



Cisco Nexus 3000 Series NX-OS Interfaces Configuration Guide, Release 9.2x

First Published: 2018-07-06

Last Modified: 2019-02-17

Americas Headquarters

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

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

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

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

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

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

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

Cisco and the Cisco logo are trademarks or registered trademarks of Cisco and/or its affiliates in the U.S. and other countries. To view a list of Cisco trademarks, go to this URL: <http://www.cisco.com/go/trademarks>. Third-party trademarks mentioned are the property of their respective owners. The use of the word partner does not imply a partnership relationship between Cisco and any other company. (1110R)

© 2019 Cisco Systems, Inc. All rights reserved.



CONTENTS

PREFACE

Preface	xi
Audience	xi
Document Conventions	xi
Related Documentation for Cisco Nexus 3000 Series Switches	xii
Documentation Feedback	xii
Communications, Services, and Additional Information	xii

CHAPTER 1

New and Changed Information	1
New and Changed Information	1

CHAPTER 2

Configuring Layer 2 Interfaces	3
Licensing Requirements	3
Information About Ethernet Interfaces	3
Interface Command	3
Unidirectional Link Detection Parameter	4
Default UDLD Configuration	4
UDLD Aggressive and Nonaggressive Modes	5
Interface Speed	5
40-Gigabit Ethernet Interface Speed	6
Port Modes	7
SVI Autostate	10
Cisco Discovery Protocol	11
Default CDP Configuration	11
Error-Disabled State	12
Default Interfaces	12
Debounce Timer Parameters	12

MTU Configuration	13
Counter Values	13
Downlink Delay	14
Default Physical Ethernet Settings	14
Configuring Ethernet Interfaces	15
Guidelines for Configuring Ethernet Interfaces	15
Configuring the UDLD Mode	15
Triggering the Link State Consistency Checker	16
Changing an Interface Port Mode	17
Configuring the Interface Speed	19
Configuring Break-Out 10-Gigabit Interface Speed Ports	20
Configuring Break-In 40-Gigabit Ethernet Interface Speed Ports	21
Switching Between QSFP and SFP+ Ports	22
Disabling Link Negotiation	23
Disabling SVI Autostate	25
Configuring a Default Interface	26
Configuring the CDP Characteristics	27
Enabling or Disabling CDP	28
Enabling the Error-Disabled Detection	29
Enabling the Error-Disabled Recovery	30
Configuring the Error-Disabled Recovery Interval	30
Disabling the Error-Disabled Recovery	31
Configuring the Debounce Timer	32
Configuring the Description Parameter	33
Disabling and Restarting Ethernet Interfaces	34
Configuring Downlink Delay	34
Displaying Interface Information	35
MIBs for Layer 2 Interfaces	37

CHAPTER 3**Configuring Layer 3 Interfaces 39**

Information About Layer 3 Interfaces	39
Routed Interfaces	39
Subinterfaces	40
VLAN Interfaces	41

Changing VRF Membership for an Interface	41
Notes About Changing VRF Membership for an Interface	42
Loopback Interfaces	42
Tunnel Interfaces	42
Guidelines and Limitations for Layer 3 Interfaces	43
Default Settings for Layer 3 Interfaces	43
SVI Autostate Disable	43
DHCP Client Discovery	43
Limitations for Using DHCP Client Discovery on Interfaces	44
MAC-Embedded IPv6 Address	44
Configuring Layer 3 Interfaces	45
Configuring a Routed Interface	45
Configuring a Subinterface	46
Configuring the Bandwidth on an Interface	47
Configuring a VLAN Interface	48
Enabling Layer 3 Retention During VRF Membership Change	49
Configuring a Loopback Interface	49
Assigning an Interface to a VRF	50
Configuring an Interface MAC Address	51
Configuring a MAC-Embedded IPv6 Address	52
Configuring SVI Autostate Disable	54
Configuring a DHCP Client on an Interface	55
Verifying the Layer 3 Interfaces Configuration	56
Triggering the Layer 3 Interface Consistency Checker	57
Monitoring Layer 3 Interfaces	58
Configuration Examples for Layer 3 Interfaces	59
Example of Changing VRF Membership for an Interface	60
Related Documents for Layer 3 Interfaces	62
MIBs for Layer 3 Interfaces	62
Standards for Layer 3 Interfaces	62
Feature History for Layer 3 Interfaces	62

