port channel.
- E port security is implemented in the port VSAN of the E port. In this case, the sWWN is used to secure authorization checks.
- Once activated, the configuration database can be modified without any effect on the active database.
- By saving the running configuration, you save the configuration database and activated entries in the active database. Learned entries in the active database are not saved.

### Adding Authorized Port Pairs

After identifying the WWN pairs that need to be bound, add those pairs to the port security database.

**Tip**

Remote switch binding can be specified at the local switch. To specify the remote interfaces, you can use either the fWWN or sWWN-interface combination.

To add authorized port pairs for port security, perform this task:

**SUMMARY STEPS**

1. switch# **configuration terminal**
2. switch(config)# **port-security database vsan vsan-id**
3. switch(config)# **no port-security database vsan vsan-id**
4. switch(config-port-security)# **swwn swwn-id interface san-port-channel 5**
5. switch(config-port-security)# **any-wwn interface fc slot/port - fc slot/port**

**DETAILED STEPS**

|               | <b>Command or Action</b>                                                           | <b>Purpose</b>                                                            |
|---------------|------------------------------------------------------------------------------------|---------------------------------------------------------------------------|
| <b>Step 1</b> | switch# <b>configuration terminal</b>                                              | Enters configuration mode.                                                |
| <b>Step 2</b> | switch(config)# <b>port-security database vsan vsan-id</b>                         | Enters the port security database mode for the specified VSAN.            |
| <b>Step 3</b> | switch(config)# <b>no port-security database vsan vsan-id</b>                      | Deletes the port security configuration database from the specified VSAN. |
| <b>Step 4</b> | switch(config-port-security)# <b>swwn swwn-id interface san-port-channel 5</b>     | Configures the specified sWWN to only log in through SAN port channel 5.  |
| <b>Step 5</b> | switch(config-port-security)# <b>any-wwn interface fc slot/port - fc slot/port</b> | Configures any WWN to log in through the specified interfaces.            |

This example enters the port security database mode for VSAN 2:

```
switch(config)# port-security database vsan 2
```

This example configures the specified sWWN to only log in through SAN port channel 5:

```
switch(config-port-security)# swwn 20:01:33:11:00:2a:4a:66 interface san-port-channel 5
```

This example configures the specified pWWN to log in through the specified interface in the specified switch:

```
switch(config-port-security)# pwwn 20:11:33:11:00:2a:4a:66 swwn 20:00:00:0c:85:90:3e:80
interface fc 3/2
```

This example configures any WWN to log in through the specified interface in any switch:

```
switch(config-port-security)# any-wwn interface fc 3/2
```

**Port Security Configuration Distribution**

The port security feature uses the Cisco Fabric Services (CFS) infrastructure to enable efficient database management, provide a single point of configuration for the entire fabric in the VSAN, and enforce the port security policies throughout the fabric.

**Related Topics**

[Using Cisco Fabric Services, on page 329](#)

**Enabling Port Security Distribution**

All the configurations performed in distributed mode are stored in a pending (temporary) database. If you modify the configuration, you need to commit or discard the pending database changes to the configurations. The fabric remains locked during this period. Changes to the pending database are not reflected in the configurations until you commit the changes.

**Note**

Port activation or deactivation and auto-learning enable or disable do not take effect until after a CFS commit if CFS distribution is enabled. Always follow any one of these operations with a CFS commit to ensure proper configuration.

For example, if you activate port security, follow up by disabling auto-learning, and finally commit the changes in the pending database, then the net result of your actions is the same as entering a **port-security activate vsan vsan-id no-auto-learn** command.

**Tip**

We recommend that you perform a commit after you activate port security and after you enable auto learning.

To enable the port security distribution, perform this task:

**SUMMARY STEPS**

1. switch# **configuration terminal**
2. switch(config)# **port-security distribute**
3. switch(config)# **no port-security distribute**

**DETAILED STEPS**

|               | <b>Command or Action</b>                           | <b>Purpose</b>             |
|---------------|----------------------------------------------------|----------------------------|
| <b>Step 1</b> | switch# <b>configuration terminal</b>              | Enters configuration mode. |
| <b>Step 2</b> | switch(config)# <b>port-security distribute</b>    | Enables distribution.      |
| <b>Step 3</b> | switch(config)# <b>no port-security distribute</b> | Disables distribution.     |

**Related Topics**

[Activation and Auto-Learning Configuration Distribution, on page 728](#)

## Locking the Fabric

The first action that modifies the existing configuration creates the pending database and locks the feature in the VSAN. Once you lock the fabric, the following situations apply:

- No other user can make any configuration changes to this feature.
- A copy of the configuration database becomes the pending database.

## Committing the Changes

If you commit the changes made to the configurations, the configurations in the pending database are distributed to other switches. On a successful commit, the configuration change is applied throughout the fabric and the lock is released.

To commit the port security configuration changes for the specified VSAN, perform this task:

### SUMMARY STEPS

1. switch# **configuration terminal**
2. switch(config)# **port-security commit vsan** *vsan-id*

### DETAILED STEPS

|               | Command or Action                                               | Purpose                                                  |
|---------------|-----------------------------------------------------------------|----------------------------------------------------------|
| <b>Step 1</b> | switch# <b>configuration terminal</b>                           | Enters configuration mode.                               |
| <b>Step 2</b> | switch(config)# <b>port-security commit vsan</b> <i>vsan-id</i> | Commits the port security changes in the specified VSAN. |

## Discarding the Changes

If you discard (abort) the changes made to the pending database, the configuration remains unaffected and the lock is released.

To discard the port security configuration changes for the specified VSAN, perform this task:

### SUMMARY STEPS

1. switch# **configuration terminal**
2. switch(config)# **port-security abort vsan** *vsan-id*

**DETAILED STEPS**

|               | <b>Command or Action</b>                                | <b>Purpose</b>                                                                                          |
|---------------|---------------------------------------------------------|---------------------------------------------------------------------------------------------------------|
| <b>Step 1</b> | switch# <b>configuration terminal</b>                   | Enters configuration mode.                                                                              |
| <b>Step 2</b> | switch(config)# <b>port-security abort vsan vsan-id</b> | Discards the port security changes in the specified VSAN and clears the pending configuration database. |

**Activation and Auto-Learning Configuration Distribution**

Activation and auto-learning configurations in distributed mode are remembered as actions to be performed when you commit the changes in the pending database.

Learned entries are temporary and do not have any role in determining if a login is authorized or not. As such, learned entries do not participate in distribution. When you disable learning and commit the changes in the pending database, the learned entries become static entries in the active database and are distributed to all switches in the fabric. After the commit, the active database on all switches are identical and learning can be disabled.

If the pending database contains more than one activation and auto-learning configuration when you commit the changes, the activation and auto-learning changes are consolidated and the resulting operation may change (see the following table).

**Table 93: Scenarios for Activation and Auto-learning Configurations in Distributed Mode**

| <b>Scenario</b>                                                                                    | <b>Actions</b>                                                       | <b>Distribution = OFF</b>                                                      | <b>Distribution = ON</b>                                                                                             |
|----------------------------------------------------------------------------------------------------|----------------------------------------------------------------------|--------------------------------------------------------------------------------|----------------------------------------------------------------------------------------------------------------------|
| A and B exist in the configuration database, activation is not done and devices C,D are logged in. | 1. You activate the port security database and enable auto-learning. | configuration database = {A,B}<br>active database = {A,B, C <sup>2</sup> , D*} | configuration database = {A,B}<br>active database = {null}<br>pending database = {A,B + activation to be enabled}    |
|                                                                                                    | 2. A new entry E is added to the configuration database.             | configuration database = {A,B, E}<br>active database = {A,B, C*, D*}           | configuration database = {A,B}<br>active database = {null}<br>pending database = {A,B, E + activation to be enabled} |
|                                                                                                    | 3. You issue a commit.                                               | Not applicable                                                                 | configuration database = {A,B, E}<br>active database = {A,B, E, C*, D*}<br>pending database = empty                  |

| Scenario                                                                                           | Actions                                                              | Distribution = OFF                                                | Distribution = ON                                                                                                                                                                     |
|----------------------------------------------------------------------------------------------------|----------------------------------------------------------------------|-------------------------------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| A and B exist in the configuration database, activation is not done and devices C,D are logged in. | 1. You activate the port security database and enable auto-learning. | configuration database = {A,B}<br>active database = {A,B, C*, D*} | configuration database = {A,B}<br>active database = {null}<br>pending database = {A,B + activation to be enabled}                                                                     |
|                                                                                                    | 2. You disable learning.                                             | configuration database = {A,B}<br>active database = {A,B, C, D}   | configuration database = {A,B}<br>active database = {null}<br>pending database = {A,B + activation to be enabled + learning to be disabled}                                           |
|                                                                                                    | 3. You issue a commit.                                               | Not applicable                                                    | configuration database = {A,B}<br>active database = {A,B} and devices C and D are logged out. This is equal to an activation with auto-learning disabled.<br>pending database = empty |

<sup>2</sup> The \* (asterisk) indicates learned entries.

## Port Security Database Merge Guidelines

A database merge refers to a union of the configuration database and static (unlearned) entries in the active database.

When merging the database between two fabrics, follow these guidelines:

- Verify that the activation status and the auto-learning status is the same in both fabrics.
- Verify that the combined number of configurations for each VSAN in both databases does not exceed 2000.



### Caution

If you do not follow these two conditions, the merge will fail. The next distribution will forcefully synchronize the databases and the activation states in the fabric.

### Related Topics

[CFS Merge Support, on page 334](#)

## Database Interaction

The following table lists the differences and interaction between the active and configuration databases.

**Table 94: Active and Configuration Port Security Databases**

| Active Database                                                                                                                                                                                       | Configuration Database                                                                                |
|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-------------------------------------------------------------------------------------------------------|
| Read-only.                                                                                                                                                                                            | Read-write.                                                                                           |
| Saving the configuration only saves the activated entries. Learned entries are not saved.                                                                                                             | Saving the configuration saves all the entries in the configuration database.                         |
| Once activated, all devices that have already logged into the VSAN are also learned and added to the active database.                                                                                 | Once activated, the configuration database can be modified without any effect on the active database. |
| You can overwrite the active database with the configured database by activating the port security database. Forcing an activation may violate the entries already configured in the active database. | You can overwrite the configuration database with the active database.                                |



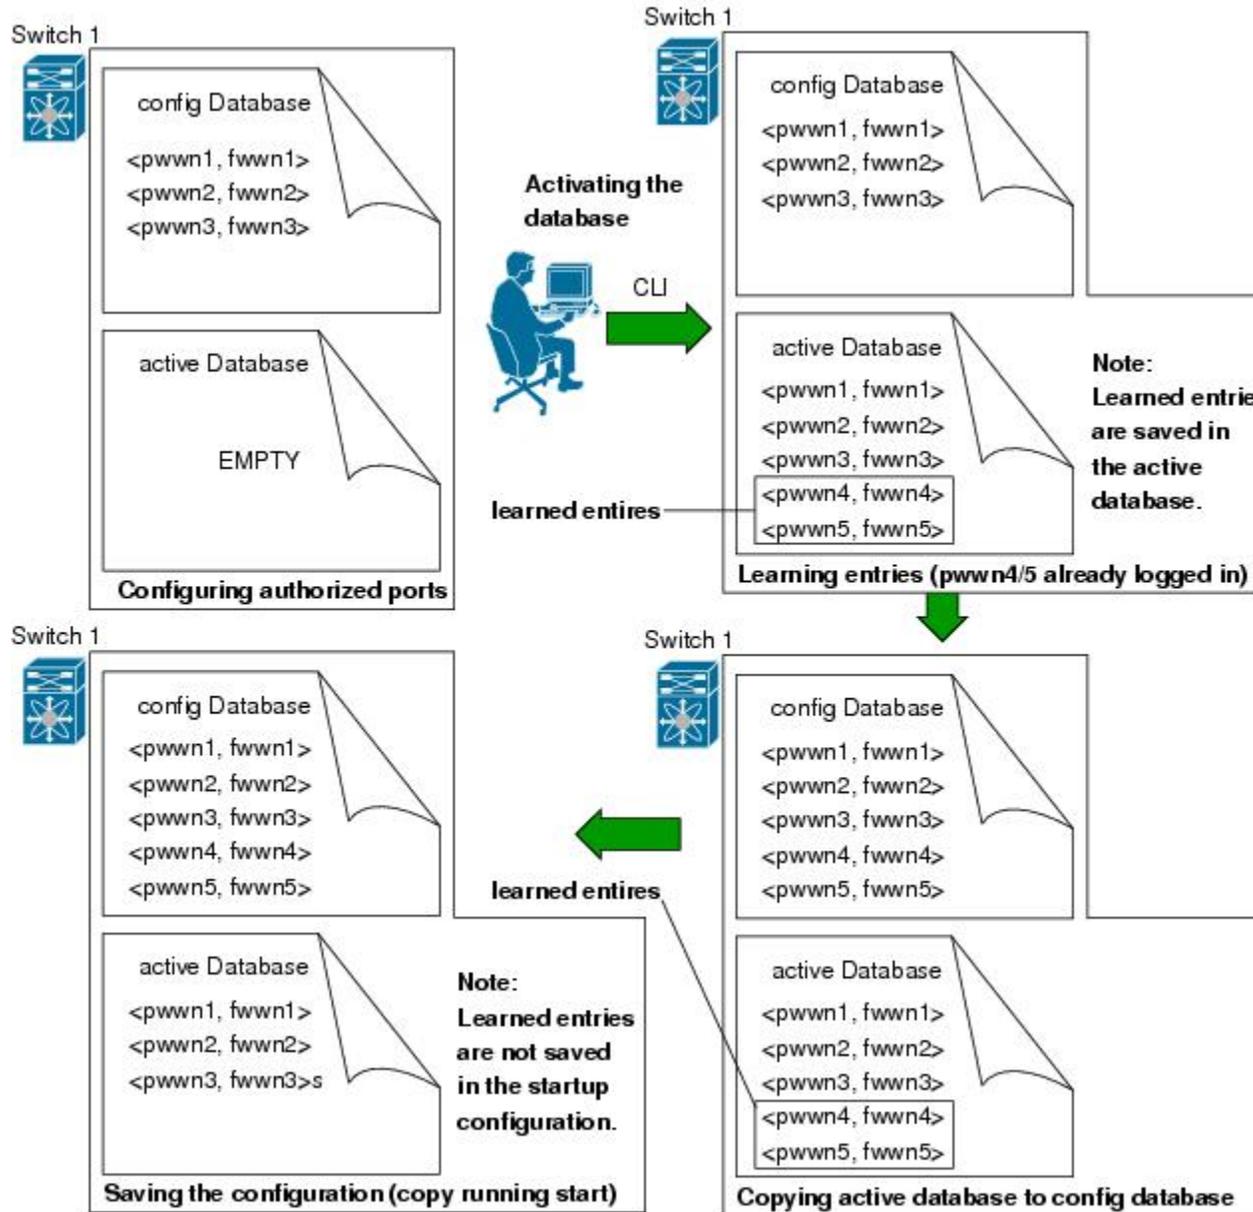
**Note**

You can overwrite the configuration database with the active database using the **port-security database copy vsan** command. The **port-security database diff active vsan** command in EXEC mode lists the differences between the active database and the configuration database.

### Database Scenarios

the following figure illustrates various scenarios showing the active database and the configuration database status based on port security configurations.

**Figure 74: Port Security Database Scenarios**



## Copying the Port Security Database



### Tip

We recommend that you copy the active database to the config database after disabling auto-learning. This action will ensure that the configuration database is in synchronization with the active database. If distribution is enabled, this command creates a temporary copy (and consequently a fabric lock) of the configuration database. If you lock the fabric, you need to commit the changes to the configuration databases in all the switches.

Use the **port-security database copy vsan** command to copy from the active to the configured database. If the active database is empty, this command is not accepted.

```
switch# port-security database copy vsan 1
```

Use the **port-security database diff active vsan** command to view the differences between the active database and the configuration database. This command can be used when resolving conflicts.

```
switch# port-security database diff active vsan 1
```

Use the **port-security database diff config vsan** command to obtain information on the differences between the configuration database and the active database.

```
switch# port-security database diff config vsan 1
```

## Deleting the Port Security Database



### Tip

If the distribution is enabled, the deletion creates a copy of the database. An explicit **port-security commit** command is required to actually delete the database.

Use the **no port-security database vsan** command in configuration mode to delete the configured database for a specified VSAN.

```
switch(config)# no port-security database vsan 1
```

## Clearing the Port Security Database

Use the **clear port-security statistics vsan** command to clear all existing statistics from the port security database for a specified VSAN.

```
switch# clear port-security statistics vsan 1
```

Use the **clear port-security database auto-learn interface** command to clear any learned entries in the active database for a specified interface within a VSAN.

```
switch# clear port-security database auto-learn interface fc2/1 vsan 1
```

Use the **clear port-security database auto-learn vsan** command to clear any learned entries in the active database for the entire VSAN.

```
switch# clear port-security database auto-learn vsan 1
```



### Note

The **clear port-security database auto-learn** and **clear port-security statistics** commands are only relevant to the local switch and do not acquire locks. Also, learned entries are only local to the switch and do not participate in distribution.

Use the **port-security clear vsan** command to clear the pending session in the VSAN from any switch in the VSAN.

```
switch# clear port-security session vsan 5
```

## Displaying Port Security Configuration

The **show port-security database** commands display the configured port security information. You can optionally specify a fWWN and a VSAN, or an interface and a VSAN in the **show port-security** command to view the output of the activated port security.

Access information for each port can be individually displayed. If you specify the fWWN or interface options, all devices that are paired in the active database (at that point) with the given fWWN or the interface are displayed.

The following example shows how to display the port security configuration database:

```
switch# show port-security database
```

The following example shows how to display the port security configuration database for VSAN 1:

```
switch# show port-security database vsan 1
```

The following example shows how to display the activated database:

```
switch# show port-security database active
```

The following example shows how to display difference between the temporary configuration database and the configuration database:

```
switch# show port-security pending-diff vsan 1
```

The following example shows how to display the configured fWWN port security in VSAN 1:

```
switch# show port-security database fwwn 20:01:00:05:30:00:95:de vsan 1
20:00:00:0c:88:00:4a:e2 (swwn)
```

The following example shows how to display the port security statistics:

```
switch# show port-security statistics
```

The following example shows how to verify the status of the active database and the auto-learning configuration:

```
switch# show port-security status
```

## Default Port Security Settings

The following table lists the default settings for all port security features in any switch.

