



Cisco Nexus 9000 Series NX-OS Interfaces Configuration Guide, Release 7.x

First Published: 2015-01-27

Last Modified: 2021-10-22

Americas Headquarters

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

THE SPECIFICATIONS AND INFORMATION REGARDING THE PRODUCTS REFERENCED IN THIS DOCUMENTATION ARE SUBJECT TO CHANGE WITHOUT NOTICE. EXCEPT AS MAY OTHERWISE BE AGREED BY CISCO IN WRITING, ALL STATEMENTS, INFORMATION, AND RECOMMENDATIONS IN THIS DOCUMENTATION ARE PRESENTED WITHOUT WARRANTY OF ANY KIND, EXPRESS OR IMPLIED.

The Cisco End User License Agreement and any supplemental license terms govern your use of any Cisco software, including this product documentation, and are located at: <http://www.cisco.com/go/softwareterms>. Cisco product warranty information is available at <http://www.cisco.com/go/warranty>. US Federal Communications Commission Notices are found here <http://www.cisco.com/c/en/us/products/us-fcc-notice.html>.

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

Any products and features described herein as in development or available at a future date remain in varying stages of development and will be offered on a when-and if-available basis. Any such product or feature roadmaps are subject to change at the sole discretion of Cisco and Cisco will have no liability for delay in the delivery or failure to deliver any products or feature roadmap items that may be set forth in this document.

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

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

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

© 2015–2021 Cisco Systems, Inc. All rights reserved.



CONTENTS

PREFACE

Preface	xvii
Audience	xvii
Document Conventions	xvii
Related Documentation for Cisco Nexus 9000 Series Switches	xviii
Documentation Feedback	xviii
Communications, Services, and Additional Information	xviii

CHAPTER 1

New and Changed Information	1
New and Changed Information	1

CHAPTER 2

Overview	9
Licensing Requirements	9
Supported Platforms	9
About Interfaces	9
Ethernet Interfaces	10
Access Ports	10
Routed Ports	10
Management Interface	10
Port-Channel Interfaces	10
Subinterfaces	11
Loopback Interfaces	11
Breakout Interfaces	11
Module Level Breakout	11
Dynamic Breakout (Per-port Level Breakout)	11
About the Lane Selector	12
Notes About Breakout Interfaces	13

Virtual Device Contexts 17

High Availability for Interfaces 17

CHAPTER 3**Configuring Basic Interface Parameters 19**

About the Basic Interface Parameters 19

Description 19

Beacon 19

Error Disabled 19

Interface Status Error Policy 20

Modifying Interface MTU Size 20

Bandwidth 22

Throughput Delay 22

Administrative Status 22

Unidirectional Link Detection Parameter 23

UDLD Overview 23

Default UDLD Configuration 24

UDLD Normal and Aggressive Modes 24

Port-Channel Parameters 25

Port Profiles 25

Cisco QSFP+ to SFP+ Adapter Module Support 27

Cisco SFP+ Adapter Module Support 27

25G Autonegotiation Overview 28

Guidelines and Limitations 28

Default Settings 30

Configuring the Basic Interface Parameters 31

Specifying the Interfaces to Configure 31

Configuring the Description 33

Configuring the Beacon Mode 34

Configuring the Error-Disabled State 35

Enabling the Error-Disable Detection 35

Enabling the Error-Disabled Recovery 36

Configuring the Error-Disabled Recovery Interval 37

Configuring the MTU Size 38

Configuring the Interface MTU Size 38

Configuring the System Jumbo MTU Size	40
Configuring the Bandwidth	41
Configuring the Throughput Delay	42
Shutting Down and Activating the Interface	44
Configuring the UDLD Mode	45
Configuring Debounce Timers	48
Configuring Port Profiles	50
Creating a Port Profile	50
Entering Port-Profile Configuration Mode and Modifying a Port Profile	51
Assigning a Port Profile to a Range of Interfaces	52
Enabling a Specific Port Profile	53
Inheriting a Port Profile	53
Removing a Port Profile from a Range of Interfaces	54
Removing an Inherited Port Profile	55
Configuring link mac-up timer	55
Verifying the Basic Interface Parameters	56
Monitoring the Interface Counters	56
Displaying Interface Statistics	56
Clearing Interface Counters	58
Configuration Example for QSA	58

CHAPTER 4

Configuring Layer 2 Interfaces	61
Information About Access and Trunk Interfaces	62
About Access and Trunk Interfaces	62
IEEE 802.1Q Encapsulation	63
Access VLANs	64
Native VLAN IDs for Trunk Ports	64
Tagging Native VLAN Traffic	65
Allowed VLANs	65
Switchport Isolated for up to 3967 VLANs on Trunk Interfaces	65
Default Interfaces	66
Switch Virtual Interface and Autostate Behavior	66
SVI Autostate Exclude	66
SVI Autostate Disable	67

High Availability	67
Virtualization Support	67
Counter Values	67
Prerequisites for Layer 2 Interfaces	68
Guidelines and Limitations for Layer 2 Interfaces	68
Default Settings for Layer 2 Interfaces	71
Configuring Access and Trunk Interfaces	72
Guidelines for Configuring Access and Trunk Interfaces	72
Configuring a VLAN Interface as a Layer 2 Access Port	72
Configuring Access Host Ports	74
Configuring Trunk Ports	75
Configuring the Native VLAN for 802.1Q Trunking Ports	76
Configuring the Allowed VLANs for Trunking Ports	78
Configuring MAC addresses Limitation on a VLAN	80
Configuring switchport isolated	81
Configuring a Default Interface	81
Configuring SVI Autostate Exclude	83
Configuring SVI Autostate Disable for the System	84
Configuring SVI Autostate Disable Per SVI	85
Configuring the Device to Tag Native VLAN Traffic	86
Configuring Interface Breakout Profile for 50-G Interfaces in a 16-Slot Chassis	88
Changing the System Default Port Mode to Layer 2	89
Verifying the Interface Configuration	90
Monitoring the Layer 2 Interfaces	91
Configuration Examples for Access and Trunk Ports	91
Related Documents	92

CHAPTER 5**Configuring Layer 3 Interfaces 93**

About Layer 3 Interfaces	93
Routed Interfaces	93
Subinterfaces	94
Limitations for Subinterfaces	94
VLAN Interfaces	95
Changing VRF Membership for an Interface	95

Notes About Changing VRF Membership for an Interface	96
Loopback Interfaces	96
IP Unnumbered	96
MAC-Embedded IPv6 Address	97
High Availability	97
Virtualization Support	97
DHCP Client	97
Limitations for Using DHCP Client on Interfaces	98
Layer 3 Static MAC Addresses	98
Prerequisites for Layer 3 Interfaces	99
Guidelines and Limitations	99
Default Settings	100
Configuring Layer 3 Interfaces	100
Configuring a Routed Interface	100
Configuring a Subinterface on a Routed Interface	102
Configuring a Subinterface on a Port-Channel Interface	103
Configuring a VLAN Interface	105
Enabling Layer 3 Retention During VRF Membership Change	106
Configuring a Static MAC Address on a Layer 3 Interface	107
Configuring a Loopback Interface	108
Configuring IP Unnumbered on an Ethernet Interface	109
Configuring OSPF for an IP Unnumbered Interface	110
Configuring ISIS for an IP Unnumbered Interface	111
Configuring PBR on SVI on the Gateway	113
Configuring IP Unnumbered on SVI Secondary VLAN on the Gateway	115
Assigning an Interface to a VRF	116
Configuring a MAC-Embedded IPv6 Address	117
Configuring a DHCP Client on an Interface	120
Verifying the Layer 3 Interfaces Configuration	121
Monitoring the Layer 3 Interfaces	122
Configuration Examples for Layer 3 Interfaces	123
Example of Changing VRF Membership for an Interface	123
Related Documents	125

CHAPTER 6	Configuring Bidirectional Forwarding Detection	127
	About BFD	127
	Asynchronous Mode	127
	BFD Detection of Failures	128
	Distributed Operation	129
	BFD Echo Function	129
	Security	129
	High Availability	129
	Virtualization Support	129
	Prerequisites for BFD	129
	Guidelines and Limitations	130
	Default Settings	132
	Configuring BFD	132
	Configuration Hierarchy	132
	Task Flow for Configuring BFD	133
	Enabling the BFD Feature	133
	Configuring Global BFD Parameters	134
	Configuring BFD on an Interface	135
	Configuring BFD on a Port Channel	137
	Configuring the BFD Echo Function	138
	Configuring Per-Member Link BFD Sessions	139
	BFD Enhancement to Address Per-link Efficiency	139
	Limitations of the IETF Bidirectional Forwarding Detection	139
	Configuring Port Channel Interface	141
	(Optional) Configuring BFD Start Timer	141
	Enabling IETF Per-link BFD	142
	Configuring BFD Destination IP Address	142
	Verifying Micro BFD Session Configurations	143
	Examples: Configuring Micro BFD Sessions	143
	Configuring BFD Support for Routing Protocols	146
	Configuring BFD on BGP	146
	Configuring BFD on EIGRP	147
	Configuring BFD on OSPF	149

Configuring BFD on IS-IS	150
Configuring BFD on HSRP	151
Configuring BFD on VRRP	152
Configuring BFD on PIM	153
Configuring BFD on Static Routes	154
Disabling BFD on an Interface	155
Configuring BFD Interoperability	156
Configuring BFD Interoperability in Cisco NX-OS Devices in a Point-to-Point Link	156
Configuring BFD Interoperability in Cisco NX-OS Devices in a Switch Virtual Interface	157
Configuring BFD Interoperability in Cisco NX-OS Devices in Logical Mode	158
Verifying BFD Interoperability in a Cisco Nexus 9000 Series Device	159
Verifying the BFD Configuration	159
Monitoring BFD	160
Configuration Examples for BFD	160
Show Example for BFD	161
Related Documents	161
RFCs	161

CHAPTER 7

Configuring Port Channels	163
About Port Channels	163
Port Channels	164
Port-Channel Interfaces	165
Basic Settings	165
Compatibility Requirements	166
Load Balancing Using Port Channels	168
Symmetric Hashing	169
Guidelines and Limitations for ECMP	169
Resilient Hashing	170
GTP Tunnel Load Balancing	170
LACP	171
LACP Overview	171
Port-Channel Modes	172
LACP ID Parameters	174
LACP System Priority	174

LACP Port Priority	174
LACP Administrative Key	174
LACP Marker Responders	174
LACP-Enabled and Static Port Channels Differences	175
LACP Compatibility Enhancements	175
LACP Port-Channel Minimum Links and MaxBundle	176
LACP Fast Timers	176
Virtualization Support	176
High Availability	177
Prerequisites for Port Channeling	177
Guidelines and Limitations	177
Default Settings	179
Configuring Port Channels	180
Creating a Port Channel	180
Adding a Layer 2 Port to a Port Channel	182
Adding a Layer 3 Port to a Port Channel	184
Applying Port Channel Configuration to all Configured Port Channels	186
Configuring the Bandwidth and Delay for Informational Purposes	187
Shutting Down and Restarting the Port-Channel Interface	188
Configuring a Port-Channel Description	189
Configuring the Speed and Duplex Settings for a Port-Channel Interface	190
Configuring Load Balancing Using Port Channels	191
Enabling LACP	192
Configuring LACP Port-Channel Port Modes	193
Configuring LACP Port-Channel Minimum Links	194
Configuring the LACP Port-Channel MaxBundle	196
Configuring the LACP Fast Timer Rate	197
Configuring the LACP System Priority	198
Configuring the LACP Port Priority	199
Configuring LACP System MAC and Role	200
Disabling LACP Graceful Convergence	201
Reenabling LACP Graceful Convergence	202
Disabling LACP Suspend Individual	203
Reenabling LACP Suspend Individual	204

Configuring Delayed LACP	205
Configuring Port Channel Hash Distribution	207
Configuring Port Channel Hash Distribution at the Global Level	207
Configuring Port Channel Hash Distribution at the Port Channel Level	208
Enabling ECMP Resilient Hashing	209
Disabling ECMP Resilient Hashing	209
Configuring ECMP Load Balancing based on Inner NVGRE Header	210
Verifying the ECMP Resilient Hashing Configuration	211
Verifying the Port-Channel Configuration	212
Monitoring the Port-Channel Interface Configuration	213
Example Configurations for Port Channels	213
Related Documents	214

CHAPTER 8
Configuring vPCs 215

Information About vPCs	215
vPC Overview	215
Hitless vPC Role Change	217
vPC Terminology	218
vPC Peer-Link Overview	219
Features That You Must Manually Configure on the Primary and Secondary Devices	221
Configuring Layer 3 Backup Routes on a vPC Peer-Link	222
Peer-Keepalive Link and Messages	222
vPC Peer-Gateway	223
vPC Domain	224
vPC Topology	224
Compatibility Parameters for vPC Interfaces	226
Configuration Parameters That Must Be Identical	226
Configuration Parameters That Should Be Identical	227
Consequences of Parameter Mismatches	228
vPC Number	229
Moving Other Port Channels into a vPC	229
vPC Object Tracking	229
vPC Interactions with Other Features	231
vPC and LACP	231

vPC Peer-Links and STP	232
vPC Peer Switch	233
vPC and ARP or ND	234
vPC Multicast—PIM, IGMP, and IGMP Snooping	234
Multicast PIM Dual DR (Proxy DR)	235
IP PIM PRE-BUILD SPT	236
vPC Peer-Links and Routing	236
CFSoSE	237
vPC and Orphan Ports	238
Virtualization Support	238
vPC Recovery After an Outage	238
Autorecovery	238
vPC Peer Roles After a Recovery	238
High Availability	239
vPC Forklift Upgrade Scenario	239
Guidelines and Limitations	241
Best Practices for Layer 3 and vPC Configuration	244
Layer 3 and vPC Configuration Overview	244
Supported Topologies for Layer 3 and vPC	245
Peering with an External Router Using Layer 3 Links	245
Peering Between vPC Devices for a Backup Routing Path	246
Direct Layer 3 Peering Between Routers	247
Peering Between Two Routers with vPC Devices as Transit Switches	248
Peering with an External Router on Parallel Interconnected Routed Ports	248
Peering between vPC Switch Pairs on Parallel Interconnected Routed Ports	249
Peering Over a PC Interconnection and Dedicated Interswitch Link Using non-vPC VLAN	249
Peering Directly Over a vPC Connection	250
Configuring Layer 3 over vPC	251
Default Settings	253
Configuring vPCs	253
Enabling vPCs	253
Disabling vPCs	254
Creating a vPC Domain and Entering vpc-domain Mode	255
Configuring a vPC Keepalive Link and Messages	256

Creating a vPC Peer-Link	258
Configuring a vPC Peer-Gateway	260
Configuring Fast Convergence	261
Configuring LACP vPC Convergence	261
Configuring a Graceful Consistency Check	262
Checking the Configuration Compatibility on a vPC Peer-Link	263
Moving Other Port Channels into a vPC	264
Manually Configuring a vPC Domain MAC Address	266
Manually Configuring the System Priority	267
Manually Configuring the vPC Peer Device Role	268
Configuring vPC Object Tracking Tracking Feature on a Single-Module vPC	269
Configuring for Recovery After an Outage	271
Configuring Reload Restore	271
Configuring an Autorecovery	273
Configuring the Suspension of Orphan Ports	274
Configuring Delay Restore on an Orphan Port	275
Configuring the vPC Peer Switch	276
Configuring a Pure vPC Peer Switch Topology	276
Configuring a Hybrid vPC Peer Switch Topology	278
Configuring Hitless vPC Role Change	279
Use Case Scenario for vPC Role Change	280
Enabling STP to Use the Cisco MAC Address	281
Verifying the vPC Configuration	281
Monitoring vPCs	282
Configuration Examples for vPCs	282
Related Documents	285

CHAPTER 9**Configuring IP Tunnels 287**

Information About IP Tunnels	287
IP Tunnel Overview	287
GRE Tunnels	288
Point-to-Point IP-in-IP Tunnel Encapsulation and Decapsulation	288
Multi-Point IP-in-IP Tunnel Decapsulation	288
Path MTU Discovery	288

High Availability	289
Prerequisites for IP Tunnels	289
Guidelines and Limitations	289
Default Settings	293
Configuring IP Tunnels	294
Enabling Tunneling	294
Creating a Tunnel Interface	294
Creating an IP-in-IP Tunnel with a Netmask	296
Configuring a Tunnel Interface	298
Configuring a GRE Tunnel	301
Configuring a GRE Tunnel	302
Enabling Path MTU Discovery	304
Assigning VRF Membership to a Tunnel Interface	305
Verifying the IP Tunnel Configuration	306
Configuration Examples for IP Tunneling	306
Related Documents	307

CHAPTER 10

Configuring Q-in-Q VLAN Tunnels	309
Information About Q-in-Q Tunnels	309
Q-in-Q Tunneling	309
Native VLAN Hazard	311
Information About Layer 2 Protocol Tunneling	312
Guidelines and Limitations	313
Configuring Q-in-Q Tunnels and Layer 2 Protocol Tunneling	315
Creating a 802.1Q Tunnel Port	315
Configuring VLAN Mapping for Selective Q-in-Q on a 802.1Q Tunnel Port	317
Changing the EtherType for Q-in-Q	318
Enabling the Layer 2 Protocol Tunnel	318
Configuring Global CoS for L2 Protocol Tunnel Ports	319
Configuring Thresholds for Layer 2 Protocol Tunnel Ports	320
Configuring Q-in-Q Double Tagging	321
Verifying the Q-in-Q Configuration	322
Configuration Examples for Q-in-Q and Layer 2 Protocol Tunneling	323

CHAPTER 11**Configuring Smart-Channel 325**

- About Smart Channel 325
 - Smart Channel Features 325
 - Benefits of Smart Channel 326
 - Examples of the Deployment Use Cases 326
- Topology Examples for Smart Channel 326
- Prerequisites for Smart Channel 328
- Guidelines and Limitations for Smart Channel 329
- Default Settings for Smart Channel 329
- Configuring Smart Channel 329
 - Enabling Smart Channel 329
 - Configuring Port Groups 330
 - Configuring the Smart Channel Service 331
- Verifying the Smart Channel Configuration 332
- Configuration Examples for Smart Channel 333

CHAPTER 12**Configuring Static and Dynamic NAT Translation 337**

- Network Address Translation Overview 337
- Information About Static NAT 338
- Dynamic NAT Overview 339
- Timeout Mechanisms 340
- NAT Inside and Outside Addresses 341
- Pool Support for Dynamic NAT 342
- Static and Dynamic Twice NAT Overview 342
- VRF Aware NAT 343
- Guidelines and Limitations for Static NAT 345
- Restrictions for Dynamic NAT 346
- Guidelines and Limitations for Dynamic Twice NAT 347
- Configuring Static NAT 348
 - Enabling Static NAT 348
 - Configuring Static NAT on an Interface 348
 - Enabling Static NAT for an Inside Source Address 349
 - Enabling Static NAT for an Outside Source Address 350

Configuring Static PAT for an Inside Source Address 350

Configuring Static PAT for an Outside Source Address 351

Configuring Static Twice NAT 352

Configuration Example for Static NAT and PAT 353

Example: Configuring Static Twice NAT 354

Verifying the Static NAT Configuration 354

Configuring Dynamic NAT 355

Configuring Dynamic Translation and Translation Timeouts 355

Configuring Dynamic NAT Pool 357

Configuring Source Lists 358

Configuring Dynamic Twice NAT for an Inside Source Address 359

Configuring Dynamic Twice NAT for an Outside Source Address 360

Clearing Dynamic NAT Translations 361

Verifying Dynamic NAT Configuration 362

Example: Configuring Dynamic Translation and Translation Timeouts 364

APPENDIX A **Configuring Layer 2 Data Center Interconnect 365**

 Overview 365

 Example of Layer 2 Data Center Interconnect 365

APPENDIX B **IETF RFCs supported by Cisco NX-OS Interfaces 367**

 IPv6 RFCs 367

APPENDIX C **Configuration Limits for Cisco NX-OS Interfaces 369**



Preface

This preface includes the following sections:

- [Audience, on page xvii](#)
- [Document Conventions, on page xvii](#)
- [Related Documentation for Cisco Nexus 9000 Series Switches, on page xviii](#)
- [Documentation Feedback, on page xviii](#)
- [Communications, Services, and Additional Information, on page xviii](#)

Audience

This publication is for network administrators who install, configure, and maintain Cisco Nexus switches.

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 you supply the values.
[x]	Square brackets enclose an optional element (keyword or argument).
[x y]	Square brackets enclosing keywords or arguments that are separated by a vertical bar indicate an optional choice.
{x y}	Braces enclosing keywords or arguments that are 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.

Convention	Description
<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 includes the quotation marks.

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 that 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.

Related Documentation for Cisco Nexus 9000 Series Switches

The entire Cisco Nexus 9000 Series switch documentation set is available at the following URL:

http://www.cisco.com/en/US/products/ps13386/tsd_products_support_series_home.html

Documentation Feedback

To provide technical feedback on this document, or to report an error or omission, please send your comments to nexus9k-docfeedback@cisco.com. We appreciate your feedback.

Communications, Services, and Additional Information

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

Cisco Bug Search Tool

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



CHAPTER 1

New and Changed Information

This chapter provides release-specific information for each new and changed feature in the *Cisco Nexus 9000 Series NX-OS Interfaces Configuration Guide*.

- [New and Changed Information, on page 1](#)

New and Changed Information

This table summarizes the new and changed features for the *Cisco Nexus 9000 Series NX-OS Interfaces Configuration Guide* and where they are documented.

Table 1: New and Changed Features

Feature	Description	Changed in Release	Where Documented
Configuring ECMP Load Balancing based on Inner NVGRE Header	Support to enable ECMP load sharing based on IPv6 inner NVGRE header for IPv6 packets on Cisco Nexus 9500-FM-E platform switches.	7.0(3)I7(8)	Configuring ECMP Load Balancing based on Inner NVGRE Header, on page 210
Command interface port-channel all	Added support to select the configuration a port channel and then apply that configuration to the member ports of all the configured port channels using interface port-channel all command.	7.0(3)I7(8)	Guidelines and Limitations, on page 177 Applying Port Channel Configuration to all Configured Port Channels, on page 186
FEC Information.	Added support to display FEC information.	7.0(3)I7(7)	Caveats, on page 13

Feature	Description	Changed in Release	Where Documented
Debounce timer for the link-up events	Added support to configure debounce timer for the link-up events on Cisco Nexus N9K-C9732C-FX, N9K-C9364C and N9K-X97160YC-EX switches.	7.0(3)I7(6)	Configuring Debounce Timers, on page 48
QSFP-100G SRBD	Added support for QSFP-100G SRBD on Cisco Nexus N9K-C93180YC-FX and N9K-C93180TC-FX switches.	7.0(3)I7(5)	Guidelines and Limitations, on page 130
Debounce timer for the link-up events	Added support to configure debounce timer for the link-up events on Cisco Nexus N9K-C9336C-FX2 and N9K-C93240YC-FX2 switches.	7.0(3)I7(5)	Configuring Debounce Timers, on page 48
Link Aggregation Control Protocol (LACP) vPC convergence	Added support for Link Aggregation Control Protocol (LACP) vPC convergence on Cisco Nexus 9500 Series Switches.	7.0(3)I7(5)	Configuring LACP vPC Convergence, on page 261
Configuring MAC addresses Limitation on a VLAN	Added support to configure MAC addresses Limitation on a VLAN	7.0(3)F3(4)	Configuring MAC addresses Limitation on a VLAN, on page 80
Configuring Custom EtherType or Tag Protocol Identifier (TPID)	Added support to Custom EtherType or Tag Protocol Identifier	7.0(3)F3(4)	Changing the EtherType for Q-in-Q, on page 318
QinQ Double Tagging	Added support to configure the switch to all multi-tagged BPDUs on a tunnel port.	7.0(3)I7(3)	Information About Layer 2 Protocol Tunneling, on page 312 Configuring Q-in-Q Double Tagging, on page 321
GPRS Tunneling Protocol	Added support to deliver mobile data on wireless networks via Cisco Nexus 9300-EX Series switches.	7.0(3)I7(3)	GTP Tunnel Load Balancing, on page 170
Tunnel Source Direct	Support added for the tunnel source direct command.	7.0(3)I7(3)	Guidelines and Limitations, on page 289
1024 x 50-G Interfaces	Added support to breakout 100-G ports into two 50-G interfaces for slots 8 to 16 in the Cisco Nexus 9516 switch for -EX line cards.	7.0(3)I7(3)	Configuring Interface Breakout Profile for 50-G Interfaces in a 16-Slot Chassis, on page 88

Feature	Description	Changed in Release	Where Documented
QSFP 40/100-G BiDi	QSFP 40/100-G BiDi comes up in the highest possible speed available on the port.	7.0(3)I7(3)	Guidelines and Limitations, on page 130
LACP System MAC and Role	Added support to configure the MAC address used by the LACP for protocol exchanges and the option role.	7.0(3)I7(3)	Configuring LACP System MAC and Role, on page 200
mac-address bpd source version 2 command.	Triggers STP to use the new Cisco MAC address (00:26:0b:xx:xx:xx).	7.0(3)I7(2)	Enabling STP to Use the Cisco MAC Address, on page 281
Breakout support for the Cisco Nexus C93180LC-EX switch	Added breakout support for the Cisco Nexus C93180LC-EX switch. (TOR switch)	7.0(3)I7(1)	Cisco Nexus 9000 C93180LC-EX Switch, on page 16
ECMP Resilient Hashing	Added support for the Cisco Nexus 9200 and Cisco Nexus 9300-EX platform switches.	7.0(3)I7(1)	#unique_27
Hitless vPC Role Change	Added support for hitless vPC role change.	7.0(3)I7(1)	Hitless vPC Role Change, on page 217
IP Unnumbered on SVI Secondary VLAN on the Gateway	Added support for IP unnumbered on the secondary SVI in the gateway.	7.0(3)I7(1)	Configuring IP Unnumbered on SVI Secondary VLAN on the Gateway, on page 115
vPC Fast Convergence Convergence	Introduced this feature.	7.0(3)I7(1)	Configuring Fast Convergence, on page 261 Configuring LACP vPC Convergence, on page 261
mac-address bpd source version 2 command.	Triggers STP to use the new Cisco MAC address (00:26:0b:xx:xx:xx).	7.0(3)I6(2)	Enabling STP to Use the Cisco MAC Address, on page 281
Configuring IP-in-IP tunnel decapsulation on directly connected IP addresses using tunnel source direct and tunnel mode ipv6ipv6 decapsulate-any	Added support for configuring IP-in-IP tunnel decapsulation on any directly connected IP addresses.	7.0(3)I6(1)	Guidelines and Limitations, on page 289 and Configuring a Tunnel Interface, on page 298
Configuring IP unnumbered interfaces for port channels	Added a guideline for configuring IP unnumbered interfaces for port channels on all Cisco Nexus 9000 Series switches.	7.0(3)I6(1)	IP Unnumbered, on page 96

Feature	Description	Changed in Release	Where Documented
Smart channel	Introduced this feature.	7.0(3)I6(1)	Configuring Smart-Channel, on page 325
Port profile support on Cisco Nexus 9500 Series switches	Updated the guideline to include port profile support on Cisco Nexus 9500 Series switches.	7.0(3)I5(2)	Guidelines and Limitations, on page 28
Support for QSA adapter	Updated the guideline to include QSA adapter support on Cisco Nexus 9200 and 9300-EX Series switches and Cisco Nexus 3232C and 3264Q Series switches.	7.0(3)I5(2)	Cisco QSFP+ to SFP+ Adapter Module Support, on page 27
Configuring Layer 3 over vPC	Added support for configuring Layer 3 over vPC.	7.0(3)I5(1)	Configuring Layer 3 over vPC, on page 251
Configuring selective Q-in-Q tunneling on a 802.1Q tunnel port	Added support for configuring selective Q-in-Q tunneling on a 802.1Q tunnel port.	7.0(3)I5(1)	Configuring VLAN Mapping for Selective Q-in-Q on a 802.1Q Tunnel Port, on page 317
Configuring BFD per-member link support (RFC 7130)	Added per-link Bidirectional Forwarding support that enables to configure individual BFD sessions on every Link Aggregation Group member interfaces (as defined in RFC 7130).	7.0(3)I5(1)	BFD Enhancement to Address Per-link Efficiency, on page 139
SFP+ adapter module	Added 25G optics support on the CVR-2QSFP28-8SFP adapter for 100G ports of the Cisco Nexus 9236C switch.	7.0(3)I4(2)	Cisco SFP+ Adapter Module Support, on page 27
Support for changing VRF membership for an SVI	Added support for changing VRF membership for an SVI.	7.0(3)I4(1)	Changing VRF Membership for an Interface
Support for port profiles.	Added support for port profiles.	7.0(3)I4(1)	Port Profiles
Support for PIM SSM over vPCs.	Added support for PIM SSM over vPCs.	7.0(3)I4(1)	vPC Multicast—PIM, IGMP, and IGMP Snooping
NAT support for In Service Software Upgrade (ISSU).	Added NAT support for In Service Software Upgrade (ISSU).	7.0(3)I4(1)	Guidelines and Limitations for Static NAT

Feature	Description	Changed in Release	Where Documented
Breakout and FEX support for the Cisco Nexus C9332PQ switch.	Added breakout and FEX support for the Cisco Nexus C9332PQ switch. (TOR switch)	7.0(3)I3(1)	Cisco Nexus C9332PQ Switch
Breakout support for the Cisco Nexus 92160YC switch.	Added breakout support for the Cisco Nexus 92160YC switch. (TOR switch)	7.0(3)I3(1)	Cisco Nexus C92160YC Switch
Breakout support for the Cisco Nexus 9272Q switch	Added breakout support for the Cisco Nexus 9272Q switch. (TOR switch)	7.0(3)I3(1)	Cisco Nexus C9272Q Switch
ip unnumbered command	Added support for IP unnumbered command.	7.0(3)I3(1)	IP Unnumbered
Layer 2 Data Center Interconnect	Added support for Layer 2 Data Center Interconnect.	7.0(3)I2(2)	Configuring Layer 2 Data Center Interconnect
shut and no shut support for vPC domain command	Added support for shut and no shut for vPC domain command.	7.0(3)I2(2)	Configuring vPCs
DHCP Client	Added support for Cisco Nexus 9500 Series switches.	7.0(3)I2(2)	DHCP Client
switchport isolated support	Added support for the switchport isolated command.	7.0(3)I2(1)	Switchport Isolated for up to 3967 VLANs on Trunk Interfaces
DHCP Client	Added support to configure the IPv4 or IPv6 address of a DHCP client on a management interface or a physical Ethernet interface.	7.0(3)I2(1)	DHCP Client
GRE tunnel enhancements	Added support for GRE v6 payload over v4 tunnel and GRE v4 payload over v6 tunnel.	7.0(3)I2(1)	Configuring a GRE Tunnel
Source interface support	Added source-interface command option support to configure an IPv4 and/or IPv6 inband or outband source IP address.	7.0(3)I2(1)	Guidelines and Limitations
regex support	Support for regular expression to address a set of interfaces.	7.0(3)I2(1)	Guidelines and Limitations
BFD startup timer	Introduced this feature.	7.0(3)I2(1)	Configuring Global BFD Parameters

Feature	Description	Changed in Release	Where Documented
Static and Dynamic NAT Translation support	Added support for static and dynamic NAT translation.	7.0(3)I2(1)	Configuring Static and Dynamic NAT Translation
Switch Virtual Interface and Autostate Behavior support	Added support for switch virtual interface and autostate behavior.	7.0(3)I2(1)	Switch Virtual Interface and Autostate Behavior
IP-in-IP tunnel mask support	Added support for IP-in-IP tunnel mask.	7.0(3)I2(1)	Creating a Tunnel Interface
Q-in-Q VLAN Tunnel support	Added support for Q-in-Q VLAN tunnels.	7.0(3)I2(1)	Configuring Q-in-Q VLAN Tunnels
MAC-embedded IPv6 (MEv6) address	Introduced this feature.	7.0(3)I2(1)	MAC-Embedded IPv6 Address
vPC forklift upgrade support	Support for upgrading from a pair of Nexus 9000 switches in a vPC topology to a different pair of Nexus 9000 series switches.	7.0(3)I1(2)	vPC Forklift Upgrade Scenario
Delayed LACP support	The delayed LACP feature enables the delay of port-channel members to come up until LACP PDUs are received.	7.0(3)I1(2)	LACP Compatibility Enhancements, on page 175
negotiate auto command support	Enables configuration of speed, duplex, and automatic flow control for an Ethernet interface.	7.0(3)I1(2)	Guidelines and Limitations
IP-in-IP tunnel support	Enables encapsulation and decapsulation of packets to create a tunnel.	7.0(3)I1(2)	Point-to-Point IP-in-IP Tunnel Encapsulation and Decapsulation
Subinterface support on port-channel interfaces	Support for one or more subinterfaces on a port-channel interface.	7.0(3)I1(2)	Configuring a Subinterface on a Port-Channel Interface
Dynamic Breakout support	Enables high bandwidth ports to be split into four breakout ports.	7.0(3)I1(1)	Dynamic Breakout (Per-port Level Breakout)

Feature	Description	Changed in Release	Where Documented
Symmetric Hashing support	When enabled on a port channel, symmetric hashing ensures bidirectional traffic uses the same physical interface.	7.0(3)I1(1)	Symmetric Hashing
Additional show interface tunnel command support.	Added support for statistics for interface counters.	7.0(3)I1(1)	Verifying the IP Tunnel Configuration
BFDv6 support	Added BFDv6 support.	7.0(3)I1(1)	Configuring Bidirectional Forwarding Detection, on page 127



CHAPTER 2

Overview

- [Licensing Requirements, on page 9](#)
- [Supported Platforms, on page 9](#)
- [About Interfaces, on page 9](#)
- [Virtual Device Contexts, on page 17](#)
- [High Availability for Interfaces, on page 17](#)

Licensing Requirements

For a complete explanation of Cisco NX-OS licensing recommendations and how to obtain and apply licenses, see the [Cisco NX-OS Licensing Guide](#) and the [Cisco NX-OS Licensing Options Guide](#).

Supported Platforms

Starting with Cisco NX-OS release 7.0(3)I7(1), use the [Nexus Switch Platform Support Matrix](#) to know from which Cisco NX-OS releases various Cisco Nexus 9000 and 3000 switches support a selected feature.

About Interfaces

Cisco NX-OS supports multiple configuration parameters for each of the interface types supported. Most of these parameters are covered in this guide but some are described in other documents.

The following table shows where to get further information on the parameters you can configure for an interface.

Table 2: Interface Parameters

Feature	Parameters	Further Information
Basic parameters	Description, duplex, error disable, flow control, MTU, beacon	“Configuring Basic Interface Parameters”
Layer 3	Medium, IPv4 and IPv6 addresses	“Configuring Layer 3 Interfaces”

Feature	Parameters	Further Information
Layer 3	Bandwidth, delay, IP routing, VRFs	<i>Cisco Nexus 9000 Series NX-OS Unicast Routing Configuration Guide</i> <i>Cisco Nexus 9000 Series NX-OS Multicast Routing Configuration Guide</i>
Port Channels	Channel group, LACP	“Configuring Port Channels”
Security	EOU	<i>Cisco Nexus 9000 Series NX-OS Security Configuration Guide</i>

Ethernet Interfaces

- Ethernet interfaces include routed ports.

Access Ports

An access port carries traffic for one VLAN. This type of port is a Layer 2 interface only.

For more information on access ports, see the “Information About Access and Trunk Interfaces” section.

Routed Ports

A routed port is a physical port that can route IP traffic to another device. A routed port is a Layer 3 interface only.

For more information on routed ports, see the “Routed Interfaces” section.

Management Interface

You can use the management Ethernet interface to connect the device to a network for remote management using a Telnet client, the Simple Network Management Protocol (SNMP), or other management agents. The management port (mgmt0) is autosensing and operates in full-duplex mode at a speed of 10/100/1000 Mb/s.

For more information on the management interface, see the [Cisco Nexus 9000 Series NX-OS Fundamentals Configuration Guide](#).

Port-Channel Interfaces

A port channel is a logical interface that is an aggregation of multiple physical interfaces. You can bundle up to 32 individual links (physical ports) into a port channel to improve bandwidth and redundancy. For more information about port-channel interfaces, see the “Configuring Port Channels” section.

Subinterfaces

You can create virtual subinterfaces using a parent interface configured as a Layer 3 interface. A parent interface can be either a physical port or a port-channel. A parent interface can be a physical port. Subinterfaces divide the parent interface into two or more virtual interfaces on which you can assign unique Layer 3 parameters such as IP addresses and dynamic routing protocols.

Loopback Interfaces

A loopback interface is a virtual interface with a single endpoint that is always up. Any packet that is transmitted over a virtual loopback interface is immediately received by that interface. Loopback interfaces emulate a physical interface. For more information about subinterfaces, see the “Loopback Interfaces” section.

Breakout Interfaces

Cisco NX-OS supports the breakout of a high bandwidth interface into one or more low bandwidth interfaces at the module level or at the per-port level.

Module Level Breakout

For module level breakout, the **interface breakout** command splits the high bandwidth 40G interface of a module into four 10G interfaces. The module is reloaded and the configuration for the interface is removed when the command is executed.

The following is an example of the command:

```
switch# configure terminal
switch(config)# interface breakout module 1
Module will be reloaded. Are you sure you want to continue(yes/no)? yes
```

The **no interface breakout module *module_number*** command undoes the breakout configuration. It puts all interfaces of the module in 40G mode and deletes the configuration for the previous 10G interfaces.

Dynamic Breakout (Per-port Level Breakout)

For dynamic breakout (per-port level breakout), the **interface breakout** command splits high bandwidth 40G ports into four 10G breakout ports and 100G ports into four 25G breakout ports. The breakout ports are identified as **Ethernet <slot>/<front-panel-port>/<breakout-port>**. For example, per-port level breakout ports can be identified as **Ethernet 1/2/1**, **Ethernet 1/2/2**, **Ethernet 1/2/3**, and **Ethernet 1/2/4**.

When one or more 40G interfaces of the module are broken-out at the per-port level, the configuration for the interfaces are removed when the command is executed.



Note A per-port level breakout does not require the module to be reloaded.

The following is an example of configuring a breakout port:

```
switch(config)# interface breakout module 1 port 1 map 10g-4x
switch(config)#
```

The following is an example of configuring a multiple breakout ports:

```
switch(config)# interface breakout module 1 port 1-4 map 10g-4x
switch(config)#
```

The following is an example of configuring 40G and 10G interfaces mixed together.

```
switch(config-if)# show int eth1/49 transceiver
Ethernet1/49
transceiver is present
type is QSFP-40G-SR-BD
name is CISCO-AVAGO
part number is AFBR-79EBPZ-CS2
revision is 01
```

```
switch(config-if)# show int eth1/52 transceiver
Ethernet1/52
transceiver is present
type is QSFP-Cazadero
name is CISCO-DNI
part number is CAZADERO-R
revision is 03
nominal bitrate is 10000 MBit/sec per channel
```

```
switch(config-if)# show int eth1/53 transceiver
Ethernet1/53
transceiver is present
type is QSFP-Cazadero
name is CISCO-DNI
part number is CAZADERO-R
revision is 03
nominal bitrate is 10000 MBit/sec per channel
```

```
switch(config)# interface breakout module 1 port 52-53 map 10g-4x
```

```
switch(config-if)# show int br | i up
mgmt0 -- up 10.122.160.192 100 1500
Eth1/49 -- eth routed up none 40G(D) -- << Running 40G
Eth1/50 -- eth routed up none 40G(D) --
Eth1/52/1 -- eth routed up none 10G(D) -- << Broken out to 10G
Eth1/53/1 -- eth routed up none 10G(D) -- << Broken out to 10G
```

You can undo the breakout ports with the **no interface breakout** command.

The following is an example to undo the breakout ports:

```
switch(config)# no interface breakout module 1 port 1 map 10g-4x
switch(config)#
```

About the Lane Selector

The lane selector is a push button switch and 4 LEDs located on the Cisco Nexus switch (left side of front panel, labeled "LS"). The push button switch and LEDs are used to indicate the status of the ports. The lane selector is supported on Cisco Nexus Series 9000 series switches and the Cisco Nexus 3164 and 3232 switches.

By default, the LEDs indicate the link/activity status of a 1 x 40G configuration. When the ports are configured as 4 x 10G, you can access the link status of each individual 10G port with the lane selector.

By pressing the lane selector push button, the port LED shows the selected lane's link/activity status. The 1st time the push button is pressed, the first LED displays the status of the first port. Pressing the push button a 2nd time displays the status of the second port, and so on. You can display the status of each of the four ports by pressing the push button in this manner.

For example, if port 60 is configured as 4 x 10G, pressing the lane selector push button once displays the link status of 60/1/1. Pressing the push button a second time displays the link status of 60/1/2.

When you press the push button after displaying the status of the last port, all four of the LEDs should extinguish to indicate that the lane selector has returned to display the status for the default 1 x 40G configuration.



Note A 10G breakout port's LED blinks when the beacon feature has been configured for it.



Note When a port is configured to be in 10G breakout mode and no lane is selected, the 40G port's LED illuminates as green even though only one of the 10G breakout ports is up.

Notes About Breakout Interfaces

Cisco Nexus 9516 switch does not support breakout on Modules 8 to 16.

Beginning with Cisco NX-OS Release 7.0(3)F2(1), the 36-port 100-Gigabit Ethernet QSFP28 line cards (N9K-X9636C-R) and 36-port 40-Gigabit Ethernet QSFP+ line cards (N9K-X9636Q-R) provide 4x10-Gigabit support.

The following table provides detailed information of the supported or not supported breakout modes. For more information, see [Cisco Nexus Data Sheets](#):

Caveats

- As of Cisco NX-OS Release 7.0(3)I7(2), manual breakout of QSA ports is not supported.

Manual breakout is supported on the following platforms because auto-breakout does not happen successfully on them—N9K-C93128TX, N9K-9332, N9K-C9396PX, N9K-C9396TX, N9K-C9372PX, N9K-C9372TX, N9K-C9332PQ, N9K-C93120TX, N9K-9432PQ, N9K-9536PQ, N9K-9636PQ, N9K-X9632PC-QSFP100, N9K-X9432C-S, N3K-C3132Q-V, N3K-C3164Q, N3K-C3132C, N3K-C3232C, N3K-C3264Q, N3K-C3264C, N3K-3064Q, N3K-3016, N3K-3172.

You need to perform manual breakout using "interface breakout module <module number> port <port range> map <breakout mapping>" command.

- When a break-out port is configured as a part of a port-channel, you need to apply the configuration twice (after write-erase/reload), to ensure the effectiveness of the port-channel.
- When you upgrade a Cisco Nexus 9000 device to Cisco NX-OS Release 7.0(3)I7(2) or later releases, if a QSFP port is configured with the manual breakout command and is using a QSA, the configuration of the interface Ethernet 1/50/1 is no longer supported and will need to be removed. To restore the configuration, you must manually configure the Ethernet 1/50 on the device.

This behaviour is not applicable to the following platforms—N9K-C93128TX, N9K-9332, N9K-C9396PX, N9K-C9396TX, N9K-C9372PX, N9K-C9372TX, N9K-C9332PQ, N9K-C93120TX, N9K-9432PQ, N9K-9536PQ, N9K-9636PQ, N9K-X9632PC-QSFP100, N9K-X9432C-S, N3K-C3132Q-V, N3K-C3164Q, N3K-C3132C, N3K-C3232C, N3K-C3264Q, N3K-C3264C, N3K-3064Q, N3K-3016, N3K-3172—because manual breakout is supported on these platforms.

- Cisco Nexus 9000 Series switches have 40G ports. When you breakout one of the 40G ports into 4x10G ports with a QSFP breakout cable, not all of the subinterfaces can be added to a port channel. The following error message is displayed:

```
switch# channel-group 99 mode active
command failed: port not compatible [Buffer boost]
```



Note As a workaround, configure **no buffer-boost** on all of the subinterfaces. This enables the channel-group configuration to go through.

Using the **force** keyword on the port-channel does not enable all the interfaces to be added to the port-channel even though the error message suggests using the keyword.

- Beginning with Cisco NX-OS Release 7.0(3)I7(3) you see two additional options to configure FEC such as **rs-cons16** and **rs-ieee** as per IEEE standards.
- Beginning with Cisco NX-OS Release 7.0(3)I7(7) you can display the admin and oper status of FEC interface information with the **show interface fec** command.

Example:

```
switch# show interface fec
-----
Name  Ifindex Admin-fec Oper-fec Status Speed Type
-----
Eth1/1  0x1a000000 auto auto connected 10G SFP-H10GB-AOC2M
Eth1/2  0x1a000200 Rs-fec notconnected auto QSFP-100G-AOC3M
Eth1/3/1 0x38014000 auto auto disabled auto QSFP-H40G-AOC3M
Eth1/3/2 0x38015000 auto auto disabled auto QSFP-H40G-AOC3M
Eth1/3/3 0x38016000 auto auto disabled auto QSFP-H40G-AOC3M
Eth1/3/4 0x38017000 auto auto disabled auto QSFP-H40G-AOC3M
```



Note Auto-FEC is not supported in Cisco NX-OS Release 7.0(3)I7(x)

High Bandwidth Interfaces

The breakout of high bandwidth interfaces (module level or per-port level) are supported only on:

- The X9636PQ, X9432PQ, X9536PQ, and X9732C-EX line cards on a Cisco Nexus 9500 Series switch.
- The Cisco Nexus 9332PQ switch.
- The Cisco Nexus 3164Q switch.

Cisco Nexus C92160YC Switch

For 7.0(3)I3(1) and later, the Cisco Nexus C92160YC switch provides two different modes of operation:

- Mode 1: 48 * 10G/25G + 4 * 40G + 2 * 100G (Default configuration)
 - Hardware profile portmode 48x25G + 2x100G + 4x40G
 - Breakout is supported in 2 * 100G ports

- Mode 2: 48 * 10G/25G + 4 * 100G
 - Hardware profile portmode 48x25G + 4x100G
 - Breakout is supported on 3 * 100G ports (Ports 50,51 & 52).

Use the **show running-config | grep portmode** command to display the current operation mode.

Example:

```
switch(config-if-range)# show running-config | grep portmode
hardware profile portmode 48x25G+2x100G+4x40G
```

See the installation guide for the Cisco Nexus C92160YC switches for more information. ([Install and Upgrade Guides for Cisco Nexus 9000 Series Switches](#))

With the Cisco Nexus C92160YC switch, there are two breakout modes:

- 40G to 4x10G breakout ports
 - Enables the breakout of 40G ports into 4 X 10G ports.
 - Use the **interface breakout module 1 port x map 10g-4x** command.
- 100G to 4x25G breakout ports
 - Enables the breakout of 100G ports into 4 X 25G ports.
 - Use the **interface breakout module 1 port x map 25g-4x** command.

Cisco Nexus C9272Q Switch

For 7.0(3)I3(1) and later, the Cisco Nexus C9272Q Switch provides 72 40G ports. Ports 37 - 71 support breakout interfaces.

To configure a breakout interface, use the **interface breakout module 1 port x map 10g-4x** command.

Example:

```
switch(config)# interface breakout module 1 port 38 map 10g-4x
switch(config)# show interface ethernet 1/38 capabilities | grep -i break

Breakout capable:      yes
```

Cisco Nexus C9332PQ Switch

For 7.0(3)I3(1) and later, the Cisco Nexus C9332PQ Switch provides 24 40G ports that support breakout mode and can be connected to four 10G NIF ports on a FEX. Ports 1 - 12 and 15 - 26 ports are supported. (Ports 13 and 14 ports are reserved and cannot be used for breakout mode.)



Note All FEXs are supported.



Note Only the Cisco Nexus 9332PQ switch provides interface breakout support for FEX fabric interfaces. (7.0(3)I3(1) and later)

Cisco Nexus 9000 C93180LC-EX Switch

For 7.0(3)I7(1) and later, Cisco Nexus 9000 C93180LC-EX switch provides three different modes of operation:

- Mode 1: 28 x 40G + 4 x 40G/100G (Default configuration)
 - Hardware profile portmode 4x100g + 28x40g.
 - 10x4 breakout is supported on the top ports from 1 to 27 (ports 1,3,5, 7...27). If any of the top port is broken out, the corresponding bottom port becomes non-operational. For example, if port 1 is broken out port 2 becomes non-operational.
 - 1 Gigabit and 10 Gigabit QSA is supported on ports 29, 30, 31, and 32. However, QSAs on the top and bottom front panel ports must be of same speed.
 - Ports 29, 30, 31, and 32 support 10x4, 25x4, and 50x2 breakout.
- Mode 2: 24 x 40G + 6 x 40G/100G
 - Hardware profile portmode 6x100g + 24x40g.
 - 10x4 breakout is supported on the top ports from 1 to 23 (ports 1,3,5, 7...23). If any of the top port is broken out the corresponding bottom port becomes non-operational.
 - Ports 25, 27, 29, 30, 31, and 32 support 10x4, 25x4, and 50x2 breakout.
 - 1 Gigabit and 10 Gigabit QSA is supported on ports 29, 30, 31, and 32. However, QSAs on the top and bottom front panel ports must be of same speed.
- Mode 3: 18 x 40G/100G
 - Hardware profile portmode 18x100g.
 - 10x4, 25x4, and 50x2 breakout is supported on top ports from 1 to 27 (ports 1,3,5, 7...27) and on ports 29,30,31,32.
 - 1 Gigabit and 10 Gigabit QSA is supported on all the 18 ports.

Changing Mode 3 to any other mode or vice versa requires **copy running-config startup-config** command followed by **reload** command to take effect. However, moving between Modes 1 and 2 is dynamic and requires only **copy running-config startup-config** command.

Use the **show running-config | grep portmode** command to display the current operation mode.

Example:

```
switch(config-if-range)# show running-config | grep portmode
hardware profile portmode 4x100G+28x40G
```

With the Cisco Nexus C93180LC-EX switch, there are three breakout modes:

- 40G to 4x10G breakout ports
 - Enables the breakout of 40G ports into 4 X 10G ports.
 - Use the **interface breakout module 1 port x map 10g-4x** command.
- 100G to 4x25G breakout ports
 - Enables the breakout of 100G ports into 4 X 25G ports.
 - Use the **interface breakout module 1 port x map 25g-4x** command.
- 100G to 2x50G breakout ports
 - Enables the breakout of 100G ports into 2 X 50G ports.
 - Use the **interface breakout module 1 port x map 50g-2x** command.

Virtual Device Contexts

Cisco NX-OS can segment operating system and hardware resources into virtual device contexts (VDCs) that emulate virtual devices. The Cisco Nexus 9000 Series switch does not support multiple VDCs. All switch resources are managed in the default VDC.

High Availability for Interfaces

Interfaces support stateful and stateless restarts. A stateful restart occurs on a supervisor switchover. After the switchover, Cisco NX-OS applies the runtime configuration.



CHAPTER 3

Configuring Basic Interface Parameters

This chapter describes how to configure the basic interface parameters on Cisco NX-OS devices.

- [About the Basic Interface Parameters, on page 19](#)
- [Guidelines and Limitations, on page 28](#)
- [Default Settings, on page 30](#)
- [Configuring the Basic Interface Parameters, on page 31](#)
- [Verifying the Basic Interface Parameters, on page 56](#)
- [Monitoring the Interface Counters, on page 56](#)
- [Configuration Example for QSA, on page 58](#)

About the Basic Interface Parameters

Description

For the Ethernet and management interfaces, you can configure the description parameter to provide a recognizable name for the interface. Using a unique name for each interface allows you to quickly identify the interface when you are looking at a listing of multiple interfaces.

For information about setting the description parameter for port-channel interfaces, see the “Configuring a Port-Channel Description” section. For information about configuring this parameter for other interfaces, see the “Configuring the Description” section.

Beacon

The beacon mode allows you to identify a physical port by flashing its link state LED with a green light. By default, this mode is disabled. To identify the physical port for an interface, you can activate the beacon parameter for the interface.

For information about configuring the beacon parameter, see the “Configuring the Beacon Mode” section.

Error Disabled

A port is in the error-disabled (err-disabled) state when the port is enabled administratively (using the **no shutdown** command) but disabled at runtime by any process. For example, if UDLD detects a unidirectional link, the port is shut down at runtime. However, because the port is administratively enabled, the port status

displays as `err-disable`. Once a port goes into the `err-disable` state, you must manually reenable it or you can configure a timeout value that provides an automatic recovery. By default, the automatic recovery is not configured, and by default, the `err-disable` detection is enabled for all causes.

When an interface is in the `err-disabled` state, use the **`errdisable detect cause`** command to find information about the error.

You can configure the automatic error-disabled recovery timeout for a particular error-disabled cause and configure the recovery period.

The **`errdisable recovery cause`** command provides an automatic recovery after 300 seconds.

You can use the **`errdisable recovery interval`** command to change the recovery period within a range of 30 to 65535 seconds. You can also configure the recovery timeout for a particular `err-disable` cause.

If you do not enable the error-disabled recovery for the cause, the interface stays in the error-disabled state until you enter the **`shutdown`** and **`no shutdown`** commands. If the recovery is enabled for a cause, the interface is brought out of the error-disabled state and allowed to retry operation once all the causes have timed out. Use the **`show interface status err-disabled`** command to display the reason behind the error.

Interface Status Error Policy

Cisco NX-OS policy servers such as Access Control List (ACL) Manager and Quality of Service (QoS) Manager, maintain a policy database. A policy is defined through the command-line interface.

Policies are pushed when you configure a policy on an interface to ensure that policies that are pushed are consistent with the hardware policies. To clear the errors and to allow the policy programming to proceed with the running configuration, enter the **`no shutdown`** command. If the policy programming succeeds, the port is allowed to come up. If the policy programming fails, the configuration is inconsistent with the hardware policies and the port is placed in an error-disabled policy state. The error-disabled policy state remains and the information is stored to prevent the same port from being brought up in the future. This process helps to avoid unnecessary disruption to the system.

Modifying Interface MTU Size

The maximum transmission unit (MTU) size specifies the maximum frame size that an Ethernet port can process. For transmissions to occur between two ports, you must configure the same MTU size for both ports. A port drops any frames that exceed its MTU size.

By default, the Cloud-Scale ASIC NX-OS system always allows an extra 166B in the MTU on top of the configured value in order to fully support/accept different types of encapsulations in the hardware.

Cisco NX-OS allows you to configure MTU on an interface, with options to configure it on different level in the protocol stack. By default, each interface has an MTU of 1500 bytes, which is the IEEE 802.3 standard for Ethernet frames. Larger MTU sizes are possible for more efficient processing of data to allow different application requirements. The larger frames, are also called jumbo frames, can be up to 9216 bytes in size.

MTU is configured per interface, where an interface can be a Layer 2 or a Layer 3 interface. For a Layer 2 interface, you can configure the MTU size with one of two values, the value system default MTU value or the system jumbo MTU value. The system default MTU value is 1500 bytes. Every Layer 2 interface is configured with this value by default. You can configure an interface with the default system jumbo MTU value, that is 9216 bytes. To allow an MTU value from 1500 through 9216, you must adjust the system jumbo MTU to an appropriate value where interface can be configured with the same value.



Note You can change the system jumbo MTU size. When the value is changed, the Layer 2 interfaces that use the system jumbo MTU value, will automatically changes to the new system jumbo MTU value.

A Layer 3 interface, can be Layer 3 physical interface (configure with no switchport), switch virtual interface (SVI), and sub-interface, you can configure an MTU size between 576 and 9216 bytes.

For the Cisco Nexus 9372 switch, the following applies:

- The 10-G interfaces are mapped to specific hardware ports where the default MTU is 1500.
- The 40-G interfaces are mapped as a HiGiG port where the default MTU is 3FFF and the MTU limit check is disabled.
- In the case of 40-G interfaces, since the MTU limit check is disabled, it ignores the packet size and traffic flows irrespective of its MTU.
- When the configured MTU of all interfaces on the switch do not match, the switch's behavior may vary depending on the specific port that is mismatched as well as the traffic flow. The following are examples of the switch's behavior in various scenarios:
 - When a Layer 3 port receives a frame whose length exceeds the port's MTU size, the port will drop the frame.
 - When a Layer 3 port receives a frame whose length is less than the ingress port's MTU size, but greater than the egress Layer 3 port's MTU size, then the frame is punted to the supervisor of the switch.
 1. If the frame is an IP packet that has the Don't Fragment (DF) bit set, then the frame will be dropped in software. Otherwise, the frame will be fragmented in software.
 2. Otherwise, the frame will be fragmented in software.
 3. This can cause performance issues (such as increased latency or packet loss for affected traffic flows) due to Control Plane Policing (CoPP) enabled by default on Cisco Nexus switches. For more information about Control Plane Policing, refer to the **Configuring Control Plane Policing** chapter of the **Cisco Nexus 9000 Series NX-OS Security Configuration Guide**.
 - When a Layer 2 port receives a frame whose length exceeds the port's MTU size, the port will drop the frame.
 - When a Layer 2 port receives a frame whose length is less than the ingress port's MTU size, but greater than the egress Layer 2 port's MTU size, and the frame is routed between VLANs by the switch, then the frame is punted to the supervisor of the switch.
 1. If the frame is an IP packet that has the Don't Fragment (DF) bit set, then the frame will be dropped in software. Otherwise, the frame will be fragmented in software.
 2. Otherwise, the frame will be fragmented in software.
 3. This can cause performance issues (such as increased latency or packet loss for affected traffic flows) due to Control Plane Policing (CoPP) enabled by default on Cisco Nexus switches. For more information about Control Plane Policing, refer to the **Configuring Control Plane Policing** chapter of the **Cisco Nexus 9000 Series NX-OS Security Configuration Guide**.

- When a Layer 2 port receives a frame whose length is less than the ingress port's MTU size, but greater than the egress Layer 2 port's MTU size, and the frame is switched within the same VLAN by the switch, then the switch will drop the frame.

For information about setting the MTU size, see the *Configuring the MTU Size* section.



Note On Cisco Nexus 9300-FX2 and 9300-GX devices, if ingress interface is configured with an MTU less than 9216, FTE does not capture input errors and does not display any events. However, if the ingress interface is configured with an MTU of 9216, FTE displays all the events.

Bandwidth

Ethernet ports have a fixed bandwidth of 1,000,000 Kb at the physical layer. Layer 3 protocols use a bandwidth value that you can set for calculating their internal metrics. The value that you set is used for informational purposes only by the Layer 3 protocols—it does not change the fixed bandwidth at the physical layer. For example, the Enhanced Interior Gateway Routing Protocol (EIGRP) uses the minimum path bandwidth to determine a routing metric, but the bandwidth at the physical layer remains at 1,000,000 Kb.

For information about configuring the bandwidth parameter for port-channel interfaces, see the “Configuring the Bandwidth and Delay for Informational Purposes” section. For information about configuring the bandwidth parameter for other interfaces, see the “Configuring the Bandwidth” section.

Throughput Delay

Specifying a value for the throughput-delay parameter provides a value used by Layer 3 protocols; it does not change the actual throughput delay of an interface. The Layer 3 protocols can use this value to make operating decisions. For example, the Enhanced Interior Gateway Routing Protocol (EIGRP) can use the delay setting to set a preference for one Ethernet link over another, if other parameters such as link speed are equal. The delay value that you set is in the tens of microseconds.

For information about configuring the bandwidth parameter for port-channel interfaces, see the “Configuring the Bandwidth and Delay for Informational Purposes” section. For information about configuring the throughput-delay parameter for other interfaces, see the “Configuring the Throughput Delay” section.

Administrative Status

The administrative-status parameter determines whether an interface is up or down. When an interface is administratively down, it is disabled and unable to transmit data. When an interface is administratively up, it is enabled and able to transmit data.

For information about configuring the administrative status parameter for port-channel interfaces, see the “Shutting Down and Restarting the Port-Channel Interface” section. For information about configuring the administrative-status parameter for other interfaces, see the “Shutting Down and Activating the Interface” section.

Unidirectional Link Detection Parameter

UDLD Overview

The Cisco-proprietary Unidirectional Link Detection (UDLD) protocol allows devices that are connected through fiber-optic 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 a device detects a unidirectional link, UDLD shuts down the affected LAN port and alerts the user. Unidirectional links can cause a variety of problems.

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 detections work 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, UDLD 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.


The Cisco Nexus 9000 Series device 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. You can configure the transmission interval for the UDLD frames, either globally or for the specified interfaces.



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

The 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 1: Unidirectional Link



Default UDLD Configuration

The following table shows the default UDLD configuration.

Table 3: UDLD Default Configuration

Feature	Default Value
UDLD global enable state	Globally 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
UDLD aggressive mode	Disabled
UDLD message interval	15 seconds

For information about configuring the UDLD for the device and its port, see the “Configuring the UDLD Mode” section.

UDLD Normal and Aggressive Modes

UDLD supports Normal and Aggressive modes of operation. By default, Normal mode is enabled.

In Normal mode, UDLD detects the following link errors by examining the incoming UDLD packets from the peer port:

- Empty echo packet
- Uni-direction
- TX/RX loop
- Neighbor mismatch

By default, UDLD aggressive mode is disabled. 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 frame, UDLD tries to re-establish the connection with the neighbor. After eight failed retries, the port is disabled.

In the following scenarios, enabling the UDLD aggressive mode disables one of the ports to prevent the discarding of traffic.

- 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



Note You enable the UDLD aggressive mode globally to enable that mode on all the fiber ports. You must enable the UDLD aggressive mode on copper ports on specified interfaces.



Tip When a line card upgrade is being performed during an in-service software upgrade (ISSU) and some of the ports on the line card are members of a Layer 2 port channel and are configured with UDLD aggressive mode, if you shut down one of the remote ports, UDLD puts the corresponding port on the local device into an error-disabled state. This behavior is correct.

To restore service after the ISSU has completed, enter the **shutdown** command followed by the **no shutdown** command on the local port.

Port-Channel Parameters

A port channel is an aggregation of physical interfaces that comprise a logical interface. You can bundle up to 32 individual interfaces into a port channel to provide increased bandwidth and redundancy. Port channeling also load balances traffic across these physical interfaces. The port channel stays operational if at least one physical interface within the port channel is operational.

You can create Layer 3 port channels by bundling compatible Layer 3 interfaces.

Any configuration changes that you apply to the port channel are applied to each interface member of that port channel.

For information about port channels and for information about configuring port channels, see Chapter 6, “Configuring Port Channels.”

Port Profiles

Beginning with 7.0(3)I4(1) on Cisco Nexus 9300 Series switches, you can create a port profile that contains many interface commands and apply that port profile to a range of interfaces. Each port profile can be applied only to a specific type of interface; the choices are as follows:

- Ethernet
- VLAN network interface
- Port channel

When you choose Ethernet or port channel as the interface type, the port profile is in the default mode which is Layer 3. Enter the **switchport** command to change the port profile to Layer 2 mode.

You inherit the port profile when you attach the port profile to an interface or range of interfaces. When you attach, or inherit, a port profile to an interface or range of interfaces, the system applies all the commands in that port profile to the interfaces. Additionally, you can have one port profile inherit the settings from another port profile. Inheriting another port profile allows the initial port profile to assume all of the commands of the second, inherited, port profile that do not conflict with the initial port profile. Four levels of inheritance are supported. The same port profile can be inherited by any number of port profiles.

The system applies the commands inherited by the interface or range of interfaces according to the following guidelines:

- Commands that you enter under the interface mode take precedence over the port profile’s commands if there is a conflict. However, the port profile retains that command in the port profile.
- The port profile’s commands take precedence over the default commands on the interface, unless the port-profile command is explicitly overridden by the default command.

- When a range of interfaces inherits a second port profile, the commands of the initial port profile override the commands of the second port profile if there is a conflict.
- After you inherit a port profile onto an interface or range of interfaces, you can override individual configuration values by entering the new value at the interface configuration level. If you remove the individual configuration values at the interface configuration level, the interface uses the values in the port profile again.
- There are no default configurations associated with a port profile.

A subset of commands are available under the port-profile configuration mode, depending on which interface type you specify.



Note You cannot use port profiles with Session Manager. See the *Cisco Nexus 9000 Series NX-OS System Management Configuration Guide* for information about Session Manager.

To apply the port-profile configurations to the interfaces, you must enable the specific port profile. You can configure and inherit a port profile onto a range of interfaces prior to enabling the port profile. You would then enable that port profile for the configurations to take effect on the specified interfaces.

If you inherit one or more port profiles onto an original port profile, only the last inherited port profile must be enabled; the system assumes that the underlying port profiles are enabled.

When you remove a port profile from a range of interfaces, the system undoes the configuration from the interfaces first and then removes the port-profile link itself. Also, when you remove a port profile, the system checks the interface configuration and either skips the port-profile commands that have been overridden by directly entered interface commands or returns the command to the default value.

If you want to delete a port profile that has been inherited by other port profiles, you must remove the inheritance before you can delete the port profile.

You can also choose a subset of interfaces from which to remove a port profile from among that group of interfaces that you originally applied the profile. For example, if you configured a port profile and configured ten interfaces to inherit that port profile, you can remove the port profile from just some of the specified ten interfaces. The port profile continues to operate on the remaining interfaces to which it is applied.

If you delete a specific configuration for a specified range of interfaces using the interface configuration mode, that configuration is also deleted from the port profile for that range of interfaces only. For example, if you have a channel group inside a port profile and you are in the interface configuration mode and you delete that port channel, the specified port channel is also deleted from the port profile as well.

Just as in the device, you can enter a configuration for an object in port profiles without that object being applied to interfaces yet. For example, you can configure a virtual routing and forward (VRF) instance without it being applied to the system. If you then delete that VRF and related configurations from the port profile, the system is unaffected.

After you inherit a port profile on an interface or range of interfaces and you delete a specific configuration value, that port-profile configuration is not operative on the specified interfaces.

If you attempt to apply a port profile to the wrong type of interface, the system returns an error.

When you attempt to enable, inherit, or modify a port profile, the system creates a checkpoint. If the port-profile configuration fails, the system rolls back to the prior configuration and returns an error. A port profile is never only partially applied.

Cisco QSFP+ to SFP+ Adapter Module Support

The Cisco QSFP+ to SFP+ Adapter (QSA) module provides 10G support for the 40G uplink ports that are a part of the Cisco Nexus M6PQ and Cisco Nexus M12PQ uplink modules of specific Cisco Nexus 9300 devices.

A group of six consecutive ports in the M6PQ or M12PQ uplink module must be operating at the same speed (40G or 10G) to use the QSA/QSFP modules.

- For Cisco Nexus 9396PX devices, 2/1-6 ports form the first port speed group and the remaining 2/7-12 ports form the second port speed group.
- For Cisco Nexus 93128PX/TX devices, 2/1-6 ports form the first port speed group and the remaining 2/7-8 ports form the second port speed group.
- For Cisco Nexus 937xPX/TX devices, 1/49-54 ports form the only port speed group.
- For Cisco Nexus 93120TX devices, 1/97-102 ports form the only port speed group.
- For Cisco Nexus 9332PQ devices, 1/27-32 ports form the only port speed group.

Use the **speed-group 10000** command to configure the first port of a port speed group for the QSA. This command specifies the administrator speed preference for the port group. (The default port speed is 40G.)

- The **speed-group 10000** command specifies a speed of 10G.
- The **no speed-group 10000** command specifies a speed of 40G.
- If you add, modify, or remove the speed-group command on any interface in the speed group, all interfaces in the speed group will change to the newly configured speed. This may cause link state changes based on the transceiver that is plugged in. This also includes the removal of an existing speed-group command using the default interface command.
- Uplink modules should not be removed from a Cisco Nexus 9300 platform switch that is running Cisco NX-OS Release 7.0(3)I7(5). The ports on uplink modules should be used only for uplinks.

After the speed has been configured, the compatible transceiver modules are enabled. The remaining transceiver modules in the port group (incompatible transceiver modules) become error disabled with a reason of "check speed-group config".



Note The Cisco QSFP+ to SFP+ Adapter (QSA) module does not provide 10G support for the 40G line cards for Cisco Nexus 9500 devices.

Beginning with Cisco NX-OS Release 7.0(3)I5(2), you can use a QSFP-to-SFP adapter on Cisco Nexus 9200 and 9300-EX Series switches and Cisco Nexus 3232C and 3264Q Series switches.

Cisco SFP+ Adapter Module Support

Beginning with Cisco NX-OS Release 7.0(3)I4(2), you can use the CVR-2QSFP28-8SFP adapter for 25-Gigabit optics support on 100-Gigabit ports of the Cisco Nexus 9236C switch.

The **interface breakout module** command can be used to split this switch's 100G interfaces into four 25G interfaces. After you enter this command, you must copy the running configuration to the startup configuration.

25G Autonegotiation Overview

Guidelines and Limitations

Basic interface parameters have the following configuration guidelines and limitations:

- MDIX is enabled by default on copper ports. It is not possible to disable it.
- **show** commands with the **internal** keyword are not supported.
- Fiber-optic Ethernet ports must use Cisco-supported transceivers. To verify that the ports are using Cisco-supported transceivers, use the **show interface transceivers** command. Interfaces with Cisco-supported transceivers are listed as functional interfaces.
- A port can be either a Layer 2 or a Layer 3 interface; it cannot be both simultaneously.

By default, each port is a Layer 3 interface.

You can change a Layer 3 interface into a Layer 2 interface by using the **switchport** command. You can change a Layer 2 interface into a Layer 3 interface by using the **no switchport** command.

- You usually configure Ethernet port speed and duplex mode parameters to auto to allow the system to negotiate the speed and duplex mode between ports. If you decide to configure the port speed and duplex modes manually for these ports, consider the following:
 - Before you configure the speed and duplex mode for an Ethernet or management interface, see the Default Settings section for the combinations of speeds and duplex modes that can be configured at the same time.
 - If you set the Ethernet port speed to auto, the device automatically sets the duplex mode to auto.
 - If you enter the **no speed** command, the device automatically sets both the speed and duplex parameters to auto (the **no speed** command produces the same results as the **speed auto** command).
 - If you configure an Ethernet port speed to a value other than auto (for example, 1G, 10G, or 40G), you must configure the connecting port to match. Do not configure the connecting port to negotiate the speed.
 - To configure speed, duplex, and automatic flow control for an Ethernet interface, you can use the **negotiate auto** command. To disable automatic negotiation, use the **no negotiate auto** command.



Note The device cannot automatically negotiate the Ethernet port speed and duplex mode if the connecting port is configured to a value other than auto.



Caution Changing the Ethernet port speed and duplex mode configuration might shut down and reenable the interface.

- When using a QSFP-40G-CR4 cable to connect between N9K-C9332PQ non-ALE ports and N9K-C9372PX ALE ports, you need to manually set the speed to 40000.

- For BASE-T copper ports, auto-negotiation is enabled even when fixed speed is configured.
- For 7.0(3)I2(1) and later, using a regular expression to address a set of interfaces is supported with the **regex** command option. The **regex** command option is an extension that is available for all interface commands.

Example:

```
switch(config-if-range)# interface ethernet regex [2]/
switch(config-if-range)# where
conf; interface Ethernet2/1-8 admin@switch%default
switch(config-if-range)# interface ethernet regex [1]/2[2-4]
switch(config-if-range)# where
conf; interface Ethernet1/22-24 admin@switch%default
```

- For 7.0(3)I2(1) and later, the source-interface command option provides support for management applications to configure an IPv4 and/or IPv6 inband or outband source IP address for the copy command and other processes (such as tacacs, ntp, ping/ping6, icmp-error and traceroute).

- Configuration commands

ip services source-interface interface vrf vrf name

Examples:

- ip ftp source-interface ethernet 8/1 vrf management
- ip http source-interface loopback 1 vrf blue
- ip ssh source-interface ethernet ethernet 5/1
/*This command executes in the VRF context.*/
- ip ping source-interface ethernet 8/1 vrf blue
- ip traceroute source-interface ethernet 8/1 vrf red
- ip icmp-errors source-interface ethernet 8/1
/*This command executes in the VRF context.*/

- Show commands:

show ip copy services source-interface interface vrf vrf name

- show ip ftp source-interface ethernet 8/1 vrf management
- show ip http source-interface loopback 1 vrf blue
- show ip ssh source-interface ethernet ethernet 5/1
/*This command executes in the VRF context.*/
- show ip ping source-interface ethernet 8/1 vrf blue
- show ip traceroute source-interface ethernet 8/1 vrf red
- show ip icmp-errors source-interface ethernet 8/1
/*This command executes in the VRF context.*/

- Service commands:

copy *service://username@hostname/path file source-interface interface name*

Examples:

- `copy ftp://username@hostname/usr/local/bin file source-interface ethernet 8/1`
- `copy scp://username@hostname/usr/local/bin file source-interface ethernet 8/1`
- `copy tftp://username@hostname/usr/local/bin file source-interface ethernet 8/1`
- `copy http://username@hostname/usr/local/bin file source-interface ethernet 8/1`
- `copy sftp://username@hostname/usr/local/bin file source-interface ethernet 8/1`

- Beginning with Cisco NX-OS Release 7.0(3)I5(2), port profiles are supported on Cisco Nexus 9300 Series switches and Cisco Nexus 9500 Series switches.
- When using a QSFP-40G-CR4 on Cisco Nexus 9000 switches, you must configure the default speed as 40G in the auto-negotiation parameters. Otherwise, the interface may not be able to bring the link up.
- When you use the QSFP ports on N9K-X9788TC-FX line card, you must set the speed appropriately and disable auto negotiation on ports 49 to 52. You must also set the speed and disable auto negotiation on the peering device.
- The following line cards do not support Link Training:
 - Nexus 9300 Modules:
 - N9K-M12PQ (C9396PX, C9396TX, C93128PX, C93128TX)
 - Nexus 9500 Modules:
 - X9536PQ
 - X9564PX
 - X9564TX
- If cable length is more than 5 meters, Auto Negotiation is not supported. This cable length limitation is applicable only to copper cables and not applicable to optical cables.

Default Settings

The following lists the default settings for the basic interface parameters.

Parameter	Default
Description	Blank
Beacon	Disabled
Bandwidth	Data rate of interface
Throughput delay	100 microseconds

Parameter	Default
Administrative status	Shutdown
MTU	1500 bytes
UDLD global	Globally disabled
UDLD per-port enable state for fiber-optic media	Enabled on all Ethernet fiber-optic LAN ports
UDLD per-port enable state for copper media	Disabled on all Ethernet 1G, 10G, or 40G LAN ports
UDLD message interval	Disabled
UDLD aggressive mode	Disabled
Error disable	Disabled
Error disable recovery	Disabled
Error disable recovery interval	300 seconds
Buffer-boost	Enabled Note Feature available on N9K-X9564TX and N9K-X9564PX line cards and Cisco Nexus 9300 series devices.

Configuring the Basic Interface Parameters

When you configure an interface, you must specify the interface before you can configure its parameters.

Specifying the Interfaces to Configure

Before you begin

Before you can configure the parameters for one or more interfaces of the same type, you must specify the type and the identities of the interfaces.

The following table shows the interface types and identities that you should use for specifying the Ethernet and management interfaces.

Table 4: Information Needed to Identify an Interface for Configurations

Interface Type	Identity
Ethernet	I/O module slot numbers and port numbers on the module
Management	0 (for port 0)

The interface range configuration mode allows you to configure multiple interfaces with the same configuration parameters. After you enter the interface range configuration mode, all command parameters you enter are attributed to all interfaces within that range until you exit out of the interface range configuration mode.

You enter a range of interfaces using dashes (-) and commas (.). Dashes separate contiguous interfaces and commas separate noncontiguous interfaces. When you enter noncontiguous interfaces, you must enter the media type for each interface.