CHAPTER 4
Configuring Port Channels 63

Information About Port Channels	63
---------------------------------	----

Understanding Port Channels	64
Compatibility Requirements	64
Load Balancing Using Port Channels	66
Resilient Hashing	68
Hashing for NVGRE Traffic	68
Symmetric Hashing	68
Understanding LACP	69
LACP Overview	69
LACP ID Parameters	70
Channel Modes	70
LACP Marker Responders	71
LACP-Enabled and Static Port Channel Differences	71
LACP Port Channel Minimum Links and MaxBundle	72
Configuring Port Channels	72
Creating a Port Channel	72
Adding a Port to a Port Channel	73
Configuring Load Balancing Using Port Channels	74
Enabling LACP	76
Configuring the Channel Mode for a Port	76
Configuring LACP Port Channel MinLinks	77
Configuring the LACP Port-Channel MaxBundle	78
Configuring the LACP Fast Timer Rate	80
Configuring the LACP System Priority and System ID	81
Configuring the LACP Port Priority	81
Verifying Port Channel Configuration	82
Triggering the Port Channel Membership Consistency Checker	83
Verifying the Load-Balancing Outgoing Port ID	84
Feature History for Port Channels	84
Port Profiles	85
Configuring Port Profiles	86
Creating a Port Profile	86
Entering Port-Profile Configuration Mode and Modifying a Port Profile	88
Assigning a Port Profile to a Range of Interfaces	88
Enabling a Specific Port Profile	89

Inheriting a Port Profile	90
Removing a Port Profile from a Range of Interfaces	91
Removing an Inherited Port Profile	92

CHAPTER 5
Configuring IP Tunnels 95

Information About IP Tunnels	95
GRE Tunnels	96
Point-to-Point IP-in-IP Tunnel Encapsulation and Decapsulation	96
Multi-Point IP-in-IP Tunnel Decapsulation	96
Prerequisites for IP Tunnels	96
Guidelines and Limitations for IP Tunnels	96
Default Settings for IP Tunneling	100
Configuring IP Tunnels	101
Enabling Tunneling	101
Creating a Tunnel Interface	101
Configuring a Tunnel Interface	103
Configuring a Tunnel Interface Based on Policy Based Routing	105
Configuring a GRE Tunnel	106
Assigning VRF Membership to a Tunnel Interface	109
Verifying the IP Tunnel Configuration	110
Configuration Examples for IP Tunneling	110
Related Documents for IP Tunnels	111
Standards for IP Tunnels	111
Feature History for Configuring IP Tunnels	111

CHAPTER 6
Configuring Virtual Port Channels 113

Information About vPCs	113
vPC Overview	113
Terminology	114
vPC Terminology	114
vPC Domain	114
Peer-Keepalive Link and Messages	115
Compatibility Parameters for vPC Peer Links	116
Configuration Parameters That Must Be Identical	116

Configuration Parameters That Should Be Identical	117
Per-VLAN Consistency Check	118
vPC Auto-Recovery	118
vPC Peer Links	118
vPC Peer Link Overview	118
vPC Number	120
vPC Interactions with Other Features	120
vPC and LACP	120
vPC Peer Links and STP	120
CFSoE	121
Guidelines and Limitations for vPCs	121
Verifying the vPC Configuration	122
Viewing the Graceful Type-1 Check Status	123
Viewing a Global Type-1 Inconsistency	123
Viewing an Interface-Specific Type-1 Inconsistency	125
Viewing a Per-VLAN Consistency Status	126
vPC Default Settings	128
Configuring vPCs	129
Enabling vPCs	129
Disabling vPCs	129
Creating a vPC Domain	130
Configuring a vPC Keepalive Link and Messages	131
Creating a vPC Peer Link	133
Checking the Configuration Compatibility	134
Enabling vPC Auto-Recovery	135
Configuring the Restore Time Delay	136
Configuring Delay Restore on an Orphan Port	137
Configuring the Suspension of Orphan Ports	138
Excluding VLAN Interfaces from Shutting Down a vPC Peer Link Fails	139
Configuring the VRF Name	140
Moving Other Port Channels into a vPC	141
Manually Configuring a vPC Domain MAC Address	142
Manually Configuring the System Priority	143
Manually Configuring a vPC Peer Switch Role	144