**Table 95: Default Security Settings**

| Parameters    | Default                                                                               |
|---------------|---------------------------------------------------------------------------------------|
| Auto-learn    | Enabled if port security is enabled.                                                  |
| Port security | Disabled.                                                                             |
| Distribution  | Disabled.<br><b>Note</b> Enabling distribution enables it on all VSANs in the switch. |





## Configuring Fabric Binding

This chapter contains the following sections:

- [Configuring Fabric Binding, page 735](#)

### Configuring Fabric Binding

#### Information About Fabric Binding

The fabric binding feature ensures that ISLs are only enabled between specified switches in the fabric. Fabric binding is configured on a per-VSAN basis.

This feature helps prevent unauthorized switches from joining the fabric or disrupting current fabric operations. It uses the Exchange Fabric Membership Data (EFMD) protocol to ensure that the list of authorized switches is identical in all switches in the fabric.

#### Licensing Requirements for Fabric Binding

Fabric Binding requires the Storage Protocol Services license.

#### Port Security Versus Fabric Binding

Port security and fabric binding are two independent features that can be configured to complement each other. The following table compares the two features.

**Table 96: Fabric Binding and Port Security Comparison**

| <b>Fabric Binding</b>                           | <b>Port Security</b>                  |
|-------------------------------------------------|---------------------------------------|
| Uses a set of sWWNs and a persistent domain ID. | Uses pWWNs/nWWNs or fWWNs/sWWNs.      |
| Binds the fabric at the switch level.           | Binds devices at the interface level. |

| Fabric Binding                                                                                                                                          | Port Security                                                                                                                                                                                                                                                                                                             |
|---------------------------------------------------------------------------------------------------------------------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Authorizes only the configured sWWN stored in the fabric binding database to participate in the fabric.                                                 | Allows a preconfigured set of Fibre Channel devices to logically connect to a SAN port. The switch port, identified by a WWN or interface number, connects to a Fibre Channel device (a host or another switch), also identified by a WWN. By binding these two devices, you lock these two ports into a group (or list). |
| Requires activation on a per VSAN basis.                                                                                                                | Requires activation on a per VSAN basis.                                                                                                                                                                                                                                                                                  |
| Allows specific user-defined switches that are allowed to connect to the fabric, regardless of the physical port to which the peer switch is connected. | Allows specific user-defined physical ports to which another device can connect.                                                                                                                                                                                                                                          |
| Does not learn about switches that are logging in.                                                                                                      | Learns about switches or devices that are logging in if learning mode is enabled.                                                                                                                                                                                                                                         |
| Cannot be distributed by CFS and must be configured manually on each switch in the fabric.                                                              | Can be distributed by CFS.                                                                                                                                                                                                                                                                                                |

Port-level checking for xE ports is as follows:

- The switch login uses both port security binding and fabric binding for a given VSAN.
- Binding checks are performed on the port VSAN as follows:
  - E port security binding check on port VSAN
  - TE port security binding check on each allowed VSAN

While port security complements fabric binding, they are independent features and can be enabled or disabled separately.

### Fabric Binding Enforcement

To enforce fabric binding, configure the switch world wide name (sWWN) to specify the xE port connection for each switch. Enforcement of fabric binding policies are done on every activation and when the port tries to come up. For a Fibre Channel VSAN, the fabric binding feature requires all sWWNs connected to a switch to be part of the fabric binding active database.

## Configuring Fabric Binding

The fabric binding feature ensures ISLs are only enabled between specified switches in the fabric binding configuration. Fabric binding is configured on a per-VSAN basis.

### Configuring Fabric Binding

To configure fabric binding in each switch in the fabric, perform this task:

## SUMMARY STEPS

1. Enable the fabric configuration feature.
2. Configure a list of sWWNs and their corresponding domain IDs for devices that are allowed to access the fabric.
3. Activate the fabric binding database.
4. Copy the fabric binding active database to the fabric binding configuration database.
5. Save the fabric binding configuration.
6. Verify the fabric binding configuration.

## DETAILED STEPS

- 
- |               |                                                                                                                 |
|---------------|-----------------------------------------------------------------------------------------------------------------|
| <b>Step 1</b> | Enable the fabric configuration feature.                                                                        |
| <b>Step 2</b> | Configure a list of sWWNs and their corresponding domain IDs for devices that are allowed to access the fabric. |
| <b>Step 3</b> | Activate the fabric binding database.                                                                           |
| <b>Step 4</b> | Copy the fabric binding active database to the fabric binding configuration database.                           |
| <b>Step 5</b> | Save the fabric binding configuration.                                                                          |
| <b>Step 6</b> | Verify the fabric binding configuration.                                                                        |
- 

### Enabling Fabric Binding

The fabric binding feature must be enabled in each switch in the fabric that participates in the fabric binding. By default, this feature is disabled in Cisco Nexus 5000 Series switches. The configuration and verification commands for the fabric binding feature are only available when fabric binding is enabled on a switch. When you disable this configuration, all related configurations are automatically discarded.

To enable fabric binding on any participating switch, perform this task:

## SUMMARY STEPS

1. switch# **configuration terminal**
2. switch(config)# **fabric-binding enable**
3. switch(config)# **no fabric-binding enable**

## DETAILED STEPS

|               | Command or Action                               | Purpose                                           |
|---------------|-------------------------------------------------|---------------------------------------------------|
| <b>Step 1</b> | switch# <b>configuration terminal</b>           | Enters configuration mode.                        |
| <b>Step 2</b> | switch(config)# <b>fabric-binding enable</b>    | Enables fabric binding on that switch.            |
| <b>Step 3</b> | switch(config)# <b>no fabric-binding enable</b> | Disables (default) fabric binding on that switch. |

## About Switch WWN Lists

A user-specified fabric binding list contains a list of switch WWNs (sWWNs) within a fabric. If an sWWN attempts to join the fabric, and that sWWN is not on the list or the sWWN is using a domain ID that differs from the one specified in the allowed list, the ISL between the switch and the fabric is automatically isolated in that VSAN and the switch is denied entry into the fabric.

## Configuring Switch WWN List

To configure a list of sWWNs and optional domain IDs for a Fibre Channel VSAN, perform this task:

### SUMMARY STEPS

1. switch# **configuration terminal**
2. switch(config)# **fabric-binding database vsan** *vsan-id*
3. switch(config)# **no fabric-binding database vsan** *vsan-id*
4. switch(config-fabric-binding)#**swwn** *swwn-id* **domain** *domain-id*
5. switch(config-fabric-binding)#**no swwn** *swwn-id* **domain** *domain-id*

### DETAILED STEPS

|               | Command or Action                                                                           | Purpose                                                                                   |
|---------------|---------------------------------------------------------------------------------------------|-------------------------------------------------------------------------------------------|
| <b>Step 1</b> | switch# <b>configuration terminal</b>                                                       | Enters configuration mode.                                                                |
| <b>Step 2</b> | switch(config)# <b>fabric-binding database vsan</b> <i>vsan-id</i>                          | Enters the fabric binding submode for the specified VSAN.                                 |
| <b>Step 3</b> | switch(config)# <b>no fabric-binding database vsan</b> <i>vsan-id</i>                       | Deletes the fabric binding database for the specified VSAN.                               |
| <b>Step 4</b> | switch(config-fabric-binding)# <b>swwn</b> <i>swwn-id</i> <b>domain</b> <i>domain-id</i>    | Adds the sWWN of another switch for a specific domain ID to the configured database list. |
| <b>Step 5</b> | switch(config-fabric-binding)# <b>no swwn</b> <i>swwn-id</i> <b>domain</b> <i>domain-id</i> | Deletes the sWWN and domain ID of a switch from the configured database list.             |

## About Fabric Binding Activation and Deactivation

The fabric binding feature maintains a configuration database (config database) and an active database. The config database is a read-write database that collects the configurations you perform. These configurations are only enforced upon activation. This activation overwrites the active database with the contents of the config database. The active database is read-only and is the database that checks each switch that attempts to log in.

By default, the fabric binding feature is not activated. You cannot activate the fabric binding database on the switch if entries existing in the config database conflict with the current state of the fabric. For example, one of the already logged in switches may be denied login by the config database. You can choose to forcefully override these situations.



**Note** After activation, any already logged in switch that violates the current active database will be logged out, and all switches that were previously denied login because of fabric binding restrictions are reinitialized.

### Activating Fabric Binding

To activate the fabric binding feature, perform this task:

#### SUMMARY STEPS

1. switch# **configuration terminal**
2. switch(config)# **fabric-binding activate vsan vsan-id**
3. switch(config)# **no fabric-binding activate vsan vsan-id**

#### DETAILED STEPS

|               | Command or Action                                              | Purpose                                                         |
|---------------|----------------------------------------------------------------|-----------------------------------------------------------------|
| <b>Step 1</b> | switch# <b>configuration terminal</b>                          | Enters configuration mode.                                      |
| <b>Step 2</b> | switch(config)# <b>fabric-binding activate vsan vsan-id</b>    | Activates the fabric binding database for the specified VSAN.   |
| <b>Step 3</b> | switch(config)# <b>no fabric-binding activate vsan vsan-id</b> | Deactivates the fabric binding database for the specified VSAN. |

### Forcing Fabric Binding Activation

If the database activation is rejected due to one or more conflicts listed in the previous section, you may decide to proceed with the activation by using the force option.

To forcefully activate the fabric binding database, perform this task:

#### SUMMARY STEPS

1. switch# **configuration terminal**
2. switch(config)# **fabric-binding activate vsan vsan-id force**
3. switch(config)# **no fabric-binding activate vsan vsan-id force**

## DETAILED STEPS

|               | Command or Action                                                        | Purpose                                                                                                               |
|---------------|--------------------------------------------------------------------------|-----------------------------------------------------------------------------------------------------------------------|
| <b>Step 1</b> | switch# <b>configuration terminal</b>                                    | Enters configuration mode.                                                                                            |
| <b>Step 2</b> | switch(config)# <b>fabric-binding activate vsan<br/>vsan-id force</b>    | Activates the fabric binding database for the specified VSAN forcefully, even if the configuration is not acceptable. |
| <b>Step 3</b> | switch(config)# <b>no fabric-binding activate vsan<br/>vsan-id force</b> | Reverts to the previously configured state or to the factory default (if no state is configured).                     |

### Copying Fabric Binding Configurations

When you copy the fabric binding configuration, the config database is saved to the running configuration.

You can use the following commands to copy to the config database:

- Use the **fabric-binding database copy vsan** command to copy from the active database to the config database. If the configured database is empty, this command is not accepted.

```
switch# fabric-binding database copy vsan 1
```

- Use the **fabric-binding database diff active vsan** command to view the differences between the active database and the config database. This command can be used when resolving conflicts.

```
switch# fabric-binding database diff active vsan 1
```

- Use the **fabric-binding database diff config vsan** command to obtain information on the differences between the config database and the active database.

```
switch# fabric-binding database diff config vsan 1
```

- Use the **copy running-config startup-config** command to save the running configuration to the startup configuration so that the fabric binding config database is available after a reboot.

```
switch# copy running-config startup-config
```

### Clearing the Fabric Binding Statistics

Use the **clear fabric-binding statistics** command to clear all existing statistics from the fabric binding database for a specified VSAN.

```
switch# clear fabric-binding statistics vsan 1
```

### Deleting the Fabric Binding Database

Use the **no fabric-binding** command in configuration mode to delete the configured database for a specified VSAN.

```
switch(config)# no fabric-binding database vsan 10
```

## Verifying Fabric Binding Information

To display fabric binding information, perform one of the following tasks

### SUMMARY STEPS

1. switch# **show fabric-binding database [active]**
2. switch# **show fabric-binding database [active] [vsan vsan-id]**
3. switch# **show fabric-binding statistics**
4. switch# **show fabric-binding status**
5. switch# **show fabric-binding violations**
6. switch# **show fabric-binding efmd [vsan vsan-id]**

### DETAILED STEPS

|               | Command or Action                                                   | Purpose                                                                                                                     |
|---------------|---------------------------------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------|
| <b>Step 1</b> | switch# <b>show fabric-binding database [active]</b>                | Displays the configured fabric binding database. Include keyword active to display only the active fabric binding database. |
| <b>Step 2</b> | switch# <b>show fabric-binding database [active] [vsan vsan-id]</b> | Displays the configured fabric binding database for the specified VSAN.                                                     |
| <b>Step 3</b> | switch# <b>show fabric-binding statistics</b>                       | Displays statistics for the fabric binding database.                                                                        |
| <b>Step 4</b> | switch# <b>show fabric-binding status</b>                           | Displays fabric binding status for all VSANs.                                                                               |
| <b>Step 5</b> | switch# <b>show fabric-binding violations</b>                       | Displays fabric binding violations.                                                                                         |
| <b>Step 6</b> | switch# <b>show fabric-binding efmd [vsan vsan-id]</b>              | Displays the configured fabric binding database for the specified VSAN.                                                     |

The following example displays the active fabric binding information for VSAN 4:

```
switch# show fabric-binding database active vsan 4
```

The following example displays fabric binding violations:

```
switch# show fabric-binding violations
```

```

VSAN Switch WWN [domain] Last-Time [Repeat count] Reason

2 20:00:00:05:30:00:4a:1e [0xeb] Nov 25 05:46:14 2003 [2] Domain mismatch
3 20:00:00:05:30:00:4a:1e [*] Nov 25 05:44:58 2003 [2] sWWN not found
4 20:00:00:05:30:00:4a:1e [*] Nov 25 05:46:25 2003 [1] Database mismatch
```

**Note**

In VSAN 3, the sWWN was not found in the list. In VSAN 2, the sWWN was found in the list, but has a domain ID mismatch.

The following example displays EFMD Statistics for VSAN 4:

```
switch# show fabric-binding efmd statistics vsan 4
```

## Default Fabric Binding Settings

The following table lists the default settings for the fabric binding feature.

**Table 97: Default Fabric Binding Settings**

| Parameters     | Default  |
|----------------|----------|
| Fabric binding | Disabled |



# CHAPTER 49

## Configuring Fabric Configuration Servers

This chapter contains the following sections:

- [Configuring Fabric Configuration Servers, page 743](#)

## Configuring Fabric Configuration Servers

### Information About FCS

The Fabric Configuration Server (FCS) provides discovery of topology attributes and maintains a repository of configuration information of fabric elements. A management application is usually connected to the FCS on the switch through an N port. The FCS views the entire fabric based on the following objects:

- Interconnect element (IE) object—Each switch in the fabric corresponds to an IE object. One or more IE objects form a fabric.
- Port object—Each physical port in an IE corresponds to a port object. This includes the switch ports (xE and F ports) and their attached N ports.
- Platform object—A set of nodes may be defined as a platform object to make it a single manageable entity. These nodes are end-devices (host systems, storage subsystems) attached to the fabric. Platform objects reside at the edge switches of the fabric.

Each object has its own set of attributes and values. A null value may also be defined for some attributes.

In the Cisco Nexus 5000 Series switch environment, a fabric may consist of multiple VSANs. One instance of the FCS is present per VSAN.

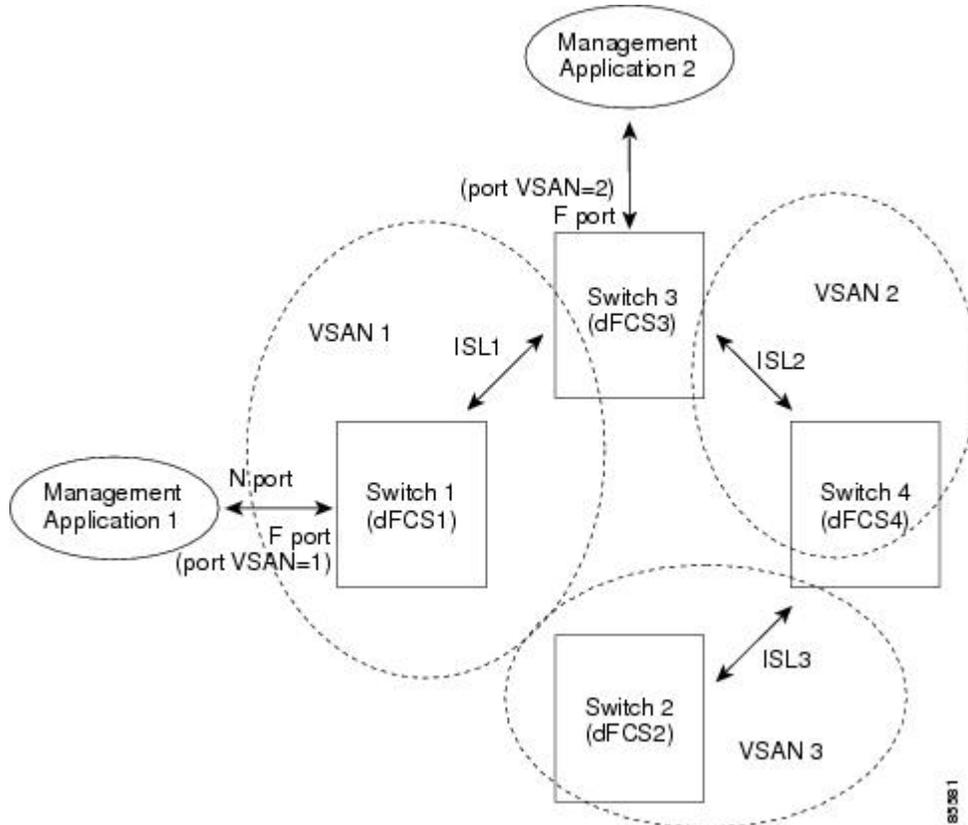
FCS supports the discovery of virtual devices. The **fcs virtual-device-add** command, entered in FCS configuration submode, allows you to discover virtual devices in a particular VSAN or in all VSANs.