This example shows how to configure a contiguous interface range:

```
switch(config)# interface ethernet 2/29-30
switch(config-if-range)#
```

This example shows how to configure a noncontiguous interface range:

```
switch(config)# interface ethernet 2/29, ethernet 2/33, ethernet 2/35
switch(config-if-range)#
```

You can specify subinterfaces in a range only when the subinterfaces are on the same port, for example, 2/29.1-2. But you cannot specify the subinterfaces in a range of ports, for example, you cannot enter 2/29.2-2/30.2. You can specify two of the subinterfaces discretely, for example, you can enter 2/29.2, 2/30.2.

This example shows how to configure a breakout cable:

```
switch(config)# interface ethernet 1/2/1
switch(config-if-range)#
```

Procedure

	Command or Action	Purpose
Step 1	configure terminal Example: <pre>switch# configure terminal switch(config)#</pre>	Enters global configuration mode.
Step 2	interface interface Example: <pre>switch(config)# interface ethernet 2/1 switch(config-if)#</pre> Example: <pre>switch(config)# interface mgmt0 switch(config-if)#</pre>	Specifies the interface that you are configuring. You can specify the interface type and identity. For an Ethernet port, use ethernet slot/port . For the management interface, use mgmt0 . Examples: <ul style="list-style-type: none"> • The 1st example shows how to specify the slot 2, port 1 Ethernet interface. • The 2nd example shows how to specify the management interface. Note You do not need to add a space between the interface type and identity (port or slot/port number). For example, for the Ethernet slot 4, port 5 interface, you can specify either “ethernet 4/5” or “ethernet4/5.” The management interface is either “mgmt0” or “mgmt 0.”

	Command or Action	Purpose
		When you are in the interface configuration mode, the commands that you enter configure the interface that you specified for this mode.

Configuring the Description

You can provide textual interface descriptions for the Ethernet and management interfaces.

Procedure

	Command or Action	Purpose
Step 1	configure terminal Example: switch# configure terminal switch(config)#	Enters global configuration mode.
Step 2	interface interface Example: switch(config)# interface ethernet 2/1 switch(config-if)# Example: switch(config)# interface mgmt0 switch(config-if)#	Specifies the interface that you are configuring. You can specify the interface type and identity. For an Ethernet port, use ethernet slot/port . For the management interface, use mgmt0 . Examples: <ul style="list-style-type: none"> • The 1st example shows how to specify the slot 2, port 1 Ethernet interface. • The 2nd example shows how to specify the management interface.
Step 3	description text Example: switch(config-if)# description Ethernet port 3 on module 1 switch(config-if)#	Specifies the description for the interface.
Step 4	show interface interface Example: switch(config)# show interface ethernet 2/1	(Optional) Displays the interface status, which includes the description parameter.
Step 5	exit Example: switch(config-if)# exit switch(config)#	Exits the interface mode.
Step 6	copy running-config startup-config Example:	(Optional) Copies the running configuration to the startup configuration.

	Command or Action	Purpose
	switch(config)# copy running-config startup-config	

Example

This example shows how to set the interface description to Ethernet port 24 on module 3:

```
switch# configure terminal
switch(config)# interface ethernet 3/24
switch(config-if)# description server1
switch(config-if)#
```

The output of the **show interface eth** command is enhanced as shown in the following example:

```
Switch# show version
Software
BIOS: version 06.26
NXOS: version 6.1(2)I2(1) [build 6.1(2)I2.1]
BIOS compile time: 01/15/2014
NXOS image file is: bootflash:///n9000-dk9.6.1.2.I2.1.bin
NXOS compile time: 2/25/2014 2:00:00 [02/25/2014 10:39:03]
```

```
switch# show interface ethernet 6/36
Ethernet6/36 is up
admin state is up, Dedicated Interface
Hardware: 40000 Ethernet, address: 0022.bdf6.bf91 (bia 0022.bdf8.2bf3)
Internet Address is 192.168.100.1/24
MTU 9216 bytes, BW 40000000 Kbit, DLY 10 usec
```

Configuring the Beacon Mode

You can enable the beacon mode for an Ethernet port to flash its LED to confirm its physical location.

Procedure

	Command or Action	Purpose
Step 1	configure terminal Example: switch# configure terminal switch(config)#	Enters global configuration mode.
Step 2	interface ethernet <i>slot/port</i> Example: switch(config)# interface ethernet 3/1 switch(config-if)#	Specifies an interface to configure, and enters interface configuration mode.
Step 3	[no] beacon Example: switch(config)# beacon switch(config-if)#	Enables the beacon mode or disables the beacon mode. The default mode is disabled.

	Command or Action	Purpose
Step 4	show interface ethernet <i>slot/port</i> Example: <pre>switch(config)# show interface ethernet 2/1 switch(config-if)#</pre>	(Optional) Displays the interface status, which includes the beacon mode state.
Step 5	exit Example: <pre>switch(config-if)# exit switch(config)#</pre>	Exits the interface mode.
Step 6	copy running-config startup-config Example: <pre>switch(config)# copy running-config startup-config</pre>	(Optional) Copies the running configuration to the startup configuration.

Example

This example shows how to enable the beacon mode for the Ethernet port 3/1:

```
switch# configure terminal
switch(config)# interface ethernet 3/1
switch(config-if)# beacon
switch(config-if)#
```

This example shows how to disable the beacon mode for the Ethernet port 3/1:

```
switch# configure terminal
switch(config)# interface ethernet 3/1
switch(config-if)# no beacon
switch(config-if)#
```

This example shows how to configure the dedicated mode for Ethernet port 4/17 in the group that includes ports 4/17, 4/19, 4/21, and 4/23:

```
switch# configure terminal
switch(config)# interface ethernet 4/17, ethernet 4/19, ethernet 4/21, ethernet 4/23
switch(config-if)# shutdown
switch(config-if)# interface ethernet 4/17
switch(config-if)# no shutdown
switch(config-if)#
```

Configuring the Error-Disabled State

You can view the reason that an interface moves to the error-disabled state and configure automatic recovery.

Enabling the Error-Disable Detection

You can enable error-disable detection in an application. As a result, when a cause is detected on an interface, the interface is placed in an error-disabled state, which is an operational state that is similar to the link-down state.

Procedure

	Command or Action	Purpose
Step 1	configure terminal Example: switch# configure terminal switch(config)#	Enters global configuration mode.
Step 2	errdisable detect cause {acl-exception all link-flap loopback} Example: switch(config)# errdisable detect cause all switch(config-if)#	Specifies a condition under which to place the interface in an error-disabled state. The default is enabled.
Step 3	shutdown Example: switch(config-if)# shutdown switch(config)#	Brings the interface down administratively. To manually recover the interface from the error-disabled state, enter this command first.
Step 4	no shutdown Example: switch(config-if)# no shutdown switch(config)#	Brings the interface up administratively and enables the interface to recover manually from the error-disabled state.
Step 5	show interface status err-disabled Example: switch(config)# show interface status err-disabled	(Optional) Displays information about error-disabled interfaces.
Step 6	copy running-config startup-config Example: switch(config)# copy running-config startup-config	(Optional) Copies the running configuration to the startup configuration.

Example

This example shows how to enable the error-disabled detection in all cases:

```
switch(config)# errdisable detect cause all
switch(config)#
```

Enabling the Error-Disabled Recovery

You can specify the application to bring the interface out of the error-disabled state and retry coming up. It retries after 300 seconds, unless you configure the recovery timer (see the **errdisable recovery interval** command).

Procedure

	Command or Action	Purpose
Step 1	configure terminal Example: switch# configure terminal switch(config)#	Enters global configuration mode.
Step 2	errdisable recovery cause {all bpduguard failed-port-state link-flap loopback miscabling psecure-violation security-violation storm-control udd vpc-peerlink} Example: switch(config)# errdisable recovery cause all switch(config-if)#	Specifies a condition under which the interface automatically recovers from the error-disabled state, and the device retries bringing the interface up. The device waits 300 seconds to retry. The default is disabled.
Step 3	show interface status err-disabled Example: switch(config)# show interface status err-disabled switch(config-if)#	(Optional) Displays information about error-disabled interfaces.
Step 4	copy running-config startup-config Example: switch(config)# copy running-config startup-config	(Optional) Copies the running configuration to the startup configuration.

Example

This example shows how to enable error-disabled recovery under all conditions:

```
switch(config)# errdisable recovery cause all
switch(config)#
```

Configuring the Error-Disabled Recovery Interval

You can configure the error-disabled recovery timer value.

Procedure

	Command or Action	Purpose
Step 1	configure terminal Example: switch# configure terminal switch(config)#	Enters global configuration mode.

	Command or Action	Purpose
Step 2	errdisable recovery interval <i>interval</i> Example: <pre>switch(config)# errdisable recovery interval 32 switch(config-if)#</pre>	Specifies the interval for the interface to recover from the error-disabled state. The range is from 30 to 65535 seconds, and the default is 300 seconds.
Step 3	show interface status err-disabled Example: <pre>switch(config)# show interface status err-disabled switch(config-if)#</pre>	(Optional) Displays information about error-disabled interfaces.
Step 4	copy running-config startup-config Example: <pre>switch(config)# copy running-config startup-config</pre>	(Optional) Copies the running configuration to the startup configuration.

Example

This example shows how to configure the error-disabled recovery timer to set the interval for recovery to 32 seconds:

```
switch(config)# errdisable recovery interval 32
switch(config)#
```

Configuring the MTU Size

MTU is configured per interface, where the interface can be a Layer 2 or a Layer 3 interface. Every interface has default MTU of 1500 bytes. This value is called system default MTU. You can configure a Layer 2 interface, with a value of 9216 bytes, which is the default value of the system jumbo MTU. To allow an MTU value that is between 1500 and 9216, system jumbo MTU needs to be adjusted to appropriate value where interface can be configured with the same value.



Note You can change the system jumbo MTU size. When the value is changed, the Layer 2 interfaces that use the system jumbo MTU value, will automatically changes to the new system jumbo MTU value.

A Layer 3 interface, can be Layer 3 physical interface switch virtual interface (SVI), and subinterface, you can configure an MTU size between 576–9216 bytes.

Configuring the Interface MTU Size

For Layer 3 interfaces, you can configure an MTU with keyword MTU and value in bytes where value is between 576–9216 bytes.

For Layer 2 interfaces, you can configure an interface using the keyword MTU with value in bytes. The value can be a system default MTU size of 1500 bytes or the system jumbo MTU value that can be adjusted to the default size of 9216 bytes.

If you need to use a different system jumbo MTU size for Layer 2 interfaces, see the *Configuring the System Jumbo MTU Size* section.

Procedure

	Command or Action	Purpose
Step 1	configure terminal Example: <pre>switch# configure terminal switch(config)#</pre>	Enters global configuration mode.
Step 2	interface ethernet slot/port, vlan vlan-id mgmt 0 Example: <pre>switch(config)# interface ethernet 3/1 switch(config-if)# switch(config)# interface vlan 100 switch(config-if)# switch(config)# interface mgmt 0 switch(config-if)#</pre>	Specifies an Ethernet interface to configure, and enters interface configuration mode.
Step 3	mtu size Example: <pre>switch(config-if)# mtu 9216 switch(config-if)#</pre>	Configure the MTU value on an interface. For a Layer 3 interface, a physical Layer 3 interface, an SVI or sub-interface, then the value can be between 576-9216 bytes. If the interface is a physical Layer 2 interface, then the value can be 1500 or system jumbo MTU value.
Step 4	exit Example: <pre>switch(config-if)# exit switch(config)#</pre>	Exits the interface mode.

Example

This example shows how to configure the Layer 2 Ethernet port 3/1 with the default MTU size (1500):

```
switch# configure terminal
switch(config)# interface ethernet 3/1
switch(config-if)# switchport
switch(config-if)# mtu 1500
switch(config-if)#
```

This example displays the output of show running-config interface command:

```
switch# show run int mgmt0
!Command: show running-config interface mgmt0
!Running configuration last done at: Fri May 31 11:32:28 2019
!Time: Fri May 31 11:32:33 2019
```

```

version 9.3(1) Bios:version 07.65
interface mgmt0
mtu 9216
vrf member management
ip address 168.51.170.73/82

```

Configuring the System Jumbo MTU Size

You can configure the system jumbo MTU size, which can be used to specify the MTU size for Layer 2 interfaces. You can specify an even number between 1500 and 9216. If you do not configure the system jumbo MTU size, it defaults to 9216 bytes.



Note To configure jumbo frames for FEX modules, configure the FEX fabric port-channel interface with the required MTU size for the FEX module.

Procedure

	Command or Action	Purpose
Step 1	configure terminal Example: switch# configure terminal switch(config)#	Enters global configuration mode.
Step 2	system jumbomtu size Example: switch(config)# system jumbomtu 8000 switch(config)#	Specifies the system jumbo MTU size. Use an even number between 1500 and 9216. Note In general accepted practice, a jumbo frame is considered to have an MTU size greater than 9000 bytes.
Step 3	show running-config all Example: switch(config)# show running-config all include jumbomtu	(Optional) Displays the current operating configuration, which includes the system jumbo MTU size.
Step 4	interface type slot/port Example: switch(config)# interface ethernet 2/1 switch(config-if)#	Specifies an interface to configure and enters interface configuration mode.
Step 5	interface type Example: switch(config if)# interface mgmt0 switch(config-if)#	Specifies the management interface to configure.

	Command or Action	Purpose
Step 6	mtu size Example: <pre>switch(config-if)# mtu 1500 switch(config-if)#</pre>	For a Layer 2 interface, specifies either the default MTU size (1500) or the system jumbo MTU size that you specified earlier. For a Layer 3 interface, specifies any even size between 576 and 9216.
Step 7	exit Example: <pre>switch(config-if)# exit switch(config)#</pre>	Exits the interface mode.
Step 8	copy running-config startup-config Example: <pre>switch(config)# copy running-config startup-config</pre>	(Optional) Copies the running configuration to the startup configuration.

Example

This example shows how to configure the system jumbo MTU as 8000 bytes and how to change the MTU specification for an interface that was configured with the previous jumbo MTU size:

```
switch# configure terminal
switch(config)# system jumbomtu 8000
switch(config)# show running-config
switch(config)# interface ethernet 2/2
switch(config-if)# switchport
switch(config-if)# mtu 1500
switch(config-if)#
```

Configuring the Bandwidth

You can configure the bandwidth for Ethernet interfaces. The physical layer uses an unchangeable bandwidth of 1G, 10G, or 40G, but you can configure a value of 1 to 100,000,000 KB for Level 3 protocols.

Procedure

	Command or Action	Purpose
Step 1	configure terminal Example: <pre>switch# configure terminal switch(config)#</pre>	Enters global configuration mode.
Step 2	interface ethernet slot/port Example: <pre>switch(config)# interface ethernet 3/1 switch(config-if)#</pre>	Specifies an Ethernet interface to configure, and enters interface configuration mode.

	Command or Action	Purpose
Step 3	bandwidth <i>kbps</i> Example: switch(config-if) # bandwidth 1000000 switch(config-if) #	Specifies the bandwidth as an informational-only value between 1 and 100,000,000.
Step 4	show interface ethernet <i>slot/port</i> Example: switch(config) # show interface ethernet 2/1	(Optional) Displays the interface status, which includes the bandwidth value.
Step 5	exit Example: switch(config-if) # exit switch(config) #	Exits the interface mode.
Step 6	copy running-config startup-config Example: switch(config) # copy running-config startup-config	(Optional) Copies the running configuration to the startup configuration.

Example

This example shows how to configure an informational value of 1,000,000 Kb for the Ethernet slot 3, port 1 interface bandwidth parameter:

```
switch# configure terminal
switch(config)# interface ethernet 3/1
switch(config-if)# bandwidth 1000000
switch(config-if)#
```

Configuring the Throughput Delay

You can configure the interface throughput delay for Ethernet interfaces. The actual delay time does not change, but you can set an informational value between 1 and 16777215, where the value represents the number of tens of microseconds.

Procedure

	Command or Action	Purpose
Step 1	configure terminal Example: switch# configure terminal switch(config) #	Enters global configuration mode.
Step 2	interface ethernet <i>slot/port</i> Example:	Specifies an Ethernet interface to configure, and enters interface configuration mode.

	Command or Action	Purpose
	<code>switch(config)# interface ethernet 3/1</code> <code>switch(config-if)#</code>	
Step 3	delay value Example: <code>switch(config-if)# delay 10000</code> <code>switch(config-if)#</code>	Specifies the delay time in tens of microseconds. You can set an informational value range between 1 and 16777215 tens of microseconds.
Step 4	show interface ethernet slot/port Example: <code>switch(config)# show interface ethernet 3/1</code> <code>switch(config-if)#</code>	(Optional) Displays the interface status, which includes the throughput-delay time.
Step 5	exit Example: <code>switch(config-if)# exit</code> <code>switch(config)#</code>	Exits the interface mode.
Step 6	copy running-config startup-config Example: <code>switch(config)# copy running-config startup-config</code>	(Optional) Copies the running configuration to the startup configuration.

Example

This example shows how to configure the throughput-delay time so that one interface is preferred over another. A lower delay value is preferred over a higher value. In this example, Ethernet 7/48 is preferred over 7/47. The default delay for 7/48 is less than the configured value on 7/47, which is set for the highest value (16777215):

```
switch# configure terminal
switch(config)# interface ethernet 7/47
switch(config-if)# delay 16777215
switch(config-if)# ip address 192.168.10.1/24
switch(config-if)# ip router eigrp 10
switch(config-if)# no shutdown
switch(config-if)# exit
switch(config)# interface ethernet 7/48
switch(config-if)# ip address 192.168.11.1/24
switch(config-if)# ip router eigrp 10
switch(config-if)# no shutdown
switch(config-if)#
```



Note You must first ensure the EIGRP feature is enabled by running the **feature eigrp** command.

Shutting Down and Activating the Interface

You can shut down and restart Ethernet or management interfaces. When you shut down interfaces, they become disabled and all monitoring displays show them as being down. This information is communicated to other network servers through all dynamic routing protocols. When the interfaces are shut down, the interface is not included in any routing updates. To activate the interface, you must restart the device.

Procedure

	Command or Action	Purpose
Step 1	configure terminal Example: <pre>switch# configure terminal switch(config)#</pre>	Enters global configuration mode.
Step 2	interface <i>interface</i> Example: <pre>switch(config)# interface ethernet 2/1 switch(config-if)# switch(config)# interface mgmt0 switch(config-if)#</pre>	Specifies the interface that you are configuring. You can specify the interface type and identity. For an Ethernet port, use <i>ethernet slot/port</i> . For the management interface, use <i>mgmt0</i> . Examples: <ul style="list-style-type: none"> • The 1st example shows how to specify the slot 2, port 1 Ethernet interface. • The 2nd example shows how to specify the management interface.
Step 3	shutdown Example: <pre>switch(config-if)# shutdown switch(config-if)#</pre>	Disables the interface.
Step 4	show interface <i>interface</i> Example: <pre>switch(config-if)# show interface ethernet 2/1 switch(config-if)#</pre>	(Optional) Displays the interface status, which includes the administrative status.
Step 5	no shutdown Example: <pre>switch(config-if)# no shutdown switch(config-if)#</pre>	Reenables the interface.
Step 6	show interface <i>interface</i> Example: <pre>switch(config-if)# show interface ethernet 2/1 switch(config-if)#</pre>	(Optional) Displays the interface status, which includes the administrative status.

	Command or Action	Purpose
Step 7	exit Example: <pre>switch(config-if) # exit switch(config) #</pre>	Exits the interface mode.
Step 8	copy running-config startup-config Example: <pre>switch(config) # copy running-config startup-config</pre>	(Optional) Copies the running configuration to the startup configuration.

Example

This example shows how to change the administrative status for Ethernet port 3/1 from disabled to enabled:

```
switch# configure terminal
switch(config)# interface ethernet 3/1
switch(config-if)# shutdown
switch(config-if)# no shutdown
switch(config-if)#
```

Configuring the UDLD Mode

You can configure normal unidirectional link detection (UDLD) modes for Ethernet interfaces on devices configured to run UDLD.

Before you can enable the aggressive UDLD mode for an interface, you must make sure that UDLD is already enabled globally on the device and on the specified interfaces.



Note If the interface is a copper port, you must use the command `enable UDLD` to enable the UDLD. If the interface is a fiber port you need not explicitly enable UDLD on the interface. However if you attempt to enable UDLD on a fiber port using the `enable UDLD` command, you may get an error message indicating that is not a valid command.

The following table lists CLI details to enable and disable UDLD on different interfaces

Table 5: CLI Details to Enable or Disable UDLD on Different Interfaces

Description	Fiber port	Copper or Nonfiber port
Default setting	Enabled	Disabled
Enable UDLD command	no udld disable	udld enable
Disable UDLD command	udld disable	no udld enable

Before you begin

You must enable UDLD for the other linked port and its device.

Procedure

	Command or Action	Purpose
Step 1	configure terminal Example: <pre>switch# configure terminal switch(config)#</pre>	Enters global configuration mode.
Step 2	[no] feature udld Example: <pre>switch(config)# feature udld switch(config)# switch(config)# no feature udld switch(config)#</pre>	Enables/Disables UDLD for the device.
Step 3	udld message-time seconds Example: <pre>switch(config)# udld message-time 30 switch(config)#</pre>	(Optional) Specifies the interval between sending UDLD messages. The range is from 7 to 90 seconds, and the default is 15 seconds.
Step 4	udld aggressive Example: <pre>switch(config)# udld aggressive switch(config)#</pre>	<p>Enables UDLD in aggressive mode by default on all fiber interfaces. Use the no form to disable aggressive mode UDLD on all fibers ports by default.</p> <p>Note Use the <code>udld aggressive</code> command to configure the ports to use a UDLD mode:</p> <ul style="list-style-type: none"> • To enable fiber interfaces for the aggressive mode, enter the udld aggressive command in the global command mode and all the fiber interfaces will be in aggressive UDLD mode. • To enable the copper interfaces for the aggressive mode, you must enter the udld aggressive command in the interface mode, specifying each interface you want in aggressive UDLD mode.

	Command or Action	Purpose
		To use the aggressive UDLD mode, you must configure the interfaces on both ends of the link for the aggressive UDLD mode.
Step 5	interface ethernet <i>slot/port</i> Example: switch(config)# interface ethernet 3/1 switch(config-if)#	(Optional) Specifies an interface to configure, and enters interface configuration mode.
Step 6	udld [enable disable] Example: switch(config-if)# udld enable switch(config-if)#	Enables UDLD in normal mode by default on all fiber interfaces. Use the no form to disable normal mode UDLD on all fibers ports by default.
Step 7	show udld [ethernet <i>slot/port</i> global neighbors] Example: switch(config)# show udld switch(config)#	(Optional) Displays the UDLD status.
Step 8	exit Example: switch(config-if-range)# exit switch(config)#	Exits the interface mode.
Step 9	copy running-config startup-config Example: switch(config)# copy running-config startup-config	(Optional) Copies the running configuration to the startup configuration.

Example

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

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

This example shows how to set the UDLD message interval to 30 seconds:

```
switch# configure terminal
switch(config)# feature udld
switch(config)# udld message-time 30
switch(config)#
```

This example shows how to disable UDLD for Ethernet port 3/1:

```
switch# configure terminal
switch(config)# interface ethernet 3/1
switch(config-if-range)# no udld enable
switch(config-if-range)# exit
```

This example shows how to disable UDLD for the device:

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

This example shows how to enable fiber interfaces for the aggressive UDLD mode:

```
switch# configure terminal
switch(config)# udld aggressive
```

This example shows how to enable the aggressive UDLD mode for the copper Ethernet interface3/1:

```
switch# configure terminal
switch(config)# interface ethernet 3
switch(config-if)# udld aggressive
```

This example shows how to check if aggressive mode is enabled.

```
switch# sh udld global

UDLD global configuration mode: enabled-aggressive
UDLD global message interval: 15
switch#
```

This example shows how to check if udld aggressive mode is operational for a given interface.

```
switch# sh udld ethernet 8/2

Interface Ethernet8/2
-----
Port enable administrative configuration setting: device-default
Port enable operational state: enabled-aggressive
Current bidirectional state: bidirectional
Current operational state: advertisement - Single neighbor detected
Message interval: 15
Timeout interval: 5
<>
```

Configuring Debounce Timers

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.



Note The link state of 10G and 100G ports may change repeatedly when connected to service provider network. As a part of *link reset* or *break-link* functionality, it is expected that the Tx power light on the SFP to change to N/A state, at an event of link state change.

However, to prevent this behavior during the link state change, you may increase the link debounce timer to start from 500ms and increase it in 500ms intervals until the link stabilizes. On the DWDM, UVN, and WAN network, it is recommended to disable automatic link suspension (ALS) whenever possible. ALS suspends the link on the WAN when the Nexus turn off the link.



Note The **link debounce time** and **link debounce link-up time** commands can only be applied to a physical Ethernet interface.

Use the **show interface debounce** command to display the debounce times for all Ethernet ports.

The **link debounce time** command is not supported on 10G and 40G ports on the Cisco Nexus 93300YC-FX and Cisco Nexus 9336C-FX switches.

The **link debounce time** command is supported on 1G, 10G, 40G, 25G and 100G SFP/QSFP ports on the Cisco Nexus 9000 series switches.

The **link debounce time** is supported on 1G, 10G, 25G, 40G and 100G ports on Cisco Nexus N9K-C9732C-FX, N9K-C9364C, N9K-X97160YC-EX, N9K-C9336C-FX2, and N9K-C93240YC-FX2 platform switches.

Procedure

	Command or Action	Purpose
Step 1	configure terminal Example: <pre>switch# configure terminal switch(config)#</pre>	Enters global configuration mode.
Step 2	interface ethernet <i>slot/port</i> Example: <pre>switch(config)# interface ethernet 3/1 switch(config-if)#</pre>	Specifies an Ethernet interface to configure, and enters interface configuration mode.
Step 3	link debounce time <i>time</i> Example: <pre>switch(config-if)# link debounce time 1000 switch(config-if)#</pre>	<p>Enables the debounce timer for the specified time (1 to 5000 milliseconds).</p> <p>If you specify 0 milliseconds, the debounce timer is disabled.</p>
Step 4	link debounce link-up <i>time</i> Example: <pre>switch(config-if)# link debounce link-up 1000 switch(config-if)#</pre>	<p>Enables the link-up timer for the specified time (1000 to 10000 milliseconds). This command applies only if the port speeds are 10G, 25G, 40G and 100G.</p> <p>The default value of the timer is 0. If the value is set to 0 the interface will be up without any delay.</p> <p>Note The no link debounce link-up command also resets the value to 0.</p>

	Command or Action	Purpose
		<p>Note This command is supported only on Cisco Nexus N9K-X9732C-FX, N9K-C93300YC-FX, N9K-C9336C-FX2, N9K-C9364C and N9K-X97160YC-EX switches.</p>

Example

- The following example enables the debounce timer and sets 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
```

- The following example disables the debounce timer for an Ethernet interface:

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

- The following example sets the debounce link-up timer to 1000 milliseconds for an Ethernet interface:

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

Configuring Port Profiles

Beginning with 7.0(3)I4(1), you can apply several configuration parameters to a range of interfaces simultaneously. All the interfaces in the range must be the same type. You can also inherit the configurations from one port profile into another port profile. The system supports four levels of inheritance.

Creating a Port Profile

You can create a port profile on the device. Each port profile must have a unique name across types and the network.



Note Port profile names can include only the following characters:

- a-z
- A-Z
- 0-9
- No special characters are allowed, except for the following:
 - .
 - -
 - _

Procedure

	Command or Action	Purpose
Step 1	<code>configure terminal</code>	Enters the global configuration mode.
Step 2	<code>port-profile [type {ethernet interface-vlan port-channel}] name</code>	Creates and names a port profile for the specified type of interface and enters the port-profile configuration mode.
Step 3	<code>exit</code>	Exits the port-profile configuration mode.
Step 4	(Optional) <code>show port-profile</code>	Displays the port-profile configuration.
Step 5	(Optional) <code>copy running-config startup-config</code>	Copies the running configuration to the startup configuration.

Example

This example shows how to create a port profile named test for ethernet interfaces:

```
switch# configure terminal
switch(config)# port-profile type ethernet test
switch(config-ppm)#
```

Entering Port-Profile Configuration Mode and Modifying a Port Profile

You can enter the port-profile configuration mode and modify a port profile. To modify the port profile, you must be in the port-profile configuration mode.

Procedure

	Command or Action	Purpose
Step 1	<code>configure terminal</code>	Enters the global configuration mode.

	Command or Action	Purpose
Step 2	port-profile [type { ethernet interface-vlan port-channel }] <i>name</i>	Enters the port-profile configuration mode for the specified port profile and allows you to add or remove configurations to the profile.
Step 3	exit	Exits the port-profile configuration mode.
Step 4	(Optional) show port-profile	Displays the port-profile configuration.
Step 5	(Optional) copy running-config startup-config	Copies the running configuration to the startup configuration.

Example

This example shows how to enter the port-profile configuration mode for the specified port profile and bring all the interfaces administratively up:

```
switch# configure terminal
switch(config)# port-profile type ethernet test
switch(config-ppm)# no shutdown
switch(config-ppm)#
```

Assigning a Port Profile to a Range of Interfaces

You can assign a port profile to an interface or to a range of interfaces. All the interfaces must be the same type.

Procedure

	Command or Action	Purpose
Step 1	configure terminal	Enters the global configuration mode.
Step 2	interface [ethernet <i>slot/port</i> interface-vlan <i>vlan-id</i> port-channel <i>number</i>]	Selects the range of interfaces.
Step 3	inherit port-profile <i>name</i>	Assigns the specified port profile to the selected interfaces.
Step 4	exit	Exits the port-profile configuration mode.
Step 5	(Optional) show port-profile	Displays the port-profile configuration.
Step 6	(Optional) copy running-config startup-config	Copies the running configuration to the startup configuration.

Example

This example shows how to assign the port profile named adam to Ethernet interfaces 7/3 to 7/5, 10/2, and 11/20 to 11/25:


```
switch# configure terminal
switch(config)# interface ethernet7/3-5, ethernet10/2, ethernet11/20-25
switch(config-if)# inherit port-profile adam
switch(config-if)#
```

Enabling a Specific Port Profile

To apply the port-profile configurations to the interfaces, you must enable the specific port profile. You can configure and inherit a port profile onto a range of interfaces before you enable that port profile. You would then enable that port profile for the configurations to take effect on the specified interfaces.

If you inherit one or more port profiles onto an original port profile, only the last inherited port profile must be enabled; the system assumes that the underlying port profiles are enabled.

You must be in the port-profile configuration mode to enable or disable port profiles.

Procedure

	Command or Action	Purpose
Step 1	configure terminal	Enters the global configuration mode.
Step 2	port-profile [type { ethernet interface-vlan port-channel }] <i>name</i>	Creates and names a port profile for the specified type of interface and enters the port-profile configuration mode.
Step 3	state enabled	Enables that port profile.
Step 4	exit	Exits the port-profile configuration mode.
Step 5	(Optional) show port-profile	Displays the port-profile configuration.
Step 6	(Optional) copy running-config startup-config	Copies the running configuration to the startup configuration.

Example

This example shows how to enter the port-profile configuration mode and enable the port profile:

```
switch# configure terminal
switch(config)# port-profile type ethernet test
switch(config-ppm)# state enabled
switch(config-ppm)#
```

Inheriting a Port Profile

You can inherit a port profile onto an existing port profile. The system supports four levels of inheritance.

Procedure

	Command or Action	Purpose
Step 1	configure terminal	Enters the global configuration mode.

	Command or Action	Purpose
Step 2	port-profile <i>name</i>	Enters the port-profile configuration mode for the specified port profile.
Step 3	inherit port-profile <i>name</i>	Inherits another port profile onto the existing one. The original port profile assumes all the configurations of the inherited port profile.
Step 4	exit	Exits the port-profile configuration mode.
Step 5	(Optional) show port-profile	Displays the port-profile configuration.
Step 6	(Optional) copy running-config startup-config	Copies the running configuration to the startup configuration.

Example

This example shows how to inherit the port profile named adam onto the port profile named test:

```
switch# configure terminal
switch(config)# port-profile test
switch(config-ppm)# inherit port-profile adam
switch(config-ppm)#
```

Removing a Port Profile from a Range of Interfaces

You can remove a port profile from some or all of the interfaces to which you have applied the profile. You do this configuration in the interfaces configuration mode.

Procedure

	Command or Action	Purpose
Step 1	configure terminal	Enters the global configuration mode.
Step 2	interface [<i>ethernet slot/port</i> interface-vlan <i>vlan-id</i> port-channel <i>number</i>]	Selects the range of interfaces.
Step 3	no inherit port-profile <i>name</i>	Un-assigns the specified port profile to the selected interfaces.
Step 4	exit	Exits the port-profile configuration mode.
Step 5	(Optional) show port-profile	Displays the port-profile configuration.
Step 6	(Optional) copy running-config startup-config	Copies the running configuration to the startup configuration.

Example

This example shows how to unassign the port profile named adam to Ethernet interfaces 7/3 to 7/5, 10/2, and 11/20 to 11/25:

```
switch# configure terminal
switch(config)# interface ethernet 7/3-5, 10/2, 11/20-25
switch(config-if)# no inherit port-profile adam
switch(config-if)#
```

Removing an Inherited Port Profile

You can remove an inherited port profile. You do this configuration in the port-profile mode.

Procedure

	Command or Action	Purpose
Step 1	<code>configure terminal</code>	Enters the global configuration mode.
Step 2	<code>port-profile name</code>	Enters the port-profile configuration mode for the specified port profile.
Step 3	<code>no inherit port-profile name</code>	Removes an inherited port profile from this port profile.
Step 4	<code>exit</code>	Exits the port-profile configuration mode.
Step 5	(Optional) <code>show port-profile</code>	Displays the port-profile configuration.
Step 6	(Optional) <code>copy running-config startup-config</code>	Copies the running configuration to the startup configuration.

Example

This example shows how to remove the inherited port profile named adam from the port profile named test:

```
switch# configure terminal
switch(config)# port-profile test
switch(config-ppm)# no inherit port-profile adam
switch(config-ppm)#
```

Configuring link mac-up timer

This procedure describes how to configure mac up timers on DWDM/Dark fiber circuits.

Procedure

	Command or Action	Purpose
Step 1	configure terminal Example: <pre>switch# configure terminal switch(config)#</pre>	Enters global configuration mode.
Step 2	interface <i>type slot/port</i> Example: <pre>switch(config)# interface ethernet1/2 switch(config-if)#</pre>	Configures an interface and enters interface configuration mode.
Step 3	link mac-up timer <i>seconds</i> Example: <pre>switch(config-if)# link mac-up timer 10</pre>	Enables modification of the link mac-up timer. The link mac-up timer range is 0-120. Note This should only be done on DWDM links.

Verifying the Basic Interface Parameters

You can verify the basic interface parameters by displaying their values. You can also clear the counters listed when you display the parameter values.

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

Command	Purpose
show cdp all	Displays the CDP status.
show interface <i>interface</i>	Displays the configured states of one or all interfaces.
show interface <i>brief</i>	Displays a table of interface states.
show interface status err-disabled	Displays information about error-disabled interfaces.
show udld <i>interface</i>	Displays the UDLD status for the current interface or all interfaces.
show udld global	Displays the UDLD status for the current device.

Monitoring the Interface Counters

You can display and clear interface counters using Cisco NX-OS.

Displaying Interface Statistics

You can set up to three sampling intervals for statistics collections on interfaces.

Procedure

	Command or Action	Purpose
Step 1	configure terminal Example: <pre>switch# configure terminal switch(config)#</pre>	Enters global configuration mode.
Step 2	interface ether <i>slot/port</i> Example: <pre>switch(config)# interface ether 4/1 switch(config)#</pre>	Specifies interface.
Step 3	load-interval counters [1 2 3] seconds Example: <pre>switch(config)# load-interval counters 1 100 switch(config)#</pre>	Sets up to three sampling intervals to collect bit-rate and packet-rate statistics. The default values for each counter is as follows: 1—30 seconds (60 seconds for VLAN) 2—300 seconds 3—not configured
Step 4	show interface <i>interface</i> Example: <pre>switch(config)# show interface ethernet 2/2 switch#</pre>	(Optional) Displays the interface status, which includes the counters.
Step 5	exit Example: <pre>switch(config-if-range)# exit switch(config)#</pre>	Exits the interface mode.
Step 6	copy running-config startup-config Example: <pre>switch(config)# copy running-config startup-config</pre>	(Optional) Copies the running configuration to the startup configuration.

Example

This example shows how to set the three sample intervals for the Ethernet port 3/1:

```
switch# configure terminal
switch(config)# interface ethernet 3/1
switch(config-if)# load-interval counter 1 60
switch(config-if)# load-interval counter 2 135
switch(config-if)# load-interval counter 3 225
switch(config-if)#
```

Clearing Interface Counters

You can clear the Ethernet and management interface counters by using the **clear counters interface** command. You can perform this task from the configuration mode or interface configuration mode.

Procedure

	Command or Action	Purpose
Step 1	clear counters interface [all ethernet <i>slot/port</i> loopback <i>number</i> mgmt <i>number</i> port channel <i>channel-number</i>] Example: <pre>switch# clear counters ethernet 2/1 switch#</pre>	Clears the interface counters.
Step 2	show interface <i>interface</i> Example: <pre>switch# show interface ethernet 2/1 switch#</pre>	(Optional) Displays the interface status.
Step 3	show interface [ethernet <i>slot/port</i> port channel <i>channel-number</i>] counters Example: <pre>switch# show interface ethernet 2/1 counters switch#</pre>	(Optional) Displays the interface counters.

Example

This example shows how to clear the counters on Ethernet port 5/5:

```
switch# clear counters interface ethernet 5/5
switch#
```

Configuration Example for QSA

For a Cisco Nexus 9396PX:

- Using the default configuration on port 2/1, all the QSFPs in port group 2/1-6 are brought up with a speed of 40G. If there are any QSA modules in port group 2/1-6, they are error disabled.
- Using the **speed-group** [**10000** | **40000**] command to configure port 2/7, all the QSAs in port group 2/7-12 are brought up with a speed of 10G or 40G. If there are any QSFP modules in port group 2/7-12, they are error disabled.

This example shows how to configure QSA for the first port in the speed group for a Cisco Nexus 9396PX:

```
switch# conf t
```

```
switch(config)# interface ethernet 2/7  
switch(config-if)# speed-group 10000
```




CHAPTER 4

Configuring Layer 2 Interfaces

This chapter describes how to configure Layer 2 switching ports as access or trunk ports on Cisco NX-OS devices.



Note A Layer 2 port can function as either one of the following:

- A trunk port
- An access port



Note See the [Cisco Nexus 9000 Series NX-OS System Management Configuration Guide](#) for information about configuring a SPAN destination interface.

You can configure Layer 2 switching ports as access or trunk ports. Trunks carry the traffic of multiple VLANs over a single link and allow you to extend VLANs across an entire network. All Layer 2 switching ports maintain media access control (MAC) address tables.



Note See the [Cisco Nexus 9000 Series NX-OS Layer 2 Switching Configuration Guide](#) for information about VLANs, MAC address tables, private VLANs, and the Spanning Tree Protocol.

- [Information About Access and Trunk Interfaces](#), on page 62
- [Prerequisites for Layer 2 Interfaces](#), on page 68
- [Guidelines and Limitations for Layer 2 Interfaces](#), on page 68
- [Default Settings for Layer 2 Interfaces](#), on page 71
- [Configuring Access and Trunk Interfaces](#), on page 72
- [Verifying the Interface Configuration](#), on page 90
- [Monitoring the Layer 2 Interfaces](#), on page 91
- [Configuration Examples for Access and Trunk Ports](#), on page 91
- [Related Documents](#), on page 92

Information About Access and Trunk Interfaces



Note See the [Cisco Nexus 9000 Series NX-OS High Availability and Redundancy Guide](#) for complete information on high-availability features.



Note The device supports only IEEE 802.1Q-type VLAN trunk encapsulation.

About Access and Trunk Interfaces

A Layer 2 port can be configured as an access or a trunk port as follows:

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


By default, all the ports on Cisco Nexus 9300-EX switches are Layer 3 ports and all the ports on Cisco Nexus 9300 switches are Layer 2 ports.

You can make all ports Layer 2 ports using the setup script or by entering the **system default switchport** command. See the [Cisco Nexus 9000 Series NX-OS Fundamentals Configuration Guide](#) for information about using the setup script. To configure the port as a Layer 2 port using the CLI, use the **switchport** command.

All ports in the same trunk must be in the same VDC, and trunk ports cannot carry VLANs from different VDCs.

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

Figure 2: Trunk and Access Ports and VLAN Traffic



Note See the [Cisco Nexus 9000 Series NX-OS Layer 2 Switching Configuration Guide](#) for information about VLANs.

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 (see the “IEEE 802.1Q Encapsulation” section for more information).



Note See the [Cisco Nexus 9000 Series NX-OS Unicast Routing Configuration Guide](#) for information about subinterfaces on Layer 3 interfaces.

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 that it takes the designated port to begin to forward packets.

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.

A Layer 2 interface can function as either an access port or a trunk port; it cannot function as both port types simultaneously.

When you change a Layer 2 interface back to a Layer 3 interface, that interface loses all the Layer 2 configuration and resumes the default VLAN configurations.

IEEE 802.1Q Encapsulation




Note For information about VLANs, see the [Cisco Nexus 9000 Series NX-OS Layer 2 Switching Configuration Guide](#).

A trunk is a point-to-point link between the switch 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, or 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. Also, the encapsulated VLAN tag allows the trunk to move traffic end-to-end through the network on the same VLAN.

Figure 3: Header Without and With 802.1Q Tag



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 shuts 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.

Native VLAN IDs for Trunk Ports

A trunk port can carry nontagged 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. That is, the native VLAN ID is the VLAN that carries untagged traffic on trunk ports.



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

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 You cannot use a Fibre Channel over Ethernet (FCoE) VLAN as a native VLAN for an Ethernet trunk switchport.

Tagging Native VLAN Traffic

The Cisco software supports the IEEE 802.1Q standard on trunk ports. In order to pass untagged traffic through the trunk ports, you must create a VLAN that does not tag any packets (or you can use the default VLAN). Untagged packets can pass through trunk ports and access ports.

However, all packets that enter the device with an 802.1Q tag that matches the value of the native VLAN on the trunk are stripped of any tagging and egress the trunk port as untagged packets. This situation can cause problems because you may want to retain the tagging on packets on the native VLAN for the trunk port.

You can configure the device to drop all untagged packets on the trunk ports and to retain the tagging of packets entering the device with 802.1Q values that are equal to that of the native VLAN ID. All control traffic still passes on the native VLAN. This configuration is global; trunk ports on the device either do or do not retain the tagging for the native VLAN.

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. Later, you can add any specific VLANs that you may want the trunk to carry traffic for back to the list.

To partition the 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.



Note See the [Cisco Nexus 9000 Series NX-OS Layer 2 Switching Configuration Guide](#) for more information about STP.



Note You can change the block of VLANs reserved for internal use. See the [Cisco Nexus 9000 Series NX-OS Layer 2 Switching Configuration Guide](#) for more information about changing the reserved VLANs.

Switchport Isolated for up to 3967 VLANs on Trunk Interfaces

Trunk interface can carry multiple VLANs, where a trunk interface that is configured in switchport isolated mode allows you to configure multiple VLANs per interface. Sometimes, you may require higher number VLANs per port. The scale of logical ports of Per-VLAN Spanning Tree (PVST) and virtual ports of Multiple Spanning Tree (MST) can be limiting. By configuring switchport isolated on a trunk interface, you can configure up to 48 interfaces with up to 3967 VLANs per port, on switches in Cisco Nexus 9000 portfolio.

When you change the member VLANs on the isolated interfaces, all VLANs on these interfaces get moved to forwarding state. Switchport isolated feature is supported only on the host interfaces, because no spanning tree is running on these ports (the switch doesn't send STP BPDUs), connecting other network devices may create a loop in the network. Switchport isolated feature is supported on physical interfaces, port channels, and vPC. Switchport isolated feature has the following restrictions:

- Supports Per-VLAN Rapid Spanning Tree (PVRST) and isolated VLANs. It allows some ports to be in isolated mode, and some ports run Rapid Per-VLAN Spanning Tree (RPVST) with the same VLANs.
- Rapid Spanning Tree (RSTP) running on other ports with same VLAN is supported.
- Is not supported on FEX HIF, FEX fabric interface, on an interface when another network device is connected.
- Supports up to 48 ports configured with up to 3967 VLANs.
- When used in vPC environment if configured inconsistently will trigger vPC type 1 constancy checks.
- Port channel members are required to have the same switchport isolated configuration.

Default Interfaces

You can use the default interface feature to clear the configured parameters for both physical and logical interfaces such as the Ethernet, loopback, VLAN network, tunnel, and the port-channel interface.



Note A maximum of eight ports can be selected for the default interface. The default interfaces feature is not supported for management interfaces because the device could go to an unreachable state.

Switch Virtual Interface and Autostate Behavior

In Cisco NX-OS, a switch virtual interface (SVI) represents a logical interface between the bridging function and the routing function of a VLAN in the device.

The operational state of this interface is governed by the state of the various ports in its corresponding VLAN. An SVI interface on a VLAN comes up when at least one port in that VLAN is in the Spanning Tree Protocol (STP) forwarding state. Similarly, this interface goes down when the last STP forwarding port goes down or goes to another STP state.

SVI Autostate Exclude

Typically, when a VLAN interface has multiple ports in the VLAN, the SVI goes to the down state when all the ports in the VLAN go down. You can use the SVI autostate exclude feature to exclude specific ports and port channels while defining the status of the SVI (up or down) even if it belongs to the same VLAN. For example, even if the excluded port or port channel is in the up state and other ports are in the down state in the VLAN, the SVI state is changed to down.



Note You can use the SVI autostate exclude feature only for switched physical Ethernet ports and port channels.

SVI Autostate Disable

You can configure the autostate disable feature to keep an SVI up even if no interface is up in the corresponding VLAN. You can configure this feature for the system (for all SVIs) or for an individual SVI.

High Availability

See the [Cisco Nexus 9000 Series NX-OS High Availability and Redundancy Guide](#) for complete information about high availability features.

Virtualization Support

All ports in the same trunk must be in the same VDC, and trunk ports cannot carry VLANs from different VDCs.

Counter Values

See the following information on the configuration, packet size, incremented counter values, and traffic.

Configuration	Packet Size	Incremented Counters	Traffic
L2 port – without any MTU configuration	6400 and 10000	Jumbo, giant, and input error	Dropped
L2 port – with jumbo MTU 9216 in network-qos configuration	6400	Jumbo	Forwarded
L2 port – with jumbo MTU 9216 in network-qos configuration	10000	Jumbo, giant, and input error	Dropped
Layer 3 port with default Layer 3 MTU and jumbo MTU 9216 in network-qos configuration	6400	Jumbo	Packets are punted to the CPU (subjected to CoPP configs), get fragmented, and then they are forwarded by the software.
Layer 3 port with default Layer 3 MTU and jumbo MTU 9216 in network-qos configuration	6400	Jumbo	Packets are punted to the CPU (subjected to CoPP configs), get fragmented, and then they are forwarded by the software.
Layer 3 port with default Layer 3 MTU and jumbo MTU 9216 in network-qos configuration	10000	Jumbo, giant, and input error	Dropped

Configuration	Packet Size	Incremented Counters	Traffic
Layer 3 port with jumbo Layer 3 MTU and jumbo MTU 9216 in network-qos configuration	6400	Jumbo	Forwarded without any fragmentation.
Layer 3 port with jumbo Layer 3 MTU and jumbo MTU 9216 in network-qos configuration	10000	Jumbo, giant, and input error	Dropped
Layer 3 port with jumbo Layer 3 MTU and default L2 MTU configuration	6400 and 10000	Jumbo, giant, and input error	Dropped

**Note**

- Under 64 bytes packet with good CRC—The short frame counter increments.
- Under 64 bytes packet with bad CRC—The runts counter increments.
- Greater than 64 bytes packet with bad CRC—The CRC counter increments.

Prerequisites for Layer 2 Interfaces

Layer 2 interfaces have the following prerequisites:

- By default, Cisco NX-OS configures Layer 3 parameters. If you want to configure Layer 2 parameters, you need to switch the port mode to Layer 2. You can change the port mode by using the **switchport** command.
- You must configure the port as a Layer 2 port before you can use the **switchport mode** command. By default, all ports on the device are Layer 3 ports. By default, all ports on the Cisco Nexus 9504 and Cisco Nexus 9508 devices are Layer 2 ports.

Guidelines and Limitations for Layer 2 Interfaces

VLAN trunking has the following configuration guidelines and limitations:

- Cisco Nexus 9000 Series switches have the **vlan dot1q tag native** command that can be configured globally. This tags the native VLAN on the configured trunk ports. However, connected switches such as Catalyst 6500 or third-party switches, probably would not have a similar configuration enabled. This could result in unexpected behaviors. Therefore, it is recommended to have the **vlan dot1q tag native** command disabled in case the connected switch does not have it configured.
- Beginning with Release 7.0(3)F3(2), auto-negotiation is not supported on Cisco Nexus 9508 switches with N9K-X9636C-R, N9K-X9636C-RX, and N9K-X9636Q-R lines cards.

- When a store-and-forwarding device receives a packet larger than the MTU size, the packet is dropped. In a cut-through switching mode, packets larger than MTU are transferred.
- Auto-negotiation is not supported on 25-G Ethernet transceiver modules on Cisco Nexus 9200 and 9300-FX platform switches, and Cisco Nexus 9500 platform switches that use N9K-X9700-EX line cards.
- On the Cisco Nexus 9364C switches, auto-negotiation might not work on ports 49-64 when bringing up 100G links using the QSFP-100G-CR4 cable. The workaround for this issue is that you must hard code the speed on ports 49-64 and disable auto-negotiation.
- Autonegotiation (40 G/100 G) and 1 GB with QSA is not supported on the following ports:
 - Cisco Nexus 9336C-FX2 switch: ports 1-6 and 33-36
 - Cisco Nexus 9364C switch: ports 49-66
 - Cisco Nexus 93240YC-FX2 switch: ports 51-54
 - Cisco Nexus 9788TC line card: ports 49-52



Note Peer speed must be set when using copper cables on these ports.

- **show** commands with the **internal** keyword are not supported.
- Starting with Release 7.0(3)I2(1) or later releases on Cisco Nexus 9300 platform switches, a unicast ARP request to SVI is flooded to the other ports within the VLAN.
- ASE2 and ASE3 based Cisco Nexus 9000 Series switches acting as transit switches do not preserve the inner tag for double-tagged packets.

The following CLI is mandatory only on LSE based Cisco Nexus 9000 Series switches. For seamless packet forwarding and preservation of all VLAN tags on pure transit boxes in the SP cloud that have no Q-in-Q encapsulation or decapsulation requirement, configure the CLI command, **system dot1q-tunnel transit** . To remove the CLI, use **no system dot1q-tunnel transit** CLI command.

The caveats with the CLI that is executed on the switches are:

- L2 frames that egress out of the trunk ports are tagged even on the native VLAN on the port.
- Any other tunneling mechanism, for example, VXLAN and MPLS does not work with the CLI configured.
- A port can be either a Layer 2 or a Layer 3 interface; it cannot be both simultaneously.
- When you change a Layer 3 port to a Layer 2 port or a Layer 2 port to a Layer 3 port, all layer-dependent configuration is lost. When you change an access or trunk port to a Layer 3 port, all information about the access VLAN, native VLAN, allowed VLANs, and so forth, is lost.
- Do not connect devices with access links because access links may partition a VLAN.
- When connecting Cisco devices through an 802.1Q trunk, make sure that the native VLAN for an 802.1Q trunk is the same on both ends of the trunk link. If the native VLAN on one end of the trunk is different from the native VLAN on the other end, spanning tree loops might result.

- Disabling spanning tree on the native VLAN of an 802.1Q trunk without disabling spanning tree on every VLAN in the network can cause spanning tree loops. You must leave spanning tree enabled on the native VLAN of an 802.1Q trunk. If you cannot leave spanning tree enabled, you must disable spanning tree on every VLAN in the network. Make sure that your network has no physical loops before you disable spanning tree.
- When you connect two Cisco devices through 802.1Q trunks, the devices exchange spanning tree bridge protocol data units (BPDUs) on each VLAN allowed on the trunks. The BPDUs on the native VLAN of the trunk are sent untagged to the reserved IEEE 802.1D spanning tree multicast MAC address (01-80-C2-00-00-00). The BPDUs on all other VLANs on the trunk are sent tagged to the reserved Cisco Shared Spanning Tree (SSTP) multicast MAC address (01-00-0c-cc-cc-cd).
- Non-Cisco 802.1Q devices maintain only a single instance of spanning tree (the Mono Spanning Tree) that defines the spanning tree topology for all VLANs. When you connect a Cisco switch to a non-Cisco switch through an 802.1Q trunk, the Mono Spanning Tree of the non-Cisco switch and the native VLAN spanning tree of the Cisco switch combine to form a single spanning tree topology known as the Common Spanning Tree (CST).
- Because Cisco devices transmit BPDUs to the SSTP multicast MAC address on VLANs other than the native VLAN of the trunk, non-Cisco devices do not recognize these frames as BPDUs and flood them on all ports in the corresponding VLAN. Other Cisco devices connected to the non-Cisco 802.1Q cloud receive these flooded BPDUs. This BPDU reception allows Cisco switches to maintain a per-VLAN spanning tree topology across a cloud of non-Cisco 802.1Q devices. The non-Cisco 802.1Q cloud that separates the Cisco devices is treated as a single broadcast segment between all devices connected to the non-Cisco 802.1Q cloud through 802.1Q trunks.
- Make certain that the native VLAN is the same on all of the 802.1Q trunks that connect the Cisco devices to the non-Cisco 802.1Q cloud.
- If you are connecting multiple Cisco devices to a non-Cisco 802.1Q cloud, all of the connections must be through 802.1Q trunks. You cannot connect Cisco devices to a non-Cisco 802.1Q cloud through access ports because doing so places the access port on the Cisco device into the spanning tree “port inconsistent” state and no traffic will pass through the port.
- You can group trunk ports into port-channel groups, but all trunks in the group must have the same configuration. When a group is first created, all ports follow the parameters set for the first port to be added to the group. If you change the configuration of one of these parameters, the device propagates that setting to all ports in the group, such as the allowed VLANs and the trunk status. For example, if one port in a port group ceases to be a trunk, all ports cease to be trunks.
- If you try to enable 802.1X on a trunk port, an error message appears, and 802.1X is not enabled. If you try to change the mode of an 802.1X-enabled port to trunk, the port mode is not changed.
- Only ingress unicast packet counters are supported for SVI counters.
- When MAC addresses are cleared on a VLAN with the clear mac address-table dynamic command, the dynamic ARP (Address Resolution Protocol) entries on that VLAN are refreshed.
- If a static ARP entry exists on the VLAN and no MAC address to port mapping is present, the supervisor may generate an ARP request to learn the MAC address. Upon learning the MAC address, the adjacency entry points to the correct physical port.
- Cisco NX-OS does not support transparent bridging between two VLANs when one of the SVIs is on the Cisco Nexus 9000 using the BIA MAC (burned-in MAC address). This occurs when the BIA MAC

is shared between SVIs/VLANs. A MAC, different from the BIA MAC, can be configured under the SVI for transparent bridging to work properly.



Note This behavior is applicable to Cisco Nexus 9300 Switches (Network Forwarding Engine) and Cisco Nexus 9500 Switches with 95xx,96xx,94xx line cards. This behavior is not applicable to Cisco Nexus 9200 Switches, Cisco Nexus 9300-EX and Cisco Nexus 9500 Switches with 9700-EX line cards.

- Port-local VLANs do not support Fabric Extenders (FEX).
- On Cisco Nexus 9364C switches, auto-negotiation may not work on ports 49-64 when bringing up 100G links using QSFP-100G-CR4 cable. To workaround this issue, you must hard-code the speed on ports 49-64 and disable auto-negotiation.
- Cisco Nexus 9000 switch may fail to connect to a checkpoint with 1G SFP where an autonegotiation enabled Cisco Nexus 93180YC-EX switch is connected. The workaround is to disable autonegotiation and to manually enter the speed and duplex values.

Default Settings for Layer 2 Interfaces

The following table lists the default settings for device access and trunk port mode parameters.

Table 6: Default Access and Trunk Port Mode Parameters (7.0(3)I1(2) and earlier)

Parameters	Default
Switchport mode	Access
Allowed VLANs	1 to 3967, 4048 to 4094
Access VLAN ID	VLAN1
Native VLAN ID	VLAN1
Native VLAN ID tagging	Disabled
Administrative state	Shut

Table 7: Default Access and Trunk Port Mode Parameters (7.0(3)I2(1) and later)

Parameters	Default
Switchport mode	Access
Allowed VLANs	1 to 3967, 4048 to 4094
Access VLAN ID	VLAN1
Native VLAN ID	VLAN1

Parameters	Default
Native VLAN ID tagging	Disabled
Administrative state	Shut
SVI autostate	Enabled

Configuring Access and Trunk Interfaces



Note If you are familiar with the Cisco IOS CLI, be aware that the Cisco NX-OS commands for this feature might differ from the Cisco IOS commands that you would use.

Guidelines for Configuring Access and Trunk Interfaces

All VLANs on a trunk must be in the same VDC.

Configuring a VLAN Interface as a Layer 2 Access Port

You can configure a Layer 2 port as an access port. An access port transmits packets on only one, untagged VLAN. You specify which VLAN traffic that the interface carries, which becomes the access VLAN. If you do not specify a VLAN for an access port, that 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.

Before you begin

Ensure that you are configuring a Layer 2 interface.

Procedure

	Command or Action	Purpose
Step 1	configure terminal Example: <pre>switch# configure terminal switch(config)#</pre>	Enters global configuration mode.
Step 2	interface ethernet <i>{{type slot/port}}</i> <i>{port-channel number}</i> Example: <pre>switch(config)# interface ethernet 3/1 switch(config-if)#</pre>	Specifies an interface to configure, and enters interface configuration mode.

	Command or Action	Purpose
Step 3	switchport mode [access trunk] Example: <pre>switch(config-if) # switchport mode access</pre>	Sets the interface as a nontrunking nontagged, single-VLAN Layer 2 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	switchport access vlan <i>vlan-id</i> Example: <pre>switch(config-if) # switchport access vlan 5</pre>	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.
Step 5	exit Example: <pre>switch(config-if) # exit switch(config) #</pre>	Exits the interface configuration mode.
Step 6	show interface Example: <pre>switch# show interface</pre>	(Optional) Displays the interface status and information.
Step 7	no shutdown Example: <pre>switch# configure terminal switch(config) # int e3/1 switch(config-if) # no shutdown</pre>	(Optional) Clears the errors on the interfaces and VLANs where policies correspond with hardware policies. This command allows policy programming to continue and the port to come up. If policies do not correspond, the errors are placed in an error-disabled policy state.
Step 8	copy running-config startup-config Example: <pre>switch(config) # copy running-config startup-config</pre>	(Optional) Copies the running configuration to the startup configuration.

Example

This example shows how to set Ethernet 3/1 as a Layer 2 access port that carries traffic for VLAN 5 only:

```
switch# configure terminal
switch(config)# interface ethernet 3/1
switch(config-if)# switchport mode access
switch(config-if)# switchport access vlan 5
switch(config-if)#
```

Configuring Access Host Ports



Note You should apply the `switchport host` command only to interfaces that are connected to an end station.

You can optimize the performance of access ports that are connected to end stations by simultaneously setting that port as an access port. An access host port handles the 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 port channeling on that interface.



Note See “Configuring Port Channels” section and the [Cisco Nexus 9000 Series NX-OS Layer 2 Switching Configuration Guide](#) for information about port-channel interfaces

Before you begin

Ensure that you are configuring the correct interface to an interface that is an end station.

Procedure

	Command or Action	Purpose
Step 1	configure terminal Example: <pre>switch# configure terminal switch(config)#</pre>	Enters global configuration mode.
Step 2	interface ethernet <i>type slot/port</i> Example: <pre>switch(config)# interface ethernet 3/1 switch(config-if)#</pre>	Specifies an interface to configure, and enters interface configuration mode.
Step 3	switchport host Example: <pre>switch(config-if)# switchport host</pre>	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.
Step 4	exit Example: <pre>switch(config-if-range)# exit switch(config)#</pre>	Exits the interface mode.
Step 5	show interface Example: <pre>switch# show interface</pre>	(Optional) Displays the interface status and information.

	Command or Action	Purpose
Step 6	no shutdown Example: <pre>switch# configure terminal switch(config)# int e3/1 switch(config-if)# no shutdown</pre>	(Optional) Clears the errors on the interfaces and VLANs where policies correspond with hardware policies. This command allows policy programming to continue and the port to come up. If policies do not correspond, the errors are placed in an error-disabled policy state.
Step 7	copy running-config startup-config Example: <pre>switch(config)# copy running-config startup-config</pre>	(Optional) Copies the running configuration to the startup configuration.

Example

This example shows how to set Ethernet 3/1 as a Layer 2 access port with PortFast enabled and port channel disabled:

```
switch# configure terminal
switch(config)# interface ethernet 3/1
switch(config-if)# switchport host
switch(config-if)#
```

Configuring Trunk Ports

You can configure a Layer 2 port as a trunk port. A trunk port transmits untagged packets for one VLAN plus encapsulated, tagged, packets for multiple VLANs. (See the “IEEE 802.1Q Encapsulation” section for information about encapsulation.)



Note The device supports 802.1Q encapsulation only.

Before you begin

Before you configure a trunk port, ensure that you are configuring a Layer 2 interface.

Procedure

	Command or Action	Purpose
Step 1	configure terminal Example: <pre>switch# configure terminal switch(config)#</pre>	Enters global configuration mode.
Step 2	interface <i>{type slot/port port-channel number}</i> Example:	Specifies an interface to configure, and enters interface configuration mode.

	Command or Action	Purpose
	<code>switch(config)# interface ethernet 3/1</code> <code>switch(config-if)#</code>	
Step 3	switchport mode [access trunk] Example: <code>switch(config-if)# switchport mode trunk</code>	Sets the interface as a Layer 2 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.
Step 4	exit Example: <code>switch(config-if)# exit</code> <code>switch(config)#</code>	Exits the interface mode.
Step 5	show interface Example: <code>switch# show interface</code>	(Optional) Displays the interface status and information.
Step 6	no shutdown Example: <code>switch# configure terminal</code> <code>switch(config)# int e3/1</code> <code>switch(config-if)# no shutdown</code>	(Optional) Clears the errors on the interfaces and VLANs where policies correspond with hardware policies. This command allows policy programming to continue and the port to come up. If policies do not correspond, the errors are placed in an error-disabled policy state.
Step 7	copy running-config startup-config Example: <code>switch(config)# copy running-config startup-config</code>	(Optional) Copies the running configuration to the startup configuration.

Example

This example shows how to set Ethernet 3/1 as a Layer 2 trunk port:

```
switch# configure terminal
switch(config)# interface ethernet 3/1
switch(config-if)# switchport mode trunk
switch(config-if)#
```

Configuring the Native VLAN for 802.1Q Trunking Ports

You can configure the native VLAN for 802.1Q trunk ports. If you do not configure this parameter, the trunk port uses the default VLAN as the native VLAN ID.



Note You cannot configure an FCoE VLAN as a native VLAN for an Ethernet interface.

Procedure

	Command or Action	Purpose
Step 1	configure terminal Example: switch# configure terminal switch(config)#	Enters global configuration mode.
Step 2	interface <i>{{type slot/port}} {port-channel number}}</i> Example: switch(config)# interface ethernet 3/1 switch(config-if)#	Specifies an interface to configure, and enters interface configuration mode.
Step 3	switchport trunk native vlan <i>vlan-id</i> Example: switch(config-if)# switchport trunk native vlan 5	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.
Step 4	exit Example: switch(config-if-range)# exit switch(config)#	Exits interface configuration mode.
Step 5	show vlan Example: switch# show vlan	(Optional) Displays the status and information of VLANs.
Step 6	no shutdown Example: switch# configure terminal switch(config)# int e3/1 switch(config-if)# no shutdown	(Optional) Clears the errors on the interfaces and VLANs where policies correspond with hardware policies. This command allows policy programming to continue and the port to come up. If policies do not correspond, the errors are placed in an error-disabled policy state.
Step 7	copy running-config startup-config Example: switch(config)# copy running-config startup-config	(Optional) Copies the running configuration to the startup configuration.

Example

This example shows how to set the native VLAN for the Ethernet 3/1, Layer 2 trunk port to VLAN 5:

```
switch# configure terminal
switch(config)# interface ethernet 3/1
switch(config-if)# switchport trunk native vlan 5
switch(config-if)#
```

Configuring the Allowed VLANs for Trunking Ports

You can specify the IDs for the VLANs that are allowed on the specific trunk port.



Note The **switchport trunk allowed vlan** *vlan-list* command replaces the current VLAN list on the specified port with the new list. You are prompted for confirmation before the new list is applied.

If you are doing a copy and paste of a large configuration, you might see some failures because the CLI is waiting for a confirmation before accepting other commands. To avoid this problem, you can disable prompting by using the **terminal dont-ask** command before you paste the configuration.

Before you begin

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.



Note You can change the block of VLANs reserved for internal use. See the [Cisco Nexus 9000 Series NX-OS Layer 2 Switching Configuration Guide](#) for more information about changing the reserved VLANs.

Procedure

	Command or Action	Purpose
Step 1	configure terminal Example: switch# configure terminal switch(config)#	Enters global configuration mode.
Step 2	interface { ethernet <i>slot/port</i> port-channel <i>number</i> } Example: switch(config)# interface ethernet 3/1	Specifies an interface to configure, and enters interface configuration mode.
Step 3	switchport trunk allowed vlan { <i>vlan-list add</i> <i>vlan-list</i> all except <i>vlan-list</i> none remove <i>vlan-list</i> }	Sets the 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

	Command or Action	Purpose
	<p>Example:</p> <pre>switch(config-if)# switchport trunk allowed vlan add 15-20</pre>	<p>3968 to 4047 are the default VLANs reserved for internal use by default. By default, all VLANs are allowed on all trunk interfaces.</p> <p>The default reserved VLANs are 3968 to 4094, and you can change the block of reserved VLANs. See the Cisco Nexus 9000 Series NX-OS Layer 2 Switching Configuration Guide for more information.</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>
Step 4	<p>exit</p> <p>Example:</p> <pre>switch(config-if)# exit switch(config)#</pre>	Exits the interface mode.
Step 5	<p>show vlan</p> <p>Example:</p> <pre>switch# show vlan</pre>	(Optional) Displays the status and information for VLANs.
Step 6	<p>no shutdown</p> <p>Example:</p> <pre>switch# configure terminal switch(config)# int e3/1 switch(config-if)# no shutdown</pre>	(Optional) Clears the errors on the interfaces and VLANs where policies correspond with hardware policies. This command allows policy programming to continue and the port to come up. If policies do not correspond, the errors are placed in an error-disabled policy state.
Step 7	<p>copy running-config startup-config</p> <p>Example:</p> <pre>switch(config)# copy running-config startup-config</pre>	(Optional) Copies the running configuration to the startup configuration.

Example

This example shows how to add VLANs 15 to 20 to the list of allowed VLANs on the Ethernet 3/1, Layer 2 trunk port:

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

Configuring MAC addresses Limitation on a VLAN

Beginning with 7.0(3)F3(4), Cisco Nexus 9500 Series switches with N9K-X9636C-R, N9K-X9636Q-R, N9K-X9636C-RX line cards provides the ability to set an upper limit for the number of MAC addresses that can reside inside MAC address table of a Line-card Expansion-module (LEM). You can configure the limitations at System, VLAN, port, trunk and tunnel levels. For instance if the specified VLAN limitation is 2000 MACs, the Layer 2 Forwarding Manager (L2FM) accepts the first 2000 MACs it receives and reject the remaining MACs. To configure MAC address limitation on VLAN, complete the following steps:



Note Though this method stops learning MAC address learning via VLAN or system after the limit is reached, the traffic passes through the interface.

Procedure

	Command or Action	Purpose
Step 1	switch# configure terminal	Enters global configuration mode.
Step 2	switch(config)# mac address-table limit system value	Specifies an upper limit for MAC learning at system level.
Step 3	switch(config)# mac address-table limit vlan value	Specifies an upper limit for MAC learning at VLAN level.
Step 4	switch(config)# exit	Exits configuration mode.

Example

This example shows how to configure the upper limit for MAC learning at system and VLAN levels:

```
switch# configure terminal
switch(config)# mac address-table limit system 10000
Configuring Mac address limit will result in flushing existing Macs in the specified
VLAN/System.Proceed(yes/no)? [no] yes
switch(config)# mac address-table limit vlan 30 3000
Configuring Mac address limit will result in flushing existing Macs in the specified
VLAN/System.Proceed(yes/no)? [no] yes
switch(config)# exit
```

This example shows how to display the MAC address limitations:

```
switch# configure terminal
switch(config)# sh mac address-table limit
```

```
System Limit: 10000
```

```
Vlan      Learning Limit
----      -
1         196000
20        196000
30        3000
```

```
100      196000
switch(config)# exit
```

Configuring switchport isolated

Beginning with 7.0(3)I2(1), the switchport isolated feature is supported. Switchport isolated can be configured on an interface to accommodate up to 3967 VLANs on an interface. The interfaces that are configured with switchport isolated do not send STP BPDUs.



Note The **switchport isolated** mode is not supported on an interface that is connected to a FEX, a switch, router or any other networking devices. Switchport Isolated is not supported on the FEX HIF ports.

Procedure

	Command or Action	Purpose
Step 1	configure terminal Example: switch# configure terminal switch(config)#	Enters global configuration mode.
Step 2	interface {{ ethernet slot/port } port-channel number }} Example: switch(config)# interface ethernet 3/1 switch(config-if)#	Specifies an interface to configure, and enters interface configuration mode.
Step 3	switchport isolated Example: switch(config-if)# switchport isolated	Enables the switchport isolated feature.
Step 4	show running-config interface port-channel port-channel-number	(Optional) Displays the interface status and information.

Configuring a Default Interface

The default interface feature allows you to clear the existing configuration of multiple interfaces such as Ethernet, loopback, VLAN network, port-channel, and tunnel interfaces. All user configuration under a specified interface will be deleted. You can optionally create a checkpoint before clearing the interface configuration so that you can later restore the deleted configuration.



Note The default interface feature is not supported for management interfaces because the device could go to an unreachable state.

If the speed group is configured, the **default interface** command displays the following error:

```
Error: default interface is not supported as speed-group is configured
```

Procedure

	Command or Action	Purpose
Step 1	configure terminal Example: switch# configure terminal switch(config)#	Enters global configuration mode.
Step 2	default interface <i>int-if</i> [<i>checkpoint name</i>] Example: switch(config)# default interface ethernet 3/1 checkpoint test8	Deletes the configuration of the interface and restores the default configuration. Use the ? keyword to display the supported interfaces. Use the checkpoint keyword to store a copy of the running configuration of the interface before clearing the configuration.
Step 3	exit Example: switch(config)# exit switch(config)#	Exits global configuration mode.
Step 4	show interface Example: switch# show interface	(Optional) Displays the interface status and information.
Step 5	no shutdown Example: switch# configure terminal switch(config)# int e3/1 switch(config-if)# no shutdown	(Optional) Clears the errors on the interfaces and VLANs where policies correspond with hardware policies. This command allows policy programming to continue and the port to come up. If policies do not correspond, the errors are placed in an error-disabled policy state.

Example

This example shows how to delete the configuration of an Ethernet interface while saving a checkpoint of the running configuration for rollback purposes:

```
switch# configure terminal
switch(config)# default interface ethernet 3/1 checkpoint test8
.....Done
switch(config)#
```

Configuring SVI Autostate Exclude

Beginning with 7.0(3)I2(1), you can configure the SVI Autostate Exclude feature on an Ethernet interface or a port channel. You can use the Autostate Exclude option to enable or disable the port from bringing up or down the SVI calculation and applying it to all VLANs that are enabled on the selected port. You can also use the SVI Autostate Exclude VLAN feature to exclude a VLAN from the autostate excluded interface.

Procedure

	Command or Action	Purpose
Step 1	configure terminal Example: switch# configure terminal switch(config)#	Enters global configuration mode.
Step 2	interface <i>{{type slot/port} {port-channel number}}</i> Example: switch(config)# interface ethernet 3/1 switch(config-if)#	Specifies an interface to configure, and enters interface configuration mode.
Step 3	switchport Example: switch(config-if)# switchport	Sets the interface as a Layer 2 interface.
Step 4	[no] switchport autostate exclude Example: switch(config-if)# switchport autostate exclude	Excludes this port from the VLAN interface link-up calculation when there are multiple ports in the VLAN. To revert to the default settings, use the no form of this command.
Step 5	[no] switchport autostate exclude vlan <i>{vlan id all except}</i> Example: switch(config-if)# switchport autostate exclude vlan 10	(Optional) Excludes a vlan or a set of vlans from the autostate-excluded interface. This will help to minimize any disruption to the system. To revert to the default settings, use the no form of this command.
Step 6	exit Example: switch(config-if)# exit switch(config)#	Exits the interface configuration mode.
Step 7	show running-config interface <i>{{type slot/port} {port-channel number}}</i> Example: switch(config)# show running-config interface ethernet 3/1	(Optional) Displays configuration information about the specified interface.

	Command or Action	Purpose
Step 8	no shutdown Example: <pre>switch# configure terminal switch(config)# int e3/1 switch(config-if)# no shutdown</pre>	(Optional) Clears the errors on the interfaces and VLANs where policies correspond with hardware policies. This command allows policy programming to continue and the port to come up. If policies do not correspond, the errors are placed in an error-disabled policy state.
Step 9	copy running-config startup-config Example: <pre>switch(config)# copy running-config startup-config</pre>	(Optional) Copies the running configuration to the startup configuration.

Example

This example shows how to exclude a port from the VLAN interface link-up calculation on the Cisco NX-OS device:

```
switch# configure terminal
switch(config)# interface ethernet 3/1
switch(config-if)# switchport
switch(config-if)# switchport autostate exclude
```

This example shows how to exclude a VLAN from the auto-excluded interface:

```
switch# configure terminal
switch(config)# interface ethernet 3/1
switch(config-if)# switchport
switch(config-if)# switchport autostate exclude
switch(config-if)# switchport autostate exclude vlan 10
```

Configuring SVI Autostate Disable for the System

You can manage an SVI with the SVI autostate feature. You can configure the SVI autostate disable feature to keep an SVI up even if no interface is up in the corresponding VLAN. (Similarly, configure the SVI autostate enable feature so an SVI goes down when no interface is up in the corresponding VLAN). Use this procedure to configure this feature for the entire system.



Note The system default `interface-vlan autostate` command enables the SVI autostate feature.

Procedure

	Command or Action	Purpose
Step 1	configure terminal Example: <pre>switch# configure terminal switch(config)#</pre>	Enters global configuration mode.

	Command or Action	Purpose
Step 2	<p>[no] system default interface-vlan autostate</p> <p>Example:</p> <pre>switch(config)# no system default interface-vlan autostate</pre>	<p>Disables the default autostate behavior for the device.</p> <p>Note Use the system default interface-vlan autostate command to enable the autostate behavior for the device.</p>
Step 3	<p>no shutdown</p> <p>Example:</p> <pre>switch# configure terminal switch(config)# int e3/1 switch(config-if)# no shutdown</pre>	<p>(Optional) Clears the errors on the interfaces and VLANs where policies correspond with hardware policies. This command allows policy programming to continue and the port to come up. If policies do not correspond, the errors are placed in an error-disabled policy state.</p>
Step 4	<p>show running-config [all]</p> <p>Example:</p> <pre>switch(config)# show running-config</pre>	<p>(Optional) Displays the running configuration.</p> <p>To display the default and configured information, use the all keyword.</p>

Example

This example shows how to disable the default autostate behavior on the Cisco NX-OS device:

```
switch# configure terminal
switch(config)# no system default interface-vlan autostate
switch(config)# show running-config
```

Configuring SVI Autostate Disable Per SVI

You can configure SVI autostate enable or disable on individual SVIs. The SVI-level setting overrides the system-level SVI autostate configuration for that particular SVI.

Procedure

	Command or Action	Purpose
Step 1	<p>configure terminal</p> <p>Example:</p> <pre>switch# configure terminal switch(config)#</pre>	<p>Enters global configuration mode.</p>
Step 2	<p>feature interface-vlan</p> <p>Example:</p> <pre>switch(config)# feature interface-vlan</pre>	<p>Enables VLAN interface mode.</p>

	Command or Action	Purpose
Step 3	interface vlan <i>vlan-id</i> Example: switch(config-if) # interface vlan10 switch(config) #	Creates a VLAN interface and enters interface configuration mode. The range is from 1 and 4094.
Step 4	[no] autostate Example: switch(config-if) # no autostate	By default, enables the SVI autostate feature on specified interface. To disable the default settings, use the no form of this command.
Step 5	exit Example: switch(config-if) # exit switch(config) #	Exits the interface configuration mode.
Step 6	show running-config interface vlan <i>vlan-id</i> Example: switch(config) # show running-config interface vlan10	(Optional) Displays the running configuration for the specified VLAN interface.
Step 7	no shutdown Example: switch# configure terminal switch(config) # int e3/1 switch(config-if) # no shutdown	(Optional) Clears the errors on the interfaces and VLANs where policies correspond with hardware policies. This command allows policy programming to continue and the port to come up. If policies do not correspond, the errors are placed in an error-disabled policy state.
Step 8	show startup-config interface vlan <i>vlan-id</i> Example: switch(config) # show startup-config interface vlan10	(Optional) Displays the VLAN configuration in the startup configuration.

Example

This example shows how to disable the default autostate behavior on an individual SVI:

```
switch# configure terminal
switch(config) # feature interface-vlan
switch(config) # interface vlan10
switch(config-if) # no autostate
```

Configuring the Device to Tag Native VLAN Traffic

When you are working with 802.1Q trunked interfaces, you can maintain the tagging for all packets that enter with a tag that matches the value of the native VLAN ID and drops all untagged traffic (you will still carry

control traffic on that interface). This feature applies to the entire device; you cannot apply it to selected VLANs on a device.

The **vlan dot1q tag native global** command changes the behavior of all native VLAN ID interfaces on all trunks on the device.



Note If you enable 802.1Q tagging on one device and disable it on another device, all traffic is dropped on the device and this feature is disabled. You must configure this feature identically on each device.

Procedure

	Command or Action	Purpose
Step 1	configure terminal Example: switch# configure terminal switch(config)#	Enters global configuration mode.
Step 2	vlan dot1q tag native Example: switch(config)# vlan dot1q tag native	Modifies the behavior of a 802.1Q trunked native VLAN ID interface. The interface maintains the taggings for all packets that enter with a tag that matches the value of the native VLAN ID and drops all untagged traffic. The control traffic is still carried on the native VLAN.
Step 3	exit Example: switch(config-if-range)# exit switch(config)#	Exits the interface configuration mode.
Step 4	show vlan Example: switch# show vlan	(Optional) Displays the status and information for VLANs.
Step 5	no shutdown Example: switch# configure terminal switch(config)# int e3/1 switch(config-if)# no shutdown	(Optional) Clears the errors on the interfaces and VLANs where policies correspond with hardware policies. This command allows policy programming to continue and the port to come up. If policies do not correspond, the errors are placed in an error-disabled policy state.
Step 6	copy running-config startup-config Example: switch(config)# copy running-config startup-config	(Optional) Copies the running configuration to the startup configuration.

Example

This example shows how to change the behavior of the native VLAN on an 802.1Q trunked interface to maintain the tagged packets and drop all untagged traffic (except control traffic):

```
switch# configure terminal
switch(config)# vlan dot1q tag native
switch#
```

Configuring Interface Breakout Profile for 50-G Interfaces in a 16-Slot Chassis

The interface breakout profile is needed to breakout high bandwidth 100-G ports into two 50-G interfaces for slot 8 to 16 in the Cisco Nexus 9516 switch for -EX line cards.

Procedure

	Command or Action	Purpose
Step 1	configure terminal Example: switch# configure terminal	Enters global configuration mode.
Step 2	(Optional) interface breakout-profile 50g-2x-only Example: switch(config)# interface breakout-profile 50g-2x-only Warning: Please save config and reload the switch for breakout-profile config to take effect Please save config and reload the switch for the configuration to take effect	This command is required to breakout slots 8 to 16. It is not required for slots 1 to 7.
Step 3	copy running-config startup-config Example: switch(config-inf)# copy running-config startup-config [#####] 100% Copy complete, now saving to disk (please wait)... Copy complete.	Copies the running configuration to the startup configuration.
Step 4	reload Example: switch(config-inf)# reload This command will reboot the system. (y/n)? [n] y	Reboots the switch. Note After the switch reloads and the modules are up, enter the following CLI for any module or ports to breakout.

	Command or Action	Purpose
Step 5	interface breakout module <i>module-number</i> port <i>port-range</i> map [10g-4x 25g-4x 50g-2x] Example: <pre>switch(config)# interface breakout module 1 port 1-32 map 50g-2x</pre>	Breaks out the 100-Gb port to 2 50-Gb ports. The range of <i>module-number</i> is 1 to 30. The range of <i>port-range</i> is 1 to 72.

Changing the System Default Port Mode to Layer 2

You can set the system default port mode to Layer 2 access ports.

Procedure

	Command or Action	Purpose
Step 1	configure terminal Example: <pre>switch# configure terminal switch(config)#</pre>	Enters global configuration mode.
Step 2	system default switchport [shutdown] Example: <pre>switch(config-if)# system default switchport</pre>	Sets the default port mode for all interfaces on the system to Layer 2 access port mode and enters interface configuration mode. By default, all the interfaces are Layer 3. Note When the system default switchport shutdown command is issued: <ul style="list-style-type: none"> Any FEX HIFs that are not configured with no shutdown are shutdown. To avoid the shutdown, configure the FEX HIFs with no shut Any Layer 2 port that is not specifically configured with no shutdown are shutdown. To avoid the shutdown, configure the Layer 2 port with no shut
Step 3	exit Example: <pre>switch(config-if)# exit switch(config)#</pre>	Exits the interface configuration mode.

	Command or Action	Purpose
Step 4	show interface brief Example: switch# show interface brief	(Optional) Displays the status and information for interfaces.
Step 5	no shutdown Example: switch# configure terminal switch(config)# int e3/1 switch(config-if)# no shutdown	(Optional) Clears the errors on the interfaces and VLANs where policies correspond with hardware policies. This command allows policy programming to continue and the port to come up. If policies do not correspond, the errors are placed in an error-disabled policy state.
Step 6	copy running-config startup-config Example: switch(config)# copy running-config startup-config	(Optional) Copies the running configuration to the startup configuration.

Example

This example shows how to set the system ports to be Layer 2 access ports by default:

```
switch# configure terminal
switch(config-if)# system default switchport
switch(config-if)#
```

Verifying the Interface Configuration

To display access and trunk interface configuration information, perform one of the following tasks.

Command	Purpose
show interface ethernet <i>slot/port</i> [brief counters debounce description flowcontrol mac-address status transceiver]	Displays the interface configuration.
show interface brief	Displays interface configuration information, including the mode.
show interface switchport	Displays information, including access and trunk interface, information for all Layer 2 interfaces.
show interface trunk [module <i>module-number</i> vlan <i>vlan-id</i>]	Displays trunk configuration information.
show interface capabilities	Displays information about the capabilities of the interfaces.

Command	Purpose
<code>show running-config [all]</code>	Displays information about the current configuration. The all command displays the default and current configurations.
<code>show running-config interface ethernet slot/port</code>	Displays configuration information about the specified interface.
<code>show running-config interface port-channel slot/port</code>	Displays configuration information about the specified port-channel interface.
<code>show running-config interface vlan vlan-id</code>	Displays configuration information about the specified VLAN interface.

Monitoring the Layer 2 Interfaces

Use the following commands to display Layer 2 interfaces:

Command	Purpose
<code>clear counters interface [interface]</code>	Clears the counters.
<code>load- interval {interval seconds {1 2 3}}</code>	Cisco Nexus 9000 Series devices set three different sampling intervals to bit-rate and packet-rate statistics.
<code>show interface counters [module module]</code>	Displays input and output octets unicast packets, multicast packets, and broadcast packets.
<code>show interface counters detailed [all]</code>	Displays input packets, bytes, and multicast as well as output packets and bytes.
<code>show interface counters errors [module module]</code>	Displays information on the number of error packets.

Configuration Examples for Access and Trunk Ports

This example shows how to configure a Layer 2 access interface and assign the access VLAN mode for that interface:

```
switch# configure terminal
switch(config)# interface ethernet 2/30
switch(config-if)# switchport
switch(config-if)# switchport mode access
switch(config-if)# switchport access vlan 5
switch(config-if)#
```

This example shows how to configure a Layer 2 trunk interface, assign the native VLAN and the allowed VLANs, and configure the device to tag the native VLAN traffic on the trunk interface:

```
switch# configure terminal
switch(config)# interface ethernet 2/35
switch(config-if)# switchport
```

```

switch(config-if) # switchport mode trunk
switch(config-if) # switchport trunk native vlan 10
switch(config-if) # switchport trunk allowed vlan 5, 10
switch(config-if) # exit
switch(config) # vlan dot1q tag native
switch(config) #

```

Related Documents

Related Documents	Document Title
Configuring Layer 3 interfaces	Configuring Layer 2 Interfaces section
Port Channels	Configuring Port Channels section
VLANs, private VLANs, and STP	<i>Cisco Nexus 9000 Series NX-OS Layer 2 Switching Configuration Guide</i>
System management	<i>Cisco Nexus 9000 Series NX-OS System Management Configuration Guide</i>
High availability	<i>Cisco Nexus 9000 Series NX-OS High Availability and Redundancy Guide</i>
Licensing	<i>Cisco NX-OS Licensing Guide</i>
Release Notes	<i>Cisco Nexus 9000 Series NX-OS Release Notes</i>



CHAPTER 5

Configuring Layer 3 Interfaces

- [About Layer 3 Interfaces, on page 93](#)
- [Prerequisites for Layer 3 Interfaces, on page 99](#)
- [Guidelines and Limitations, on page 99](#)
- [Default Settings, on page 100](#)
- [Configuring Layer 3 Interfaces, on page 100](#)
- [Verifying the Layer 3 Interfaces Configuration, on page 121](#)
- [Monitoring the Layer 3 Interfaces, on page 122](#)
- [Configuration Examples for Layer 3 Interfaces, on page 123](#)
- [Related Documents, on page 125](#)

About Layer 3 Interfaces

Layer 3 interfaces forward IPv4 and IPv6 packets to another device using static or dynamic routing protocols. You can use Layer 3 interfaces for IP routing and inter-VLAN routing of Layer 2 traffic.

Routed Interfaces

You can configure a port as a Layer 2 interface or a Layer 3 interface. A routed interface is a physical port that can route IP traffic to another device. A routed interface is a Layer 3 interface only and does not support Layer 2 protocols, such as the Spanning Tree Protocol (STP).

All Ethernet ports are routed interfaces by default. You can change this default behavior with the CLI setup script.



Note The default behavior varies based on the type of switch (Cisco Nexus 9300, Cisco Nexus 9500, or Cisco Nexus 3164).



Note Cisco Nexus 9300 Series switches (except Cisco Nexus 9332 switch) have a Layer 2 default mode.

You can assign an IP address to the port, enable routing, and assign routing protocol characteristics to this routed interface.

You can also create a Layer 3 port channel from routed interfaces. For more information about port channels, see the “Configuring Port Channels” section.

Routed interfaces support exponentially decayed rate counters. Cisco NX-OS tracks the following statistics with these averaging counters:

- Input packets/sec
- Output packets/sec
- Input bytes/sec
- Output bytes/sec

Subinterfaces

You can create virtual subinterfaces on a parent interface configured as a Layer 3 interface. A parent interface can be a physical port.

Subinterfaces divide the parent interface into two or more virtual interfaces on which you can assign unique Layer 3 parameters such as IP addresses and dynamic routing protocols. The IP address for each subinterface should be in a different subnet from any other subinterface on the parent interface.


You create a subinterface with a name that consists of the parent interface name (for example, Ethernet 2/1) followed by a period and then by a number that is unique for that subinterface. For example, you could create a subinterface for Ethernet interface 2/1 named Ethernet 2/1.1 where .1 indicates the subinterface.

Cisco NX-OS enables subinterfaces when the parent interface is enabled. You can shut down a subinterface independent of shutting down the parent interface. If you shut down the parent interface, Cisco NX-OS shuts down all associated subinterfaces as well.

One use of subinterfaces is to provide unique Layer 3 interfaces to each virtual local area network (VLAN) supported by the parent interface. In this scenario, the parent interface connects to a Layer 2 trunking port on another device. You configure a subinterface and associate the subinterface to a VLAN ID using 802.1Q trunking.

The following figure shows a trunking port from a switch that connects to router B on interface E 2/1. This interface contains three subinterfaces that are associated with each of the three VLANs carried by the trunking port.

Figure 4: Subinterfaces for VLANs



For more information about VLANs, see the [Cisco Nexus 9000 Series NX-OS Layer 2 Switching Configuration Guide](#).

Limitations for Subinterfaces

The following are the limitations for subinterfaces:

- Statistics for subinterfaces are not supported.

VLAN Interfaces

A VLAN interface, or switch virtual interface (SVI), is a virtual routed interface that connects a VLAN on the device to the Layer 3 router engine on the same device. Only one VLAN interface can be associated with a VLAN, but you need to configure a VLAN interface for a VLAN only when you want to route between VLANs or to provide IP host connectivity to the device through a virtual routing and forwarding (VRF) instance that is not the management VRF. When you enable VLAN interface creation, Cisco NX-OS creates a VLAN interface for the default VLAN (VLAN 1) to permit remote switch administration.

You must enable the VLAN network interface feature before you can see configure it. The system automatically takes a checkpoint prior to disabling the feature, and you can roll back to this checkpoint. See the [Cisco Nexus 9000 Series NX-OS System Management Configuration Guide](#) for information on rollbacks and checkpoints.




Note You cannot delete the VLAN interface for VLAN 1.

You can route across VLAN interfaces to provide Layer 3 inter-VLAN routing by configuring a VLAN interface for each VLAN that you want to route traffic to and assigning an IP address on the VLAN interface. For more information about IP addresses and IP routing, see the [Cisco Nexus 9000 Series NX-OS Unicast Routing Configuration Guide](#).

The following figure shows two hosts connected to two VLANs on a device. You can configure VLAN interfaces for each VLAN that allows Host 1 to communicate with Host 2 using IP routing between the VLANs. VLAN 1 communicates at Layer 3 over VLAN interface 1 and VLAN 10 communicates at Layer 3 over VLAN interface 10.

Figure 5: Connecting Two VLANs with VLAN interfaces



Changing VRF Membership for an Interface

When you enter the **vrf member** command under an interface, you receive an alert regarding the deletion of interface configurations and to notify the clients/listeners (such as CLI-Server) to delete configurations with respect to the interface.

Entering the **system vrf-member-change retain-l3-config** command (7.0(3)I4(1) and later) enables the retention of the Layer 3 configuration when the VRF member changes on the interface. It does this by sending notification to the clients/listeners to store (buffer) the existing configurations, delete the configurations from the old vrf context, and reapply the stored configurations under the new VRF context.



Note When the **system vrf-member-change retain-l3-config** command is enabled, the Layer 3 configuration is not deleted and remains stored (buffered). When this command is not enabled (default mode), the Layer 3 configuration is not retained when the VRF member changes.

You can disable the retention of the Layer 3 configuration with the **no system vrf-member-change retain-l3-config** command. In this mode, the Layer 3 configuration is not retained when the VRF member changes.

Notes About Changing VRF Membership for an Interface

- Momentary traffic loss may occur when changing the VRF name.
- Only the configurations under the interface level are processed when the **system vrf-member-change retain-l3-config** command is enabled. You must manually process any configurations at the router level to accommodate routing protocols after a VRF change.
- The **system vrf-member-change retain-l3-config** command supports interface level configurations with:
 - Layer 3 configurations maintained by the CLI Server, such as **ip address** and **ipv6 address** (secondary) and all OSPF/ISIS/EIGRP CLIs available under the interface configuration.
 - HSRP
 - DHCP Relay Agent CLIs, such as **ip dhcp relay address [use-vrf]** and **ipv6 dhcp relay address [use-vrf]**.
- For DHCP:
 - As a best practice, the client and server interface VRF should be changed one at a time. Otherwise, the DHCP packets cannot be exchanged on the relay agent.
 - When the client and server are in different VRFs, use the **ip dhcp relay address [use-vrf]** command to exchange the DHCP packets in the relay agent over the different VRFs.

Loopback Interfaces

A loopback interface is a virtual interface with a single endpoint that is always up. Any packet transmitted over a loopback interface is immediately received by this interface. Loopback interfaces emulate a physical interface. You can configure up to 1024 loopback interfaces, numbered 0 to 1023.

You can use loopback interfaces for performance analysis, testing, and local communications. Loopback interfaces can act as a termination address for routing protocol sessions. This loopback configuration allows routing protocol sessions to stay up even if some of the outbound interfaces are down.

IP Unnumbered

The IP unnumbered feature enables the processing of IP packets on a point to point (p2p) interface without explicitly configuring a unique IP address on it. This approach borrows an IP address from another interface and conserves address space on point to point links.

Any interface which conforms to the point to point mode can be used as an IP unnumbered interface. For 7.0(3)I3(1) and later, the IP unnumbered feature is supported only on Ethernet interfaces and sub-interfaces. The borrowed interface can only be a loopback interface and is known as the numbered interface.

A loopback interface is ideal as a numbered interface in that it is always functionally up. However, because loopback interfaces are local to a switch/router, the reachability of unnumbered interfaces first needs to be established through static routes or by using an interior gateway protocol, such as OSPF or ISIS.

Beginning with Cisco NX-OS Release 7.0(3)I6(1), configuring IP unnumbered interfaces for port channels is supported on all Cisco Nexus 9000 Series switches.

MAC-Embedded IPv6 Address

BGP allows an IPv4 prefix to be carried over an IPv6 next hop. The IPv6 next hop is leveraged to remove neighbor discovery (ND)-related traffic from the network. To do this (7.0(3)I2(1) and later), the MAC address is embedded in the IPv6 address. Such an address is called a MAC-embedded IPv6 (MEv6) address. The router extracts the MAC address directly from the MEv6 address instead of going through ND. Local interface and next-hop MAC addresses are extracted from the IPv6 addresses.

On MEv6-enabled IPv6 interfaces, the same MEv6-extracted MAC address is used for IPv4 traffic as well. MEv6 is supported on all Layer 3-capable interfaces except switch virtual interfaces (SVIs).



Important When MEv6 is enabled on an interface, ping6 to the IPv6 link local address, OSPFv3, and BFDv6 are not supported on that interface.

High Availability

Layer 3 interfaces support stateful and stateless restarts. After the switchover, Cisco NX-OS applies the runtime configuration after the switchover.

See the [Cisco Nexus 9000 Series NX-OS High Availability and Redundancy Guide](#) for complete information about high availability.

Virtualization Support

Layer 3 interfaces support Virtual Routing and Forwarding instances (VRFs). VRFs exist within virtual device contexts (VDCs). By default, Cisco NX-OS places you in the default VDC and default VRF .



Note You must assign an interface to a VRF before you configure the IP address for that interface.

DHCP Client

Beginning with 7.0(3)I2(1), Cisco NX-OS supports DHCP client for IPv4 and IPv6 addresses on SVIs, physical Ethernet, and management interfaces. You can configure the IP address of a DHCP client by using the **ip address dhcp** or **ipv6 address dhcp** command. These commands send a request from the DHCP client to the DHCP server soliciting an IPv4 or IPv6 address from the DHCP server. The DHCP client on the Cisco Nexus

switch identifies itself to the DHCP server. The DHCP server uses this identifier to send the IP address back to the DHCP client.

When a DHCP client is configured on the SVI with the DHCP server sending router and DNS options, the **ip route 0.0.0.0/0 router-ip** and **ip name-server dns-ip** commands are configured on the switch automatically.

Limitations for Using DHCP Client on Interfaces

The following are the limitations for using DHCP client on interfaces:

- This feature is supported only on physical Ethernet interfaces, management interfaces, and SVIs.
- This feature is supported on non-default virtual routing and forwarding (VRF) instances.
- The DNS server and default router option-related configurations are saved in the startup configuration when you enter the **copy running-config startup-config** command. When you reload the switch, if this configuration is not applicable, you might have to remove it.
- You can configure a maximum of six DNS servers on the switch, which is a switch limitation. This maximum number includes the DNS servers configured by the DHCP client and the DNS servers configured manually.

If the number of DNS servers configured on the switch is more than six, and if you get a DHCP offer for an SVI with DNS option set, the IP address is not assigned to the SVI.

- A Cisco Nexus 9000 Series switch supports a maximum of 10 IPv4 and 10 IPv6 DHCP clients (7.0(3)I4(1) and later).
- DHCP relay and DHCP client configurations are incompatible and are not supported on the same switch. You must remove the DHCP relay configuration before configuring the DHCP Client on an interface.
- When DHCP snooping is enabled on the VLAN whose SVI is configured with the DHCP client, the DHCP snooping is not enforced on the SVI DHCP client.
- When configuring the IPv6 DHCP client, you must configure with the **ipv6 address use-link-local-only** command before the **ipv6 address dhcp** command.

Layer 3 Static MAC Addresses

You can configure a static MAC address for the following Layer 3 interfaces:

- Layer 3 interfaces
- Layer 3 subinterfaces
- Layer 3 port channels
- VLAN network interface



Note You cannot configure static MAC address on tunnel interfaces.

Prerequisites for Layer 3 Interfaces

Layer 3 interfaces have the following prerequisites:

- You are familiar with IP addressing and basic configuration. See the [Cisco Nexus 9000 Series NX-OS Unicast Routing Configuration Guide](#) for more information about IP addressing.

Guidelines and Limitations

Layer 3 interfaces have the following configuration guidelines and limitations:

- **show** commands with the **internal** keyword are not supported.
- Configuring a subinterface on a physical interface that is configured to be a member of a port-channel is not supported. One must configure the subinterface under the port-channel interface itself.
- If you change a Layer 3 interface to a Layer 2 interface, Cisco NX-OS shuts down the interface, reenables the interface, and removes all configuration specific to Layer 3. (7.0(3)I1(2) and later)
- If you change a Layer 2 interface to a Layer 3 interface, Cisco NX-OS shuts down the interface, reenables the interface, and deletes all configuration specific to Layer 2. (7.0(3)I1(2) and later)
- The Dynamic Host Configuration Protocol (DHCP) option is not supported when configuring a subinterface on a port-channel interface.
- When an IP unnumbered interface is configured, a loopback interface should be in the same VRF as the IP unnumbered interface. (7.0(3)I3(1) and later)
- An **admin-shutdown** command on a loopback interface that is a numbered interface does not bring down the IP unnumbered interface. This means that the routing protocols running over the IP unnumbered interface continue to be up. (7.0(3)I3(1) and later)
- Static routes running over the IP unnumbered interface should use pinned static routes. (7.0(3)I3(1) and later)



Note The IP unnumbered interface through which the route is resolved needs to be specified.

- An IP unnumbered interface is supported only on physical and sub-interfaces. (7.0(3)I3(1) and later)
- Only loopback interfaces can use unnumbered interfaces as numbered interfaces. (7.0(3)I3(1) and later)
- OSPF over an IP unnumbered interface is supported. (7.0(3)I3(1) and later)
- ISIS over an IP unnumbered interface is supported. (7.0(3)I3(1) and later)
- BGP over a loopback interface with an IP unnumbered interface as an overlay interface is supported. (7.0(3)I5(1) and later)
- The default and non-default VRF is supported by IP unnumbered interfaces. (7.0(3)I3(1) and later)

- The switch has a limit of 16 user-defined MAC addresses (MEv6/static). Configuring beyond this limit might result in issues documented in [CSCux84428](#).
- Layer 3 routed traffic for missing Layer 2 adjacency information is not flooded back onto VLAN members of ingress units when the source MAC address of routed traffic is a non-VDC (Virtual Device Context) MAC address. This limitation is for hardware flood traffic and can occur when the SVI (Switched Virtual Interface) has a user-configured MAC address.



Note If you are familiar with the Cisco IOS CLI, be aware that the Cisco NX-OS commands for this feature might differ from the Cisco IOS commands that you would use.

Default Settings

The following table lists the default settings for Layer 3 interface parameters.

Table 8: Default Layer 3 Interface Parameters

Parameters	Default
Admin state	Shut

Configuring Layer 3 Interfaces

Configuring a Routed Interface

You can configure any Ethernet port as a routed interface.

Procedure

	Command or Action	Purpose
Step 1	configure terminal Example: <code>switch# configure terminal</code> <code>switch(config)#</code>	Enters global configuration mode.
Step 2	interface ethernet <i>slot/port</i> Example: <code>switch(config)# interface ethernet 2/1</code> <code>switch(config-if)#</code>	Enters interface configuration mode.
Step 3	no switchport Example: <code>switch(config-if)# no switchport</code>	Configures the interface as a Layer 3 interface.

	Command or Action	Purpose
Step 4	<p>[ip address <i>ip-address/length</i> ipv6 address <i>ipv6-address/length</i>]</p> <p>Example:</p> <pre>switch(config-if) # ip address 192.0.2.1/8</pre> <p>Example:</p> <pre>switch(config-if) # ipv6 address 2001:0DB8::1/8</pre>	<ul style="list-style-type: none"> Configures an IP address for this interface. See the Cisco Nexus 9000 Series NX-OS Unicast Routing Configuration Guide for more information about IP addresses. Configures an IPv6 address for this interface. See the Cisco Nexus 9000 Series NX-OS Unicast Routing Configuration Guide for more information about IPv6 addresses.
Step 5	<p>show interfaces</p> <p>Example:</p> <pre>switch(config-if) # show interfaces ethernet 2/1</pre>	(Optional) Displays the Layer 3 interface statistics.
Step 6	<p>no shutdown</p> <p>Example:</p> <pre>switch# switch(config-if) # int e2/1 switch(config-if) # no shutdown</pre>	(Optional) Clears the errors on the interfaces where policies correspond with hardware policies. This command allows policy programming to continue and the port to come up. If policies do not correspond, the errors are placed in an error-disabled policy state.
Step 7	<p>copy running-config startup-config</p> <p>Example:</p> <pre>switch(config) # copy running-config startup-config</pre>	(Optional) Saves the configuration change.

Example

- Use the **medium** command to set the interface medium to either point to point or broadcast.

Command	Purpose
<p>medium {broadcast p2p}</p> <p>Example:</p> <pre>switch(config-if) # medium p2p medium p2p</pre>	Configures the interface medium as either point to point or broadcast.



Note The default setting is **broadcast**, and this setting does not appear in any of the **show** commands. However, if you do change the setting to **p2p**, you will see this setting when you enter the **show running config** command.

- Use the **switchport** command to convert a Layer 3 interface into a Layer 2 interface.

Command	Purpose
switchport Example: <pre>switch(config-if)# switchport switchport</pre>	Configures the interface as a Layer 2 interface and deletes any configuration specific to Layer 3 on this interface.

- This example shows how to configure a routed interface:

```
switch# configure terminal
switch(config)# interface ethernet 2/1
switch(config-if)# no switchport
switch(config-if)# ip address 192.0.2.1/8
switch(config-if)# copy running-config startup-config
```

The default setting for interfaces is routed. If you want to configure an interface for Layer 2, enter the **switchport** command. Then, if you change a Layer 2 interface to a routed interface, enter the **no switchport** command.

Configuring a Subinterface on a Routed Interface

You can configure one or more subinterfaces on a routed interface made from routed interfaces.

Before you begin

Configure the parent interface as a routed interface.

See the “Configuring a Routed Interface” section.

Procedure

	Command or Action	Purpose
Step 1	configure terminal Example: <pre>switch# configure terminal switch(config)#</pre>	Enters global configuration mode.
Step 2	interface ethernet <i>slot/port.number</i> Example: <pre>switch(config)# interface ethernet 2/1.1 switch(config-subif)#</pre>	Creates a subinterface and enters subinterface configuration mode. The number range is from 1 to 4094.
Step 3	[ip address <i>ip-address/length</i> ipv6 address <i>ipv6-address/length</i>] Example: <pre>switch(config-subif)# ip address 192.0.2.1/8</pre> Example: <pre>switch(config-subif)# ipv6 address 2001:0DB8::1/8</pre>	<ul style="list-style-type: none"> • Configures an IP address for this subinterface. See the Cisco Nexus 9000 Series NX-OS Unicast Routing Configuration Guide for more information on IP addresses. • Configures an IPv6 address for this subinterface. See the Cisco Nexus 9000 Series NX-OS Unicast Routing

	Command or Action	Purpose
		Configuration Guide for more information on IPv6 addresses.
Step 4	encapsulation dot1Q <i>vlan-id</i> Example: switch(config-subif)# encapsulation dot1Q 33	Configures IEEE 802.1Q VLAN encapsulation on the subinterface. The range is from 2 to 4093.
Step 5	show interfaces Example: switch(config-subif)# show interfaces ethernet 2/1.1	(Optional) Displays the Layer 3 interface statistics.
Step 6	copy running-config startup-config Example: switch(config)# copy running-config startup-config	(Optional) Saves the configuration change.

Example

- This example shows how to create a subinterface:

```
switch# configure terminal
switch(config)# interface ethernet 2/1.1
switch(config-if)# ip address 192.0.2.1/8
switch(config-if)# encapsulation dot1Q 33
switch(config-if)# copy running-config startup-config
```

- The output of the **show interface eth** command is enhanced for the subinterfaces as shown in the following :

```
switch# show interface ethernet 1/2.1
Ethernet1/2.1 is down (Parent Interface Admin down)
admin state is down, Dedicated Interface, [parent interface is Ethernet1/2]
Hardware: 40000 Ethernet, address: 0023.ac67.9bc1 (bia 4055.3926.61d4)
Internet Address is 10.10.10.1/24
MTU 1500 bytes, BW 40000000 Kbit, DLY 10 usec
reliability 255/255, txload 1/255, rxload 1/255
Auto-mdix is turned off
EtherType is 0x8100
L3 in Switched:
  ucast: 0 pkts, 0 bytes - mcast: 0 pkts, 0 bytes
L3 out Switched:
  ucast: 0 pkts, 0 bytes - mcast: 0 pkts, 0 bytes
```

Configuring a Subinterface on a Port-Channel Interface

You can configure one or more subinterfaces on a port-channel interface.



Note Subinterfaces on a port-channel interface do not support multicast routing, router ACLs, QoS, policy-based routing (PBR), SPAN, or ERSPAN.

Before you begin

Configure the parent interface as a port-channel interface. (7.0(3)I1(2) and later)

See the “Configuring Port Channels” chapter.

Procedure

	Command or Action	Purpose
Step 1	configure terminal Example: <pre>switch# configure terminal switch(config)#</pre>	Enters global configuration mode.
Step 2	interface port-channel <i>channel-id.number</i> Example: <pre>switch(config)# interface port-channel 100.1 switch(config-subif)#</pre>	Creates a subinterface and enters subinterface configuration mode.
Step 3	[ip address <i>ip-address/length</i> ipv6 address <i>ipv6-address/length</i>] Example: <pre>switch(config-subif)# ip address 192.0.2.1/8</pre> Example: <pre>switch(config-subif)# ipv6 address 2001:0DB8::1/8</pre>	<ul style="list-style-type: none"> Configures an IP address for this subinterface. See the Cisco Nexus 9000 Series NX-OS Unicast Routing Configuration Guide for more information on IP addresses. Configures an IPv6 address for this subinterface. See the Cisco Nexus 9000 Series NX-OS Unicast Routing Configuration Guide for more information on IPv6 addresses.
Step 4	encapsulation dot1Q <i>vlan-id</i> Example: <pre>switch(config-subif)# encapsulation dot1Q 33</pre>	Configures IEEE 802.1Q VLAN encapsulation on the subinterface. The range is from 2 to 4093.
Step 5	show interfaces Example: <pre>switch(config-subif)# show interfaces ethernet 2/1.1</pre>	(Optional) Displays the Layer 3 interface statistics.
Step 6	copy running-config startup-config Example:	(Optional) Saves the configuration change.

	Command or Action	Purpose
	<code>switch(config)# copy running-config startup-config</code>	

Example

This example shows how to create a subinterface:

```
switch# configure terminal
switch(config)# interface port-channel 115.3
switch(config-subif)# ip address 141.143.101.2/24
switch(config-subif)# encapsulation dot1q 3
switch(config-subif)# copy running-config startup-config
```

Configuring a VLAN Interface

You can create VLAN interfaces to provide inter-VLAN routing.

Procedure

	Command or Action	Purpose
Step 1	configure terminal Example: <code>switch# configure terminal</code> <code>switch(config)#</code>	Enters configuration mode.
Step 2	feature interface-vlan Example: <code>switch(config)# feature interface-vlan</code>	Enables VLAN interface mode.
Step 3	interface vlan <i>number</i> Example: <code>switch(config)# interface vlan 10</code> <code>switch(config-if)#</code>	Creates a VLAN interface. The number range is from 1 to 4094.
Step 4	[ip address <i>ip-address/length</i> ipv6 address <i>ipv6-address/length</i>] Example: <code>switch(config-if)# ip address 192.0.2.1/8</code> Example: <code>switch(config-if)# ipv6 address 2001:0DB8::1/8</code>	<ul style="list-style-type: none"> Configures an IP address for this VLAN interface. See the Cisco Nexus 9000 Series NX-OS Unicast Routing Configuration Guide for more information on IP addresses. Configures an IPv6 address for this VLAN interface. See the Cisco Nexus 9000 Series NX-OS Unicast Routing Configuration Guide for more information on IPv6 addresses.

	Command or Action	Purpose
Step 5	show interface vlan <i>number</i> Example: <pre>switch(config-if)# show interface vlan 10</pre>	(Optional) Displays the Layer 3 interface statistics.
Step 6	no shutdown Example: <pre>switch(config)# int e3/1 switch(config)# no shutdown</pre>	(Optional) Clears the errors on the interfaces where policies correspond with hardware policies. This command allows policy programming to continue and the port to come up. If policies do not correspond, the errors are placed in an error-disabled policy state.
Step 7	copy running-config startup-config Example: <pre>switch(config-if)# copy running-config startup-config</pre>	(Optional) Saves the configuration change.

Example

This example shows how to create a VLAN interface:

```
switch# configure terminal
switch(config)# feature interface-vlan
switch(config)# interface vlan 10
switch(config-if)# ip address 192.0.2.1/8
switch(config-if)# copy running-config startup-config
```

Enabling Layer 3 Retention During VRF Membership Change

The following steps enable the retention of the Layer 3 configuration when changing the VRF membership on the interface.

Procedure

	Command or Action	Purpose
Step 1	configure terminal Example: <pre>switch# configure terminal switch(config)#</pre>	Enters configuration mode.
Step 2	system vrf-member-change retain-l3-config Example: <pre>switch(config)# system vrf-member-change retain-l3-config</pre>	Enables Layer 3 configuration retention during VRF membership change. Note To disable the retention of the Layer 3 configuration, use the no system vrf-member-change retain-l3-config command.

	Command or Action	Purpose
	Warning: Will retain L3 configuration when vrf member change on interface.	

Configuring a Static MAC Address on a Layer 3 Interface

You can configure static MAC addresses on Layer 3 interfaces. You cannot configure broadcast or multicast addresses as static MAC addresses.



Note You cannot configure static MAC addresses on tunnel interfaces.



Note This configuration is limited to 16 VLAN interfaces. Applying the configuration to additional VLAN interfaces results in a down state for the interface with a `Hardware prog failed. status`.

Procedure

	Command or Action	Purpose
Step 1	config t Example: switch# config t switch(config)#	Enters configuration mode.
Step 2	interface [ethernet slot/port ethernet slot/port.number port-channel number vlan vlan-id] Example: switch(config)# interface ethernet 7/3	Specifies the Layer 3 interface and enters the interface configuration mode. Note You must create the Layer 3 interface before you can assign the static MAC address.
Step 3	mac-address mac-address Example: switch(config-if)# mac-address 22ab.47dd.ff89 switch(config-if)#	Specifies a static MAC address to add to the Layer 3 interface.
Step 4	exit Example: switch(config-if)# exit switch(config)#	Exits the interface mode.
Step 5	(Optional) show interface [ethernet slot/port ethernet slot/port.number port-channel number vlan vlan-id] Example:	Displays information about the Layer 3 interface.

	Command or Action	Purpose
	<code>switch# show interface ethernet 7/3</code>	
Step 6	(Optional) copy running-config startup-config Example: <code>switch# copy running-config startup-config</code>	Copies the running configuration to the startup configuration.

Example

This example shows how to configure the Layer 3 interface on slot 7, port 3 with a static MAC address:

```
switch# config t
switch(config)# interface ethernet 7/3
switch(config-if)# mac-address 22ab.47dd.ff89
switch(config-if)#
```

Configuring a Loopback Interface

You can configure a loopback interface to create a virtual interface that is always up.

Before you begin

Ensure that the IP address of the loopback interface is unique across all routers on the network.

Procedure

	Command or Action	Purpose
Step 1	configure terminal Example: <code>switch# configure terminal</code> <code>switch(config)#</code>	Enters configuration mode.
Step 2	interface loopback <i>instance</i> Example: <code>switch(config)# interface loopback 0</code> <code>switch(config-if)#</code>	Creates a loopback interface. The range is from 0 to 1023.
Step 3	[ip address <i>ip-address/length</i> ipv6 address <i>ipv6-address/length</i>] Example: <code>switch(config-if)# ip address 192.0.2.1/8</code> Example: <code>switch(config-if)# ipv6 address 2001:0DB8::1/8</code>	<ul style="list-style-type: none"> Configures an IP address for this interface. See the Cisco Nexus 9000 Series NX-OS Unicast Routing Configuration Guide for more information about IP addresses. Configures an IPv6 address for this interface. See the Cisco Nexus 9000 Series NX-OS Unicast Routing Configuration

	Command or Action	Purpose
		Guide for more information about IPv6 addresses.
Step 4	show interface loopback <i>instance</i> Example: <pre>switch(config-if)# show interface loopback 0</pre>	(Optional) Displays the loopback interface statistics.
Step 5	copy running-config startup-config Example: <pre>switch(config-if)# copy running-config startup-config</pre>	(Optional) Saves the configuration change.

Example

This example shows how to create a loopback interface:

```
switch# configure terminal
switch(config)# interface loopback 0
switch(config-if)# ip address 192.0.2.1/8
switch(config-if)# copy running-config startup-config
```

Configuring IP Unnumbered on an Ethernet Interface

You can configure the IP unnumbered feature on an ethernet interface.

Procedure

	Command or Action	Purpose
Step 1	configure terminal Example: <pre>switch# configure terminal switch(config)#</pre>	Enters global configuration mode.
Step 2	interface ethernet <i>slot/port</i> Example: <pre>switch(config)# interface ethernet 1/1 switch(config-if)#</pre>	Enters interface configuration mode.
Step 3	medium p2p Example: <pre>switch(config-if)# medium p2p</pre>	Configures the interface medium as point to point.
Step 4	ip unnumbered <i>type number</i> Example:	Enables IP processing on an interface without assigning an explicit IP address to the interface.

	Command or Action	Purpose
	<pre>switch(config-if)# ip unnumbered loopback 100</pre>	<p><i>type</i> and <i>number</i> specify another interface on which the router has an assigned IP address. The interface specified cannot be another unnumbered interface.</p> <p>Note <i>type</i> is limited to loopback. (7.0(3)I3(1) and later)</p>

Configuring OSPF for an IP Unnumbered Interface

You can configure OSPF for an IP unnumbered loopback interface.

Procedure

	Command or Action	Purpose
Step 1	<p>configure terminal</p> <p>Example:</p> <pre>switch# configure terminal switch(config)#</pre>	Enters global configuration mode.
Step 2	<p>interface ethernet <i>slot/port</i></p> <p>Example:</p> <pre>switch(config)# interface ethernet 1/20.1 switch(config-if)#</pre>	Enters interface configuration mode.
Step 3	<p>encapsulation dot1Q <i>vlan-id</i></p> <p>Example:</p> <pre>switch(config-if)# encapsulation dot1Q 100</pre>	Configures IEEE 802.1Q VLAN encapsulation on the subinterface. The range is from 2 to 4093.
Step 4	<p>medium p2p</p> <p>Example:</p> <pre>switch(config-if)# medium p2p</pre>	Configures the interface medium as point to point.
Step 5	<p>ip unnumbered <i>type number</i></p> <p>Example:</p> <pre>switch(config-if)# ip unnumbered loopback 101</pre>	<p>Enables IP processing on an interface without assigning an explicit IP address to the interface.</p> <p><i>type</i> and <i>number</i> specify another interface on which the router has an assigned IP address. The interface specified cannot be another unnumbered interface.</p> <p>Note <i>type</i> is limited to loopback. (7.0(3)I3(1) and later)</p>

	Command or Action	Purpose
Step 6	(Optional) ip ospf authentication Example: switch(config-if)# ip ospf authentication	Specifies the authentication type for interface.
Step 7	(Optional) ip ospf authentication-key password Example: switch(config-if)# ip ospf authentication 3 b7bdf15f62bbd250	Specifies the authentication password for OSPF authentication.
Step 8	ip router ospf instance area area-number Example: switch(config-if)# ip router ospf 100 area 0.0.0.1	Configures routing process for IP on an interface and specifies an area. Note The ip router ospf command is required for both the unnumbered and the numbered interface.
Step 9	no shutdown Example: switch(config-if)# no shutdown	Brings up the interface (administratively).
Step 10	interface loopback instance Example: switch(config)# interface loopback 101	Creates a loopback interface. The range is from 0 to 1023.
Step 11	ip address ip-address/length Example: switch(config-if)# 192.168.101.1/32	Configures an IP address for the interface.
Step 12	ip router ospf instance area area-number Example: switch(config-if)# ip router ospf 100 area 0.0.0.1	Configures routing process for IP on an interface and specifies an area. Note The ip router ospf command is required for both the unnumbered and the numbered interface.

Configuring ISIS for an IP Unnumbered Interface

You can configure ISIS for an IP unnumbered loopback interface.

Procedure

	Command or Action	Purpose
Step 1	configure terminal Example: switch# configure terminal switch(config)#	Enters global configuration mode.
Step 2	feature isis Example: Switch(config)# feature isis	Enables ISIS.
Step 3	router isis area-tag Example: Switch(config)# router isis 100	Assigns a tag to an IS-IS process and enters router configuration mode.
Step 4	net network-entity-title Example: Switch(config-router)# net 49.0001.0100.0100.1001.00	Configures the network entity title (NET) on the device.
Step 5	end Example: Switch(config-router)# end	Exit router configuration mode.
Step 6	interface ethernet slot/port Example: switch(config)# interface ethernet 1/20.1	Enters interface configuration mode.
Step 7	encapsulation dot1Q vlan-id Example: switch(config-subif)# encapsulation dot1Q 100	Configures IEEE 802.1Q VLAN encapsulation on the subinterface. The range is from 2 to 4093.
Step 8	medium p2p Example: switch(config-subif)# medium p2p	Configures the interface medium as point to point.
Step 9	ip unnumbered type number Example: switch(config-if)# ip unnumbered loopback 101	Enables IP processing on an interface without assigning an explicit IP address to the interface. <i>type</i> and <i>number</i> specify another interface on which the router has an assigned IP address. The interface specified cannot be another unnumbered interface.

	Command or Action	Purpose
		Note <i>type</i> is limited to loopback . (7.0(3)I3(1) and later)
Step 10	ip router isis <i>area-tag</i> Example: switch(config-subif) # ip router isis 100	Enables ISIS on the unnumbered interface.
Step 11	no shutdown Example: switch(config-subif) # no shutdown	Brings up the interface (administratively).

Configuring PBR on SVI on the Gateway

This procedure configures PBR on the primary SVI interface in the gateway.



Note Steps 2 through 6 are needed if you want to configure a PBR policy on the unnumbered Primary/Secondary VLAN interfaces. This is not mandatory for IP unnumbered on the SVI feature.

Procedure

	Command or Action	Purpose
Step 1	configure terminal Example: switch# configure terminal	Enter global configuration mode.
Step 2	ip access-list <i>list-name</i> Example: switch(config)# ip access-list pbr-sample	Configure access list.
Step 3	permit tcp host <i>ipaddr host ipaddr eq</i> <i>port-number</i> Example: switch(config-acl)# permit tcp host 10.1.1.1 host 192.168.2.1 eq 80	Specify the packets to forward on a specific port.
Step 4	exit Example: switch(config-acl) # exit	Exit configuration mode.

	Command or Action	Purpose
Step 5	route-map <i>route-map-name</i> Example: switch(config)# route-map pbr-sample	Create a route-map or enter route-map command mode.
Step 6	match ip address <i>access-list-name</i> Example: switch(config-route-map)# match ip address pbr-sample	Match values from the routing table.
Step 7	set ip next-hop <i>addr1</i> Example: switch(config-route-map)# set ip next-hop 192.168.1.1	Set IP address of the next hop.
Step 8	exit Example: switch(config-route-map)# exit	Exit command mode.
Step 9	interface vlan <i>vlan-id</i> Example: switch(config)# interface vlan 2003	Creates a VLAN interface and enters interface configuration mode. The range is from 1 and 4094. This is the primary VLAN.
Step 10	ip address <i>ip-addr</i> Example: switch(config-if)# ip address 10.0.0.1/8	Configures an IP address for the interface.
Step 11	no ip redirects Example: switch(config-if)# no ip redirects	Needs to be configured on all unnumbered primary and secondary VLAN interfaces.
Step 12	(Optional) ip policy route-map <i>pbr-sample</i> Example: switch(config-if)# ip policy route-map pbr-sample	Enter this command if you want to apply a PBR policy on the unnumbered Primary/Secondary VLAN interface.
Step 13	exit Example: switch(config-if)# exit	Exit command mode.
Step 14	hsrp version 2 Example: switch(config-if)# hsrp version 2	Set the HSRP version.
Step 15	hsrp group-num Example:	Set the HSRP group number.

	Command or Action	Purpose
	<code>switch(config-if) # hsrp 200</code>	
Step 16	name <i>name-val</i> Example: <code>switch(config-if-hsrp) # name primary</code>	Configure the redundancy name string.
Step 17	ip <i>ip-addr</i> Example: <code>switch(config-if-hsrp) # ip 10.0.0.100</code>	Configures an IP address.
Step 18	no shutdown Example: <code>switch(config-if-hsrp) # no shutdown</code>	Negates shutdown.

Configuring IP Unnumbered on SVI Secondary VLAN on the Gateway

This procedure configures IP unnumbered on the secondary SVI in the gateway.

Beginning with Cisco NX-OS Release 7.0(3)I7(1), this feature is supported on Cisco Nexus 93108TC-EX, 93180LC-EX and 93180YC-EX switches; and Cisco Nexus 9500 series switches with X9736C-EX, X97160YC-EX and X9732C-EX line cards.

Procedure

	Command or Action	Purpose
Step 1	configure terminal Example: <code>switch# configure terminal</code>	Enter configuration mode.
Step 2	interface vlan <i>vlan-list</i> Example: <code>switch(config) # interface vlan 2001</code>	Creates a VLAN interface and enters interface configuration mode. The range is from 1 to 4094. This is the secondary VLAN.
Step 3	ip unnumbered vlan <i>primary-vlan-id</i> Example: <code>switch(config-if) # ip unnumbered vlan 2003</code>	Enables IP processing on an interface without assigning an explicit IP address to an interface.
Step 4	(Optional) ip policy route-map <i>pbr-sample</i> Example: <code>switch(config-if) # ip policy route-map pbr-sample</code>	Enter this command if you want to apply a PBR policy on the unnumbered Primary/Secondary VLAN interface.
Step 5	no ip redirects Example:	Needs to be configured on all unnumbered primary and secondary VLAN interfaces.

	Command or Action	Purpose
	<code>switch(config-if)# no ip redirects</code>	
Step 6	hsrp version 2 Example: <code>switch(config-if)# hsrp version 2</code>	Set the HSRP version.
Step 7	hsrp group-num Example: <code>switch(config-if)# hsrp 200</code>	Set the HSRP group number.
Step 8	follow name Example: <code>switch(config-if-hsrp)# follow primary</code>	Configure the group to be followed.
Step 9	ip ip-addr Example: <code>switch(config-if-hsrp)# ip 10.0.0.100</code>	Enters HSRP IPv4 and sets the virtual IP address.
Step 10	no shutdown Example: <code>switch(config-if-hsrp)# no shutdown</code>	Negate shutdown.

Assigning an Interface to a VRF

You can add a Layer 3 interface to a VRF.

Procedure

	Command or Action	Purpose
Step 1	configure terminal Example: <code>switch# configure terminal</code> <code>switch(config)#</code>	Enters configuration mode.
Step 2	interface interface-type number Example: <code>switch(config)# interface loopback 0</code> <code>switch(config-if)#</code>	Enters interface configuration mode.
Step 3	vrf member vrf-name Example: <code>switch(config-if)# vrf member RemoteOfficeVRF</code>	Adds this interface to a VRF.

	Command or Action	Purpose
Step 4	ip address <i>ip-prefix/length</i> Example: switch(config-if)# ip address 192.0.2.1/16	Configures an IP address for this interface. You must do this step after you assign this interface to a VRF.
Step 5	show vrf [<i>vrf-name</i>] interface <i>interface-type number</i> Example: switch(config-vrf)# show vrf Enterprise interface loopback 0	(Optional) Displays VRF information.
Step 6	copy running-config startup-config Example: switch(config-if)# copy running-config startup-config	(Optional) Saves the configuration change.

Example

This example shows how to add a Layer 3 interface to the VRF:

```
switch# configure terminal
switch(config)# interface loopback 0
switch(config-if)# vrf member RemoteOfficeVRF
switch(config-if)# ip address 209.0.2.1/16
switch(config-if)# copy running-config startup-config
```

Configuring a MAC-Embedded IPv6 Address

Beginning with 7.0(3)I2(1), you can configure a MAC-embedded IPv6 (MEv6) address.

Procedure

	Command or Action	Purpose
Step 1	configure terminal Example: switch# configure terminal switch(config)#	Enters global configuration mode.
Step 2	interface <i>type slot/port</i> Example: switch(config)# interface ethernet 1/3 switch(config-if)#	Enters the interface configuration mode for the specified interface.
Step 3	no switchport Example: switch(config-if)# no switchport	Configures the interface as a Layer 3 interface and deletes any configuration specific to Layer 2 on this interface.

	Command or Action	Purpose
		Note To convert a Layer 3 interface back into a Layer 2 interface, use the switchport command.
Step 4	mac-address ipv6-extract Example: switch(config-if)# mac-address ipv6-extract	Extracts the MAC address embedded in the IPv6 address configured on the interface. Note The MEv6 configuration is currently not supported with the EUI-64 format of IPv6 address.
Step 5	ipv6 address ip-address/length Example: switch(config-if)# ipv6 address 2002:1::10/64	Configures an IPv6 address for this interface.
Step 6	ipv6 nd mac-extract [exclude nud-phase] Example: switch(config-if)# ipv6 nd mac-extract	Extracts the next-hop MAC address embedded in a next-hop IPv6 address. The exclude nud-phase option blocks packets during the ND phase only. When the exclude nud-phase option is not specified, packets are blocked during both ND and neighbor unreachability detection (NUD) phases.
Step 7	(Optional) show ipv6 icmp interface type slot/port Example: switch(config-if)# show ipv6 icmp interface ethernet 1/3	Displays IPv6 Internet Control Message Protocol version 6 (ICMPv6) interface information.
Step 8	(Optional) copy running-config startup-config Example: switch(config-if)# copy running-config startup-config	Copies the running configuration to the startup configuration.

Example

This example shows how to configure a MAC-embedded IPv6 address with ND mac-extract enabled:

```
switch# configure terminal
switch(config)# interface ethernet 1/3
switch(config-if)# no switchport
switch(config-if)# mac-address ipv6-extract
switch(config-if)# ipv6 address 2002:1::10/64
switch(config-if)# ipv6 nd mac-extract
switch(config-if)# show ipv6 icmp interface ethernet 1/3
ICMPv6 Interfaces for VRF "default"
Ethernet1/3, Interface status: protocol-up/link-up/admin-up
IPv6 address: 2002:1::10
IPv6 subnet: 2002:1::/64
```

```

IPv6 interface DAD state:  VALID
ND mac-extract :  Enabled
ICMPv6 active timers:
    Last Neighbor-Solicitation sent:  00:01:39
    Last Neighbor-Advertisement sent:  00:01:40
    Last Router-Advertisement sent:    00:01:41
    Next Router-Advertisement sent in:  00:03:34
Router-Advertisement parameters:
    Periodic interval: 200 to 600 seconds
    Send "Managed Address Configuration" flag: false
    Send "Other Stateful Configuration" flag: false
    Send "Current Hop Limit" field: 64
    Send "MTU" option value: 1500
    Send "Router Lifetime" field: 1800 secs
    Send "Reachable Time" field: 0 ms
    Send "Retrans Timer" field: 0 ms
    Suppress RA: Disabled
    Suppress MTU in RA: Disabled
Neighbor-Solicitation parameters:
    NS retransmit interval: 1000 ms
ICMPv6 error message parameters:
    Send redirects: true
    Send unreachable: false
ICMPv6-nd Statistics (sent/received):
    RAs: 3/0, RSs: 0/0, NAs: 2/0, NSs: 7/0, RDs: 0/0
    Interface statistics last reset: never

```

This example shows how to configure a MAC-embedded IPv6 address with ND mac-extract (excluding NUD phase) enabled:

```

switch# configure terminal
switch(config)# interface ethernet 1/5
switch(config-if)# no switchport
switch(config-if)# mac-address ipv6-extract
switch(config-if)# ipv6 address 2002:2::10/64
switch(config-if)# ipv6 nd mac-extract exclude nud-phase
switch(config-if)# show ipv6 icmp interface ethernet 1/5
ICMPv6 Interfaces for VRF "default"
Ethernet1/5, Interface status: protocol-up/link-up/admin-up
IPv6 address: 2002:2::10
IPv6 subnet: 2002:2::/64
IPv6 interface DAD state:  VALID
ND mac-extract :  Enabled (Excluding NUD Phase)
ICMPv6 active timers:
    Last Neighbor-Solicitation sent:  00:06:45
    Last Neighbor-Advertisement sent:  00:06:46
    Last Router-Advertisement sent:    00:02:18
    Next Router-Advertisement sent in:  00:02:24
Router-Advertisement parameters:
    Periodic interval: 200 to 600 seconds
    Send "Managed Address Configuration" flag: false
    Send "Other Stateful Configuration" flag: false
    Send "Current Hop Limit" field: 64
    Send "MTU" option value: 1500
    Send "Router Lifetime" field: 1800 secs
    Send "Reachable Time" field: 0 ms
    Send "Retrans Timer" field: 0 ms
    Suppress RA: Disabled
    Suppress MTU in RA: Disabled
Neighbor-Solicitation parameters:
    NS retransmit interval: 1000 ms
ICMPv6 error message parameters:
    Send redirects: true

```

```

Send unreachable: false
ICMPv6-nd Statistics (sent/received):
  RAs: 6/0, RSs: 0/0, NAs: 2/0, NSs: 7/0, RDs: 0/0
Interface statistics last reset: never

```

Configuring a DHCP Client on an Interface

You can configure the DHCP client on an SVI, a management interface, or a physical Ethernet interface for IPv4 or IPv6 address

Procedure

	Command or Action	Purpose
Step 1	switch# configure terminal	Enters global configuration mode.
Step 2	switch(config)# interface ethernet type slot/port mgmt mgmt-interface-number vlan vlan id	Creates a physical Ethernet interface, a management interface, or a VLAN interface. The range of <i>vlan id</i> is from 1 to 4094.
Step 3	switch(config-if)# [no] ipv6 address use-link-local-only	Prepares for request to the DHCP server. Note This command is only required for an IPv6 address.
Step 4	switch(config-if)# [no] [ip ipv6] address dhcp	Requests the DHCP server for an IPv4 or IPv6 address. The no form of this command removes any address that was acquired.
Step 5	(Optional) switch(config)# copy running-config startup-config	Saves the change persistently through reboots and restarts by copying the running configuration to the startup configuration.

Example

This example shows how to configure the IP address of a DHCP client on an SVI:

```

switch# configure terminal
switch(config)# interface vlan 15
switch(config-if)# ip address dhcp

```

This example shows how to configure an IPv6 address of a DHCP client on a management interface:

```

switch# configure terminal
switch(config)# interface mgmt 0
switch(config-if)# ipv6 address use-link-local-only
switch(config-if)# ipv6 address dhcp

```

Verifying the Layer 3 Interfaces Configuration

To display the Layer 3 configuration, perform one of the following tasks:

Command	Purpose
show interface ethernet <i>slot/port</i>	Displays the Layer 3 interface configuration, status, and counters (including the 5-minute exponentially decayed moving average of inbound and outbound packet and byte rates).
show interface ethernet <i>slot/port</i> brief	Displays the Layer 3 interface operational status.
show interface ethernet <i>slot/port</i> capabilities	Displays the Layer 3 interface capabilities, including port type, speed, and duplex.
show interface ethernet <i>slot/port</i> description	Displays the Layer 3 interface description.
show interface ethernet <i>slot/port</i> status	Displays the Layer 3 interface administrative status, port mode, speed, and duplex.
show interface ethernet <i>slot/port.number</i>	Displays the subinterface configuration, status, and counters (including the f-minute exponentially decayed moving average of inbound and outbound packet and byte rates).
show interface port-channel <i>channel-id.number</i>	Displays the port-channel subinterface configuration, status, and counters (including the 5-minute exponentially decayed moving average of inbound and outbound packet and byte rates).
show interface loopback <i>number</i>	Displays the loopback interface configuration, status, and counters.
show interface loopback <i>number</i> brief	Displays the loopback interface operational status.
show interface loopback <i>number</i> description	Displays the loopback interface description.
show interface loopback <i>number</i> status	Displays the loopback interface administrative status and protocol status.
show interface vlan <i>number</i>	Displays the VLAN interface configuration, status, and counters.
show interface vlan <i>number</i> brief	Displays the VLAN interface operational status.
show interface vlan <i>number</i> description	Displays the VLAN interface description.
show interface vlan <i>number</i> status	Displays the VLAN interface administrative status and protocol status.
show ip interface brief	Displays interface address and interface status (numbered/unnumbered).

Command	Purpose
show ip route	Displays routes learned via OSPF or ISIS. (Includes addresses for best unicast and multicast next-hop.)

Monitoring the Layer 3 Interfaces

Use the following commands to display Layer 3 statistics:

Command	Purpose
load- interval { <i>interval seconds</i> { 1 2 3 }}	Cisco Nexus 9000 Series devices set three different sampling intervals to bit-rate and packet-rate statistics. The range for VLAN network interface is 60 to 300 seconds, and the range for Layer interfaces is 30 to 300 seconds.
show interface ethernet <i>slot/port</i> counters	Displays the Layer 3 interface statistics (unicast, multicast, and broadcast).
show interface ethernet <i>slot/port</i> counters brief	Displays the Layer 3 interface input and output counters.
show interface ethernet errors <i>slot/port</i> detailed [all]	Displays the Layer 3 interface statistics. You can optionally include all 32-bit and 64-bit packet and byte counters (including errors).
show interface ethernet errors <i>slot/port</i> counters errors	Displays the Layer 3 interface input and output errors.
show interface ethernet errors <i>slot/port</i> counters snmp	Displays the Layer 3 interface counters reported by SNMP MIBs.
show interface ethernet <i>slot/port.number</i> counters	Displays the subinterface statistics (unicast, multicast, and broadcast).
show interface port-channel <i>channel-id.number</i> counters	Displays the port-channel subinterface statistics (unicast, multicast, and broadcast).
show interface loopback <i>number</i> counters	Displays the loopback interface input and output counters (unicast, multicast, and broadcast).
show interface loopback <i>number</i> detailed [all]	Displays the loopback interface statistics. You can optionally include all 32-bit and 64-bit packet and byte counters (including errors).
show interface loopback <i>number</i> counters errors	Displays the loopback interface input and output errors.
show interface vlan <i>number</i> counters	Displays the VLAN interface input and output counters (unicast, multicast, and broadcast).

Command	Purpose
show interface vlan <i>number</i> counters detailed [all]	Displays the VLAN interface statistics. You can optionally include all Layer 3 packet and byte counters (unicast and multicast).
show interface vlan <i>number</i> counters snmp	Displays the VLAN interface counters reported by SNMP MIBs.

Configuration Examples for Layer 3 Interfaces

This example shows how to configure Ethernet subinterfaces:

```
interface ethernet 2/1.10
description Layer 3
ip address 192.0.2.1/8
```

This example shows how to configure a loopback interface:

```
interface loopback 3
ip address 192.0.2.2/32
```

Example of Changing VRF Membership for an Interface

- Enable Layer 3 configuration retention when changing VRF membership.

```
switch# configure terminal
switch(config)# system vrf-member-change retain-l3-config
```

Warning: Will retain L3 configuration when vrf member change on interface.

- Verify Layer 3 retention.

```
switch# show running-config | include vrf-member-change
system vrf-member-change retain-l3-config
```

- Configure the SVI interface with Layer 3 configuration as VRF "blue".

```
switch# configure terminal
switch(config)# show running-config interface vlan 2002
```

```
interface Vlan2002
description TESTSVI
no shutdown
mtu 9192
vrf member blue
no ip redirects
ip address 192.168.211.2/27
ipv6 address 2620:10d:c041:12::2/64
ipv6 link-local fe80::1
ip router ospf 1 area 0.0.0.0
ipv6 router ospfv3 1 area 0.0.0.0
hsrp version 2
hsrp 2002
preempt delay minimum 300 reload 600
```

```

priority 110 forwarding-threshold lower 1 upper 110
ip 192.168.211.1
hsrp 2002 ipv6
preempt delay minimum 300 reload 600
priority 110 forwarding-threshold lower 1 upper 110
ip 2620:10d:c041:12::1

```

- Change the SVI interface VRF to "red".

```
switch# configure terminal
```

Enter configuration commands, one per line. End with CNTL/Z.

```
switch(config)# interface vlan 2002
switch(config-if)# vrf member red
```

Warning: Retain-L3-config is on, deleted and re-added L3 config on interface Vlan2002

- Verify SVI interface after VRF change.