Configuring Layer 3 over vPC 145

CHAPTER 7

Configuring Static and Dynamic NAT Translation 149

Network Address Translation Overview 149

Information About Static NAT 150

Dynamic NAT Overview 151

Timeout Mechanisms 151

NAT Inside and Outside Addresses 153

Pool Support for Dynamic NAT 153

Static and Dynamic Twice NAT Overview 154

Guidelines and Limitations for Static NAT 154

Restrictions for Dynamic NAT 155

Guidelines and Limitations for Dynamic Twice NAT 156

Configuring Static NAT 156

Enabling Static NAT 156

Configuring Static NAT on an Interface 157

Enabling Static NAT for an Inside Source Address 157

Enabling Static NAT for an Outside Source Address 158

Configuring Static PAT for an Inside Source Address 159

Configuring Static PAT for an Outside Source Address 160

Configuring Static Twice NAT 160

Configuration Example for Static NAT and PAT 162

Example: Configuring Static Twice NAT 163

Verifying the Static NAT Configuration 163

Configuring Dynamic NAT 164

Configuring Dynamic Translation and Translation Timeouts 164

Configuring Dynamic NAT Pool 167

Configuring Source Lists 168

Configuring Dynamic Twice NAT for an Inside Source Address 169

Configuring Dynamic Twice NAT for an Outside Source Address 170

Clearing Dynamic NAT Translations 172

Verifying Dynamic NAT Configuration 172

Example: Configuring Dynamic Translation and Translation Timeouts 173

Information About VRF Aware NAT 174

Configuring VRF Aware NAT 174

CHAPTER 8

Information About Q-in-Q Tunnels 177

Native VLAN Hazard 179

Information About Layer 2 Protocol Tunneling 180

Guidelines and Limitations for Q-in-Q Tunneling 182

Configuring Q-in-Q Tunnels and Layer 2 Protocol Tunneling 183

 Creating a 802.1Q Tunnel Port 183

 Enabling the Layer 2 Protocol Tunnel 184

 Configuring Thresholds for Layer 2 Protocol Tunnel Ports 185

 Configuring VLAN Mapping for Selective Q-in-Q on a 802.1Q Tunnel Port 186

Verifying the Q-in-Q Configuration 187

Configuration Example for Q-in-Q and Layer 2 Protocol Tunneling 188

Feature History for Q-in-Q Tunnels and Layer 2 Protocol Tunneling 188



Preface

This preface includes the following sections:

- [Audience, on page xi](#)
- [Document Conventions, on page xi](#)
- [Related Documentation for Cisco Nexus 3000 Series Switches, on page xii](#)
- [Documentation Feedback, on page xii](#)
- [Communications, Services, and Additional Information, on page xii](#)

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 the user supplies the values.
[x]	Square brackets enclose an optional element (keyword or argument).
[x y]	Square brackets enclosing keywords or arguments separated by a vertical bar indicate an optional choice.
{x y}	Braces enclosing keywords or arguments separated by a vertical bar indicate a required choice.
[x {y z}]	Nested set of square brackets or braces indicate optional or required choices within optional or required elements. Braces and a vertical bar within square brackets indicate a required choice within an optional element.

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 will include 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 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 3000 Series Switches

The entire Cisco Nexus 3000 Series switch documentation set is available at the following URL:

<https://www.cisco.com/c/en/us/support/switches/nexus-3000-series-switches/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 nexus3k-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

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

New and Changed Information

The following table provides an overview of the significant changes made to this configuration guide. The table does not provide an exhaustive list of all the changes made to this guide or all new features in a particular release.

Feature	Description	Added or Changed in Release	Where Documented
No updates since Cisco NX-OS Release 7x	First 9x Release	N/A	N/A



CHAPTER 2

Configuring Layer 2 Interfaces

- [Licensing Requirements, on page 3](#)
- [Information About Ethernet Interfaces, on page 3](#)
- [Default Physical Ethernet Settings , on page 14](#)
- [Configuring Ethernet Interfaces, on page 15](#)
- [Displaying Interface Information, on page 35](#)
- [MIBs for Layer 2 Interfaces, on page 37](#)

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](#).

Information About Ethernet Interfaces

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

The Ethernet interfaces are enabled by default.

Interface Command

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

- Interface type—All physical Ethernet interfaces use the **ethernet** keyword.
- Slot number:
 - Slot 1 includes all the fixed ports.
 - Slot 2 includes the ports on the upper expansion module (if populated).
 - Slot 3 includes the ports on the lower expansion module (if populated).
 - Slot 4 includes the ports on the lower expansion module (if populated).
- Port number— Port number within the group.

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

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

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

Unidirectional Link Detection Parameter

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


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

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

A Cisco Nexus 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.

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

Figure 1: Unidirectional Link



Default UDLD Configuration

The following table shows the default UDLD configuration.

Table 1: UDLD Default Configuration

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

UDLD Aggressive and Nonaggressive Modes

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

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

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

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

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

Interface Speed

Cisco Nexus 3000 Series switches have a number of fixed 10-Gigabit ports; each is equipped with SFP+ interface adapters. Cisco Nexus 3100 Series switches have 32 Quad Same Factor Pluggable (QSFP) ports and 4 SFP+ interface adapters. The default speed for these 32 ports is 40 Gbps.