If you have attached a management application to a switch, all the frames directed towards the FCS in the switch are part of the port VSAN in the switch port (F port). Your view of the management application is limited only to this VSAN. However, information about other VSANs that this switch is part of can be obtained either through the SNMP or CLI.

In the following figure, Management Application 1 (M1) is connected through an F port with port VSAN ID 1, and Management Application 2 (M2) is connected through an F port with port VSAN ID 2. M1 can query the FCS information of switches S1 and S3, and M2 can query switches S3 and S4. Switch S2 information is

not known to both of them. FCS operations can be done only on those switches that are visible in the VSAN. M2 can send FCS requests only for VSAN 2 even though S3 is also a part of VSAN 1.

**Figure 75: FCSs in a VSAN Environment**



## FCS Characteristics

FCSs have the following characteristics:

- Support network management including the following:
  - N port management application can query and obtain information about fabric elements.
  - SNMP manager can use the FCS management information base (MIB) to start discovery and obtain information about the fabric topology.
- Support TE ports in addition to the standard F and E ports.
- Can maintain a group of nodes with a logical name and management address when a platform registers with it. FCSs maintain a backup of all registrations in secondary storage and update it with every change. When a restart or switchover happens, FCSs retrieve the secondary storage information and rebuild its database.
- SNMP manager can query FCSs for all IEs, ports, and platforms in the fabric.

## FCS Name Specification

You can specify if the unique name verification is for the entire fabric (globally) or only for locally (default) registered platforms.



### Note

Set this command globally only if every switch in the fabric belong to the Cisco MDS 9000 Family or Cisco Nexus 5000 Series of switches.

To enable global checking of the platform name, perform this task:

To register platform attributes, perform this task:

## SUMMARY STEPS

1. switch# **configuration terminal**
2. switch(config)# **fcs plat-check-global vsan vsan-id**
3. switch(config)# **no fcs plat-check-global vsan vsan-id**

## DETAILED STEPS

|               | Command or Action                                            | Purpose                                                  |
|---------------|--------------------------------------------------------------|----------------------------------------------------------|
| <b>Step 1</b> | switch# <b>configuration terminal</b>                        | Enters configuration mode.                               |
| <b>Step 2</b> | switch(config)# <b>fcs plat-check-global vsan vsan-id</b>    | Enables global checking of the platform name.            |
| <b>Step 3</b> | switch(config)# <b>no fcs plat-check-global vsan vsan-id</b> | Disables (default) global checking of the platform name. |

## Displaying FCS Information

You can use the **show fcs** commands to display the status of the WWN configuration.

The following example shows how to display the FCS local database:

```
switch# show fcs database
```

The following example shows how to display a list of all interconnect elements for VSAN 1:

```
switch# show fcs ie vsan 1
```

The following example shows how to display information for a specific platform:

```
switch# show fcs platform name SamplePlatform vsan 1
```

The following example shows how to display port information for a specific pWWN:

```
switch# show fcs port pwn 20:51:00:05:30:00:16:de vsan 24
```

## Default FCS Settings

The following table lists the default FCS settings.

**Table 98: Default FCS Settings**

| <b>Parameters</b>                    | <b>Default</b> |
|--------------------------------------|----------------|
| Global checking of the platform name | Disabled       |
| Platform node type                   | Unknown        |



# Configuring Port Tracking

---

This chapter contains the following sections:

- [Configuring Port Tracking, page 747](#)

## Configuring Port Tracking

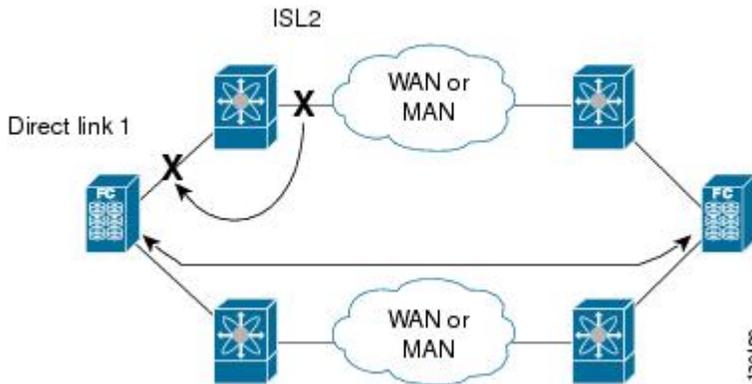
Cisco Nexus 5000 Series switches offer the port tracking feature on physical Fibre Channel interfaces (but not on virtual Fibre Channel interfaces). This feature uses information about the operational state of the link to initiate a failure in the link that connects the edge device. This process of converting the indirect failure to a direct failure triggers a faster recovery process towards redundant links. When enabled, the port tracking feature brings down the configured links based on the failed link and forces the traffic to be redirected to another redundant link.

### Information About Port Tracking

Generally, hosts can instantly recover from a link failure on a link that is immediately (direct link) connected to a switch. However, recovering from an indirect link failure between switches in a WAN or MAN fabric with a keepalive mechanism is dependent on several factors such as the timeout values (TOVs) and on registered state change notification (RSCN) information.

In the following figure, when the direct link 1 to the host fails, recovery can be immediate. However, when the ISL 2 fails between the two switches, recovery depends on TOVs, RSCNs, and other factors.

**Figure 76: Traffic Recovery Using Port Tracking**



The port tracking feature monitors and detects failures that cause topology changes and brings down the links connecting the attached devices. When you enable this feature and explicitly configure the linked and tracked ports, the switch software monitors the tracked ports and alters the operational state of the linked ports on detecting a link state change.

The following terms are used in this chapter:

- **Tracked ports**—A port whose operational state is continuously monitored. The operational state of the tracked port is used to alter the operational state of one or more ports. Fibre Channel, VSAN, SAN port channel, or a Gigabit Ethernet port can be tracked. Generally, ports in E and TE port modes can also be F ports.
- **Linked ports**—A port whose operational state is altered based on the operational state of the tracked ports. Only physical Fibre Channel ports can be linked ports.

Port tracking has the following features:

- The application brings the linked port down when the tracked port goes down. When the tracked port recovers from the failure and comes back up again, the linked port is also brought up automatically (unless otherwise configured).
- You can forcefully continue to keep the linked port down, even though the tracked port comes back up. In this case, you must explicitly bring up the linked port when required.

### Related Topics

[About RSCN Information, on page 667](#)

[Fibre Channel Timeout Values, on page 679](#)

## Configuring Port Tracking

Before configuring port tracking, consider the following guidelines:

- Verify that the tracked ports and the linked ports are on the same Cisco switch.

- Be aware that the linked port is automatically brought down when the tracked port goes down.
- Do not track a linked port back to itself (for example, Port fc2/2 to Port fc2/4 and back to Port fc2/2) to avoid recursive dependency.

### Enabling Port Tracking

The port tracking feature is disabled by default in Cisco Nexus 5000 Series switches. When you enable this feature, port tracking is globally enabled for the entire switch.

To configure port tracking, enable the port tracking feature and configure the linked ports for the tracked port.

To enable port tracking, perform this task:

### SUMMARY STEPS

1. switch# **configuration terminal**
2. switch(config)# **port-track enable**
3. switch(config)# **no port-track enable**

### DETAILED STEPS

|               | Command or Action                           | Purpose                                                                               |
|---------------|---------------------------------------------|---------------------------------------------------------------------------------------|
| <b>Step 1</b> | switch# <b>configuration terminal</b>       | Enters configuration mode.                                                            |
| <b>Step 2</b> | switch(config)# <b>port-track enable</b>    | Enables port tracking.                                                                |
| <b>Step 3</b> | switch(config)# <b>no port-track enable</b> | Removes the currently applied port tracking configuration and disables port tracking. |

### About Configuring Linked Ports

You can link ports using one of two methods:

- Operationally binding the linked ports to the tracked port (default).
- Continuing to keep the linked port down forcefully, even if the tracked port has recovered from the link failure.

### Operationally Binding a Tracked Port

When you configure the first tracked port, operational binding is automatically in effect. When you use this method, you have the option to monitor multiple ports or monitor ports in one VSAN.

To operationally bind a tracked port, perform this task:

**SUMMARY STEPS**

1. switch# **configuration terminal**
2. switch(config)# **interface fc slot/port**
3. switch(config-if)# **port-track interface fc slot/port | san-port-channel port**
4. switch(config-if)# **no port-track interface fc slot/port | san-port-channel port**

**DETAILED STEPS**

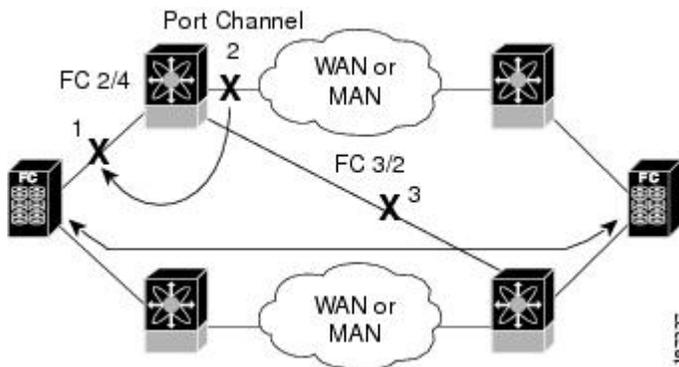
|               | Command or Action                                                                      | Purpose                                                                                               |
|---------------|----------------------------------------------------------------------------------------|-------------------------------------------------------------------------------------------------------|
| <b>Step 1</b> | switch# <b>configuration terminal</b>                                                  | Enters configuration mode.                                                                            |
| <b>Step 2</b> | switch(config)# <b>interface fc slot/port</b>                                          | Enters the interface configuration mode for the linked port. You can now configure the tracked ports. |
| <b>Step 3</b> | switch(config-if)# <b>port-track interface fc slot/port   san-port-channel port</b>    | Specifies the tracked port. When the tracked port goes down, the linked port is also brought down.    |
| <b>Step 4</b> | switch(config-if)# <b>no port-track interface fc slot/port   san-port-channel port</b> | Removes the port tracking configuration that is currently applied to the interface.                   |

**About Tracking Multiple Ports**

You can control the operational state of the linked port based on the operational states of multiple tracked ports. When more than one tracked port is associated with a linked port, the operational state of the linked port will be set to down only if all the associated tracked ports are down. Even if one tracked port is up, the linked port will stay up.

In the following figure, only if both ISLs 2 and 3 fail, will the direct link 1 be brought down. Direct link 1 will not be brought down if either 2 or 3 are still functioning as desired.

*Figure 77: Traffic Recovery Using Port Tracking*



## Tracking Multiple Ports

To track multiple ports, perform this task:

### SUMMARY STEPS

1. switch# **configuration terminal**
2. switch(config)# **interface fc slot/port**
3. switch(config-if)# **port-track interface interface fc slot/port | san-port-channel port**

### DETAILED STEPS

|               | Command or Action                                                                             | Purpose                                                                                                                     |
|---------------|-----------------------------------------------------------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------|
| <b>Step 1</b> | switch# <b>configuration terminal</b>                                                         | Enters configuration mode.                                                                                                  |
| <b>Step 2</b> | switch(config)# <b>interface fc slot/port</b>                                                 | Configures the specified interface and enters the interface configuration mode. You can now configure tracked ports.        |
| <b>Step 3</b> | switch(config-if)# <b>port-track interface interface fc slot/port   san-port-channel port</b> | Tracks the linked port with the specified interface. When the tracked port goes down, the linked port is also brought down. |

## About Monitoring Ports in a VSAN

You can optionally configure one VSAN from the set of all operational VSANs on the tracked port with the linked port by specifying the required VSAN. This level of flexibility provides higher granularity in tracked ports. In some cases, when a tracked port is a TE port, the set of operational VSANs on the port can change dynamically without bringing down the operational state of the port. In such cases, the port VSAN of the linked port can be monitored on the set of operational VSANs on the tracked port.

If you configure this feature, the linked port is up only when the VSAN is up on the tracked port.

The specified VSAN does not have to be the same as the port VSAN of the linked port.

## Monitoring Ports in a VSAN

To monitor a tracked port in a specific VSAN, perform this task :

### SUMMARY STEPS

1. switch# **configuration terminal**
2. switch(config)# **interface fc slot/port**
3. switch(config-if)# **port-track interface san-port-channel 1 vsan 2**
4. switch(config-if)# **no port-track interface san-port-channel 1 vsan 2**

**DETAILED STEPS**

|               | <b>Command or Action</b>                                                    | <b>Purpose</b>                                                                                                       |
|---------------|-----------------------------------------------------------------------------|----------------------------------------------------------------------------------------------------------------------|
| <b>Step 1</b> | switch# <b>configuration terminal</b>                                       | Enters configuration mode.                                                                                           |
| <b>Step 2</b> | switch(config)# <b>interface fc slot/port</b>                               | Configures the specified interface and enters the interface configuration mode. You can now configure tracked ports. |
| <b>Step 3</b> | switch(config-if)# <b>port-track interface san-port-channel 1 vsan 2</b>    | Enables tracking of the SAN port channel in VSAN 2.                                                                  |
| <b>Step 4</b> | switch(config-if)# <b>no port-track interface san-port-channel 1 vsan 2</b> | Removes the VSAN association for the linked port. The SAN port channel link remains in effect.                       |

**About Forceful Shutdown**

If a tracked port flaps frequently, then tracking ports using the operational binding feature may cause frequent topology change. In this case, you may choose to keep the port in the down state until you are able to resolve the reason for these frequent flaps. Keeping the flapping port in the down state forces the traffic to flow through the redundant path until the primary tracked port problems are resolved. When the problems are resolved and the tracked port is back up, you can explicitly enable the interface.

If you configure this feature, the linked port continues to remain in the shutdown state even after the tracked port comes back up. You must explicitly remove the forced shut state (by administratively bringing up this interface) of the linked port once the tracked port is up and stable.

**Forcefully Shutting Down a Tracked Port**

To forcefully shut down a tracked port, perform this task:

**SUMMARY STEPS**

1. switch# **configuration terminal**
2. switch(config)# **interface fc slot/port**
3. switch(config-if)# **port-track force-shut**
4. switch(config-if)# **no port-track force-shut**

**DETAILED STEPS**

|               | <b>Command or Action</b>                      | <b>Purpose</b>                                                                                                       |
|---------------|-----------------------------------------------|----------------------------------------------------------------------------------------------------------------------|
| <b>Step 1</b> | switch# <b>configuration terminal</b>         | Enters configuration mode.                                                                                           |
| <b>Step 2</b> | switch(config)# <b>interface fc slot/port</b> | Configures the specified interface and enters the interface configuration mode. You can now configure tracked ports. |

|               | Command or Action                                  | Purpose                                                       |
|---------------|----------------------------------------------------|---------------------------------------------------------------|
| <b>Step 3</b> | switch(config-if)# <b>port-track force-shut</b>    | Forcefully shuts down the tracked port.                       |
| <b>Step 4</b> | switch(config-if)# <b>no port-track force-shut</b> | Removes the port shutdown configuration for the tracked port. |

## Displaying Port Tracking Information

The **show** commands display the current port tracking settings for the switch.

The following example shows how to display tracked port configuration for a specific interface:

```
switch# show interface fc2/1
fc2/1 is down (Administratively down)
 Hardware is Fibre Channel, FCOT is short wave laser w/o OFC (SN)
 Port WWN is 20:01:00:05:30:00:0d:de
 Admin port mode is FX
 Port vsan is 1
 Receive data field Size is 2112
 Beacon is turned off
 Port tracked with interface fc2/2 (down)
 Port tracked with interface san-port-channel 1 vsan 2 (down)
 5 minutes input rate 0 bits/sec, 0 bytes/sec, 0 frames/sec
 ...
```

The following example shows how to display tracked port configuration for a SAN port channel:

```
switch# show interface san-port-channel 1
port-channel 1 is down (No operational members)
 Hardware is Fibre Channel
 Port WWN is 24:01:00:05:30:00:0d:de
 Admin port mode is auto, trunk mode is on
 Port vsan is 2
 Linked to 1 port(s)
 Port linked to interface fc2/1
 ...
```

The following example shows how to display the port track mode:

```
switch# show interface fc 2/4
fc2/4 is up
 Hardware is Fibre Channel, FCOT is short wave laser
 ...
 Transmit B2B Credit is 64
 Receive B2B Credit is 16
 Receive data field Size is 2112
 Beacon is turned off
 Port track mode is force_shut <-- this port remains shut even if the tracked port is
 back up
```

## Default Port Tracking Settings

The following table lists the default settings for port tracking parameters.

**Table 99: Default Port Tracking Parameters**

| Parameters    | Default  |
|---------------|----------|
| Port tracking | Disabled |

| Parameters          | Default                          |
|---------------------|----------------------------------|
| Operational binding | Enabled along with port tracking |



# PART **VIII**

## **Troubleshooting**

- [Configuring SPAN, page 757](#)
- [Troubleshooting, page 767](#)





# CHAPTER 51

## Configuring SPAN

---

This chapter contains the following sections:

- [Configuring SPAN, page 757](#)

### Configuring SPAN

The Switched Port Analyzer (SPAN) feature (sometimes called port mirroring or port monitoring) selects network traffic for analysis by a network analyzer. The network analyzer can be a Cisco SwitchProbe, a Fibre Channel Analyzer, or other Remote Monitoring (RMON) probes.

#### SPAN Sources

SPAN sources refer to the interfaces from which traffic can be monitored. The Cisco Nexus 5000 Series switch supports Ethernet, Fibre Channel, virtual Fibre Channel, port channels, SAN port channels, VLANs, and VSANs as SPAN sources. With VLANs or VSANs, all supported interfaces in the specified VLAN or VSAN are included as SPAN sources. You can choose the SPAN traffic in the ingress direction, the egress direction, or both directions for Ethernet, Fibre Channel, and virtual Fibre Channel source interfaces:

- Ingress source (Rx)—Traffic entering the switch through this source port is copied to the SPAN destination port.
- Egress source (Tx)—Traffic exiting the switch through this source port is copied to the SPAN destination port.

#### Characteristics of Source Ports

A source port, also called a monitored port, is a switched interface that you monitor for network traffic analysis. The switch supports any number of ingress source ports (up to the maximum number of available ports on the switch) and any number of source VLANs or VSANs.