```
switch# configure terminal
```

```
switch(config)# show running-config interface vlan 2002
```

```

interface Vlan2002
description TESTSVI
no shutdown
mtu 9192
vrf member red
no ip redirects
ip address 192.168.211.2/27
ipv6 address 2620:10d:c041:12::2/64
ipv6 link-local fe80::1
ip router ospf 1 area 0.0.0.0
ipv6 router ospfv3 1 area 0.0.0.0
hsrp version 2
hsrp 2002
preempt delay minimum 300 reload 600
priority 110 forwarding-threshold lower 1 upper 110
ip 192.168.211.1
hsrp 2002 ipv6
preempt delay minimum 300 reload 600
priority 110 forwarding-threshold lower 1 upper 110
ip 2620:10d:c041:12::1

```


**Note**

- When changing the VRF, the Layer 3 configuration retention affects:
 - Physical Interface
 - Loopback Interface
 - SVI Interface
 - Sub-interface
 - Tunnel Interface
 - Port-Channel
- When changing the VRF, the existing Layer 3 configuration is deleted and reapplied. All routing protocols, such as OSPF/ISIS/EIGRP/HSRP, go down in the old VRF and come up in the new VRF.
- Direct/Local IPv4/IPv6 addresses are removed from the old VRF and installed in the new VRF.
- Some traffic loss might occur during the VRF change.

Related Documents

Related Documents	Document Title
IP	<i>Cisco Nexus 9000 Series NX-OS Unicast Routing Configuration Guide</i>
VLANs	<i>Cisco Nexus 9000 Series NX-OS Layer 2 Switching Configuration Guide</i>



CHAPTER 6

Configuring Bidirectional Forwarding Detection

- [About BFD, on page 127](#)
- [Prerequisites for BFD, on page 129](#)
- [Guidelines and Limitations, on page 130](#)
- [Default Settings, on page 132](#)
- [Configuring BFD, on page 132](#)
- [Configuring BFD Support for Routing Protocols, on page 146](#)
- [Configuring BFD Interoperability, on page 156](#)
- [Verifying the BFD Configuration, on page 159](#)
- [Monitoring BFD, on page 160](#)
- [Configuration Examples for BFD, on page 160](#)
- [Related Documents, on page 161](#)
- [RFCs, on page 161](#)

About BFD

BFD is a detection protocol designed to provide fast forwarding-path failure detection times for media types, encapsulations, topologies, and routing protocols. You can use BFD to detect forwarding path failures at a uniform rate, rather than the variable rates for different protocol hello mechanisms. BFD makes network profiling and planning easier and reconvergence time consistent and predictable.

BFD provides subsecond failure detection between two adjacent devices and can be less CPU-intensive than protocol hello messages because some of the BFD load can be distributed onto the data plane on supported modules.

Asynchronous Mode


Cisco NX-OS supports the BFD asynchronous mode, which sends BFD control packets between two adjacent devices to activate and maintain BFD neighbor sessions between the devices. You configure BFD on both devices (or BFD neighbors). Once BFD has been enabled on the interfaces and on the appropriate protocols, Cisco NX-OS creates a BFD session, negotiates BFD session parameters, and begins to send BFD control packets to each BFD neighbor at the negotiated interval. The BFD session parameters include the following:

- **Desired minimum transmit interval**—The interval at which this device wants to send BFD hello messages.
- **Required minimum receive interval**—The minimum interval at which this device can accept BFD hello messages from another BFD device.

- Detect multiplier—The number of missing BFD hello messages from another BFD device before this local device detects a fault in the forwarding path.

The following figure shows how a BFD session is established. The figure shows a simple network with two routers running Open Shortest Path First (OSPF) and BFD. When OSPF discovers a neighbor (1), it sends a request to the local BFD process to initiate a BFD neighbor session with the OSPF neighbor router (2). The BFD neighbor session with the OSPF neighbor router is now established (3).

Figure 6: Establishing a BFD Neighbor Relationship



BFD Detection of Failures

Once a BFD session has been established and timer negotiations are complete, BFD neighbors send BFD control packets that act in the same manner as an IGP hello protocol to detect liveness, except at a more accelerated rate. BFD detects a failure, but the protocol must take action to bypass a failed peer.


BFD sends a failure detection notice to the BFD-enabled protocols when it detects a failure in the forwarding path. The local device can then initiate the protocol recalculation process and reduce the overall network convergence time.

The following figure shows what happens when a failure occurs in the network (1). The BFD neighbor session with the OSPF neighbor router is torn down (2). BFD notifies the local OSPF process that the BFD neighbor is no longer reachable (3). The local OSPF process tears down the OSPF neighbor relationship (4). If an alternative path is available, the routers immediately start converging on it.



Note Note The BFD failure detection occurs in less than a second, which is much faster than OSPF Hello messages could detect the same failure.

Figure 7: Tearing Down an OSPF Neighbor Relationship



Distributed Operation

Cisco NX-OS can distribute the BFD operation to compatible modules that support BFD. This process offloads the CPU load for BFD packet processing to the individual modules that connect to the BFD neighbors. All BFD session traffic occurs on the module CPU. The module informs the supervisor when a BFD failure is detected.

BFD Echo Function

Echo packets are defined and processed only by the transmitting system. For IPv4 and IPv6, the echo packets' destination address is that of the transmitting device. It is chosen in such a way as to cause the remote system to forward the packet back to the local system. This bypasses the routing lookup on the remote system and relies on the forwarding information base (FIB) instead. BFD can use the slow timer to slow down the asynchronous session when the echo function is enabled and reduce the number of BFD control packets that are sent between two BFD neighbors. The Echo function tests only the forwarding path of the remote system by having the remote (neighbor) system loop them back, so there is less inter-packet delay variability and faster failure detection times.

Security

Cisco NX-OS uses the packet Time to Live (TTL) value to verify that the BFD packets came from an adjacent BFD peer. For all asynchronous and echo request packets, the BFD neighbor sets the TTL value to 255 and the local BFD process verifies the TTL value as 255 before processing the incoming packet. For the echo response packet, BFD sets the TTL value to 254.

You can configure SHA-1 authentication of BFD packets.

High Availability

BFD supports stateless restarts. After a reboot or supervisor switchover, Cisco NX-OS applies the running configuration and BFD immediately sends control packets to the BFD peers.

Virtualization Support

BFD supports virtual routing and forwarding instances (VRFs). VRFs exist within virtual device contexts (VDCs). By default, Cisco NX-OS places you in the default VDC and default VRF.

Prerequisites for BFD

BFD has the following prerequisites:

- You must enable the BFD feature.
- Disable Internet Control Message Protocol (ICMP) redirect messages on BFD-enabled interfaces.
- Disable the IP packet verification check for identical IP source and destination addresses.
- See other detailed prerequisites that are listed with the configuration tasks.

Guidelines and Limitations

BFD has the following configuration guidelines and limitations:

- Beginning with Cisco NX-OS Release 7.0(3)I7(3), the QSFP 40/100-G BiDi comes up in the highest possible speed available on the port. For example, in the Cisco Nexus 93180LC-EX switch it comes up as 40 G in the first 28 ports and 100 G in the last 4 ports. If you need to connect to 40-G SR4 BiDi, the speed on the 40/100-G BiDi needs to be set to 40 G.
- BFD is not supported on unnumbered interfaces. Beginning with Cisco NX-OS Release 7.0(3)I7(6), the BFD process is prevented from running on unnumbered interfaces.
- Beginning with Cisco NX-OS Release 7.0(3)I7(5), the QSFP-40/100-SRBD comes up in the speed of 100-G on Cisco Nexus N9K-C93180YC-FX and N9K-C93180TC-FX switches.
- **show** commands with the **internal** keyword are not supported.
- Starting with Release 7.0(3)I5(1), BFD per-member link support is added on Cisco Nexus 9000 Series switches.
- BFD supports BFD version 1.
- BFD supports IPv4 and IPv6.
- BFD supports OSPFv3.
- BFD supports IS-ISv6.
- BFD supports BGPv6.
- BFD supports EIGRPv6.
- BFD supports only one session per address family, per Layer 3 interface.
- BFD supports only sessions which have unique (src_ip, dst_ip, interface/vrf) combination.
- BFD supports single-hop BFD.
- BFD for BGP supports single-hop EBGp and iBGP peers.
- BFD supports keyed SHA-1 authentication.
- BFD supports the following Layer 3 interfaces—physical interfaces, port channels, subinterfaces, and VLAN interfaces.
- BFD depends on a Layer 3 adjacency information to discover topology changes, including Layer 2 topology changes. A BFD session on a VLAN interface (SVI) may not be up after the convergence of the Layer 2 topology if there is no Layer 3 adjacency information available.
- For BFD on a static route between two devices, both devices must support BFD. If one or both of the devices do not support BFD, the static routes are not programmed in the Routing Information Base (RIB).
- Port channel configuration limitations:
 - For Layer 3 port channels used by BFD, you must enable LACP on the port channel.
 - For Layer 2 port channels used by SVI sessions, you must enable LACP on the port channel.

- SVI limitations:
 - An ASIC reset causes traffic disruption for other ports and it can cause the SVI sessions on the other ports to flap. For example, if the carrier interface is a virtual port channel (vPC), BFD is not supported over the SVI interface and it could cause a trigger for an ASIC reset. When a BFD session is over SVI using virtual port channel (vPC) Peer-Link, the BFD echo function is not supported. You must disable the BFD echo function for all sessions over SVI between vPC peer nodes.

An SVI on the Cisco Nexus series switches should not be configured to establish a BFD neighbor adjacency with a device connected to it via a vPC. This is because the BFD keepalives from the neighbour, if sent over the vPC member link connected to the vPC peer-switch, do not reach this SVI causing the BFD adjacency to fail.
 - When you change the topology (for example, add or delete a link into a VLAN, delete a member from a Layer 2 port channel, and so on), the SVI session could be affected. It may go down first and then come up after the topology discovery is finished.
 - BFD over FEX HIF interfaces is not supported.
 - When a BFD session is over SVI using virtual port-channel (vPC) Peer-Link (either BCM or GEM based ports), the BFD echo function is not supported. You must disable the BFD echo function for all sessions over SVI between vPC peer nodes using the **no bfd echo** command at the SVI configuration level.



Tip If you do not want the SVI sessions to flap and you need to change the topology, you can disable the BFD feature before making the changes and reenale BFD after the changes have been made. You can also configure the BFD timer to be a large value (for example, 5 seconds), and change it back to a fast timer after the above events complete.

- When you configure the BFD Echo function on the distributed Layer 3 port channels, reloading a member module flaps the BFD session hosted on that module, which results in a packet loss.

If you connect the BFD peers directly without a Layer 2 switch in between, you can use the BFD per-link mode as an alternative solution.



Note Using BFD per-link mode and subinterface optimization simultaneously on a Layer 3 port channel is not supported.

- When you specify a BFD neighbor prefix in the **clear {ip | ipv6} route *prefix*** command, the BFD echo session will flap.
- The **clear {ip | ipv6} route *** command causes BFD echo sessions to flap.
- HSRP for IPv4 is supported with BFD.
- BFD packets generated by the Cisco NX-OS device linecards are sent with COS 6/DSCP CS6. The DSCP/COS values for BFD packets are not user configurable.
- When configuring BFDv6 in no-bfd-echo mode, it is recommended to run with timers of 150 ms with a multiplier of 3.

- BFDv6 is not supported for VRRPv3 and HSRP for v6.
- IPv6 **igrp bfd** cannot be disabled on an interface.
- Port channel configuration notes:
 - When the BFD per-link mode is configured, the BFD echo function is not supported. You must disable the BFD echo function using the **no bfd echo** command before configuring the **bfd per-link** command.
 - Configuring BFD per-link with link-local is not supported.

Default Settings

The following table lists the default settings for BFD parameters.

Table 9: Default BFD Parameters

Parameters	Default
BFD feature	Disabled
Required minimum receive interval	50 milliseconds
Desired minimum transmit interval	50 milliseconds
Detect multiplier	3
Echo function	Enabled
Mode	Asynchronous
Port-channel	Logical mode (one session per source-destination pair address)
Slow timer	2000 milliseconds
Startup timer (7.0(3)I2(1) and later)	5 seconds

Configuring BFD

Configuration Hierarchy

You can configure BFD at the global level and at the interface level. The interface configuration overrides the global configuration.

For physical ports that are members of a port channel, the member port inherits the primary port channel BFD configuration.

Task Flow for Configuring BFD

Follow these steps in the following sections to configure BFD:

- Enabling the BFD Feature.
- Configuring Global BFD Parameters or Configuring BFD on an Interface.

Enabling the BFD Feature

You must enable the BFD feature before you can configure BFD on an interface and protocol.



Note Use the **no feature bfd** command to disable the BFD feature and remove all associated configuration.

Command	Purpose
no feature bfd Example: switch(config)# no feature bfd	Disables the BFD feature and removes all associated configuration.

Procedure

	Command or Action	Purpose
Step 1	configure terminal Example: switch# configure terminal switch(config)#	Enters configuration mode.
Step 2	feature bfd Example: switch(config)# feature bfd	Enables the BFD feature.
Step 3	show feature include bfd Example: switch(config)# show feature include bfd	(Optional) Displays enabled and disabled features.
Step 4	copy running-config startup-config Example: switch(config)# copy running-config startup-config	(Optional) Saves the configuration change.

Configuring Global BFD Parameters

You can configure the BFD session parameters for all BFD sessions on the device. The BFD session parameters are negotiated between the BFD peers in a three-way handshake.

See the Configuring BFD on an Interface section to override these global session parameters on an interface.

Before you begin

Enable the BFD feature.

Procedure

	Command or Action	Purpose
Step 1	configure terminal Example: <pre>switch# configure terminal switch(config)#</pre>	Enters configuration mode.
Step 2	bfd interval <i>mintx</i> <i>min_rx</i> <i>msec</i> <i>multiplier</i> <i>value</i> Example: <pre>switch(config)# bfd interval 50 min_rx 50 multiplier 3</pre>	Configures the BFD session parameters for all BFD sessions on the device. This command overrides these values by configuring the BFD session parameters on an interface. The <i>mintx</i> and <i>msec</i> range is from 50 to 999 milliseconds and the default is 50. The multiplier range is from 1 to 50. The multiplier default is 3.
Step 3	bfd slow-timer [<i>interval</i>] Example: <pre>switch(config)# bfd slow-timer 2000</pre>	Configures the slow timer used in the echo function. This value determines how fast BFD starts up a new session and at what speed the asynchronous sessions use for BFD control packets when the echo function is enabled. The slow-timer value is used as the new control packet interval, while the echo packets use the configured BFD intervals. The echo packets are used for link failure detection, while the control packets at the slower rate maintain the BFD session. The range is from 1000 to 30000 milliseconds. The default is 2000.
Step 4	[no] bfd startup-timer [<i>seconds</i>] Example: <pre>switch(config)# bfd startup-timer 20</pre>	Configures the BFD startup timer, which delays the startup time for BFD sessions in order to give the routes that are being used by local and remote routers time to settle down in the hardware. Using this feature can prevent BFD flaps in higher scale scenarios. The range is from 0 to 30 seconds. The default is 5 seconds. The bfd startup-timer 0 command disables the BFD startup timer.

	Command or Action	Purpose
		The no bfd startup-timer command sets the BFD startup timer to 5 seconds (the default value). Important bfd startup-timer command applicable for 7.0(3)I2(1) and later.
Step 5	bfd echo-interface loopback <i>interface number</i> Example: switch(config)# bfd echo-interface loopback 1 3	Configures the interface used for Bidirectional Forwarding Detection (BFD) echo frames. This command changes the source address for the echo packets to the one configured on the specified loopback interface. The interface number range is from 0 to 1023.
Step 6	show running-config bfd Example: switch(config)# show running-config bfd	(Optional) Displays the BFD running configuration.
Step 7	copy running-config startup-config Example: switch(config)# copy running-config startup-config	(Optional) Saves the configuration change.

Configuring BFD on an Interface

You can configure the BFD session parameters for all BFD sessions on an interface. The BFD session parameters are negotiated between the BFD peers in a three-way handshake.

This configuration overrides the global session parameters for the configured interface.

Before you begin

Ensure that Internet Control Message Protocol (ICMP) redirect messages are disabled on BFD-enabled interfaces. Use the **no ip redirects** command or the **no ipv6 redirects** command on the interface.

Enable the BFD feature. See the Enabling the BFD Feature section.

Procedure

	Command or Action	Purpose
Step 1	configure terminal Example: switch# configure terminal switch(config)#	Enters configuration mode.

	Command or Action	Purpose
Step 2	interface <i>int-if</i> Example: <pre>switch(config)# interface ethernet 2/1 switch(config-if)#</pre>	Enters interface configuration mode. Use the ? keyword to display the supported interfaces.
Step 3	bfd interval <i>mintx min_rx msec multiplier value</i> Example: <pre>switch(config-if)# bfd interval 50 min_rx 50 multiplier 3</pre>	<p>Configures the BFD session parameters for all BFD sessions on the device. This command overrides these values by configuring the BFD session parameters on an interface. The <i>mintx</i> and <i>msec</i> range is from 50 to 999 milliseconds and the default is 50. The multiplier range is from 1 to 50. The multiplier default is 3.</p> <p>Beginning with Cisco NX-OS Release 9.3(5), configuring BFD session parameters under interface with default timer values using the bfd interval 50 min_rx 50 multiplier 3 command is functionally equivalent to no bfd interval command.</p> <p>Once BFD session parameters under interface are set to default values, those BFD sessions running on that interface will inherit global session parameters, if present.</p>
Step 4	bfd authentication keyed-sha1 keyid <i>id key ascii_key</i> Example: <pre>switch(config-if)# bfd authentication keyed-sha1 keyid 1 ascii_key cisco123</pre>	<p>(Optional) Configures SHA-1 authentication for all BFD sessions on the interface. The <i>ascii_key</i> string is a secret key shared among BFD peers. The <i>id</i> value, a number between 0 and 255, is assigned to this particular <i>ascii_key</i>. BFD packets specify the key by <i>id</i>, allowing the use of multiple active keys.</p> <p>To disable SHA-1 authentication on the interface, use the no form of the command.</p>
Step 5	show running-config bfd Example: <pre>switch(config-if)# show running-config bfd</pre>	(Optional) Displays the BFD running configuration.
Step 6	copy running-config startup-config Example: <pre>switch(config-if)# copy running-config startup-config</pre>	(Optional) Saves the configuration change.

Configuring BFD on a Port Channel

You can configure the BFD session parameters for all BFD sessions on a port channel. If per-link mode is used for Layer 3 port channels, BFD creates a session for each link in the port channel and provides an aggregate result to client protocols. For example, if the BFD session for one link on a port channel is up, BFD informs client protocols, such as OSPF, that the port channel is up. The BFD session parameters are negotiated between the BFD peers in a three-way handshake.

This configuration overrides the global session parameters for the configured port channel. The member ports of the port channel inherit the port channel BFD session parameters.

Before you begin

Ensure that you enable LACP on the port channel before you enable BFD.

Ensure that Internet Control Message Protocol (ICMP) redirect messages are disabled on BFD-enabled interfaces. Use the **no ip redirects** command on the interface.

Enable the BFD feature. See the Enabling the BFD Feature section.

Procedure

	Command or Action	Purpose
Step 1	configure terminal Example: switch# configure terminal switch(config)#	Enters global configuration mode.
Step 2	interface port-channel <i>number</i> Example: switch(config)# interface port-channel 2 switch(config-if)#	Enters port-channel configuration mode. Use the ? keyword to display the supported number range.
Step 3	bfd per-link Example: switch(config-if)# bfd per-link	Configures the BFD sessions for each link in the port channel.
Step 4	bfd interval <i>mintx min_rx msec multiplier value</i> Example: switch(config-if)# bfd interval 50 min_rx 50 multiplier 3	(Optional) Configures the BFD session parameters for all BFD sessions on the port channel. This command overrides these values by configuring the BFD session parameters. The <i>mintx</i> and <i>msec</i> range is from 50 to 999 milliseconds and the default is 50. The multiplier range is from 1 to 50. The multiplier default is 3.
Step 5	bfd authentication keyed-sha1 <i>keyid id key ascii_key</i> Example:	(Optional) Configures SHA-1 authentication for all BFD sessions on the interface. The <i>ascii_key</i> string is a secret key shared among BFD peers. The <i>id</i> value, a number between 0

	Command or Action	Purpose
	<code>switch(config-if) # bfd authentication keyed-sha1 keyid 1 ascii_key cisco123</code>	and 255, is assigned to this particular <i>ascii_key</i> . BFD packets specify the key by <i>id</i> , allowing the use of multiple active keys. To disable SHA-1 authentication on the interface, use the no form of the command.
Step 6	show running-config bfd Example: <code>switch(config-if) # show running-config bfd</code>	(Optional) Displays the BFD running configuration.
Step 7	copy running-config startup-config Example: <code>switch(config-if) # copy running-config startup-config</code>	(Optional) Saves the configuration change.

Configuring the BFD Echo Function

You can configure the BFD echo function on one or both ends of a BFD-monitored link. The echo function slows down the required minimum receive interval, based on the configured slow timer. The RequiredMinEchoRx BFD session parameter is not set to zero if the echo function is disabled in compliance with RFC 5880. The slow timer becomes the required minimum receive interval if the echo function is enabled.

Before you begin

Enable the BFD feature. See the Enabling the BFD Feature section.

Configure the BFD session parameters. See the Configuring Global BFD Parameters section on or the Configuring BFD on an Interface section.

Ensure that Internet Control Message Protocol (ICMP) redirect messages are disabled on BFD-enabled interfaces. Use the **no ip redirects** command on the interface.

Procedure

	Command or Action	Purpose
Step 1	configure terminal Example: <code>switch# configure terminal switch(config)#</code>	Enters configuration mode.
Step 2	bfd slow-timer echo-interval Example: <code>switch(config)# bfd slow-timer 2000</code>	Configures the slow timer used in the echo function. This value determines how fast BFD starts up a new session and is used to slow down the asynchronous sessions when the BFD echo function is enabled. This value overwrites the required minimum receive interval when the echo function is enabled. The range is from

	Command or Action	Purpose
		1000 to 30000 milliseconds. The default is 2000.
Step 3	interface <i>int-if</i> Example: <pre>switch(config)# interface ethernet 2/1 switch(config-if)#</pre>	Enters interface configuration mode. Use the ? keyword to display the supported interfaces.
Step 4	bfd echo Example: <pre>switch(config-if)# bfd echo</pre>	Enables the echo function. The default is enabled.
Step 5	show running-config bfd Example: <pre>switch(config-if)# show running-config bfd</pre>	(Optional) Displays the BFD running configuration.
Step 6	copy running-config startup-config Example: <pre>switch(config-if)# copy running-config startup-config</pre>	(Optional) Saves the configuration change.

Configuring Per-Member Link BFD Sessions

Starting with Release 7.0(3)I5(1), BFD per-member link support is added on Cisco Nexus 9000 Series switches. See the following sections for more information.

BFD Enhancement to Address Per-link Efficiency

The Bidirectional Forwarding (BFD) enhancement to address per-link efficiency, called as IETF Micro BFD, lets you configure the individual BFD sessions on every Link Aggregation Group (LAG) member interfaces (as defined in RFC 7130).

With this enhancement, the BFD sessions run on each member link of the port-channel. If BFD detects a link failure, the member link is removed from the forwarding table. This mechanism delivers faster failure detection as the BFD sessions are created on an individual port-channel interface.

The BFD sessions running on member links of the port-channel are called as Micro BFD sessions. You can configure RFC 7130 BFD over main port-channel interface, that performs bandwidth monitoring over LAG by having one Micro BFD session over each member. If any of the member port goes down, the port is removed from the forwarding table and this prevents traffic disruption on that member.

Micro BFD sessions are supported for both LACP and non-LACP based-port channels. For more information on how to configure Micro BFD sessions, see *Configuring Micro BFD Sessions*.

Limitations of the IETF Bidirectional Forwarding Detection

See the following limitations of the IETF Bidirectional Forwarding Detection:

- BFD Limitations

- It cannot co-exist with BFD over logical port-channels or proprietary BFD per-member links. BFD IPv6 logical/proprietary per-link session is also not supported when BFD IETF IPv4 is configured on PC.
 - When you configure logical BFD session under any routing protocol, make sure that is not applied to any IETF port-channel. Having both logical and IETF configuration for same port-channel results in undefined behavior during ISSU/reloads.
 - IETF BFD IPv6 is not supported.
 - Echo functionality is not supported for Micro-BFD sessions.
 - Port-channel interfaces should be directly connected between two switches that are running the BFD sessions. No intermediate Layer 2 switches are expected.
- EthPCM/LACP Limitations
 - If a LACP port-channel has members in hot-standby state, BFD failure in one of the active links may not cause the hot-standby link to come up directly. Once the active link with BFD failure goes down, the hot-standby member becomes active. However, it may not be able to prevent the port-channel from going down before the hot-standby link comes up, in cases where port-channel min-link condition is hit.
- General Limitations:
 - It is supported only on Layer 3 port-channels.
 - It is not supported on the following:
 - vPC
 - Layer 3 sub-interfaces
 - Layer 2 port-channels/Layer 2 Fabric Path
 - FPC/HIF PC
 - Layer 3 sub-interfaces
 - SVI over port-channels

Guidelines for Migration/Configuration of IETF Per-Member Sessions:

See the following guidelines for migration/configuration of IETF per-member sessions:

- The logical BFD sessions that are created using the routing protocols over port-channel sub-interfaces (where RFC 7130 cannot run) are still supported. The main port-channel interface however does not support both logical and RFC 7130 sessions that co-exist. It can support only either of them.
- You can configure RFC 7130 BFD over the main port-channel interface that perform bandwidth monitoring over the LAG by having one Micro-BFD session over each member. If any of the member port goes down, BFD notifies it to the port-channel manager that removes the port from the LTL, thereby preventing blackholing of the traffic on that member.
- If the minimum number of links required to have the port-channel operationally *up* is not met in the above case, the port-channel is brought down by the port-channel manager. This in turn brings down the

port-channel sub-interfaces if they are configured and thereby the logical BFD session also comes down notifying the routing protocol.

- When you are using RFC 7130 on the main port-channel and logical BFD on the sub-interfaces, the logical BFD session should be run with lesser aggressive timers than the RFC 7130 BFD session. You can have RFC 7130 configured on the port-channel interface or you can have it configured in conjunction with the logical BFD sessions on the port-channel sub-interfaces.
- When a proprietary per-link is configured, enabling IETF Micro-BFD sessions is not allowed on a port channel and vice-versa. You have to remove the proprietary per-link configuration. Current implementation of proprietary per-link does not allow changing the configuration (no per-link), if there is any BFD session that is bootstrapped by the applications. You need to remove the BFD tracking on the respective applications and remove per-link configuration. The migration path from the proprietary per-link to IETF Micro-BFD is as follows:
 - Remove the BFD configuration on the applications.
 - Remove the per-link configuration.
 - Enable the IETF Micro-BFD command.
 - Enable BFD on the applications.

The same migration path can be followed for proprietary BFD to IETF Micro-BFD on the main port-channel interface.

Configuring Port Channel Interface

Before you begin

Ensure that the BFD feature is enabled.

Procedure

Step 1 switch(config)# **interface port-channel** *port-number*

Configures interface port-channel.

Step 2 switch(config-if)# **no switchport**

Configures interface as Layer 3 port-channel.

What to do next

- Configuring BFD Start Timer
- Enabling IETF Per-link BFD

(Optional) Configuring BFD Start Timer

Complete the following steps to configure the BFD start timer:

Procedure

```
switch(config-if)# port-channel bfd start 60
```

Configures the BFD start timer for a port-channel.

Note The default value is infinite (that is no timer is running). The range of BFD Start Timer value for port-channel is from 60 to 3600 seconds. For start timer to work, configure start timer value before completing the port-channel BFD configurations (that is before port-channel bfd track-member-link and port-channel bfd destination are configured for Layer 3 port-channel interface with the active members).

What to do next

- Enabling IETF Per-link BFD
- Configuring BFD Destination IP Address

Enabling IETF Per-link BFD

Procedure

```
switch(config-if)# port-channel bfd track-member-link
```

Enables IETF BFD on port-channel interface.

What to do next

- Configuring BFD Destination IP Address
- Verifying Micro BFD Session Configurations

Configuring BFD Destination IP Address

Complete the following steps to configure the BFD destination IP address:

Procedure

```
switch(config-if)# port-channel bfd destination ip-address
```

Configures an IPv4 address to be used for the BFD sessions on the member links.

What to do next

- Verifying Micro BFD Sessions Configuration

Verifying Micro BFD Session Configurations

Use the following commands to verify the Micro BFD session configurations.

Procedure

- Step 1** Displays the port-channel and port-channel member operational state.
switch# **show port-channel summary**
- Step 2** switch# **show bfd neighbors**
Displays Micro BFD sessions on port-channel members.
- Step 3** switch# **show bfd neighbors details**
Displays BFD session for a port channel interface and the associated Micro BFD sessions on members.
- Step 4** switch# **show tech-support bfd**
Displays the technical support information for BFD.
- Step 5** switch# **show tech-support lacp all**
Displays the technical support information for Ethernet Port Manager, Ethernet Port-channel Manager, and LACP.
- Step 6** switch# **show running-config interface port-channel** *port-channel-number*
Displays the running configuration information of the port-channel interface.
-


Examples: Configuring Micro BFD Sessions

See the following examples for configuring Micro BFD sessions.

Configuring Micro BFD Sessions

In this example, the following topology is used.

Figure 8: Configuring Micro BFD Session



The sample configuration of switch 1 is as follows:

```
feature bfd
configure terminal
  interface port-channel 10
    port-channel bfd track-member-link
    port-channel bfd destination 10.1.1.2
    port-channel bfd start 60
    ip address 10.1.1.1/24
```

The sample configuration of switch 2 is as follows:

```
feature bfd
configure terminal
  interface port-channel 10
    port-channel bfd track-member-link
    port-channel bfd destination 10.1.1.1
    port-channel bfd start 60
    ip address 10.1.1.2/24
```

Verifying Micro BFD Sessions Configuration

The following example displays the show output of the **show running-config interface port-channel** <port-channel>, **show port-channel summary**, **show bfd neighbors vrf internet_routes**, and **show bfd neighbors interface port-channel** <port-channel> **vrf internet_routes details** commands.

```
switch# show running-config interface port-channel 1001

!Command: show running-config interface port-channel1001
!Time: Fri Oct 21 09:08:00 2016

version 7.0(3)I5(1)

interface port-channel1001
  no switchport
  vrf member internet_routes
  port-channel bfd track-member-link
  port-channel bfd destination 40.4.1.2
  ip address 40.4.1.1/24
  ipv6 address 2001:40:4:1::1/64

switch# show por
port-channel port-profile
switch# show port-channel summary
```

```

Flags: D - Down          P - Up in port-channel (members)
       I - Individual    H - Hot-standby (LACP only)
       s - Suspended    r - Module-removed
       b - BFD Session Wait
       S - Switched     R - Routed
       U - Up (port-channel)
       p - Up in delay-lacp mode (member)
       M - Not in use. Min-links not met
    
```

```

-----
Group Port-      Type      Protocol  Member Ports
Channel
-----
1001 Po1001(RU)  Eth       LACP      Eth1/11/1(P) Eth1/11/2(P) Eth1/12/1(P)
                                     Eth1/12/2(P)
    
```

switch# show bfd neighbors vrf internet_routes

OurAddr	NeighAddr	LD/RD	RH/RS	Holdown(mult)	
State	Int	Vrf			
40.4.1.1	40.4.1.2	1090519041/0	Up	N/A (3)	Up
	Po1001	internet_routes			
40.4.1.1	40.4.1.2	1090519042/1090519051	Up	819 (3)	Up
	Eth1/12/1	internet_routes			
40.4.1.1	40.4.1.2	1090519043/1090519052	Up	819 (3)	Up
	Eth1/12/2	internet_routes			
40.4.1.1	40.4.1.2	1090519044/1090519053	Up	819 (3)	Up
	Eth1/11/1	internet_routes			
40.4.1.1	40.4.1.2	1090519045/1090519054	Up	819 (3)	Up
	Eth1/11/2	internet_routes			

switch#

switch# show bfd neighbors interface port-channel 1001 vrf internet_routes details

OurAddr	NeighAddr	LD/RD	RH/RS	Holdown(mult)	
State	Int	Vrf			
40.4.1.1	40.4.1.2	1090519041/0	Up	N/A (3)	Up
	Po1001	internet_routes			

Session state is Up

Local Diag: 0

Registered protocols: eth_port_channel

Uptime: 1 days 11 hrs 4 mins 8 secs

Hosting LC: 0, Down reason: None, Reason not-hosted: None

Parent session, please check port channel config for member info

switch#

switch# show bfd neighbors interface ethernet 1/12/1 vrf internet_routes details

OurAddr	NeighAddr	LD/RD	RH/RS	Holdown(mult)	
State	Int	Vrf			
40.4.1.1	40.4.1.2	1090519042/1090519051	Up	604 (3)	Up
	Eth1/12/1	internet_routes			

Session state is Up and not using echo function

Local Diag: 0, Demand mode: 0, Poll bit: 0, Authentication: None

MinTxInt: 100000 us, MinRxInt: 100000 us, Multiplier: 3

Received MinRxInt: 300000 us, Received Multiplier: 3

Holdown (hits): 900 ms (0), Hello (hits): 300 ms (458317)

Rx Count: 427188, Rx Interval (ms) min/max/avg: 19/1801/295 last: 295 ms ago

Tx Count: 458317, Tx Interval (ms) min/max/avg: 275/275/275 last: 64 ms ago

Registered protocols: eth_port_channel

Uptime: 1 days 11 hrs 4 mins 24 secs

Last packet: Version: 1 - Diagnostic: 0

State bit: Up - Demand bit: 0

Poll bit: 0 - Final bit: 0

```

Multiplier: 3          - Length: 24
My Discr.: 1090519051  - Your Discr.: 1090519042
Min tx interval: 300000 - Min rx interval: 300000
Min Echo interval: 300000 - Authentication bit: 0
Hosting LC: 1, Down reason: None, Reason not-hosted: None
Member session under parent interface Po1001

switch# show bfd neighbors interface ethernet 1/12/2 vrf internet_routes details

OurAddr      NeighAddr      LD/RD          RH/RS          Holdown(mult)
State        Int            Vrf
40.4.1.1     40.4.1.2       1090519043/1090519052 Up              799(3)          Up
              Eth1/12/2      internet_routes

Session state is Up and not using echo function
Local Diag: 0, Demand mode: 0, Poll bit: 0, Authentication: None
MinTxInt: 100000 us, MinRxInt: 100000 us, Multiplier: 3
Received MinRxInt: 300000 us, Received Multiplier: 3
Holdown (hits): 900 ms (0), Hello (hits): 300 ms (458336)
Rx Count: 427207, Rx Interval (ms) min/max/avg: 19/1668/295 last: 100 ms ago
Tx Count: 458336, Tx Interval (ms) min/max/avg: 275/275/275 last: 251 ms ago
Registered protocols: eth_port_channel
Uptime: 1 days 11 hrs 4 mins 30 secs
Last packet: Version: 1          - Diagnostic: 0
              State bit: Up      - Demand bit: 0
              Poll bit: 0       - Final bit: 0
              Multiplier: 3     - Length: 24
              My Discr.: 1090519052 - Your Discr.: 1090519043
              Min tx interval: 300000 - Min rx interval: 300000
              Min Echo interval: 300000 - Authentication bit: 0
Hosting LC: 1, Down reason: None, Reason not-hosted: None
Member session under parent interface Po1001
switch#

```

Configuring BFD Support for Routing Protocols

Configuring BFD on BGP

You can configure BFD for the Border Gateway Protocol (BGP).

Before you begin

Enable the BFD feature. See the Enabling the BFD Feature section.

Configure the BFD session parameters. See the Configuring Global BFD Parameters section or the Configuring BFD on an Interface section.

Enable the BGP feature. See the [Cisco Nexus 9000 Series NX-OS Unicast Routing Configuration Guide](#) for more information.

Procedure

	Command or Action	Purpose
Step 1	configure terminal Example: switch# configure terminal switch(config)#	Enters configuration mode.
Step 2	router bgp <i>as-number</i> Example: switch(config)# router bgp 64496 switch(config-router)#	Enables BGP and assigns the AS number to the local BGP speaker. The AS number can be a 16-bit integer or a 32-bit integer in the form of a higher 16-bit decimal number and a lower 16-bit decimal number in xx.xx format.
Step 3	neighbor (<i>ip-address</i> <i>ipv6-address</i>) remote-as <i>as-number</i> Example: switch(config-router)# neighbor 209.165.201.1 remote-as 64497 switch(config-router-neighbor)#	Configures the IPv4 or IPv6 address and AS number for a remote BGP peer. The <i>ip-address</i> format is x.x.x.x. The <i>ipv6-address</i> format is A::B::C:D.
Step 4	bfd Example: switch(config-router-neighbor)# bfd	Configures BFD on the neighbor.
Step 5	update-source <i>interface</i> Example: switch(config-router-neighbor)# update-source ethernet 2/1	Allows BGP sessions to use the primary IP address from a particular interface as the local address when forming a BGP session with a neighbor and enables BGP to register as a client with BFD.
Step 6	show running-config bgp Example: switch(config-router-neighbor)# show running-config bgp	(Optional) Displays the BGP running configuration.
Step 7	copy running-config startup-config Example: switch(config-router-neighbor)# copy running-config startup-config	(Optional) Saves the configuration change.

Configuring BFD on EIGRP

You can configure BFD for the Enhanced Interior Gateway Routing Protocol (EIGRP).

Before you begin

Enable the BFD feature. See the Enabling the BFD Feature section.

Configure the BFD session parameters. See the Configuring Global BFD Parameters section or the Configuring BFD on an Interface section.

Enable the EIGRP feature. See the [Cisco Nexus 9000 Series NX-OS Unicast Routing Configuration Guide](#) for more information.

Procedure

	Command or Action	Purpose
Step 1	configure terminal Example: <pre>switch# configure terminal switch(config)#</pre>	Enters configuration mode.
Step 2	router eigrp <i>instance-tag</i> Example: <pre>switch(config)# router eigrp Test1 switch(config-router)#</pre>	<p>Creates a new EIGRP process with the configured instance tag. The instance tag can be any case-sensitive, alphanumeric string up to 20 characters.</p> <p>If you configure an instance-tag that does not qualify as an AS number, you must use the autonomous-system to configure the AS number explicitly or this EIGRP instance will remain in the shutdown state.</p>
Step 3	bfd [ipv4 ipv6] Example: <pre>switch(config-router-neighbor)# bfd ipv4</pre>	(Optional) Enables BFD for all EIGRP interfaces.
Step 4	interface <i>int-if</i> Example: <pre>switch(config-router-neighbor)# interface ethernet 2/1 switch(config-if)#</pre>	Enters interface configuration mode. Use the ? keyword to display the supported interfaces.
Step 5	ip eigrp <i>instance-tag</i> bfd Example: <pre>switch(config-if)# ip eigrp Test1 bfd</pre>	<p>(Optional) Enables or disables BFD on an EIGRP interface. The instance tag can be any case-sensitive, alphanumeric string up to 20 characters.</p> <p>The default is disabled.</p>
Step 6	show ip eigrp [<i>vrf vrf-name</i>] [<i>interfaces if</i>] Example: <pre>switch(config-if)# show ip eigrp</pre>	(Optional) Displays information about EIGRP. The <i>vrf-name</i> can be any case-sensitive, alphanumeric string up to 32 characters.
Step 7	copy running-config startup-config Example: <pre>switch(config-if)# copy running-config startup-config</pre>	(Optional) Saves the configuration change.

Configuring BFD on OSPF

You can configure BFD for the Open Shortest Path First.

Before you begin

Enable the BFD feature. See the Enabling the BFD Feature section.

Configure the BFD session parameters. See the Configuring Global BFD Parameters section or the Configuring BFD on an Interface section.

Enable the OSPF feature. See the [Cisco Nexus 9000 Series NX-OS Unicast Routing Configuration Guide](#) for more information.

Procedure

	Command or Action	Purpose
Step 1	configure terminal Example: switch# configure terminal switch(config)#	Enters global configuration mode.
Step 2	router ospf instance-tag Example: switch(config)# router ospf 200 switch(config-router)#	Creates a new OSPF instance with the configured instance tag. The instance tag can be any case-sensitive, alphanumeric string up to 20 characters.
Step 3	bfd [ipv4 ipv6] Example: switch(config-router)# bfd	(Optional) Enables BFD for all OSPF interfaces.
Step 4	interface int-if Example: switch(config-router)# interface ethernet 2/1 switch(config-if)#	Enters interface configuration mode. Use the ? keyword to display the supported interfaces.
Step 5	ip ospf bfd Example: switch(config-if)# ip ospf bfd	(Optional) Enables or disables BFD on an OSPF interface. The default is disabled.
Step 6	show ip ospf [vrf vrf-name] [interfaces if] Example: switch(config-if)# show ip ospf	(Optional) Displays information about OSPF. The <i>vrf-name</i> can be any case-sensitive, alphanumeric string up to 32 characters.
Step 7	copy running-config startup-config Example: switch(config-if)# copy running-config startup-config	(Optional) Saves the configuration change.

Example Configurations for BFD on OSPF

Example configuration where BFD is enabled under a non-default VRF (OSPFv3 neighbors in vrf3).

```
configure terminal
router ospfv3 10
vrf vrf3
bfd
```

Configuring BFD on IS-IS

You can configure BFD for the Intermediate System-to-Intermediate System (IS-IS) protocol.

Before you begin

Enable the BFD feature. See the Enabling the BFD Feature section.

Configure the BFD session parameters. See the Configuring Global BFD Parameters section or the Configuring BFD on an Interface section.

Enable the IS-IS feature. See the [Cisco Nexus 9000 Series NX-OS Unicast Routing Configuration Guide](#) for more information.

Procedure

	Command or Action	Purpose
Step 1	configure terminal Example: <pre>switch# configure terminal switch(config)#</pre>	Enters global configuration mode.
Step 2	router isis instance-tag Example: <pre>switch(config)# router isis 100 switch(config-router)# net 49.0001.1720.1600.1001.00 switch(config-router)# address-family ipv6 unicast</pre>	Creates a new IS-IS instance with the configured <i>instance tag</i> .
Step 3	bfd [ipv4 ipv6] Example: <pre>switch(config-router)# bfd</pre>	(Optional) Enables BFD for all OSPF interfaces.
Step 4	interface int-if Example: <pre>switch(config-router)# interface ethernet 2/1 switch(config-if)#</pre>	Enters interface configuration mode. Use the ? keyword to display the supported interfaces.

	Command or Action	Purpose
Step 5	isis bfd Example: switch(config-if)# isis bfd	(Optional) Enables or disables BFD on an IS-IS interface. The default is disabled.
Step 6	show isis [vrf vrf-name] [interface if] Example: switch(config-if)# show isis	(Optional) Displays information about IS-IS. The <i>vrf-name</i> can be any case-sensitive, alphanumeric string up to 32 characters.
Step 7	copy running-config startup-config Example: switch(config-if)# copy running-config startup-config	(Optional) Saves the configuration change.

Example Configurations for BFD on IS-IS

Example configuration for IS-IS where BFD is enabled under IPv4 and an IPv6 address family.

```
configure terminal
router isis isis-1
  bfd
  address-family ipv6 unicast
  bfd
```

Configuring BFD on HSRP

You can configure BFD for the Hot Standby Router Protocol (HSRP). The active and standby HSRP routers track each other through BFD. If BFD on the standby HSRP router detects that the active HSRP router is down, the standby HSRP router treats this event as an active time expiry and takes over as the active HSRP router.

The **show hsrp detail** command shows this event as BFD@Act-down or BFD@Sby-down.

Before you begin

Enable the BFD feature. See the Enabling the BFD Feature section.

Configure the BFD session parameters. See the Configuring Global BFD Parameters section or the Configuring BFD on an Interface section.

Enable the HSRP feature. See the [Cisco Nexus 9000 Series NX-OS Unicast Routing Configuration Guide](#) for more information.

Procedure

	Command or Action	Purpose
Step 1	configure terminal Example:	Enters global configuration mode.

	Command or Action	Purpose
	<pre>switch# configure terminal switch(config)#</pre>	
Step 2	hsrp bfd all-interfaces Example: <pre>switch# hsrp bfd all-interfaces</pre>	(Optional) Enables or disables BFD on all HSRP interfaces. The default is disabled.
Step 3	interface int-if Example: <pre>switch(config-router)# interface ethernet 2/1 switch(config-if)#</pre>	Enters interface configuration mode. Use the ? keyword to display the supported interfaces.
Step 4	hsrp bfd Example: <pre>switch(config-if)# hsrp bfd</pre>	(Optional) Enables or disables BFD on an HSRP interface. The default is disabled.
Step 5	show running-config hsrp Example: <pre>switch(config-if)# show running-config hsrp</pre>	(Optional) Displays the HSRP running configuration.
Step 6	copy running-config startup-config Example: <pre>switch(config-if)# copy running-config startup-config</pre>	(Optional) Saves the configuration change.

Configuring BFD on VRRP

You can configure BFD for the Virtual Router Redundancy Protocol (VRRP). The active and standby VRRP routers track each other through BFD. If BFD on the standby VRRP router detects that the active VRRP router is down, the standby VRRP router treats this event as an active time rexpirt and takes over as the active VRRP router.

The **show vrrp detail** command shows this event as BFD@Act-down or BFD@Sby-down.

Before you begin

Enable the BFD feature. See the Enabling the BFD Feature section.

Configure the BFD session parameters. See the Configuring Global BFD Parameters section or the Configuring BFD on an Interface section.

Enable the VRRP feature. See the [Cisco Nexus 9000 Series NX-OS Unicast Routing Configuration Guide](#) for more information.

Procedure

	Command or Action	Purpose
Step 1	configure terminal Example: switch# configure terminal switch(config)#	Enters global configuration mode.
Step 2	interface <i>int-if</i> Example: switch(config)# interface ethernet 2/1 switch(config-if)#	Enters interface configuration mode. Use the ? keyword to display the supported interfaces.
Step 3	vrrp <i>group-no</i> Example: switch(config-if)# vrrp 2	Specifies the VRRP group number.
Step 4	vrrp bfd <i>address</i> Example: switch(config-if)# vrrp bfd	Enables or disables BFD on a VRRP interface. The default is disabled.
Step 5	show running-config vrrp Example: switch(config-if)# show running-config vrrp	(Optional) Displays the VRRP running configuration.
Step 6	copy running-config startup-config Example: switch(config-if)# copy running-config startup-config	(Optional) Saves the configuration change.

Configuring BFD on PIM

You can configure BFD for the Protocol Independent Multicast (PIM) protocol.

Before you begin

Enable the BFD feature. See the Enabling the BFD Feature section.

Enable the PIM feature. See the [Cisco Nexus 9000 Series NX-OS Unicast Routing Configuration Guide](#) for more information.

Procedure

	Command or Action	Purpose
Step 1	configure terminal Example: switch# configure terminal switch(config)#	Enters global configuration mode.
Step 2	ip pim bfd Example: switch(config)# ip pim bfd	Enables BFD for PIM.
Step 3	interface int-if Example: switch(config)# interface ethernet 2/1 switch(config-if)#	Enters interface configuration mode. Use the ? keyword to display the supported interfaces.
Step 4	ip pim bfd-instance [disable] Example: switch(config-if)# ip pim bfd-instance	(Optional) Enables or disables BFD on a PIM interface. The default is disabled.
Step 5	show running-config pim Example: switch(config)# show running-config pim	(Optional) Displays the PIM running configuration.
Step 6	copy running-config startup-config Example: switch(config)# copy running-config startup-config	(Optional) Saves the configuration change.

Configuring BFD on Static Routes

You can configure BFD for static routes on an interface. You can optionally configure BFD on a static route within a virtual routing and forwarding (VRF) instance.

Before you begin

Enable the BFD feature. See the Enabling the BFD Feature section.

Procedure

	Command or Action	Purpose
Step 1	configure terminal Example:	Enters global configuration mode.

	Command or Action	Purpose
	<code>switch# configure terminal</code> <code>switch(config)#</code>	
Step 2	vrf context <i>vrf-name</i> Example: <code>switch(config)# vrf context Red</code> <code>switch(config-vrf)#</code>	(Optional) Enters VRF configuration mode.
Step 3	ip route <i>route interface {nh-address nh-prefix}</i> Example: <code>switch(config-vrf)# ip route 192.0.2.1</code> <code>ethernet 2/1 192.0.2.4</code>	Creates a static route Use the ? keyword to display the supported interfaces.
Step 4	ip route static bfd <i>interface {nh-address nh-prefix}</i> Example: <code>switch(config-vrf)# ip route static bfd</code> <code>ethernet 2/1 192.0.2.4</code>	Enables BFD for all static routes on an interface. Use the? keyword to display the supported interfaces.
Step 5	show ip route static [<i>vrf vrf-name</i>] Example: <code>switch(config-vrf)# show ip route static</code> <code>vrf Red</code>	(Optional) Displays the static routes.
Step 6	copy running-config startup-config Example: <code>switch(config-vrf)# copy</code> <code>running-config startup-config</code>	(Optional) Saves the configuration change.

Disabling BFD on an Interface

You can selectively disable BFD on an interface for a routing protocol that has BFD enabled at the global or VRF level.

To disable BFD on an interface, use one of the following commands in interface configuration mode:

Command	Purpose
ip eigrp <i>instance-tag</i> bfd disable Example: <code>switch(config-if)# ip eigrp Test1 bfd</code> <code>disable</code>	Disables BFD on an EIGRP interface. The instance tag can be any case-sensitive, alphanumeric string up to 20 characters.
ip ospf bfd disable Example: <code>switch(config-if)# ip ospf bfd disable</code>	Disables BFD on an OSPFv2 interface.

Command	Purpose
isis bfd disable Example: switch(config-if)# isis bfd disable	Disables BFD on an IS-IS interface.

Disabling BFD on an Interface

Example configuration where BFD is disabled per interface.

```
configure terminal
interface port-channel 10
no ip redirects
ip address 22.1.10.1/30
ipv6 address 22:1:10::1/120
no ipv6 redirects
ip router ospf 10 area 0.0.0.0
ip ospf bfd disable          /*** disables IPv4 BFD session for OSPF
ospfv3 bfd disable          /*** disables IPv6 BFD session for OSPFv3
```

Configuring BFD Interoperability

Configuring BFD Interoperability in Cisco NX-OS Devices in a Point-to-Point Link

Procedure

	Command or Action	Purpose
Step 1	configure terminal Example: switch# configure terminal switch(config)#	Enters global configuration mode.
Step 2	interface port-channel <i>int-if</i> Example: switch(config-if)# interface ethernet 2/1	Enters interface configuration mode. Use the ? keyword to display the supported interfaces.
Step 3	ip ospf bfd Example: switch(config-if)# ip ospf bfd	Enables BFD on an OSPFv2 interface. The default is disabled. OSPF is used as an example. You can enable BFD of any of the supported protocols.
Step 4	no ip redirects Example:	Prevents the device from sending redirects.

	Command or Action	Purpose
	<code>switch(config-if)# no ip redirects</code>	
Step 5	bfd interval <i>mintx min_rx msec multiplier value</i> Example: <code>switch(config-if)# bfd interval 50 min_rx 50 multiplier 3</code>	Configures the BFD session parameters for all BFD sessions on the port channel. This command overrides these values by configuring the BFD session parameters. The <i>mintx</i> and <i>msec</i> range is from 50 to 999 milliseconds and the default is 50. The multiplier range is from 1 to 50. The multiplier default is 3.
Step 6	exit Example: <code>switch(config-if)# exit</code>	Exits interface configuration mode and returns to EXEC mode.

Configuring BFD Interoperability in Cisco NX-OS Devices in a Switch Virtual Interface

Procedure

	Command or Action	Purpose
Step 1	configure terminal Example: <code>switch# configure terminal</code> <code>switch(config)#</code>	Enters global configuration mode.
Step 2	interface port-channel <i>vlan vlan-id</i> Example: <code>switch(config)# interface vlan 998</code> <code>switch(config-if)#</code>	Creates a dynamic Switch Virtual Interface (SVI).
Step 3	bfd interval <i>mintx min_rx msec multiplier value</i> Example: <code>switch(config-if)# bfd interval 50 min_rx 50 multiplier 3</code>	Configures the BFD session parameters for all BFD sessions on the device. The <i>mintx</i> and <i>msec</i> range is from 50 to 999 milliseconds and the default is 50. The multiplier range is from 1 to 50. The multiplier default is 3.
Step 4	no ip redirects Example: <code>switch(config-if)# no ip redirects</code>	Prevents the device from sending redirects.
Step 5	ip address <i>ip-address/length</i> Example: <code>switch(config-if)# ip address 10.1.0.253/24</code>	Configures an IP address for this interface.

	Command or Action	Purpose
Step 6	ip ospf bfd Example: switch(config-if) # ip ospf bfd	Enables BFD on an OSPFv2 interface. The default is disabled.
Step 7	exit Example: switch(config-if) # exit	Exits interface configuration mode and returns to EXEC mode.

Configuring BFD Interoperability in Cisco NX-OS Devices in Logical Mode

Procedure

	Command or Action	Purpose
Step 1	configure terminal Example: switch# configure terminal switch(config)#	Enters global configuration mode.
Step 2	interface port-channel type number.subinterface-id Example: switch(config-if) # interface port-channel 50.2	Enters port channel configuration mode. Use the ? keyword to display the supported number range.
Step 3	bfd interval mintx min_rx msec multiplier value Example: switch(config-if) # bfd interval 50 min_rx 50 multiplier 3	Configures the BFD session parameters for all BFD sessions on the port channel. The <i>mintx</i> and <i>msec</i> range is from 50 to 999 milliseconds and the default is 50. The multiplier range is from 1 to 50. The multiplier default is 3.
Step 4	no ip redirects Example: switch(config-if) # no ip redirects	Prevents the device from sending redirects.
Step 5	ip ospf bfd Example: switch(config-if) # ip ospf bfd	Enables BFD on an OSPFv2 interface. The default is disabled. OSPF is used as an example. You can enable BFD of any of the supported protocols.
Step 6	exit Example: switch(config-if) # exit	Exits interface configuration mode and returns to EXEC mode.

Verifying BFD Interoperability in a Cisco Nexus 9000 Series Device

The following example shows how to verify BFD interoperability in a Cisco Nexus 9000 Series device.

```
switch# show bfd neighbors details
OurAddr NeighAddr LD/RD RH/RS Holdown(mult) State Int
Vrf
10.1.1.1 10.1.1.2 1140850707/2147418093 Up 6393(4) Up Vlan2121
default
Session state is Up and using echo function with 50 ms interval
Local Diag: 0, Demand mode: 0, Poll bit: 0, Authentication: None
MinTxInt: 50000 us, MinRxInt: 2000000 us, Multiplier: 3
Received MinRxInt: 2000000 us, Received Multiplier: 4
Holdown (hits): 8000 ms (0), Hello (hits): 2000 ms (108)
Rx Count: 92, Rx Interval (ms) min/max/avg: 347/1996/1776 last: 1606 ms ago
Tx Count: 108, Tx Interval (ms) min/max/avg: 1515/1515/1515 last: 1233 ms ago
Registered protocols: ospf
Uptime: 0 days 0 hrs 2 mins 44 secs
Last packet: Version: 1 - Diagnostic: 0
State bit: Up - Demand bit: 0
Poll bit: 0 - Final bit: 0
Multiplier: 4 - Length: 24
My Discr.: 2147418093 - Your Discr.: 1140850707
Min tx interval: 2000000 - Min rx interval: 2000000
Min Echo interval: 1000 - Authentication bit: 0
Hosting LC: 10, Down reason: None, Reason not-hosted: None
```

```
switch# show bfd neighbors details
OurAddr NeighAddr LD/RD RH/RS Holdown(mult) State Int
Vrf
10.0.2.1 10.0.2.2 1140850695/131083 Up 270(3) Up Po14.121
default
Session state is Up and not using echo function
Local Diag: 0, Demand mode: 0, Poll bit: 0, Authentication: None
MinTxInt: 50000 us, MinRxInt: 50000 us, Multiplier: 3
Received MinRxInt: 100000 us, Received Multiplier: 3
Holdown (hits): 300 ms (0), Hello (hits): 100 ms (3136283)
Rx Count: 2669290, Rx Interval (ms) min/max/avg: 12/1999/93 last: 29 ms ago
Tx Count: 3136283, Tx Interval (ms) min/max/avg: 77/77/77 last: 76 ms ago
Registered protocols: ospf
Uptime: 2 days 21 hrs 41 mins 45 secs
Last packet: Version: 1 - Diagnostic: 0
State bit: Up - Demand bit: 0
Poll bit: 0 - Final bit: 0
Multiplier: 3 - Length: 24
My Discr.: 131083 - Your Discr.: 1140850695
Min tx interval: 100000 - Min rx interval: 100000
Min Echo interval: 0 - Authentication bit: 0
Hosting LC: 8, Down reason: None, Reason not-hosted: None
```

Verifying the BFD Configuration

To display BFD configuration information, perform one of the following:

Command	Purpose
<code>show running-config bfd</code>	Displays the running BFD configuration.

Command	Purpose
<code>show startup-config bfd</code>	Displays the BFD configuration that will be applied on the next system startup.

Monitoring BFD

Use the following commands to display BFD:

Command	Purpose
<code>show bfd neighbors [application name] [details]</code>	Displays information about BFD for a supported application, such as BGP or OSPFv2.
<code>show bfd neighbors [interface int-if] [details]</code>	Displays information about BGP sessions on an interface.
<code>show bfd neighbors [dest-ip ip-address] [src-ip ip-address][details]</code>	Displays information about the specified BGP session on an interface.
<code>show bfd neighbors [vrf vrf-name] [details]</code>	Displays information about BFD for a VRF.
<code>show bfd [ipv4 ipv6] [neighbors]</code>	Displays information about IPv4 neighbors or IPv6 neighbors.

Configuration Examples for BFD

This example shows how to configure BFD for OSPFv2 on Ethernet 2/1, using the default BFD session parameters:

```
feature bfd
feature ospf
router ospf Test1
interface ethernet 2/1
ip ospf bfd
no shutdown
```

This example shows how to configure BFD for all EIGRP interfaces, using the default BFD session parameters:

```
feature bfd
feature eigrp
bfd interval 100 min_rx 100 multiplier 4
router eigrp Test2
bfd
```

This example shows how to configure BFDv6:

```
feature bfd
feature ospfv3
router ospfv3 Test1
```

```
interface Ethernet2/7
  ipv6 router ospfv3 Test1 area 0.0.0.0
  ospfv3 bfd
  no shutdown
```

Show Example for BFD

This example shows results of the **show bfd ipv6 neighbors details** command.

```
#show bfd ipv6 neighbors details
```

```
OurAddr          RH/RS          NeighAddr      Holdown(mult)  State  Int
LD/RD
Vrf
cc:10::2         1090519335/1090519260 Up             5692 (3)      Up     Po1
default
```

```
Session state is Up and using echo function with 250 ms interval
Local Diag: 0, Demand mode: 0, Poll bit: 0, Authentication: None
MinTxInt: 250000 us, MinRxInt: 2000000 us, Multiplier: 3
Received MinRxInt: 2000000 us, Received Multiplier: 3
Holdown (hits): 6000 ms (4), Hello (hits): 2000 ms (205229)
Rx Count: 227965, Rx Interval (ms) min/max/avg: 124/1520/1510 last: 307 ms ago
Tx Count: 205229, Tx Interval (ms) min/max/avg: 1677/1677/1677 last: 587 ms ago
Registered protocols:  bgp
Uptime: 3 days 23 hrs 31 mins 13 secs
Last packet: Version: 1          - Diagnostic: 0
                State bit: Up    - Demand bit: 0
                Poll bit: 0      - Final bit: 0
                Multiplier: 3    - Length: 24
                My Discr.: 1090519260 - Your Discr.: 1090519335
                Min tx interval: 250000 - Min rx interval: 2000000
                Min Echo interval: 250000 - Authentication bit: 0
Hosting LC: 1, Down reason: None, Reason not-hosted: None
```

Related Documents

Related Topic	Document Title
BFD commands	<i>Cisco Nexus 9000 Series NX-OS Unicast Routing Configuration Guide</i>

RFCs

RFC	Title
RFC 5880	<i>Bidirectional Forwarding Detection (BFD)</i>
RFC 5881	<i>BFD for IPv4 and IPv6 (Single Hop)</i>

RFC	Title
RFC 7130	<i>Bidirectional Forwarding Detection (BFD) on Link Aggregation Group (LAG) Interfaces</i>



CHAPTER 7

Configuring Port Channels

This chapter describes how to configure port channels and to apply and configure the Link Aggregation Control Protocol (LACP) for more efficient use of port channels in the Cisco NX-OS devices.

On a single switch, the port-channel compatibility parameters must be the same among all the port-channel members on the physical switch.

- [About Port Channels, on page 163](#)
- [Port Channels, on page 164](#)
- [Port-Channel Interfaces, on page 165](#)
- [Basic Settings, on page 165](#)
- [Compatibility Requirements, on page 166](#)
- [Load Balancing Using Port Channels, on page 168](#)
- [Symmetric Hashing, on page 169](#)
- [Guidelines and Limitations for ECMP, on page 169](#)
- [Resilient Hashing, on page 170](#)
- [GTP Tunnel Load Balancing, on page 170](#)
- [LACP, on page 171](#)
- [Prerequisites for Port Channeling, on page 177](#)
- [Guidelines and Limitations, on page 177](#)
- [Default Settings, on page 179](#)
- [Configuring Port Channels, on page 180](#)

About Port Channels

A port channel is an aggregation of multiple physical interfaces that creates a logical interface. You can bundle up to 32 individual active links into a port channel to provide increased bandwidth and redundancy. Port channeling also load balances traffic across these physical interfaces. The port channel stays operational as long as at least one physical interface within the port channel is operational.

You can create a Layer 2 port channel by bundling compatible Layer 2 interfaces, or you can create Layer 3 port channels by bundling compatible Layer 3 interfaces. You cannot combine Layer 2 and Layer 3 interfaces in the same port channel.

You can also change the port channel from Layer 3 to Layer 2. See the [Configuring Layer 2 Interfaces](#) chapter for information about creating Layer 2 interfaces.

A Layer 2 port channel interface and its member ports can have different STP parameters. Changing the STP parameters of the port channel does not impact the STP parameters of the member ports because a port channel interface takes precedence if the member ports are bundled.



Note After a Layer 2 port becomes part of a port channel, all switchport configurations must be done on the port channel; you can no longer apply switchport configurations to individual port-channel members. You cannot apply Layer 3 configurations to an individual port-channel member either; you must apply the configuration to the entire port channel.

You can use static port channels, with no associated aggregation protocol, for a simplified configuration.

For more flexibility, 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. You cannot configure LACP on shared interfaces.

See the LACP Overview section for information about LACP.

Port Channels

A port channel bundles physical links into a channel group to create a single logical link that provides the aggregate bandwidth of up to 32 physical links. If a member port within a port channel fails, the traffic previously carried over the failed link switches to the remaining member ports within the port channel.

However, you can enable the LACP to use port channels more flexibly. Configuring port channels with LACP and static port channels require a slightly different procedure (see the “Configuring Port Channels” section).



Note The device does not support Port Aggregation Protocol (PAgP) for port channels.

Each port can be in only one port channel. All the ports in a port channel must be compatible; they must use the same speed and duplex mode (see the “Compatibility Requirements” section). When you run static port channels with no aggregation protocol, the physical links are all in the on channel mode; you cannot change this mode without enabling LACP (see the “Port-Channel Modes” section).

You can create port channels 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, the software creates a matching port channel automatically if the port channel does not already exist. In this instance, the port channel assumes the Layer 2 or Layer 3 configuration of the first interface. You can also create the port channel first. In this instance, the Cisco NX-OS software creates an empty channel group with the same channel number as the port channel and takes the default Layer 2 or Layer 3 configuration, as well as the compatibility configuration (see the “Compatibility Requirements” section).

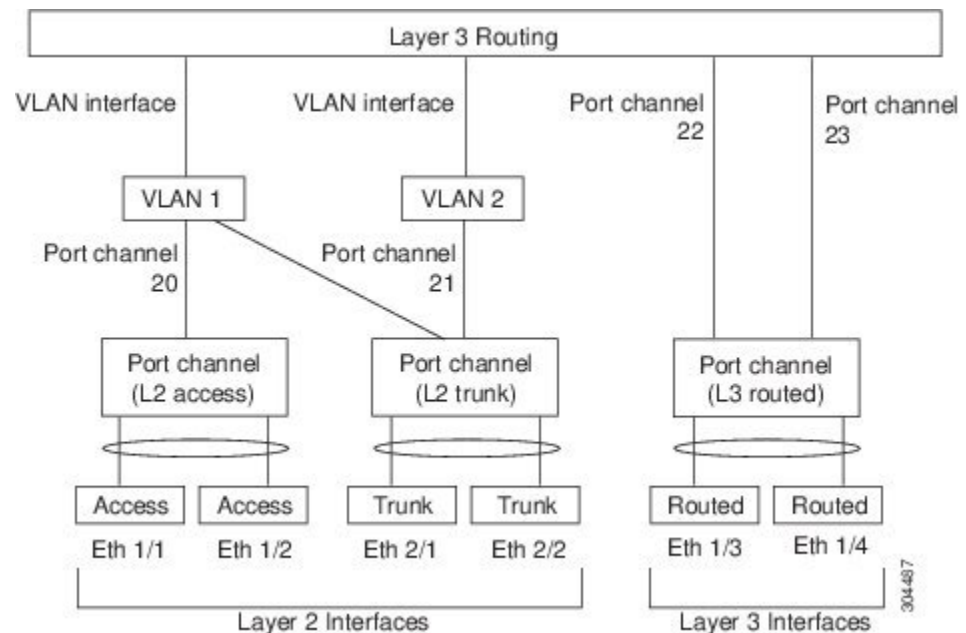


Note The port channel is operationally up when at least one of the member ports is up and that port’s status is channeling. The port channel is operationally down when all member ports are operationally down.

Port-Channel Interfaces

The following shows port-channel interfaces.

Figure 9: Port-Channel Interfaces



You can classify port-channel interfaces as Layer 2 or Layer 3 interfaces. In addition, you can configure Layer 2 port channels in either access or trunk mode. Layer 3 port-channel interfaces have routed ports as channel members.

You can configure a Layer 3 port channel with a static MAC address. If you do not configure this value, the Layer 3 port channel uses the router MAC of the first channel member to come up. See the [Cisco Nexus 9000 Series NX-OS Layer 2 Switching Configuration Guide](#) for information about configuring static MAC addresses on Layer 3 port channels.

See the "Configuring Layer 2 Interfaces" chapter for information about configuring Layer 2 ports in access or trunk mode and the "Configuring Layer 3 Interfaces" chapter for information about configuring Layer 3 interfaces and subinterfaces.

Basic Settings

You can configure the following basic settings for the port-channel interface:

- Bandwidth—Use this setting for informational purposes only; this setting is to be used by higher-level protocols.
- Delay—Use this setting for informational purposes only; this setting is to be used by higher-level protocols.
- Description
- Duplex

- IP addresses
- Maximum Transmission Unit (MTU)
- Shutdown
- Speed

Compatibility Requirements

When you add an interface to a channel group, the software checks certain interface attributes to ensure that the interface is compatible with the channel group. For example, you cannot add a Layer 3 interface to a Layer 2 channel group. The Cisco NX-OS software 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:

- Network layer
- (Link) speed capability
- Speed configuration
- Duplex capability
- Duplex configuration
- Port mode
- Access VLAN
- Trunk native VLAN
- Tagged or untagged
- Allowed VLAN list
- MTU size
- SPAN—Cannot be a SPAN source or a destination port
- Storm control
- Flow-control capability
- Flow-control configuration
- Media type, either copper or fiber

Use the **show port-channel compatibility-parameters** command to see the full list of compatibility checks that the Cisco NX-OS uses.

You can only add interfaces configured with the channel mode set to on to static port channels, and you can only add interfaces configured with the channel mode as active or passive to port channels that are running LACP. You can configure these attributes on an individual member port. If you configure a member port with an incompatible attribute, the software suspends that port in the port channel.

Alternatively, you can force ports with incompatible parameters to join the port channel if the following parameters are the same:

- (Link) speed capability
- Speed configuration
- Duplex capability
- Duplex configuration
- Flow-control capability
- Flow-control configuration

When the interface joins a port channel, some of its individual parameters are removed and replaced with the values on the port channel as follows:

- Bandwidth
- Delay
- Extended Authentication Protocol over UDP
- VRF
- IP address
- MAC address
- Spanning Tree Protocol
- NAC
- Service policy
- Access control lists (ACLs)

Many interface parameters remain unaffected when the interface joins or leaves a port channel as follows:

- Beacon
- Description
- CDP
- LACP port priority
- Debounce
- UDLD
- MDIX
- Rate mode
- Shutdown
- SNMP trap



Note When you delete the port channel, the software sets all member interfaces as if they were removed from the port channel.

See the “LACP Marker Responders” section for information about port-channel modes.

Load Balancing Using Port Channels

The Cisco NX-OS software load balances traffic across all operational interfaces in a port channel by hashing the addresses in the frame to a numerical value that selects one of the links in the channel. Port channels provide load balancing by default. Port-channel load balancing uses MAC addresses, IP addresses, or Layer 4 port numbers to select the link. Port-channel load balancing uses either source or destination addresses or ports, or both source and destination addresses or ports.

You can configure the load-balancing mode to apply to all port channels that are configured on the entire device. You can configure one load-balancing mode for the entire device. You cannot configure the load-balancing method per port channel.

You can configure the type of load-balancing algorithm used. You can choose the load-balancing algorithm that determines which member port to select for egress traffic by looking at the fields in the frame.

The default load-balancing mode for Layer 3 interfaces is the source and destination IP L4 ports, and the default load-balancing mode for non-IP traffic is the source and destination MAC address. Use the **port-channel load-balance** command to set the load-balancing method among the interfaces in the channel-group bundle. The default method for Layer 2 packets is `src-dst-mac`. The default method for Layer 3 packets is `src-dst-ip-l4port`.

You can configure the device to use one of the following methods to load balance across the port channel:

- Destination MAC address
- Source MAC address
- Source and destination MAC address
- Destination IP address
- Source IP address
- Source and destination IP address
- Source TCP/UDP port number
- Destination TCP/UDP port number
- Source and destination TCP/UDP port number
- GRE inner IP headers with source, destination and source-destination

Non-IP and Layer 3 port channels both follow the configured load-balancing method, using the source, destination, or source and destination parameters. For example, when you configure load balancing to use the source IP address, all non-IP traffic uses the source MAC address to load balance the traffic while the Layer 3 traffic load balances the traffic using the source IP address. Similarly, when you configure the destination

MAC address as the load-balancing method, all Layer 3 traffic uses the destination IP address while the non-IP traffic load balances using the destination MAC address.

The unicast and multicast traffic is load-balanced across port-channel links based on configured load-balancing algorithm displayed in **show port-channel load-balancing** command output.

The multicast traffic uses the following methods for load balancing with port channels:

- Multicast traffic with Layer 4 information—Source IP address, source port, destination IP address, destination port
- Multicast traffic without Layer 4 information—Source IP address, destination IP address
- Non-IP multicast traffic—Source MAC address, destination MAC address



Note Devices that run Cisco IOS can optimize the behavior of the member ports ASICs if a failure of a single member occurred by running the port-channel hash-distribution command. The Cisco Nexus 9000 Series device performs this optimization by default and does not require or support this command. Cisco NX-OS does support the customization of the load-balancing criteria on port channels through the port-channel load-balance command for the entire device.

Symmetric Hashing

To be able to effectively monitor traffic on a port channel, it is essential that each interface connected to a port channel receives both forward and reverse traffic flows. Normally, there is no guarantee that the forward and reverse traffic flows will use the same physical interface. However, when you enable symmetric hashing on the port channel, bidirectional traffic is forced to use the same physical interface and each physical interface in the port channel is effectively mapped to a set of flows.

When symmetric hashing is enabled, the parameters used for hashing, such as the source and destination IP address, are normalized before they are entered into the hashing algorithm. This process ensures that when the parameters are reversed (the source on the forward traffic becomes the destination on the reverse traffic), the hash output is the same. Therefore, the same interface is chosen.

Only the following load-balancing algorithms support symmetric hashing:

- src-dst ip
- src-dst ip-l4port

Guidelines and Limitations for ECMP

You might observe that load balancing with Layer 2/Layer 3 GW flows are not load balanced equally among all links when the switch comes up initially after reload. There are two CLIs to change the ECMP hash configuration in the hardware. The two CLI commands are mutually exclusive.

- Enter the **port-channel load-balance [src | src-dst | dst] mac** command for MAC-based only hash.
- For hash based on IP/Layer 4 ports, enter either the **ip load-share** or **port-channel load-balance** command.

- The **port-channel load-balance** command can overwrite the **ip load-share** command. It is better to enter the **port-channel load-balance** command which helps to set both the IP and MAC parameters.
- There are no options to force the hashing algorithm based on the IP/Layer 4 port. The default MAC configuration is always programmed as a part of the port channel configuration.
- ECMP resilient hashing is not supported for traffic flows over tunnel.

Resilient Hashing

With the exponential increase in the number of physical links used in data centers, there is also the potential for an increase in the number of failed physical links. In static hashing systems that are used for load balancing flows across members of port channels or Equal Cost Multipath (ECMP) groups, each flow is hashed to a link. If a link fails, all flows are rehashed across the remaining working links. This rehashing of flows to links results in some packets being delivered out of order even for those flows that were not hashed to the failed link.

This rehashing also occurs when a link is added to the port channel or Equal Cost Multipath (ECMP) group. All flows are rehashed across the new number of links, which results in some packets being delivered out of order.

Resilient hashing maps flows to physical ports and it is supported for both ECMP groups and port channel interfaces.

If a physical link fails, the flows originally assigned to the failed link are redistributed uniformly among the remaining working links. The existing flows through the working links are not rehashed and hence are not impacted.

Resilient hashing supports IPv4 and IPv6 unicast traffic, but it does not support IPv4 multicast traffic.

Resilient hashing is supported on all the Cisco Nexus 9000 Series platforms. (NX-OS 7.0(3)I3(1) release and later). Beginning Cisco NX-OS Release 9.3(3), resilient hashing is supported on Cisco Nexus 92160YC-X, 92304QC, 9272Q, 9232C, 9236C, 92300YC switches.

GTP Tunnel Load Balancing

GPRS Tunneling Protocol (GTP) is used mainly to deliver mobile data on wireless networks via Cisco Nexus 9000 Series switches as the core router. When two routers carrying GTP traffic are connected with link bundling, the traffic is required to be distributed evenly between all bundle members.

To achieve load balancing, Cisco Nexus 9000 Series switches use 5-tuple load balancing mechanism. The load balancing mechanism takes into account the source IP, destination IP, protocol, Layer 4 resource and destination port (if traffic is TCP or UDP) fields from the packet. In the case of GTP traffic, a limited number of unique values for these fields restrict the equal distribution of traffic load on the tunnel.

In order to avoid polarization for GTP traffic in load balancing, a tunnel endpoint identifier (TEID) in the GTP header is used instead of a UDP port number. Since the TEID is unique per tunnel, traffic can be evenly load balanced across multiple links in the bundle.

Beginning Cisco Nexus Release 7.0(3)I7(3), GTP Tunnel Load Balancing is supported on Cisco Nexus 9300-EX platform switches.

Beginning Cisco Nexus Release 7.0(3)I7(4), GTP Tunnel Load Balancing is supported on Cisco Nexus 9300-FX and 9364C platform switches.

This feature overrides the source and destination port information with the 32-bit TEID value that is present in GTPU packets.

GTP tunnel load balancing feature adds support for:

- GTP with IPv4/IPv6 transport header on physical interface
- GTP traffic over TE tunnel
- GTPU with UDP port 2152

The **ip load-sharing address source-destination gtpu** command enables the GTP tunnel load balancing.

To know the egress interface for GTP traffic after load balancing, use **show cef {ipv4 | ipv6} exact-route** command with TEID in place of L4 protocol source and destination port number. Use 16MSBist of TEID in source port and 16LSBits of TEID in destination port.

LACP

LACP allows you to configure up to 16 interfaces into a port channel.

LACP Overview

The Link Aggregation Control Protocol (LACP) for Ethernet is defined in IEEE 802.1AX and IEEE 802.3ad. This protocol controls how physical ports are bundled together to form one logical channel.




Note You must enable LACP before you can use LACP. By default, LACP is disabled. See the “Enabling LACP” section for information about enabling LACP.

The system automatically takes a checkpoint before disabling the feature, and you can roll back to this checkpoint. See the [Cisco Nexus 9000 Series NX-OS System Management Configuration Guide](#) for information about rollbacks and checkpoints.

The following figure shows how individual links can be combined into LACP port channels and channel groups as well as function as individual links.

Figure 10: Individual Links Combined into a Port Channel



With LACP, you can bundle up to 32 interfaces in a channel group.



Note When you delete the port channel, the software automatically deletes the associated channel group. All member interfaces revert to their original configuration.



Note If you downgrade a Cisco Nexus 9500 series switch that is configured to use LACP vPC convergence feature, that runs Cisco NX-OS Release 7.0(3)I7(5) to a lower release, the configuration is removed. You must configure the LACP vPC convergence feature again when you upgrade the switch.

You cannot disable LACP while any LACP configurations are present.

Port-Channel Modes

Individual interfaces in port channels are configured with channel modes. When you run static port channels with no aggregation 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 either **active** or **passive**. You can configure channel mode for individual links in the LACP channel group when you are adding the links to the 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 10: Channel Modes for Individual Links in a Port Channel

Channel Mode	Description
passive	The LACP is enabled on this port channel and the ports are in a passive negotiating state. Ports responds to LACP packets that it receives but does not initiate LACP negotiation.
active	The LACP is enabled on this port channel and the ports are in an active negotiating state. Ports initiate negotiations with other ports by sending LACP packets.
on	The LACP is disabled on this port channel and the ports are in a non-negotiating state. The on state of the port channel represents the static mode. The port will not verify or negotiate port channel memberships. If you attempt to change the channel mode to active or passive before enabling LACP, the device displays an error message. 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. The on state is the default port-channel mode

Both the passive and active modes allow LACP to negotiate between ports to determine if they can form a port channel 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.

Two devices can form an LACP port channel when their ports are in different LACP modes if the modes are compatible as in the following example:

Table 11: Channel Modes Compatibility

Device 1 > Port-1	Device 2 > Port-2	Result
Active	Active	Can form a port channel.
Active	Passive	Can form a port channel.
Passive	Passive	Cannot form a port channel because no ports can initiate negotiation.
On	Active	Cannot form a port channel because LACP is enabled only on one side.
On	Passive	Cannot form a port channel because LACP is not enabled.

LACP ID Parameters

This section describes the LACP 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 that is 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 and the duplex capability
- Configuration restrictions that you establish

LACP Marker Responders

You can dynamically redistribute the data traffic by using port channels. This redistribution might result from a removed or added link or a change in the load-balancing scheme. Traffic redistribution that occurs in the middle of a traffic flow can cause misordered frames.

LACP uses the Marker Protocol to ensure that frames are not duplicated or reordered due to this redistribution. The Marker Protocol detects when all the frames of a given traffic flow are successfully received at the remote end. LACP sends Marker PDUs on each of the port-channel links. The remote system responds to the Marker PDU once it receives all the frames received on this link prior to the Marker PDU. The remote system then sends a Marker Responder. Once the Marker Responders are received by the local system on all member links of the port channel, the local system can redistribute the frames in the traffic flow with no chance of misordering. The software supports only Marker Responders.

LACP-Enabled and Static Port Channels Differences

The following table summarizes the major differences between port channels with LACP enabled and static port channels.

Table 12: Port Channels with LACP Enabled and Static Port Channels

Configurations	Port Channels with LACP Enabled	Static Port Channels
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
Maximum number of links in channel	32	32

LACP Compatibility Enhancements

When a Cisco Nexus 9000 Series device is connected to a non-Nexus peer, its graceful failover defaults may delay the time that is taken to bring down a disabled port or cause traffic from the peer to be lost. To address these conditions, the **lACP graceful-convergence** command was added.

By default, LACP sets a port to suspended state if it does not receive an LACP PDU from the peer. **lACP suspend-individual** is a default configuration on Cisco Nexus 9000 series switches. This command puts the port in suspended state if it does not receive any LACP PDUs. In some cases, although this feature helps in preventing loops created due to misconfigurations, it can cause servers fail to boot up because they require LACP to logically bring up the port. You can put a port into an individual state by using the **no lACP suspend-individual**. Port in individual state takes attributes of the individual port based on the port configuration.

LACP port-channels exchange LACP PDUs for quick bundling of links when connecting a server and a switch. However, the links go into suspended state when the PDUs are not received.

The **delayed LACP** feature enables one port-channel member, the delayed-LACP port, to come up first as a member of a regular port-channel before LACP PDUs are received. After it is connected in LACP mode, other members, the auxiliary LACP ports, are brought up. This avoids having the links becoming suspended when PDUs are not received.

Which port in the port-channel comes up first depends on the port-priority value of the ports. A member link in a port channel with lowest priority value, will come up first as a LACP delayed port. Regardless of the operational status of the links, the configured priority of a LACP port is used to select the delayed-lACP port

This feature supports Layer 2 port channels, trunk mode spanning tree, and vPC and has the following limitations:

- Using **no lACP suspend-individual** and **lACP mode delay** on a same port channel is not recommended because it can put non-lACP delayed ports in individual state. As a best practice, you must avoid combining these two configurations.
- Not supported on Layer 3 port channels.

- Not supported on Cisco Nexus 9500 Switches and FEX HIF and FEX fabric ports.

LACP Port-Channel Minimum Links and MaxBundle

A port channel aggregates similar ports to provide increased bandwidth in a single manageable interface.

The introduction of the minimum links and maxbundle feature further refines LACP port-channel operation and provides increased bandwidth in one manageable interface.

The LACP port-channel minimum links feature does the following:

- Configures the minimum number of ports that must be linked up and bundled in the LACP port channel.
- Prevents the low-bandwidth LACP port channel from becoming active.
- Causes the LACP port channel to become inactive if there are few active members ports to supply the required minimum bandwidth.

The LACP MaxBundle defines the maximum number of bundled ports allowed in a LACP port channel.

The LACP MaxBundle feature does the following:

- Defines an upper limit on the number of bundled ports in an LACP port channel.
- Allows hot-standby ports with fewer bundled ports. (For example, in an LACP port channel with five ports, you can designate two of those ports as hot-standby ports.)



Note The minimum links and maxbundle feature works only with LACP port channels. However, the device allows you to configure this feature in non-LACP port channels, but the feature is not operational.

LACP Fast Timers

You can change the LACP timer rate to modify the duration of the LACP timeout. Use the `lacp rate` command to set the rate at which LACP control packets are sent to an LACP-supported interface. You can change the timeout rate from the default rate (30 seconds) to the fast rate (1 second). This command is supported only on LACP-enabled interfaces. To configure the LACP fast time rate, see the “Configuring the LACP Fast Timer Rate” section.

ISSU and ungraceful switchovers are not supported with LACP fast timers.

Virtualization Support

You must configure the member ports and other port channel-related configuration from the virtual device context (VDC) that contains the port channel and member ports. You can use the numbers from 1 to 4096 in each VDC to number the port channels.

All ports in one port channel must be in the same VDC. When you are using LACP, all possible 8 active ports and all possible 8 standby ports must be in the same VDC.



Note You must configure load balancing using port channels in the default VDC. See the “Load Balancing Using Port Channels” section for more information about load balancing.

High Availability

Port channels provide high availability by load balancing traffic across multiple ports. If a physical port fails, the port channel is still operational if there is an active member in the port channel. You can bundle ports from different modules and create a port channel that remains operational even if a module fails because the settings are common across the module.

Port channels support stateful and stateless restarts. A stateful restart occurs on a supervisor switchover. After the switchover, the Cisco NX-OS software applies the runtime configuration after the switchover.

The port channel goes down if the operational ports fall below the configured minimum links number.



Note See the *Cisco Nexus 9000 Series NX-OS High Availability and Redundancy Guide* for complete information about high-availability features.

Prerequisites for Port Channeling

Port channeling has the following prerequisites:

- You must be logged onto the device.
- All ports for a single port channel must be either Layer 2 or Layer 3 ports.
- All ports for a single port channel must meet the compatibility requirements. See the “Compatibility Requirements” section for more information about the compatibility requirements.
- You must configure load balancing from the default VDC.

Guidelines and Limitations

Port channeling has the following configuration guidelines and limitations:

- For scaled port-channel deployments on Cisco Nexus 9516 switch with Gen 1 line cards, you need to use the **port-channel scale-fanout** command followed by **copy run start** and **reload** commands.
- **show** commands with the **internal** keyword are not supported.
- The LACP port-channel minimum links and maxbundle feature is not supported for host interface port channels.
- You must enable LACP before you can use that feature.
- You can configure multiple port channels on a device.

- Do not put shared and dedicated ports into the same port channel. (See the “Configuring Basic Interface Parameters” chapter for information about shared and dedicated ports.)
- For Layer 2 port channels, ports with different STP port path costs can form a port channel if they are compatibly configured with each other. See the “Compatibility Requirements” section for more information about the compatibility requirements.
- In STP, the port-channel cost is based on the aggregated bandwidth of the port members.
- After you configure a port channel, the configuration that you apply to the port channel interface affects the port channel member ports. The configuration that you apply to the member ports affects only the member port where you apply the configuration.
- LACP does not support half-duplex mode. Half-duplex ports in LACP port channels are put in the suspended state.
- Do not configure ports that belong to a port channel group as private VLAN ports. While a port is part of the private VLAN configuration, the port channel configuration becomes inactive.
- Channel member ports cannot be a source or destination SPAN port.
- Port-channels are not supported on generation 1 100G line cards (N9K-X9408PC-CFP2) or generic expansion modules (N9K-M4PC-CFP2).
- Port-channels are supported on devices with generation 2 (and later) 100G interfaces. (7.0(3)I3(1) and later)
- The port channel might be affected by the limitations of the Application Leaf Engine (ALE) uplink ports on Cisco Nexus 9300 and 9500 Series devices: [Limitations for ALE Uplink Ports](#)
- Resilient hashing for port channels is not supported on Cisco Nexus 9200, Cisco Nexus 9300-EX, and Cisco Nexus 9500 switches with 9700-EX line cards.
- In Cisco NX-OS Release 7.0(3)I4(1), resilient hashing (port-channel load-balancing resiliency) and VXLAN configurations are not compatible with VTEPs using ALE uplink ports.



Note Resilient hashing is disabled by default.

- The maximum number of subinterfaces for a port is 511. The maximum number of subinterfaces for a satellite/FEX port is 63.
- On a Cisco Nexus 92300YC switch, the first 24 ports that are part of the same quadrant. All the ports in the same quadrant must have same speed. Having different speed on ports in a quadrant is not supported. Following are the first 24 ports on the Cisco Nexus 92300YC switch that share same quadrant:
 - 1,4,7,10
 - 2,5,8,11
 - 3,6,9,12
 - 13,16,19,22
 - 14,17,20,23
 - 15,18,21,24

- On a Cisco Nexus 9500 switch with a X96136YC-R line card, the ports 17–48 are part of the same quadrant. Ports in the same quadrant must have same speed (1/10G or 25G) on all ports. Having different speed on ports in a quadrant is not supported. If you set different speed in any of the ports in a quadrant, the ports go into error disable state. Interfaces in same quadrant are:
 - 17–20
 - 21–24
 - 25–28
 - 29–32
 - 33–36
 - 37–40
 - 41–44
 - 45–48
- Beginning Cisco NX-OS Release 7.0(3)I7(8), you can use **interface port-channel all** command to select all the available parent port channel interfaces in the system and then apply any subcommand to all the member ports.



Note The subcommands impacts the FEX and VPC ports as well. For example, a **shut** subcommand may shutdown any connected FEX or vPC ports.

- If you have layer 2 and layer 3 port channels in a switch some commands may not be applicable to both the layer 2 and layer 3 port channels. In those instances, the configuration may not be applied properly.
- Use the **interface port-channel all** command only on similar type of port channels in a switch. This command works on parent interfaces. However the command may not work if port channel sub interfaces are present.

Default Settings

The following table lists the default settings for port-channel parameters.

Table 13: Default Port-Channel Parameters

Parameters	Default
Port channel	Admin up
Load balancing method for Layer 3 interfaces	Source and destination IP address
Load balancing method for Layer 2 interfaces	Source and destination MAC address
Load balancing per module	Disabled
LACP	Disabled

Parameters	Default
Channel mode	on
LACP system priority	32768
LACP port priority	32768
Minimum links for LACP	1
Maxbundle	32
Minimum links for FEX fabric port channel	1

Configuring Port Channels



Note See the "Configuring Basic Interface Parameters" chapter for information about configuring the maximum transmission unit (MTU) for the port-channel interface. See the "Configuring Layer 3 Interfaces" chapter for information about configuring IPv4 and IPv6 addresses on the port-channel interface.



Note If you are familiar with the Cisco IOS CLI, be aware that the Cisco NX-OS commands for this feature might differ from the Cisco IOS commands that you would use.

Creating a Port Channel

You can create a port channel before you create a channel group. The software automatically creates the associated channel group.



Note When the port channel is created before the channel group, the port channel should be configured with all of the interface attributes that the member interfaces are configured with. Use the **switchport mode trunk** *{allowed vlan vlan-id | native vlan-id}* command to configure the members.

This is required only when the channel group members are Layer 2 ports (switchport) and trunks (switchport mode trunk).



Note Use the **no interface port-channel** command to remove the port channel and delete the associated channel group.

Command	Purpose
no interface port-channel <i>channel-number</i> Example: switch(config)# no interface port-channel 1	Removes the port channel and deletes the associated channel group.

Before you begin

Enable LACP if you want LACP-based port channels.

Procedure

	Command or Action	Purpose
Step 1	configure terminal Example: switch# configure terminal switch(config)#	Enters global configuration mode.
Step 2	interface port-channel <i>channel-number</i> Example: switch(config)# interface port-channel 1 switch(config-if)	Specifies the port-channel interface to configure, and enters the interface configuration mode. The range is from 1 to 4096. The Cisco NX-OS software automatically creates the channel group if it does not already exist.
Step 3	show port-channel summary Example: switch(config-router)# show port-channel summary	(Optional) Displays information about the port channel.
Step 4	no shutdown Example: switch# configure terminal switch(config)# int e3/1 switch(config-if)# no shutdown	(Optional) Clears the errors on the interfaces and VLANs where policies correspond with hardware policies. This command allows policy programming to continue and the port to come up. If policies do not correspond, the errors are placed in an error-disabled policy state.
Step 5	copy running-config startup-config Example: switch(config)# copy running-config startup-config	(Optional) Copies the running configuration to the startup configuration.

Example

This example shows how to create a port channel:

```
switch# configure terminal
switch (config)# interface port-channel 1
```

See the “Compatibility Requirements” section for details on how the interface configuration changes when you delete the port channel.

Adding a Layer 2 Port to a Port Channel

You can add a Layer 2 port to a new channel group or to a channel group that already contains Layer 2 ports. The software creates the port channel associated with this channel group if the port channel does not already exist.



Note Use the **no channel-group** command to remove the port from the channel group.

Command	Purpose
no channel-group Example: switch(config)# no channel-group	Removes the port from the channel group.

Before you begin

Enable LACP if you want LACP-based port channels.

All Layer 2 member ports must run in full-duplex mode and at the same speed

Procedure

	Command or Action	Purpose
Step 1	configure terminal Example: switch# configure terminal switch(config)#	Enters global configuration mode.
Step 2	interface type slot/port Example: switch(config)# interface ethernet 1/4 switch(config-if)#	Specifies the interface that you want to add to a channel group, and enters the interface configuration mode.
Step 3	switchport Example: switch(config)# switchport	Configures the interface as a Layer 2 access port.

	Command or Action	Purpose
Step 4	switchport mode trunk Example: <pre>switch(config)# switchport mode trunk</pre>	(Optional) Configures the interface as a Layer 2 trunk port.
Step 5	switchport trunk {allowed vlan <i>vlan-id</i> native <i>vlan-id</i>} Example: <pre>switch(config)# switchport trunk native 3 switch(config-if)#</pre>	(Optional) Configures necessary parameters for a Layer 2 trunk port.
Step 6	channel-group <i>channel-number</i> [force] [mode {on active passive}] Example: <ul style="list-style-type: none"> • <pre>switch(config-if)# channel-group 5</pre> • <pre>switch(config-if)# channel-group 5 force</pre> 	<p>Configures the port in a channel group and sets the mode. The channel-number range is from 1 to 4096. This command creates the port channel associated with this channel group if the port channel does not already exist. All static port-channel interfaces are set to mode on. You must set all LACP-enabled port-channel interfaces to active or passive. The default mode is on.</p> <p>(Optional) Forces an interface with some incompatible configurations to join the channel. The forced interface must have the same speed, duplex, and flow control settings as the channel group.</p> <p>Note The force option fails if the port has a QoS policy mismatch with the other members of the port channel.</p>
Step 7	show interface <i>type slot/port</i> Example: <pre>switch# show interface port channel 5</pre>	(Optional) Displays interface information.
Step 8	no shutdown Example: <pre>switch# configure terminal switch(config)# int e3/1 switch(config-if)# no shutdown</pre>	(Optional) Clears the errors on the interfaces and VLANs where policies correspond with hardware policies. This command allows policy programming to continue and the port to come up. If policies do not correspond, the errors are placed in an error-disabled policy state.
Step 9	copy running-config startup-config Example: <pre>switch(config)# copy running-config startup-config</pre>	(Optional) Copies the running configuration to the startup configuration.

Example

This example shows how to add a Layer 2 Ethernet interface 1/4 to channel group 5:

```
switch# configure terminal
switch (config)# interface ethernet 1/4
switch(config-if)# switchport
switch(config-if)# channel-group 5
```

Adding a Layer 3 Port to a Port Channel

You can add a Layer 3 port to a new channel group or to a channel group that is already configured with Layer 3 ports. The software creates the port channel associated with this channel group if the port channel does not already exist.

If the Layer 3 port that you are adding has a configured IP address, the system removes that IP address before adding the port to the port channel. After you create a Layer 3 port channel, you can assign an IP address to the port-channel interface.



Note Use the **no channel-group** command to remove the port from the channel group. The port reverts to its original configuration. You must reconfigure the IP addresses for this port.

Command	Purpose
no channel-group Example: <pre>switch(config)# no channel-group</pre>	Removes the port from the channel group.

Before you begin

Enable LACP if you want LACP-based port channels.

Remove any IP addresses configured on the Layer 3 interface.

Procedure

	Command or Action	Purpose
Step 1	configure terminal Example: <pre>switch# configure terminal switch(config)#</pre>	Enters global configuration mode.
Step 2	interface type slot/port Example: <pre>switch(config)# interface ethernet 1/4 switch(config-if)#</pre>	Specifies the interface that you want to add to a channel group, and enters the interface configuration mode.

	Command or Action	Purpose
Step 3	no switchport Example: switch(config-if)# no switchport	Configures the interface as a Layer 3 port.
Step 4	channel-group <i>channel-number</i> [force] [mode { on active passive }] Example: <ul style="list-style-type: none"> • switch(config-if)# channel-group 5 • switch(config-if)# channel-group 5 force 	Configures the port in a channel group and sets the mode. The channel-number range is from 1 to 4096. The Cisco NX-OS software creates the port channel associated with this channel group if the port channel does not already exist. (Optional) Forces an interface with some incompatible configurations to join the channel. The forced interface must have the same speed, duplex, and flow control settings as the channel group.
Step 5	show interface <i>type slot/port</i> Example: switch# show interface ethernet 1/4	(Optional) Displays interface information.
Step 6	no shutdown Example: switch# configure terminal switch(config)# int e3/1 switch(config-if)# no shutdown	(Optional) Clears the errors on the interfaces and VLANs where policies correspond with hardware policies. This command allows policy programming to continue and the port to come up. If policies do not correspond, the errors are placed in an error-disabled policy state.
Step 7	copy running-config startup-config Example: switch(config)# copy running-config startup-config	(Optional) Copies the running configuration to the startup configuration.

Example

This example shows how to add a Layer 3 Ethernet interface 1/5 to channel group 6 in on mode:

```
switch# configure terminal
switch (config)# interface ethernet 1/5
switch(config-if)# switchport
switch(config-if)# channel-group 6
```

This example shows how to create a Layer 3 port-channel interface and assign the IP address:

```
switch# configure terminal
switch (config)# interface port-channel 4
switch(config-if)# ip address 192.0.2.1/8
```

Applying Port Channel Configuration to all Configured Port Channels

You can select all the available port channel's parent interfaces that are already configured in the system and then apply that configuration to the member ports of all the configured port channels.

Before you begin

Procedure

	Command or Action	Purpose
Step 1	configure terminal Example: <pre>switch# configure terminal switch(config)#</pre>	Enters global configuration mode.
Step 2	interface port-channel all Example: <pre>switch(config)# interface port-channel all switch(config-if-range)#</pre>	Changes the config mode to IF range.
Step 3	where Example: <pre>switch(config-if-range)# where switch(config-if-range)#</pre>	Apply the configuration specified to all port channels in the running-configuration.
Step 4	shutdown Example: <pre>switch(config-if-range)# shutdown switch(config-if-range)#</pre>	Shuts down the local traffic on all those interfaces selected by the command interface port-channel all . To return the interface to its default operational state, use the no form of this command.
Step 5	show running-config interface port-channel all Example: <pre>switch(config-if-range)# show running-config interface port-channel all switch(config-if-range)#</pre>	Displays the running configuration of the port channels for which the new configuration is applied.

Example

This example shows how to apply a port channel configuration to all member port channels:

```
switch(config)# interface port-channel all
switch(config-if-range)# where
conf; interface port-channell1-2 admin@switch%default
switch(config-if-range)#
switch(config-if-range)# shutdown
```

```
switch(config-if-range) #
switch(config-if-range) # sh run int po all
```

See the “Compatibility Requirements” section for details on how the interface configuration changes when you delete the port channel.

Configuring the Bandwidth and Delay for Informational Purposes

The bandwidth of the port channel is determined by the number of total active links in the channel.

You configure the bandwidth and delay on port-channel interfaces for informational purposes.

Procedure

	Command or Action	Purpose
Step 1	configure terminal Example: switch# configure terminal switch(config) #	Enters global configuration mode.
Step 2	interface port-channel <i>channel-number</i> Example: switch(config) # interface port-channel 2 switch(config-if) #	Specifies the port-channel interface that you want to configure, and enters the interface mode.
Step 3	bandwidth <i>value</i> Example: switch(config-if) # bandwidth 60000000 switch(config-if) #	Specifies the bandwidth, which is used for informational purposes. The range is from 1 to 3,200,000,000 kbs. The default value depends on the total active interfaces in the channel group.
Step 4	delay <i>value</i> Example: switch(config-if) # delay 10000 switch(config-if) #	Specifies the throughput delay, which is used for informational purposes. The range is from 1 to 16,777,215 tens of microseconds. The default value is 10 microseconds.
Step 5	exit Example: switch(config-if) # exit switch(config) #	Exits the interface mode and returns to the configuration mode.
Step 6	show interface port-channel <i>channel-number</i> Example: switch# show interface port-channel 2	(Optional) Displays interface information for the specified port channel.
Step 7	copy running-config startup-config Example:	(Optional) Copies the running configuration to the startup configuration.

	Command or Action	Purpose
	<code>switch(config)# copy running-config startup-config</code>	

Example

This example shows how to configure the informational parameters of the bandwidth and delay for port channel 5:

```
switch# configure terminal
switch (config)# interface port-channel 5
switch(config-if)# bandwidth 60000000
switch(config-if)# delay 10000
switch(config-if)#
```

Shutting Down and Restarting the Port-Channel Interface

You can shut down and restart the port-channel interface. When you shut down a port-channel interface, no traffic passes and the interface is administratively down.

Procedure

	Command or Action	Purpose
Step 1	configure terminal Example: <pre>switch# configure terminal switch(config)#</pre>	Enters global configuration mode.
Step 2	interface port-channel <i>channel-number</i> Example: <pre>switch(config)# interface port-channel 2 switch(config-if)#</pre>	Specifies the port-channel interface that you want to configure, and enters the interface mode.
Step 3	shutdown Example: <pre>switch(config-if)# shutdown switch(config-if)#</pre>	<p>Shuts down the interface. No traffic passes and the interface displays as administratively down. The default is no shutdown.</p> <p>Note Use the no shutdown command to open the interface.</p> <p>The interface displays as administratively up. If there are no operational problems, traffic passes. The default is no shutdown.</p>
Step 4	exit Example:	Exits the interface mode and returns to the configuration mode.

	Command or Action	Purpose
	<pre>switch(config-if)# exit switch(config)#</pre>	
Step 5	<p>show interface port-channel <i>channel-number</i></p> <p>Example:</p> <pre>switch(config-router)# show interface port-channel 2</pre>	(Optional) Displays interface information for the specified port channel.
Step 6	<p>no shutdown</p> <p>Example:</p> <pre>switch# configure terminal switch(config)# int e3/1 switch(config-if)# no shutdown</pre>	(Optional) Clears the errors on the interfaces and VLANs where policies correspond with hardware policies. This command allows policy programming to continue and the port to come up. If policies do not correspond, the errors are placed in an error-disabled policy state.
Step 7	<p>copy running-config startup-config</p> <p>Example:</p> <pre>switch(config)# copy running-config startup-config</pre>	(Optional) Copies the running configuration to the startup configuration.

Example

This example shows how to bring up the interface for port channel 2:

```
switch# configure terminal
switch (config)# interface port-channel 2
switch(config-if)# no shutdown
```

Configuring a Port-Channel Description

You can configure a description for a port channel.

Procedure

	Command or Action	Purpose
Step 1	<p>configure terminal</p> <p>Example:</p> <pre>switch# configure terminal switch(config)#</pre>	Enters global configuration mode.
Step 2	<p>interface port-channel <i>channel-number</i></p> <p>Example:</p> <pre>switch(config)# interface port-channel 2 switch(config-if)#</pre>	Specifies the port-channel interface that you want to configure, and enters the interface mode.

	Command or Action	Purpose
Step 3	description Example: <pre>switch(config-if)# description engineering switch(config-if)#</pre>	Allows you to add a description to the port-channel interface. You can use up to 80 characters in the description. By default, the description does not display; you must configure this parameter before the description displays in the output.
Step 4	exit Example: <pre>switch(config-if)# exit switch(config)#</pre>	Exits the interface mode and returns to the configuration mode.
Step 5	show interface port-channel <i>channel-number</i> Example: <pre>switch# show interface port-channel 2</pre>	(Optional) Displays interface information for the specified port channel.
Step 6	copy running-config startup-config Example: <pre>switch(config)# copy running-config startup-config</pre>	(Optional) Copies the running configuration to the startup configuration.

Example

This example shows how to add a description to port channel 2:

```
switch# configure terminal
switch (config)# interface port-channel 2
switch(config-if)# description engineering
```

Configuring the Speed and Duplex Settings for a Port-Channel Interface

You can configure the speed and duplex settings for a port-channel interface.

Procedure

	Command or Action	Purpose
Step 1	configure terminal Example: <pre>switch# configure terminal switch(config)#</pre>	Enters global configuration mode.
Step 2	interface port-channel <i>channel-number</i> Example: <pre>switch(config)# interface port-channel 2 switch(config-if)#</pre>	Specifies the port-channel interface that you want to configure, and enters the interface mode.

	Command or Action	Purpose
Step 3	speed {10 100 1000 auto} Example: switch(config-if)# speed auto switch(config-if)#	Sets the speed for the port-channel interface. The default is auto for autonegotiation.
Step 4	duplex {auto full half} Example: switch(config-if)# speed auto switch(config-if)#	Sets the duplex for the port-channel interface. The default is auto for autonegotiation.
Step 5	exit Example: switch(config-if)# exit switch(config)#	Exits the interface mode and returns to the configuration mode.
Step 6	show interface port-channel <i>channel-number</i> Example: switch# show interface port-channel 2	(Optional) Displays interface information for the specified port channel.
Step 7	copy running-config startup-config Example: switch(config)# copy running-config startup-config	(Optional) Copies the running configuration to the startup configuration.

Example

This example shows how to set port channel 2 to 100 Mb/s:

```
switch# configure terminal
switch (config)# interface port-channel 2
switch(config-if)# speed 100
```

Configuring Load Balancing Using Port Channels

You can configure the load-balancing algorithm for port channels that applies to the entire device.



Note Use the **no port-channel load-balance** command to restore the default load-balancing algorithm of source-dest-mac for non-IP traffic and source-dest-ip for IP traffic.

Command	Purpose
no port-channel load-balance Example: switch(config)# no port-channel load-balance	Restores the default load-balancing algorithm.

Before you begin

Enable LACP if you want LACP-based port channels.

Procedure

	Command or Action	Purpose
Step 1	configure terminal Example: <pre>switch# configure terminal switch(config)#</pre>	Enters global configuration mode.
Step 2	port-channel load-balance <i>method</i> { dst ip dst ip-gre dst ip-l4port dst ip-l4port-vlan dst ip-vlan dst l4port dst mac src ip src ip-gre src ip-l4port src ip-l4port-vlan src ip-vlan src l4port src mac src-dst ip src-dst ip-gre src-dst ip-l4port [symmetric] src-dst ip-l4port-vlan src-dst ip-vlan src-dst l4port src-dst mac } [<i>fex {fex-range all}</i>] [<i>rotate rotate</i>] Example: <ul style="list-style-type: none"> • <pre>switch(config)# port-channel load-balance src-dst mac switch(config)#</pre> • <pre>switch(config)# no port-channel load-balance src-dst mac switch(config)#</pre> 	Specifies the load-balancing algorithm for the device. The range depends on the device. The default for Layer 3 is src-dst ip-l4port for both IPv4 and IPv6, and the default for non-IP is src-dst mac . Note Only the following load-balancing algorithms support symmetric hashing: <ul style="list-style-type: none"> • src-dst ip • src-dst ip-l4port
Step 3	show port-channel load-balance Example: <pre>switch(config-router)# show port-channel load-balance</pre>	(Optional) Displays the port-channel load-balancing algorithm.
Step 4	copy running-config startup-config Example: <pre>switch(config)# copy running-config startup-config</pre>	(Optional) Copies the running configuration to the startup configuration.

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 group the links into a port channel. The port channel is then added to the spanning tree as a single bridge port.

To configure LACP, you must do the following:

- Enable LACP globally by using the **feature lacp** command.
- You can use different modes for different interfaces within the same LACP-enabled port channel. You can change the mode between **active** and **passive** for an interface only if it is the only interface that is designated to the specified channel group.

Procedure

	Command or Action	Purpose
Step 1	configure terminal Example: switch# configure terminal switch(config)#	Enters global configuration mode.
Step 2	feature lacp Example: switch(config)# feature lacp	Enables LACP on the device.
Step 3	copy running-config startup-config Example: switch(config)# copy running-config startup-config	(Optional) Copies the running configuration to the startup configuration.

Example

This example shows how to enable LACP:

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

Configuring LACP Port-Channel Port Modes

After you enable LACP, you can configure the channel mode for each individual link in the LACP port channel as **active** or **passive**. This channel configuration mode allows the link to operate with LACP.

When you configure port channels with no associated aggregation protocol, all interfaces on both sides of the link remain in the **on** channel mode.

Procedure

	Command or Action	Purpose
Step 1	configure terminal Example: switch# configure terminal switch(config)#	Enters global configuration mode.

	Command or Action	Purpose
Step 2	interface <i>type slot/port</i> Example: <pre>switch(config)# interface ethernet 1/4 switch(config-if)#</pre>	Specifies the interface that you want to add to a channel group, and enters the interface configuration mode.
Step 3	channel-group <i>number mode {active on passive}</i> Example: <pre>switch(config-if)# channel-group 5 mode active</pre>	<p>Specifies the port mode for the link in a port channel. After LACP is enabled, you configure each link or the entire channel as active or passive.</p> <p>When you run port channels with no associated aggregation protocol, the port-channel mode is always on.</p> <p>The default port-channel mode is on.</p>
Step 4	show port-channel summary Example: <pre>switch(config-if)# show port-channel summary</pre>	(Optional) Displays summary information about the port channels.
Step 5	copy running-config startup-config Example: <pre>switch(config)# copy running-config startup-config</pre>	(Optional) Copies the running configuration to the startup configuration.

Example

This example shows how to set the LACP-enabled interface to the 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 LACP Port-Channel Minimum Links

You can configure the LACP minimum links feature. Although minimum links and maxbundles work only in LACP, you can enter the CLI commands for these features for non-LACP port channels, but these commands are nonoperational.



Note Use the **no lacp min-links** command to restore the default port-channel minimum links configuration.

Command	Purpose
no lacp min-links Example: switch(config)# no lacp min-links	Restores the default port-channel minimum links configuration.

Before you begin

Ensure that you are in the correct port-channel interface.

Procedure

	Command or Action	Purpose
Step 1	configure terminal Example: switch# configure terminal switch(config)#	Enters global configuration mode.
Step 2	interface port-channel <i>number</i> Example: switch(config)# interface port-channel 3 switch(config-if)#	Specifies the interface to configure, and enters the interface configuration mode.
Step 3	lacp min-links <i>number</i> Example: switch(config-if)# lacp min-links 3	Specifies the port-channel interface to configure the number of minimum links. The range is from 1 to 16.
Step 4	show running-config interface port-channel <i>number</i> Example: switch(config-if)# show running-config interface port-channel 3	(Optional) Displays the port-channel minimum links configuration.

Example

This example shows how to configure the minimum number of port-channel member interfaces to be up/active for the port-channel to be up/active:

```
switch# configure terminal
switch(config)# interface port-channel 3
switch(config-if)# lacp min-links 3
```

Configuring the LACP Port-Channel MaxBundle

You can configure the LACP maxbundle feature. Although minimum links and maxbundles work only in LACP, you can enter the CLI commands for these features for non-LACP port channels, but these commands are nonoperational.



Note Use the `no lacp max-bundle` command to restore the default port-channel max-bundle configuration.

Command	Purpose
no lacp max-bundle Example: <code>switch(config)# no lacp max-bundle</code>	Restores the default port-channel max-bundle configuration.

Before you begin

Ensure that you are in the correct port-channel interface.

Procedure

	Command or Action	Purpose
Step 1	configure terminal Example: <code>switch# configure terminal</code> <code>switch(config)#</code>	Enters global configuration mode.
Step 2	interface port-channel <i>number</i> Example: <code>switch(config)# interface port-channel</code> <code>3</code> <code>switch(config-if)#</code>	Specifies the interface to configure, and enters the interface configuration mode.
Step 3	lacp max-bundle <i>number</i> Example: <code>switch(config-if)# lacp max-bundle</code>	Specifies the port-channel interface to configure max-bundle. The default value for the port-channel max-bundle is 16. The allowed range is from 1 to 32. Note Even if the default value is 16, the number of active members in a port channel is the minimum of the <code>pc_max_links_config</code> and <code>pc_max_active_members</code> that is allowed in the port channel.

	Command or Action	Purpose
Step 4	show running-config interface port-channel <i>number</i> Example: <pre>switch(config-if)# show running-config interface port-channel 3</pre>	(Optional) Displays the port-channel max-bundle configuration.

Example

This example shows how to configure the port channel interface max-bundle:

```
switch# configure terminal
switch(config)# interface port-channel 3
switch(config-if)# lacp max-bundle 3
```

Configuring the LACP Fast Timer Rate

You can change the LACP timer rate to modify the duration of the LACP timeout. Use the **lacp rate** command to set the rate at which LACP control packets are sent to an LACP-supported interface. You can change the timeout rate from the default rate (30 seconds) to the fast rate (1 second). This command is supported only on LACP-enabled interfaces.



Note We do not recommend changing the LACP timer rate. HA and SSO are not supported when the LACP fast rate timer is configured.



Note Configuring **lacp rate fast** is not recommended on the vPC Peer-Links. When **lacp rate fast** is configured on the vPC Peer-Link member interfaces, an alert is displayed in the syslog messages only when the LACP logging level is set to 5.

Before you begin

Ensure that you have enabled the LACP feature.

Procedure

	Command or Action	Purpose
Step 1	configure terminal Example: <pre>switch# configure terminal switch(config)#</pre>	Enters global configuration mode.
Step 2	interface <i>type slot/port</i> Example:	Specifies the interface to configure and enters the interface configuration mode.

	Command or Action	Purpose
	switch(config)# interface ethernet 1/4 switch(config-if)#	
Step 3	lacp rate fast Example: switch(config-if)# lacp rate fast	Configures the fast rate (one second) at which LACP control packets are sent to an LACP-supported interface. To reset the timeout rate to its default, use the no form of the command.

Example

This example shows how to configure the LACP fast rate on Ethernet interface 1/4:

```
switch# configure terminal
switch (config)# interface ethernet 1/4
switch(config-if)# lacp rate fast
```

This example shows how to restore the LACP default rate (30 seconds) on Ethernet interface 1/4.

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

Configuring the LACP System Priority

The LACP system ID is the combination of the LACP system priority value and the MAC address.

Before you begin

Enable LACP.

Procedure

	Command or Action	Purpose
Step 1	configure terminal Example: switch# configure terminal switch(config)#	Enters global configuration mode.
Step 2	lacp system-priority <i>priority</i> Example: switch(config)# lacp system-priority 40000	Configures the system priority for use with LACP. Valid values are from 1 through 65535, and higher numbers have a lower priority. The default value is 32768. Note Each VDC has a different LACP system ID because the software adds the MAC address to this configured value.

	Command or Action	Purpose
Step 3	show lacp system-identifier Example: <pre>switch(config-if)# show lacp system-identifier</pre>	(Optional) Displays the LACP system identifier.
Step 4	copy running-config startup-config Example: <pre>switch(config)# copy running-config startup-config</pre>	(Optional) Copies the running configuration to the startup configuration.

Example

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

When you enable LACP, you can configure each link in the LACP port channel for the port priority.

Before you begin

Enable LACP.

Procedure

	Command or Action	Purpose
Step 1	configure terminal Example: <pre>switch# configure terminal switch(config)#</pre>	Enters global configuration mode.
Step 2	interface <i>type slot/port</i> Example: <pre>switch(config)# interface ethernet 1/4 switch(config-if)#</pre>	Specifies the interface that you want to add to a channel group, and enters the interface configuration mode.
Step 3	lacp port-priority <i>priority</i> Example: <pre>switch(config-if)# lacp port-priority 40000</pre>	Configures the port priority for use with LACP. Valid values are from 1 through 65535, and higher numbers have a lower priority. The default value is 32768.
Step 4	copy running-config startup-config Example:	(Optional) Copies the running configuration to the startup configuration.

	Command or Action	Purpose
	<code>switch(config-if) # copy running-config startup-config</code>	

Example

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
```

Configuring LACP System MAC and Role

You can configure the MAC address used by the LACP for protocol exchanges and the optional role. By default, the LACP uses the VDC MAC address. By default, the role is primary.

Use the **no lacp system-mac** command to make LACP use the default (VDC) MAC address and default role.

This procedure is supported on the Cisco Nexus 9336C-FX2, 93300YC-FX2, and 93240YC-FX2-Z switches.

Before you begin

LACP must be enabled.

Procedure

	Command or Action	Purpose
Step 1	configure terminal Example: <code>switch# configure terminal</code>	Enter global configuration mode.
Step 2	lacp system-mac mac-address role role-value Example: <code>switch(config)# lacp system-mac 000a.000b.000c role primary</code> <code>switch(config)# lacp system-mac 000a.000b.000c role secondary</code>	Specifies the MAC address to use in the LACP protocol exchanges. The role is optional. Primary is the default.
Step 3	(Optional) show lacp system-identifier Example: <code>switch(config)# show lacp system-identifier</code>	Displays the configured MAC address.
Step 4	copy running-config startup-config Example: <code>switch(config)# copy running-config startup-config</code>	Copies the running configuration to the startup configuration.

Example

The following example shows how to configure the role of a switch as primary.

```
Switch1# sh lacp system-identifier
32768,0-b-0-b-0-b
Switch1# sh run | grep lacp
feature lacp
lacp system-mac 000b.000b.000b role primary
```

The following example shows how to configure the role of a switch as secondary.

```
Switch2# sh lacp system-identifier
32768,0-b-0-b-0-b
Switch2# sh run | grep lacp
feature lacp
lacp system-mac 000b.000b.000b role secondary
```

Disabling LACP Graceful Convergence

By default, LACP graceful convergence is enabled. In situations where you need to support LACP interoperability with devices where the graceful failover defaults may delay the time taken for a disabled port to be brought down or cause traffic from the peer to be lost, you can disable convergence. If the downstream access switch is not a Cisco Nexus device, disable the LACP graceful convergence option.



Note The port channel has to be in the administratively down state before the command can be run.

Before you begin

Enable LACP.

Procedure

	Command or Action	Purpose
Step 1	configure terminal Example: switch# configure terminal switch(config)#	Enters global configuration mode.
Step 2	interface port-channel <i>number</i> Example: switch(config)# interface port-channel 1 switch(config-if)#	Specifies the port channel interface to configure and enters the interface configuration mode.
Step 3	shutdown Example: switch(config-if) shutdown	Administratively shuts down the port channel.

	Command or Action	Purpose
Step 4	no lacp graceful-convergence Example: <pre>switch(config-if)# no lacp graceful-convergence</pre>	Disables LACP graceful convergence on the port channel.
Step 5	no shutdown Example: <pre>switch(config-if) no shutdown</pre>	Brings the port channel administratively up.
Step 6	copy running-config startup-config Example: <pre>switch(config)# copy running-config startup-config</pre>	(Optional) Copies the running configuration to the startup configuration.

Example

This example shows how to disable LACP graceful convergence on a port channel:

```
switch# configure terminal
switch (config)# interface port-channel 1
switch(config-if)# shutdown
switch(config-if)# no lacp graceful-convergence
switch(config-if)# no shutdown
```

Reenabling LACP Graceful Convergence

If the default LACP graceful convergence is once again required, you can reenabling convergence.

Procedure

	Command or Action	Purpose
Step 1	configure terminal Example: <pre>switch# configure terminal switch(config)#</pre>	Enters global configuration mode.
Step 2	interface port-channel <i>number</i> Example: <pre>switch(config)# interface port-channel 1 switch(config-if)#</pre>	Specifies the port channel interface to configure and enters the interface configuration mode.
Step 3	shutdown Example: <pre>switch(config-if) shutdown</pre>	Administratively shuts down the port channel.

	Command or Action	Purpose
Step 4	lACP graceful-convergence Example: switch(config-if)# lACP graceful-convergence	Enables LACP graceful convergence on the port channel.
Step 5	no shutdown Example: switch(config-if) no shutdown	Brings the port channel administratively up.
Step 6	copy running-config startup-config Example: switch(config)# copy running-config startup-config	(Optional) Copies the running configuration to the startup configuration.

Example

This example shows how to enable LACP graceful convergence on a port channel:

```
switch# configure terminal
switch (config)# interface port-channel 1
switch(config-if)# shutdown
switch(config-if)# lACP graceful-convergence
switch(config-if)# no shutdown
```

Disabling LACP Suspend Individual

LACP sets a port to the suspended state if it does not receive an LACP PDU from the peer. This process can cause some servers to fail to boot up as they require LACP to logically bring up the port.



Note You should only enter the **lACP suspend-individual** command on edge ports. The port channel has to be in the administratively down state before you can use this command.

Before you begin

Enable LACP.

Procedure

	Command or Action	Purpose
Step 1	configure terminal Example: switch# configure terminal switch(config)#	Enters global configuration mode.

	Command or Action	Purpose
Step 2	interface port-channel <i>number</i> Example: <pre>switch(config)# interface port-channel 1 switch(config-if)#</pre>	Specifies the port channel interface to configure and enters the interface configuration mode.
Step 3	shutdown Example: <pre>switch(config-if) shutdown</pre>	Administratively shuts down the port channel.
Step 4	no lacp suspend-individual Example: <pre>switch(config-if)# no lacp suspend-individual</pre>	Disables LACP individual port suspension behavior on the port channel.
Step 5	no shutdown Example: <pre>switch(config-if) no shutdown</pre>	Brings the port channel administratively up.
Step 6	copy running-config startup-config Example: <pre>switch(config)# copy running-config startup-config</pre>	(Optional) Copies the running configuration to the startup configuration.

Example

This example shows how to disable LACP individual port suspension on a port channel:

```
switch# configure terminal
switch (config)# interface port-channel 1
switch(config-if)# shutdown
switch(config-if)# no lacp suspend-individual
switch(config-if)# no shutdown
```

Reenabling LACP Suspend Individual

You can reenable the default LACP individual port suspension.

Procedure

	Command or Action	Purpose
Step 1	configure terminal Example: <pre>switch# configure terminal switch(config)#</pre>	Enters global configuration mode.

	Command or Action	Purpose
Step 2	interface port-channel <i>number</i> Example: switch(config)# interface port-channel 1 switch(config-if)#	Specifies the port channel interface to configure and enters the interface configuration mode.
Step 3	shutdown Example: switch(config-if) shutdown	Administratively shuts down the port channel.
Step 4	lACP suspend-individual Example: switch(config-if)# lACP suspend-individual	Enables LACP individual port suspension behavior on the port channel.
Step 5	no shutdown Example: switch(config-if) no shutdown	Brings the port channel administratively up.
Step 6	copy running-config startup-config Example: switch(config)# copy running-config startup-config	(Optional) Copies the running configuration to the startup configuration.

Example

This example shows how to reenble the LACP individual port suspension on a port channel:

```
switch# configure terminal
switch (config)# interface port-channel 1
switch(config-if)# shutdown
switch(config-if)# lACP suspend-individual
switch(config-if)# no shutdown
```

Configuring Delayed LACP

The delayed LACP feature enables one port channel member, the delayed LACP port, to come up first as a member of a regular port channel before LACP PDUs are received. You configure the delayed LACP feature using the **lACP mode delay** command on a port channel followed by configuring the LACP port priority on a one member port of the port channel.



Note For vPC, you must enable the delayed LACP on both vPC switches.



Note For vPC, when the delayed LACP port is on the primary switch and the primary switch fails to boot, you need to remove the vPC configuration on the delayed LACP port-channel of the acting primary switch and flap the port-channel for a new port to be chosen as the delayed LACP port on the existing port-channel.

Procedure

	Command or Action	Purpose
Step 1	configure terminal	Enters global configuration mode.
Step 2	interface port-channel <i>number</i>	Specifies the port channel interface to configure and enters the interface configuration mode.
Step 3	lacp mode delay	<p>Enables delayed LACP.</p> <p>Note To disable delayed LACP, use the no lacp mode delay command.</p> <p>Complete the configuration of the delayed LACP by configuring the LACP port priority. See the "Configuring the LACP Port Priority" section for details.</p> <p>The priority of a LACP port determines the election of the delayed LACP port. The port with the lowest numerical priority is elected.</p> <p>When two or more ports have the same best priority, the VDC system MAC is used to determine which vPC is used. Then within a non-vPC switch or the elected vPC switch, the smallest of the ethernet port names is used.</p> <p>When the delayed LACP feature is configured and made effective with a port channel flap, the delayed LACP port operates as a member of a regular port channel, allowing data to be exchanged between the server and switch. After receiving the first LACP PDU, the delayed LACP port transitions from a regular port member to a LACP port member.</p> <p>Note The election of the delayed LACP port is not complete or effective until the port channel flaps on the switch or at a remote server.</p>

Example

The following example configures delayed LACP.

```
switch# config terminal
switch(config)# interface po 1
switch(config-if)# lacp mode delay
```

```
switch# config terminal
switch(config)# interface ethernet 1/1
switch(config-if)# lacp port-priority 1
switch(config-if)# channel-group 1 mode active
```

The following example disables delayed LACP.

```
switch# config terminal
switch(config)# interface po 1
switch(config-if)# no lacp mode delay
```

Configuring Port Channel Hash Distribution

Cisco NX-OS supports the adaptive and fixed hash distribution configuration for both global and port-channel levels. This option minimizes traffic disruption by minimizing Result Bundle Hash (RBH) distribution changes when members come up or go down so that flows that are mapped to unchange RBH values continue to flow through the same links. The port-channel level configuration overrules the global configuration. The default configuration is adaptive globally, and there is no configuration for each port channel, so there is no change during an ISSU. No ports are flapped when the command is applied, and the configuration takes effect at the next member link change event. Both modes work with RBH module or non-module schemes.

During an ISSD to a lower version that does not support this feature, you must disable this feature if the fixed mode command is being used globally or if there is a port-channel level configuration.

Configuring Port Channel Hash Distribution at the Global Level

Procedure

	Command or Action	Purpose
Step 1	configure terminal Example: switch# configure terminal switch(config)#	Enters global configuration mode.
Step 2	no port-channel hash-distribution {adaptive fixed} Example: switch(config)# port-channel hash-distribution adaptive switch(config)#	Specifies the port-channel hash distribution at the global level. The default is adaptive mode. The command does not take effect until the next member link event (link down/up/no shutdown/shutdown). (Do you still want to continue (y/n)? [yes])
Step 3	copy running-config startup-config Example:	(Optional) Copies the running configuration to the startup configuration.

	Command or Action	Purpose
	<code>switch(config)# copy running-config startup-config</code>	

Example

This example shows how to configure hash distribution at the global level:

```
switch# configure terminal
switch(config)# no port-channel hash-distribution fixed
```

Configuring Port Channel Hash Distribution at the Port Channel Level

Procedure

	Command or Action	Purpose
Step 1	configure terminal Example: <pre>switch# configure terminal switch(config)#</pre>	Enters global configuration mode.
Step 2	interface port-channel { <i>channel-number</i> <i>range</i> } Example: <pre>switch# interface port-channel 4 switch(config-if)#</pre>	Specifies the interface to configure, and enters the interface configuration mode.
Step 3	no port-channel port hash-distribution { <i>adaptive</i> <i>fixed</i> } Example: <pre>switch(config-if)# port-channel port hash-distribution adaptive switch(config-if)</pre>	<p>Specifies the port-channel hash distribution at the port channel level.</p> <p>There is no default.</p> <p>The command does not take effect until the next member link event (link down/up/no shutdown/shutdown). (Do you still want to continue (y/n)? [yes])</p>
Step 4	copy running-config startup-config Example: <pre>switch(config)# copy running-config startup-config</pre>	(Optional) Copies the running configuration to the startup configuration.

Example

This example shows how to configure hash distribution as a global-level command:

```
switch# configure terminal
switch(config)# no port-channel hash-distribution fixed
```

Enabling ECMP Resilient Hashing

Resilient ECMP ensures minimal impact to the existing flows when members are deleted from an ECMP group. This is achieved by replicating the existing members in a round-robin fashion at the indices that were previously occupied by the deleted members.

Procedure

	Command or Action	Purpose
Step 1	configure terminal Example: switch# <code>configure terminal</code>	Enters global configuration mode.
Step 2	hardware profile ecmp resilient Example: switch(config)# <code>hardware profile ecmp resilient</code>	Enables ECMP resilient hashing and displays the following: Warning: The command will take effect after next reload. Note This command is not supported on Cisco Nexus 9808 platform switches.
Step 3	copy running-config startup-config Example: switch(config)# <code>copy running-config startup-config</code>	Copies the running configuration to the startup configuration.
Step 4	reload Example: switch(config)# <code>reload</code>	Reboots the switch.

Disabling ECMP Resilient Hashing

Before you begin

ECMP resilient hashing is enabled.

Procedure

	Command or Action	Purpose
Step 1	configure terminal Example: switch# <code>configure terminal</code>	Enters global configuration mode.

	Command or Action	Purpose
Step 2	no hardware profile ecmp resilient Example: <pre>switch(config)# no hardware profile ecmp resilient</pre>	Disables ECMP resilient hashing and displays the following: Warning: The command will take effect after next reload.
Step 3	copy running-config startup-config Example: <pre>switch(config)# copy running-config startup-config</pre>	Copies the running configuration to the startup configuration.
Step 4	reload Example: <pre>switch(config)# reload</pre>	Reboots the switch.

Configuring ECMP Load Balancing based on Inner NVGRE Header

Configuring the ECMP load-sharing algorithm based on inner NVGRE header enables hashing for Layer 3 ECMP paths and port-channels. To configure the ECMP load-sharing algorithm based on inner NVGRE header, use the following command in global configuration mode:



Note ECMP load balancing based on inner NVGRE header is supported only on N9K-C9504-FM-E, N9K-C9508-FM-E, N9K-C9516-FM-E, N9K-C9516-FM-E2 devices.

Procedure

	Command or Action	Purpose
Step 1	ip load-sharing address { source-destination [port source-destination inner all infra-mode] } [universal-id seed] [rotate rotate] [concatenation] Example: <pre>ip load-sharing address source-destination port source-destination inner all infra-mode</pre>	Configures the ECMP load-sharing algorithm for data traffic. The following options are available for IP load sharing configuration based on inner NVGRE header: <ul style="list-style-type: none"> • The universal-id option sets the random seed for the hash algorithm and shifts the flow from one link to another. You do not need to configure the universal ID. Cisco NX-OS chooses the universal ID if you do not configure it. The <i>universal-id</i> range is from 1 to 4294967295. • The rotate option causes the hash algorithm to rotate the link picking selection so that it does not continually

	Command or Action	Purpose
		<p>choose the same link across all nodes in the network. It does so by influencing the bit pattern for the hash algorithm. This option shifts the flow from one link to another and load balances the already load-balanced (polarized) traffic from the first ECMP level across multiple links.</p> <p>If you specify a <i>rotate</i> value, the 64-bit stream is interpreted starting from that bit position in a cyclic rotation. The <i>rotate</i> range is from 1 to 63, and the default is 32.</p> <ul style="list-style-type: none"> The concatenation option ties together the hash tag values for ECMP and the hash tag values for port channels in order to use a stronger 64-bit hash. If you do not use this option, you can control ECMP load-balancing and port-channel load-balancing independently. The default is disabled. <p>The <i>infra-mode</i> option is not needed for Cisco Nexus N9K-C9516-FM-E2 based switch.</p> <p>Caution The VXLAN and infra-mode option cannot coexist. You must disable feature nv overlay before enabling this option.</p> <p>Caution This configuration enables inner header based hashing for GRE traffic on port channel and ECMP.</p>
<p>Step 2</p>	<p>(Optional) show ip load-sharing</p> <p>Example:</p> <pre>switch(config)# show ip load-sharing IPv4/IPv6 ECMP load sharing: Universal-id (Random Seed): 398176901 Load-share mode : address source-destination port source-destination inner all infra-mode GRE-Outer hash is disabled Concatenation is disabled Rotate: 32</pre>	<p>Displays the ECMP load-sharing algorithm for data traffic.</p>

Verifying the ECMP Resilient Hashing Configuration

To display ECMP Resilient Hashing configuration information, perform one of the following tasks:

Command	Purpose
<pre>switch(config)# show running-config grep "hardware profile ecmp resilient hardware profile ecmp resilient switch(config)#</pre>	Displays the enabled status.
<pre>switch(config)# show running-config grep "hardware profile ecmp resilient switch(config)#</pre>	Displays the disabled status.

Verifying the Port-Channel Configuration

To display port-channel configuration information, perform one of the following tasks:

Command	Purpose
show interface port-channel <i>channel-number</i>	Displays the status of a port-channel interface.
show feature	Displays enabled features.
load- interval {interval <i>seconds</i> {1 2 3}}	Sets three different sampling intervals to bit-rate and packet-rate statistics.
show port-channel compatibility-parameters	Displays the parameters that must be the same among the member ports in order to join a port channel.
show port-channel database [interface port-channel <i>channel-number</i>]	Displays the aggregation state for one or more port-channel interfaces.
show port-channel load-balance	Displays the type of load balancing in use for port channels.
show port-channel summary	Displays a summary for the port-channel interfaces.
show port-channel traffic	Displays the traffic statistics for port channels.
show port-channel usage	Displays the range of used and unused channel numbers.
show lacp {counters [interface port-channel <i>channel-number</i>] [interface type/slot] neighbor [interface port-channel <i>channel-number</i>] port-channel [interface port-channel <i>channel-number</i>] system-identifier]}}	Displays information about LACP.
show running-config interface port-channel <i>channel-number</i>	Displays information about the running configuration of the port-channel.

Monitoring the Port-Channel Interface Configuration

Use the following commands to display port-channel interface configuration information.

Command	Purpose
clear counters interface port-channel <i>channel-number</i>	Clears the counters.
clear lacp counters [interface port-channel <i>channel-number</i>]	Clears the LACP counters.
load- interval { interval seconds { 1 2 3 }}	Sets three different sampling intervals to bit-rate and packet-rate statistics.
show interface counters [module <i>module</i>]	Displays input and output octets unicast packets, multicast packets, and broadcast packets.
show interface counters detailed [all]	Displays input packets, bytes, and multicast and output packets and bytes.
show interface counters errors [module <i>module</i>]	Displays information about the number of error packets.
show lacp counters	Displays statistics for LACP.

Example Configurations for Port Channels

This example shows how to create an LACP port channel and add two Layer 2 interfaces to that port channel:

```
switch# configure terminal
switch (config)# feature lacp
switch (config)# interface port-channel 5
switch (config-if)# interface ethernet 1/4
switch(config-if)# switchport
switch(config-if)# channel-group 5 mode active
switch(config-if)# lacp port priority 40000
switch(config-if)# interface ethernet 1/7
switch(config-if)# switchport
switch(config-if)# channel-group 5 mode
```

This example shows how to add two Layer 3 interfaces to a channel group. The Cisco NX-OS software automatically creates the port channel:

```
switch# configure terminal
switch (config)# interface ethernet 1/5
switch(config-if)# no switchport
switch(config-if)# no ip address
switch(config-if)# channel-group 6 mode active
switch (config)# interface ethernet 2/5
switch(config-if)# no switchport
switch(config-if)# no ip address
switch(config-if)# channel-group 6 mode active
switch (config)# interface port-channel 6
```

```
switch(config-if) # ip address 192.0.2.1/8
```

Related Documents

Related Topic	Document Title
System management	<i>Cisco Nexus 9000 Series NX-OS System Management Configuration Guide</i>
High availability	<i>Cisco Nexus 9000 Series NX-OS High Availability and Redundancy Guide</i>
Licensing	<i>Cisco NX-OS Licensing Guide</i>



CHAPTER 8

Configuring vPCs

This chapter describes how to configure virtual port channels (vPCs) on Cisco NX-OS devices.


- [Information About vPCs, on page 215](#)
- [Guidelines and Limitations, on page 241](#)
- [Best Practices for Layer 3 and vPC Configuration, on page 244](#)
- [Default Settings, on page 253](#)
- [Configuring vPCs, on page 253](#)
- [Verifying the vPC Configuration, on page 281](#)
- [Monitoring vPCs, on page 282](#)
- [Configuration Examples for vPCs, on page 282](#)
- [Related Documents, on page 285](#)

Information About vPCs

vPC Overview

A virtual port channel (vPC) allows links that are physically connected to two Cisco Nexus 9000 Series devices to appear as a single port channel by a third device (see figure). The third device can be a switch, server, or any other networking device that supports port channels. A vPC can provide Layer 2 multipathing, which allows you to create redundancy and increase the bisectional bandwidth by enabling multiple parallel paths between nodes and allowing load balancing traffic.

Figure 11: vPC Architecture



You can use only Layer 2 port channels in the vPC. You configure the port channels by using one of the following:

- No protocol
- Link Aggregation Control Protocol (LACP)

When you configure the port channels in a vPC—including the vPC Peer-Link channel—without using LACP, each device can have up to 32 active links in a single port channel. When you configure the port channels in a vPC—including the vPC Peer-Link channels—using LACP, each device can have 32 active links and eight standby links in a single port channel. (See the “vPC Interactions with Other Features” section for more information on using LACP and vPCs.)



Note You must enable the vPC feature before you can configure or run the vPC functionality.

After you enable the vPC functionality, you create the peer-keepalive link, which sends heartbeat messages between the two vPC peer devices.

You can create a vPC Peer-Link by configuring a port channel on one Cisco Nexus 9000 Series chassis by using two or more Ethernet ports higher speed than 1-Gigabit Ethernet. To ensure that you have the correct hardware to enable and run a vPC, enter the **show hardware feature-capability** command. If you see an X across from the vPC in your command output, your hardware cannot enable the vPC feature.

We recommend that you configure the vPC Peer-Link Layer 2 port channels as trunks. On another Cisco Nexus 9000 Series chassis, you configure another port channel again using two or more Ethernet ports with speed higher than 1-Gigabit in the dedicated port mode. Connecting these two port channels creates a vPC Peer-Link in which the two linked Cisco Nexus devices appear as one device to a third device. The third device, or downstream device, can be a switch, server, or any other networking device that uses a regular port channel to connect to the vPC.

For modular Cisco Nexus 9500 switches, we recommend that you configure the vPC Peer-Links on dedicated ports of different modules to reduce the possibility of a failure. For the best resiliency scenario, use at least two modules.

You can use any of the interfaces of the Nexus 9000 device for the vPC Peer-Link. If you must configure all the vPC Peer-Links and core-facing interfaces on a single module, you should configure a track object that is associated with the Layer 3 link to the core and on all the links on the vPC Peer-Link on both vPC peer devices.

The vPC domain includes both vPC peer devices, the vPC peer-keepalive link, the vPC Peer-Link, and all of the port channels in the vPC domain connected to the downstream device. You can have only one vPC domain ID on each device.

In this version, you can connect each downstream device to a single vPC domain ID using a single port channel.




Note Devices attached to a vPC domain using port channels should be connected to both of vPC peers.

A vPC (see figure) provides the following benefits:

- Allows a single device to use a port channel 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 device fails
- Provides link-level resiliency
- Assures high availability

Figure 12: vPC Interfaces



Hitless vPC Role Change

A virtual port channel (vPC) allows links that are physically connected to two different Cisco Nexus 9000 Series devices to appear as a single port channel. The vPC role change feature enables you switch vPC roles

between vPC peers without impacting traffic flow. The vPC role switching is done based on the role priority value of the device under the vPC domain. A vPC peer device with lower role priority is selected as the primary vPC device during the vPC Role switch. You can use the `vpc role preempt` command to switch vPC role between peers.

For information about how to configure Hitless vPC Role Change, see [Configuring Hitless vPC Role Change, on page 279](#).


vPC Terminology

The terminology used in vPCs is as follows:

- vPC—The combined port channel between the vPC peer devices and the downstream device.
- vPC peer device—One of a pair of devices that are connected with the special port channel known as the vPC Peer-Link.
- vPC Peer-Link—The link used to synchronize state between the vPC peer devices. This link must use a 10-Gigabit Ethernet interface at a minimum. Higher-bandwidth interfaces (such as 25-Gigabit Ethernet, 40-Gigabit Ethernet, 100-Gigabit Ethernet, and so on) may also be used.
- vPC member port—An interface that belongs to a vPC.
- Host vPC port—A Fabric Extender host interfaces that belongs 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.
- vPC peer-keepalive link—The peer-keepalive link monitors the vitality of a vPC peer Cisco Nexus 9000 Series device. The peer-keepalive link sends configurable, periodic keepalive messages between vPC peer devices.

We recommend that you associate a peer-keepalive link to a separate virtual routing and forwarding (VRF) instance that is mapped to a Layer 3 interface in each vPC peer device. If you do not configure a separate VRF, the system uses the management VRF by default. However, if you use the management interfaces for the peer-keepalive link, you must put a management switch connected to both the active and standby management ports on each vPC peer device (see figure).

Figure 13: Separate Switch Required to Connect Management Ports for vPC Peer-Keepalive Link



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 a vPC.


- vPC member port—Interfaces that belong to the vPCs.
- Dual-active— Both vPC peers act as primary. This situation occurs when the peer-keepalive and vPC Peer-Link go down when both the peers are still active. In this case, the secondary vPC assumes that the primary vPC is inactive and acts as the primary vPC.
- Recovery—When the peer-keepalive and the vPC Peer-Link come up, one switch becomes the secondary vPC. On the switch that becomes the secondary vPC, the vPC links go down and come back up.

vPC Peer-Link Overview

You can have only two devices as vPC peers; each device can serve as a vPC peer to only one other vPC peer. The vPC peer devices can also have non-vPC links to other devices.

See the following figure for invalid vPC peer configurations.

Figure 14: vPC Peer Configurations That Are Not Allowed



To make a valid configuration, you first configure a port channel on each device and then configure the vPC domain. You assign the port channel on each device as a vPC Peer-Link, using the same vPC domain ID. For redundancy, we recommend that you should configure at least two of the dedicated ports into the port channel because if one of the interfaces in the vPC Peer-Link fails, the device automatically falls back to use another interface in the vPC Peer-Link.



Note We recommend that you configure the Layer 2 port channels in trunk mode.

Many operational parameters and configuration parameters must be the same in each device connected by a vPC Peer-Link (see the [Compatibility Parameters for vPC Interfaces](#) section). Because each device is completely independent on the management plane, you must ensure that the devices are compatible on the critical parameters. vPC peer devices have separate control planes. After configuring the vPC Peer-Link, you should display the configuration on each vPC peer device to ensure that the configurations are compatible.



Note You must ensure that the two devices connected by the vPC Peer-Link have certain identical operational and configuration parameters. For more information on required configuration consistency, see the [Compatibility Parameters for vPC Interfaces](#) section.

When you configure the vPC Peer-Link, the vPC peer devices negotiate that one of the connected devices is the primary device and the other connected device is the secondary device (see the “Configuring vPCs” section). By default, the Cisco NX-OS software uses the lowest MAC address to elect the primary device. However, if the role priority is set, then the device with the lowest priority will be elected as the primary device. The software takes different actions on each device—that is, the primary and secondary—only in certain failover conditions. If the primary device fails, the secondary device becomes the new primary device when the system recovers, and the previously primary device is now the secondary device.

You can also configure which of the vPC devices is the primary device. Changing the priority of the vPC peer devices can cause the interfaces in your network to go up and down. If you want to configure the role priority again to make one vPC device the primary device, configure the role priority on both the primary vPC device with a lower priority value and the secondary vPC device with the higher value. Then, shut down the port channel that is the vPC Peer-Link on both devices by entering the **shutdown** command, and finally reenables the port channel on both devices by entering the **no shutdown** command.



Note We recommend that you use two different modules for redundancy on each vPC peer device on each vPC Peer-Link.

The software keeps all traffic that forwards across the vPC peer devices as local traffic. A packet that ingresses the port channel uses one of the local links rather than moving across the vPC Peer-Link. Unknown unicast, multicast, and broadcast traffic (including STP BPDUs) are flooded across the vPC Peer-Link. The software keeps the multicast forwarding state synchronized on both of the vPC peer devices.

You can configure any of the standard load-balancing schemes on both the vPC Peer-Link devices and the downstream device (see the *Configuring Port Channels* chapter for information about load balancing).

Configuration information flows across the vPC Peer-Links using the Cisco Fabric Services over Ethernet (CFSoE) protocol. (See the [CFSoE, on page 237](#) section for more information about CFSoE.)

All MAC addresses for those VLANs configured on both devices are synchronized between vPC peer devices. The software uses CFSoE for this synchronization. (See the [CFSoE, on page 237](#) section for information about CFSoE.)

If the vPC Peer-Link fails, the software checks the status of the remote vPC peer device using the peer-keepalive link, which is a link between vPC peer devices that ensures that both devices are up. If the vPC peer device is up, the secondary vPC device disables all vPC ports on its device, to prevent loops and disappearing or flooding traffic. The data then forwards down the remaining active links of the port channel.

The software learns of a vPC peer device 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 devices. 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 device. The keepalive messages are used only when all the links in the vPC Peer-Link fail. See the “Peer-Keepalive Link and Messages” section for information about the keepalive message.

Features That You Must Manually Configure on the Primary and Secondary Devices

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 STP primary root device and configure the vPC secondary device to be the STP secondary root device. See the “vPC Peer-Links and STP” section for more information about vPCs and STP.
 - 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 per VLAN Spanning Tree plus (PVST+) so that the primary device is the root for all VLANs and configure Multiple Spanning Tree (MST) so that the primary device is the root for all instances.
- Layer 3 VLAN network interface—Configure Layer 3 connectivity from each vPC peer device by configuring a VLAN network interface for the same VLAN from both devices.
- HSRP active—If you want to use Hot Standby Router Protocol (HSRP) and VLAN interfaces on the vPC peer devices, configure the primary vPC peer device with the HSRP active highest priority. Configure the secondary device to be the HSRP standby and ensure that you have VLAN interfaces on each vPC

device that are in the same administrative and operational mode. (See the “vPC Peer-Links and Routing” section for more information on vPC and HSRP.)

While you configure Unidirectional Link Detection (UDLD), note the following recommendations:

- If LACP is used as port-channel aggregation protocol, UDLD is not required in a vPC domain.
- If LACP is not used as the port-channel aggregation protocol (static port-channel), use UDLD in normal mode on vPC member ports.
- If STP is used without Bridge Assurance and if LACP is not used, use UDLD in normal mode on vPC orphan ports.

Configuring Layer 3 Backup Routes on a vPC Peer-Link

You can use VLAN network interfaces on the vPC peer devices to link to Layer 3 of the network for such applications as HSRP and PIM. Ensure that you have a VLAN network interface configured on each peer device and that the interface is connected to the same VLAN on each device. Also, each VLAN interface must be in the same administrative and operational mode. For more information about configuring VLAN network interfaces, see the “Configuring Layer 3 Interfaces” chapter.

If a failover occurs on the vPC Peer-Link, the VLAN interfaces on the vPC peer devices are also affected. If a vPC Peer-Link fails, the system brings down associated VLAN interfaces on the secondary vPC peer device.

You can ensure that specified VLAN interfaces do not go down on the vPC secondary device when the vPC Peer-Link fails.

Peer-Keepalive Link and Messages

The Cisco NX-OS software uses the peer-keepalive link between the vPC peers to transmit periodic, configurable keepalive messages. You must have Layer 3 connectivity between the peer devices to transmit these messages; the system cannot bring up the vPC Peer-Link unless the peer-keepalive link is already up and running.



Note We recommend that you associate the vPC peer-keepalive link to a separate VRF mapped to a Layer 3 interface in each vPC peer device. If you do not configure a separate VRF, the system uses the management VRF and management ports by default. Do not use the vPC Peer-Link itself to send and receive vPC peer-keepalive messages.

If one of the vPC peer devices fails, the vPC peer device on the other side of the vPC Peer-Link senses the failure by not receiving any peer-keepalive messages. The default interval time for the vPC peer-keepalive message is 1 second, and you can configure the interval between 400 milliseconds and 10 seconds.

You can configure a hold-timeout value with a range of 3 to 10 seconds; the default hold-timeout value is 3 seconds. This timer starts when the vPC Peer-Link goes down. During this hold-timeout period, the secondary vPC peer device ignores vPC peer-keepalive messages, which ensures that network convergence occurs before a vPC action takes place. The purpose of the hold-timeout period is to prevent false-positive cases.

You can also configure a timeout value with a range of 3 to 20 seconds; the default timeout value is 5 seconds. This timer starts at the end of the hold-timeout interval. During the timeout period, the secondary vPC peer device checks for vPC peer-keepalive hello messages from the primary vPC peer device. If the secondary

vPC peer device receives a single hello message, that device disables all vPC interfaces on the secondary vPC peer device.

The difference between the hold-timeout and the timeout parameters is as follows:

- During the hold-timeout, the vPC secondary device does not take any action based on any keepalive messages received, which prevents the system taking action when the keepalive might be received just temporarily, such as if a supervisor fails a few seconds after the vPC Peer-Link goes down.
- During the timeout, the vPC secondary device takes action to become the vPC primary device if no keepalive message is received by the end of the configured interval.

See the “Configuring vPC Keepalive Link and Messages” section for information about configuring the timer for the keepalive messages.



Note 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. Peer-keepalive IP addresses must be global unicast addresses. Link-local addresses are not supported.

Use the command-line interface (CLI) to configure the interfaces you are using the vPC peer-keepalive messages as trusted ports. Leave the precedence at the default (6) or configure it higher.

vPC Peer-Gateway

You can configure vPC peer devices to act as the gateway even for packets that are destined to the vPC peer device's MAC address.

Use the **peer-gateway** command to configure this feature.



Note The **peer-gateway exclude-vlan** command that is used when configuring a VLAN interface for Layer 3 backup routing on vPC peer devices is not supported.

Some network-attached storage (NAS) devices or load balancers might have features that help to optimize the performances of particular applications. These features enable the device to avoid a routing-table lookup when responding to a request that originated from a host that is not locally attached to the same subnet. Such devices might reply to traffic using the MAC address of the sender Cisco Nexus 9000 Series device rather than the common HSRP gateway. This behavior is noncompliant with some basic Ethernet RFC standards. Packets that reach a vPC device for the nonlocal router MAC address are sent across the vPC Peer-Link and could be dropped by the built in vPC loop avoidance mechanism if the final destination is behind another vPC.

The vPC peer-gateway capability allows a vPC switch to act as the active gateway for packets that are addressed to the router MAC address of the vPC peer. This feature enables local forwarding of packets without the need to cross the vPC Peer-Link. In this scenario, the feature optimizes use of the vPC Peer-Link and avoids potential traffic loss.

Configuring the peer-gateway feature must be done on both primary and secondary vPC peers and is nondisruptive to the operations of the device or to the vPC traffic. The vPC peer-gateway feature can be configured globally under the vPC domain submode.

When you enable this feature, Cisco NX-OS automatically disables IP redirects on all interface VLANs mapped over a vPC VLAN to avoid generation of IP redirect messages for packets switched through the peer gateway router.

Packets that arrive at the peer-gateway vPC device have their Time to Live (TTL) decremented, so that packets carrying a TTL of 1 might get dropped in transit due to TTL expiration. You should take this situation into account when the peer-gateway feature is enabled and particular network protocols that source packets with a TTL of 1 operate on a vPC VLAN.

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 devices.

The vPC domain is also a configuration mode that you use to configure the keepalive messages and other vPC Peer-Link parameters rather than accept the default values. See the “Configuring vPCs” section for more information about configuring these parameters.

To create a vPC domain, you must first create a vPC domain ID on each vPC peer device using a number from 1 to 1000. You can have only one vPC domain per vPC peer.

You must explicitly configure the port channel that you want to act as the vPC Peer-Link on each device. You associate the port channel that you made a vPC Peer-Link on each device with the same vPC domain ID to form a single vPC domain. Within this domain, the system provides a loop-free topology and Layer 2 multipathing.

You can only configure these port channels and vPC Peer-Links statically. You can configure the port channels and vPC Peer-Links either using LACP or no protocol. We recommend that you use LACP with the interfaces in active mode to configure port channels in each vPC, which ensures an optimized, graceful recovery in a port-channel failover scenario and provides configuration checks against configuration mismatches among the port channels themselves.

The vPC peer devices 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 devices 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 Layer 2 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.

See the “vPC and Orphan Ports” section for more information about displaying the vPC MAC table.

After you create a vPC domain, the Cisco NX-OS software creates a system priority for the vPC domain. You can also configure a specific system priority for the vPC domain.




Note When manually configuring the system priority, you must ensure that you assign the same priority value on both vPC peer devices. If the vPC peer devices have different system priority values, vPC does not come up.

vPC Topology

The following figure shows a basic configuration in which the Cisco Nexus 9000 Series device ports are directly connected to another switch or host and are configured as part of a port channel that becomes part of a vPC.


Figure 15: Switch vPC Topology



In the figure, vPC 20 is configured on port channel 20, which has Eth1/10 on the first device and Eth2/1 on the second as member ports.

You can configure a vPC from the peer devices through Fabric Extenders (FEXs) as shown in the figure.

Figure 16: FEX Straight-Through Topology (Host vPC)



In the figure, each FEX is single-homed (straight-through FEX topology) with a Cisco Nexus 9000 Series device. The host interfaces on this FEX are configured as port channels and those port channels are configured as vPCs. Eth101/1/1 and Eth102/1/5 are configured as members of PO200, and PO200 is configured for vPC 200.

In both topologies, port channels P020 and P0200 must be configured identically on the peer switches and configuration synchronization is used to synchronize the configurations of the vPC switches.

See the [Cisco Nexus 2000 Series NX-OS Fabric Extender Configuration Guide for Cisco Nexus 9000 Series Switches](#) for more information about configuring FEX ports.

Compatibility Parameters for vPC Interfaces

Many configuration and operational parameters must be identical on all interfaces in the vPC. We recommend that you configure the Layer 2 port channels that you use for the vPC Peer-Link in trunk mode.

After you enable the vPC feature and configure the vPC Peer-Link on both vPC peer devices, Cisco Fabric Services (CFS) messages provide a copy of the configuration on the local vPC peer device configuration to the remote vPC peer device. The system then determines whether any of the crucial configuration parameters differ on the two devices. (See the “vPC and Orphan Ports” section for more information about CFS.)



Note 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.



Note The port channel compatibility parameters must be the same for all the port channel members on the physical switch. You cannot configure shared interfaces to be part of a vPC.

The compatibility check process for vPCs differs from the compatibility check for regular port channels.

See the “Configuring Port Channels” chapter for information about regular port channels.

Configuration Parameters That Must Be Identical

The configuration parameters in this section must be configured identically on both devices of the vPC Peer-Link; otherwise, the vPC moves fully or partially into a suspended mode.



Note You must ensure that all interfaces in the vPC have the identical operational and configuration parameters listed in this section.



Note 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 devices automatically check for compatibility for some 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 (port-channel mode can, however, be active/passive on each side of the vPC peer)

- 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
- Enable/disable state per VLAN
- STP global settings:
 - Bridge Assurance setting
 - Port type setting
 - Loop Guard settings
- STP interface settings:
 - Port type setting
 - Loop Guard
 - Root Guard
- Maximum Transmission Unit (MTU)

If any of these parameters are not enabled or defined on either device, 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.

In the output of **show vpc** or **show vpc brief** command, after every 50th configured vPC port-channel the following message will be displayed:

```
Please check "show vpc consistency-parameters vpc <vpc-num>" for the consistency reason of down vpc and for type-2 consistency reasons for any vpc.
```

Configuration Parameters That Should Be Identical

When any of the following parameters are not configured identically on both vPC peer devices, a misconfiguration might cause undesirable behavior in the traffic flow:

- MAC aging timers

- Static MAC entries
- VLAN interface—Each device 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 device of the vPC Peer-Link do not pass traffic using the vPC or vPC Peer-Link. You must create all VLANs on both the primary and secondary vPC devices, or the VLAN will be suspended.
- All ACL configurations and parameters
- Quality of Service (QoS) configuration and parameters
- STP interface settings:
 - BPDU Filter
 - BPDU Guard
 - Cost
 - Link type
 - Priority
 - VLANs (Rapid PVST+)
- Port security
- Cisco Trusted Security (CTS)
- Dynamic Host Configuration Protocol (DHCP) snooping
- Network Access Control (NAC)
- Dynamic ARP Inspection (DAI)
- IP source guard (IPSG)
- Internet Group Management Protocol (IGMP) snooping
- Hot Standby Routing Protocol (HSRP)
- Protocol Independent Multicast (PIM)
- All routing protocol configurations

To ensure that all the configuration parameters are compatible, we recommend that you display the configurations for each vPC peer device once you configure the vPC.

Consequences of Parameter Mismatches

You can configure the graceful consistency check feature, which suspends only the links on the secondary peer device when a mismatch is introduced in a working vPC. This feature is configurable only in the CLI and is enabled by default.

The graceful consistency-check command is configured by default.

As part of the consistency check of all parameters from the list of parameters that must be identical, the system checks the consistency of all VLANs.

The vPC remains operational, and only the inconsistent VLANs are brought down. This per-VLAN consistency check feature cannot be disabled and does not apply to Multiple Spanning Tree (MST) VLANs.

vPC Number

Once you have created the vPC domain ID and the vPC Peer-Link, you create port channels to attach the downstream device to each vPC peer device. That is, you create one port channel to the downstream device from the primary vPC peer device and you create another port channel to the downstream device from the secondary peer device.



Note We recommend that you configure the ports on the downstream devices that connect to a host or a network device that is not functioning as a switch or a bridge as STP edge ports.

On each vPC peer device, you assign a vPC number to the port channel that connects to the downstream device. You will experience minimal traffic disruption when you are creating vPCs. To simplify the configuration, you can assign the vPC ID number to every port channel to be the same as the port channel itself (that is, vPC ID 10 for port channel 10).



Note The vPC number that you assign to the port channel that connects to the downstream device from the vPC peer device must be identical on both vPC peer devices.

Moving Other Port Channels into a vPC



Note You must attach a downstream device using a port channel to both vPC peer devices.

To connect to the downstream device, you create a port channel to the downstream device from the primary vPC peer device and you create another port channel to the downstream device from the secondary peer device. On each vPC peer device, you assign a vPC number to the port channel that connects to the downstream device. You will experience minimal traffic disruption when you are creating vPCs.

vPC Object Tracking



Note We recommend that you configure the vPC Peer-Links on dedicated ports of different modules on Cisco Nexus 9500 devices. This is recommended to reduce the possibility of a failure. For the best resiliency scenario, use at least two modules.

vPC object tracking is used to prevent traffic black-holing in case of failure of a module where both vPC Peer-Link and uplinks to the core resides. By tracking interface feature can suspend vPC on affected switch and prevent traffic black-holing.

If you must configure all the vPC Peer-Links and core-facing interfaces on a single module, you should configure, using the command-line interface, a track object and a track list that is associated with the Layer 3 link to the core and on all vPC Peer-Links on both vPC peer devices. You use this configuration to avoid dropping traffic if that particular module goes down because when all the tracked objects on the track list go down, the system does the following:

- Stops the vPC primary peer device sending peer-keepalive messages, which forces the vPC secondary peer device to take over.
- Brings down all the downstream vPCs on that vPC peer device, which forces all the traffic to be rerouted in the access switch toward the other vPC peer device.

Once you configure this feature and if the module fails, the system automatically suspends all the vPC links on the primary vPC peer device and stops the peer-keepalive messages. This action forces the vPC secondary device to take over the primary role and all the vPC traffic to go to this new vPC primary device until the system stabilizes.

You should create a track list that contains all the links to the core and all the vPC Peer-Links as its object. Enable tracking for the specified vPC domain for this track list. Apply this same configuration to the other vPC peer device. See the [Cisco Nexus 9000 Series NX-OS Unicast Routing Configuration Guide](#) for information about configuring object tracking and track lists.



Note This example uses Boolean OR in the track list and forces all traffic to the vPC peer device only for a complete module failure. If you want to trigger a switchover when any core interface or vPC Peer-Link goes down, use a Boolean AND in the track list below.

To configure a track list to switch over a vPC to the remote peer when all related interfaces on a single module fail, follow these steps:

1. Configure track objects on an interface (Layer 3 to core) and on a port channel (vPC Peer-Link).

```
switch(config-if)# track 35 interface ethernet 8/35 line-protocol
switch(config-track)# track 23 interface ethernet 8/33 line-protocol
switch(config)# track 55 interface port-channel 100 line-protocol
```

2. Create a track list that contains all the interfaces in the track list using the Boolean OR to trigger when all objects fail.

```
switch(config)# track 44 list boolean OR
switch(config-track)# object 23
switch(config-track)# object 35
switch(config-track)# object 55
switch(config-track)# end
```

3. Add this track object to the vPC domain:

```
switch(config)# vpc domain 1
switch(config-vpc-domain)# track 44
```

4. Display the track object:

```

switch# show vpc brief
Legend:
(*) - local vPC is down, forwarding via vPC peer-link
vPC domain id : 1
Peer status : peer adjacency formed ok
vPC keep-alive status : peer is alive
Configuration consistency status: success
vPC role : secondary
Number of vPCs configured : 52
Track object : 44
vPC Peer-link status
-----
id Port Status Active vlans
-----
1 Po100 up 1-5,140
vPC status
-----
id Port Status Consistency Reason Active vlans
-----
1 Po1 up success success 1-5,140

```

This example shows how to display information about the track objects:

```

switch# show track brief
Track Type Instance Parameter State Last
Change
23 Interface Ethernet8/33 Line Protocol UP 00:03:05
35 Interface Ethernet8/35 Line Protocol UP 00:03:15
44 List ----- Boolean
or UP 00:01:19
55 Interface port-channel100 Line Protocol UP 00:00:34

```

vPC Interactions with Other Features

vPC and LACP

LACP uses the system MAC address of the vPC domain to form the LACP Aggregation Group (LAG) ID for the vPC. (See the “Configuring Port Channels” chapter for information about LAG-ID and LACP.)

You can use LACP on all the vPC port channels, including those channels from the downstream device. We recommend that you configure LACP with active mode on the interfaces on each port channel on the vPC peer devices. This configuration allows you to more easily detect compatibility between devices, unidirectional links, and multihop connection, and provides dynamic reaction to run-time changes and link failures.

We recommend that you manually configure the system priority on the vPC Peer-Link devices to ensure that the vPC Peer-Link devices have a higher LACP priority than the downstream connected devices. 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 devices. If the vPC peer devices have different system priority values, vPC does not come up.

vPC Peer-Links and STP

Although vPCs provide a loop-free Layer 2 topology, STP is still required to provide a fail-safe mechanism to protect against any incorrect or defective cabling or possible misconfiguration. When you first bring up a vPC, 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. If the STP enhancements are already configured, they do not cause any problems for the vPC Peer-Links..

When you are running both MST and Rapid PVST+, ensure that the PVST simulation feature is correctly configured.

See the [Cisco Nexus 9000 Series NX-OS Layer 2 Switching Configuration Guide](#) for information about STP enhancement features and PVST simulation.



Note You must configure a list of parameters to be identical on the vPC peer devices on both sides of the vPC Peer-Link. See the “Compatibility Parameters for vPC Interfaces” section for information about these required matched settings.

STP is distributed; that is, the protocol continues running on both vPC peer devices. However, the configuration on the vPC peer device elected as the primary device controls the STP process for the vPC interfaces on the secondary vPC peer device.

The primary vPC device synchronizes the STP state on the vPC secondary peer device using Cisco Fabric Services over Ethernet (CFS over Ethernet). See the “vPC and Orphan Ports” section for information about CFS over Ethernet.

The STP process for vPC also relies on the periodic keepalive messages to determine when one of the connected devices on the vPC Peer-Link fails. See the “Peer-Keepalive Link and Messages” section for information about these messages.

The vPC manager performs a proposal/handshake agreement between the vPC peer devices that set the primary and secondary devices and coordinates the two devices for STP. The primary vPC peer device then controls the STP protocol on both the primary and secondary devices. We recommend that you configure the primary vPC peer device as the STP primary root device and configure the secondary vPC device to be the STP secondary root device.

If the primary vPC peer device fails over to the secondary vPC peer device, there is no change in the STP topology.

The BPDUs use the MAC address set for the vPC for the STP bridge ID in the designated bridge ID field. The vPC primary device sends these BPDUs on the vPC interfaces.