Note If you set a port configuration that does not use all of the ports, the unused ports are left in the removed state. For example, if you configure 96 x 25G + 32 x 100G on a Cisco Nexus 3264C-E platform switch, the configuration uses 56 ports and leaves 8 ports in the removed state.

```
switch(config)# hardware profile portmode ?
 128x25g          128x25G port mode
 64x100g         64 100G ports with 2x50G, 1x100G, 1x40G capability
 96x25g+32x100g 96x25G+32x100G port mode
```

Where:

- 128x25g: Only 32 QSFP ports are usable
- 64x100g: All 64 ports are usable (default port mode)
- 96x25g+32x100g: Only 56 ports are usable

40-Gigabit Ethernet Interface Speed

You can operate QSFP ports as either 40-Gigabit Ethernet or 4 x 10-Gigabit Ethernet modes on Cisco Nexus 3132 and Cisco Nexus 3172 switches. By default, there are 32 ports in the 40-Gigabit Ethernet mode. These 40-Gigabit Ethernet ports are numbered in a 2-tuple naming convention. For example, the second 40-Gigabit Ethernet port is numbered as 1/2. The process of changing the configuration from 40-Gigabit Ethernet to 10-Gigabit Ethernet is called breakout and the process of changing the configuration from 10-Gigabit Ethernet to Gigabit Ethernet is called breakin. When you break out a 40-Gigabit Ethernet port into 10-Gigabit Ethernet ports, the resulting ports are numbered using a 3-tuple naming convention. For example, the break-out ports of the second 40-Gigabit Ethernet port are numbered as 1/2/1, 1/2/2, 1/2/3, 1/2/4.



Note The breakout ports are in administratively enabled state after the breakout of the ports into 4x10G mode or the breakin of the ports into 40G mode. On upgrade from the earlier releases, the configuration restored takes care of restoring the appropriate administrative state of the ports.

You can break out the 40-Gigabit Ethernet port into four 10-Gigabit Ethernet ports by using the **speed 10000** command and using a splitter cable to connect to multiple peer switches. You can break in four 10-Gigabit Ethernet ports to a 40-Gigabit Ethernet port by using the **speed 40000** command. The configuration change from 40-Gigabit Ethernet to 10-Gigabit Ethernet and from 10-Gigabit Ethernet to 40-Gigabit Ethernet takes effect immediately. You do not need to reload the switch. A QSFP transceiver security check is also performed.



Note When you break out from 40-Gigabit Ethernet to 10-Gigabit Ethernet, or break in from 10-Gigabit Ethernet to 40-Gigabit Ethernet, all interface configurations are reset, and the affected ports are administratively unavailable. To make these ports available, use the **no shut** command.



Note Starting with Release 6.0(2)U5(1), a new QSFP+ 40-Gb transceiver is now supported on the Cisco Nexus 3000 Series switches. The new QSFP+ (40-Gb) transceiver has a cable that splits into four 10Gb SFP-10G-LR transceivers. To use it, you need the port to be in 4x10G mode. If you are using the breakout cable, you need to run that 40G port in 4x10G mode.

The Cisco Nexus 3000 Series switches will have *auto* as default speed on all port types. Based on the transceiver type, the client auto-configures the default speeds 40000 and 10000 on QSFP and SFP+ ports respectively. The `show running-config` command treats *speed auto* as default port speed. However, you can explicitly configure different speed on these ports.

The ability to break out a 40-Gigabit Ethernet port into four 10-Gigabit Ethernet ports and break in four 10-Gigabit Ethernet ports into a 40-Gigabit Ethernet port dynamically allows you to use any of the breakout-capable ports to work in the 40-Gigabit Ethernet or 10-Gigabit Ethernet modes without permanently defining them.

For Cisco Nexus 3132Q switches, when the Ethernet interface 1/1 is in the 40-Gigabit Ethernet mode, the first QSFP port is active. After breakout, when the Ethernet interface 1/1/1-4 is in the 10-Gigabit Ethernet mode, you can choose to use either QSFP ports or SFP+ ports. However, both the first QSFP port and the four SFP+ ports cannot be active at the same time.

When using a QSFP-40G-CR4 on Cisco Nexus 3000 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.

Port Modes

Cisco Nexus 3100 Series switches have various port modes. Breakout port modes are supported on Cisco Nexus 3132Q, 3132Q-V, 31108PC-V, and 31108TC-V switches.



Note The default port mode on Cisco Nexus 3132Q and Cisco Nexus 3132CR Series switches after write erase is 32x40G mode. Breakout port modes are supported on Cisco Nexus 3132Q-V, 31108PC-V, and 31108TC-V switches.



Note One QSFP to SFP adapter fits in two QSFP ports (up and down) and provides 8 SFP+ interfaces. If you insert the QSFP to SFP adapter into the first two ports of N3K-C3132Q-40GX or N3K-C3132Q-V platforms and if you use the **hardware profile front portmode sfp-plus** command, it makes the first QSFP port inactive and an access to the adapter is disabled. Therefore, if the QSFP to SFP adapter is present on ports 1 and 2, do not use the **hardware profile front portmode sfp-plus** command.