A source port has these characteristics:

- Can be of any port type: Ethernet, Fibre Channel, virtual Fibre Channel, port channel, SAN port channel, VLAN, and VSAN.
- Cannot be monitored in multiple SPAN sessions.

- Cannot be a destination port.
- Each source port can be configured with a direction (ingress, egress, or both) to monitor. For VLAN and VSAN sources, the monitored direction can only be ingress and applies to all physical ports in the group. The RX/TX option is not available for VLAN or VSAN SPAN sessions.
- Beginning with Cisco NX-OS Release 5.0(2)N1(1), Port Channel and SAN Port Channel interfaces can be configured as ingress or egress source ports.
- Source ports can be in the same or different VLANs or VSANs.
- For VLAN or VSAN SPAN sources, all active ports in the source VLAN or VSAN are included as source ports.
- For Cisco NX-OS Release 4.2(1)N2(1) and earlier, the Cisco Nexus 5010 Switch and the Cisco Nexus 5020 Switch supports a maximum of two egress SPAN source ports.

## SPAN Destinations

SPAN destinations refer to the interfaces that monitors source ports. The Cisco Nexus 5000 Series switch supports Ethernet and Fibre Channel interfaces as SPAN destinations.

| Source SPAN           | Dest SPAN       |
|-----------------------|-----------------|
| Ethernet              | Ethernet        |
| Fibre Channel         | Fibre Channel   |
| Fibre Channel         | Ethernet (FCoE) |
| Virtual Fibre Channel | Fibre Channel   |
| Virtual Fibre Channel | Ethernet (FCoE) |

## Characteristics of Destination Ports

Each local SPAN session must have a destination port (also called a monitoring port) that receives a copy of traffic from the source ports, VLANs, or VSANs. A destination port has these characteristics:

- Can be any physical port, Ethernet, Ethernet (FCoE), or Fibre Channel, and virtual Fibre Channel ports cannot be destination ports.
- Cannot be a source port.
- Cannot be a port channel or SAN port channel group.
- Does not participate in spanning tree while the SPAN session is active.
- Is excluded from the source list and is not monitored if it belongs to a source VLAN of any SPAN session.
- Receives copies of sent and received traffic for all monitored source ports. If a destination port is oversubscribed, it can become congested. This congestion can affect traffic forwarding on one or more of the source ports.

## Configuring SPAN

### Creating and Deleting a SPAN Session

You create a SPAN session by assigning a session number using the monitor command. If the session already exists, any additional configuration is added to that session.

#### SUMMARY STEPS

1. switch# **configure terminal**
2. switch(config)# **monitor session** *session-number*

#### DETAILED STEPS

|               | Command or Action                                            | Purpose                                                                                                          |
|---------------|--------------------------------------------------------------|------------------------------------------------------------------------------------------------------------------|
| <b>Step 1</b> | switch# <b>configure terminal</b>                            | Enters configuration mode.                                                                                       |
| <b>Step 2</b> | switch(config)# <b>monitor session</b> <i>session-number</i> | Enters the monitor configuration mode. New session configuration is added to the existing session configuration. |

### Configuring the Destination Port

*Configuring an Ethernet Destination Port*



#### Note

The SPAN destination port can only be a physical port on the switch.

You can configure an Ethernet interface as a SPAN destination port.

#### SUMMARY STEPS

1. switch# **configure terminal**
2. switch(config)# **interface ethernet** *slot/port*
3. switch(config-if)# **switchport monitor**
4. switch(config-if)# **exit**
5. switch(config)# **monitor session** *session-number*
6. switch(config-monitor)# **destination interface ethernet** *slot/port*

## DETAILED STEPS

|               | Command or Action                                                       | Purpose                                                                                                                  |
|---------------|-------------------------------------------------------------------------|--------------------------------------------------------------------------------------------------------------------------|
| <b>Step 1</b> | switch# <b>configure terminal</b>                                       | Enters configuration mode.                                                                                               |
| <b>Step 2</b> | switch(config)# <b>interface ethernet slot/port</b>                     | Enters interface configuration mode for the specified Ethernet interface selected by the slot and port values.           |
| <b>Step 3</b> | switch(config-if)# <b>switchport monitor</b>                            | Sets the interface to monitor mode. Priority flow control is disabled when the port is configured as a SPAN destination. |
| <b>Step 4</b> | switch(config-if)# <b>exit</b>                                          | Reverts to global configuration mode.                                                                                    |
| <b>Step 5</b> | switch(config)# <b>monitor session session-number</b>                   | Enters the monitor configuration mode.                                                                                   |
| <b>Step 6</b> | switch(config-monitor)# <b>destination interface ethernet slot/port</b> | Configures the Ethernet destination port.                                                                                |

The following example shows configuring an Ethernet SPAN destination port:

```
switch# configure terminal
switch(config)# interface ethernet 1/3
switch(config-if)# switchport monitor
switch(config-if)# exit
switch(config)# monitor session 2
switch(config-monitor)# destination interface ethernet 1/3
```

## Configuring Fibre Channel Destination Port

**Note**

The SPAN destination port can only be a physical port on the switch.

You can configure a Fibre Channel port as a SPAN destination port.

## SUMMARY STEPS

1. switch# **configure terminal**
2. switch(config)# **interface fc slot/port**
3. switch(config-if)# **switchport mode SD**
4. switch(config-if)# **switchport speed 1000**
5. switch(config-if)# **exit**
6. switch(config)# **monitor session session-number**
7. switch(config-monitor)# **destination interface fc slot/port**

## DETAILED STEPS

|               | Command or Action                                                 | Purpose                                                                                                             |
|---------------|-------------------------------------------------------------------|---------------------------------------------------------------------------------------------------------------------|
| <b>Step 1</b> | switch# <b>configure terminal</b>                                 | Enters configuration mode.                                                                                          |
| <b>Step 2</b> | switch(config)# <b>interface fc slot/port</b>                     | Enters interface configuration mode for the specified Fibre Channel interface selected by the slot and port values. |
| <b>Step 3</b> | switch(config-if)# <b>switchport mode SD</b>                      | Sets the interface to SPAN destination (SD) mode.                                                                   |
| <b>Step 4</b> | switch(config-if)# <b>switchport speed 1000</b>                   | Sets the interface speed to 1000. The auto speed option is not allowed.                                             |
| <b>Step 5</b> | switch(config-if)# <b>exit</b>                                    | Reverts to global configuration mode.                                                                               |
| <b>Step 6</b> | switch(config)# <b>monitor session session-number</b>             | Enters the monitor configuration mode.                                                                              |
| <b>Step 7</b> | switch(config-monitor)# <b>destination interface fc slot/port</b> | Configures the Fibre Channel destination port.                                                                      |

The following example shows configuring an Ethernet SPAN destination port:

```
switch# configure terminal
switch(config)# interface fc 2/4
switch(config-if)# switchport mode SD
switch(config-if)# switchport speed 1000
switch(config-if)# exit
switch(config)# monitor session 2
switch(config-monitor)# destination interface fc 2/4
```

### Configuring Source Ports

You can configure the source ports for a SPAN session. The source ports can be Ethernet, Fibre Channel, or virtual Fibre Channel ports.

## SUMMARY STEPS

1. switch(config-monitor)# **source interface type slot/port [rx | tx | both]**

## DETAILED STEPS

|               | Command or Action                                                               | Purpose                                                                                                                                                                                                                                                                                    |
|---------------|---------------------------------------------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Step 1</b> | switch(config-monitor)# <b>source interface type slot/port [rx   tx   both]</b> | Configures sources and the traffic direction in which to duplicate packets. You can enter a range of Ethernet, Fibre Channel, or virtual Fibre Channel ports. You can specify the traffic direction to duplicate as ingress (rx), egress (tx), or both. By default, the direction is both. |

The following example shows configuring an Ethernet SPAN source port:

```
switch# configure terminal
switch(config)# monitor session 2
switch(config-monitor)# source interface ethernet 1/16
```

The following example shows configuring a Fibre Channel SPAN source port:

```
switch(config-monitor)# source interface fc 2/1
```

The following example shows configuring a virtual Fibre Channel SPAN source port:

```
switch(config-monitor)# source interface vfc 129
```

### Configuring Source Port Channels, VLANs, or VSANs

You can configure the source channels for a SPAN session. These ports can be port channels, SAN port channels, VLANs, and VSANs. The monitored direction can only be ingress and applies to all physical ports in the group.



**Note**

The Cisco Nexus 5000 Series switch supports two active SPAN sessions. When you configure more than two SPAN sessions, the first two sessions are active. During startup, the order of active sessions is reversed; the last two sessions are active. For example, if you configured ten sessions 1 to 10 where 1 and 2 are active, after a reboot, sessions 9 and 10 will be active. To enable deterministic behavior, explicitly suspend the sessions 3 to 10 with the **monitor session session-number shut** command. See *Suspending a SPAN Session*.

### SUMMARY STEPS

1. switch(config-monitor)# source {interface {port-channel | san-port-channel} channel-number rx | vlan vlan-range | vsan vsan-range }

### DETAILED STEPS

|        | Command or Action                                                                                                                   | Purpose                                                                                                                          |
|--------|-------------------------------------------------------------------------------------------------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------|
| Step 1 | switch(config-monitor)# source {interface {port-channel   san-port-channel} channel-number rx   vlan vlan-range   vsan vsan-range } | Configures port channel, SAN port channel, VLAN, or VSAN sources. For VLAN or VSAN sources, the monitored direction is implicit. |

The following example shows configuring a port channel SPAN source:

```
switch# configure terminal
switch(config)# monitor session 2
switch(config-monitor)# source interface port-channel 1 rx
```

The following example shows configuring a SAN port channel SPAN source:

```
switch(config-monitor)# source interface san-port-channel 3 rx
```

The following example shows configuring a VLAN SPAN source:

```
switch(config-monitor)# source vlan 1
```

The following example shows configuring a VSAN SPAN source:

```
switch(config-monitor)# source vsan 1
```

## Configuring the Description of a SPAN Session

You can provide a descriptive name of the SPAN session for ease of reference.

### SUMMARY STEPS

1. `switch(config-monitor)# description description`

### DETAILED STEPS

|        | Command or Action                                                   | Purpose                                         |
|--------|---------------------------------------------------------------------|-------------------------------------------------|
| Step 1 | <code>switch(config-monitor)# description <i>description</i></code> | Applies a descriptive name to the SPAN session. |

The following example shows configuring a description of a SPAN session:

```
switch# configure terminal
switch(config)# monitor session 2
switch(config-monitor)# description monitoring ports fc2/2-fc2/4
```

## Activating a SPAN Session

The default is to keep the session state shut. You can open a session that duplicates packets from sources to destinations.

### SUMMARY STEPS

1. `switch(config)# no monitor session {all | session-number} shut`

### DETAILED STEPS

|        | Command or Action                                                                  | Purpose                                           |
|--------|------------------------------------------------------------------------------------|---------------------------------------------------|
| Step 1 | <code>switch(config)# no monitor session {all   <i>session-number</i>} shut</code> | Opens the specified SPAN session or all sessions. |

The following example shows activating a SPAN session:

```
switch(config)# no monitor session 3 shut
```

## Suspending a SPAN Session

The default is to keep the session state shut. You can suspend a SPAN session.

### SUMMARY STEPS

1. `switch(config)# monitor session {all | session-number} shut`

## DETAILED STEPS

|        | Command or Action                                                                 | Purpose                                              |
|--------|-----------------------------------------------------------------------------------|------------------------------------------------------|
| Step 1 | switch(config)# <b>monitor session</b> {all   <i>session-number</i> } <b>shut</b> | Suspends the specified SPAN session or all sessions. |

The following example shows suspending a SPAN session:

```
switch(config)# monitor session 3 shut
```

**Note**

The Cisco Nexus 5000 Series switch supports two active SPAN sessions. When you configure more than two SPAN sessions, the first two sessions are active. During startup, the order of active sessions is reversed; the last two sessions are active. For example, if you configured ten sessions 1 to 10 where 1 and 2 are active, after a reboot, sessions 9 and 10 will be active. To enable deterministic behavior, explicitly suspend the sessions 3 to 10 with the **monitor session *session-number* shut** command.

## Displaying SPAN Information

To display SPAN information, perform this task:

## SUMMARY STEPS

1. switch# **show monitor** [session {all | *session-number* | range *session-range*} [brief]]

## DETAILED STEPS

|        | Command or Action                                                                                         | Purpose                          |
|--------|-----------------------------------------------------------------------------------------------------------|----------------------------------|
| Step 1 | switch# <b>show monitor</b> [session {all   <i>session-number</i>   range <i>session-range</i> } [brief]] | Displays the SPAN configuration. |

This example shows how to display SPAN session information:

```
switch# show monitor
SESSION STATE REASON DESCRIPTION
----- -
2 up The session is up
3 down Session suspended
4 down No hardware resource
```

This example shows how to display SPAN session details:

```
switch# show monitor session 2
 session 2

type : local
state : up
source intf :
 rx : fc3/1
 tx : fc3/1
 both : fc3/1
source VLANs :
 rx :
source VSANs :
 rx : 1
destination ports : Eth3/1
```





# Troubleshooting

---

- [Troubleshooting, page 767](#)

## Troubleshooting

### Recovering a Lost Password

This section describes how to recover a lost network administrator password using the console port of the switch.

You can recover the network administrator password using one of two methods:

- From the CLI with a username that has network-admin privileges
- By power cycling the switch

#### Using the CLI with Network-Admin Privileges

If you are logged in to, or can log into, the switch with a username that has network-admin privileges, follow these steps:

#### SUMMARY STEPS

1. Verify that your username has network-admin privileges.
2. Assign a new network administrator password if your username has network-admin privileges.
3. Save the configuration.

#### DETAILED STEPS

---

**Step 1** Verify that your username has network-admin privileges.

**Example:**

```
switch# show user-account
user:admin
 this user account has no expiry date
 roles:network-admin
user:dbgusr
 this user account has no expiry date
 roles:network-admin network-operator
```

**Step 2** Assign a new network administrator password if your username has network-admin privileges.

**Example:**

```
switch# configure terminal
switch(config)# username admin password <new password>
switch(config)# exit
```

**Step 3** Save the configuration.

**Example:**

```
switch# copy running-config startup-config
```

### Power Cycling the Switch

If you cannot start a session on the switch that has network-admin privileges, you must recover the network administrator password by power cycling the switch.



**Caution** This procedure disrupts all traffic on the switch.



**Note** You cannot recover the administrator password from a Telnet or SSH session. You must have access to the local console connection.

To recover the network administrator password by power cycling the switch, follow these steps:

Establish a terminal session on the console port of the supervisor module.

### SUMMARY STEPS

1. Power cycle the switch.
2. Press the **Ctrl-J** key sequence from the console port session when the switch begins the Cisco NX-OS software boot sequence to enter the boot prompt mode.
3. Reset the network administrator password.
4. Display the bootflash: contents to locate the Cisco NX-OS software image file.
5. Load the Cisco NX-OS system software image.
6. Log in to the switch using the new administrator password.
7. Reset the new password to ensure that it is also the SNMP password.
8. Save the configuration.

## DETAILED STEPS

---

**Step 1** Power cycle the switch.

**Step 2** Press the **Ctrl-J** key sequence from the console port session when the switch begins the Cisco NX-OS software boot sequence to enter the boot prompt mode.

**Note** In releases of Cisco NX-OS prior to 4.0(1a) the key sequence to enter the boot prompt mode was **Ctrl-Shift-B**.

**Example:**

```
Ctrl-J
switch (boot) #
```

**Step 3** Reset the network administrator password.

**Example:**

```
switch (boot) # configure terminal
switch (boot-config) # admin-password <new password>
switch (boot-config) # exit
```

**Step 4** Display the bootflash: contents to locate the Cisco NX-OS software image file.

**Example:**

```
switch (boot) # dir bootflash:
```

**Step 5** Load the Cisco NX-OS system software image.

**Example:**

In the following example, the system image filename is nx-os.bin:

```
switch (boot) # load bootflash:nx-os.bin
```

**Step 6** Log in to the switch using the new administrator password.

**Example:**

```
switch login: admin
Password: <new password>
```

**Step 7** Reset the new password to ensure that it is also the SNMP password.

**Example:**

```
switch# configure terminal
switch (config) # username admin password <new password>
switch (config) # exit
```

**Step 8** Save the configuration.

**Example:**

```
switch# copy running-config startup-config
```

---

## Using Ethalyzer

Ethalyzer is a Cisco NX-OS protocol analyzer tool based on the Wireshark (formerly Ethereal) open source code. Ethalyzer is a command-line version of Wireshark that captures and decodes packets. You can use Ethalyzer to troubleshoot your network and analyze the control-plane traffic.

To configure Ethalyzer, use one or more of the following commands:

### SUMMARY STEPS

1. switch# **ethalyzer local interface** *interface*
2. switch# **ethalyzer local interface** *interface* **brief**
3. switch# **ethalyzer local interface** *interface* **limit-captured-frames**
4. switch# **ethalyzer local interface** *interface* **limit-frame-size**
5. switch# **ethalyzer local interface** *interface* **capture-filter**
6. switch# **ethalyzer local interface** *interface* **display-filter**
7. switch# **ethalyzer local interface** *interface* **write**
8. switch# **ethalyzer local read** *file*

### DETAILED STEPS

|               | Command or Action                                                                      | Purpose                                                                                                                                                                                                                                                                                     |
|---------------|----------------------------------------------------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Step 1</b> | switch# <b>ethalyzer local interface</b> <i>interface</i>                              | Captures packets sent or received by the supervisor and provides detailed protocol information.<br><br><b>Note</b> For all commands in this table, interface is inbound-hi (Inbound high-priority interface), inbound-low (Inbound low-priority interface), or mgmt (management interface). |
| <b>Step 2</b> | switch# <b>ethalyzer local interface</b> <i>interface</i> <b>brief</b>                 | Captures packets sent or received by the supervisor and provides a summary of protocol information.                                                                                                                                                                                         |
| <b>Step 3</b> | switch# <b>ethalyzer local interface</b> <i>interface</i> <b>limit-captured-frames</b> | Limits the number of frames to capture.                                                                                                                                                                                                                                                     |
| <b>Step 4</b> | switch# <b>ethalyzer local interface</b> <i>interface</i> <b>limit-frame-size</b>      | Limits the length of the frame to capture.                                                                                                                                                                                                                                                  |
| <b>Step 5</b> | switch# <b>ethalyzer local interface</b> <i>interface</i> <b>capture-filter</b>        | Filters the types of packets to capture.                                                                                                                                                                                                                                                    |
| <b>Step 6</b> | switch# <b>ethalyzer local interface</b> <i>interface</i> <b>display-filter</b>        | Filters the types of captured packets to display.                                                                                                                                                                                                                                           |
| <b>Step 7</b> | switch# <b>ethalyzer local interface</b> <i>interface</i> <b>write</b>                 | Saves the captured data to a file.                                                                                                                                                                                                                                                          |
| <b>Step 8</b> | switch# <b>ethalyzer local read</b> <i>file</i>                                        | Opens a captured data file and analyzes it.                                                                                                                                                                                                                                                 |

Ethalyzer does not capture data traffic that Cisco NX-OS forwards in the hardware.