You must configure both ends of vPC Peer-Link with the identical STP configuration for the following parameters:

- STP global settings:
 - STP mode
 - STP region configuration for MST
 - Enable/disable state per VLAN

- Bridge Assurance setting
- Port type setting
- Loop Guard settings
- STP interface settings:
 - Port type setting
 - Loop Guard
 - Root Guard



Note If any of these parameters are misconfigured, the Cisco NX-OS software suspends all interfaces in the vPC. Check the syslog and enter the **show vpc brief** command to see if the vPC interfaces are suspended.

Ensure that the following STP interface configurations are identical on both sides of the vPC Peer-Links or you may see unpredictable behavior in the traffic flow:

- BPDU Filter
- BPDU Guard
- Cost
- Link type
- Priority
- VLANs (PVRST+)



Note Display the configuration on both sides of the vPC Peer-Link to ensure that the settings are identical.

You can use the **show spanning-tree** command to display information about the vPC when that feature is enabled. See the [Cisco Nexus 9000 Series NX-OS Layer 2 Switching Configuration Guide](#) for an example.



Note We recommend that you configure the ports on the downstream devices as STP edge ports. You should configure all host ports connected to a switch as STP edge ports. See the [Cisco Nexus 9000 Series NX-OS Layer 2 Switching Configuration Guide](#) for more information about STP port types.

vPC Peer Switch

The vPC peer switch feature was added to Cisco NX-OS to address performance concerns around STP convergence. This feature allows a pair of Cisco Nexus 9000 Series devices to appear as a single STP root in the Layer 2 topology. This feature eliminates the need to pin the STP root to the vPC primary switch and improves vPC convergence if the vPC primary switch fails.

To avoid loops, the vPC Peer-Link is excluded from the STP computation. In vPC peer switch mode, STP BPDUs are sent from both vPC peer devices to avoid issues related to STP BPDU timeout on the downstream switches, which can cause traffic disruption.

This feature can be used with the pure peer switch topology in which the devices all belong to the vPC.



Note Peer-switch feature is supported on networks that use vPC and STP-based redundancy is not supported. If the vPC Peer-Link fail in a hybrid peer-switch configuration, you can lose traffic. In this scenario, the vPC peers use the same STP root ID as well as the same bridge ID. The access switch traffic is split in two with half going to the first vPC peer and the other half to the second vPC peer. With vPC Peer-Link failure, there is no impact to the north/south traffic but the east/west traffic is lost.

See the [Cisco Nexus 9000 Series NX-OS Layer 2 Switching Configuration Guide](#) for information about STP enhancement features and Rapid PVST+.

vPC and ARP or ND

A feature was added to Cisco NX-OS to address table synchronization across vPC peers using the reliable transport mechanism of the Cisco Fabric Service over Ethernet (CFS over Ethernet) protocol. You must enable the **ip arp synchronize** and **ipv6 nd synchronize** commands to support faster convergence of address tables between the vPC peers. This convergence overcomes the delay that occurs in ARP table restoration for IPv4 or ND table restoration for IPv6 when the vPC Peer-Link port channel flaps or when a vPC peer comes back online.

vPC Multicast—PIM, IGMP, and IGMP Snooping

The Cisco NX-OS software for the Nexus 9000 Series devices supports the following on a vPC:

- PIM Any Source Multicast (ASM).
- PIM Source-Specific Multicast (SSM) (7.0(3)I4(1) and later).



Note The Cisco NX-OS software does not support Bidirectional (BIDR) on a vPC.

The software keeps the multicast forwarding state synchronized on both of the vPC peer devices. The IGMP snooping process on a vPC peer device shares the learned group information with the other vPC peer device through the vPC Peer-Link; the multicast states are always synchronized on both vPC peer devices. The PIM process in vPC mode ensures that only one of the vPC peer devices forwards the multicast traffic to the receivers.

Each vPC peer is a Layer 2 or Layer 3 device. Multicast traffic flows from only one of the vPC peer devices. You might see duplicate packets in the following scenarios:

- Orphan hosts
- When the source and receivers are in the Layer 2 vPC cloud in different VLANs with multicast routing enabled and a vPC member link goes down.

You might see negligible traffic loss in the following scenarios:

- When you reload the vPC peer device that is forwarding the traffic.

- When you restart PIM on the vPC peer device that is forwarding the traffic.

Overall multicast convergence times are scale and vPC role change / PIM restart duration dependent.

Ensure that you dual-attach all Layer 3 devices to both vPC peer devices. If one vPC peer device goes down, the other vPC peer device continues to forward all multicast traffic normally.

The following outlines vPC PIM and vPC IGMP/IGMP snooping:

- vPC PIM—The PIM process in vPC mode ensures that only one vPC peer device forwards multicast traffic. The PIM process in vPC mode synchronizes the source state with both vPC peer devices and elects which vPC peer device forwards the traffic.
- vPC IGMP/IGMP snooping—The IGMP process in vPC mode synchronizes the designated router (DR) information on both vPC peer devices. Dual DRs are available for IGMP when you are in vPC mode. Dual DRs are not available when you are not in vPC mode, because both vPC peer devices maintain the multicast group information between the peers.



Note A PIM adjacency between a Switched Virtual Interface (SVI) on a vPC VLAN (a VLAN that is carried on a vPC Peer-Link) and a downstream device is not supported; this configuration can result in dropped multicast packets. If a PIM neighbor relationship is required with a downstream device, a physical Layer 3 interface must be used on the Nexus switches instead of a vPC SVI.

For SVIs on vPC VLANs, only one PIM adjacency is supported, which is with the vPC peer switch. PIM adjacencies over the vPC Peer-Link with devices other than the vPC peer switch for the vPC-SVI are not supported.

You should enable or disable IGMP snooping identically on both vPC peer devices, and all the feature configurations should be identical. IGMP snooping is on by default.



Note The following commands are not supported in vPC mode:

- **ip pim spt-threshold infinity**
- **ip pim use-shared-tree-only**

See the *Cisco Nexus 9000 Series NX-OS Multicast Routing Configuration Guide* for more information about multicasting.

Multicast PIM Dual DR (Proxy DR)

By default, a multicast router sends PIM joins upstream only if it has interested receivers. These interested receivers can either be IGMP hosts (they communicate through IGMP reports) or other multicast routers (they communicate through PIM joins).

In the Cisco NX-OS vPC implementation, PIM works in dual designated router (DR) mode. That is, if a vPC device is a DR on a vPC SVI outgoing interface (OIF), its peer automatically assumes the proxy DR role. IGMP adds an OIF (the report is learned on that OIF) to the forwarding if the OIF is a DR. With dual DRs, both vPC devices have an identical (*,G) entry with respect to the vPC SVI OIFs as shown in this example:

```
VPC Device1:
-----
(*,G)
oif1 (igmp)
VPC Device2:
-----
(*,G)
oif1 (igmp)
```

IP PIM PRE-BUILD SPT

When the multicast source is in a Layer 3 cloud (outside the vPC domain), one vPC peer is elected as the forwarder for the source. This forwarder election is based on the metrics to reach the source. If there is a tie, the vPC primary is chosen as the forwarder. Only the forwarder has the vPC OIFs in its associated (S,G) and the nonforwarder (S,G) has 0 OIFs. Therefore, only the forwarder sends PIM (S,G) joins toward the source as shown in this example:

```
VPC Device1 (say this is Forwarder for Source 'S'):
-----
(*,G)
oif1 (igmp)
(S,G)
oif1 (mrib)
VPC Device2:
-----
(*,G)
oif1 (igmp)
(S,G)
NULL
```

In the case of a failure (for example, a Layer 3 Reverse Path Forwarding (RPF) link on the forwarder becomes inoperable or the forwarder gets reloaded), if the current nonforwarder ends up becoming the forwarder, it has to start sending PIM joins for (S,G) toward the source to pull the traffic. Depending upon the number of hops to reach the source, this operation might take some time (PIM is a hop-by-hop protocol).

To eliminate this issue and get better convergence, use the **ip pim pre-build-spt** command. This command enables PIM send joins even if the multicast route has 0 OIFs. In a vPC device, the nonforwarder sends PIM (S,G) joins upstream toward the source. The downside is that the link bandwidth upstream from the nonforwarder gets used for the traffic that is ultimately dropped by it. The benefits that result with better convergence far outweigh the link bandwidth usage. Therefore, we recommend that you use this command if you use vPCs.

vPC Peer-Links and Routing

The First Hop Redundancy Protocols (FHRPs) interoperate with vPCs. The Hot Standby Routing Protocol (HSRP), and Virtual Router Redundancy Protocol (VRRP) all interoperate with vPCs. We recommend that you dual-attach all Layer 3 devices to both vPC peer devices.

The primary FHRP device responds to ARP requests, even though the secondary vPC device forwards the data traffic.

To simplify initial configuration verification and vPC/HSRP troubleshooting, you can configure the primary vPC peer device with the FHRP active router highest priority.

In addition, you can use the `priority` command in the `if-hsrp` configuration mode to configure failover thresholds for when a group state enabled on a vPC Peer-Link is in standby or in listen state. You can configure lower and upper thresholds to prevent the interface from going up and down.

VRRP acts similarly to HSRP when running on vPC peer devices. You should configure VRRP the same way that you configure HSRP.

When the primary vPC peer device fails over to the secondary vPC peer device, the FHRP traffic continues to flow seamlessly.

We recommend that you configure routing adjacency between the two vPC peer devices to act as a backup routing path. If one vPC peer device loses Layer 3 uplinks, the vPC can redirect the routed traffic to the other vPC peer device and leverage its active Layer 3 uplinks.

You can configure the inter-switch link for a backup routing path in the following ways:

- Create a Layer 3 link between the two vPC peer devices.
- Use the non-VPC VLAN trunk with a dedicated VLAN interface.
- Use a vPC Peer-Link with a dedicated VLAN interface.

We do not recommend that you configure the burnt-in MAC address option (`use-bia`) for HSRP or manually configure virtual MAC addresses for any FHRP protocol in a vPC environment because these configurations can adversely affect vPC load balancing. The HSRP `use-bia` option is not supported on vPCs. When you are configuring custom MAC addresses, you must configure the same MAC address on both vPC peer devices.

You can use the `delay restore` command to configure a restore timer that delays the vPC coming back up until after the peer adjacency forms and the VLAN interfaces are back up. This feature enables you to avoid packet drops when the routing tables might not be converged before the vPC is once again passing traffic. Use the `delay restore` command to configure this feature.

To delay the VLAN interfaces on the restored vPC peer device from coming up, use the `interfaces-vlan` option of the `delay restore` command.

See the [Cisco Nexus 9000 Series NX-OS Unicast Routing Configuration Guide](#) for more information about FHRPs and routing.

CFSoE

The Cisco Fabric Services over Ethernet (CFSoE) is a reliable state transport mechanism that is used 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.

The CFSoE transport is local to each VDC.

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. You must enable CFSoE for vPC functionality. If you do enter either of these commands with vPC 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 CFS over Ethernet.

CFS also transports data over TCP/IP. See the [Cisco Nexus 9000 Series NX-OS System Management Configuration Guide](#) for more information about CFS over IP.



Note The software does not support CFS regions.

vPC and Orphan Ports

When a device that is not vPC-capable connects to each peer, the connected ports are known as orphan ports because they are not members of a vPC. The device’s link to one peer will be active (forwarding) and the other link will be standby (blocking) due to STP.

If a vPC Peer-Link failure or restoration occurs, an orphan port’s connectivity might be bound to the vPC failure or restoration process. For example, if a device’s active orphan port connects to the secondary vPC peer, the device loses any connections through the primary peer if a vPC Peer-Link failure occurs and the vPC ports are suspended by the secondary peer. If the secondary peer were to also suspend the active orphan port, the device’s standby port becomes active, provides a connection to the primary peer, and restores connectivity. You can configure in the CLI that specific orphan ports are suspended by the secondary peer when it suspends its vPC ports and are restored when the vPC is restored.

Virtualization Support

All ports in a given vPC must be in the same VDC. This version of the software supports only one vPC domain per VDC. You can use the numbers from 1 to 4096 in each VDC to number the vPC.

vPC Recovery After an Outage

In a data center outage, both the vPC peer in vPC domain get reloaded. Occasionally only one peer can be restored. With no functioning peer-keepalive or vPC Peer-Link, the vPC cannot function normally, a method might be available to allow vPC services to use only the local ports of the functional peer.

Autorecovery

You can configure the Cisco Nexus 9000 Series device to restore vPC services when its peer fails to come online by using the **auto-recovery** command. You must save this setting in the startup configuration. On reload, if the vPC Peer-Link is down and three consecutive peer-keepalive messages are lost, the secondary device assumes the primary STP role and the primary LACP role. The software reinitializes the vPCs, bringing up its local ports. Because there are no peers, the consistency check is bypassed for the local vPC ports. The device elects itself to be the STP primary regardless of its role priority and also acts as the primary device for LACP port roles.

vPC Peer Roles After a Recovery

When the other peer device completes its reload and adjacency forms, the following process occurs:

1. The first vPC peer maintains its current role to avoid any transition reset to other protocols. The peer accepts the other available role.

2. When an adjacency forms, consistency checks are performed and appropriate actions are taken.

High Availability

During an In-Service Software Upgrade (ISSU), the software reload process on the first vPC device locks its vPC peer device by using CFS messaging over the vPC communications channel. Only one device at a time is upgraded. When the first device completes its upgrade, it unlocks its peer device. The second device then performs the upgrade process, locking the first device as it does so. During the upgrade, the two vPC devices temporarily run different releases of Cisco NX-OS, however the system functions correctly because of its backward compatibility support.



Note See the [Cisco Nexus 9000 Series NX-OS High Availability and Redundancy Guide](#) for complete information about high-availability features.

vPC Forklift Upgrade Scenario

The following procedure describes a scenario of migrating pair of Cisco Nexus 9500 switches in a vPC domain to a different pair of Cisco Nexus 9500 switches with a same type of line cards. Migrating from Cisco Nexus 9504 switches to Cisco Nexus 9508 switches for the need of more interfaces is a typical example of such migration. The following migration scenarios are not supported:

- Migration of Cisco Nexus 9500 switches with a different set of line cards. For example, from a Cisco Nexus 9500 switches with N9K-X94xx line card to Cisco Nexus 9500 switches with N9K-X97xx line card.
- Migration between different generations of Cisco Nexus 9300 switches. For example, migration from Cisco Nexus N9K-C9372PX to Cisco Nexus N9K-93180YC-EX switches
- Having different generations of Cisco Nexus 9000 switches in a vPC domain is not supported

Considerations for a vPC forklift upgrade:

- vPC Role Election and Sticky-bit

By default, the Cisco NX-OS software uses the lowest MAC address to elect the primary device. However, if the role priority is set, then the device with the lowest priority will be elected as the primary device. When the primary device is reloaded, the system comes back online and connectivity to the vPC secondary device (now the operational primary) is restored. The operational role of the secondary device (operational primary) does not change (to avoid unnecessary disruptions). This behavior is achieved with a sticky-bit, where the sticky information is not saved in the startup configuration. This method makes the device that is up and running win over the reloaded device. Hence, the vPC primary becomes the vPC operational secondary. Sticky-bit is also set when a vPC node comes up with vPC Peer-Link and peer-keepalive down and it becomes primary after the auto recovery period.

- vPC Delay Restore

The delay restore timer is used to delay the vPC from coming up on the restored vPC peer device after a reload when the peer adjacency is already established.

To delay the VLAN interfaces on the restored vPC peer device from coming up, use the **interfaces-vlan** option of the **delay restore** command.

- vPC Auto-Recovery

During a data center power outage when both vPC peer switches go down, if only one switch is restored, the auto-recovery feature allows that switch to assume the role of the primary switch and the vPC links come up after the auto-recovery time period. The default auto-recovery period is 240 seconds.

The following example is a migration scenario that replaces vPC peer nodes Node1 and Node2 with New_Node1 and New_Node2.

	Migration Step	Expected Behavior	Node1 Configured role (Ex: role priority 100)	Node1 Operational role	Node2 Configured role (Ex: role priority 200)	Node2 Operational role
1	Initial state	Traffic is forwarded by both vPC peers – Node1 and Node2. Node1 is primary and Node2 is secondary.	primary	Primary Sticky bit: False	secondary	Secondary Sticky bit: False
2	Node2 replacement – Shut all vPCs and uplinks on Node2. vPC Peer-Link and vPC peer-keepalive are in administrative up state.	Traffic converged on Primary vPC peer Node1.	primary	Primary Sticky bit: False	secondary	Secondary Sticky bit: False
3	Remove Node2.	Node1 will continue to forward traffic.	primary	Primary Sticky bit: False	n/a	n/a
4	Configure New_Node2. Copy the configuration to startup config. vPC vPC Peer-Link and peer-keepalive in administrative up state. Power off New_Node2. Make all connections. Power on New_Node2.	New_Node2 will come up as secondary. Node1 continue to be primary. Traffic will continue to be forwarded on Node01.	primary	Primary Sticky bit: False	secondary	Secondary Sticky bit: False

	Migration Step	Expected Behavior	Node1 Configured role (Ex: role priority 100)	Node1 Operational role	Node2 Configured role (Ex: role priority 200)	Node2 Operational role
5	Bring up all vPCs and uplink ports on New_Node2.	Traffic will be forwarded by both Node 1 and New_Node2.	primary	Primary Sticky bit: False	secondary	Secondary Sticky bit: False
6	Node1 replacement - Shut vPCs and uplinks on Node1.	Traffic will converge on New_Node2.	primary	Primary Sticky bit: False	secondary	Secondary Sticky bit: False
7	Remove Node1.	New_Node2 will become secondary, operational primary and sticky bit will be set to True.	n/a	n/a	secondary	Primary Sticky bit: True
8	Configure New_Node1. Copy running to startup. Power off the new Node1. Make all connections. Power on New_Node1.	New_Node1 will come up as primary, operational secondary.	primary	Secondary Sticky bit: False	secondary	Primary Sticky bit: True
9	Bring up all vPCs and uplink ports on New_Node1.	Traffic will be forwarded by both New Node1 and new Node2.	primary	Secondary Sticky bit: False	secondary	Primary Sticky bit: True



Note If you prefer to have the configured secondary node as the operational secondary and the configured primary as the operational primary, then Node2 can be reloaded at the end of the migration. This is optional and does not have any functional impact.

Guidelines and Limitations

vPCs have the following configuration guidelines and limitations:

- Make sure that both vPC peers are in the same mode (regular mode or enhanced mode) before performing a non-disruptive upgrade.



Note vPC peering between an enhanced ISSU mode (boot mode lxc) configured switch and a non-enhanced ISSU mode switch is not supported.

- **show** commands with the **internal** keyword are not supported.
- Cisco Nexus 9000 Series switches do not support NAT on vPC topology.
- The **spanning-tree pseudo-information** command is not available on Cisco Nexus 92160 and Cisco Nexus 93180 switches starting from Release 7.0(3)I4(1).
- vPC peers must run the same Cisco NX-OS release. During a software upgrade, make sure to upgrade the primary vPC peer first.
- All ports for a given vPC must be in the same VDC.
- You must enable vPCs before you can configure them.
- You must configure the peer-keepalive link and messages before the system can form the vPC Peer-Link.
- Only Layer 2 port channels can be in vPCs.
- You must configure both vPC peer devices; the configuration is not sent from one device to the other.
- To configure multilayer (back-to-back) vPCs, you must assign unique vPC domain ID for each respective vPC.
- Check that the necessary configuration parameters are compatible on both sides of the vPC Peer-Link. See the “Compatibility Parameters for vPC Interfaces” section for information about compatibility recommendations.
- You might experience minimal traffic disruption while configuring vPCs.
- The software does not support BIDR PIM on vPCs.
- The software does not support DHCP snooping, DAI, or IPSG in a vPC environment; DHCP Relay is supported.
- The software does not support CFS regions.
- Port security is not supported on port channels.
- When **peer-switch** features are configured under **vpc domain** configuration mode on two Cisco Nexus 9000 Series switches, the spanning-tree root changes even for VLANs that are not enabled on the vPC Peer-Link. Both the switches act as one system with one MAC address as the bridge address. This is true even for non-vPC mst-instance or VLANs. Therefore, a non vPC Peer-Link between the two switches gets blocked as a backup link. This is an expected behavior.
- We recommend that you configure all the port channels in the vPC using LACP with the interfaces in active mode.
- Having the same Hot Standby Router Protocol (HSRP)/Virtual Router Redundancy Protocol (VRRP) group on all nodes on a double sided vPC is supported on Cisco NX-OS 7.0(3)I2(1) and later releases.
- When migrating from a pair of spine nodes to a pair of Cisco Nexus 9000 devices, the HSRP priority should be configured so that the Cisco Nexus 9000 vPC peers are in Active/Standby state. There is no

support for Cisco Nexus 9000 vPC peers in HSRP state to be in Active/Listen state, or Standby/Listen state. (7.0)I2(2) or later)

- When using vPCs, we recommend that you use default timers for FHRP (HSRP, VRRP), and PIM configurations. There is no advantage in convergence times when using aggressive timers in vPC configurations.
- If you configure open shortest path first (OSPF) in a vPC environment, use the following timer commands in router configuration mode on the core switch to ensure fast OSPF convergence when a vPC Peer-Link is shut down:

```
switch (config-router)# timers throttle spf 1 50 50
switch (config-router)# timers lsa-arrival 10
```

See the *Cisco Nexus 9000 Series NX-OS Unicast Routing Configuration Guide* for further details about OSPF.

- BFD for VRRP/HSRP is not supported in a vPC environment.
- The STP port cost is fixed to 200 in a vPC environment.
- Jumbo frames are enabled by default on the vPC Peer-Link.
- To accommodate increased traffic when the vPC goes down and traffic needs to cross the vPC Peer-Link, it is a best practice to use multiple high bandwidth interfaces (such as the 40G interfaces for the Cisco Nexus 9000) across linecards for the vPC Peer-Link.
- The **vpc orphan-ports suspend** command also applies to ports in non-vPC VLANs and Layer 3 ports. However, it is recommended to be used with ports in VPC VLANs.
- Starting in NX-OS 7.0(3)I5(2), FEX-AA (dual-homed FEX) and FEX-ST (FEX straight-thru) topologies (FEX-AA and FEX-ST) are supported. The following mixing is not supported as the parent switches:
 - Cisco Nexus 9300-EX and 9300 switches
 - Cisco Nexus 9300 and 9500 switches
 - Cisco Nexus 9300-EX and 9500 switches
- Starting in NX-OS 7.0(3)I2(2), when configuring vPCs, the behavior previously provided by using the `ip pim pre-build-spt` command has now been enabled automatically by default and cannot be disabled.
- Starting in NX-OS 7.0(3)I2(2), a vPC port channel member link that is operating in Individual state will be flapped while checking for VLAN inconsistencies. To avoid having the link flapped during server provisioning, disable the VPC graceful consistency check with the **no graceful consistency-check** command.

The following example disables the VPC graceful consistency check:

```
switch# conf t
Enter configuration commands, one per line. End with CNTL/Z.

switch(config)# vpc domain 1
switch(config-vpc-domain)# no graceful consistency-check
```

- vPC STP hitless role change feature is supported only from Cisco Nexus 9000 Release 7.0(3)I7(1) onwards.

- vPC role change can be performed from either of the peer devices.
- When forming a vPC domain between two Cisco Nexus 9300 Series switches, both switches must be the exact same model to form a supported vPC domain. When forming a vPC domain between two Cisco Nexus 9500 Series switches, both switches must consist of the same models of line cards, fabric modules, supervisor modules, and system controllers inserted in the same slots of the chassis to form a supported vPC domain.
- If the original secondary device has higher role priority value than the original primary device, role swapping cannot be performed. Change the role priority on either vPC device so that the value of the original secondary device is lower than the original primary one. To view the existing role of a device, use the `show vpc role` command on local and peer switch.
- Always check the existing configured role priority before configuring vPC hitless role change feature
- In a vPC domain, enable the `peer-switch` command, where both vPC peers have same STP priorities, and ensure it is operational before issuing a role change. If you do not enable the `peer-switch` command, it can lead to convergence issues. Use `show spanning-tree summary | grep peer` command to verify whether the peer vPC switch is operational or not.
- All the devices that are attached to a vPC domain through a vPC must be dual homed.
- The first generation Broadcom based Nexus 9300 series switches and Nexus 9500 series line-cards does not support policy based routing route map with a `set ip next-hop` statement where the egress interface is the vPC Peer-Link while the vPC convergence TCAM region is allocated. This limitation does not apply to cloud scale based Nexus 9000 series devices such as Cisco Nexus 9200 switches, 9300 switches with EX/FX/FX2 line-cards and Nexus 9500 platform switches with 9700-EX/FX line-cards.
- Beginning with Cisco NX-OS Release 7.0(3)I5(1), Layer 3 over vPC is supported on Cisco Nexus 9000 Series switches for Layer 3 unicast communication only. Layer 3 over vPC is not supported for Layer 3 multicast traffic. For more information see the *Best Practices for Layer 3 and vPC Configuration* section
- The default behavior with Layer 3 peer-router and TTL=1 packet destined to IP of vPC peer is to punt packet to CPU and then forward the software to vPC peer. This is applicable to the Cloud Scale based EOR switches.
- Starting with Cisco NX-OS Release 7.0(3)I7(9) Cloud Scale based TOR switches can forward TTL=1 packet destined to vPC peer in hardware/data plane. It is recommended to use one of these releases or later releases for a seamless operation of the feature.

Best Practices for Layer 3 and vPC Configuration

This section describes best practices for using and configuring Layer 3 with vPC.

Layer 3 and vPC Configuration Overview

When a Layer 3 device is connected to a vPC domain through a vPC, it has the following views:


- At Layer 2, the Layer 3 device sees a unique Layer 2 switch presented by the vPC peer devices.
- At Layer 3, the Layer 3 device sees two distinct Layer 3 devices (one for each vPC peer device).

vPC is a Layer 2 virtualization technology, so at Layer 2, both vPC peer devices present themselves as a unique logical device to the rest of the network.

There is no virtualization technology at Layer 3, so each vPC peer device is seen as a distinct Layer 3 device by the rest of the network.

The following figure illustrates the two different Layer 2 and Layer 3 views with vPC.

Figure 17: Different Views for vPC Peer Devices

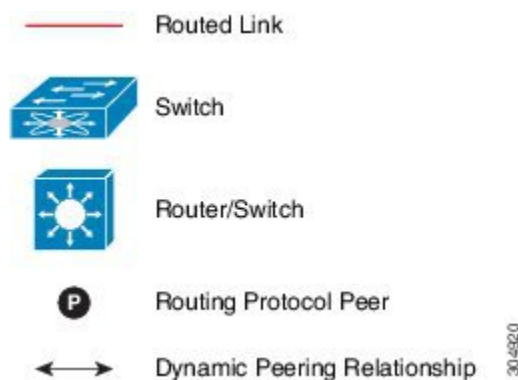


Supported Topologies for Layer 3 and vPC

This section contains examples of Layer 3 and vPC network topologies.

There are two approaches for Layer 3 and vPC interactions. The first one is by using dedicated Layer 3 links to connect the Layer 3 devices to each vPC peer device. The second one is by allowing the Layer 3 devices to peer with the SVIs defined on each of the vPC peer device, on a dedicated VLAN that is carried on the vPC connection. The following sections describe all the supported topologies leveraging the elements that are described in the legends in the following figure.

Figure 18: Legend




Peering with an External Router Using Layer 3 Links

This example shows a topology that uses Layer 3 links to connect a Layer 3 device to the Cisco Nexus 9000 switches that are part of the a vPC domain



Note Interconnecting the two entities together in this way allows to support Layer 3 unicast and multicast communication.

Figure 19: Peering with an External Router Using Layer 3 Links



Layer 3 devices can initiate Layer 3 routing protocol adjacencies with both vPC peer devices.

One or multiple Layer 3 links can be used to connect a Layer 3 device to each vPC peer device. Cisco Nexus 9000 series devices support Layer 3 Equal Cost Multipathing (ECMP) with up to 16 hardware load-sharing paths per prefix. Traffic from a vPC peer device to a Layer 3 device can be load-balanced across all the Layer 3 links interconnecting the two devices together.

Using Layer 3 ECMP on the Layer 3 device can effectively use all Layer 3 links from the device to the vPC domain. Traffic from a Layer 3 device to the vPC domain can be load-balanced across all the Layer 3 links interconnecting the two entities together.

Follow these guidelines when connecting a Layer 3 device to the vPC domain using Layer 3 links:


- Use separate Layer 3 links to connect Layer 3 devices to the vPC domain. Each link represents a point-to-point Layer 3 connection and should get assigned an IP address taken from a small IP subnet (/30 or /31).
- If the Layer 3 peering is required for multiple VRFs, it is recommended to define multiple sub-interfaces, each mapped to an individual VRF.

Peering Between vPC Devices for a Backup Routing Path

This example shows peering between the two vPC peer devices with a Layer 3 backup routed path. If the Layer 3 uplinks on vPC peer device 1 or vPC peer device 2 fail, the path between the two peer devices is used to redirect traffic to the switch that has the Layer 3 uplinks in the up state.

The Layer 3 backup routing path can be implemented using a dedicated interface VLAN (such as SVI) over the vPC Peer-Link or by using dedicated Layer 2 or Layer 3 links across the two vPC peer devices.


Figure 20: Peering Between vPC Devices for a Backup Routing Path



Direct Layer 3 Peering Between Routers

In this scenario, the Nexus 9000 devices part of the vPC domain are simply used as a Layer 2 transit path to allow the routers connected to them to establish Layer 3 peering and communication.

Figure 21: Peering Between Routers



The Layer 3 devices can peer with each other in following two methods. Peering also depends on the specific device deployed for this role.

- Defining a VLAN network interface (SVI) for a VLAN that is extended between the Layer 3 devices through the intermediate Cisco Nexus 9000 vPC peer switches.
- Defining a Layer 3 port-channel interface on each Layer 3 device and establishing a point-to-point Layer 3 peering.




Note In deployments where the Layer 3 peering must be established for multiple VRFs, the first method require the definition on the Layer 3 devices of a VLAN (and SVI) per VRF. For the second method, it is possible to create a Layer 3 port-channel subinterface per VRF

Peering Between Two Routers with vPC Devices as Transit Switches

This example is similar to the peering between routers topology. In this case also, the Cisco Nexus 9000 devices that are part of the same vPC domain are only used as Layer 2 transit paths. The difference here is that there are two pairs of Cisco Nexus 9000 switches. Each switch that is connected with a Layer 3 device using a vPC connection, also establishes a back-to-back vPC connection between them. The difference is that the vPC domains are only used as Layer 2 transit paths.

Figure 22: Peering Between Two Routers with vPC Devices as Transit Switches



This topology is commonly used when you want to establish connectivity between separate data centers that are interconnected with direct links (dark fibers or DWDM circuits). The two pairs of Cisco Nexus 9000 switches, in this case, provide only Layer 2 extension services, allowing the Layer 3 devices to peer with each other at Layer 3.

Peering with an External Router on Parallel Interconnected Routed Ports


When you require both routed and bridged traffic, use individual Layer 3 links for routed traffic and a separate Layer 2 port-channel for bridged traffic, as shown in following figure.

The Layer 2 links are used for bridged traffic (traffic staying in the same VLAN) or inter-VLAN traffic (assuming vPC domain hosts the interface VLAN and associated HSRP configuration).

The Layer 3 links are used for routing protocol peering adjacency with each vPC peer device.

The purpose of this topology is to attract specific traffic to go through the Layer 3 device. Layer 3 links are also used to carry routed traffic from a Layer 3 device to the vPC domain.


Figure 23: Peering with an External Router on Parallel Interconnected Routed Ports



Peering between vPC Switch Pairs on Parallel Interconnected Routed Ports

An alternative design to what is shown in the previous section (Peering Between Two Routers with vPC Devices as Transit Switches), uses two pairs of Cisco Nexus 9000 switches that are deployed in each data center for providing both Layer 2 and Layer 3 extension services. When routing protocol peering adjacency is required to be established between the two pairs of Cisco Nexus 9000 devices, the best practice is to add dedicated Layer 3 links between the two sites as shown in the following example.

Figure 24: Peering Over a vPC Interconnection on Parallel Interconnected Routed Ports



The back-to-back vPC connection between the two data centers carry bridged traffic or inter-VLAN traffic while the dedicated Layer 3 links carry the routed traffic across the two sites.


Peering Over a PC Interconnection and Dedicated Interswitch Link Using non-vPC VLAN

This example shows when the Layer 3 device is single-attached to the vPC domain, you can use a non-vPC VLAN with a dedicated inter-switch link to establish the routing protocol peering adjacency between the Layer 3 device and each vPC peer device. However, the non-vPC VLAN must be configured to use a static MAC that is different than the vPC VLAN.



Note Configuring the vPC VLAN (and vPC Peer-Link) for this purpose is not supported.

Figure 25: Peering Over a PC Interconnection and Dedicated Interswitch Link Using non-vPC VLAN




Peering Directly Over a vPC Connection

Beginning with Cisco NX-OS Release 7.0(3)I5(1), an alternative method has been introduced to establish Layer 3 peering between a Layer 3 router and a pair of Cisco Nexus 9000 vPC switches.



Note Peering directly over a vPC connection is supported only for Layer 3 unicast communication but not for Layer 3 multicast traffic. If you require Layer 3 multicast, you must establish peering over dedicated Layer 3 links

Figure 26: Supported: Peering Over a vPC Interconnection Where the Router Peers with Both the vPC Peers.




501182

In this scenario, the Layer 3 peering between the external router and the Cisco Nexus 9000 switches that are part of a same vPC domain is established directly on a VLAN carried on the vPC connection. The external router in this case peers with SVI interfaces defined on each vPC device. As for the scenario shown in previous figure 12, the external router could use an SVI or a Layer 3 Port-Channel to peer with the vPC devices (multiple SVIs or Port-Channel subinterfaces could be used for a multi-VRF deployment).

This deployment model requires configuring **layer3 peer-router** command as part of the vPC domain. You can adopt the same approach for establishing Layer 2 and Layer 3 connectivity on a vPC back-to-back connection established between two separate pairs of vPC switches.

Figure 27: Supported: Peering Over a vPC Interconnection Where Each Nexus Device Peers with Two vPC Peers.



In this deployment model, SVI interfaces in the same VLAN is configured on all the four Cisco Nexus 9000 switches to establish routing peering and connectivity between them.

Configuring Layer 3 over vPC

Beginning with Cisco NX-OS Release 7.0(3)I5(1), Layer 3 over vPC is supported on Cisco Nexus 9000 Series switches.

Before you begin

Ensure that the peer-gateway feature is enabled and it is configured on both the peers and both the peers run an image that supports Layer 3 over vPC. If you enter the **layer3 peer-router** command without enabling the peer-gateway feature, a syslog message is displayed recommending you to enable the peer-gateway feature.

Ensure that the vPC Peer-Link is up.

Procedure

	Command or Action	Purpose
Step 1	<pre>switch# configure terminal</pre> <p>Example:</p> <pre>switch# configure terminal switch(config)#</pre>	Enters global configuration mode.
Step 2	<pre>switch(config)# vpc domain domain-id</pre> <p>Example:</p> <pre>switch(config)# vpc domain 5 switch(config-vpc-domain)#</pre>	Creates a vPC domain if it does not already exist, and enters the vpc-domain configuration mode. There is no default; the range is from <1 to 1000>.

	Command or Action	Purpose
Step 3	switch(config-vpc-domain)# layer3 peer-router	Enables the Layer 3 device to form peering adjacency with both the peers. Note Configure this command in both the peers. If you configure this command only on one of the peers or you disable it on one peer, the operational state of layer 3 peer-router gets disabled. You get a notification when there is a change in the operational state.
Step 4	switch(config-vpc-domain)# exit	Exits the vpc-domain configuration mode.
Step 5	(Optional) switch# show vpc brief	Displays brief information about each vPC domain.
Step 6	(Optional) switch# copy running-config startup-config	Copies the running configuration to the startup configuration.

Example

The following example shows how to configure Layer 3 over vPC feature:

```
switch# configure terminal
switch(config)# vpc domain 5
switch(config-vpc-domain)# layer3 peer-router

switch(config-vpc-domain)# exit

switch(config)#
```

This example shows how to verify if the Layer 3 over vPC feature is configured. The **Operational Layer3 Peer** is enabled or disabled depending up on how the operational state of Layer 3 over vPC is configured.

```
switch# show vpc brief

vPC domain id : 5
Peer status : peer adjacency formed ok
vPC keep-alive status : peer is alive
Configuration consistency status : success
Per-vlan consistency status : failed
Type-2 consistency status : success
vPC role : secondary
Number of vPCs configured : 2
Peer Gateway : Enabled
Peer gateway excluded VLANs : -
Dual-active excluded VLANs : -
Graceful Consistency Check : Enabled
Auto-recovery status : Enabled (timeout = 240 seconds)
Operational Layer3 Peer : Enabled
```


Default Settings

The following table lists the default settings for vPC parameters.

Table 14: 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 vPCs



Note You must use these procedures on both devices on both sides of the vPC Peer-Link. You configure both of the vPC peer devices using these procedures.

This section describes how to configure vPCs using the command-line interface (CLI).



Note If you are familiar with the Cisco IOS CLI, be aware that the Cisco NX-OS commands for this feature might differ from the Cisco IOS commands that you would use.

Enabling vPCs

You must enable the feature vPC before you can configure and use vPCs.

Procedure

	Command or Action	Purpose
Step 1	configure terminal Example: switch# <code>configure terminal</code> switch(config)#	Enters global configuration mode.
Step 2	feature vpc Example:	Enables vPCs on the device.

	Command or Action	Purpose
	<code>switch(config)# feature vpc</code>	
Step 3	exit Example: <code>switch(config)# exit</code> <code>switch#</code>	Exits global configuration mode.
Step 4	show feature Example: <code>switch# show feature</code>	(Optional) Displays which features are enabled on the device.
Step 5	copy running-config startup-config Example: <code>switch# copy running-config startup-config</code>	(Optional) Copies the running configuration to the startup configuration.

Example

This example shows how to enable the vPC feature:

```
switch# configure terminal
switch(config)# feature vpc
switch(config)# exit
switch(config)#
```

Disabling vPCs



Note When you disable the vPC functionality, the device clears all the vPC configurations.

Procedure

	Command or Action	Purpose
Step 1	configure terminal Example: <code>switch# configure terminal</code> <code>switch(config)#</code>	Enters global configuration mode.
Step 2	no feature vpc Example: <code>switch(config)# no feature vpc</code>	Disables vPCs on the device.
Step 3	exit Example:	Exits global configuration mode.

	Command or Action	Purpose
	switch(config)# exit switch#	
Step 4	show feature Example: switch# show feature	(Optional) Displays which features are enabled on the device.
Step 5	copy running-config startup-config Example: switch# copy running-config startup-config	(Optional) Copies the running configuration to the startup configuration.

Example

This example shows how to disable the vPC feature:

```
switch# configure terminal
switch(config)# no feature vpc
switch(config)# exit
switch#
```

Creating a vPC Domain and Entering vpc-domain Mode

You can create a vPC domain and put the vPC Peer-Link port channels into the identical vPC domain on both vPC peer devices. Use a unique vPC domain number throughout a single vPC domain . This domain ID is used to automatically to form the vPC system MAC address.

You can also use this command to enter vpc-domain command mode.

Procedure

	Command or Action	Purpose
Step 1	configure terminal Example: switch# configure terminal switch(config)#	Enters global configuration mode.
Step 2	vpc domain domain-id [shut no shut] Example: switch(config)# vpc domain 5 switch(config-vpc-domain)#	Creates a vPC domain on the device, and enters vpc-domain configuration mode for configuration purposes. There is no default; the range is from 1 to 1000.
Step 3	exit Example: switch(config)# exit switch#	Exits vpc-domain configuration mode.

	Command or Action	Purpose
Step 4	show vpc brief Example: switch# show vpc brief	(Optional) Displays brief information about each vPC domain.
Step 5	copy running-config startup-config Example: switch# copy running-config startup-config	(Optional) Copies the running configuration to the startup configuration.

Example

This example shows how to enter the vpc-domain command mode to configure an existing vPC domain:

```
switch# configure terminal
switch(config)# vpc domain 5
switch(config-vpc-domain)# exit
switch(config)#
```

Configuring a vPC Keepalive Link and Messages

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.



Note You must configure the vPC peer-keepalive link before the system can form the vPC Peer-Link.



Note We recommend that you configure a separate VRF instance and put a Layer 3 port from each vPC peer device into that VRF for the vPC peer-keepalive link. Do not use the vPC Peer-Link itself to send vPC peer-keepalive messages. For information about creating and configuring VRFs, see the [Cisco Nexus 9000 Series NX-OS Unicast Routing Configuration Guide](#). Ensure that both the source and destination IP addresses used for the peer-keepalive message are unique in your network. The management port and management VRF are the defaults for these keepalive messages.

Before you begin

Ensure that you have enabled the vPC feature.

Procedure

	Command or Action	Purpose
Step 1	configure terminal Example: switch# configure terminal switch(config)#	Enters global configuration mode.
Step 2	vpc domain domain-id [shut no shut] Example: switch(config)# vpc domain 5 switch(config-vpc-domain)#	Creates a vPC domain on the device, and enters vpc-domain configuration mode.
Step 3	peer-keepalive destination ipaddress [hold-timeout secs interval msec {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}] Example: switch(config-vpc-domain)# peer-keepalive destination 172.28.230.85 switch(config-vpc-domain)#	Configures the IPv4 and IPv6 addresses 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. Note You may get the following error message if you do not specify the source IP address when you configure an IPv6 address for the remote end of the vPC peer-keepalive link. <pre>Cannot configure IPV6 peer-keepalive without source IPV6 address</pre> The management ports and VRF are the defaults. Note We recommend that you configure a separate VRF and use a Layer 3 port from each vPC peer device in that VRF for the vPC peer-keepalive link. For more information about creating and configuring VRFs, see the Cisco Nexus 9000 Series NX-OS Unicast Routing Configuration Guide .
Step 4	exit Example: switch(config)# exit switch#	Exits global configuration mode.

	Command or Action	Purpose
Step 5	show vpc statistics Example: switch# show vpc statistics	(Optional) Displays information about the configuration for the keepalive messages.
Step 6	copy running-config startup-config Example: switch# copy running-config startup-config	(Optional) Copies the running configuration to the startup configuration.

Example

For more information about configuring VRFs, see the [Cisco Nexus 9000 Series NX-OS Unicast Routing Configuration Guide](#).

This example shows how to configure the destination and source IP address and VRF for the vPC-peer-keepalive link:

```
switch# configure terminal
switch(config)# vpc domain 100
switch(config-vpc-domain)# peer-keepalive destination 172.168.1.2 source 172.168.1.1 vrf vpc-keepalive
switch(config-vpc-domain)# exit
switch#
```

Creating a vPC Peer-Link

You create the vPC Peer-Link by designating the port channel that you want on each device as the vPC Peer-Link for the specified vPC domain. We recommend that you configure the Layer 2 port channels 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 device for redundancy.

Before you begin

Ensure that you have enabled the vPC feature.

Procedure

	Command or Action	Purpose
Step 1	configure terminal Example: switch# configure terminal switch(config)#	Enters global configuration mode.

	Command or Action	Purpose
Step 2	interface port-channel <i>channel-number</i> Example: switch(config)# interface port-channel 20 switch(config-if)#	Selects the port channel that you want to use as the vPC Peer-Link for this device, and enters interface configuration mode.
Step 3	switchport mode trunk Example: switch(config-if)# switchport mode trunk	(Optional) Configures this interface in trunk mode.
Step 4	switchport trunk allowed vlan <i>vlan-list</i> Example: switch(config-if)# switchport trunk allowed vlan 1-120,201-3967	(Optional) Configures the permitted VLAN list.
Step 5	vpc peer-link Example: switch(config-if)# vpc peer-link switch(config-vpc-domain)#	Configures the selected port channel as the vPC Peer-Link, and enters vpc-domain configuration mode.
Step 6	exit Example: switch(config)# exit switch#	Exits vpc-domain configuration mode.
Step 7	show vpc brief Example: switch# show vpc brief	(Optional) Displays information about each vPC, including information about the vPC Peer-Link.
Step 8	copy running-config startup-config Example: switch# copy running-config startup-config	(Optional) Copies the running configuration to the startup configuration.

Example

This example shows how to configure a vPC Peer-Link:

```
switch# configure terminal
switch(config)# interface port-channel 20
switch(config-if)# switchport mode
switch(config-if)# switchport mode trunk
switch(config-if)# switchport trunk allowed vlan 1-120,201-3967
switch(config-if)# vpc peer-link
switch(config-vpc-domain)# exit
switch(config)#
```

Configuring a vPC Peer-Gateway

You can configure vPC peer devices to act as the gateway for packets that are destined to the vPC peer device's MAC address.

Before you begin

Ensure that you have enabled the vPC feature.

Procedure

	Command or Action	Purpose
Step 1	configure terminal Example: <pre>switch# configure terminal switch(config)#</pre>	Enters global configuration mode.
Step 2	vpc domain <i>domain-id</i> [shut no shut] Example: <pre>switch(config-if)# vpc domain 5 switch(config-vpc-domain)#</pre>	Creates a vPC domain if it does not already exist, and enters vpc-domain configuration mode.
Step 3	peer-gateway Example: <pre>switch(config-vpc-domain)# peer-gateway</pre> Note Disable IP redirects on all interface-vlans of this vPC domain for correct operation of this feature.	Enables Layer 3 forwarding for packets destined to the peer's gateway MAC address.
Step 4	exit Example: <pre>switch(config)# exit switch#</pre>	Exits vpc-domain configuration mode.
Step 5	show vpc brief Example: <pre>switch# show vpc brief</pre>	(Optional) Displays information about each vPC, including information about the vPC Peer-Link.
Step 6	copy running-config startup-config Example: <pre>switch# copy running-config startup-config</pre>	(Optional) Copies the running configuration to the startup configuration.

Configuring Fast Convergence

Beginning with Cisco NX-OS Release 7.0(3)I7(1), fast convergence feature is supported on Cisco Nexus 9000 Series platforms. You can enable or disable the vPC optimizations using this command. To achieve faster convergence, you must enable **[no] fast-convergence** on both vPC peers to achieve fast-convergence. The optimization is archived on secondary switch, vPC member ports and orphan ports with **vpc orphan-ports suspend** command is configured. In case the vPC Peer-Link fails, these ports will be suspended immediately and traffic will be forwarded to primary vPC peer only to improve convergence.

Procedure

	Command or Action	Purpose
Step 1	<code>configure terminal</code>	Enters global configuration mode.
Step 2	<code>switch(config) # vpc domain <domain></code>	Configure the VPC domain number.
Step 3	<code>switch(config) # peer-switch</code>	Define the peer switch.
Step 4	<code>switch(config) # show vpc peer-keepalive</code>	Displays information about the peer keepalive messages
Step 5	<code>switch(config) # delay restore { time }</code>	Number of seconds to delay bringing up the restored vPC peer device. The range is from 1 to 3600.
Step 6	<code>switch(config) # peer-gateway</code>	To enable Layer 3 forwarding for packets destined to the gateway MAC address of the virtual Port Channel (vPC), use the peer-gateway command. To disable Layer 3 forwarding packets, use the no form of this command.
Step 7	<code>switch(config) # delay restore orphan-port</code>	Number of seconds to delay bringing up the restored device's orphan port
Step 8	<code>switch(config-vpc-domain)# fast-convergence</code>	Configure vPC fast convergence.

Configuring LACP vPC Convergence

Beginning with Cisco NX-OS Release 7.0(3)I7(1), Link Aggregation Control Protocol (LACP) vPC convergence feature is supported on Cisco Nexus 9200 and 9300 Series Switches. You can configure LACP vPC convergence feature for more efficient use of port channels by reducing convergence time of vPC port channel for member link going down and first member bring up.

Beginning with Cisco NX-OS Release 7.0(3)I7(5), Link Aggregation Control Protocol (LACP) vPC convergence feature is supported on Cisco Nexus 9500 Series Switches with 9700-EX and 9700-FX line cards. This feature is not supported on Nexus 9500 with 9400, 9500, and 9600 and 9600-R line cards.

When you configure LACP vPC convergence on a Cisco Nexus 9000 switch, it waits until all the VLANs are initialized and programmed and then send LACP sync PDU, which will start sending traffic to the VPC domain

without drops. You may configure the **lacp vpc-convergence** command in a VXLAN and non-VXLAN environments that have vPC port-channels to hosts that support LACP.

Procedure

	Command or Action	Purpose
Step 1	configure terminal	Enters global configuration mode.
Step 2	switch(config) # interface <i>{type/slot portchannel number}</i>	Specifies an interface to configure, and enters interface configuration mode.
Step 3	switch(config-if) # lacp vpc-convergence	<p>Configure LACP convergence. Reduce the convergence time of the vPC port channel for member link going down and first member bring up.</p> <p>Note You must enable this command on both the vPC peer switches. This command must be configured only on PortFast ports (vPC port channels on which the spanning-tree port type edge [trunk] is enabled).</p> <p>Note In a vPC environment, when this command is not configured on vPC port-channel interfaces to devices that support LACP, and if one of the vPC peers is reloaded or if one of the links is brought up, the link(s) connected to the vPC peer switch in the 'up' state will remain active and forward traffic. The other link(s) may go down and will transition to the 'up' state. The links transitioning to the 'up' state starts initializing VLANs. When the VLANs are initialized, LACP sync PDUs are sent per initialized VLAN, which will bring port-channel into 'up' state and that leads to traffic black-hole for non-idealized VLANs.</p>

Configuring a Graceful Consistency Check

You can configure the graceful consistency check feature, which is enabled by default. Unless this feature is enabled, the vPC is completely suspended when a mismatch in a mandatory compatibility parameter is introduced in a working vPC. When this feature is enabled, only the links on the secondary peer device are

suspended. See the “Compatibility Parameters for vPC Interfaces” section for information about consistent configurations on the vPCs.

Procedure

	Command or Action	Purpose
Step 1	configure terminal Example: switch# configure terminal switch(config)#	Enters global configuration mode.
Step 2	vpc domain <i>domain-id</i> [shut no shut] Example: switch(config-if)# vpc domain 5 switch(config-vpc-domain)#	Creates a vPC domain if it does not already exist, and enters vpc-domain configuration mode.
Step 3	graceful consistency-check Example: switch(config-vpc-domain)# graceful consistency-check	Specifies that only the links on the secondary peer device are suspended when a mismatch is detected in a mandatory compatibility parameter. Use the no form of this command to disable the feature.
Step 4	exit Example: switch(config)# exit switch#	Exits vpc-domain configuration mode.
Step 5	show vpc brief Example: switch# show vpc brief	(Optional) Displays information on the vPCs.

Example

This example shows how to enable the graceful consistency check feature:

```
switch# configure terminal
switch(config)# vpc domain 5
switch(config-vpc-domain)# graceful consistency-check
switch(config-vpc-domain)# exit
switch(config)#
```

Checking the Configuration Compatibility on a vPC Peer-Link

After you have configured the vPC Peer-Link on both vPC peer devices, check that the configurations are consistent on all vPC interfaces. See the “Compatibility Parameters for vPC Interfaces” section for information about consistent configurations on the vPCs.

Procedure

	Command or Action	Purpose
Step 1	configure terminal Example: <pre>switch# configure terminal switch(config)#</pre>	Enters global configuration mode.
Step 2	show vpc consistency-parameters {global interface port-channel <i>channel-number</i>} Example: <pre>switch(config)# show vpc consistency-parameters global switch(config)#</pre>	(Optional) Displays the status of those parameters that must be consistent across all vPC interfaces.

Example

This example shows how to check that the required configurations are compatible across all the vPC interfaces:

```
switch# configure terminal
switch(config)# show vpc consistency-parameters global
switch(config)#
```



Note Messages regarding the vPC interface configuration compatibility are also logged to the syslog.

Moving Other Port Channels into a vPC

We recommend that you attach the vPC domain downstream port channel to two devices for redundancy.

To connect to the downstream device, you create a port channel from the downstream device to the primary vPC peer device and you create another port channel from the downstream device to the secondary peer device. On each vPC peer device, you assign a vPC number to the port channel that connects to the downstream device. You will experience minimal traffic disruption when you are creating vPCs.

Before you begin

Ensure that you have enabled the vPC feature.

Ensure that you are using a Layer 2 port channel.

Procedure

	Command or Action	Purpose
Step 1	configure terminal Example:	Enters global configuration mode.

	Command or Action	Purpose
	<pre>switch# configure terminal switch(config)#</pre>	
Step 2	<p>interface port-channel <i>channel-number</i></p> <p>Example:</p> <pre>switch(config)# interface port-channel 20 switch(config-if)#</pre>	Selects the port channel that you want to put into the vPC to connect to the downstream device, and enters interface configuration mode.
Step 3	<p>vpc <i>number</i></p> <p>Example:</p> <pre>switch(config-if)# vpc 5 switch(config-vpc-domain)#</pre>	<p>Configures the selected port channel into the vPC to connect to the downstream device. You can use any module in the device for these port channels. The range is from 1 and 4096.</p> <p>Note The vPC number that you assign to the port channel connecting to the downstream device from the vPC peer device must be identical on both vPC peer devices.</p>
Step 4	<p>exit</p> <p>Example:</p> <pre>switch(config)# exit switch#</pre>	Exits vpc-domain configuration mode.
Step 5	<p>show vpc brief</p> <p>Example:</p> <pre>switch# show vpc brief</pre>	(Optional) Displays information on the vPCs.
Step 6	<p>copy running-config startup-config</p> <p>Example:</p> <pre>switch# copy running-config startup-config</pre>	(Optional) Copies the running configuration to the startup configuration.

Example

This example shows how to configure a port channel to connect to the downstream device:

```
switch# configure terminal
switch(config)# interface port-channel 20
switch(config-if)# vpc 5
switch(config-if)# exit
switch(config)#
```

Manually Configuring a vPC Domain MAC Address

When you create a vPC domain, the Cisco NX-OS software automatically creates a vPC system MAC address, which is used for operations that are confined to the link-scope, such as LACP. However, you might choose to configure the vPC domain MAC address manually.

Before you begin

Ensure that you have enabled the vPC feature.

Procedure

	Command or Action	Purpose
Step 1	configure terminal Example: switch# configure terminal switch(config)#	Enters global configuration mode.
Step 2	vpc domain <i>domain-id</i> [shut no shut] Example: switch(config)# vpc domain 5 switch(config-vpc-domain)#	Enters the vPC domain number that you want to configure. The system enters vpc-domain configuration mode.
Step 3	system-mac <i>mac-address</i> Example: switch(config-vpc-domain)# system-mac 23fb.4ab5.4c4e switch(config-vpc-domain)#	Enters the MAC address that you want for the specified vPC domain in the following format: aaaa.bbbb.cccc.
Step 4	exit Example: switch(config-vpc-domain)# exit switch#	Exits vpc-domain configuration mode.
Step 5	show vpc role Example: switch# show vpc brief	(Optional) Displays the vPC system MAC address.
Step 6	copy running-config startup-config Example: switch# copy running-config startup-config	(Optional) Copies the running configuration to the startup configuration.

Example

This example shows how to manually configure a vPC domain MAC address:

```
switch# configure terminal
switch(config)# vpc domain 5
```

```
switch(config-vpc-domain) # system-mac 13gb.4ab5.4c4e
switch(config-vpc-domain) # exit
switch(config) #
```

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 devices are the primary devices on LACP. When you manually configure the system priority, ensure that you configure the same priority value on both vPC peer devices. If these values do not match, vPC does not come up.

Before you begin

Ensure that you have enabled the vPC feature.

Procedure

	Command or Action	Purpose
Step 1	configure terminal Example: switch# configure terminal switch(config) #	Enters global configuration mode.
Step 2	vpc domain <i>domain-id</i> [shut no shut] Example: switch(config) # vpc domain 5 switch(config-vpc-domain) #	Enters the vPC domain number that you want to configure. The system enters vpc-domain configuration mode.
Step 3	system-priority <i>priority</i> Example: switch(config-vpc-domain) # system-priority 4000 switch(config-vpc-domain) #	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	exit Example: switch(config-vpc-domain) # exit switch#	Exits vpc-domain configuration mode.
Step 5	show vpc role Example: switch# show vpc role	(Optional) Displays the vPC system priority.

	Command or Action	Purpose
Step 6	copy running-config startup-config Example: <pre>switch# copy running-config startup-config</pre>	(Optional) Copies the running configuration to the startup configuration.

Example

This example shows how to manually configure the vPC domain system priority:

```
switch# configure terminal
switch(config)# vpc domain 5
switch(config-vpc-domain)# system-priority 4000
switch(config-vpc-domain)# exit
switch(config)#
```

Manually Configuring the vPC Peer Device Role

By default, the Cisco NX-OS software elects a primary and secondary vPC peer device after you configure the vPC domain and both sides of the vPC Peer-Link. However, you might want to elect a specific vPC peer device as the primary device for the vPC. Then, you would manually configure the role value for the vPC peer device that you want as the primary device to be lower than the other vPC peer device.

vPCs do not support role preemption. If the primary vPC peer device fails, the secondary vPC peer device takes over to become operationally the vPC primary device. However, the original operational roles are not restored if the formerly primary vPC comes up again.

Before you begin

Ensure that you have enabled the vPC feature.

Procedure

	Command or Action	Purpose
Step 1	configure terminal Example: <pre>switch# configure terminal switch(config)#</pre>	Enters global configuration mode.
Step 2	vpc domain <i>domain-id</i> [shut no shut] Example: <pre>switch(config)# vpc domain 5 switch(config-vpc-domain)#</pre>	Enters the vPC domain number that you want to configure. The system enters vpc-domain configuration mode.
Step 3	role priority <i>priority</i> Example:	Enters the role priority that you want for the vPC system priority. The range of values is from 1 to 65636, and the default value is 32667.

	Command or Action	Purpose
	<pre>switch(config-vpc-domain)# role priority 4 switch(config-vpc-domain)#</pre>	A lower value means that this switch has a better chance of being the primary vPC.
Step 4	<p>exit</p> <p>Example:</p> <pre>switch(config)# exit switch#</pre>	Exits vpc-domain configuration mode.
Step 5	<p>show vpc role</p> <p>Example:</p> <pre>switch# show vpc role</pre>	(Optional) Displays the vPC system priority.
Step 6	<p>copy running-config startup-config</p> <p>Example:</p> <pre>switch# copy running-config startup-config</pre>	(Optional) Copies the running configuration to the startup configuration.

Example

This example shows how to manually configure the role priority of the vPC peer device:

```
switch# configure terminal
switch(config)# vpc domain 5
switch(config-vpc-domain)# role priority 4
switch(config-vpc-domain)# exit
switch(config)#
```

Configuring vPC Object Tracking Tracking Feature on a Single-Module vPC

If you must configure all the vPC Peer-Links and core-facing interfaces on a single module, you should configure a track object and a track list that is associated with the Layer 3 link to the core and on all the links on the vPC Peer-Link on both primary vPC peer devices. Once you configure this feature and if the primary vPC peer device fails, the system automatically suspends all the vPC links on the primary vPC peer device. This action forces all the vPC traffic to the secondary vPC peer device until the system stabilizes.

You must put this configuration on both vPC peer devices. Additionally, you should put the identical configuration on both vPC peer devices because either device can become the operationally primary vPC peer device.

Before you begin

Ensure that you have enabled the vPC feature.

Ensure that you have configured the track object and the track list. Ensure that you assign all interfaces that connect to the core and to the vPC Peer-Link to the track-list object on both vPC peer devices.

Procedure

	Command or Action	Purpose
Step 1	configure terminal Example: switch# configure terminal switch(config)#	Enters global configuration mode.
Step 2	vpc domain <i>domain-id</i> [shut no shut] Example: switch(config)# vpc domain 5 switch(config-vpc-domain)#	Enters the vPC domain number that you want to configure, and enters vpc-domain configuration mode.
Step 3	track <i>track-object-id</i> Example: switch(config-vpc-domain)# track object 23 switch(config-vpc-domain)#	Adds the previously configured track-list object with its associated interfaces to the vPC domain. See the Cisco Nexus 9000 Series NX-OS Unicast Routing Configuration Guide for information about configuring object tracking and track lists.
Step 4	exit Example: switch(config-vpc-domain)# exit switch#	Exits vpc-domain configuration mode.
Step 5	show vpc brief Example: switch# show vpc brief	(Optional) Displays information about the tracked objects.
Step 6	copy running-config startup-config Example: switch# copy running-config startup-config	(Optional) Copies the running configuration to the startup configuration.

Example

This example shows how to put the previously configured track-list object into the vPC domain on the vPC peer device:

```
switch# configure terminal
switch(config)# vpc domain 5
switch(config-vpc-domain)# track object 5
switch(config-vpc-domain)# exit
switch(config)#
```

Configuring for Recovery After an Outage

If an outage occurs, the vPC waits for a peer adjacency to form on a switch reload. This situation can result in an unacceptably long service disruption. You can configure the Cisco Nexus 9000 Series device to restore vPC services when its peer fails to come on line.

Configuring Reload Restore

The **reload restore** command and procedure described in this section is deprecated. We recommend that you use the **auto-recovery** command and procedure described in the “Configuring an Autorecovery” section.

You can configure the Cisco Nexus 9000 Series device to restore vPC services when its peer fails to come online by using the **reload restore** command.

Before you begin

Ensure that you have enabled the vPC feature.

Procedure

	Command or Action	Purpose
Step 1	configure terminal Example: switch# configure terminal switch(config)#	Enters global configuration mode.
Step 2	vpc domain <i>domain-id</i> [shut no shut] Example: switch(config)# vpc domain 5 switch(config-vpc-domain)#	Enters the vPC domain number that you want to configure, and enters vpc-domain configuration mode.
Step 3	reload restore [delay <i>time-out</i>] Example: switch(config-vpc-domain)# reload restore	Configures the vPC to assume its peer is not functional and to bring up the vPC. The default delay is 240 seconds. You can configure a time-out delay from 240 to 3600 seconds. Use the no form of the command to reset the vPC to its default settings.
Step 4	exit Example: switch(config-vpc-domain)# exit switch#	Exits vpc-domain configuration mode.
Step 5	show running-config vpc Example: switch# show running-config vpc	(Optional) Displays information about the vPC, specifically the reload status.

	Command or Action	Purpose
Step 6	show vpc consistency-parameters interface port-channel <i>number</i> Example: <pre>switch# show vpc consistency-parameters interface port-channel 1</pre>	(Optional) Displays information about the vPC consistency parameters for the specified interface.
Step 7	copy running-config startup-config Example: <pre>switch# copy running-config startup-config</pre>	(Optional) Copies the running configuration to the startup configuration. Note To ensure the reload feature is enabled, you should perform this step.

Example

This example shows how to set the vPC reload restore feature and save it in the switch startup configuration:

```
switch# configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
switch(config)# vpc domain 5
switch(config-vpc-domain)# reload restore
```

Warning:

Enables restoring of vPCs in a peer-detached state after reload, will wait for 240 seconds (by default) to determine if peer is un-reachable

```
switch(config-vpc-domain)# exit
switch(config)# exit
switch# copy running-config startup-config
switch# show running-config vpc
```

```
!Command: show running-config vpc
!Time: Wed Mar 24 18:43:54 2010
version 5.0(2)
feature vpc
logging level vpc 6
vpc domain 5
reload restore
```

This example shows how to examine the consistency parameters:

```
switch# show vpc consistency-parameters interface port-channel 1
```

Legend:

Type 1 : vPC will be suspended in case of mismatch
Name Type Local Value Peer Value

```
-----
STP Port Type 1 Default -
STP Port Guard 1 None -
STP MST Simulate PVST 1 Default -
mode 1 on -
Speed 1 1000 Mb/s -
Duplex 1 full -
Port Mode 1 trunk -
Native Vlan 1 1 -
```

```
MTU 1 1500 -
Allowed VLANs - 1-3967,4048-4093
Local suspended VLANs
```

Configuring an Autorecovery

You can configure the Cisco Nexus 9000 Series device to restore vPC services when its peer fails to come online by using the auto-recovery command.

You can configure the Cisco Nexus 9000 Series device to restore vPC services on the secondary vPC peer when its vPC primary peer fails and bringing down peer-keepalive and vPC Peer-Link, by using the **auto-recovery** command. In case of failure of primary switch where both peer-keepalive and vPC Peer-Links are down secondary switch will suspend vPC member. However, after 3 missed keepalive heartbeats secondary switch resumes the role of a primary switch and bring up vPC member ports. The **auto-recovery reload restore** command can be used in scenarios when vPC primary switch reloads, where secondary switch resumes the role of the vPC primary and bring ip VPC member ports.

Before you begin

Ensure that you have enabled the vPC feature.

Procedure

	Command or Action	Purpose
Step 1	configure terminal Example: switch# configure terminal switch(config)#	Enters global configuration mode.
Step 2	vpc domain <i>domain-id</i> [shut no shut] Example: switch(config)# vpc domain 5 switch(config-vpc-domain)#	Enters the vPC domain number that you want to configure, and enters vpc-domain configuration mode.
Step 3	auto-recovery [reload-delay <i>time</i>] Example: switch(config-vpc-domain)# auto-recovery	Configures the vPC to assume its peer is not functional and to bring up the vPC, and specifies the time to wait after a reload to restore the vPC. The default delay is 240 seconds. You can configure a delay from 240 to 3600 seconds. Use the no form of the command to reset the vPC to its default settings.
Step 4	exit Example: switch(config-vpc-domain)# exit switch#	Exits vpc-domain configuration mode.
Step 5	show running-config vpc Example: switch# show running-config vpc	(Optional) Displays information about the vPC, specifically the reload status.

	Command or Action	Purpose
Step 6	show vpc consistency-parameters interface port-channel <i>number</i> Example: <pre>switch# show vpc consistency-parameters interface port-channel 1</pre>	(Optional) Displays information about the vPC consistency parameters for the specified interface.
Step 7	copy running-config startup-config Example: <pre>switch# copy running-config startup-config</pre>	(Optional) Copies the running configuration to the startup configuration. Note To ensure the autorecovery feature is enabled, you should perform this step.

Example

This example shows how to set the vPC autorecovery feature and save it in the switch startup configuration:

```
switch# configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
switch(config)# vpc domain 5
switch(config-vpc-domain)# auto-recovery
switch(config-vpc-domain)# auto-recovery auto-recovery reload-delay 100
```

Warning:

Enables restoring of vPCs in a peer-detached state after reload, will wait for 240 seconds to determine if peer is un-reachable

```
switch(config-vpc-domain)# exit
switch(config)# exit
switch# copy running-config startup-config
```

Configuring the Suspension of Orphan Ports

When a device that is not vPC-capable connects to each peer, the connected ports are known as orphan ports because they are not members of a vPC. You can explicitly declare physical interfaces as orphan ports to be suspended (shut down) by the secondary peer when it suspends its vPC ports in response to a vPC Peer-Link or peer-keepalive failure. The orphan ports are restored when the vPC is restored.



Note You can configure vPC orphan port suspension only on physical ports, portchannels. However, you cannot configure the same on individual port channel member ports.

Before you begin

Ensure that you have enabled the vPC feature.

Procedure

	Command or Action	Purpose
Step 1	configure terminal Example: switch# configure terminal switch(config)#	Enters global configuration mode.
Step 2	show vpc orphan-ports Example: switch# show vpc orphan-ports	(Optional) Displays a list of the orphan ports.
Step 3	interface type slot/port Example: switch(config)# interface ethernet 3/1 switch(config-if)#	Specifies an interface to configure, and enters interface configuration mode.
Step 4	vpc orphan-port suspend Example: switch(config-if)# vpc orphan-ports suspend	Configures the selected interface as a vPC orphan port to be suspended by the secondary peer in the case of a vPC failure.
Step 5	exit Example: switch(config-if)# exit switch#	Exits interface configuration mode.
Step 6	copy running-config startup-config Example: switch# copy running-config startup-config	(Optional) Copies the running configuration to the startup configuration.

Example

This example shows how to configure an interface as a vPC orphan port to be suspended by the secondary peer in the case of a vPC failure:

```
switch# configure terminal
switch(config)# interface ethernet 3/1
switch(config-if)# vpc orphan-ports suspend
switch(config-if)# exit
switch(config)#
```

Configuring Delay Restore on an Orphan Port

You can configure **delay restore orphan-port** command on Cisco Nexus 9000 Series switches to configure a restore timer that delays the bringing up of restored device's orphan port starting from Cisco NX-OS Release 7.0(3)I7(1).



Note The delay restore orphan-port command applies only to interfaces that has vpc orphan-port suspend command configured. Other orphan ports may not delay bringing up devices.

Procedure

	Command or Action	Purpose
Step 1	<code>configure terminal</code>	Enters global configuration mode.
Step 2	<code>switch(config) # vpc domain <domain></code>	Configure the VPC domain number.
Step 3	<code>switch(config) # peer-switch</code>	Define the peer switch.
Step 4	<code>switch(config) # show vpc peer-keepalive</code>	Displays information about the peer keepalive messages
Step 5	<code>switch(config) # delay restore { time }</code>	Number of seconds to delay bringing up the restored vPC peer device. The range is from 1 to 3600.
Step 6	<code>switch(config) # peer-gateway</code>	To enable Layer 3 forwarding for packets destined to the gateway MAC address of the virtual Port Channel (vPC), use the peer-gateway command. To disable Layer 3 forwarding packets, use the no form of this command.
Step 7	<code>switch(config) # delay restore orphan-port</code>	Number of seconds to delay bringing up the restored device's orphan port

Configuring the vPC Peer Switch

You can configure the Cisco Nexus 9000 Series device to make a pair of vPC devices appear as a single STP root in the Layer 2 topology.

Configuring a Pure vPC Peer Switch Topology

You can configure a pure vPC peer switch topology by using the peer-switch command and then setting the best possible (lowest) spanning tree bridge priority value.

Before you begin

Ensure that you have enabled the vPC feature.



Note When using a non-VPC dedicated trunk link between the VPC peers, the non-VPC VLANs should have a different global priority on the peers to prevent STP from blocking the VLANs.

Procedure

	Command or Action	Purpose
Step 1	configure terminal Example: switch# configure terminal switch(config)#	Enters global configuration mode.
Step 2	vpc domain <i>domain-id</i> [shut no shut] Example: switch(config)# vpc domain 5 switch(config-vpc-domain)#	Enters the vPC domain number that you want to configure, and enters vpc-domain configuration mode.
Step 3	peer-switch Example: switch(config-vpc-domain)# peer-switch	Enables the vPC switch pair to appear as a single STP root in the Layer 2 topology. Use the no form of the command to disable the peer switch vPC topology.
Step 4	spanning-tree vlan <i>vlan-range</i> priority <i>value</i> Example: switch(config)# spanning-tree vlan 1 priority 8192	Configures the bridge priority of the VLAN. Valid values are multiples of 4096. The default value is 32768.
Step 5	exit Example: switch(config-vpc-domain)# exit switch#	Exits vpc-domain configuration mode.
Step 6	show spanning-tree summary Example: switch# show spanning-tree summary	(Optional) Displays a summary of the spanning tree port states including the vPC peer switch.
Step 7	copy running-config startup-config Example: switch# copy running-config startup-config	(Optional) Copies the running configuration to the startup configuration.

Example

This example shows how to configure a pure vPC peer switch topology:

```
switch# configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
switch(config)# vpc domain 5
switch(config-vpc-domain)# peer-switch
```

```
2010 Apr 28 14:44:44 switch %STP-2-VPC_PEERSWITCH_CONFIG_ENABLED: vPC peer-switch
configuration is enabled. Please make sure to configure spanning tree "bridge" priority as
per recommended guidelines to make vPC peer-switch operational.
```

```
switch(config-vpc-domain)# spanning-tree vlan 1 priority 8192
switch(config-vpc-domain)# exit
switch(config)#
```

Configuring a Hybrid vPC Peer Switch Topology

You can configure a hybrid vPC and non-vPC peer switch topology by using the spanning-tree pseudo-information command to change the designated bridge ID so that it meets the STP VLAN-based load-balancing criteria and then change the root bridge ID priority to a value that is better than the best bridge priority. You then enable the peer switch.

Before you begin

Ensure that you have enabled the vPC feature.

When using a non-VPC dedicated trunk link between the VPC peers, the non-VPC VLANs should have a different pseudo root priority on the peers to prevent STP from blocking the VLANs.

Procedure

	Command or Action	Purpose
Step 1	configure terminal Example: switch# configure terminal switch(config)#	Enters global configuration mode.
Step 2	spanning-tree pseudo-information Example: switch(config)# spanning-tree pseudo-information switch(config-pseudo)#	Configures the spanning tree pseudo information.
Step 3	vlan <i>vlan-id</i> designated priority <i>priority</i> Example: switch(config-pseudo)# vlan 1 designated priority 8192	Configures the designated bridge priority of the VLAN. Valid values are multiples of 4096 from 0 to 61440.
Step 4	vlan <i>vlan-id</i> root priority <i>priority</i> Example: switch(config-pseudo)# vlan 1 root priority 4096	Configures the root bridge priority of the VLAN. Valid values are multiples of 4096 from 0 to 61440.
Step 5	vpc domain <i>domain-id</i> [shut no shut] Example: switch(config)# vpc domain 5 switch(config-vpc-domain)#	Enters the vPC domain number that you want to configure, and enters vpc-domain configuration mode.
Step 6	peer-switch Example:	Enables the vPC switch pair to appear as a single STP root in the Layer 2 topology.

	Command or Action	Purpose
	<code>switch(config-vpc-domain) # peer-switch</code>	Use the no form of the command to disable the peer switch vPC topology.
Step 7	exit Example: <code>switch(config-vpc-domain) # exit</code> <code>switch#</code>	Exits vpc-domain configuration mode.
Step 8	show spanning-tree summary Example: <code>switch# show spanning-tree summary</code>	(Optional) Displays a summary of the spanning tree port states including the vPC peer switch.
Step 9	copy running-config startup-config Example: <code>switch# copy running-config startup-config</code>	(Optional) Copies the running configuration to the startup configuration.

Example

This example shows how to configure a hybrid vPC peer switch topology:

```
switch# configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
switch(config)# spanning-tree pseudo-information
switch(config-pseudo)# vlan 1 designated priority 8192
switch(config-pseudo)# vlan 1 root priority 4096
switch(config-pseudo)# vpc domain 5
switch(config-vpc-domain)# peer-switch
switch(config-vpc-domain)# exit
switch(config)#
```

Configuring Hitless vPC Role Change

Complete these steps to enable hitless vPC role change.

Before you begin

- Ensure that the vPC feature is enabled.
- Ensure that the vPC Peer-Link is up
- Verify the role priority of devices

Procedure

	Command or Action	Purpose
Step 1	vpc role preempt Example:	Enable hitless vPC role change.

	Command or Action	Purpose
	switch# vpc role preempt switch(config)#	
Step 2	show vpc role Example: switch(config)# show vpc role	(Optional) Verify hitless vPC role change feature.

Example

This example on how to configure hitless vPC role change:

```
switch# show vpc role
vPC Role status
-----
vPC role                : secondary
vPC system-mac          : 00:23:04:ee:be:01
vPC system-priority     : 32667
vPC local system-mac    : 8c:60:4f:03:84:41
vPC local role-priority : 32668
vPC peer system-mac     : 8c:60:4f:03:84:43
vPC peer role-priority  : 32667

! Configure vPC hitless role change on the device!

switch(config)# vpc role preempt
! The following is an output from the show vpc role command after the
vPC hitless feature is configured
switch(config)# show vpc role
vPC Role status
-----
vPC role                : primary
vPC system-mac          : 00:00:00:00:00:00
vPC system-priority     : 32667
vPC local system-mac    : 8c:60:4f:03:84:41
vPC local role-priority : 32666
vPC peer system-mac     : 8c:60:4f:03:84:43
vPC peer role-priority  : 32667

switch(config)#
```

Use Case Scenario for vPC Role Change

The hitless vPC role change feature can be used in the following scenarios:

- Role change request—When you want to change the roles of the peer devices in a vPC domain.
- Primary switch reload—When the devices comes up after a reload and roles are defined, you can use the hitless vPC role change feature to restore the roles. For example, after a reload if the primary device takes the role of operational secondary and the secondary device takes the role of primary operational, you can change the vPC peer roles to their original defined roles using the **vpc role preempt** command.