Nexus 3100 Series Switches	Ports	Port Modes
Cisco Nexus 3132Q	32 x QSFP ports and 4 SFP+ ports	<p>The following port modes support breakout:</p> <ul style="list-style-type: none"> • 32x40G—This is an oversubscribed port mode. All 32 ports are oversubscribed and the first 24 QSFP ports are break-out capable. You cannot enter the speed 10000 command on ports 25 through 32. 32x40G breakout mode is the default port mode. • 26x40G—This is an oversubscribed port mode. Of the 26 ports, 12 ports are nonoversubscribed (cut-through). These ports are 2,4 to 8,14,and 16 to 20. The remaining 14 ports are oversubscribed. All available QSFP ports are break-out capable. • 24x40G—This is the only nonoversubscribed (cut-through) mode. All available QSFP ports are break-out capable. <p>The Fixed32x40G port mode does not support breakout.</p>
Cisco Nexus 3132Q-V	32 x 40G QSFP ports	<ul style="list-style-type: none"> • 32x40G—This is the default port mode. Of the 32 ports, first 24 QSFP ports are 10Gx4 break-out capable and the last 8 QSFP ports has a fixed speed of 40G. The maximum port counts are: 96x10G + 8x40G. All ports are oversubscribed equally. The 10G ports does not support cut-through switching. • 26x40G—This port mode supports the maximum number of 10G ports. The first 26 QSFP ports are 10Gx4 break-out capable and the last 6 QSFP ports are not usable. The maximum port counts are: 104x10G. All available QSFP ports are break-out capable. Of the 26 ports, 12 ports are non-oversubscribed (cut-through). These ports are 2,4 to 8,14,and 16 to 20. The remaining 14 ports are oversubscribed. • 24x40G—This is a non-oversubscribed, line rate port mode. Of the 32 ports, first 24 QSFP ports are 10Gx4 break-out capable. The maximum port counts can be 96x10G. The 10G ports support cut-through switching.

Nexus 3100 Series Switches	Ports	Port Modes
Cisco Nexus 31108PC-V	48 x 10G SFP+ ports and 6 x 100G QSFP ports	<p>The following two port modes are supported:</p> <ul style="list-style-type: none"> • 48x10G SFP+ ports + 6x100G QSFP ports. • 48x10G SFP+ ports + 4x100G QSFP ports + 2x40G QSFP ports. <p>The following features are specific to port mode-1:</p> <ul style="list-style-type: none"> • SFP+ Ports 1 through 8 always delivers at line-rate. • SFP+ Ports 9 through 48 are always over subscribed. • QSFP Ports 49 through 52 are capable of 100G and/or 40G and 10Gx4 breakout. • QSFP Ports 53 through 54 are capable of 100G and/or 40G. These ports do not support 10Gx4 breakout. • The maximum supported port count is 64x10G + 2x40G/100G. <p>The following features are specific to port mode-2:</p> <ul style="list-style-type: none"> • SFP+ Ports 1 through 48 always delivers at line-rate and support cut-through switching. • Line-rate SFP+ ports support cut-through switching and because of that it will have less latency. • QSFP Ports 49 through 52 are capable of 100G/40G and 10Gx4 breakout. • QSFP Ports 53 through 54 are capable of 40G and 10Gx4 breakout. These ports do not support 100G. • The maximum supported port count is 72x10G.

Nexus 3100 Series Switches	Ports	Port Modes
Cisco Nexus 31108TC-V	48 x 10GBase-T and 6 x 100G ports	<p>The following two port modes are supported:</p> <ul style="list-style-type: none"> • 48x10GBASE-T ports + 6x100G QSFP ports. • 48x10GBASE-T ports + 4x100G QSFP ports + 2x40G QSFP ports. <p>The following features are specific to port mode-1:</p> <ul style="list-style-type: none"> • 10GBASE-T Ports 1 through 8 always delivers at line-rate. • 10GBASE-T Ports 9 through 48 are always over subscribed. • QSFP Ports 49 through 52 are capable of 100G and/or 40G and 10Gx4 breakout. • QSFP Ports 53 through 54 are capable of 100G and/or 40G. These ports do not support 10Gx4 breakout. • The maximum supported port count is 64x10G + 2x40G/100G. <p>The following features are specific to port mode-2:</p> <ul style="list-style-type: none"> • 10GBASE-T Ports 1 through 48 always delivers at line-rate and support cut-through switching. • Line-rate SFP+ ports support cut-through switching and because of that it will have less latency. • QSFP Ports 49 through 52 are capable of 100G/40G and 10Gx4 breakout. • QSFP Ports 53 through 54 are capable of 40G and 10Gx4 breakout. These ports do not support 100G. • The maximum supported port count is 72x10G.
Cisco Nexus 3172PQ	6 x QSFP ports and 48 SFP+ ports	<p>The following is the default port mode and supports breakout:</p> <ul style="list-style-type: none"> • 48x10G+breakout6x40G <p>The following are the fixed port modes that do not support breakout:</p> <ul style="list-style-type: none"> • 48x10G+6x40G • 72x10G

SVI Autostate

The Switch Virtual Interface (SVI) represents a logical interface between the bridging function and the routing function of a VLAN in the device. By default, 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.