Ethalyzer uses the same capture filter syntax as tcpdump. For more information, see the following URL: [http://www.tcpdump.org/tcpdump\\_man.html](http://www.tcpdump.org/tcpdump_man.html)

For information on the syntax of the display filter, see the following URL: <http://wiki.wireshark.org/DisplayFilters>

This example shows captured data (limited to four packets) on the management interface:

```
switch# ethanalyzer local interface mgmt brief limit-captured-frames 4
Capturing on eth0
2005-01-25 07:18:08.997132 10.193.24.42 -> 10.200.0.103 TELNET Telnet Data ...
2005-01-25 07:18:09.166266 10.200.0.103 -> 10.193.24.42 TCP 1235 > telnet [ACK] Seq=0 Ack=19
 Win=64129 Len=0
2005-01-25 07:18:09.166830 10.193.24.42 -> 10.200.0.103 TELNET Telnet Data ...
2005-01-25 07:18:09.376250 10.200.0.103 -> 10.193.24.42 TCP 1235 > telnet [ACK] Seq=0 Ack=99
 Win=64049 Len=0
4 packets captured
```

This example shows detailed captured data for one HSRP packet:

```
switch(config)# ethanalyzer local interface mgmt capture-filter "tcp port 23"
limit-captured-frames 1
Capturing on eth0
Frame 1 (60 bytes on wire, 60 bytes captured)
 Arrival Time: Jan 25, 2005 08:49:49.250719000
 [Time delta from previous captured frame: 1106642989.250719000 seconds]
 [Time delta from previous displayed frame: 1106642989.250719000 seconds]
 [Time since reference or first frame: 1106642989.250719000 seconds]
 Frame Number: 1
 Frame Length: 60 bytes
 Capture Length: 60 bytes
 [Frame is marked: False]
 [Protocols in frame: eth:ip:tcp]
Ethernet II, Src: 00:1a:a2:d2:d7:00 (00:1a:a2:d2:d7:00), Dst: 00:0d:ec:6d:81:00
(00:0d:ec:6d:81:00)
 Destination: 00:0d:ec:6d:81:00 (00:0d:ec:6d:81:00)
 Address: 00:0d:ec:6d:81:00 (00:0d:ec:6d:81:00)
 0. = IG bit: Individual address (unicast)
 0. = LG bit: Globally unique address (factory default)
 Source: 00:1a:a2:d2:d7:00 (00:1a:a2:d2:d7:00)
 Address: 00:1a:a2:d2:d7:00 (00:1a:a2:d2:d7:00)
 0. = IG bit: Individual address (unicast)
 0. = LG bit: Globally unique address (factory default)
 Type: IP (0x0800)
 Trailer: 000000000000
Internet Protocol, Src: 10.200.0.103 (10.200.0.103), Dst: 10.193.24.42 (10.193.24.42)
 Version: 4
 Header length: 20 bytes
 Differentiated Services Field: 0x00 (DSCP 0x00: Default; ECN: 0x00)
 0000 00.. = Differentiated Services Codepoint: Default (0x00)
 0. = ECN-Capable Transport (ECT): 0
 0. = ECN-CE: 0
 Total Length: 40
 Identification: 0xa651 (42577)
 Flags: 0x04 (Don't Fragment)
 0... = Reserved bit: Not set
 .1.. = Don't fragment: Set
 ..0. = More fragments: Not set
 Fragment offset: 0
 Time to live: 127
 Protocol: TCP (0x06)
 Header checksum: 0x2765 [correct]
 [Good: True]
 [Bad : False]
 Source: 10.200.0.103 (10.200.0.103)
 Destination: 10.193.24.42 (10.193.24.42)
Transmission Control Protocol, Src Port: 1288 (1288), Dst Port: telnet (23), Seq: 0, Ack:
0, Len: 0
 Source port: 1288 (1288)
 Destination port: telnet (23)
 Sequence number: 0 (relative sequence number)
 Acknowledgement number: 0 (relative ack number)
 Header length: 20 bytes
 Flags: 0x10 (ACK)
 0... = Congestion Window Reduced (CWR): Not set
 .0.. = ECN-Echo: Not set
 ..0. = Urgent: Not set
 ...1 = Acknowledgment: Set
 0... = Push: Not set
 0... = Reset: Not set
 0. = Syn: Not set
 0.0 = Fin: Not set
 Window size: 64334
 Checksum: 0x934f [correct]
 [Good Checksum: True]
 [Bad Checksum: False]
```

1 packets captured  
 For more information on Wireshark, see the following URL: <http://www.wireshark.org/docs/>.

## Troubleshooting Fibre Channel

### fctrace

The fctrace feature provides the following capabilities:

- Trace the route followed by data traffic.
- Compute inter-switch (hop-to-hop) latency.

You can invoke fctrace by providing the FC ID, the N port WWN, or the device alias of the destination.

The trace frame is routed normally through the network until it reaches the far edge of the fabric. When the frame reaches the edge of the fabric (the F port connected to the end node with the given port WWN or the FC ID), the frame is looped back (swapping the source ID and the destination ID) to the originator.

If the destination cannot be reached, the path discovery starts, which traces the path up to the point of failure.



#### Note

The fctrace feature works only on TE ports. Make sure that only TE ports exist in the path to the destination. If there is an E port in the path, the fctrace frame is dropped by that switch. Also, fctrace times out in the originator, and path discovery does not start.

To perform the fctrace operation, perform this task:

### SUMMARY STEPS

1. switch# **fctrace** {**device-alias** *aliasname* | **fcid** *fcid*} **vsan** *vsan-id* [**timeout** *seconds*]

### DETAILED STEPS

```
switch# fctrace {device-alias aliasname | fcid fcid} vsan vsan-id [timeout seconds]
```

The device-alias option specifies the device alias name. The fcid specifies the FCID of the destination N port, with the format 0xhhhhhh. The pwwn specifies the PWWN of the destination N port, with the format hh:hh:hh:hh:hh:hh:hh:hh. The vsan option specifies a VSAN ID.

**Note** By default the period to wait before a time out is 5 seconds and the range is from one through 10 seconds.

This example shows invoking fctrace for the specified FC ID of the destination N port:

```
switch# fctrace fcid 0xd70000 vsan 1
Route present for : 0xd70000

20:00:00:0b:46:00:02:82 (0xffffcd5)

Timestamp Invalid.
20:00:00:05:30:00:18:db (0xffffcd7)

Timestamp Invalid.
20:00:00:05:30:00:18:db (0xffffcd7)
```

This example shows invoking fctrace using the pWWN of the destination N port.

```
switch# fctrace pwn 21:00:00:e0:8b:06:d9:1d vsan 1 timeout 5
Route present for : 21:00:00:e0:8b:06:d9:1d
20:00:00:0b:46:00:02:82(0xffffcd5)

Timestamp Invalid.
20:00:00:05:30:00:18:db(0xffffcd7)
```

```
Timestamp Invalid.
20:00:00:05:30:00:18:db(0xffffcd7)
```

This example shows invoking fctrace using the device alias of the destination N port.

```
switch# fctrace device-alias disk1 vsan 1
Route present for : 22:00:00:0c:50:02:ce:f8
20:00:00:05:30:00:31:1e(0xffffca9)
```

## fcping

The fcping feature verifies reachability of a node by checking its end-to-end connectivity. You can invoke the fcping feature by providing the FC ID, the destination port WWN, or the device alias information.

To perform a fcping operation, perform this task:

### SUMMARY STEPS

1. switch# **fcping** {**device-alias** *aliasname* | **fcid** {*fc-port* | *domain-controller-id*} | **pwn** *pwn-id*} vsan *vsan-id* [[**count**] [ *number* ] [[**timeout**] [ *value* ] [[**usr-priority**] [ *priority* ]]]]

### DETAILED STEPS

---

```
switch# fcping {device-alias aliasname | fcid {fc-port | domain-controller-id} | pwn pwn-id} vsan vsan-id [[count] [number] [[timeout] [value] [[usr-priority] [priority]]]]
```

The device-alias option specifies the device alias name. The fcid specifies the FCID of the destination N port, with the format 0xhhhhhh. The domain-controller-id option verifies connection to the destination switch. The pwn specifies the PWWN of the destination N port, with the format hh:hh:hh:hh:hh:hh:hh:hh. The vsan option specifies a VSAN ID.

The last three are optional: The count option specifies the number of frames to send in a range of 0 to 2147483647. A value of 0 sends forever. By default, five frames are sent. The timeout option specifies the timeout value in seconds. The range is 1 to 10. The usr-priority option specifies the priority the frame receives in the switch fabric.

---

This example shows invoking fcping for the specified FCID of the destination:

```
switch# fcping fcid 0xd70000 vsan 1
28 bytes from 0xd70000 time = 730 usec
28 bytes from 0xd70000 time = 165 usec
28 bytes from 0xd70000 time = 262 usec
28 bytes from 0xd70000 time = 219 usec
28 bytes from 0xd70000 time = 228 usec
5 frames sent, 5 frames received, 0 timeouts
Round-trip min/avg/max = 165/270/730 usec
```

This example shows invoking `fcping` using the `count` option:

```
switch# fcping fcid 0xd70000 vsan 1 count 10
28 bytes from 0xd70000 time = 730 usec
28 bytes from 0xd70000 time = 165 usec
28 bytes from 0xd70000 time = 262 usec
28 bytes from 0xd70000 time = 219 usec
28 bytes from 0xd70000 time = 228 usec
28 bytes from 0xd70000 time = 230 usec
28 bytes from 0xd70000 time = 230 usec
28 bytes from 0xd70000 time = 225 usec
28 bytes from 0xd70000 time = 229 usec
28 bytes from 0xd70000 time = 183 usec
10 frames sent, 10 frames received, 0 timeouts
Round-trip min/avg/max = 165/270/730 usec
```

This example shows invoking `fcping` with a timeout value:

```
switch# fcping fcid 0xd500b4 vsan 1 timeout 10
28 bytes from 0xd500b4 time = 1345 usec
...
5 frames sent, 5 frames received, 0 timeouts
Round-trip min/avg/max = 340/581/1345 usec
```

This example shows invoking `fcping` for the specified device alias of the destination:

```
switch# fcping device-alias disk1 vsan 1
28 bytes from 22:00:00:0c:50:02:ce:f8 time = 1883 usec
28 bytes from 22:00:00:0c:50:02:ce:f8 time = 493 usec
28 bytes from 22:00:00:0c:50:02:ce:f8 time = 277 usec
28 bytes from 22:00:00:0c:50:02:ce:f8 time = 391 usec
28 bytes from 22:00:00:0c:50:02:ce:f8 time = 319 usec
5 frames sent, 5 frames received, 0 timeouts
Round-trip min/avg/max = 277/672/1883 usec
```

This example shows invoking the `fcping` command when there is resource exhaustion at the N port:

```
switch# fcping fcid 0x010203 vsan 1
No response from the N port.
switch# fcping pwnn 21:00:00:20:37:6f:db:dd vsan 1
28 bytes from 21:00:00:20:37:6f:db:dd time = 1454 usec
...
5 frames sent, 5 frames received, 0 timeouts
Round-trip min/avg/max = 364/784/1454 usec
```



**Note**

The command returns a "No response from the N port" message even when the N port is active. Retry the command a few seconds later.

### Verifying Switch Connectivity

You can verify connectivity to a destination switch.



**Note**

The FC ID variable used in this procedure is the domain controller address; it is not a duplication of the domain ID.

To verify connectivity to a destination switch, perform this task:

## SUMMARY STEPS

1. switch# **show fcdomain domain-list vsan 200 0xda(218) 20:c8:00:05:30:00:87:9f [Local]**
2. switch# **fcping fcid 0xFFFCDA vsan 200**

DETAILED STEPS

|        | Command or Action                                                                                                                                                                                                                                                                                                                                                                                                                   | Purpose                                                                                                                                                                                                                |
|--------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Step 1 | <pre>switch# show fcdomain domain-list vsan 200 0xda(218) 20:c8:00:05:30:00:87:9f [Local]  Example: Number of domains: 7 Domain ID          WWN ----- 0x01(1)           20:c8:00:05:30:00:59:df [Principal] 0x02(2)           20:c8:00:0b:5f:d5:9f:c1 0x6f(111)         20:c8:00:05:30:00:60:df 0x06(6)           20:c8:00:0b:46:79:f2:41 0x04(4)           20:c8:00:05:30:00:86:5f 0x6a(106)         20:c8:00:05:30:00:f8:e3</pre> | <p>Displays the destination switch's domain ID.</p> <p>To obtain the domain controller address, concatenate the domain ID with FFFC. For example, if the domain ID is 0xda(218), the concatenated ID is 0xffffcda.</p> |
| Step 2 | <pre>switch# fcping fcid 0xFFFFCDA vsan 200  Example: 28 bytes from 0xFFFFCDA time = 298 usec 28 bytes from 0xFFFFCDA time = 260 usec 28 bytes from 0xFFFFCDA time = 298 usec 28 bytes from 0xFFFFCDA time = 294 usec 28 bytes from 0xFFFFCDA time = 292 usec 5 frames sent, 5 frames received, 0 timeouts Round-trip min/avg/max = 260/288/298 usec</pre>                                                                          | <p>Verifies reachability of the destination switch by checking its end-to-end connectivity.</p>                                                                                                                        |

show tech-support Command

The **show tech-support** command is useful when collecting a large amount of information about the switch for troubleshooting purposes. The output of this command can be provided to technical support representatives when reporting a problem.

The **show tech-support** command displays the output of several **show** commands at once. The output from this command varies depending on your configuration. Use the **show tech-support** command in EXEC mode to display general information about the switch when reporting a problem.

You can choose to have detailed information for each command. You can specify the output for a particular interface, module, or VSAN. Each command output is separated by line and the command precedes the output.



Note

Explicitly set the **terminal length** command to 0 (zero) to disable auto-scrolling and enable manual scrolling. Use the **show terminal** command to view the configured the terminal size. After obtaining the output of this command, remember to reset your terminal length as required.

You can save the output of this command to a file by appending > (left arrow) and the filename to the **show tech-support** command. If you save this file, verify you have sufficient space to do so—each of these files may take about 1.8 MB. However, you can zip this file using the **gzip filename** command. Copy the zipped file to the required location using the **copy** command and unzip the file using the **gunzip** command.

The default output of the **show tech-support** command includes the output of the following commands:

- **show switchname**
- **show system uptime**
- **show interface mgmt0**
- **show interface mgmt1**
- **show system resources**
- **show version**
- **dir bootflash:**
- **show inventory**
- **show diagnostic result all**
- **show logging log**
- **show module**
- **show environment**
- **show sprom backplane**
- **show clock**
- **show callhome**
- **show cfs application**
- **show cfs lock**
- **show snmp**
- **show interface brief**
- **show interface**
- **show running-config**
- **show startup-config**
- **show ip route**
- **show arp**
- **show monitor session all**
- **show accounting log**
- **show process**
- **show process cpu**
- **show process log**
- **show process memory**
- **show processes log details**
- **show logging log**
- **show license host-id**

- **show license**
- **show license usage**
- **show system reset-reason**
- **show logging nvram**
- **show install all status**
- **show install all failure-reason**
- **show system internal log install**
- **show system internal log install details**
- **show cores**
- **show topology**
- **show kernel internal aipc**
- **show tech-support acl**
- **show vlan**
- **show vlan access-map**
- **show mac-address-table**
- **show spanning-tree summary**
- **show spanning-tree active**
- **show interface trunk**
- **show aclmgr status**
- **show aclmgr internal dictionaries**
- **show aclmgr internal log**
- **show aclmgr internal ppf**
- **show aclmgr internal state-cache**
- **show access-lists**
- **show platform software ethpm internal info all**
- **show object-group**
- **show logging onboard obfl-logs**

### **show tech-support brief Command**

Use the **show tech-support brief** command to obtain a quick, condensed review of the switch configurations. This command provides a summary of the current running state of the switch (see the following example).

The **show tech-support brief** command is useful when collecting information about the switch for troubleshooting purposes. The output of this command can be provided to technical support representatives when reporting a problem.

You can save the output of this command to a file by appending > (left arrow) and the filename to the **show tech-support brief** command.