Note Always check the existing device role priority before switching vPC role.

- Dual-active recovery—In a dual-active recovery scenario, the vPC primary switch continues to be (operational) primary, but the vPC secondary switch becomes the targeted primary switch and keeps its vPC member ports up. You can use the vPC hitless feature and restore the device roles. After the Dual-active recovery, if one side is operational primary and the other side operational secondary, then you can use the **vpc role preempt** command to restore the device roles to be primary and secondary

Enabling STP to Use the Cisco MAC Address

This procedure enables STP to use the Cisco MAC address (00:26:0b:xx:xx:xx).

Before you begin

Ensure that you have enabled the vPC feature.

Procedure

	Command or Action	Purpose
Step 1	configure terminal Example: switch# configure terminal	Enters global configuration mode.
Step 2	vpc domain <i>domain-id</i> Example: switch(config)# vpc domain 5	Creates a vPC domain if it does not already exist, and enters vpc-domain configuration mode.
Step 3	[no] mac-address bpdud source version 2 Example: switch(config-vpc-domain)# mac-address bpdud source version 2	Enables STP to use the Cisco MAC address (00:26:0b:xx:xx:xx) as the source address of BPDUs generated on vPC ports.
Step 4	exit Example: switch(config-vpc-domain)# exit	Exits vpc-domain configuration mode.
Step 5	(Optional) copy running-config startup-config Example: switch(config)# copy running-config startup-config	Copies the running configuration to the startup configuration.

Verifying the vPC Configuration

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

Command	Purpose
show feature	Displays whether the vPC is enabled or not.

Command	Purpose
show vpc brief	Displays brief information about the vPCs.
show vpc consistency-parameters	Displays the status of those parameters that must be consistent across all vPC interfaces.
show running-config vpc	Displays running configuration information for vPCs.
show port-channel capacity	Displays how many port channels are configured and how many are still available on the device.
show vpc statistics	Displays statistics about the vPCs.
show vpc peer-keepalive	Displays information about the peer-keepalive messages.
show vpc role	Displays the peer status, the role of the local device, the vPC system MAC address and system priority, and the MAC address and priority for the local vPC device.

Monitoring vPCs

Use the **show vpc statistics** command to display vPC statistics.




Note This command displays the vPC statistics only for the vPC peer device that you are working on.

Configuration Examples for vPCs

The following example shows how to configure vPC on device A as shown in the figure:

Figure 28: vPC Configuration Example



1. Enable vPC and LACP.

```
switch# configure terminal
switch(config)# feature vPC
switch(config)# feature lACP
```

2. (Optional) Configure one of the interfaces that you want to be a vPC Peer-Link in the dedicated port mode.

```
switch(config)# interface ethernet 7/1,
ethernet 7/3, ethernet 7/5. ethernet 7/7
switch(config-if)# shutdown
switch(config-if)# exit
switch(config)# interface ethernet 7/1
```

```
switch(config-if)# no shutdown
switch(config-if)# exit
switch(config)#
```

3. (Optional) Configure the second, redundant interface that you want to be a vPC Peer-Link in the dedicated port mode.

```
switch(config)# interface ethernet 7/2, ethernet 7/4,
ethernet 7/6. ethernet 7/8
switch(config-if)# shutdown
switch(config-if)# exit
switch(config)# interface ethernet 7/2
```

```
switch(config-if)# no shutdown
switch(config-if)# exit
switch(config)#
```

4. Configure the two interfaces (for redundancy) that you want to be in the vPC Peer-Link to be an active Layer 2 LACP port channel.

```
switch(config)# interface ethernet 7/1-2
switch(config-if)# switchport
```

```
switch(config-if) # switchport mode trunk
switch(config-if) # switchport trunk allowed vlan 1-50
switch(config-if) # switchport trunk native vlan 20
switch(config-if) # channel-group 20 mode active
switch(config-if) # exit
```

5. Create and enable the VLANs.

```
switch(config) # vlan 1-50
switch(config-vlan) # no shutdown
switch(config-vlan) # exit
```

6. Create a separate VRF for the vPC peer-keepalive link and add a Layer 3 interface to that VRF.

```
switch(config) # vrf context pkal
switch(config-vrf) # exit
switch(config) # interface ethernet 8/1
switch(config-if) # vrf member pkal
switch(config-if) # ip address 172.23.145.218/24
switch(config-if) # no shutdown
switch(config-if) # exit
```

7. Create the vPC domain and add the vPC peer-keepalive link.

```
switch(config) # vpc domain 1
switch(config-vpc-domain) # peer-keepalive
switch(config-vpc-domain) # destination 172.23.145.217 source 172.23.145.218 vrf pkal
switch(config-vpc-domain) # exit
```

8. Configure the vPC vPC Peer-Link.

```
switch(config) # interface port-channel 20
switch(config-if) # switchport mode trunk
switch(config-if) # switchport trunk allowed vlan 1-50
switch(config-if) # vpc peer-link
switch(config-if) # exit
switch(config) #
```

9. Configure the interface for the port channel to the downstream device of the vPC.

```
switch(config) # interface ethernet 7/9
switch(config-if) # switchport mode trunk
switch(config-if) # allowed vlan 1-50
switch(config-if) # native vlan 20
switch(config-if) # channel-group 50 mode active
switch(config-if) # exit
switch(config) # interface port-channel 50
switch(config-if) # vpc 50
switch(config-if) # exit
switch(config) #
```

10. Save the configuration.

```
switch(config) # copy running-config startup-config
```



Note If you configure the port channel first, ensure that it is a Layer 2 port channel.

Related Documents

Related Topic	Related Topic
System management	System management
High availability	High availability
Release Notes	Release Notes



CHAPTER 9

Configuring IP Tunnels

This chapter describes how to configure IP tunnels using Generic Route Encapsulation (GRE) on Cisco NX-OS devices.

- [Information About IP Tunnels, on page 287](#)
- [Prerequisites for IP Tunnels, on page 289](#)
- [Guidelines and Limitations, on page 289](#)
- [Default Settings, on page 293](#)
- [Configuring IP Tunnels, on page 294](#)
- [Verifying the IP Tunnel Configuration, on page 306](#)
- [Configuration Examples for IP Tunneling, on page 306](#)
- [Related Documents, on page 307](#)

Information About IP Tunnels

IP tunnels can encapsulate a same-layer or higher layer protocol and transport the result over IP through a tunnel created between two devices.

IP Tunnel Overview

IP tunnels consists of the following three main components:

- **Passenger protocol**—The protocol that needs to be encapsulated. IPv4 is an example of a passenger protocol.
- **Carrier protocol**—The protocol that is used to encapsulate the passenger protocol. Cisco NX-OS supports GRE as a carrier protocol.
- **Transport protocol**—The protocol that is used to carry the encapsulated protocol. IPv4 is an example of a transport protocol. An IP tunnel takes a passenger protocol, such as IPv4, and encapsulates that protocol within a carrier protocol, such as GRE. The device then transmits this carrier protocol over a transport protocol, such as IPv4.

You configure a tunnel interface with matching characteristics on each end of the tunnel.


You must enable the tunnel feature before you can configure it. The system automatically takes a checkpoint prior to disabling the feature, and you can roll back to this checkpoint. See the [Cisco Nexus 9000 Series NX-OS System Management Configuration Guide](#) for information about rollbacks and checkpoints.

GRE Tunnels

You can use generic routing encapsulation (GRE) as the carrier protocol for a variety of passenger protocols.

The following figure shows the IP tunnel components for a GRE tunnel. The original passenger protocol packet becomes the GRE payload and the device adds a GRE header to the packet. The device then adds the transport protocol header to the packet and transmits it.

Figure 29: GRE PDU



Point-to-Point IP-in-IP Tunnel Encapsulation and Decapsulation

For 6.1(2)I3(4), 7.0(3)I1(2), and later, point-to-point IP-in-IP encapsulation and decapsulation is a type of tunnel that you can create to send encapsulated packets from a source tunnel interface to a destination tunnel interface. This type of tunnel will carry both inbound and outbound traffic.



Note The selection of GRE or IP-in-IP tunnel destination based on the PBR policy is not supported.

Multi-Point IP-in-IP Tunnel Decapsulation

For 6.1(2)I3(4), 7.0(3)I1(2), and later, multi-point IP-in-IP decapsulate-any is a type of tunnel that you can create to decapsulate packets from any number of IP-in-IP tunnels to one tunnel interface. This tunnel will not carry any outbound traffic. However, any number of remote tunnel endpoints can use a tunnel configured this way as their destination.

Path MTU Discovery

Path maximum transmission unit (MTU) discovery (PMTUD) prevents fragmentation in the path between two endpoints by dynamically determining the lowest MTU along the path from the packet's source to its destination. PMTUD reduces the send MTU value for the connection if the interface receives information that the packet would require fragmentation.

When you enable PMTUD, the interface sets the Don't Fragment (DF) bit on all packets that traverse the tunnel. If a packet that enters the tunnel encounters a link with a smaller MTU than the MTU value for the packet, the remote link drops the packet and sends an ICMP message back to the sender of the packet. This message indicates that fragmentation was required (but not permitted) and provides the MTU of the link that dropped the packet.



Note PMTUD on a tunnel interface requires that the tunnel endpoint can receive ICMP messages generated by devices in the path of the tunnel. Check that ICMP messages can be received before using PMTUD over firewall connections.

High Availability

IP tunnels support stateful restarts. A stateful restart occurs on a supervisor switchover. After the switchover, Cisco NX-OS applies the runtime configuration after the switchover.

Prerequisites for IP Tunnels

IP tunnels have the following prerequisites:

- You must be familiar with TCP/IP fundamentals to configure IP tunnels.
- You are logged on to the switch.
- You must enable the tunneling feature in a device before you can configure and enable any IP tunnels.

Guidelines and Limitations

IP tunnels have the following configuration guidelines and limitations:

- Guidelines for **source-direct** and **ipv6ipv6-decapsulate-any** options for tunnels:
 - Beginning with Cisco NX-OS Release 7.0(3)I7(3), the **source-direct** command is supported on the Cisco Nexus 9500 platform switches with the Application Spine Engine (ASE) and the Leaf Spine Engine (LSE).

Cisco Nexus 9500 platform switches with the Network Forwarding Engine (NFE) do not support the **tunnel source direct** command.
 - The **tunnel source direct** command with the **tunnel mode ipv6ipv6 decapsulate-any** command on the Cisco Nexus 9500 platform switches is only supported in the MPLS heavy routing template.
- The IP tunnel supports the **tunnel source** CLI command with interface, IPv4 address, IPv6 address, or IPv4 prefix. Beginning with Cisco NX-OS Release 7.0(3)I6(1), you can configure IP-in-IP tunnel decapsulation on directly connected IP addresses (for example, physical interface, port-channel, loopback, and SVI) using the new **tunnel source direct** CLI command. You can select the IP ECMP links when there are multiple IP links between the two switches. A single tunnel interface can decapsulate the tunneled packets whose outer destination IP is any of the IPv4 or IPv6 address that is locally configured and it is operationally *Up* in the switch.
- Currently, **tunnel mode ipip decapsulate-any** is supported for decapsulating IPv4 payload over IPv4 transport (IPv4inIPv4 packets). Beginning with Cisco NX-OS Release 7.0(3)I6(1), the new CLI **tunnel mode ipv6ipv6 decapsulate-any** command is introduced to support IPv6 payload over IPv6 transport (IPv6inIPv6 packets).

- The **tunnel source direct** and **tunnel mode ipv6ipv6 decapsulate-any** CLI commands are not supported on Cisco Nexus 9500 platform switches with the Network Formation Engine (NFE).
 - The **tunnel source direct** CLI command is supported only when an administrator uses the IP-in-IP decapsulation to source route the packets through the network. The source-direct tunnel is always operationally *Up* unless it is administratively shut down. The directly connected interfaces are identified using the **show ip route direct** CLI command.
 - The **tunnel source direct** CLI command is supported only on decapsulate-any tunnel modes, for example, **tunnel mode ipip decapsulate-any** and **tunnel mode ipv6ipv6 decapsulate-any**.
 - Auto-recovery for source-direct is not supported.
 - For **ipv6ipv6 decapsulate-any**, inter-VRF is not supported. The tunnel interface VRF (iVRF) and tunnel transport or forwarding VRF (fVRF) must be the same. Only one decapsulate-any tunnel (irrespective of VRF) can be present in Cisco Nexus 9200 and 9300-EX Series switches.
 - To enable IPv6 on ipv6ipv6 decap-any tunnel interface, configure a valid IPv6 address or configure **ipv6 address use-link-local-only** under the tunnel interface.
- See the following hardware limitations on the maximum sources that can be accommodated on a source direct tunnel and the related behavior:
 - Source direct tunnel is now supported for Cisco Nexus 9000 Series switches with Network Forwarding Engine (NFE), Application Spine Engine (ASE), and Leaf Spine Engine (LSE). Most of the limitations are only in case of scaled SIP (number of total IP/IPv6 addresses on the interfaces (L3, sub-interface, PC, PC-sub interfaces, loopback, SVI, and any secondary IP/IPv6 addresses.)

See the following sample use cases.

- Use Case 1: Non-deterministic behavior of which SIP gets installed if the number of IP/IPv6 interface scale is more.

Both the switches have 512 entries for tunnel SIP. With tunnel source, direct any IP or IPv6 address w.r.t **ipip or ipv6ipv6 decap any** with tunnel source gets installed in the above table.

The insertion of these entries is on a first come first serve basis without any CLI command to control which interface IP addresses get installed. If the system has more number of IP/IPv6 interfaces to be installed, the behavior is non-deterministic (The behavior can change across reload with interface flaps.)

- Use Case 2: The scale numbers are different in both switches and each has its own advantages and disadvantages.

IPv4 individual scale can be more (up to 512) in case of switches with NFE but it is shared with IPv6. In the switches with ASE and LSE, the IPv4 individual scale can be 256 but it is not shared with IPv6.

Whenever the tunnel decap table gets filled, the TABLE_FULL error is displayed. If some entry gets deleted after the table gets full, the table full error is cleared.

Table 15: Scale Numbers

Commands	Switches with NFE: Table size 512, v4 takes 1 entry, v6 takes 4 entries	Switches with ASE and LSE: Table size 512, v4 takes 1 entry, v6 takes 2 entries (paired index)
IPIP decap any with tunnel source direct	Shared between v4 and v6, v6 takes 4 entries $v4 + 4 * v6 = 512$ Maximum entries can be 512 with no v6 entries	Dedicated 256
IPv6IPv6 decap any with tunnel source direct	Shared between v4 and v6, v6 takes 4 entries $v4 + 4 * v6 = 512$ Maximum entries can be 128 with no v4 entries	Dedicated 128

- Use Case 3: Auto-recovery is not supported.

If any entry does not get installed in the hardware due to exhaustion of above table, removal of an already installed IP/IPv6 from interfaces does not automatically trigger the addition of the failed SIP in the table though the table has space now. You need to flap the tunnel interface or IP interface to get them installed.

However, if an entry does not get installed in the hardware due to a duplicate entry (if there was already a **decap-any** with one source present and now the **source direct tunnel** CLI command is configured, there is a duplicate entry for the prior source configured) that was taken care of by removing the entry only when both the tunnels get deleted.

- For Cisco Nexus 9000 Series switches with Network Forwarding Engine (NFE) and Application Spine Engine (ASE), the syslog is different as the dedicated IPv4 and IPv6 decap antries are carved in the syslog. If the **tunnel-decap-table** is full, the user gets a syslog as follows:

```
2017 Apr 6 12:18:04 switch %$ VDC-1 %$ %IPFIB-2-FIB_HW_IPV4_TUNNEL_DECAP_TABLE_FULL:
  IPv4 tunnel decap hardware table full.
IP tunnel decapsulation may not work for some GRE/IPinIP traffic
```

```
2017 Apr 6 12:18:11 switch %$ VDC-1 %$ %IPFIB-2-FIB_HW_IPV6_TUNNEL_DECAP_TABLE_FULL:
  IPv6 tunnel decap hardware table full.
IP tunnel decapsulation may not work for some GRE/IPinIP traffic
```

If the table is full and once some entry gets deleted from the table (due to an interface being operationally down or removal of IP address), the clear syslog for the table is displayed. Note that the deletion of the tunnel removes all the entries that are added as part of that tunnel.

```
2017 Apr 5 13:29:25 switch %$ VDC-1 %$
%IPFIB-2-FIB_HW_IPV4_TUNNEL_DECAP_TABLE_FULL_CLR: IPv4 tunnel decap hardware table
full exception cleared
```

```
2017 Apr 4 19:41:22 switch %$ VDC-1 %$
```

```
%IPFIB-2-FIB_HW_IPV6_TUNNEL_DECAP_TABLE_FULL_CLRD: IPv6 tunnel decap hardware table
full exception cleared
```

- Beginning with Cisco NX-OS Release 7.0(3)I6(1), IP-in-IP tunnel decapsulation is supported on IPv6 enabled networks.

```
!
interface tunnel 1
  ipv6 address use-link-local-only          <<< enable IPv6
  tunnel mode ipv6ip6 decapsulate-any
  tunnel source direct
  description IPinIP Decapsulation Interface
  mtu 1476
  no shutdown
```

- The **show** commands with the **internal** keyword are not supported.
- Cisco NX-OS supports only the following protocols:
 - IPv4 passenger protocol.
 - GRE carrier protocol.
- Cisco NX-OS supports the following maximum number of tunnels:
 - IP tunnels - 8 tunnels.
 - GRE and IP-in-IP regular tunnels - 8 tunnels. (7.0(3)I1(2) and later)
- IP tunnels do not support access control lists (ACLs) or QoS policies.
- Cisco NX-OS supports the GRE header defined in IETF RFC 2784. Cisco NX-OS does not support tunnel keys and other options from IETF RFC 1701.
- Cisco NX-OS does not support GRE tunnel keepalives.
- All unicast routing protocols are supported by IP tunnels.
- The IP tunnel interface cannot be configured to be a span source or destination.
- IP tunnels do not support PIM or other Multicast features and protocols. (7.0(3)I1(2) and later)
- The selection of GRE or IP-in-IP tunnel based on the PBR policy is not supported. (7.0(3)I1(2) and later)
- The **feature tunnel** feature on Cisco Nexus 9000 switches cannot co-exist with the VXLAN feature **feature nv overlay**.
- IP tunnels are supported only in the default **system routing** mode and not in other modes. (7.0(3)I1(2) and later)
- When configuring a tunnel interface to **ipip mode**, the maximum mtu value is 9192 (7.0(3)I2(2) and earlier) or 9196 (7.0(3)I3(1) and later).

When downgrading from NX-OS 7.0(3)I3(1) (or later release) to an earlier release, with a tunnel interface in **ipip mode** that has an mtu value of 9196, the mtu configuration is lost as a result of the downgrade operation. As a best practice, adjust the mtu value to 9192 before commencing the downgrade to avoid losing the mtu configuration.

- When configuring a tunnel interface to **ipip mode**, the default mtu value is 1476 (7.0(3)I3(1) and earlier) or 1480 (7.0(3)I4(1) and later).

When downgrading from NX-OS 7.0(3)I4(1) (or later release) to an earlier release, with a tunnel interface in **ipip mode** with no explicit mtu configuration, the mtu value changes as a result of the downgrade operation from 1480 to 1476. As a best practice, adjust the mtu value to 1476 before commencing the downgrade to avoid any changes to the mtu value.

When upgrading from NX-OS 7.0(3)I3(1) (or earlier release) to NX-OS 7.0(3)I4(1) (or later release), with a tunnel interface in **ipip mode** with no explicit mtu configuration, the mtu value changes as a result of the upgrade operation from 1476 to 1480. As a best practice, adjust the mtu value to 1480 before commencing the upgrade to avoid any changes to the mtu value.

- On Cisco Nexus 9200 Series switches, GRE packets that are received on an IP-in-IP tunnel are not dropped as expected and are instead forwarded to the packet destination.
- Tx packets originating from the switch, such as control pkts, are not included in Tx statistics.
- Tunnel destinations that are reachable over another tunnel are not supported.
- The consistency checker is not supported for routes over a tunnel.
- Non-IP routing protocols, such as isis, are not supported over IP-in-IP tunnels.
- RFC5549 is not supported over tunnels.
- BGP adjacency over tunnel is not supported in a scenario where the tunnel interface and tunnel source are in same VRF (example: VRF-A) and tunnel destination is reachable with route-leak from opposite end (example: via VRF-B)
- Configuring two tunnel interfaces with the same source and destination address is not supported. Loopback interfaces may be configured as the source addresses instead.
- Cisco Nexus 9200, 9300-EX, 9300-FX, 9300-FX2 series switches and Cisco Nexus 9500 platform switches with 9700-EX/FX line cards may not have multiple tunnel interfaces in a single VRF that are sourced from or destined to the same IP address. For example, a device may not have tunnel 0 and tunnel 1 interfaces in the default VRF that are sourced from the same IP address or interface.

Default Settings

The following table lists the default settings for IP tunnel parameters.

Table 16: Default IP Tunnel Parameters

Parameters	Default
Path MTU discovery age timer	10 minutes
Path MTU discovery minimum MTU	64
Tunnel feature	Disabled

Configuring IP Tunnels



Note If you are familiar with the Cisco IOS CLI, be aware that the Cisco NX-OS commands for this feature might differ from the Cisco IOS commands that you would use.

Enabling Tunneling

You must enable the tunneling feature before you can configure any IP tunnels.

Procedure

	Command or Action	Purpose
Step 1	configure terminal Example: switch# configure terminal switch(config)#	Enters global configuration mode.
Step 2	feature tunnel Example: switch(config)# feature tunnel switch(config-if)#	Allows the creation of a new tunnel interface. To disable the tunnel interface feature, use the no form of this command. Note The feature tunnel command may break the multicast functionality if multicast heavy template is enforced.
Step 3	exit Example: switch(config-if)# exit switch#	Exits the interface mode and returns to the configuration mode.
Step 4	show feature Example: switch(config-if)# show feature	(Optional) Displays information about the features enabled on the device.
Step 5	copy running-config startup-config Example: switch(config-if)# copy running-config startup-config	(Optional) Saves this configuration change.

Creating a Tunnel Interface

You can create a tunnel interface and then configure this logical interface for your IP tunnel.



Note Cisco NX-OS supports a maximum of 8 IP tunnels.



Note Use the **no interface tunnel** command to remove the tunnel interface and all associated configuration.

Command	Purpose
no interface tunnel <i>number</i> Example: <pre>switch(config)# no interface tunnel 1</pre>	Deletes the tunnel interface and the associated configuration.
description <i>string</i> Example: <pre>switch(config-if)# description GRE tunnel</pre>	Configures a description for the tunnel.
mtu <i>value</i> Example: <pre>switch(config-if)# mtu 1400</pre>	Sets the MTU of IP packets sent on an interface.
tunnel ttl <i>value</i> Example: <pre>switch(config-if)# tunnel ttl 100</pre>	Sets the tunnel time-to-live value. The range is from 1 to 255.

Before you begin

You can configure the tunnel source and the tunnel destination in different VRFs. Ensure that you have enabled the tunneling feature.

Procedure

	Command or Action	Purpose
Step 1	configure terminal Example: <pre>switch# configure terminal switch(config)#</pre>	Enters global configuration mode.
Step 2	interface tunnel <i>number</i> Example: <pre>switch(config)# interface tunnel 1 switch(config-if)#</pre>	Creates a new tunnel interface.
Step 3	tunnel mode {gre ip ipip {ip decapsulate-any}}	Sets this tunnel mode to GRE, ipip, or ipip decapsulate-only.

	Command or Action	Purpose
		The gre and ip keywords specify that GRE encapsulation over IP will be used. The ipip keyword specifies that IP-in-IP encapsulation will be used. The optional decapsulate-any keyword terminates IP-in-IP tunnels at one tunnel interface. This keyword creates a tunnel that will not carry any outbound traffic. However, remote tunnel endpoints can use a tunnel configured as their destination.
Step 4	tunnel source <i>{ip-address interface-name}</i> Example: <pre>switch(config-if)# tunnel source ethernet 1/2</pre>	Configures the source address for this IP tunnel. The source can be specified by IP address or logical interface name.
Step 5	tunnel destination <i>{ip-address host-name}</i> Example: <pre>switch(config-if)# tunnel destination 192.0.2.1</pre>	Configures the destination address for this IP tunnel. The destination can be specified by IP address or logical host name.
Step 6	tunnel use-vrf <i>vrf-name</i> Example: <pre>switch(config-if)# tunnel use-vrf blue</pre>	(Optional) Uses the configured VRF to look up the tunnel IP destination address.
Step 7	show interfaces tunnel <i>number</i> Example: <pre>switch# show interfaces tunnel 1</pre>	(Optional) Displays the tunnel interface statistics.
Step 8	copy running-config startup-config Example: <pre>switch(config-if)# copy running-config startup-config</pre>	(Optional) Saves this configuration change.

Example

This example shows how to create a tunnel interface

```
switch# configure terminal
switch(config)# interface tunnel 1
switch(config-if)# tunnel source ethernet 1/2
switch(config-if)# tunnel destination 192.0.2.1
switch(config-if)# copy running-config startup-config
```

Creating an IP-in-IP Tunnel with a Netmask

Creating an IP-in-IP tunnel with a netmask allows you to specify a tunnel source subnet and a tunnel destination subnet, and decap the packet if it matches. (7.0(3)I2(1) and later)

- The IP-in-IP decap-any tunnel receives encapsulated packets from any number of IP-in-IP tunnels.
- With the netmask feature, the switch receives packets from IP addresses which comply with the netmasks.

Notes for the netmask feature:

- Routing protocols are not supported on an IP-in-IP tunnel created with a netmask.
- Encap is not supported with the netmask feature; only decap from a set of sources in the same subnet is supported.

Procedure

	Command or Action	Purpose
Step 1	configure terminal Example: switch# configure terminal switch(config)#	Enters global configuration mode.
Step 2	interface tunnel <i>number</i> Example: switch(config)# interface tunnel 5 switch(config-if)#	Creates a new tunnel interface.
Step 3	tunnel mode ipip [<i>ip</i>]	Sets this tunnel mode to ipip. The ipip keyword specifies that IP-in-IP encapsulation will be used.
Step 4	tunnel source <i>ip-address / mask_length</i> Example: switch(config-if)# tunnel source 33.1.1.1 255.255.255.0	Configures the source address for this IP tunnel. The source is specified by IP address and length of mask.
Step 5	tunnel destination <i>ip-address / mask_length</i> Example: switch(config-if)# tunnel destination 33.1.1.2 255.255.255.0	Configures the destination address for this IP tunnel. The destination is specified by IP address and length of mask.
Step 6	(Optional) no shut	Clears the interface.
Step 7	ip address <i>ip-prefix/length</i> Example: switch(config-if)# ip address 50.1.1.1/24	Configures an IP address for this interface.

Example

The following example shows how to create an IP-in-IP tunnel with a netmask.

```

switch(config)# interface tunnel 10
switch(config-if)# tunnel mode ipip
switch(config-if)# tunnel source 33.1.1.2/24
switch(config-if)# tunnel destination 33.1.1.1/24
switch(config-if)# no shut
switch(config-if)# ip address 10.10.10.10/24
switch(config-if)# end
switch# show interface tunnel 10
Tunnel10 is up
  Admin State: up
  Internet address is 10.10.10.10/24
  MTU 1476 bytes, BW 9 Kbit
  Tunnel protocol/transport IPIP/IP
  Tunnel source 33.1.1.2, destination 33.1.1.1
  Transport protocol is in VRF "default"
  Last clearing of "show interface" counters never
  Tx
  0 packets output, 0 bytes
  Rx
  0 packets input, 0 bytes

switch# show run interface tunnel 10

!Command: show running-config interface Tunnel10
!Time: Wed Aug 26 13:50:01 2015

version 7.0(3)I2(1)

interface Tunnel10
 ip address 10.10.10.10/24
 tunnel mode ipip ip
 tunnel source 33.1.1.2 255.255.255.0
 tunnel destination 33.1.1.1 255.255.255.0
 no shutdown

```

Configuring a Tunnel Interface

You can set a tunnel interface to GRE tunnel mode, ipip mode, or ipip decapsulate-only mode. GRE mode is the default tunnel mode. (7.0(3)I1(2) and later).

Beginning with Cisco NX-OS Release 7.0(3)I6(1), the **tunnel source direct** and **tunnel mode ipv6ip6 decapsulate-any** CLI commands are supported on Cisco Nexus 9000 Series switches.



Note The **tunnel source direct** and **tunnel mode ipv6ip6 decapsulate-any** CLI commands are not supported on Cisco Nexus 9500 platform switches with the Network Forwarding Engine (NFE).

The new CLI **tunnel mode ipv6ip6 decapsulate-any** command is introduced to support IPv6 payload over IPv6 transport (IPv6inIPv6 packets). You can configure IP-in-IP tunnel decapsulation on directly connected IP addresses (for example, physical interface, port-channel, loopback, and SVI) using the new **tunnel source direct** CLI command.

Before you begin

Ensure that you have enabled the tunneling feature.

Procedure

	Command or Action	Purpose
Step 1	configure terminal Example: switch# configure terminal switch(config)#	Enters global configuration mode.
Step 2	interface tunnel <i>number</i> Example: switch(config)# interface tunnel 1 switch(config-if)#	Creates a new tunnel interface.
Step 3	tunnel mode {gre ip ipip {ip decapsulate-any}}	Sets this tunnel mode to GRE, ipip, or ipip decapsulate-only. The gre and ip keywords specify that GRE encapsulation over IP will be used. The ipip keyword specifies that IP-in-IP encapsulation will be used. The optional decapsulate-any keyword terminates IP-in-IP tunnels at one tunnel interface. This keyword creates a tunnel that will not carry any outbound traffic. However, remote tunnel endpoints can use a tunnel configured as their destination.
Step 4	(Optional) tunnel mode ipv6ipv6 decapsulate-any	Supports IPv6 payload over IPv6 transport (IPv6inIPv6 packets) (7.0(3)I6(1) and later). This step is applicable for IPv6 networks only. Note This command is not supported on Cisco Nexus 9500 platform switches.
Step 5	tunnel source direct	Configures IP-in-IP tunnel decapsulation on any directly connected IP addresses. Beginning with Cisco NX-OS Release 7.0(3)I6(1), this option is now supported only when the IP-in-IP decapsulation is used to source route the packets through the network. Note This command is not supported on Cisco Nexus 9500 platform switches with the Network Forwarding Engine (NFE).
Step 6	show interfaces tunnel <i>number</i> Example: switch(config-if)# show interfaces tunnel 1	(Optional) Displays the tunnel interface statistics.

	Command or Action	Purpose
Step 7	<code>mtu value</code>	<p>Sets the maximum transmission unit (MTU) of IP packets sent on an interface.</p> <p>The range is from 64 to 9192 units.</p> <p>Note When configuring tunnel mode ipip, the range is dependent on NX-OS release:</p> <ul style="list-style-type: none"> • 64 to 9192 units (7.0(3)I2(2) and earlier) • 64 to 9196 units (7.0(3)I3(1) and later)
Step 8	<p>copy running-config startup-config</p> <p>Example:</p> <pre>switch(config-if)# copy running-config startup-config</pre>	(Optional) Saves this configuration change.

Example

This example shows how to create the tunnel interface to GRE:

```
switch# configure terminal
switch(config)# interface tunnel 1
switch(config-if)# tunnel mode gre ip
switch(config-if)# copy running-config startup-config
```

This example shows how to create an ipip tunnel:

```
switch# configure terminal
switch(config)# interface tunnel 1
switch(config-if)# tunnel mode ipip
switch(config-if)# mtu 1400
switch(config-if)# copy running-config startup-config
switch(config-if)# no shut
```

This example shows how to configure IP-in-IP tunnel decapsulation on directly connected IP addresses:

```
switch# configure terminal
switch(config)# interface tunnel 0
switch(config-if)# tunnel mode ipip ip
switch(config-if)# tunnel source direct
switch(config-if)# description IPinIP Decapsulation Interface
switch(config-if)# no shut
```

This example shows how to configure IP-in-IP tunnel decapsulation on IPv6 enabled networks:

```
!
interface tunnel 1
  ipv6 address use-link-local-only          <<< enable IPv6
  tunnel mode ipv6ip6 decapsulate-any
  tunnel source direct
  description IPinIP Decapsulation Interface
  mtu 1476
```



```

no shutdown

show running-config interface tunnel 1
interface Tunnell
  tunnel mode ipv6ip6 decapsulate-any
  tunnel source direct
  no shutdown

show interface tunnel 1
Tunnell is up    Admin State: up
MTU 1460 bytes, BW 9 Kbit
Tunnel protocol/transport IPv6/DECAPANY/IPv6
Tunnel source - direct
Transport protocol is in VRF "default"
Tunnel interface is in VRF "default"
Last clearing of "show interface" counters never
Tx    0 packets output, 0 bytes    Rx    0 packets input, 0 bytes

```

Configuring a GRE Tunnel

You can set a tunnel interface to GRE tunnel mode. (6.1(2)I3(3) and 7.0(3)I1(1))



Note Cisco NX-OS supports only the GRE protocol for IPV4 over IPV4.

Before you begin

Ensure that you have enabled the tunneling feature.

Procedure

	Command or Action	Purpose
Step 1	configure terminal Example: switch# configure terminal switch(config)#	Enters global configuration mode.
Step 2	interface tunnel <i>number</i> Example: switch(config)# interface tunnel 1 switch(config-if)#	Creates a new tunnel interface.
Step 3	tunnel mode gre ip Example: switch(config-if)# tunnel mode gre ip	Sets this tunnel mode to GRE.
Step 4	show interfaces tunnel <i>number</i> Example: switch(config-if)# show interfaces tunnel 1	(Optional) Displays the tunnel interface statistics.

	Command or Action	Purpose
Step 5	copy running-config startup-config Example: <pre>switch(config-if)# copy running-config startup-config</pre>	(Optional) Saves this configuration change.

Configuring a GRE Tunnel

GRE v6 tunnel is used to carry different types of packets over IPv6 transport. GREv6 tunnel carries only IPv4 payload. The tunnel CLIs are enhanced to select IPv6 tunnel and configure v6 tunnel source and destination. (7.0(3)I2(1) and later)

You can set a tunnel interface to GRE tunnel mode, ipip mode, or ipip decapsulate-only mode. GRE mode is the default tunnel mode.

Before you begin

Ensure that you have enabled the tunneling feature.

Procedure

	Command or Action	Purpose
Step 1	switch# configure terminal	Enters global configuration mode.
Step 2	switch(config)# interface tunnel number	Enters a tunnel interface configuration mode.
Step 3	switch(config-if)# tunnel mode {gre ip ipip {ip decapsulate-any} }	<p>Sets this tunnel mode to GRE, ipip, or ipip decapsulate-only.</p> <p>The gre and ip keywords specify that GRE encapsulation over IP will be used.</p> <p>The ipip keyword specifies that IP-in-IP encapsulation will be used. The optional decapsulate-any keyword terminates IP-in-IP tunnels at one tunnel interface. This keyword creates a tunnel that will not carry any outbound traffic. However, remote tunnel endpoints can use a tunnel configured as their destination.</p>
Step 4	Required: switch(config-if)# tunnel use-vrf vrf-name	Configures tunnel VRF name.
Step 5	Required: switch(config-if)# ipv6 address IPv6 address	<p>Configures the IPv6 address.</p> <p>Note The tunnel source and the destination addresses are still the same (IPv4 address.)</p>
Step 6	(Optional) switch(config-if)# show interface tunnel number	Displays the tunnel interface statistics.

	Command or Action	Purpose
Step 7	switch(config-if)# <i>mtu value</i>	Sets the maximum transmission unit (MTU) of IP packets sent on an interface.
Step 8	(Optional) switch(config-if)# copy running-config startup-config	Saves the change persistently through reboots and restarts by copying the running configuration to the startup configuration.

Example

This example displays how to configure IPv6 Payload over GRE v4 tunnel. Configure the tunnel source, destination, IPv4 address, IPv6 address, and perform the **no shut** command. Once the GREv4 tunnel is created, you can configure v4 or v6 route via the tunnel:

```
switch# configure terminal
switch(config)# interface tunnel 10
switch(config)# tunnel source 11.1.1.1
switch(config)# tunnel destination 11.1.1.2
switch(config-if)# tunnel mode gre ip
switch(config-if)# tunnel use-vrf red
switch(config-if)# ip address 10.1.1.1/24
switch(config-if)# ipv6 address 2:2::2/64
switch(config-if)# no shut

switch(config)# ip route 50.1.1.0/24 tunnel 10
switch(config)# ipv6 route 2000:100::/64 tunnel 10
```

This example shows how to view the GRE v4 tunnel interface 10 and display IPv4 and IPv6 routes:

```
switch(config)# show int tunnel 10
Tunnel 10 is up
  Admin State: up
  Internet address(es):
    10.1.1.1/24
    1010::1/64
  MTU 1476 bytes, BW 9 Kbit
  Tunnel protocol/transport GRE/IP
  Tunnel source 11.1.1.1, destination 11.1.1.2
  Transport protocol is in VRF "default"

switch#show ipv6 route
...
2000:100::/64, ubest/mbest: 1/0, attached
  *via Tunnel10, [1/0], 00:00:16, static

#show ip route
...
50.1.1.0/24, ubest/mbest: 1/0
  *via Tunnel10, [1/0], 00:03:33, static
```

This example displays how to configure IPv4 payload over GRE v6 tunnel. Configure the tunnel mode as GRE IPv6, tunnel v6 source and destination, IPv4 address, and perform the **no shut** command. Once the GREv6 tunnel is created, you can configure v4 route via the tunnel:

```
switch# configure terminal
switch(config)# interface tunnel 20
switch(config-if)# tunnel mode gre ipv6
```

```

switch(config)# tunnel source 1313::1
switch(config)# tunnel destination 1313::2
switch(config-if)# tunnel use-vrf red
switch(config-if)# ip address 20.1.1.1/24
switch(config-if)# no shut

switch(config)# ip route 100.1.1.0/24 tunnel 20

```

This example displays how to view the GREv6 tunnel interface 20:

```

show interface tunnel 20
Tunnel 20 is up
  Admin State: up
  Internet address is 20.1.1.1/24
  MTU 1456 bytes, BW 9 Kbit
  Tunnel protocol/transport GRE/IPv6
  Tunnel source 1313::1, destination 1313::2
  Transport protocol is in VRF "default"

```

```

#show ip route
...
100.1.1.0/24, ubest/mbest: 1/0
  *via Tunnel20, [1/0], 00:01:00, static

```

```

red10# show interface brief | grep Tunnel
Tunnel10          up          10.1.1.1/24      GRE/IP          1476
Tunnel20          up          20.1.1.1/24      GRE/IPv6        1456

```

This example shows how to create an ipip tunnel:

```

switch# configure terminal
switch(config)# interface tunnel 1
switch(config-if)# tunnel mode ipip
switch(config-if)# mtu 1400
switch(config-if)# copy running-config startup-config
switch(config-if)# no shut

```

Enabling Path MTU Discovery

Use the `tunnel path-mtu discovery` command to enable path MTU discovery on a tunnel.

Procedure

	Command or Action	Purpose
Step 1	tunnel path-mtu-discovery age-timer <i>min</i> Example: <pre>switch(config-if)# tunnel path-mtu-discovery age-timer 25</pre>	Enables Path MTU Discovery (PMTUD) on a tunnel interface. <ul style="list-style-type: none"> • <code>min</code>—Number of minutes. The range is from 10 to 30. The default is 10.
Step 2	tunnel path-mtu-discovery min-mtu <i>bytes</i> Example: <pre>switch(config-if)# tunnel path-mtu-discovery min-mtu 1500</pre>	Enables Path MTU Discovery (PMTUD) on a tunnel interface. <ul style="list-style-type: none"> • <code>bytes</code>—Minimum MTU recognized. The range is from 64 to 9192. The default is 64.

Assigning VRF Membership to a Tunnel Interface

You can add a tunnel interface to a VRF.

Before you begin

Ensure that you have enabled the tunneling feature.

Assign the IP address for a tunnel interface after you have configured the interface for a VRF.

Procedure

	Command or Action	Purpose
Step 1	configure terminal Example: switch# configure terminal switch(config)#	Enters global configuration mode.
Step 2	interface tunnel <i>number</i> Example: switch(config)# interface tunnel 0 switch(config-if)#	Enters interface configuration mode.
Step 3	vrf member <i>vrf-name</i> Example: switch(config-if)# vrf member RemoteOfficeVRF	Adds this interface to a VRF.
Step 4	ip address <i>ip-prefix/length</i> Example: switch(config-if)# ip address 192.0.2.1/16	Configures an IP address for this interface. You must do this step after you assign this interface to a VRF.
Step 5	show vrf [<i>vrf-name</i>] interface <i>interface-type number</i> Example: switch(config-vrf)# show vrf Enterprise interface tunnel 0	(Optional) Displays VRF information.
Step 6	copy running-config startup-config Example: switch# copy running-config startup-config	(Optional) Saves this configuration change.

Example

This example shows how to add a tunnel interface to the VRF:

```

switch# configure terminal
switch(config)# interface tunnel 0
switch(config-if)# vrf member RemoteOfficeVRF
switch(config-if)# ip address 209.0.2.1/16
switch(config-if)# copy running-config startup-config

```

Verifying the IP Tunnel Configuration

To verify the IP tunnel configuration information, perform one of the following tasks:

Command	Purpose
<code>show interface tunnel <i>number</i></code>	Displays the configuration for the tunnel interface (MTU, protocol, transport, and VRF). Displays input and output packets, bytes, and packet rates.
<code>show interface tunnel <i>number</i> brief</code>	Displays the operational status, IP address, encapsulation type, and MTU of the tunnel interface.
<code>show interface tunnel <i>number</i> counters</code>	Displays interface counters of input/output packets. Note The byte count displayed with the interface counters include the internal header size.
<code>show interface tunnel <i>number</i> description</code>	Displays the configured description of the tunnel interface.
<code>show interface tunnel <i>number</i> status</code>	Displays the operational status of the tunnel interface.
<code>show interface tunnel <i>number</i> status err-disabled</code>	Displays the error disabled status of the tunnel interface.

Configuration Examples for IP Tunneling

The following example shows a simple GRE tunnel. Ethernet 1/2 is the tunnel source for router A and the tunnel destination for router B. Ethernet interface 2/1 is the tunnel source for router B and the tunnel destination for router A.

Router A:

```

feature tunnel
interface tunnel 0
ip address 209.165.20.2/8
tunnel source ethernet 1/2
tunnel destination 192.0.2.2
tunnel mode gre ip
tunnel path-mtu-discovery 25 1500

interface ethernet 1/2
ip address 192.0.2.55/8

```

Router B:

```
feature tunnel
interface tunnel 0
ip address 209.165.20.1/8
tunnel source ethernet 2/1
tunnel destination 192.0.2.55
tunnel mode gre ip

interface ethernet 2/1
ip address 192.0.2.2/8
```

Related Documents

Related Topic	Document Title
IP Tunnel commands	<i>Cisco Nexus 9000 Series NX-OS Interfaces Command Reference</i>



CHAPTER 10

Configuring Q-in-Q VLAN Tunnels

- [Information About Q-in-Q Tunnels, on page 309](#)
- [Guidelines and Limitations, on page 313](#)
- [Configuring Q-in-Q Tunnels and Layer 2 Protocol Tunneling, on page 315](#)
- [Configuring Q-in-Q Double Tagging, on page 321](#)
- [Verifying the Q-in-Q Configuration, on page 322](#)
- [Configuration Examples for Q-in-Q and Layer 2 Protocol Tunneling, on page 323](#)

Information About Q-in-Q Tunnels

This chapter describes how to configure IEEE 802.1Q-in-Q VLAN tunnels and Layer 2 protocol tunneling on Cisco NX-OS devices (7.0(3)F3(4), 7.0(3)I2(1) and later).

A Q-in-Q VLAN tunnel enables a service provider to segregate the traffic of different customers in their infrastructure, while still giving the customer a full range of VLANs for their internal use by adding a second 802.1Q tag to an already tagged frame.

Q-in-Q Tunneling

Business customers of service providers often have specific requirements for VLAN IDs and the number of VLANs to be supported. The VLAN ranges required by different customers in the same service-provider network might overlap, and the traffic of customers through the infrastructure might be mixed. Assigning a unique range of VLAN IDs to each customer would restrict customer configurations and could easily exceed the VLAN limit of 4096 of the 802.1Q specification.




Note Q-in-Q is supported on port channels. To configure a port channel as an asymmetrical link, all ports in the port channel must have the same tunneling configuration.

Using the 802.1Q tunneling feature, service providers can use a single VLAN to support customers who have multiple VLANs. Customer VLAN IDs are preserved and the traffic from different customers is segregated within the service-provider infrastructure even when they appear to be on the same VLAN. The 802.1Q tunneling expands the VLAN space by using a VLAN-in-VLAN hierarchy and tagging the tagged packets. A port configured to support 802.1Q tunneling is called a tunnel port. When you configure tunneling, you assign a tunnel port to a VLAN that is dedicated to tunneling. Each customer requires a separate VLAN, but that VLAN supports all of the customer's VLANs.

Customer traffic that is tagged in the normal way with appropriate VLAN IDs come from an 802.1Q trunk port on the customer device and into a tunnel port on the service-provider edge switch. The link between the customer device and the edge switch is an asymmetric link because one end is configured as an 802.1Q trunk port and the other end is configured as a tunnel port. You assign the tunnel port interface to an access VLAN ID that is unique to each customer. See the figure below.


Figure 30: 802.1Q-in-Q Tunnel Ports



Packets that enter the tunnel port on the service-provider edge switch, which are already 802.1Q-tagged with the appropriate VLAN IDs, are encapsulated with another layer of an 802.1Q tag that contains a VLAN ID that is unique to the customer. The original 802.1Q tag from the customer is preserved in the encapsulated packet. Therefore, packets that enter the service-provider infrastructure are double-tagged.

The outer tag contains the customer’s access VLAN ID (as assigned by the service provider), and the inner VLAN ID is the VLAN of the incoming traffic (as assigned by the customer). This double tagging is called tag stacking, Double-Q, or Q-in-Q as shown in the figure below.

Figure 31: Untagged, 802.1Q-Tagged, and Double-Tagged Ethernet Frames



By using this method, the VLAN ID space of the outer tag is independent of the VLAN ID space of the inner tag. A single outer VLAN ID can represent the entire VLAN ID space for an individual customer. This technique allows the customer's Layer 2 network to extend across the service provider network, potentially creating a virtual LAN infrastructure over multiple sites.



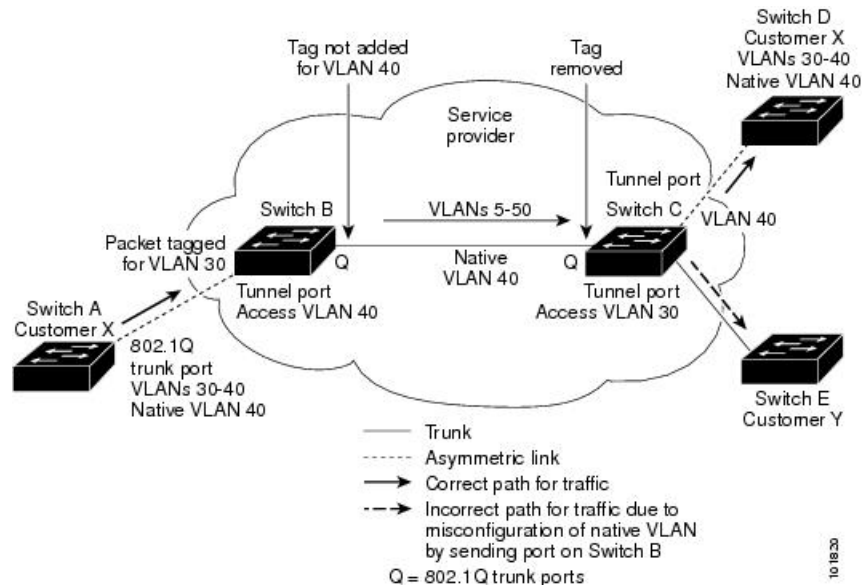
Note Hierarchical tagging, or multi-level dot1q tagging Q-in-Q, is not supported.

Native VLAN Hazard

When configuring 802.1Q tunneling on an edge switch, you must use 802.1Q trunk ports for sending out packets into the service-provider network. However, packets that go through the core of the service-provider network might be carried through 802.1Q trunks, ISL trunks, or nontrunking links. When 802.1Q trunks are used in these core switches, the native VLANs of the 802.1Q trunks must not match any native VLAN of the dot1q-tunnel port on the same switch because traffic on the native VLAN is not tagged on the 802.1Q transmitting trunk port.

In the figure below, VLAN 40 is configured as the native VLAN for the 802.1Q trunk port from Customer X at the ingress edge switch in the service-provider network (Switch B). Switch A of Customer X sends a tagged packet on VLAN 30 to the ingress tunnel port of Switch B in the service-provider network that belongs to access VLAN 40. Because the access VLAN of the tunnel port (VLAN 40) is the same as the native VLAN of the edge-switch trunk port (VLAN 40), the 802.1Q tag is not added to tagged packets that are received from the tunnel port. The packet carries only the VLAN 30 tag through the service-provider network to the trunk port of the egress-edge switch (Switch C) and is misdirected through the egress switch tunnel port to Customer Y.

Figure 32: Native VLAN Hazard



These are a couple ways to solve the native VLAN problem:

- Configure the edge switch so that all packets going out an 802.1Q trunk, including the native VLAN, are tagged by using the `vlan dot1q tag native` command. If the switch is configured to tag native VLAN packets on all 802.1Q trunks, the switch accepts untagged packets but sends only tagged packets.



Note The `vlan dot1q tag native` command is a global command that affects the tagging behavior on all trunk ports.

- Ensure that the native VLAN ID on the edge switch trunk port is not within the customer VLAN range. For example, if the trunk port carries traffic of VLANs 100 to 200, assign the native VLAN a number outside that range.

Information About Layer 2 Protocol Tunneling

Customers at different sites connected across a service-provider network need to run various Layer 2 protocols to scale their topology to include all remote sites, as well as the local sites. The Spanning Tree Protocol (STP) must run properly, and every VLAN should build a proper spanning tree that includes the local site and all remote sites across the service-provider infrastructure. The Cisco Discovery Protocol (CDP) must be able to discover neighboring Cisco devices from local and remote sites, and the VLAN Trunking Protocol (VTP) must provide consistent VLAN configuration throughout all sites in the customer network.

Starting with Cisco NX-OS Release 7.0(3)I7(3), you can configure the switch to allow multi-tagged BPDUs on a tunnel port. If you enable the `l2protocol tunnel allow-double-tag` command, when a multi-tagged customer BPDU enters the tunnel port, the original 802.1Q tags from the customer traffic is preserved and an outer VLAN tag (customer's access VLAN ID, as assigned by the service-provider) is added in the encapsulated packet. Therefore, BPDU packets that enter the service-provider infrastructure are multi tagged. When the BPDUs leave the service-provider network, the outer tag is removed and the original multi-tagged BPDU is sent to the customer network.

When protocol tunneling is enabled, edge switches on the inbound side of the service-provider infrastructure encapsulate Layer 2 protocol packets with a special MAC address and send them across the service-provider network. Core switches in the network do not process these packets, but forward them as normal packets. Bridge protocol data units (BPDUs) for CDP, STP, or VTP cross the service-provider infrastructure and are delivered to customer switches on the outbound side of the service-provider network. Identical packets are received by all customer ports on the same VLANs.


If protocol tunneling is not enabled on 802.1Q tunneling ports, remote switches at the receiving end of the service-provider network do not receive the BPDUs and cannot properly run STP, CDP, 802.1X, and VTP. When protocol tunneling is enabled, Layer 2 protocols within each customer's network are totally separate from those running within the service-provider network. Customer switches on different sites that send traffic through the service-provider network with 802.1Q tunneling achieve complete knowledge of the customer's VLAN.



Note Layer 2 protocol tunneling works by tunneling BPDUs in the software. A large number of BPDUs that come into the supervisor will cause the CPU load to go up. You might need to make use of software rate limiters to reduce the load on the supervisor CPU. See [Configuring Thresholds for Layer 2 Protocol Tunnel Ports, on page 320](#).

For example, in the figure below, Customer X has four switches in the same VLAN that are connected through the service-provider network. If the network does not tunnel BPDUs, switches on the far ends of the network cannot properly run the STP, CDP, 802.1X, and VTP protocols.


Figure 33: Layer 2 Protocol Tunneling



In the preceding example, STP for a VLAN on a switch in Customer X, Site 1 will build a spanning tree on the switches at that site without considering convergence parameters based on Customer X's switch in Site 2.

The figure below shows the resulting topology on the customer's network when BPDU tunneling is not enabled.

Figure 34: Virtual Network Topology Without BPDU Tunneling



Guidelines and Limitations

Q-in-Q tunnels and Layer 2 tunneling have the following configuration guidelines and limitations:

- Q-in-Q should be configured on the customer-facing interface of the service provider's edge device. If an Ethernet frame ingresses a Cisco Nexus 9000 series switch, the switch cannot encapsulate the frame with two 802.1Q headers within a single forwarding decision. Similarly, if a Q-in-Q-encapsulated Ethernet frame needs to egress a Cisco Nexus 9000 series switch without any 802.1Q headers, the switch cannot decapsulate two 802.1Q headers from the Ethernet frame within a single forwarding decision.
- Mapping multiple VLANs is supported from Cisco NX-OS Release 7.0(3)I5(1) and later versions.
- Multi-tagged BPDUs are supported on the Cisco Nexus 93108TC-EX and 93180YC-EX switches. We support up to three tags.
- Selective Q-in-Q tunneling is not supported with multi-tagged BPDU.
- Only multi-tagged CDP and STP BPDUs are supported.
- The inner-most tag must always be 0x8100.
- Multiple selective Q-in-Q tags are not supported. That is, Q-in-Q does not support multiple SP tags on a single interface.
- **show** commands with the **internal** keyword are not supported.
- Switches in the service-provider network must be configured to handle the increase in MTU size due to Q-in-Q tagging.
- MAC address learning for Q-in-Q tagged packets is based on the outer VLAN (Service Provider VLAN) tag. Packet forwarding issues might occur in deployments where a single MAC address is used across multiple inner (customer) VLANs.
- Layer 3 and higher parameters cannot be identified in tunnel traffic (for example, Layer 3 destination and source addresses). Tunneled traffic cannot be routed.
- The **system dot1q-tunnel transit** command have the following limitations:
 - It is required that you configure the **system dot1q-tunnel transit** command on ToR or modular devices, if the device is configured with Q-in-Q, Selective Q-in-Q or Selective Q-in-Q with multiple provider VLAN features.
 - It is required that you configure the **system dot1q-tunnel transit** command on vPC switches or non-vPC switches.
 - Layer 2 frames that exit trunk ports will always be tagged, even with the native VLAN of the port if these commands have been configured.
 - The MPLS, GRE, and IP-in-IP functionalities will not function effectively in conjunction with the Q-in-Q tunneling features if this command is configured on the switch.
- Cisco Nexus 9000 Series devices can provide only MAC-layer ACL/QoS for tunnel traffic (VLAN IDs and src/dest MAC addresses).
- You should use MAC address-based frame distribution.
- Asymmetrical links do not support the Dynamic Trunking Protocol (DTP) because only one port on the link is a trunk. You must configure the 802.1Q trunk port on an asymmetrical link to trunk unconditionally.
- You cannot configure the 802.1Q tunneling feature on ports that are configured to support private VLANs. Private VLAN are not required in these deployments.

- You must disable IGMP snooping on the tunnel VLANs.
- You should enter the `vlan dot1q tag native` command to maintain the tagging on the native VLAN and drop untagged traffic. This command prevents native VLAN misconfigurations.
- You must manually configure the 802.1Q interfaces to be edge ports.
- IGMP snooping is not supported on the inner VLAN.
- Q-in-Q is not supported on the uplink ports of Cisco Nexus 9332PQ, 9372PX, 9372TX, and 93120TX switches and Cisco Nexus 9396PX, 9396TX, and 93128TX switches with the N9K-M6PQ or N9K-M12PQ generic expansion module (GEM).
- Q-in-Q tunnels might be affected by the limitations of the Application Leaf Engine (ALE) uplink ports on Cisco Nexus 9300 and 9500 Series devices: [Limitations for ALE Uplink Ports](#)
- Q-in-Q tunneling is not supported on the following Application Spine Engine 2 (ASE2) and Application Spine Engine 3 (ASE3) based Cisco Nexus switches:
 - ASE2 - N9236C, N9272Q, N92304QC, and N92300Y
 - ASE3 - N92160YC-X
- Q-in-Q tagging is not supported in Cisco NX-OS Release 7.0(3)F3(4).
- Layer 2 protocol tunneling is not supported on Cisco Nexus 9500 Series switches with N9K-X9636C-R, N9K-X9636Q-R, N9K-X9636C-RX line cards.
- Beginning with 7.0(3)F3(4), on Cisco Nexus 9500 Series switches with N9K-X9636C-R, N9K-X9636Q-R, N9K-X9636C-RX line cards, Q-in-Q is supported only on port or port-channel Layer 2 Access VLAN Edge devices.

Configuring Q-in-Q Tunnels and Layer 2 Protocol Tunneling

Creating a 802.1Q Tunnel Port

You create the dot1q-tunnel port using the **switchport mode** command.



Note You must set the 802.1Q tunnel port to an edge port with the **spanning-tree port type edge** command. The provider VLAN membership of the port is changed using the **switchport access vlan *vlan-id*** command.

You should disable IGMP snooping on the access VLAN allocated for the dot1q-tunnel port to allow multicast packets to traverse the Q-in-Q tunnel.

For seamless packet forwarding and preservation of all VLAN tags on pure transit boxes in the SP cloud that have no Q-in-Q encapsulation or decapsulation requirement, configure the system-wide **system dot1q-tunnel transit** command. To remove the configuration, use the **no system dot1q-tunnel transit** command.

For the supported platforms and limitations of the **system dot1q-tunnel transit** command, see [Guidelines and Limitations, on page 313](#) section.

Before you begin

You must first configure the interface as a switchport.

Procedure

	Command or Action	Purpose
Step 1	switch# configure terminal	Enters global configuration mode.
Step 2	switch(config)# interface ethernet slot/port	Specifies an interface to configure, and enters interface configuration mode.
Step 3	switch(config-if)# switchport	Sets the interface as a Layer 2 switching port.
Step 4	switch(config-if)# switchport mode dot1q-tunnel	Creates a 802.1Q tunnel on the port. The port will go down and reinitialize (port flap) when the interface mode is changed. BPDU filtering is enabled and CDP is disabled on tunnel interfaces.
Step 5	switch(config-if)# spanning-tree port type edge	Designates the port as a spanning-tree edge port.
Step 6	switch(config-if)# switchport access vlan vlan-id	Configures the Provider access VLAN value.
Step 7	(Optional) switch(config-if)# no switchport mode dot1q-tunnel	Disables the 802.1Q tunnel on the port.
Step 8	switch(config-if)# exit	Exits configuration mode.
Step 9	(Optional) switch(config)# show dot1q-tunnel [interface if-range]	Displays all ports that are in dot1q-tunnel mode. Optionally, you can specify an interface or range of interfaces to display.
Step 10	(Optional) switch(config)# no shutdown	Clears the errors on the interfaces and VLANs where policies correspond with hardware policies. This command allows policy programming to continue and the port to come up. If policies do not correspond, the errors are placed in an error-disabled policy state.
Step 11	(Optional) switch(config)# copy running-config startup-config	Copies the running configuration to the startup configuration.

Example

This example shows how to create an 802.1Q tunnel port:

```
switch# configure terminal
switch(config)# interface ethernet 7/1
switch(config-if)# switchport
switch(config-if)# switchport mode dot1q-tunnel
```



```

switch(config-if)# spanning-tree port type edge
switch(config-if)# switchport access vlan vlan 10
switch(config-if)# exit
switch(config)# exit
switch# show dot1q-tunnel

```

Configuring VLAN Mapping for Selective Q-in-Q on a 802.1Q Tunnel Port

To configure VLAN mapping for selective Q-in-Q on a 802.1Q tunnel port, complete the following steps.



Note You cannot configure one-to-one mapping and selective Q-in-Q on the same interface.

You must set the 802.1Q tunnel port to an edge port with the **spanning-tree port type edge** command. The provider VLAN membership of the port is changed using the **switchport access vlan *vlan-id*** command.

You should disable IGMP snooping on the access VLAN allocated for the dot1q-tunnel port to allow multicast packets to traverse the Q-in-Q tunnel.

Procedure

	Command or Action	Purpose
Step 1	switch# configure terminal	Enters global configuration mode.
Step 2	switch(config)# interface <i>interface-id</i>	Enters interface configuration mode for the interface connected to the service provider network. You can enter a physical interface or an EtherChannel port channel.
Step 3	switch(config-if)# switchport mode dot1q-tunnel	Configures the interface as a tunnel port.
Step 4	switch(config-if)# spanning-tree port type edge	Designates the port as a spanning-tree edge port.
Step 5	switch(config-if)# switchport access vlan <i>vlan-id</i>	Configures the Provider access VLAN value.
Step 6	switch(config-if)# switchport vlan mapping <i>vlan-id-range</i> dot1q-tunnel <i>outer vlan-id</i>	Enters the VLAN IDs to be mapped: <ul style="list-style-type: none"> vlan-id-range—The customer VLAN ID range (C-VLAN) entering the switch from the customer network. The range is from 1 to 4094. You can enter a string of VLAN-IDs. outer vlan-id—Enter the outer VLAN ID (S-VLAN) of the service provider network. The range is from 1 to 4094.
Step 7	switch(config-if)# exit	Exits the configuration mode.

	Command or Action	Purpose
Step 8	switch# show interfaces <i>interface-id</i> vlan mapping	Verifies the configuration.
Step 9	switch# copy running-config startup-config	(Optional) Saves your entries in the configuration file.

Use the **no switchport vlan mapping** *vlan-id-range* **dot1q-tunnel** *outer vlan-id* command to remove the VLAN mapping configuration.

The following example shows how to configure selective Q-in-Q mapping on the port so that traffic with a C-VLAN ID of 1 to 5 enters the switch with an S-VLAN ID of 100. The traffic of any other VLAN IDs is dropped.

Example

```
switch(config)# interface gigabitethernet0/1
switch(config-if)# switchport vlan mapping 1-5 dot1q-tunnel 100
switch(config-if)# spanning-tree port type edge
switch(config-if)# switchport access vlan 10
Switch(config-if)# exit
```

Changing the EtherType for Q-in-Q

The switch default EtherType is 0x8100 for 802.1Q and Q-in-Q encapsulations. EtherType cannot be configured to 0x9100, 0x9200 and 0x88a8 on the switchport interface.

Enabling the Layer 2 Protocol Tunnel

You can enable protocol tunneling on the 802.1Q tunnel port.

Procedure

	Command or Action	Purpose
Step 1	switch# configure terminal	Enters global configuration mode.
Step 2	switch(config)# interface ethernet <i>slot/port</i>	Specifies an interface to configure, and enters interface configuration mode.
Step 3	switch(config-if)# switchport	Sets the interface as a Layer 2 switching port.
Step 4	switch(config-if)# switchport mode dot1q-tunnel	Creates a 802.1Q tunnel on the port. The port will go down and reinitialize (port flap) when the interface mode is changed. BPDU filtering is enabled and CDP is disabled on tunnel interfaces.

	Command or Action	Purpose
Step 5	switch(config-if)# l2protocol tunnel [cdp stp lacp lldp vtp]	Enables Layer 2 protocol tunneling. Optionally, you can enable CDP, STP, LACP, LLDP, or VTP tunneling.
Step 6	(Optional) switch(config-if)# no l2protocol tunnel [cdp stp lacp lldp vtp]	Disables protocol tunneling.
Step 7	switch(config-if)# exit	Exits configuration mode.
Step 8	(Optional) switch(config)# no shutdown	Clears the errors on the interfaces and VLANs where policies correspond with hardware policies. This command allows policy programming to continue and the port to come up. If policies do not correspond, the errors are placed in an error-disabled policy state.
Step 9	(Optional) switch(config)# copy running-config startup-config	Copies the running configuration to the startup configuration.

Example

This example shows how to enable protocol tunneling on an 802.1Q tunnel port:

```
switch# configure terminal
switch(config)# interface ethernet 7/1
switch(config-if)# switchport
switch(config-if)# switchport mode dot1q-tunnel
switch(config-if)# l2protocol tunnel stp
switch(config-if)# exit
switch(config)# exit
```

Configuring Global CoS for L2 Protocol Tunnel Ports

You can specify a Class of Service (CoS) value globally so that ingress BPDUs on the tunnel ports are encapsulated with the specified class.

Procedure

	Command or Action	Purpose
Step 1	switch# configure terminal	Enters global configuration mode.
Step 2	switch(config)# l2protocol tunnel cos value	Specifies a global CoS value on all Layer 2 protocol tunneling ports. The default cos-value is 5.
Step 3	(Optional) switch(config)# no l2protocol tunnel cos	Sets the global CoS value to default.
Step 4	switch(config)# exit	Exits configuration mode.

	Command or Action	Purpose
Step 5	(Optional) switch# no shutdown	Clears the errors on the interfaces and VLANs where policies correspond with hardware policies. This command allows policy programming to continue and the port to come up. If policies do not correspond, the errors are placed in an error-disabled policy state.
Step 6	(Optional) switch# copy running-config startup-config	Copies the running configuration to the startup configuration.

Example

This example shows how to specify a global CoS value for the purpose of Layer 2 protocol tunneling:

```
switch# configure terminal
switch(config)# l2protocol tunnel cos 6
switch(config)# exit
```

Configuring Thresholds for Layer 2 Protocol Tunnel Ports

You can specify the port drop and shutdown value for a Layer 2 protocol tunneling port.

Procedure

	Command or Action	Purpose
Step 1	switch# configure terminal	Enters global configuration mode.
Step 2	switch(config)# interface ethernet slot/port	Specifies an interface to configure, and enters interface configuration mode.
Step 3	switch(config-if)# switchport	Sets the interface as a Layer 2 switching port.
Step 4	switch(config-if)# switchport mode dot1q-tunnel	Creates a 802.1Q tunnel on the port.
Step 5	switch(config-if)# l2protocol tunnel drop-threshold [cdp stp vtp] packets-per-sec	Specifies the maximum number of packets that can be processed on an interface before being dropped. Optionally, you can specify CDP, STP, or VTP. Valid values for the packets are from 1 to 4096.
Step 6	(Optional) switch(config-if)# no l2protocol tunnel drop-threshold [cdp stp vtp]	Resets the threshold values to 0 and disables the drop threshold.
Step 7	switch(config-if)# l2protocol tunnel shutdown-threshold [cdp stp vtp] packets-per-sec	Specifies the maximum number of packets that can be processed on an interface. When the number of packets is exceeded, the port is put in error-disabled state. Optionally, you can

	Command or Action	Purpose
		specify CDP, STP, or VTP. Valid values for the packets is from 1 to 4096.
Step 8	(Optional) <code>switch(config-if)# no l2protocol tunnel shutdown-threshold [cdp stp vtp]</code>	Resets the threshold values to 0 and disables the shutdown threshold.
Step 9	<code>switch(config-if)# exit</code>	Exits configuration mode.
Step 10	(Optional) <code>switch(config)# no shutdown</code>	Clears the errors on the interfaces and VLANs where policies correspond with hardware policies. This command allows policy programming to continue and the port to come up. If policies do not correspond, the errors are placed in an error-disabled policy state.
Step 11	(Optional) <code>switch(config)# copy running-config startup-config</code>	Copies the running configuration to the startup configuration.

Configuring Q-in-Q Double Tagging

Enable multi-tagging for STP and CDP BPDUs.

Procedure

	Command or Action	Purpose
Step 1	<code>configure terminal</code> Example: <code>switch# configure terminal</code>	Enters global configuration mode.
Step 2	<code>interface interface</code> Example: <code>switch(config)# interface ethernet 7/1</code>	Specifies the interface that you are configuring.
Step 3	<code>switchport</code> Example: <code>switch(config-if)# switchport</code>	Sets the interface as a Layer 2 switching port.
Step 4	<code>switchport mode dot1q-tunnel</code> Example: <code>switch(config-if)# switchport mode dot1q-tunnel</code>	Creates an 802.1Q tunnel on the port. The port goes down and reinitializes (port flap) when the interface mode is changed. BPDU filtering is enabled and CDP is disabled on tunnel interfaces.
Step 5	<code>l2protocol tunnel [cdp stp]</code> Example: <code>switch(config-if)# l2protocol tunnel cdp</code>	Enables Layer 2 protocol tunneling. Optionally, you can enable CDP or STP.

	Command or Action	Purpose
Step 6	(Optional) no l2protocol tunnel [cdp stp] Example: switch(config-if)# no l2protocol tunnel stp	Disables protocol tunneling.
Step 7	l2protocol tunnel allow-double-tag Example: switch(config-if)# l2protocol tunnel allow-double-tag	Enables multi-tagging for STP and CDP BPDUs on the interface.
Step 8	(Optional) no l2protocol tunnel allow-double-tag Example: switch(config-if)# no l2protocol tunnel allow-double-tag	Disables multi-tagging for STP and CDP BPDUs on the interface.
Step 9	exit Example: switch(config-if)# exit	Exits configuration mode.

Example

This example shows how to enable multi-tagging for STP and CDP BPDUs:

```
switch# configure terminal
switch(config)# interface ethernet 7/1
switch(config-if)# switchport
switch(config-if)# switchport mode dot1q-tunnel
switch(config-if)# l2protocol tunnel cdp
switch(config-if)# l2protocol tunnel stp
switch(config-if)# l2protocol tunnel allow-double-tag
switch(config-if)# exit
switch(config)# exit
switch#
```

Verifying the Q-in-Q Configuration

Command	Purpose
clear l2protocol tunnel counters [interface <i>if-range</i>]	Clears all the statistics counters. If no interfaces are specified, the Layer 2 protocol tunnel statistics are cleared for all interfaces.
show dot1q-tunnel [interface <i>if-range</i>]	Displays a range of interfaces or all interfaces that are in dot1q-tunnel mode.

Command	Purpose
show l2protocol tunnel [interface <i>if-range</i> vlan <i>vlan-id</i>]	Displays Layer 2 protocol tunnel information for a range of interfaces, for all dot1q-tunnel interfaces that are part of a specified VLAN or all interfaces.
show l2protocol tunnel summary	Displays a summary of all ports that have Layer 2 protocol tunnel configurations.
show running-config l2pt	Displays the current Layer 2 protocol tunnel running configuration.

Configuration Examples for Q-in-Q and Layer 2 Protocol Tunneling

This example shows a service provider switch that is configured to process Q-in-Q for traffic coming in on Ethernet 7/1. A Layer 2 protocol tunnel is enabled for STP BPDUs. The customer is allocated VLAN 10 (outer VLAN tag).

```
switch# configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
switch(config)# vlan 10
switch(config-vlan)# no shutdown
switch(config-vlan)# no ip igmp snooping
switch(config-vlan)# exit
switch(config)# interface ethernet 7/1
switch(config-if)# switchport
switch(config-if)# switchport mode dot1q-tunnel
switch(config-if)# switchport access vlan 10
switch(config-if)# spanning-tree port type edge
switch(config-if)# l2protocol tunnel stp
switch(config-if)# no shutdown
switch(config-if)# exit
switch(config)# exit
switch#
```




CHAPTER 11

Configuring Smart-Channel

This chapter describes how to configure smart-channel on Cisco NX-OS devices.

This chapter includes the following sections:

- [About Smart Channel, on page 325](#)
- [Topology Examples for Smart Channel , on page 326](#)
- [Prerequisites for Smart Channel, on page 328](#)
- [Guidelines and Limitations for Smart Channel, on page 329](#)
- [Default Settings for Smart Channel, on page 329](#)
- [Configuring Smart Channel, on page 329](#)
- [Verifying the Smart Channel Configuration, on page 332](#)
- [Configuration Examples for Smart Channel, on page 333](#)

About Smart Channel

Smart channel is a hardware-based, multi-terabit solution for the Layer 2 traffic distribution, load balancing, and redirection on the Cisco Nexus switches. This feature is supported on the Cisco Nexus 9372PX, 93108TC-EX, and the Cisco Nexus 9516 switches.



Note Smart channel feature is not supported on Cisco 9500 EX / FX line cards.

Smart channel is an aggregation of multiple physical links that creates a single logical link. You can bundle up multiple physical links into a port group to provide an increased bandwidth (an aggregate of the multiple physical links) and redundancy.

If one port within a smart channel fails, the traffic switches to the remaining ports in the smart channel.

Smart channel allows you to create a cluster of transparent mode appliances.

Smart Channel Features

The smart channel features are as follows:

- Multi-terabit solution at line rate
- Simplified provisioning and ease of deployment

- Transparency to end device and stateless protocol benefits
- Removes the requirement for an expensive external load balancer

Benefits of Smart Channel

The benefits of smart channel are as follows:

- Simultaneous redirection and load balancing
- IP-stickiness and resiliency
- Health monitoring
- Removes the requirement for an expensive external load balancer
- Hashing does not depend on the wiring or the port numbering
- Every port on the switch is used for load balancing and traffic redirection
- Automatic failure handling of servers or appliances

Examples of the Deployment Use Cases

Examples of the deployment use cases for the smart channel feature are as follows:

- Load balances to a pool of firewalls.
- Scales the VDS-TC (video-caching) solution.
- Scales the transparent mode devices.


Topology Examples for Smart Channel

This section displays the following examples:

- Basic topology for smart channel
- Use case of a smart channel configuration
- Fail-action for resilient hashing


You can use the smart channel feature to load balance traffic to appliances used in a monitoring network. The following figure shows the basic topology, where the traffic is sent to the appliances where you need to load balance the traffic towards, such as the IPS or the IDS devices.

Figure 35: Standard Topology for Smart Channel




The following example shows a typical use case of smart channel in a network where the traffic is spanned from the production environment to the monitoring environment. In this example, we are using the Cisco Nexus Data Broker to send copy of the monitoring traffic and scale monitoring networks.

Figure 36: Use Case for a Smart Channel Configuration




The following example shows the fail-action of a smart channel configuration:

Figure 37: Fail-Action of a Smart Channel Configuration



The following example shows the fail-action of a smart channel configuration:

Figure 38: Fail-Action of a Smart Channel Configuration



Prerequisites for Smart Channel

Smart channel has the following prerequisite:

- You must ensure that an enough TCAM size has been allocated to the VACL. To verify the TCAM size, use the **sh hardware access-list tcam region** command. If the appropriate TCAM size is not allocated, use the **hardware access-list tcam region VACL size additional of 256** command to allocate the appropriate TCAM size.

Guidelines and Limitations for Smart Channel

Smart channel has the following configuration guidelines and limitations:

- Cisco Nexus 9372PX, 93108TC-EX, and the Cisco Nexus 9516 switches support smart channel servicing.
- This feature is supported for the Cisco Nexus 9372PX, 93108TC-EX, and the Cisco Nexus 9516 switches.