Autostate behavior is the operational state of an interface that is governed by the state of the various ports in its corresponding VLAN. An SVI interface on a VLAN comes up when there is at least one port in that vlan that is in STP forwarding state. Similarly, this interface goes down when the last STP forwarding port goes down or goes to another STP state.

By default, Autostate calculation is enabled. You can disable Autostate calculation for an SVI interface and change the default value.



Note Nexus 3000 Series switches do not support bridging between two VLANs when an SVI for one VLAN exists on the same device as the bridging link. Traffic coming into the device and bound for the SVI is dropped as a IPv4 discard. This is because the BIA MAC address is shared across VLANs/SVIs with no option to modify the MAC of the SVI.

Cisco Discovery Protocol

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

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

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

The switch supports both CDP Version 1 and Version 2.

Default CDP Configuration

The following table shows the default CDP configuration.

Table 2: Default CDP Configuration

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

Error-Disabled State

An interface is in the error-disabled (err-disabled) state when the interface is enabled administratively (using the **no shutdown** command) but disabled at runtime by any process. For example, if UDLD detects a unidirectional link, the interface is shut down at runtime. However, because the interface is administratively enabled, the interface status displays as err-disabled. Once an interface goes into the err-disabled state, you must manually reenble it or you can configure an automatic timeout recovery value. The err-disabled detection is enabled by default for all causes. The automatic recovery is not configured by default.

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 err-disabled recovery timeout for a particular err-disabled cause by changing the time variable.

The **errdisable recovery cause** command provides automatic recovery after 300 seconds. To change the recovery period, use the **errdisable recovery interval** command to specify the timeout period. You can specify 30 to 65535 seconds.

To disable recovery of an interface from the err-disabled state, use the **no errdisable recovery cause** command.

The various options for the **errdisable recover cause** command are as follows:

- **all**—Enables a timer to recover from all causes.
- **bpduguard**—Enables a timer to recover from the bridge protocol data unit (BPDU) Guard error-disabled state.
- **failed-port-state**—Enables a timer to recover from a Spanning Tree Protocol (STP) set port state failure.
- **link-flap**—Enables a timer to recover from linkstate flapping.
- **pause-rate-limit**—Enables a timer to recover from the pause rate limit error-disabled state.
- **udld**—Enables a timer to recover from the Unidirectional Link Detection (UDLD) error-disabled state.
- **loopback**—Enables a timer to recover from the loopback error-disabled state.

If you do not enable the err-disabled recovery for the cause, the interface stays in the err-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 err-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.

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, management, VLAN, and the port-channel interface.

Debounce Timer Parameters

The debounce timer delays notification of a link change, which can decrease traffic loss due to network reconfiguration. You can configure the debounce timer separately for each Ethernet port and specify the delay time in milliseconds. The delay time can range from 0 milliseconds to 5000 milliseconds. By default, this parameter is set for 100 milliseconds, which results in the debounce timer not running. When this parameter is set to 0 milliseconds, the debounce timer is disabled.



Caution Enabling the debounce timer causes the link-down detections to be delayed, which results in a loss of traffic during the debounce period. This situation might affect the convergence and reconvergence of some Layer 2 and Layer 3 protocols.

MTU Configuration

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



Note When you show the interface settings, a default MTU of 1500 is displayed for physical Ethernet interfaces.

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.

Downlink Delay

You can operationally enable uplink SFP+ ports before downlink RJ-45 ports after a reload on a Cisco Nexus 3048 switch. You must delay enabling the RJ-45 ports in the hardware until the SFP+ ports are enabled.

You can configure a timer that during reload enables the downlink RJ-45 ports in hardware only after the specified timeout. This process allows the uplink SFP+ ports to be operational first. The timer is enabled in the hardware for only those ports that are admin-enable.

Downlink delay is disabled by default and must be explicitly enabled. When enabled, if the delay timer is not specified, it is set for a default delay of 20 seconds.

Default Physical Ethernet Settings

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

Parameter	Default Setting
Duplex	Auto (full-duplex)
Encapsulation	ARPA
MTU ¹	1500 bytes
Port Mode	Access
Speed	Auto (10000)

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

Configuring Ethernet Interfaces

Guidelines for Configuring Ethernet Interfaces

There is a behavior change in configuring the interface Ethernet commands on Cisco Nexus 3000 Series switches. For example, the command **sh int ethernet Eth1/1 transceiver** does not work anymore. You have to configure the command as **sh int ethernet 1/1 transceiver**.