This example shows how to display a condensed view of the switch configurations:

```
switch# show tech-support brief
Switch Name : switch
Switch Type :
Kickstart Image : 4.0(0) bootflash:///nuova-or-kickstart-nsg.4.0.0.001.bin
System Image : 4.0(0) bootflash:/nuova-or-system-nsg.4.0.0.001.binms-or-47
IP Address/Mask : 172.16.24.47/24
Switch WWN : 20:00:00:0d:ec:6b:cd:c0
No of VSANs : 1
Configured VSANs : 1
VSAN 1: name:VSAN0001, state:active, interop mode:default
 domain id:0xa6(166), WWN:20:01:00:0d:ec:6b:cd:c1 [Principal]
 active-zone:<NONE>, default-zone:deny
```

| Interface | Vsan | Admin Mode | Admin Trunk Mode | Status    | SFP | Oper Mode | Oper Speed (Gbps) | Port Channel |
|-----------|------|------------|------------------|-----------|-----|-----------|-------------------|--------------|
| fc3/1     | 1    | auto       | on               | down      | swl | --        | --                | --           |
| fc3/2     | 1    | auto       | on               | sfpAbsent | --  | --        | --                | --           |
| fc3/3     | 1    | auto       | on               | down      | swl | --        | --                | --           |
| fc3/4     | 1    | auto       | on               | sfpAbsent | --  | --        | --                | --           |
| fc3/5     | 1    | auto       | on               | down      | swl | --        | --                | --           |
| fc3/6     | 1    | auto       | on               | sfpAbsent | --  | --        | --                | --           |
| fc3/7     | 1    | auto       | on               | down      | swl | --        | --                | --           |
| fc3/8     | 1    | auto       | on               | down      | swl | --        | --                | --           |

| Interface    | Status     | IP Address | Speed | MTU  | Port Channel |
|--------------|------------|------------|-------|------|--------------|
| Ethernet1/1  | sfpIsAbsen | --         | --    | 1500 | --           |
| Ethernet1/2  | sfpIsAbsen | --         | --    | 1500 | --           |
| Ethernet1/3  | up         | --         | 10000 | 1500 | --           |
| Ethernet1/4  | sfpIsAbsen | --         | --    | 1500 | --           |
| Ethernet1/5  | sfpIsAbsen | --         | --    | 1500 | --           |
| Ethernet1/6  | sfpIsAbsen | --         | --    | 1500 | --           |
| Ethernet1/7  | sfpIsAbsen | --         | --    | 1500 | --           |
| Ethernet1/8  | sfpIsAbsen | --         | --    | 1500 | --           |
| Ethernet1/9  | sfpIsAbsen | --         | --    | 1500 | --           |
| Ethernet1/10 | sfpIsAbsen | --         | --    | 1500 | --           |
| Ethernet1/11 | sfpIsAbsen | --         | --    | 1500 | --           |
| Ethernet1/12 | sfpIsAbsen | --         | --    | 1500 | --           |
| Ethernet1/13 | sfpIsAbsen | --         | --    | 1500 | --           |
| Ethernet1/14 | sfpIsAbsen | --         | --    | 1500 | --           |
| Ethernet1/15 | notConnect | --         | --    | 1500 | --           |
| Ethernet1/16 | sfpIsAbsen | --         | --    | 1500 | --           |
| Ethernet1/17 | sfpIsAbsen | --         | --    | 1500 | --           |
| Ethernet1/18 | sfpIsAbsen | --         | --    | 1500 | --           |
| Ethernet1/19 | notConnect | --         | --    | 1500 | --           |
| Ethernet1/20 | sfpIsAbsen | --         | --    | 1500 | --           |
| Ethernet1/21 | sfpIsAbsen | --         | --    | 1500 | --           |
| Ethernet1/22 | sfpIsAbsen | --         | --    | 1500 | --           |
| Ethernet1/23 | sfpIsAbsen | --         | --    | 1500 | --           |
| Ethernet1/24 | sfpIsAbsen | --         | --    | 1500 | --           |
| Ethernet1/25 | sfpIsAbsen | --         | --    | 1500 | --           |
| Ethernet1/26 | sfpIsAbsen | --         | --    | 1500 | --           |
| Ethernet1/27 | sfpIsAbsen | --         | --    | 1500 | --           |
| Ethernet1/28 | sfpIsAbsen | --         | --    | 1500 | --           |
| Ethernet1/29 | sfpIsAbsen | --         | --    | 1500 | --           |
| Ethernet1/30 | sfpIsAbsen | --         | --    | 1500 | --           |
| Ethernet1/31 | sfpIsAbsen | --         | --    | 1500 | --           |
| Ethernet1/32 | sfpIsAbsen | --         | --    | 1500 | --           |
| Ethernet1/33 | sfpIsAbsen | --         | --    | 1500 | --           |
| Ethernet1/34 | sfpIsAbsen | --         | --    | 1500 | --           |
| Ethernet1/35 | up         | --         | 10000 | 1500 | --           |
| Ethernet1/36 | sfpIsAbsen | --         | --    | 1500 | --           |
| Ethernet1/37 | sfpIsAbsen | --         | --    | 1500 | --           |
| Ethernet1/38 | sfpIsAbsen | --         | --    | 1500 | --           |
| Ethernet1/39 | sfpIsAbsen | --         | --    | 1500 | --           |
| Ethernet1/40 | sfpIsAbsen | --         | --    | 1500 | --           |

| Interface | Status | IP Address | Speed | MTU |
|-----------|--------|------------|-------|-----|
|-----------|--------|------------|-------|-----|

```

mgmt0 up 172.16.24.47 100 1500
```

### show tech-support fc Command

Use the **show tech-support fc** command to obtain information about the FC configuration on your switch.

The output of the **show tech-support fc** command includes the output of the following commands:

- **show interface brief**
- **show interface**
- **show port internal info all**
- **show port internal event-history lock**
- **show port internal event-history msgs**
- **show port internal event-history errors**
- **show port internal mem-stats detail**
- **show san-port-channel internal event-history all**
- **show san-port-channel internal event-history errors**
- **show san-port-channel internal event-history msgs**
- **show san-port-channel internal event-history lock**
- **show san-port-channel internal mem-stats detail**
- **show san-port-channel usage**
- **show san-port-channel summary**
- **show san-port-channel consistency detail**
- **show tech-support device-alias**
- **show fcdomain domain-list**
- **show tech-support fcns**
- **show fcns database vsan 1-4093**
- **show fcns database detail vsan 1-4093**
- **show fcns database local vsan 1-4093**
- **show fcns database local detail vsan 1-4093**
- **show fcns statistics vsan 1-4093**
- **show fcns statistics detail vsan 1-4093**
- **show fcns internal info vsan 1-4093**
- **show fcns internal event-history**
- **show fcns internal event-log**
- **show fcroute unicast**

- **show fcs database**
- **show fcs ie**
- **show fetimer**
- **show flogi database**
- **show flogi internal info**
- **show fspf**
- **show fspf database**
- **show tech-support rscn**
- **show rscn internal vsan 1-4093**
- **show rscn internal event-history**
- **show rscn internal mem-stats detail**
- **show rscn internal session-history vsan 1-4093**
- **show rscn internal merge-history vsan 1-4093**
- **show rscn statistics vsan 1-4093**
- **show rscn scr-table vsan 1-4093**
- **show rscn session status vsan 1-4093**
- **show vsan**
- **show vsan membership**
- **show tech-support zone**
- **show zone status vsan 1-4093**
- **show zoneset active vsan 1-4093**
- **show zoneset vsan 1-4093**
- **show zone vsan 1-4093**
- **show fcalias vsan 1-4093**
- **show zone-attribute-group vsan 1-4093**
- **show zone policy vsan 1-4093**
- **show zoneset pending active vsan 1-4093**
- **show zoneset pending vsan 1-4093**
- **show zone pending vsan 1-4093**
- **show zone pending active vsan 1-4093**
- **show fcalias pending vsan 1-4093**
- **show zone policy pending vsan 1-4093**
- **show zone pending-diff vsan 1-4093**

- **show zone analysis active vsan 1-4093**
- **show zone analysis vsan 1-4093**
- **show zone ess vsan 1-4093**
- **show zone internal vsan 1-4093**
- **show zone internal change event-history vsan 1-4093**
- **show zone internal ifindex-table vsan 1-4093**
- **show zone internal merge event-history vsan 1-4093**
- **show zone internal event-history**
- **show zone internal event-history errors**
- **show zone internal tcam event-history vsan 1-4093**
- **show zone statistics vsan 1-4093**
- **show system default zone**
- **show zone internal ddas-table**
- **show zone internal sdv-table vsan 1-4093**
- **show zone internal mem-stats**
- **show zone internal mem-stats detail**
- **show zone internal transit-table received vsan 1-4093**
- **show zone internal transit-table forwarded vsan 1-4093**
- **show zone internal transit-table rejected vsan 1-4093**

You can save the output of this command to a file by appending > (left arrow) and the filename to the **show tech-support zone** command.

### **show tech-support platform Command**

Use the **show tech-support platform** command to obtain information about the platform configuration of your switch.

The output of the **show tech-support platform** command includes the output of the following commands:

- **show platform fwm mem-stats detail**
- **show platform fwm info global**
- **show platform fwm info pif all verbose**
- **show platform fwm info lif all verbose**
- **show platform fwm info vlan all verbose**
- **show platform fwm info error stats**
- **show platform fwm info error history**
- **show platform fwm info stm-stats**

- **show platform fwm info pc all verbose**
- **show platform fwm info ppf**
- **show platform fwm info pss all**
- **show platform hardware fwm info vlan all**
- **show platform hardware fwm info pif all**
- **show platform hardware fwm info lif all**
- **show platform hardware fwm info global**
- **show platform software zschk internal info**
- **show platform software zschk internal msgs**
- **show platform software statsclient msgs**
- **show hardware internal gatos detail**
- **show hardware internal gatos all-ports detail**
- **show hardware internal altos detail**
- **show hardware internal altos event-history errors**
- **show hardware internal altos event-history messages**
- **show platform fcfib fcflow**
- **show platform fcfib event-history all**
- **show platform fcfib unicasts**
- **show platform fcfib unicasts forwarding-configuration**
- **show platform fcfib vsan**
- **show platform fcfib san-port-channel**
- **show platform software fcfib devices**
- **show platform software fcfib multipath**
- **show platform software fcfib vsanidxtable**
- **show platform software fcfib domainidxtable**
- **show platform hardware fcfib pathselecttable**
- **show platform hardware fcfib pathselecttable all**
- **show platform software fcfib fctable-check**
- **show fc2 internal event-history errors**
- **show system internal liod liod\_db**
- **show system internal liod queues**
- **show system internal liod state**
- **show system internal liod time\_db**

- **show system internal rib domain**
- **show system internal rib system-attributes**
- **show system internal rib unicast**
- **show system internal rib vsan-attributes**
- **show system internal fcfwd fwidxmap if\_index**
- **show system internal fcfwd idxmap interface-to-port**
- **show system internal fcfwd pemap**
- **show platform afm info global**
- **show platform afm info attachment brief**
- **show platform afm info group-cfg all**
- **show platform afm info lop all**
- **show platform software altos detail**
- **show platform software altos event-history errors**
- **show platform software altos event-history msgs**
- **show platform software altos ports all**
- **show platform hardware altos counters all**
- **show platform hardware altos counters interrupts all**
- **show platform hardware altos interrupts all detail**

## Default Settings for Troubleshooting Features

The following table lists the default settings for the features included in this chapter.

**Table 100: Default Settings for Troubleshooting Features**

| <b>Parameters</b>                          | <b>Default</b> |
|--------------------------------------------|----------------|
| Timeout period to invoke fctrace           | 5 seconds      |
| Number of frame sent by the fcping feature | 5 frames       |
| Remote capture connection protocol         | TCP            |
| Remote capture connection mode             | Passive        |
| Local capture frame limits                 | 10 frames      |
| FC ID allocation mode                      | Auto mode      |





## APPENDIX A

# Configuration Limits

- [Cisco Nexus 5000 Series Configuration Limits, page 787](#)

## Cisco Nexus 5000 Series Configuration Limits

The features supported by the Cisco Nexus 5000 Series switch have maximum configuration limits. Some of these limits apply only when one or more Cisco Nexus 2000 Series Fabric Extender units are attached to the switch.

The following tables list the Cisco verified limits for Cisco Nexus 5000 Series switches running Cisco NX-OS Release 4.1.x.

**Table 101: Ethernet Environments**

| Parameter                                    | Limit                                                            |
|----------------------------------------------|------------------------------------------------------------------|
| Active VLANs/VSANs per switch                | 512. 31 are set aside for VSANs and the remaining are for VLANs. |
| VLAN/VSAN ID Space                           | 4,096                                                            |
| STP Logical Interfaces                       | 3,600 <sup>3</sup>                                               |
| MST Instances per bridge topology per switch | 64 (IEEE Standard)                                               |
| VLAN ACLs (VACLs) per switch                 | 1,024                                                            |
| Port ACLs (PACLs) per switch                 | 256 <sup>4</sup>                                                 |
| ACL Accounting                               | 32                                                               |
| Member interfaces per EtherChannel           | 16                                                               |

<sup>3</sup> The 3600 logical interfaces apply to both PVRST and MST. If MST is used, the scale applies to the maximum number of MST instances (64).

<sup>4</sup> There can be a maximum of 50 ACEs per PAACL.

**Table 102: Fibre Channel Environments**

| Parameter                                                 | Limit                                        |
|-----------------------------------------------------------|----------------------------------------------|
| Device Aliases per fabric                                 | 8,000                                        |
| Switches per physical fabric or VSAN                      | 50 <sup>5</sup>                              |
| Domains per VSAN                                          | 40 <sup>6</sup>                              |
| Native FC Links per switch                                | 16—Requires two N5K-M1008 expansion modules. |
| FLOGIs or FDISCs per NPV port group                       | 255                                          |
| Zones per virtual or physical F port (includes all VSANs) | 32                                           |
| Zone sets per switch (includes all VSANs)                 | 500 <sup>7</sup>                             |
| Zone members per physical fabric (includes all VSANs)     | 8,000 <sup>8</sup>                           |
| Zones per switch (includes all VSANs)                     | 8,000                                        |
| Maximum diameter of a SAN Fabric                          | 3 hops <sup>9</sup>                          |
| FSPF interface instances per switch                       | 512 <sup>10</sup>                            |
| ISL instances per switch                                  | 256 <sup>11</sup>                            |
| Virtual Fibre Channel interfaces                          | 160 <sup>12</sup>                            |
| Max FCIDs allocated                                       | 2,048                                        |
| Fibre Channel Flows                                       | 32                                           |

<sup>5</sup> The switch is capable of supporting up to 239 switches per fabric but this is not currently implemented.

<sup>6</sup> The switch is capable of supporting up to 239 domains per VSAN but the scale is not currently implemented.

<sup>7</sup> The switch is capable of 1000 Zone Sets but this is not currently implemented.

<sup>8</sup> The switch is capable of supporting up to 20,000 zone members per fabric but the scale is not currently implemented.

<sup>9</sup> The switch is capable of supporting up to 12 hops but the scale is not currently implemented.

<sup>10</sup> The switch is capable of supporting 4096 (the number of Extended ISLs (16) times the number of VSANs (256)) but the scale is not currently implemented.

<sup>11</sup> Each ISL instance can support up to 32 VSAN instances.

<sup>12</sup> The switch is capable of supporting 480 virtual Fibre Channel interfaces but the scale is not currently implemented.

**Table 103: General Parameters**

| Parameter                                                                                           | Limit                                                                            |
|-----------------------------------------------------------------------------------------------------|----------------------------------------------------------------------------------|
| Maximum Fabric Extenders per Cisco Nexus 5000 Series switch                                         | 12 units <sup>13</sup>                                                           |
| Maximum Fabric Extenders dual-homed to a vPC Cisco Nexus 5000 Series switch pair                    | 12 units <sup>14</sup>                                                           |
| Maximum number of hosts connected to Fabric Extenders connected to Cisco Nexus 5000 Series switches | 480 hosts <sup>15</sup>                                                          |
| MAC Table Size                                                                                      | 16,000 entries <sup>16</sup>                                                     |
| Event Traps - forward via Email                                                                     | 4 destinations <sup>17</sup>                                                     |
| QoS System Classes                                                                                  | 5 all user-configurable classes                                                  |
| Port channels                                                                                       | 4 SAN port channels and 16 EtherChannels (with the combination not exceeding 16) |
| SPAN Sessions                                                                                       | 2 active sessions                                                                |
| Egress SPAN sources                                                                                 | 2                                                                                |

- <sup>13</sup> The number of Fabric Extenders is limited by the total number hosts connected through the single homed Fabric Extenders to the parent Cisco Nexus 5000 Series switch. This applies both in the case where the parent switch is running vPC or not.
- <sup>14</sup> The number of Fabric Extenders is limited by the total number hosts connected through dual homed Fabric Extenders to a pair of Cisco Nexus 5000 Series switches running vPC. In dual homed Fabric Extender vPC mode, the Fabric Extenders provide a redundant network path to the hosts in case of a direct failure in the primary path.
- <sup>15</sup> The limit of 480 applies equally to single homed and dual homed Fabric Extender topologies. In single homed topologies, the maximum number of Fabric Extenders is 12 which implies up to 40 hosts may be connected to each Fabric Extender for a total of 480 hosts per Cisco Nexus 5000 Series switch. In a dual homed Fabric Extender topology, the total number of hosts supported is also 480.
- <sup>16</sup> 1,000 entries are reserved multicast MAC addresses and the remainder are for unicast MAC addresses.
- <sup>17</sup> The switch is capable of up to 50 different destinations but this is not currently implemented.