Note Smart channel feature is not supported on Cisco 9500 EX / FX / R line cards.

- This feature is supported for the Cisco Nexus C93108TC-EX beginning Cisco NX-OS Release 9.2 (x).
- Smart channel does not support the vPC, port channel, and the L3 interfaces.
- Only the port group interfaces in a trunk or access mode are supported.
- You must not share the smart port-group to more than one service when the services have the access configuration.
- Ensure that the TCAM size is equal to the sum of the number of the configured VLANs on the service by the number of buckets.
- Ensure that the smart channel service does not have the same load balancing method (load-balance method src ip) and the configuration of the (source filter ip any any).
- Ensure that the port-group to be added to the smart channel service has been configured.

Default Settings for Smart Channel

The following table lists the default settings for the smart channel parameter.

Table 17: Default Smart Channel Parameter

Parameters	Default
Smart channel	Disabled

Configuring Smart Channel

Enabling Smart Channel

You must enable the smart channel feature before you can configure the smart channel services on the device.

Procedure

	Command or Action	Purpose
Step 1	configure terminal Example: switch# configure terminal switch(config)#	Enters global configuration mode.
Step 2	[no] feature smart-channel Example: switch(config)# feature smart-channel	Enables or disables smart channel. By default, smart channel is disabled.
Step 3	(Optional) show feature grep smart-channel Example: switch(config-if)# show feature grep smart-channel	Displays the status of the smart channel configuration.
Step 4	(Optional) copy running-config startup-config Example: switch(config)# copy running-config startup-config	Copies the running configuration to the startup configuration.

Configuring Port Groups

After you enable smart channel, you must create a port group and configure active interfaces on that group.

Procedure

	Command or Action	Purpose
Step 1	configure terminal Example: switch# configure terminal switch(config)#	Enters global configuration mode.
Step 2	[no] smart-channel port-group <i>port-group-name</i> Example: switch(config)# smart-channel port-group Webservers switch(config-port-group)#	Creates or deletes a port group.
Step 3	Required: interface ethernet <i>type slot/port</i> Example: switch(config)# interface ethernet 3/1 switch(config-if)#	Configures the active interfaces on the port group.

	Command or Action	Purpose
Step 4	(Optional) copy running-config startup-config Example: switch(config-port-group) # copy running-config startup-config	Copies the running configuration to the startup configuration.

Configuring the Smart Channel Service

To configure the smart channel service, you must do the following:

1. Configure the service name
2. Configure the VLAN/source/destination filters
3. Associate the port group to the smart channel service
4. Specify the load distribution scheme
5. Activate the smart channel service

Before you begin

You must enable the smart channel feature before you configure the smart channel service.



Note Beginning from Cisco Nexus NX-OS Release 9.3(3), feature Smart-channel is not supported. It is recommended to take necessary actions while upgrading from any previous release to 9.3(2) or any newer versions.

Procedure

	Command or Action	Purpose
Step 1	configure terminal Example: switch# configure terminal switch(config) #	Enters global configuration mode.
Step 2	Required: [no] smart-channel service-name Example: switch(config) # smart-channel WebTraffic switch(config-smart-channel) #	Configures or disables the smart channel service.
Step 3	[no] access vlan access-vlan vlan vlan-range Example: switch(config-smart-channel) # access vlan 10-20 switch(config-port-group) #	Configures a list of VLANs for the smart channel service. While the access VLANs create the smart channel in an access mode, the VLANs in the VLAN range creates the smart channel in the trunk mode.

	Command or Action	Purpose
Step 4	[no] port group <i>port-group-name</i> Example: switch(config-smart-channel)# port group WEBSERVERS switch(config-smart-channel)#	Associates a port-group with the smart channel service.
Step 5	[no] load-balance method [<i>src</i> <i>dst</i>] Example: switch(config-smart-channel)# load-balance method src-ip switch(config-smart-channel)#	Configures the load balancing method.
Step 6	[no] destination filter ip <i>any</i> [<i>any</i>] Example: switch(config-smart-channel)# destination filter ip any any switch(config-smart-channel)#	Configures the selected destination subnets.
Step 7	[no] source filter [<i>ipanyany</i>] Example: switch(config-smart-channel)# source filter ip any any switch(config-smart-channel)#	Configures the selected source subnets.
Step 8	[no] load-balance method [<i>src</i> <i>dst</i>] Bucket <i>B</i> Example: switch(config-smart-channel)# load-balance method src Bucket 16 switch(config-port-group)#	Configures the load balancing method.
Step 9	no shut Example: switch(config-smart-channel)# no shut switch(config-smart-channel)#	Activates the smart channel service.

Verifying the Smart Channel Configuration

Smart Channel Show Commands

To display the smart channel configuration, perform one of the following tasks:

Procedure

	Command or Action	Purpose
Step 1	show smart-channel <i>service-name</i>	Displays the smart channel configuration status.

	Command or Action	Purpose
Step 2	<code>show vlan access-list <i>vlan access-map name</i></code>	Displays the statistics for a smart channel service.
Step 3	<code>show running-config smart-channel [all]</code>	Displays the running configuration for smart channel.

Configuration Examples for Smart Channel

This example shows how to configure smart channel:

Enabling the smart channel feature.

```
switch# configure terminal
switch(config)# feature smart-channel
```

Configuring port-group.

```
switch(config)# smart-channel port-group pg1
switch(config-port-group)# interface ethernet 10/23
switch(config-pg-node)# interface ethernet 10/24
```

Configuring services.



Note Ensure that the port-group to be added to the smart channel service has been configured.

Service with Trunk configuration

```
switch(config-smart-channel)# port-group pg1
switch(config-smart-channel)# vlan 67
switch(config-smart-channel)# load-balance method src ip
switch(config-smart-channel)# destination filter ip any any
switch(config-smart-channel)# no shut
```

Service with Access configuration

```
switch(config-smart-channel)# port-group pg1
switch(config-smart-channel)# access vlan 68
switch(config-smart-channel)# load-balance method src ip
switch(config-smart-channel)# destination filter ip any any
switch(config-smart-channel)# no shut
```

Verifying the smart channel configuration.

```
switch(config-smart-channel)# show running-config smart-channel

version 7.0(3)I6(1)

smart-channel port-group pg1
  interface Eth10/23
  interface Eth10/24
```

```

smart-channel Service1
  port-group pgl
  vlan 67
  destination filter ip any any
  load-balance method src ip
  no shutdown

```

Checking the status of the service and the port-group interfaces.

```
switch(config-smart-channel)# show smart-channel Service1
```

Legend:

ST(Status): ST-Standby,LF-Link Failed,PF-Probe Failed,PD-Peer Down,IA-Inactive

Name	LB Scheme	Status	Buckets
Service1	src-ip	ACTIVE	2

Port Group

pgl

Pool

Service1_smartc_pool	Source/Destination Filter	Protocol	Port
	0.0.0.0 / 255.255.255.255	IP	0

Member-Interface	Cfg-STS	Weight	Op-STS
1 Ethernet10/23	OK	1	

Bucket List

Service1_smartc_ip_1_bucket_1

Member-Interface	Cfg-STS	Weight	Op-STS
2 Ethernet10/24	OK	1	

Bucket List

Service1_smartc_ip_1_bucket_2

=====

```

VACL config
switch# show running-config aclmgr

```

Checking the statistics.

```
switch# show vlan access-list Service1_smartc_pool
```

Checking the fail-action.



Note This is a default configuration and you are not required to perform any additional configurations.

```
switch# show smart-channel Service1
```

Legend:

ST(Status): ST-Standby, LF-Link Failed, PF-Probe Failed, PD-Peer Down, IA-Inactive

Name	LB Scheme	Status	Buckets
Service1	src-ip	ACTIVE	2

Port Group

pg1

Pool

Source/Destination Filter	Protocol	Port
0.0.0.0 / 255.255.255.255	IP	0

Member-Interface	Cfg-STS	Weight	Op-STS
1 Ethernet10/23	LF	1	

Bucket List

Service1_smartc_ip_1_bucket_1

Member-Interface	Cfg-STS	Weight	Op-STS
2 Ethernet10/24	OK	1	

Bucket List

Service1_smartc_ip_1_bucket_2

```
switch# show running-config aclmgr
```

```
switch# show vlan access-list Service1_smartc_pool
```




CHAPTER 12

Configuring Static and Dynamic NAT Translation

This chapter includes the following sections:

- [Network Address Translation Overview, on page 337](#)
- [Information About Static NAT, on page 338](#)
- [Dynamic NAT Overview, on page 339](#)
- [Timeout Mechanisms, on page 340](#)
- [NAT Inside and Outside Addresses, on page 341](#)
- [Pool Support for Dynamic NAT, on page 342](#)
- [Static and Dynamic Twice NAT Overview, on page 342](#)
- [VRF Aware NAT, on page 343](#)
- [Guidelines and Limitations for Static NAT, on page 345](#)
- [Restrictions for Dynamic NAT, on page 346](#)
- [Guidelines and Limitations for Dynamic Twice NAT, on page 347](#)
- [Configuring Static NAT, on page 348](#)
- [Configuring Dynamic NAT, on page 355](#)

Network Address Translation Overview

Network Address Translation (NAT) enables private IP internetworks that use nonregistered IP addresses to connect to the Internet. NAT operates on a device, usually connecting two networks, and translates private (not globally unique) IP addresses in the internal network into legal IP addresses before packets are forwarded to another network. You can configure NAT to advertise only one IP address for the entire network to the outside world. This ability provides additional security, effectively hiding the entire internal network behind one IP address.

A device configured with NAT has at least one interface to the inside network and one to the outside network. In a typical environment, NAT is configured at the exit router between a stub domain and a backbone. When a packet leaves the domain, NAT translates the locally significant source IP address into a globally unique IP address. When a packet enters the domain, NAT translates the globally unique destination IP address into a local IP address. If more than one exit point exists, NAT configured at each point must have the same translation table.

NAT is described in RFC 1631.

Beginning with 7.0(3)I2(1), static and dynamic NAT translation is supported.

Information About Static NAT


Static Network Address Translation (NAT) allows the user to configure one-to-one translations of the inside local addresses to the outside global addresses. It allows both IP addresses and port number translations from the inside to the outside traffic and the outside to the inside traffic. The Cisco Nexus device supports Hitless NAT, which means that you can add or remove a NAT translation in the NAT configuration without affecting the existing NAT traffic flows.

Static NAT creates a fixed translation of private addresses to public addresses. Because static NAT assigns addresses on a one-to-one basis, you need an equal number of public addresses as private addresses. Because the public address is the same for each consecutive connection with static NAT, and a persistent translation rule exists, static NAT enables hosts on the destination network to initiate traffic to a translated host if an access list exists that allows it.

With dynamic NAT and Port Address Translation (PAT), each host uses a different address or port for each subsequent translation. The main difference between dynamic NAT and static NAT is that static NAT allows a remote host to initiate a connection to a translated host if an access list exists that allows it, while dynamic NAT does not.

The figure shows a typical static NAT scenario. The translation is always active so both translated and remote hosts can originate connections, and the mapped address is statically assigned by the **static** command.

Figure 39: Static NAT



These are key terms to help you understand static NAT:

- NAT inside interface—The Layer 3 interface that faces the private network.
- NAT outside interface—The Layer 3 interface that faces the public network.
- Local address—Any address that appears on the inside (private) portion of the network.
- Global address—Any address that appears on the outside (public) portion of the network.
- Legitimate IP address—An address that is assigned by the Network Information Center (NIC) or service provider.

- Inside local address—The IP address assigned to a host on the inside network. This address does not need to be a legitimate IP address.
- Outside local address—The IP address of an outside host as it appears to the inside network. It does not have to be a legitimate address, because it is allocated from an address space that can be routed on the inside network.
- Inside global address—A legitimate IP address that represents one or more inside local IP addresses to the outside world.
- Outside global address—The IP address that the host owner assigns to a host on the outside network. The address is a legitimate address that is allocated from an address or network space that can be routed.

Dynamic NAT Overview

Dynamic Network Address Translation (NAT) translates a group of real IP addresses into mapped IP addresses that are routable on a destination network. Dynamic NAT establishes a one-to-one mapping between unregistered and registered IP addresses; however, the mapping can vary depending on the registered IP address that is available at the time of communication.

A dynamic NAT configuration automatically creates a firewall between your internal network and outside networks or the Internet. Dynamic NAT allows only connections that originate inside the stub domain—a device on an external network cannot connect to devices in your network, unless your device has initiated the contact.

Dynamic NAT translations do not exist in the NAT translation table until a device receives traffic that requires translation. Dynamic translations are cleared or timed out when not in use to make space for new entries. Usually, NAT translation entries are cleared when the ternary content addressable memory (TCAM) entries are limited. The default minimum timeout for dynamic NAT translations is 30 minutes.



Note For Cisco NX-OS 7.0(3)I2(3) and earlier releases, the minimum value of the `sampling-timeout` in the `ip nat translation sampling-timeout` command was reduced from 30 minutes to 15 minutes.



Note For Cisco NX-OS 7.0(3)I3(1) and later releases, the `ip nat translation sampling-timeout` command is not supported. Statistics are collected every 60 seconds for the installed NAT policies. These statistics are used to determine if the flow is active or not.

For Cisco NX-OS 7.0(3)I2(3) and earlier releases, the timeout of a dynamic NAT translation involves both the `sampling-timeout` value and the TCP or UDP timeout value. The `sampling-timeout` specifies the time after which the device checks for dynamic translation activity. It has a default value of 12 hours. All the other timeouts start only after the `sample-timeout` times out. After the `sampling-timeout`, the device inspects the packets that are hitting this translation. The checking happens for the TCP or UDP timeout period. If there are no packets for the TCP or UDP timeout period, the translation is cleared. If activity is detected on the translation, then the checking is stopped immediately and a `sampling-timeout` period begins.

For Cisco NX-OS 7.0(3)I2(3) and earlier releases, after waiting for this new `sampling-timeout` period, the device checks for dynamic translation activity again. During an activity check the TCAM sends a copy of the packet that matches the dynamic NAT translation to the CPU. If the Control Plane Policing (CoPP) is configured

at a low threshold, the TCP or UDP packets might not reach the CPU, and the CPU considers this as inactivity of the NAT translation.

Dynamic NAT supports Port Address Translation (PAT) and access control lists (ACLs). PAT, also known as overloading, is a form of dynamic NAT that maps multiple unregistered IP addresses to a single registered IP address by using different ports. Your NAT configuration can have multiple dynamic NAT translations with same or different ACLs. However, for a given ACL, only one interface can be specified.

Timeout Mechanisms

For Cisco NX-OS 7.0(3)I2(3) and earlier releases, after dynamic NAT translations are created, they must be cleared when not in use so that newer translations can be created, especially because the number of TCAM entries is limited. The following NAT translation timeout timers are supported on the switch:

The following NAT translation timeout timers are supported on the switch:

- **syn-timeout** - Timeout value for TCP data packets that send the SYN request, but do not receive a SYN-ACK reply.

The timeout value ranges from 1 second to 172800 seconds. The default value is 60 seconds.



Note The **syn-timeout** option is supported only on Cisco Nexus 9200 and 9300-EX, -FX, -FX2, -FX3, -FXP, -GX platform switches.

- **finrst-timeout** - Timeout value for the flow entries when a connection is terminated by receiving RST or FIN packets. Use the same keyword to configure the behavior for both RST and FIN packets.

The timeout value ranges from 1 second to 172800 seconds. The default value is 60 seconds.

- If a FIN packet is received after the connection is established, SYN-->SYN-ACK-->FIN, the finrst timer starts.
- If a FIN-ACK is received from the other side, the translation entry is cleared immediately, else it clears after the timeout value completes.
- If an RST packet is received after the connection is established, SYN-->SYN-ACK-->RST, the translation entry is cleared immediately.



Note The **finrst-timeout** option is supported only on Cisco Nexus 9200 and 9300-EX, -FX, -FX2, -FX3, -FXP, -GX platform switches.

- **tcp-timeout** - Timeout value for TCP translations for which connections have been established after a three-way handshake (SYN, SYN-ACK, ACK). If no active flow occurs after the connection has been established, the translations expire as per the configured timeout value.

The timeout value ranges from 60 seconds to 172800 seconds. The default value is 3600 seconds.

- **udp-timeout** - Timeout value for all NAT UDP packets.

The timeout value ranges from 60 seconds to 172800 seconds. The default value is 3600 seconds.

- **timeout** - Timeout value for dynamic NAT translations.

The timeout value ranges from 60 seconds to 172800 seconds. The default value is 3600 seconds.

- **icmp-timeout** - Timeout value for ICMP packets.

The timeout value ranges from 60 seconds to 172800 seconds. The default value is 3600 seconds.

- **sampling-timeout** - Time after which the device checks for dynamic translation activity.

The timeout value ranges from 900 seconds to 172800 seconds.



Note When you create dynamic entries without timeouts configured, they take the default timeout of 3600 seconds. After you change the default timeout values to the new values, the translation entries created after will pick up the latest timeout values.

The **udp-timeout** and the **timeout** value timers are triggered after the timeout configured for the **ip nat translation sampling-timeout** command expires.



Note For Cisco NX-OS 7.0(3)I3(1) and later releases, there are three different options that can be configured for aging:

- Time-out: This is applicable for all type of flows(both TCP and UDP).
 - TCP TIME-OUT: This is applicable for only TCP flows.
 - UDP TIME-OUT: This is applicable for only UDP flows.
-



Note For Cisco NX-OS 7.0(3)I3(1) and later releases, when you create dynamic entries without timeouts configured, they take the default timeout of one hour (60 minutes). If you enter the **clear ip nat translations all** command after configuring timeouts, the configured timeout take effect. A timeout can be configured from 60 to 172800 seconds.



Note When you create dynamic entries without timeouts configured, they take the default timeout of one hour (60 minutes). If you enter the **clear ip nat translations all** command after configuring timeouts, the configured timeout take effect. A timeout can be configured from 60 to 172800 seconds.

NAT Inside and Outside Addresses

NAT inside refers to networks owned by an organization that must be translated. When NAT is configured, hosts within this network will have addresses in one space (known as the local address space) that will appear to those outside the network as being in another space (known as the global address space).

Similarly, NAT outside refers to those networks to which the stub network connects. They are not generally under the control of the organization. Hosts in outside networks can be subject to translation and can have local and global addresses.

NAT uses the following definitions:

- Local address—A local IP address that appears on the inside of a network.
- Global address—A global IP address that appears on the outside of a network.
- Inside local address—The IP address that is assigned to a host on the inside network. The address is probably not a legitimate IP address assigned by the Internet Network Information Center (InterNIC) or a service provider.
- Inside global address—A legitimate IP address (assigned by InterNIC or a service provider) that represents one or more inside local IP addresses to the outside world.
- Outside local address—The IP address of an outside host as it appears to the inside network. The address is not necessarily legitimate; it was allocated from the address space that is routable on the inside.
- Outside global address—The IP address that is assigned to a host on the outside network by the owner of the host. The address was allocated from a globally routable address or a network space.

Pool Support for Dynamic NAT

Cisco NX-OS provides pool support for dynamic NAT. Dynamic NAT allows the configuration of a pool of global addresses that can be used to dynamically allocate a global address from the pool for every new translation. The addresses are returned to the pool after the session ages out or is closed. This allows for a more efficient use of addresses based on requirements.

Support for PAT includes the use of the global address pool. This further optimizes IP address utilization. PAT exhausts one IP address at a time with the use of port numbers. If no port is available from the appropriate group and more than one IP address is configured, PAT moves to the next IP address and gets the allocation based on the user defined pool (ignoring the source port or attempting to preserve it).

With dynamic NAT and PAT, each host uses a different address or port for each subsequent translation. The main difference between dynamic NAT and static NAT is that static NAT allows a remote host to initiate a connection to a translated host if an access list exists that allows it, while dynamic NAT does not.

When dynamic NAT is configured to use a pool of IP addresses, that are not locally available or configured locally, the out-to-in traffic is considered as DEST MISS. Due to this behavior, the `show system internal access-list dest-miss stats` command output displays increment in DEST MISS counters. The DEST MISS statistics is supported from Cisco NX-OS Release 9.3(5) onwards.

Static and Dynamic Twice NAT Overview

When both the source IP address and the destination IP address are translated as a single packet that goes through a Network Address Translation (NAT) device, it is referred to as twice NAT. Twice NAT is supported for static and dynamic translations.

Twice NAT allows you to configure two NAT translations (one inside and one outside) as part of a group of translations. These translations can be applied to a single packet as it flows through a NAT device. When you

add two translations as part of a group, both the individual translations and the combined translation take effect.

A NAT inside translation modifies the source IP address and port number when a packet flows from inside to outside. It modifies the destination IP address and port number when the packet returns from outside to inside. NAT outside translation modifies the source IP address and port number when the packet flows from outside to inside, and it modifies the destination IP address and port number when the packet returns from inside to outside.

Without twice NAT, only one of the translation rules is applied on a packet, either the source IP address and port number or the destination IP address and port number.

Static NAT translations that belong to the same group are considered for twice NAT configuration. If a static configuration does not have a configured group ID, the twice NAT configuration will not work. All inside and outside NAT translations that belong to a single group that is identified by the group ID are paired to form twice NAT translations.

Dynamic twice NAT translations dynamically select the source IP address and port number information from pre-defined **ip nat pool** or **interface overload** configurations. Packet filtration is done by configuring ACLs, and traffic must originate from the dynamic NAT translation rule direction such that source translation is done by using dynamic NAT rules.

Dynamic twice NAT allows you to configure two NAT translations (one inside and one outside) as part of a group of translations. One translation must be dynamic and other translation must be static. When these two translations are part of a group of translations, both the translations can be applied on a single packet as it goes through the NAT device either from inside to outside or from outside to inside.

VRF Aware NAT

The VRF aware NAT feature enables a switch to understand an address space in a VRF (virtual routing and forwarding instances) and to translate the packet. This allows the NAT feature to translate traffic in an overlapping address space that is used between two VRFs.

Notes for VRF aware NAT:

- The VRF aware NAT feature is supported on N9K-9408PC-CFP2, N9K-X9564PX, N9K-C9272Q, N9K-C9272Q, N9K-X9464TX, N9K-X9464TX2, N9K-X9564TX, N9K-X9464PX, N9K-X9536PQ, N9K-X9636PQ, N9K-X9432PQ, N9K-C9332PQ, N9K-C9372PX, N9K-C9372PX-E, N9K-C9372TX, N9K-C9372TX-E, N9K-C93120TX.
- The VRF aware NAT feature is not supported on the Cisco Nexus 9300-EX, 9300-FX, 9300-FX2 and 9300-GX platform switches.



Note This is a NAT TCAM limitation for the Cisco Nexus 9300-EX and 9300-FX platform switches. NAT TCAM is not VRF aware. NAT does not work with overlapping IP addresses on Cisco Nexus 9300-EX, 9300-FX, 9300-FX2 and 9300-GX platform switches.

- Traffic flowing from one non-default-vrf to another non-default-vrf is not translated. (For example, vrfA to vrfB.)

- For traffic flowing from a VRF to a global-VRF, a nat-outside configuration is not supported on a non-default VRF interface.
- VRF aware NAT is supported by static and dynamic NAT configurations.
 - When traffic is configured to flow from a non-default VRF (inside) to a default VRF (outside), the **match-in-vrf** option of the **ip nat** command cannot be specified.
 - When traffic is configured to flow from a non-default VRF (inside) to the same non-default VRF (outside), the **match-in-vrf** option of the **ip nat** command must be specified.

The following is an example configuration:

```
Switch(config)# ip nat inside source {list <acl-name>} {pool <pool-name> [vrf
<vrf-name> [match-in-vrf]] [overload] | interface <globalAddrInterface> [vrf
<vrf-name> [match-in-vrf]] overload} [group <group-id> dynamic]
```

```
Switch(config)#ip nat outside source list <acl-name> pool <pool-name> [vrf <vrf-name>
[match-in-vrf]] [group <group-id> dynamic]
```

- VRF aware NAT does not support fragmented packets.
- VRF aware NAT does not support application layer translations.

Therefore, Layer 4 and other embedded IPs are not translated and the following will fail:

- FTP
 - ICMP failures
 - IPSec
 - HTTPS
- VRF aware NAT supports NAT or VACL on an interface. (However, both features cannot be supported at the same time on an interface.)
 - VRF aware NAT supports egress ACLs that are applied to the original packet, not on the NAT translated packet.
 - VRF aware NAT supports only the default VRF.
 - VRF aware NAT does not provide MIB support.
 - VRF aware NAT does not provide DCNM support.
 - VRF aware NAT supports only a single global VDC.
 - VRF aware NAT does not support the active/standby supervisor model.
 - VRFs with overlapping subnets cannot go to a common destination without NAT. However, you can achieve this functionality with inter-VRF NAT on dynamic NAT rule configuration. Static NAT configuration is not supported for overlapping address.

Guidelines and Limitations for Static NAT

Static NAT has the following configuration guidelines and limitations:

- For Broadcom-based Cisco Nexus 9000 Series switches, if the route to your inside global address on the translating device is reachable via the outside interface, packets for Network Address Translated flows coming from outside to inside get software forwarded, duplicated, and looped in the network. For this situation, you must enter the **add-route** CLI argument on the end of the NAT configuration for this flow. For example, **ip nat inside source static 192.168.1.1 172.16.1.1 add-route**.
- The static NAT feature over vPC is not supported on Cisco Nexus 9300 platform switches.
- **show** commands with the **internal** keyword are not supported.
- The static NAT feature is supported on Cisco Nexus 9300 platform switches.
- The static NAT feature is supported on Cisco Nexus 9200 platform switches.
- On Cisco Nexus 9200 and 9300-EX platform switches, the **add-route** option is required for both inside and outside policies.



Note NAT is not supported on Cisco Nexus 9500 platform switches.

- NAT supports up to 1024 translations which include both static and dynamic NAT.
- If the translated IP is part of the outside interface subnet, then use the **ip proxy-arp** command on the NAT outside interface. If the **add-route** keyword is used, **ip proxy-arp** should be enabled.
- NAT and sFlow are not supported on the same port.
- The Cisco Nexus device supports NAT on the following interface types:
 - Switch Virtual Interfaces (SVIs)
 - Routed ports
 - Layer 3 and Layer 3 subinterfaces.
- NAT is supported on the default Virtual Routing and Forwarding (VRF) table only.
- NAT is supported for IPv4 Unicast only.
- The Cisco Nexus device does not support the following:
 - Software translation. All translations are done in the hardware.
 - NAT and VXLAN routing.
 - Application layer translation. Layer 4 and other embedded IPs are not translated, including FTP, ICMP failures, IPSec, and HTTPs.
 - NAT and VLAN Access Control Lists (VACLs) that are configured on an interface at the same time.
 - PAT translation of fragmented IP packets.

- NAT translation on software forwarded packets. For example, packets with IP-options are not NAT translated.
- By default no TCAM entries are allocated for the NAT feature. You allocate the TCAM size for the NAT feature by adjusting the TCAM size of other features. The TCAM can be allocated with the **hardware access-list tcam region nat tcam-size** command.
- HSRP and VRRP are supported only on Static NAT.
- If an IP address is used for Static NAT or PAT translations, it cannot be used for any other purpose. For example, it cannot be assigned to an interface.
- For Static NAT, the outside global IP address should be different from the outside interface IP address.
- NAT statistics are not available.
- When configuring a large number of translations (more than 100), it is faster to configure the translations before configuring the NAT interfaces.
- For 7.0(3)I4(1) and later, NAT supports (non-disruptive) In Service Software Upgrade (ISSU).
- NAT configurations such as **ip nat inside** or **ip nat outside** are not supported on loopback interfaces.

Restrictions for Dynamic NAT

The following restrictions apply to dynamic Network Address Translation (NAT):

- For Broadcom-based Cisco Nexus 9000 Series switches, if the route to your inside global address on the translating device is reachable via the outside interface, packets for Network Address Translated flows coming from outside to inside get software forwarded, duplicated, and looped in the network. For this situation, you must enter the **add-route** CLI argument on the end of the NAT configuration for this flow. For example, **ip nat inside source static 192.168.1.1 172.16.1.1 add-route**.
- VRF aware NAT is not supported for overlapping inside/outside IP subnet addresses on Cisco Nexus 9200 and 9300-EX platform switches.
- **show** commands with the **internal** keyword are not supported.
- The dynamic NAT feature is supported on Cisco Nexus 9300 platform switches.
- The dynamic NAT feature is supported on Cisco Nexus 9200 platform switches.
- On Cisco Nexus 9200 and 9300-EX, -FX, -FX2, -FX3, -FXP, -GX platform switches, the **add-route** option is required for both inside and outside policies.
-
- VXLAN routing is not supported on Cisco Nexus devices.
- Fragmented packets are not supported.
- Application layer gateway (ALG) translations are not supported. ALG, also known as application-level gateway, is an application that translates IP address information inside the payload of an application packet.
- Egress ACLs are not applied to translated packets.

- Nondefault virtual routing and forwarding (VRF) instances are not supported.
- MIBs are not supported.
- Cisco Data Center Network Manager (DCNM) is not supported.
- Multiple global virtual device contexts (VDCs) are not supported on Cisco Nexus devices.
- Dynamic NAT translations are not synchronized with active and standby devices.
- Stateful NAT is not supported. However, NAT and Hot Standby Router Protocol (HSRP) can coexist.
- The timeout value for take up to the configured time-out + 119 seconds.
- Normally, ICMP NAT flows time out after the expiration of the configured sampling-timeout and translation-timeout. However, when ICMP NAT flows present in the switch become idle, they time out immediately after the expiration of the sampling-timeout configured.
- Beginning with Cisco NX-OS Release 7.0(3)I5(2), hardware programming is introduced for ICMP on Cisco Nexus 9300 platform switches. Therefore, the ICMP entries consume the TCAM resources in the hardware. Because ICMP is in the hardware, the maximum limit for NAT translation in Cisco Nexus platform Series switches is changed to 1024. Maximum of 100 ICMP entries are allowed to make the best usage of the resources.
- When creating a new translation on a Cisco Nexus 9000 Series switch, the flow is software forwarded until the translation is programmed in the hardware, which might take a few seconds. During this period, there is no translation entry for the inside global address. Therefore, returning traffic is dropped. To overcome this limitation, create a loopback interface and give it an IP address that belongs to the NAT pool.
- For dynamic NAT, pool overload and interface overload are not supported for the outside NAT.
- Because the NAT overload uses PBR (Policy-Based Routing), the maximum number of available next-hop entries in the PBR table determines NAT scale. If the number of NAT inside interfaces are within the range of available next-hops entries in the PBR table, the maximum NAT translation scale remains same. Otherwise, the maximum number of supported translations may reduce. PBR and NAT-overload are not mutually exclusive; they are mutually limiting.
- The Cisco Nexus devices does not support NAT and VLAN Access Control Lists (VACLs) that are configured on an interface at the same time.
- NAT configurations such as **ip nat inside** or **ip nat outside** are not supported on loopback interfaces.
- The dynamic NAT feature over vPC is not supported.
- If traffic ingresses a PBR enabled interface, and has a NAT entry, the traffic will be routed via PBR but the IP address will not be translated.

Guidelines and Limitations for Dynamic Twice NAT

For Broadcom-based Cisco Nexus 9000 Series switches, if the route to your inside global address on the translating device is reachable via the outside interface, packets for Network Address Translated flows coming from outside to inside get software forwarded, duplicated, and looped in the network. For this situation, you must enter the **add-route** CLI argument on the end of the NAT configuration for this flow. For example, **ip nat inside source static 192.168.1.1 172.16.1.1 add-route**.

IP packets without TCP/UDP/ICMP headers are not translated with dynamic NAT.

In dynamic twice NAT, if dynamic NAT flows are not created before creating static NAT flows, dynamic twice NAT flows are not created correctly.

When an empty ACL is created, the default rule of **permit ip any any** is configured. The NAT-ACL does not match further ACL entries if the first ACL is blank.

Configuring Static NAT

Enabling Static NAT

Procedure

	Command or Action	Purpose
Step 1	switch# configure terminal	Enters global configuration mode.
Step 2	switch(config)# feature nat	Enables the static NAT feature on the device.
Step 3	switch(config)# copy running-config startup-config	Saves the change persistently through reboots and restarts by copying the running configuration to the startup configuration.

Configuring Static NAT on an Interface

Procedure

	Command or Action	Purpose
Step 1	switch# configure terminal	Enters global configuration mode.
Step 2	switch(config)# interface <i>type slot/port</i>	Specifies an interface to configure, and enters interface configuration mode.
Step 3	switch(config-if)# ip nat { inside outside }	Specifies the interface as inside or outside. Note Only packets that arrive on a marked interface can be translated. This configuration is not supported on loopback interface.
Step 4	(Optional) switch(config)# copy running-config startup-config	Saves the change persistently through reboots and restarts by copying the running configuration to the startup configuration.

Example

This example shows how to configure an interface with static NAT from the inside:

```
switch# configure terminal
switch(config)# interface ethernet 1/4
switch(config-if)# ip nat inside
```

Enabling Static NAT for an Inside Source Address

For inside source translation, the traffic flows from inside interface to the outside interface. NAT translates the inside local IP address to the inside global IP address. On the return traffic, the destination inside global IP address gets translated back to the inside local IP address.



Note When the Cisco Nexus device is configured to translate an inside source IP address (Src:ip1) to an outside source IP address (newSrc:ip2), the Cisco Nexus device implicitly adds a translation for an outside destination IP address (Dst: ip2) to an inside destination IP address (newDst: ip1).

Procedure

	Command or Action	Purpose
Step 1	switch# configure terminal	Enters global configuration mode.
Step 2	switch(config)# ip nat inside source static <i>local-ip-address global-ip-address [vrf</i> <i>vrf-name] [match-in-vrf] [group group-id]</i>	Configures static NAT to translate the inside local address to the inside global address or to translate the opposite (the inside global traffic to the inside local traffic). Specifying group specifies the group to which this translation belongs on the static twice NAT. Note While performing twice NAT configuration in Cisco Nexus 9000 Series switches, you cannot use the same group ID across different VRFs. A unique group ID should be used for unique twice NAT rules.
Step 3	(Optional) switch(config)# copy running-config startup-config	Saves the change persistently through reboots and restarts by copying the running configuration to the startup configuration.

Example

This example shows how to configure static NAT for an inside source address:

```
switch# configure terminal
switch(config)# ip nat inside source static 1.1.1.1 5.5.5.5
switch(config)# copy running-config startup-config
```

Enabling Static NAT for an Outside Source Address

For outside source translation, the traffic flows from the outside interface to the inside interface. NAT translates the outside global IP address to the outside local IP address. On the return traffic, the destination outside local IP address gets translated back to outside global IP address.

Procedure

	Command or Action	Purpose
Step 1	switch# configure terminal	Enters global configuration mode.
Step 2	switch(config)# ip nat outside source static <i>outsideGlobalIP outsideLocalIP</i> [vrf vrf-name] [match-in-vrf] [group group-id] [dynamic] [add-route]	Configures static NAT to translate the outside global address to the outside local address or to translate the opposite (the outside local traffic to the outside global traffic). Specifying group specifies the group to which this translation belongs on the static twice NAT. When an inside translation without ports is configured, an implicit add route is performed. The original add route functionality is an option while configuring an outside translation.
Step 3	(Optional) switch(config)# copy running-config startup-config	Saves the change persistently through reboots and restarts by copying the running configuration to the startup configuration.

Example

This example show how to configure static NAT for an outside source address:

```
switch# configure terminal
switch(config)# ip nat outside source static 2.2.2.2 6.6.6.6
switch(config)# copy running-config startup-config
```

Configuring Static PAT for an Inside Source Address

You can map services to specific inside hosts using Port Address Translation (PAT).

Procedure

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

	Command or Action	Purpose
Step 2	switch(config)# ip nat inside source static { <i>inside-local-address</i> <i>inside-global-address</i> { tcp udp } <i>inside-local-address</i> { <i>local-tcp-port</i> <i>local-udp-port</i> } <i>inside-global-address</i> { <i>global-tcp-port</i> <i>global-udp-port</i> }} { vrf <i>vrf-name</i> { match-in-vrf } { group <i>group-id</i> }	Maps static NAT to an inside local port to an inside global port.
Step 3	(Optional) switch(config)# copy running-config startup-config	Saves the change persistently through reboots and restarts by copying the running configuration to the startup configuration.

Example

This example shows how to map UDP services to a specific inside source address and UDP port:

```
switch# configure terminal
switch(config)# ip nat inside source static udp 20.1.9.2 63 35.48.35.48 130
switch(config)# copy running-config startup-config
```

Configuring Static PAT for an Outside Source Address

You can map services to specific outside hosts using Port Address Translation (PAT).

Procedure

	Command or Action	Purpose
Step 1	switch# configure terminal	Enters global configuration mode.
Step 2	switch(config)# ip nat outside source static { <i>outside-global-address</i> <i>outside-local-address</i> { tcp udp } <i>outside-global-address</i> { <i>global-tcp-port</i> <i>global-udp-port</i> } <i>outside-local-address</i> { <i>global-tcp-port</i> <i>global-udp-port</i> }} { group <i>group-id</i> } { add-route } { vrf <i>vrf-name</i> { match-in-vrf }}	Maps static NAT to an outside global port to an outside local port. Specifying group specifies the group to which this translation belongs on the static twice NAT. When an inside translation without ports is configured, an implicit add route is performed. The original add route functionality is an option while configuration an outside translation.
Step 3	(Optional) switch(config)# copy running-config startup-config	Saves the change persistently through reboots and restarts by copying the running configuration to the startup configuration.

Example

This example shows how to map TCP services to a specific outside source address and TCP port:

```
switch# configure terminal
switch(config)# ip nat outside source static tcp 20.1.9.2 63 35.48.35.48 130
switch(config)# copy running-config startup-config
```

Configuring Static Twice NAT

All translations within the same group are considered for creating static twice Network Address Translation (NAT) rules.

Procedure

	Command or Action	Purpose
Step 1	enable Example: <pre>switch> enable</pre>	Enables privileged EXEC mode. <ul style="list-style-type: none"> • Enter your password if prompted.
Step 2	configure terminal Example: <pre>switch# configure terminal</pre>	Enters privileged EXEC mode.
Step 3	ip nat inside source static <i>inside-local-ip-address</i> <i>inside-global-ip-address</i> [group <i>group-id</i>] [add-route] Example: <pre>switch(config)# ip nat inside source static 10.1.1.1 192.168.34.4 group 4</pre>	Configures static twice NAT to translate an inside local IP address to the corresponding inside global IP address. <ul style="list-style-type: none"> • The group keyword determines the group to which a translation belongs.
Step 4	ip nat outside source static <i>outside-global-ip-address</i> <i>outside-local-ip-address</i> [group <i>group-id</i>] [add-route] Example: <pre>switch(config)# ip nat outside source static 209.165.201.1 10.3.2.42 group 4 add-route</pre>	Configures static twice NAT to translate an outside global IP address to the corresponding outside local IP address. <ul style="list-style-type: none"> • The group keyword determines the group to which a translation belongs.
Step 5	interface <i>type number</i> Example: <pre>switch(config)# interface ethernet 1/2</pre>	Configures an interface and enters interface configuration mode.
Step 6	ip address <i>ip-address mask</i> Example: <pre>switch(config-if)# ip address 10.2.4.1 255.255.255.0</pre>	Sets a primary IP address for an interface.
Step 7	ip nat inside Example: <pre>switch(config-if)# ip nat inside</pre>	Connects the interface to an inside network, which is subject to NAT. <p>Note Configuration not supported on loopback interface.</p>

	Command or Action	Purpose
Step 8	exit Example: switch(config-if)# exit	Exits interface configuration mode and returns to global configuration mode.
Step 9	interface type number Example: switch(config)# interface ethernet 1/1	Configures an interface and enters interface configuration mode.
Step 10	ip address ip-address mask Example: switch(config-if)# ip address 10.5.7.9 255.255.255.0	Sets a primary IP address for an interface.
Step 11	ip nat outside Example: switch(config-if)# ip nat outside	Connects the interface to an outside network, which is subject to NAT. Note Configuration not supported on loopback interface.
Step 12	end Example: switch(config-if)# end	Exits interface configuration mode and returns to privileged EXEC mode.

Configuration Example for Static NAT and PAT

This example shows the configuration for static NAT:

```
ip nat inside source static 103.1.1.1 11.3.1.1
ip nat inside source static 139.1.1.1 11.39.1.1
ip nat inside source static 141.1.1.1 11.41.1.1
ip nat inside source static 149.1.1.1 95.1.1.1
ip nat inside source static 149.2.1.1 96.1.1.1
ip nat outside source static 95.3.1.1 95.4.1.1
ip nat outside source static 96.3.1.1 96.4.1.1
ip nat outside source static 102.1.2.1 51.1.2.1
ip nat outside source static 104.1.1.1 51.3.1.1
ip nat outside source static 140.1.1.1 51.40.1.1
```

This example shows the configuration for static PAT:

```
ip nat inside source static tcp 10.11.1.1 1 210.11.1.1 101
ip nat inside source static tcp 10.11.1.1 2 210.11.1.1 201
ip nat inside source static tcp 10.11.1.1 3 210.11.1.1 301
ip nat inside source static tcp 10.11.1.1 4 210.11.1.1 401
ip nat inside source static tcp 10.11.1.1 5 210.11.1.1 501
ip nat inside source static tcp 10.11.1.1 6 210.11.1.1 601
ip nat inside source static tcp 10.11.1.1 7 210.11.1.1 701
ip nat inside source static tcp 10.11.1.1 8 210.11.1.1 801
ip nat inside source static tcp 10.11.1.1 9 210.11.1.1 901
ip nat inside source static tcp 10.11.1.1 10 210.11.1.1 1001
```

Example: Configuring Static Twice NAT

```
ip nat inside source static tcp 10.11.1.1 11 210.11.1.1 1101
ip nat inside source static tcp 10.11.1.1 12 210.11.1.1 1201
```

Example: Configuring Static Twice NAT

The following example shows how to configure the inside source and outside source static twice NAT configurations:

```
Switch> enable
Switch# configure terminal
Switch(config)# ip nat inside source static 10.1.1.1 192.168.34.4 group 4
Switch(config)# ip nat outside source static 209.165.201.1 10.3.2.42 group 4
Switch(config)# interface ethernet 1/2
Switch(config-if)# ip address 10.2.4.1 255.255.255.0
Switch(config-if)# ip nat inside
switch(config-if)# exit
switch(config)# interface ethernet 1/1
switch(config-if)# ip address 10.5.7.9 255.255.255.0
switch(config-if)# ip nat outside
Switch(config-if)# end
```

Verifying the Static NAT Configuration

To display the static NAT configuration, perform this task:

Procedure

	Command or Action	Purpose
Step 1	switch# show ip nat translations	Shows the translations for the inside global, inside local, outside local, and outside global IP addresses.

Example

This example shows how to display the static NAT configuration:

```
switch# sh ip nat translations
Pro Inside global      Inside local      Outside local      Outside global
--- ---              ---              ---              ---
--- ---              ---              51.3.1.1          104.1.1.1
--- ---              ---              95.4.1.1          95.3.1.1
--- ---              ---              96.4.1.1          96.3.1.1
--- ---              ---              51.40.1.1         140.1.1.1
--- ---              ---              51.42.1.1         142.1.2.1
--- ---              ---              51.1.2.1          102.1.2.1
--- 11.1.1.1          101.1.1.1         ---              ---
--- 11.3.1.1          103.1.1.1         ---              ---
--- 11.39.1.1         139.1.1.1         ---              ---
--- 11.41.1.1         141.1.1.1         ---              ---
--- 95.1.1.1          149.1.1.1         ---              ---
--- 96.1.1.1          149.2.1.1         ---              ---
--- 130.1.1.1:590     30.1.1.100:5000   ---              ---
--- 130.2.1.1:590     30.2.1.100:5000   ---              ---
```

```

130.3.1.1:590      30.3.1.100:5000    ---      ---
130.4.1.1:590      30.4.1.100:5000    ---      ---
130.1.1.1:591      30.1.1.101:5000    ---      ---

```

```

switch# sh ip nat translations verbose
Pro Inside global      Inside local      Outside local      Outside global
any ---                ---                22.1.1.3           22.1.1.2
  Flags:0x200009 time-left(secs):-1 id:0 state:0x0 grp_id:10
any 11.1.1.130         11.1.1.3          ---                ---
  Flags:0x1 time-left(secs):-1 id:0 state:0x0 grp_id:0
any 11.1.1.133         11.1.1.33        ---                ---
  Flags:0x1 time-left(secs):-1 id:0 state:0x0 grp_id:10
any 11.1.1.133         11.1.1.33        22.1.1.3           22.1.1.2
  Flags:0x200009 time-left(secs):-1 id:0 state:0x0 grp_id:0
tcp 10.1.1.100:64490    10.1.1.2:0        20.1.1.2:0         20.1.1.2:0
  Flags:0x82 time-left(secs):43192 id:31 state:0x3 grp_id:0 vrf: default
N9300-1#

```

Configuring Dynamic NAT

Configuring Dynamic Translation and Translation Timeouts

Procedure

	Command or Action	Purpose
Step 1	enable Example: Switch> enable	Enables privileged EXEC mode. <ul style="list-style-type: none">• Enter your password if prompted.
Step 2	configure terminal Example: Switch# configure terminal	Enters global configuration mode.
Step 3	ip access-list <i>access-list-name</i> Example: Switch(config)# ip access-list acl1	Defines an access list and enters access-list configuration mode.
Step 4	permit <i>protocol source source-wildcard any</i> Example: Switch(config-acl)# permit ip 10.111.11.0/24 any	Sets conditions in an IP access list that permit traffic matching the conditions.
Step 5	deny <i>protocol source source-wildcard any</i> Example: Switch(config-acl)# deny udp 10.111.11.100/32 any	Sets conditions in an IP access list that deny packets from entering a network.

	Command or Action	Purpose
Step 6	exit Example: Switch(config-acl)# exit	Exits access-list configuration mode and returns to global configuration mode.
Step 7	ip nat inside source list <i>access-list-name</i> interface <i>type number</i> [vrf <i>vrf-name</i>] [match-in-vrf] [add-route] [overload] Example: Switch(config)# ip nat inside source list acl1 interface ethernet 1/1 overload	Establishes dynamic source translation by specifying the access list defined in Step 3.
Step 8	interface <i>type number</i> Example: Switch(config)# interface ethernet 1/4	Configures an interface and enters interface configuration mode.
Step 9	ip address <i>ip-address mask</i> Example: Switch(config-if)# ip address 10.111.11.39 255.255.255.0	Sets a primary IP address for the interface.
Step 10	ip nat inside Example: Switch(config-if)# ip nat inside	Connects the interface to an inside network, which is subject to NAT. Note Configuration not supported on loopback interface.
Step 11	exit Example: Switch(config-if)# exit	Exits interface configuration mode and returns to global configuration mode.
Step 12	interface <i>type number</i> Example: Switch(config)# interface ethernet 1/1	Configures an interface and enters interface configuration mode.
Step 13	ip address <i>ip-address mask</i> Example: Switch(config-if)# ip address 172.16.232.182 255.255.255.240	Sets a primary IP address for an interface.
Step 14	ip nat outside Example: Switch(config-if)# ip nat outside	Connects the interface to an outside network. Note Configuration not supported on loopback interface.

	Command or Action	Purpose
Step 15	exit Example: Switch(config-if)# exit	Exits interface configuration mode and returns to global configuration mode.
Step 16	ip nat translation max-entries <i>number-of-entries</i> Example: Switch(config)# ip nat translation max-entries 300	Specifies the maximum number of dynamic NAT translations. The number of entries can be between 1 and 1023.
Step 17	ip nat translation timeout <i>seconds</i> Example: switch(config)# ip nat translation timeout 13000	Specifies the timeout value for dynamic NAT translations.
Step 18	end Example: Switch(config)# end	Exits global configuration mode and returns to privileged EXEC mode.

Configuring Dynamic NAT Pool

You can create a NAT pool by either defining the range of IP addresses in a single **ip nat pool** command or by using the **ip nat pool** and **address** commands

Procedure

	Command or Action	Purpose
Step 1	switch# configure terminal	Enters global configuration mode.
Step 2	switch(config)# feature nat	Enables the NAT feature on the device.
Step 3	switch(config)# ip nat pool <i>pool-name</i> [<i>startip endip</i>] { prefix <i>prefix-length</i> netmask <i>network-mask</i> }	Creates a NAT pool with a range of global IP addresses. The IP addresses are filtered by using either a prefix length or a network mask.
Step 4	(Optional) switch(config-ipnat-pool)# address <i>startip endip</i>	Specifies the range of global IP addresses if they were not specified during creation of the pool.
Step 5	(Optional) switch(config)# no ip nat pool <i>pool-name</i>	Deletes the specified NAT pool.

Example

This example shows how to create a NAT pool with a prefix length:

```
switch# configure terminal
switch(config)# ip nat pool pool1 30.1.1.1 30.1.1.2 prefix-length 24
switch(config)#
```

This example shows how to create a NAT pool with a network mask:

```
switch# configure terminal
switch(config)# ip nat pool pool5 20.1.1.1 20.1.1.5 netmask 255.0.255.0
switch(config)#
```

This example shows how to create a NAT pool and define the range of global IP addresses using the **ip nat pool** and **address** commands:

```
switch# configure terminal
switch(config)# ip nat pool pool7 netmask 255.255.0.0
switch(config-ipnat-pool)# address 40.1.1.1 40.1.1.5
switch(config-ipnat-pool)#
```

This example shows how to delete a NAT pool:

```
switch# configure terminal
switch(config)# no ip nat pool pool4
switch(config)#
```

Configuring Source Lists

You can configure a source list of IP addresses for the inside interface and the outside interface.

Before you begin

Ensure that you configure a pool before configuring the source list for the pool.

Procedure

	Command or Action	Purpose
Step 1	switch# configure terminal	Enters global configuration mode.
Step 2	(Optional) switch# ip nat inside source list <i>list-name pool pool-name [overload]</i>	Creates a NAT inside source list with pool with or without overloading.
Step 3	(Optional) switch# ip nat outside source list <i>list-name pool pool-name [add-route]</i>	Creates a NAT outside source list with pool without overloading.

Example

This example shows how to create a NAT inside source list with pool without overloading:

```
switch# configure terminal
switch(config)# ip nat inside source list list1 pool pool1
switch(config)#
```

This example shows how to create a NAT inside source list with pool with overloading:

```
switch# configure terminal
switch(config)# ip nat inside source list list2 pool pool2 overload
switch(config)#
```

This example shows how to create a NAT outside source list with pool without overloading:

```
switch# configure terminal
switch(config)# ip nat outside source list list3 pool pool3
switch(config)#
```

Configuring Dynamic Twice NAT for an Inside Source Address

For an inside source address translation, the traffic flows from the inside interface to the outside interface. You can configure dynamic twice NAT for an inside source address.

Before you begin

Ensure that you enable NAT on the switch.

Procedure

	Command or Action	Purpose
Step 1	switch# configure terminal	Enters global configuration mode.
Step 2	switch(config)# ip nat outside source static <i>outside-global-ip-address</i> <i>outside-local-ip-address</i> [tcp udp] <i>outside-global-ip-address outside-global-port</i> <i>outside-local-ip-address outside-local-port</i> [group group-id] [dynamic] [add-route]	Configures static NAT to translate an outside global address to an inside local address or to translate inside local traffic to inside global traffic. The group keyword determines the group to which a translation belongs.
Step 3	switch(config)# ip nat inside source list <i>access-list-name</i> [interface type slot/port overload pool pool-name overload] [group <i>group-id</i>] [dynamic] [add-route]	Establishes dynamic source translation by creating a NAT inside source list with pool with or without overloading. The group keyword determines the group to which a translation belongs.
Step 4	switch(config)# ip nat pool <i>pool-name</i> [<i>startip</i> <i>endip</i>] { prefix prefix-length netmask <i>network-mask</i> }	Creates a NAT pool with a range of global IP addresses. The IP addresses are filtered by using either a prefix length or a network mask.
Step 5	switch(config)# interface <i>type slot/port</i>	Configures an interface and enters interface configuration mode.
Step 6	switch(config-if)# ip nat outside	Connects the interface to an outside network.

	Command or Action	Purpose
Step 7	switch(config-if)# exit	Exits interface configuration mode and returns to global configuration mode.
Step 8	switch(config)# interface <i>type slot/port</i>	Configures an interface and enters interface configuration mode.
Step 9	switch(config-if)# ip nat inside	Connects the interface to an inside network, which is subject to NAT.

Example

This example shows how to configure dynamic twice NAT for an inside source address:

```
switch# configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
switch(config)# ip nat outside source static 2.2.2.2 4.4.4.4 group 20 dynamic
switch(config)# ip nat inside source list acl_1 pool pool_1 overload group 20 dynamic
switch(config)# ip nat pool pool_1 3.3.3.3 3.3.3.10 prefix-length 24
switch(config)# interface Ethernet1/8
switch(config-if)# ip nat outside
switch(config-if)# exit
switch(config)# interface Ethernet1/15
switch(config-if)# ip nat inside
```

Configuring Dynamic Twice NAT for an Outside Source Address

For an outside source address translation, the traffic flows from the outside interface to the inside interface. You can configure dynamic twice NAT for an outside source address.

Before you begin

Ensure that you enable NAT on the switch.

Procedure

	Command or Action	Purpose
Step 1	switch# configure terminal	Enters global configuration mode.
Step 2	switch(config)# ip nat inside source static <i>inside-local-ip-address inside-global-ip-address</i> [tcp udp] <i>inside-local-ip-address local-port inside-global-ip-address global-port</i> [group <i>group-id</i>] [dynamic] [add-route]	Configures static NAT to translate an inside global address to an inside local address or to translate inside local traffic to inside global traffic. The group keyword determines the group to which a translation belongs.
Step 3	switch(config)# ip nat outside source list <i>access-list-name pool pool-name</i> [group <i>group-id</i>] dynamic [add-route]	Establishes dynamic source translation by creating a NAT outside source list with pool with or without overloading.

	Command or Action	Purpose
Step 4	switch(config)# ip nat pool <i>pool-name</i> [<i>startip endip</i>] { prefix <i>prefix-length</i> netmask <i>network-mask</i> }	Creates a NAT pool with a range of global IP addresses. The IP addresses are filtered by using either a prefix length or a network mask.
Step 5	switch(config)# interface <i>type slot/port</i>	Configures an interface and enters interface configuration mode.
Step 6	switch(config-if)# ip nat outside	Connects the interface to an outside network.
Step 7	switch(config-if)# exit	Exits interface configuration mode and returns to global configuration mode.
Step 8	switch(config)# interface <i>type slot/port</i>	Configures an interface and enters interface configuration mode.
Step 9	switch(config-if)# ip nat inside	Connects the interface to an inside network, which is subject to NAT.

Example

This example shows how to configure dynamic twice NAT for an outside source address:

```
switch# configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
switch(config)# ip nat inside source static 7.7.7.7 5.5.5.5 group 30 dynamic
switch(config)# ip nat outside source list acl_1 pool pool_1 group 30 dynamic
switch(config)# ip nat pool pool_2 4.4.4.4 4.4.4.10 prefix-length 24
switch(config)# interface Ethernet1/6
switch(config-if)# ip nat outside
switch(config-if)# exit
switch(config)# interface Ethernet1/11
switch(config-if)# ip nat inside
```

Clearing Dynamic NAT Translations

To clear dynamic translations, perform the following task:

Command	Purpose
clear ip nat translation [all inside <i>global-ip-address local-ip-address</i> [outside <i>local-ip-address global-ip-address</i>] outside <i>local-ip-address global-ip-address</i>]	Deletes all or specific dynamic NAT translations.

Example

This example shows how to clear all dynamic translations:

```
switch# clear ip nat translation all
```

This example shows how to clear dynamic translations for inside and outside addresses:

```
switch# clear ip nat translation inside 2.2.2.2 4.4.4.4 outside 5.5.5.5 7.7.7.7
```

Verifying Dynamic NAT Configuration

To display dynamic NAT configuration, perform the following tasks:

Command	Purpose
show ip nat translations	Displays active Network Address Translation (NAT) translations. Displays additional information for each translation table entry, including when an entry was created and used.
show run nat	Displays NAT configuration.
show ip nat max	Displays active Network Address Translation (NAT) maximum values.
show ip nat statistics	Monitor NAT statistics.

Example

This example shows how to display IP NAT Max values:

```
switch# show ip nat max

IP NAT Max values
=====
Max Dyn Translations:80
Max all-host:0
No.Static:0
No.Dyn:1
No.Dyn-ICMP:1
=====
Switch(config)#
```

This example shows how to display NAT Statistics:

```
switch# show ip nat statistics

IP NAT Statistics
=====
Stats Collected since: Mon Feb 24 18:27:34 2020
-----
Total active translations: 1
No.Static: 0
No.Dyn: 1
No.Dyn-ICMP: 1
-----
Total expired Translations: 0
SYN timer expired: 0
FIN-RST timer expired: 0
Inactive timer expired: 0
-----
Total Hits: 2          Total Misses: 2
In-Out Hits: 0        In-Out Misses: 2
```

```

Out-In Hits: 2                Out-In Misses: 0
-----
Total SW Translated Packets: 2
In-Out SW Translated: 2
Out-In SW Translated: 0
-----
Total SW Dropped Packets: 0
In-Out SW Dropped: 0
Out-In SW Dropped: 0

Address alloc. failure drop:    0
Port alloc. failure drop:      0
Dyn. Translation max limit drop: 0
ICMP max limit drop:           0
Allhost max limit drop:        0
-----
Total TCP session established: 0
Total TCP session closed:      0
-----
NAT Inside Interfaces: 1
Ethernet1/34

NAT Outside Interfaces: 1
Ethernet1/32
-----
Inside source list:
+++++++

Access list: T2
RefCount: 1
Pool: T2      Overload
Total addresses: 10
Allocated: 1   percentage: 10%
Missed: 0

Outside source list:
+++++++
=====
Switch(config)#
Switch(config)#

```

**No.Dyn-ICMP field is to display the no of icmp dynamic translations , its a subset of "No.Dyn" field.

This example shows how to display running configuration for NAT:

```

switch# show run nat

!Command: show running-config nat
!Time: Wed Apr 23 11:17:43 2014

version 6.0(2)A3(1)
feature nat

ip nat inside source list list1 pool pool1
ip nat inside source list list2 pool pool2 overload
ip nat inside source list list7 pool pool7 overload
ip nat outside source list list3 pool pool3
ip nat pool pool1 30.1.1.1 30.1.1.2 prefix-length 24
ip nat pool pool2 10.1.1.1 10.1.1.2 netmask 255.0.255.0
ip nat pool pool3 30.1.1.1 30.1.1.8 prefix-length 24

```

Example: Configuring Dynamic Translation and Translation Timeouts

```
ip nat pool pool5 20.1.1.1 20.1.1.5 netmask 255.0.255.0
ip nat pool pool7 netmask 255.255.0.0
  address 40.1.1.1 40.1.1.5
```

This example shows how to display active NAT translations:

Inside pool with overload

```
switch# show ip nat translation
Pro Inside global      Inside local      Outside local      Outside global
icmp 20.1.1.3:64762    10.1.1.2:133     20.1.1.1:0        20.1.1.1:0
icmp 20.1.1.3:64763    10.1.1.2:134     20.1.1.1:0        20.1.1.1:0
```

Outside pool without overload

```
switch# show ip nat translation
Pro Inside global      Inside local      Outside local      Outside global
any ---               ---              177.7.1.1:0       77.7.1.64:0
any ---               ---              40.146.1.1:0      40.46.1.64:0
any ---               ---              10.4.146.1:0      10.4.46.64:0
```

Example: Configuring Dynamic Translation and Translation Timeouts

The following example shows how to configure dynamic overload Network Address Translation (NAT) by specifying an access list:

```
Switch> enable
Switch# configure terminal
Switch(config)# ip access-list acl1
Switch(config-acl)# permit ip 10.111.11.0/24 any
Switch(config-acl)# deny udp 10.111.11.100/32 any
Switch(config-acl)# exit
Switch(config)# ip nat inside source list acl1 interface ethernet 1/1 overload
Switch(config)# interface ethernet 1/4
Switch(config-if)# ip address 10.111.11.39 255.255.255.0
Switch(config-if)# ip nat inside
Switch(config-if)# exit
Switch(config)# interface ethernet 1/1
Switch(config-if)# ip address 172.16.232.182 255.255.255.240
Switch(config-if)# ip nat outside
Switch(config-if)# exit
Switch(config)# ip nat translation max-entries 300
Switch(config)# ip nat translation timeout 13000
Switch(config)# end
```




APPENDIX **A**

Configuring Layer 2 Data Center Interconnect

This section contains an example of how to configure a Layer 2 Data Center Interconnect (DCI) with the use of a Virtual Port-Channel (vPC).

- [Overview, on page 365](#)
- [Example of Layer 2 Data Center Interconnect, on page 365](#)

Overview

The purpose of a Data Center Interconnect (DCI) is to extend specific VLANs between different data centers. DCI offers Layer 2 adjacency for servers and Network Attached Storage (NAS) devices that are separated by large distances.

Beginning with Cisco NX-OS release 7.0(3)I2(2), Cisco Nexus 9000 series switches support DCI with FHRP isolation. However DCI with FHRP isolation is not supported on Cisco Nexus 9500 switches with N9K-X9636C-R and N9K-X9636Q-R line cards. Creating a single logical link between multiple sites with vPC allows you to take advantage of the benefits of STP isolation using BPDU filtering across the DCI vPC port-channel. With this configuration, Bridge Protocol Data Unit (BPDU) does not cross between data centers, effectively isolating the STP fault domain between sites.



Note vPC is to interconnect a maximum of two data centers.

Example of Layer 2 Data Center Interconnect

The following is an example configuration of a Layer 2 Data Center Interconnect (DCI) with use of vPC. The example allows for First Hop Redundancy Protocol (FHRP) isolation.




Note vPC and Hot Standby Routing Protocol (HSRP) have already been configured.



Note Link Aggregation Control Protocol (LACP) should be used on the vPC link, which acts as the DCI.

Figure 40: Dual Layer 2/Layer 3 POD Interconnect



In this example, the Layer 3 (L3) gateway is configured on the same vPC pair and acts as the DCI. In order to isolate the Hot Standby Routing Protocol (HSRP), you must configure a Port Access Control List (PACL) on the DCI port-channel and disable HSRP Gratuitous Address Resolution Protocols (ARPs) (GARPs) on the Switched Virtual Interfaces (SVIs) for the VLANs that move across the DCI.

```
ip access-list DENY_HSRP_IP
 10 deny udp any 224.0.0.2/32 eq 1985
 20 deny udp any 224.0.0.102/32 eq 1985
 30 permit ip any any

interface <DCI-Port-Channel>
 ip port access-group DENY_HSRP_IP in

interface Vlan <x>
 no ip arp gratuitous hsrp duplicate
```



APPENDIX **B**


IETF RFCs supported by Cisco NX-OS Interfaces

This appendix lists the IETF RFCs for interfaces supported by Cisco NX-OS.

- [IPv6 RFCs, on page 367](#)

IPv6 RFCs

RFCs	Title
RFC 1981 (7.0(3)I1(1) and later)	<i>Path MTU Discovery for IP version 6</i>
RFC 2373	<i>IP Version 6 Addressing Architecture</i>
RFC 2374	<i>An Aggregatable Global Unicast Address Format</i>
RFC 2460	<i>Internet Protocol, Version 6 (IPv6) Specification</i>
RFC 2462	<i>IPv6 Stateless Address Autoconfiguration</i>
RFC 2464	<i>Transmission of IPv6 Packets over Ethernet Networks</i>
RFC 2467	<i>Transmission of IPv6 Packets over FDDI Networks</i>
RFC 2472	<i>IP Version 6 over PPP</i>
RFC 2492	<i>IPv6 over ATM Networks</i>
RFC 2590	<i>Transmission of IPv6 Packets over Frame Relay Networks Specification</i>
RFC 3021	<i>Using 31-Bit Prefixes on IPv4 Point-to-Point Links</i>
RFC 3152	<i>Delegation of IP6.ARPA</i>
RFC 3162	<i>RADIUS and IPv6</i>
RFC 3513	<i>Internet Protocol Version 6 (IPv6) Addressing Architecture</i>
RFC 3596	<i>DNS Extensions to Support IP version 6</i>
RFC 4193	<i>Unique Local IPv6 Unicast Addresses</i>



APPENDIX **C**

Configuration Limits for Cisco NX-OS Interfaces

The configuration limits are documented in the *Cisco Nexus 9000 Series NX-OS Verified Scalability Guide*.



INDEX

A

address [357–358](#)
admin-shutdown [99](#)
auto-recovery [238, 271, 273](#)
autonomous-system [148](#)

B

bandwidth [42, 187](#)
bfd [148–150](#)
bfd authentication keyed-sha1 keyid [136–137](#)
bfd echo [139](#)
bfd echo-interface loopback [135](#)
bfd interval [134, 136–137, 157–158](#)
bfd per-link [137](#)
bfd slow-timer [134, 138](#)
bfd startup-timer [135](#)
broadcast [101](#)

C

channel-group [183, 185, 194](#)
checkpoint [82](#)
clear counters interface [58, 91](#)
clear counters interface port-channel [213](#)
clear ip nat translation [361](#)
clear ip route [131](#)
clear ipv6 route [131](#)
clear l2protocol tunnel counters [322](#)
clear lacp counters [213](#)
config t [107](#)
copy [30, 32](#)

D

default interface [82](#)
delay [43, 187](#)
delay restore [237, 239](#)
deny [355](#)
description [33, 190, 295](#)
duplex [191](#)
duplex auto [191](#)
duplex full [191](#)

duplex half [191](#)

E

ecmp load balancing [210](#)
enable [352, 355](#)
encapsulation dot1Q [103–104, 110, 112](#)
end [112, 357](#)
errdisable detect cause [20, 36](#)
errdisable detect cause acl-exception [36](#)
errdisable detect cause all [36](#)
errdisable detect cause link-flap [36](#)
errdisable detect cause loopback [36](#)
errdisable recovery cause [20, 37](#)
errdisable recovery cause all [37](#)
errdisable recovery cause bpdguard [37](#)
errdisable recovery cause failed-port-state [37](#)
errdisable recovery cause link-flap [37](#)
errdisable recovery cause loopback [37](#)
errdisable recovery cause miscabling [37](#)
errdisable recovery cause psecure-violation [37](#)
errdisable recovery cause security-violation [37](#)
errdisable recovery cause storm-control [37](#)
errdisable recovery cause uddi [37](#)
errdisable recovery cause vpc-peerlink [37](#)
errdisable recovery interval [20, 38](#)
ethernet [33](#)
Ethernet [11](#)

F

feature bfd [133](#)
feature eigrp [43](#)
feature interface-vlan [85, 105](#)
feature isis [112](#)
feature lacp [193](#)
feature nat [348, 357](#)
feature tunnel [294](#)
feature vpc [253](#)

G

graceful consistency-check [263](#)

H

hardware access-list team region nat **346**
 how l2protocol tunnel summary **323**
 hsrp bfd **152**
 hsrp bfd all-interfaces **152**

I

include bfd **133**
 interface **33, 40, 44, 83, 107, 116–117, 136, 275, 348, 352–353, 356, 359–361**
 interface breakout **11**
 interface ether **57**
 interface ethernet **34, 39, 41–42, 47, 49, 72, 74, 78, 81, 100, 102, 107, 109–110, 112, 120, 316, 318, 320**
 interface loopback **108, 111**
 interface overload **343**
 interface port-channel **78, 81, 83, 104, 107, 137, 156–158, 181, 186–190, 195–196, 201–202, 204–206, 208, 259, 265**
 interface tunnel **295, 297, 299, 301–302, 305**
 interface vlan **86, 105, 107**
 interfaces-vlan **237, 239**
 ip **29**
 ip access-list **355**
 ip address **101–102, 104–105, 108, 111, 117, 157, 297, 305, 352–353, 356**
 ip address dhcp **97**
 ip arp synchronize **234**
 ip eigrp **148, 155**
 ip name-server **98**
 ip nat **344**
 ip nat inside **348, 352, 356, 360–361**
 ip nat inside source list **356, 358–359**
 ip nat inside source static **349, 351–352, 360**
 ip nat outside **348, 353, 356, 359, 361**
 ip nat outside source list **358, 360**
 ip nat outside source static **350–352, 359**
 ip nat pool **343, 357–359, 361**
 ip nat translation mas-entries **357**
 ip nat translation sampling-timeout **339, 341**
 ip nat translation timeout **357**
 ip ospf authentication **111**
 ip ospf authentication-key **111**
 ip ospf bfd **149, 156, 158**
 ip ospf bfd disable **155**
 ip pim bfd **154**
 ip pim bfd-instance **154**
 ip pim pre-build-spt **236**
 ip pim spt-threshold infinity **235**
 ip pim use-shared-tree-only **235**
 ip route **98, 155**
 ip route static bfd **155**
 ip router isis **113**
 ip router ospf **111**
 ip unnumbered **5, 109–110, 112**
 ipv6 address **101–102, 104–105, 108, 118, 302**

ipv6 address dhcp **97–98**
 ipv6 address use-link-local-only **98**
 ipv6 nd mac-extract **118**
 ipv6 nd synchronize **234**
 isis bfd **151**
 isis bfd disable **156**

L

l2protocol tunnel **319**
 l2protocol tunnel cos **319**
 l2protocol tunnel drop-threshold **320**
 l2protocol tunnel shutdown-threshold **320–321**
 lacp graceful-convergence **175, 203**
 lacp max-bundle **196**
 lacp min-links **195**
 lacp mode delay **206**
 lacp port-priority **199**
 lacp rate **197**
 lacp rate fast **198**
 lacp suspend-individual **203, 205**
 lacp system-priority **198**
 link debounce link-up **49**
 link debounce time **49**
 load- interval **91, 122, 212–213**
 load-interval counters **57**
 loopback **110, 113**

M

mac-address **107**
 mac-address ipv6-extract **118**
 match-in-vrf **344**
 medium **101**
 medium broadcast **101**
 medium p2p **101, 109–110, 112**
 mgmt0 **33**
 mtu **39, 41, 295, 300, 303**

N

negotiate auto **6, 28**
 neighbor **147**
 net **112**

P

p2p **101**
 peer-gateway **223, 260**
 peer-gateway exclude-vlan **223**
 peer-keepalive destination **257**
 peer-switch **277–278**
 permit **355**
 permit ip any any **348**
 port-channel load-balance **168, 192**

R

regex 29
 reload restore 271
 role priority 268
 router bgp 147
 router eigrp 148
 router isis 112, 150
 router ospf 149

S

sampling-timeout 341
 show 101
 show bfd 160
 show bfd neighbors 160
 show cdp all 56
 show cfs application 238
 show dot1q-tunnel 316, 322
 show feature 133, 212, 254–255, 281, 294
 show hardware feature-capability 216
 show hsrp detail 151
 show interface 33, 44, 56–58, 73–74, 76, 82, 107, 183, 185
 show interface brief 56, 90
 show interface capabilities 90
 show interface counters 91, 213
 show interface counters detailed 91, 213
 show interface counters errors 91, 213
 show interface eth 34, 103
 show interface ethernet 35, 42–43, 90, 107, 121–122
 show interface ethernet errors 122
 show interface fec 14
 show interface loopback 109, 121–122
 show interface port-channel 107, 121–122, 187, 189–191, 212
 show interface status err-disabled 20, 36–38, 56
 show interface switchport 90
 show interface transceivers 28
 show interface trunk 90
 show interface tunnel 7, 302, 306
 show interface vlan 106–107, 121–122
 show interfaces 103
 show interfaces tunnel 296, 299, 301
 show ip copy 29
 show ip eigrp 148
 show ip interface brief 121
 show ip load-sharing 211
 show ip nat max 362
 show ip nat statistics 362
 show ip nat translations 354, 362
 show ip ospf 149
 show ip route 122
 show ip route static 155
 show ipv6 icmp interface 118
 show isis 151
 show l2protocol tunnel 323
 show lacp 212
 show lacp counters 213
 show lacp system-identifier 199
 show mac address-table 237
 show port-channel capacity 282
 show port-channel compatibility-parameters 166, 212
 show port-channel database 212
 show port-channel load-balance 192, 212
 show port-channel summary 181, 194, 212
 show port-channel traffic 212
 show port-channel usage 212
 show run nat 362
 show running config 101
 show running-config 85, 91
 show running-config all 40
 show running-config bfd 135–136, 138–139, 159
 show running-config bgp 147
 show running-config hsrp 152
 show running-config interface ethernet 91
 show running-config interface port-channel 81, 83, 91, 197
 show running-config interface vlan 86, 91
 show running-config l2pt 323
 show running-config pim 154
 show running-config vpc 271, 273, 282
 show running-config vrrp 153
 show spanning-tree 233
 show spanning-tree summary 277, 279
 show startup-config bfd 160
 show startup-config interface vlan 86
 show udd 47, 56
 show udd global 56
 show vlan 77, 79
 show vpc brief 227, 233, 256, 259–260, 263, 265, 270, 282
 show vpc consistency-parameters 226–227, 264, 282
 show vpc consistency-parameters global 264
 show vpc consistency-parameters interface port-channel 264, 272, 274
 show vpc orphan-ports 275
 show vpc peer-keepalive 282
 show vpc role 266–267, 269, 282
 show vpc statistics 258, 282
 show vrf 117, 305
 show vrrp detail 152
 shut 5
 shut vPC domain 5
 shutdown 20, 36, 44, 186, 188, 201–202, 204–205, 220
 spanning-tree pseudo-information 278
 spanning-tree vlan 277
 speed 191
 speed 10 191
 speed 100 191
 speed 1000 191
 speed auto 28, 191
 speed-group 58
 speed-group 10000 27
 static 338
 switchport 28, 62, 83, 101–102, 182, 316, 318, 320
 switchport access vlan 73

- switchport host [74](#)
- switchport isolated [81](#)
- switchport mode [68, 73, 76](#)
- switchport mode dot1q-tunnel [316, 318, 320](#)
- switchport mode trunk [180, 183, 259](#)
- switchport trunk [183](#)
- switchport trunk allowed vlan [76, 78, 183, 259](#)
- switchport trunk native [183](#)
- switchport trunk native vlan [77](#)
- system default interface-vlan autostate [84–85](#)
- system default switchport [62, 89](#)
- system default switchport shutdown [89](#)
- system jumbomtu [40](#)
- system-mac [266](#)
- system-priority [267](#)

T

- terminal dont-ask [78](#)
- track [270](#)
- tunnel destination [296–297](#)
- tunnel mode [295, 299, 302](#)
- tunnel mode gre ip [295, 301–302](#)
- tunnel mode ipip [295, 297, 299, 302](#)
- tunnel mode ipv6ipv6 decapsulate-any [299](#)
- tunnel path-mtu discovery [304](#)

- tunnel path-mtu discovery age-timer [304](#)
- tunnel path-mtu discovery min-mtu [304](#)
- tunnel source [296–297](#)
- tunnel ttl [295](#)
- tunnel use-vrf [296, 302](#)

U

- udld [47](#)
- udld aggressive [46](#)
- udld message-time [46](#)
- update-source [147](#)

V

- vlan [278](#)
- vlan dot1q tag native [312](#)
- vpc [265](#)
- vpc domain [255, 257, 260, 263, 266–268, 270–271, 273, 277–278](#)
- vPC domain [5](#)
- vpc orphan-ports suspend [243, 275](#)
- vpc peer-link [259](#)
- vrf context [155](#)
- vrf member [116, 305](#)
- vrrp [153](#)
- vrrp bfd [153](#)