Configuring the UDLD Mode

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

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



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

SUMMARY STEPS

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

DETAILED STEPS

	Command or Action	Purpose
Step 1	switch# configure terminal	Enters global configuration mode.
Step 2	switch(config)# feature udld	Enables UDLD for the device.
Step 3	switch(config)# no feature udld	Disables UDLD for the device.
Step 4	switch(config)# show udld global	Displays the UDLD status for the device.

	Command or Action	Purpose
Step 5	switch(config)# interface <i>type slot/port</i>	Specifies an interface to configure, and enters interface configuration mode.
Step 6	switch(config-if)# udld { enable disable aggressive }	Enables the normal UDLD mode, disables UDLD, or enables the aggressive UDLD mode.
Step 7	switch(config-if)# show udld <i>interface</i>	Displays the UDLD status for the interface.

Example

This example shows how to enable UDLD for the switch:

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

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

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

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

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

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

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

This example shows how to disable UDLD for the switch:

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

Triggering the Link State Consistency Checker

You can manually trigger the link state consistency checker to compare the hardware and software link status of an interface and display the results. To manually trigger the link state consistency checker and display the results, use the following command in any mode:

SUMMARY STEPS

1. switch# **show consistency-checker link-state module** *slot*

DETAILED STEPS

	Command or Action	Purpose
Step 1	switch# show consistency-checker link-state module <i>slot</i>	Starts a link state consistency check on the specified module and displays its results.

Example

This example shows how to trigger a Link State consistency check and display its results:

```
switch# show consistency-checker link-state module 1
Link State Checks: Link state only
Consistency Check: FAILED
No inconsistencies found for:
  Ethernet1/1
  Ethernet1/2
  Ethernet1/3
  Ethernet1/4
  Ethernet1/5
  Ethernet1/6
  Ethernet1/7
  Ethernet1/8
  Ethernet1/9
  Ethernet1/10
  Ethernet1/12
  Ethernet1/13
  Ethernet1/14
  Ethernet1/15
Inconsistencies found for following interfaces:
  Ethernet1/11
```

Changing an Interface Port Mode

You can configure a Quad small form-factor pluggable (QSFP+) port by using the **hardware profile portmode** command. To restore the defaults, use the **no** form of these commands. The Cisco Nexus 3172PQ switch has 48x10g+breakout6x40g as the default port mode.

SUMMARY STEPS

1. switch# **configure terminal**
2. switch(config)# **copy running-config bootflash:** *my-config.cfg*
3. switch(config)# **write erase**
4. switch(config)# **reload**
5. switch(config)# [**no**] **hardware profile portmode** *portmode*
6. (Optional) switch(config)# **hardware profile portmode** *portmode 2-tuple*
7. (Optional) switch(config)# **copy running-config startup-config**
8. switch(config)# **reload**

DETAILED STEPS

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

	Command or Action	Purpose
Step 2	switch(config)# copy running-config bootflash: <i>my-config.cfg</i>	Copies the running configuration to the bootflash. You can use this file to configure your device later.
Step 3	switch(config)# write erase	Removes all the interface configurations.
Step 4	switch(config)# reload	Reloads the Cisco NX-OS software.
Step 5	switch(config)# [no] hardware profile portmode <i>portmode</i>	Changes the interface port mode.
Step 6	(Optional) switch(config)# hardware profile portmode <i>portmode 2-tuple</i>	Displays the port names in 2-tuple mode instead of the default 3-tuple convention mode.
Step 7	(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.
Step 8	switch(config)# reload	Reloads the Cisco NX-OS software. Manually apply all the interface configuration. You can refer to the configuration file that you saved earlier. Note The interface numbering changes if the ports are changed from 40G mode to 4x10G mode or vice versa.

Example

This example shows how to change the port mode to 48x10g+breakout6x40g for QSFP+ ports:

```
switch# configure terminal
switch(config)# copy running-config bootflash:my-config.cfg
switch(config)# write erase
switch(config)# reload
WARNING: This command will reboot the system
Do you want to continue? (y/n) [n] y
switch(config)# hardware profile portmode 48x10g+breakout6x40g
Warning: This command will take effect only after saving the configuration and reload!
Port configurations could get lost when port mode is changed!
switch(config)# copy running-config startup-config
switch(config)# reload
WARNING: This command will reboot the system
Do you want to continue? (y/n) [n] y
```

This example shows how to change the port mode to 48x10g+4x40g for QSFP+ ports:

```
switch# configure terminal
switch(config)# copy running-config bootflash:my-config.cfg
switch(config)# write erase
switch(config)# reload
WARNING: This command will reboot the system
Do you want to continue? (y/n) [n] y
switch(config)# hardware profile portmode 48x10g+4x40g
Warning: This command will take effect only after saving the configuration and reload!
Port configurations could get lost when port mode is changed!
switch(config)# copy running-config startup-config
switch(config)# reload
```



```
WARNING: This command will reboot the system
Do you want to continue? (y/n) [n] y
```