## INDEX

\* (asterisk) [586](#)

first operational port[asterisk (asterisk)] [586](#)

first operational port] [586](#)

### A

AAA [243, 244, 246, 247, 248, 251, 256, 708](#)

accounting [243](#)

authentication [243](#)

benefits [244](#)

configuring console login [248](#)

default settings [256](#)

DHCHAP authentication [708](#)

enabling MSCHAP authentication [251](#)

example configuration [256](#)

guidelines [248](#)

limitations [248](#)

prerequisites [247](#)

user login process [246](#)

verifying configurations [256](#)

AAA accounting [253](#)

configuring default methods [253](#)

AAA accounting logs [255](#)

clearing [255](#)

displaying [255](#)

AAA logins [251](#)

enabling authentication failure messages [251](#)

AAA protocols [243](#)

RADIUS [243](#)

TACACS+ [243](#)

AAA server groups [245](#)

description [245](#)

AAA servers [253, 255](#)

specifying SNMPv3 parameters [253, 255](#)

specifying user roles [255](#)

specifying user roles in VSAs [253](#)

AAA services [244, 245](#)

configuration options [245](#)

remote [244](#)

accounting [243](#)

description [243](#)

active zone sets [606, 616](#)

considerations [606](#)

enabling distribution [616](#)

address allocation cache [551](#)

description [551](#)

administrative speeds [521](#)

configuring [521](#)

administrative states [514](#)

description [514](#)

authentication [243, 245, 246, 699](#)

description [243](#)

fabric security [699](#)

local [243](#)

methods [245](#)

remote [243](#)

user login [246](#)

authorization [246](#)

user login [246](#)

auto mode [519](#)

configuring [519](#)

auto port mode [514](#)

description [514](#)

autosensing speed [521](#)

### B

BB\_credits [517, 528](#)

description [517](#)

displaying information [528](#)

reason codes [517](#)

bit error thresholds [522](#)

configuring [522](#)

description [522](#)

bit errors [522](#)

reasons [522](#)

blocking state, STP [166](#)

BPDU guard [212](#)

bridge ID [156](#)

broadcast storms [237](#)

Brocade [688](#)

native interop mode [688](#)

buffer-to-buffer credits [517](#)  
 build fabric frames [532](#)  
   description [532](#)

## C

call home [389](#)  
   smart call home feature [389](#)  
 Call Home [385](#)  
   description [385](#)  
   message format options [385](#)  
 Call Home messages [385, 388](#)  
   configuring levels [388](#)  
   format options [385](#)  
 call home notifications [409](#)  
   full-txt format for syslog [409](#)  
   XML format for syslog [409](#)  
 Cisco [254, 261](#)  
   vendor ID [254, 261](#)  
 cisco-av-pair [253, 255](#)  
   specifying AAA user parameters [253, 255](#)  
 CIST regional root [186](#)  
 CIST root [187](#)  
 community ports [86](#)  
 community VLANs [86, 87](#)  
 company IDs [686](#)  
   FC ID allocations [686](#)  
 configuring LACP [121](#)  
 configuring NPV [559](#)  
 Contiguous Domain ID Assignments [545](#)  
   About [545](#)

## D

dead time intervals [650](#)  
   configuring for FSPF [650](#)  
   description [650](#)  
 debounce timer [65](#)  
 default settings [256, 362](#)  
   AAA [256](#)  
   rollback [362](#)  
 default users [38](#)  
   description [38](#)  
 default VSANs [596](#)  
   description [596](#)  
 default zones [612, 688](#)  
   description [612](#)  
   interoperability [688](#)  
   policies [612](#)  
 destination IDs [572, 599, 654](#)  
   exchange based [572](#)

destination IDs (*continued*)  
   flow based [572](#)  
   in-order delivery [654](#)  
   path selection [599](#)  
 device alias databases [636, 637, 639](#)  
   disabling distribution [637](#)  
   discarding changes [636](#)  
   enabling distribution [637](#)  
   locking the fabric [636](#)  
   merging [639](#)  
 device aliases [631, 632, 633, 634, 638, 640, 641](#)  
   comparison with zones [632](#)  
   creating [633](#)  
   default settings [641](#)  
   description [631](#)  
   displaying information [640](#)  
   displaying zone set information [640](#)  
   enhanced mode [634](#)  
   features [631](#)  
   modifying databases [633](#)  
   requirements [632](#)  
   zone alias conversion [638](#)  
 device IDs [402](#)  
   call home format [402](#)  
 DHCPAP [699, 700, 701, 702, 704, 705, 708, 710](#)  
   AAA authentication [708](#)  
   authentication modes [702](#)  
   compatibility with other NX-OS features [701](#)  
   configuring [700](#)  
   configuring AAA authentication [708](#)  
   default settings [710](#)  
   description [700](#)  
   displaying security information [708](#)  
   enabling [701](#)  
   group settings [705](#)  
   hash algorithms [704](#)  
   passwords for local switches [705](#)  
   sample configuration [708](#)  
 diagnostics [363, 364, 365, 367](#)  
   configuring [365](#)  
   default settings [367](#)  
   expansion modules [365](#)  
   health monitoring [364](#)  
   runtime [363](#)  
 Diffie-Hellman Challenge Handshake Authentication Protocol [699](#)  
 documentation [xlv, xlvi](#)  
   additional publications [xlv](#)  
   obtaining [xlvi](#)  
 domain IDs [515, 531, 538, 541, 542, 545, 546, 613, 688](#)  
   allowed lists [541](#)  
   assignment failures [515](#)  
   configuring allowed lists [541](#)  
   configuring CFS distribution [542](#)  
   configuring fcalias members [613](#)

- domain IDs (*continued*)
    - contiguous assignments [545](#)
    - description [538](#)
    - distributing [531](#)
    - enabling contiguous assignments [545, 546](#)
    - interoperability [688](#)
    - preferred [538](#)
    - static [538](#)
  - domain manager [515, 533](#)
    - fast restart feature [533](#)
    - isolation [515](#)
  - drop latency time [657, 658](#)
    - configuring [657](#)
    - configuring for FSPF in-order delivery [657](#)
    - displaying information [658](#)
- ## E
- E port mode [513](#)
    - classes of service [513](#)
    - description [513](#)
  - E ports [515, 519, 565, 618, 643, 735, 743](#)
    - configuring [519](#)
    - fabric binding checking [735](#)
    - FCS support [743](#)
    - FSPF topologies [643](#)
    - isolation [515](#)
    - recovering from link isolations [618](#)
    - trunking configuration [565](#)
  - e-mail notifications [385](#)
    - Call Home [385](#)
  - EFMD [735, 736, 741](#)
    - displaying statistics [741](#)
    - fabric binding [735](#)
    - fabric binding initiation [736](#)
  - EISLs [571](#)
    - SAN port channel links [571](#)
  - ELP [515](#)
  - enabling NPV [558](#)
  - enhanced zones [622, 623, 625, 628](#)
    - advantages over basic zones [622](#)
    - changing from basic zones [623](#)
    - configuring default full database distribution [628](#)
    - configuring default policies [628](#)
    - configuring default switch-wide zone policies [628](#)
    - description [622](#)
    - modifying database [625](#)
  - ethanalyzer [769](#)
  - EtherChannel [111](#)
    - STP [111](#)
  - examples [256](#)
    - AAA configurations [256](#)
  - Exchange Fabric Membership Data [735](#)
  - exchange IDs [599, 654, 773](#)
    - in-order delivery [654](#)
    - load balancing [773](#)
    - path selection [599](#)
  - exchange link parameter [515](#)
  - executing a session [361](#)
  - expansion port mode [513](#)
  - extended range VLANs [78](#)
- ## F
- F port mode [513](#)
    - classes of service [513](#)
    - description [513](#)
  - F ports [513, 519](#)
    - configuring [519](#)
    - description [513](#)
  - fabric binding [701, 735, 736, 737, 739, 740, 742](#)
    - checking for E ports [735](#)
    - checking for TE ports [735](#)
    - clearing statistics [740](#)
    - compatibility with DHCPAP [701](#)
    - copying to config database [739](#)
    - copying to configuration file (procedure) [740](#)
    - creating config database (procedure) [740](#)
    - default settings [742](#)
    - deleting databases [740](#)
    - deleting from config database (procedure) [740](#)
    - description [735](#)
    - disabling [737](#)
    - EFMD [735](#)
    - enabling [737](#)
    - enforcement [736](#)
    - forceful activation [739](#)
    - forceful deactivation [739](#)
    - initiation process [736](#)
    - licensing requirements [735](#)
    - port security comparison [735](#)
    - saving to config database [739](#)
    - verifying status [737](#)
    - viewing active databases (procedure) [740](#)
    - viewing EFMD statistics (procedure) [740](#)
    - viewing violations (procedure) [740](#)
  - Fabric Configuration Servers [743](#)
  - fabric login [663](#)
  - fabric port mode [513](#)
  - fabric pWWNs [603](#)
    - zone membership [603](#)
  - fabric reconfiguration [531](#)
    - fcdomain phase [531](#)

- fabric security [699, 710](#)
    - authentication [699](#)
    - default settings [710](#)
  - Fabric Shortest Path First [643](#)
    - routing services [643](#)
  - Fabric-Device Management Interface [666](#)
  - fabrics [532](#)
  - fault tolerant fabrics [644](#)
    - example (figure) [644](#)
  - FC IDs [531, 545, 546, 613, 686](#)
    - allocating [531](#)
    - allocating default company ID lists [686](#)
    - configuring fcalias members [613](#)
    - description [545](#)
    - persistent [546](#)
  - FC-SP [699, 701, 708](#)
    - authentication [699](#)
    - enabling [701](#)
    - enabling on ISLs [708](#)
  - fcaliases [613, 620, 621](#)
    - cloning [621](#)
    - configuring for zones [613](#)
    - creating [613](#)
    - renaming [620](#)
  - fcdomains [515, 531, 533, 534, 535, 536, 537, 538, 542, 551, 552](#)
    - autoreconfigured merged fabrics [537](#)
    - configuring CFS distribution [542](#)
    - default settings [552](#)
    - description [531](#)
    - disabling [535](#)
    - displaying information [551](#)
    - domain IDs [538](#)
    - domain manager fast restart [533](#)
    - displaying statistics [551](#)
    - enabling [535](#)
    - enabling autoreconfiguration [537](#)
    - incoming RCFs [536](#)
    - initiation [535](#)
    - overlap isolation [515](#)
    - restarts [531](#)
    - switch priorities [534](#)
  - FCoE [3, 449](#)
    - disable LAN traffic [449](#)
  - fcping [775, 785](#)
    - default settings [785](#)
    - verifying switch connectivity [775](#)
  - FCSs [743, 744, 745](#)
    - characteristics [743](#)
    - configuring names [744](#)
    - default settings [745](#)
    - description [743](#)
    - displaying information [745](#)
  - ftimers [684](#)
    - displaying configured values [684](#)
  - fttrace [773, 785](#)
    - default settings [785](#)
    - invoking [773](#)
  - FDMI [666](#)
    - description [666](#)
    - displaying database information [666](#)
  - Fibre Channel [679, 738](#)
    - sWWNs for fabric binding [738](#)
    - timeout values [679](#)
    - TOV [679](#)
  - Fibre Channel domains [531](#)
  - Fibre Channel interfaces [514, 515, 517, 518, 519, 520, 521, 522, 528, 596](#)
    - administrative states [514](#)
    - BB\_credits [517](#)
    - configuring [518](#)
    - configuring auto port mode [519](#)
    - configuring bit error thresholds [522](#)
    - configuring descriptions [520](#)
    - configuring frame encapsulation [521](#)
    - configuring port modes [519](#)
    - configuring range [518](#)
    - configuring speeds [521](#)
    - default settings [528](#)
    - displaying VSAN membership [596](#)
    - operational states [514](#)
    - reason codes [515](#)
    - states [514](#)
  - Fibre Channel over Ethernet [3](#)
  - Fibre Channel Security Protocol [699](#)
  - FLOGI [663](#)
    - description [663](#)
  - flow statistics [658, 659, 660](#)
    - clearing [659](#)
    - counting [658](#)
    - description [658](#)
    - displaying [660](#)
  - frame encapsulation [521](#)
    - configuring [521](#)
  - FSCN [677](#)
    - displaying databases [677](#)
  - FSPF [643, 644, 645, 646, 647, 648, 649, 650, 651, 652, 654, 660, 688](#)
    - clearing counters [652](#)
    - clearing VSAN counters [647](#)
    - computing link cost [648](#)
    - configuring globally [645](#)
    - configuring Hello time intervals [649](#)
    - configuring link cost [648](#)
    - configuring on a VSAN [646](#)
    - configuring on interfaces [648](#)
    - dead time intervals [650](#)
    - default settings [660](#)
    - description [643](#)
    - disabling [647](#)

FSPF (*continued*)

- disabling on interfaces [651](#)
  - disabling routing protocols [647](#)
  - displaying database information [660](#)
  - displaying global information [660](#)
  - enabling [647](#)
  - fault tolerant fabrics [643](#)
  - in-order delivery [654](#)
  - interoperability [688](#)
  - link state record defaults [645](#)
  - reconvergence times [643](#)
  - redundant links [644](#)
  - resetting configuration [647](#)
  - resetting to defaults [647](#)
  - retransmitting intervals [650](#)
  - routing services [643](#)
  - topology examples [643](#)
- FSPF routes [652, 653](#)
- configuring [653](#)
  - description [652](#)
- full zone sets [606, 616](#)
- considerations [606](#)
  - enabling distribution [616](#)
- fWWNs [613](#)
- configuring fcalias members [613](#)
- Fx ports [513, 592](#)
- VSAN membership [592](#)

**G**

- GOLD diagnostics [363, 364, 365](#)
- configuring [365](#)
  - expansion modules [365](#)
  - health monitoring [364](#)
  - runtime [363](#)

**H**

- hard zoning [616](#)
- description [616](#)
- HBA ports [548](#)
- configuring area FCIDs [548](#)
- health monitoring diagnostics [364](#)
- information [364](#)
- Hello time intervals [649](#)
- configuring for FSPF [649](#)
  - description [649](#)
- host ports [86](#)
- kinds of [86](#)

**I**

- ICMPv2 [230](#)
- IDs [254, 261, 402](#)
- Cisco vendor ID [254, 261](#)
  - serial IDs [402](#)
- IEEE 802.1p [3](#)
- IEEE 802.1w [183](#)
- IEEE 802.3x [3](#)
- IGMP forwarding [231](#)
- MAC address [231](#)
- IGMP snooping [231](#)
- queries [231](#)
- IGMPv1 [230](#)
- IGMPv3 [231](#)
- in-order delivery [654, 655, 656, 657](#)
- configuring drop latency time [657](#)
  - displaying status [657](#)
  - enabling for VSANs [656](#)
  - enabling globally [656](#)
  - guidelines [655](#)
  - reordering network frames [654](#)
  - reordering port channel frames [655](#)
- indirect link failures [747](#)
- recovering [747](#)
- interface speed [64](#)
- interfaces [61, 62, 520, 522, 526, 579, 594, 595, 613](#)
- adding to SAN port channels [579](#)
  - assigning to VSANs [595](#)
  - chassis ID [61](#)
  - configuring descriptions [520](#)
  - configuring fcalias members [613](#)
  - configuring receive data field size [522](#)
  - displaying SFP information [526](#)
  - isolated states [579](#)
  - options [61](#)
  - SFP types [526](#)
  - suspended states [579](#)
  - UDLD [62](#)
  - VSAN membership [594](#)
- Interfaces [514](#)
- interop modes [688, 697](#)
- configuring mode 1 [688](#)
  - default settings [697](#)
  - description [688](#)
- interoperability [600, 688, 692](#)
- configuring interop mode 1 [688](#)
  - description [688](#)
  - verifying status [692](#)
  - VSANs [600](#)
- IOD [654](#)
- ISLs [571](#)
- SAN port channel links [571](#)
- isolated port [86](#)

isolated VLANs [86, 87](#)  
 isolated VSANs [597](#)  
   description [597](#)  
   displaying membership [597](#)

## L

LACP [111, 115, 121](#)  
   system ID [115](#)  
 license key files [47](#)  
   description [47](#)  
 licenses [47, 48, 50, 51, 52, 53, 56, 57](#)  
   claim certificates [47](#)  
   displaying information [51](#)  
   evaluation [47](#)  
   grace period expiration [56](#)  
   grace periods [47](#)  
   host IDs [47](#)  
   identifying features in use [52](#)  
   incremental [47](#)  
   installation options [48](#)  
   installing key files [51](#)  
   missing [47](#)  
   node-locked [47](#)  
   obtaining key files [50](#)  
   PAK [47](#)  
   permanent [47](#)  
   terminology [47](#)  
   transferring between switches [57](#)  
   uninstalling [53](#)  
 Link Aggregation Control Protocol [111](#)  
 link costs [648](#)  
   configuring for FSPF [648](#)  
   description [648](#)  
 Link Failure [168, 188](#)  
   detecting unidirectional [168](#)  
 link failures [747](#)  
   recovering [747](#)  
 linkDown notifications [424, 425](#)  
 linkUp notifications [424, 425](#)  
 load balancing [571, 572, 593, 599](#)  
   attributes [599](#)  
   attributes for VSANs [593](#)  
   configuring [599](#)  
   description [572, 599](#)  
   guarantees [599](#)  
   SAN port channels [571](#)  
 logical unit numbers [675](#)  
 LUNs [677](#)  
   displaying discovered SCSI targets [677](#)

## M

MAC addresses [685](#)  
   configuring secondary [685](#)  
 management access [44](#)  
   description [44](#)  
 management interfaces [44, 46](#)  
   displaying information [46](#)  
   using force option during shutdown [46](#)  
 McData [688](#)  
   native interop mode [688](#)  
 merged fabrics [537](#)  
   autoreconfigured [537](#)  
 mgmt0 interfaces [44](#)  
   description [44](#)  
 MSCHAP [251](#)  
   enabling authentication [251](#)  
 MST [186, 195](#)  
   CIST regional root [186](#)  
   setting to default values [195](#)  
 MSTP [183, 184, 185, 186, 187, 188, 195](#)  
   boundary ports [188](#)  
     described [188](#)  
   CIST regional root [186](#)  
   CIST root [187](#)  
   CIST, described [185](#)  
   CST [185, 186](#)  
     defined [185](#)  
     operations between regions [186](#)  
 IEEE 802.1s [186](#)  
   terminology [186](#)  
 IST [185, 186](#)  
   operations within a region [185](#)  
 mapping VLANs to MST instance [195](#)  
 MST region [183, 184, 185, 187](#)  
   CIST [185](#)  
   described [183](#)  
   hop-count mechanism [187](#)  
   supported spanning-tree instances [184](#)  
 multicast storms [237](#)

## N

N port identifier virtualization [524](#)  
 N ports [603, 616, 743, 773](#)  
   FCS support [743](#)  
   fctrace [773](#)  
   hard zoning [616](#)  
   zone enforcement [616](#)  
   zone membership [603](#)  
 N5K-M1008 expansion module [511](#)  
 N5K-M1404 expansion module [511](#)

- name servers [664, 665, 675, 688](#)
    - displaying database entries [665](#)
    - interoperability [688](#)
    - LUN information [675](#)
    - proxy feature [664](#)
    - registering proxies [664](#)
  - Node Proxy port mode [513](#)
  - NP links [555](#)
  - NP port mode [513](#)
  - NP-ports [553](#)
  - NPIV [524, 525](#)
    - description [524](#)
    - enabling [525](#)
  - NPV [558, 559, 561](#)
    - configuring NP interface [559](#)
    - configuring server interface [559](#)
    - enabling [558](#)
    - verifying [561](#)
- O**
- operational states [514, 519](#)
    - configuring on Fibre Channel interfaces [519](#)
    - description [514](#)
- P**
- passwords [38, 39, 347, 705](#)
    - administrator [38](#)
    - DHCHAP [705](#)
    - setting administrator default [39](#)
    - strong characteristics [347](#)
  - persistent FC IDs [546, 547, 549, 551](#)
    - configuring [547](#)
    - description [546](#)
    - displaying [551](#)
    - enabling [547](#)
    - purging [549](#)
  - PLOGI [665](#)
    - name server [665](#)
  - port channeling [111](#)
  - port channels [515, 653, 655, 688, 701](#)
    - administratively down [515](#)
    - compatibility with DHCHAP [701](#)
    - configuring Fibre Channel routes [653](#)
    - interoperability [688](#)
    - link changes [655](#)
  - port modes [514](#)
    - auto [514](#)
  - port security [701, 713, 714, 717, 718, 719, 724, 733, 735](#)
    - activating [718](#)
    - port security (*continued*)
      - activation [714](#)
      - activation rejection [718](#)
      - adding authorized pairs [724](#)
      - auto-learning [714](#)
      - compatibility with DHCHAP [701](#)
      - configuring manually without auto-learning [724](#)
      - deactivating [718](#)
      - default settings [733](#)
      - disabling [717](#)
      - displaying configuration [733](#)
      - displaying settings (procedure) [719](#)
      - displaying statistics (procedure) [719](#)
      - displaying violations (procedure) [719](#)
      - enabling [717](#)
      - enforcement mechanisms [713](#)
      - fabric binding comparison [735](#)
      - forcing activation [719](#)
      - license requirement [713](#)
      - preventing unauthorized accesses [713](#)
    - port security auto-learning [714, 715, 716, 720, 721, 726](#)
      - description [714](#)
      - device authorization [721](#)
      - disabling [721](#)
      - distributing configuration [726](#)
      - enabling [720](#)
      - guidelines for configuring with CFS [715](#)
      - guidelines for configuring without CFS [716](#)
    - port security databases [717, 719, 729, 730, 731, 732, 733](#)
      - cleaning up [732](#)
      - copying [732](#)
      - copying active to config (procedure) [719](#)
      - deleting [732](#)
      - displaying configuration [733](#)
      - interactions [730](#)
      - manual configuration guidelines [717](#)
      - merge guidelines [729](#)
      - reactivating [719](#)
      - scenarios [731](#)
    - port speeds [521](#)
      - configuring [521](#)
    - port tracking [747, 748, 749, 752, 753](#)
      - default settings [753](#)
      - description [747](#)
      - displaying information [753](#)
      - enabling [749](#)
      - guidelines [748](#)
      - shutting down ports forcefully [752](#)
    - port world wide names [603](#)
    - PortChannels [781](#)
      - show tech-support port-channel command [781](#)
    - PortFast BPDU filtering [213](#)
    - ports [594](#)
      - VSAN membership [594](#)

- preshared keys [278](#)
    - TACACS+ [278](#)
  - primary VLANs [86](#)
  - principal switches [538, 541](#)
    - assigning domain ID [538](#)
    - configuring [541](#)
  - private VLANs [86, 87, 89, 90](#)
    - community VLANs [86, 87](#)
    - end station access to [90](#)
    - isolated trunk [89](#)
    - isolated VLANs [86, 87](#)
    - ports [86](#)
      - community [86](#)
      - isolated [86](#)
      - promiscuous [86](#)
    - primary VLANs [86](#)
    - promiscuous trunk [89](#)
    - secondary VLANs [86](#)
  - promiscuous ports [86](#)
  - proxies [664](#)
    - registering for name servers [664](#)
  - pWWNs [603, 613](#)
    - configuring fcalias members [613](#)
    - zone membership [603](#)
- R**
- RADIUS [259, 262, 267, 275](#)
  - configuring servers [262](#)
  - configuring timeout intervals [267](#)
  - configuring transmission retry counts [267](#)
  - default settings [275](#)
  - example configurations [275](#)
  - network environments [259](#)
- RADIUS servers [268, 272, 273, 274, 275](#)
  - configuring timeout interval [268](#)
  - configuring transmission retry count [268](#)
  - deleting hosts [272](#)
  - displaying statistics [274](#)
  - example configurations [275](#)
  - manually monitoring [273](#)
- Rapid Spanning Tree Protocol [183](#)
- RCFs [532, 536](#)
  - description [532](#)
  - incoming [536](#)
  - rejecting incoming [536](#)
- reason codes [515](#)
  - description [515](#)
- reconfigure fabric frames [532](#)
- reduced MAC address [156](#)
- redundancy [592](#)
  - VSANs [592](#)
- Registered State Change Notifications [666](#)
- related documents [xlv](#)
- reserved-range VLANs [78](#)
- retransmitting intervals [650, 651](#)
  - configuring for FSPF [651](#)
  - description [650](#)
- roles [347](#)
  - authentication [347](#)
- rollback [359, 362](#)
  - checkpoint copy [359](#)
  - creating a checkpoint copy [359](#)
  - default settings [362](#)
  - deleting a checkpoint file [359](#)
  - description [359](#)
  - example configuration [359](#)
  - guidelines [359](#)
  - high availability [359](#)
  - implementing a rollback [359](#)
  - limitations [359](#)
  - reverting to checkpoint file [359](#)
  - verifying configuration [362](#)
- root guard [214](#)
- route costs [648](#)
  - computing [648](#)
- RSCN [666, 667, 668, 672](#)
  - default settings [672](#)
  - description [666](#)
  - displaying information [667](#)
  - multiple port IDs [667](#)
  - suppressing domain format SW-RSCNs [668](#)
  - switch RSCN [667](#)
- RSCN timers [669, 670](#)
  - configuration distribution using CFS [670](#)
  - configuring [669](#)
- RSTP [160, 164, 168, 183](#)
  - active topology [164](#)
  - BPDU [168](#)
    - processing [168](#)
  - designated port, defined [164](#)
  - designated switch, defined [164](#)
  - proposal-agreement handshake process [160](#)
  - rapid convergence [160](#)
    - point-to-point links [160](#)
    - root ports [160](#)
    - root port, defined [164](#)
- runtime checks [653](#)
  - static routes [653](#)
- runtime diagnostics [363](#)
  - information [363](#)

## S

- SAN port channel [586](#)
  - verifying configurations [586](#)
- SAN port Channel [587](#)
  - default settings [587](#)
- SAN port channel protocol [584](#)
  - configuring autocreation [584](#)
  - enabling autocreation [584](#)
- SAN port channel Protocol [582](#)
  - autocreation [582](#)
  - creating channel group [582](#)
- SAN port channels [571, 572, 575, 579](#)
  - adding interfaces [579](#)
  - comparison with trunking [572](#)
  - compatibility checks [579](#)
  - configuration guidelines [575](#)
  - description [571](#)
  - interface states [579](#)
  - load balancing [572](#)
  - misconfiguration error detection [575](#)
- scalability [592](#)
  - VSANs [592](#)
- SCR [666](#)
  - request [666](#)
- SCSI [677](#)
  - displaying LUN discovery results [677](#)
- SCSI LUNs [675, 676, 677](#)
  - customized discovery [676](#)
  - discovering targets [675](#)
  - displaying information [677](#)
  - starting discoveries [675](#)
- SD port mode [514](#)
  - description [514](#)
  - interface modes [514](#)
- SD ports [519](#)
  - configuring [519](#)
- secondary MAC addresses [685](#)
  - configuring [685](#)
- secondary VLANs [86](#)
- serial IDs [402](#)
  - description [402](#)
- server groups [245](#)
- server IDs [402](#)
  - description [402](#)
- service requests [xlvi](#)
- session manager [359, 361, 362](#)
  - committing a session [361](#)
  - configuring an ACL session (example) [362](#)
  - description [359](#)
  - discarding a session [362](#)
  - guidelines [359](#)
  - limitations [359](#)
  - saving a session [361](#)
- session manager (*continued*)
  - verifying configuration [362](#)
  - verifying the session [361](#)
- SFP+ transceiver [64](#)
- SFPs [526](#)
  - displaying transmitter types [526](#)
  - transmitter types [526](#)
- small computer system interface [675](#)
- Small form-factor pluggable (plus) transceiver [64](#)
- smart call home [389](#)
  - description [389](#)
  - registration requirements [389](#)
  - Transport Gateway (TG) aggregation point [389](#)
- SMARTnet [389](#)
  - smart call home registration [389](#)
- SNMP [389, 416, 418, 419](#)
  - access groups [419](#)
  - group-based access [419](#)
  - server contact name [389](#)
  - user synchronization with CLI [418](#)
  - Version 3 security features [416](#)
- SNMP (Simple Network Management Protocol) [416](#)
  - versions [416](#)
- SNMPv3 [253, 255, 416, 420](#)
  - assigning multiple roles [420](#)
  - security features [416](#)
  - specifying AAA parameters [253](#)
  - specifying parameters for AAA servers [255](#)
- soft zoning [616](#)
  - description [616](#)
- source IDs [402, 572, 599, 654](#)
  - call home event format [402](#)
  - exchange based [572](#)
  - flow based [572](#)
  - in-order delivery [654](#)
  - path selection [599](#)
- SPAN [757](#)
  - egress sources [757](#)
  - ingress sources [757](#)
  - sources for monitoring [757](#)
- SPAN destination port mode [514](#)
- SPAN sources [757](#)
  - egress [757](#)
  - ingress [757](#)
- SPF [645](#)
  - computational hold times [645](#)
- static routes [653](#)
  - runtime checks [653](#)
- statistics [292](#)
  - TACACS+ [292](#)
- storage devices [603](#)
  - access control [603](#)
- STP [111, 160, 166, 167, 211, 212](#)
  - edge ports [160, 211](#)

STP (*continued*)

- EtherChannel [111](#)
- network ports [212](#)
- normal ports [212](#)
- port types [211](#)
- PortFast [160, 211](#)
- understanding [166, 167](#)
  - Blocking State [166](#)
  - disabled state [167](#)
  - forwarding state [166](#)
  - learning state [166](#)
- STP bridge ID [156](#)
- STP root guard [214](#)
- switch ports [524](#)
  - configuring attribute default values [524](#)
- switch priorities [534](#)
  - default [534](#)
  - description [534](#)
- Switched Port Analyzer [757](#)
- sWWNs [738](#)
  - configuring for fabric binding [738](#)

## T

- TACACS+ [277, 278, 279, 280, 286, 292, 293](#)
  - advantages over RADIUS [277](#)
  - configuring [280](#)
  - configuring global timeout interval [286](#)
  - description [277](#)
  - displaying statistics [292](#)
  - example configurations [293](#)
  - field descriptions [293](#)
  - global preshared keys [278](#)
  - limitations [279](#)
  - prerequisites [279](#)
  - preshared key [278](#)
  - user login operation [278](#)
  - verifying configuration [292](#)
- TACACS+ servers [280, 287, 288, 291, 292, 293](#)
  - configuring hosts [280](#)
  - configuring TCP ports [288](#)
  - configuring timeout interval [287](#)
  - displaying statistics [292](#)
  - field descriptions [293](#)
  - manually monitoring [291](#)
  - verifying configuration [292](#)
- TCP ports [288](#)
  - TACACS+ servers [288](#)
- TE port mode [513](#)
  - classes of service [513](#)
  - description [513](#)
- TE ports [564, 618, 643, 688, 735, 743, 744, 773](#)
  - fabric binding checking [735](#)
  - FCS support [743](#)
  - fctrace [773](#)
  - FSPF topologies [643](#)
  - interoperability [688](#)
  - recovering from link isolations [618](#)
  - trunking restrictions [564](#)
- timeout values [679](#)
- TOV [679, 680, 688, 697](#)
  - configuring across all VSANs [679](#)
  - configuring for a VSAN [680](#)
  - default settings [697](#)
  - interoperability [688](#)
  - ranges [679](#)
- tracked ports [749](#)
  - binding operationally [749](#)
- traffic isolation [592](#)
  - VSANs [592](#)
- trap notifications [416](#)
- troubleshooting [773, 774, 775, 776](#)
  - collecting output for technical support [776](#)
  - fcping [774](#)
  - fctrace [773](#)
  - show tech-support command [776](#)
  - verifying switch connectivity [775](#)
- trunk mode [524, 565, 566, 569](#)
  - administrative default [524](#)
  - configuring [565, 566](#)
  - default settings [569](#)
- trunk ports [569](#)
  - displaying information [569](#)
- trunk-allowed VSAN lists [567](#)
  - description [567](#)
- trunking [563, 564, 565, 569, 572, 688](#)
  - comparison with port channels [572](#)
  - configuration guidelines [564](#)
  - configuring modes [565](#)
  - default settings [569](#)
  - description [563](#)
  - displaying information [569](#)
  - interoperability [688](#)
  - link state [565](#)
  - merging traffic [564](#)
  - restrictions [563](#)
- trunking E port mode [513](#)
- trunking ports [595](#)
  - associated with VSANs [595](#)
- trunking protocol [564, 565, 569](#)
  - default settings [569](#)
  - default state [565](#)
  - description [564](#)
  - detecting port isolation [564](#)

## U

- UDLD [62, 63](#)
  - aggressive mode [63](#)
  - defined [62](#)
  - nonaggressive mode [63](#)
- unicast storms [237](#)
- Unidirectional Link Detection [62](#)
- unique area FC IDs [548](#)
  - configuring [548](#)
  - description [548](#)
- user accounts [347](#)
  - password characteristics [347](#)
- user login [246](#)
  - authentication process [246](#)
  - authorization process [246](#)
- user roles [253, 255](#)
  - specifying on AAA servers [253, 255](#)
- users [347](#)
  - description [347](#)

## V

- VDC [8](#)
  - defined [8](#)
- vendor-specific attributes [254](#)
- verifying NPV [561](#)
- virtual device contexts [8](#)
- Virtual Fibre Channel interfaces [528](#)
  - default settings [528](#)
- VLANs [78](#)
  - extended range [78](#)
  - reserved range [78](#)
- VSAN IDs [513, 569, 592, 593](#)
  - allowed list [569](#)
  - description [593](#)
  - multiplexing traffic [513](#)
  - range [592](#)
  - VSAN membership [592](#)
- VSANs [513, 515, 538, 551, 564, 568, 589, 592, 593, 594, 595, 596, 597, 599, 600, 606, 643, 645, 646, 658, 664, 679, 688, 701, 743](#)
  - advantages [589](#)
  - allowed-active [564](#)
  - cache contents [551](#)
  - comparison with zones (table) [592](#)
  - compatibility with DHCHAP [701](#)
  - configuring [594](#)
  - configuring allowed-active lists [568](#)
  - configuring FSPF [645](#)
  - configuring trunk-allowed lists [568](#)
  - default settings [600](#)
  - default VSANs [596](#)
  - deleting [597](#)

## VSANs (continued)

- description [589](#)
  - displaying configuration [600](#)
  - displaying membership [595](#)
  - displaying usage [600](#)
  - domain ID automatic reconfiguration [538](#)
  - FC IDs [589](#)
  - FCS support [743](#)
  - features [589](#)
  - flow statistics [658](#)
  - FSPF [646](#)
  - FSPF connectivity [643](#)
  - interop mode [688](#)
  - isolated [597](#)
  - load balancing [599](#)
  - load balancing attributes [593](#)
  - mismatches [515](#)
  - multiple zones [606](#)
  - name server [664](#)
  - names [593](#)
  - operational states [597](#)
  - port membership [594](#)
  - states [593](#)
  - TE port mode [513](#)
  - timer configuration [679](#)
  - TOV [679](#)
  - traffic isolation [589](#)
  - trunk-allowed [564](#)
  - trunking ports [595](#)
- VSAs [254](#)
- format [254](#)
  - protocol options [254](#)
  - support description [254](#)

## W

- world wide names [684](#)
- WWNs [515, 684, 685](#)
  - description [684](#)
  - displaying information [685](#)
  - link initialization [685](#)
  - secondary MAC addresses [685](#)
  - suspended connections [515](#)

## Z

- zone aliases [638](#)
  - conversion to device aliases [638](#)
- zone attribute groups [621](#)
  - cloning [621](#)

- zone databases [621, 626](#)
  - migrating a non-MDS database [621](#)
  - release locks [626](#)
- zone members [612](#)
  - displaying information [612](#)
- zone server databases [621](#)
  - clearing [621](#)
- zone sets [603, 606, 611, 616, 617, 618, 620, 621, 622, 630](#)
  - activating [611](#)
  - analyzing [630](#)
  - cloning [621](#)
  - considerations [606](#)
  - creating [611](#)
  - displaying information [622](#)
  - distributing configuration [616](#)
  - enabling distribution [616](#)
  - exporting [618](#)
  - exporting databases [618](#)
  - features [603](#)
  - importing [618](#)
  - importing databases [618](#)
  - one-time distribution [617](#)
  - recovering from link isolations [618](#)
  - renaming [620](#)
  - viewing information [622](#)
- zones [515, 592, 603, 605, 611, 613, 618, 619, 620, 621, 622, 629, 630, 632, 778](#)
  - access control [611](#)
  - analyzing [630](#)
  - backing up (procedure) [619](#)
  - cloning [621](#)
  - compacting for downgrading [629](#)
  - comparison with device aliases [632](#)
  - comparison with VSANs (table) [592](#)
  - configuring aliases [613](#)
  - configuring fcaliases [613](#)
  - default policies [603](#)
  - displaying information [622](#)
  - exporting databases [618](#)
  - features [603, 605](#)
  - importing databases [618](#)
  - membership using pWWNs [592](#)
  - merge failures [515](#)
  - renaming [620](#)
  - restoring (procedure) [619](#)
  - show tech-support zone command [778](#)
  - viewing information [622](#)
- zoning [603, 605](#)
  - description [603](#)
  - example [605](#)
  - implementation [605](#)