This example shows how to change the port mode to 48x10g+4x40g for QSFP+ ports and verify the changes:

```
switch# configure terminal
switch(config)# hardware profile portmode 48x10g+4x40g
Warning: This command will take effect only after saving the configuration and r
eload! Port configurations could get lost when port mode is changed!
switch(config)# show running-config
!Command: show running-config
!Time: Thu Aug 25 07:39:37 2011
version 5.0(3)U2(1)
feature telnet
no feature ssh
feature lldp
username admin password 5 $1$0OV4Mdom$BAb5RkD22YanT4empqgSM0 role network-admin
ip domain-lookup
switchname BLR-QG-5
ip access-list my-acl
10 deny ip any 10.0.0.1/32
20 deny ip 10.1.1.1/32 any
class-map type control-plane match-any copp-arp
class-map type control-plane match-any copp-bpdu
:
:
control-plane
service-policy input copp-system-policy
hardware profile tcam region arpacl 128
hardware profile tcam region ifacl 256
hardware profile tcam region racl 256
hardware profile tcam region vacl 512
hardware profile portmode 48x10G+4x40G
snmp-server user admin network-admin auth md5 0xdd1d21ee42e93106836cdefd1a60e062
<--Output truncated-->
switch#
```

This example shows how to restore the default port mode for QSFP+ ports:

```
switch# configure terminal
switch(config)# no hardware profile portmode
Warning: This command will take effect only after saving the configuration and r
eload! Port configurations could get lost when port mode is changed!
switch(config)#
```

Configuring the Interface Speed



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

SUMMARY STEPS

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

- switch(config-if)# **speed** *speed*

DETAILED STEPS

	Command or Action	Purpose
Step 1	switch# configure terminal	Enters global configuration mode.
Step 2	switch(config)# interface <i>type slot/port</i>	Enters interface configuration mode for the specified interface. This interface must have a 1-Gigabit Ethernet SFP transceiver inserted into it.
Step 3	switch(config-if)# speed <i>speed</i>	<p>Sets the speed on the interface.</p> <p>This command can only be applied to a physical Ethernet interface. The <i>speed</i> argument can be set to one of the following:</p> <ul style="list-style-type: none"> • 10 Mbps • 100 Mbps • 1 Gbps • 10 Gbps • automatic

Example

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

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

Configuring Break-Out 10-Gigabit Interface Speed Ports

By default, all ports on Cisco Nexus 3132 switches are 40-Gigabit Ethernet. You can break out a 40-Gigabit Ethernet port to four x10-Gigabit Ethernet ports.

SUMMARY STEPS

- switch# **configure terminal**
- switch(config)# **interface** *type slot/port-range*
- switch(config-if)# **speed 10000**

DETAILED STEPS

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

	Command or Action	Purpose
Step 2	switch(config)# interface <i>type slot/port-range</i>	Enters interface configuration mode for the specified interface. Note Interface range is not supported for 40-Gigabit Ethernet interfaces. For example, Eth 1/2-5 is not supported.
Step 3	switch(config-if)# speed 10000	Sets the speed on the interface to 10-Gigabit per second. Note Configuring breakout on QSFP ports using speed 10000 adds interface breakout module <module number> port <port range> map 10g-4x in the running-config output.

Example

This example shows how to set the speed to 10-Gigabit per second on Ethernet interface 1/2:

```
switch# configure terminal
switch(config)# interface ethernet 1/49
switch(config-if)# speed 10000

switch(config-if)# sh running-config | grep port
  limit-resource port-channel minimum 0 maximum 511
interface breakout module 1 port 49 map 10g-4x ----->
  interface breakout is added on "speed" config
hardware profile portmode 48x10g+breakout6x40g
(config-if)#

(config)# int ethernet 1/49/1
(config-if)#no speed 10000 -----> on "no speed", the interface
breakout cmd is removed.

(config-if)# sh running-config | grep port
  limit-resource port-channel minimum 0 maximum 511
hardware profile portmode 48x10g+breakout6x40g
```

Configuring Break-In 40-Gigabit Ethernet Interface Speed Ports

You can break in four x 10-Gigabit Ethernet ports to a 40-Gigabit Ethernet port.

SUMMARY STEPS

1. switch# **configure terminal**
2. switch(config)# **interface** *type slot/port*
3. switch(config-if)# **speed 40000**

DETAILED STEPS

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

	Command or Action	Purpose
Step 2	switch(config)# interface <i>type slot/port</i>	Enters interface configuration mode for the specified interface. Note The Interface range is supported for 10-Gigabit Ethernet interfaces. For example, Eth 1/2/1-4 is supported.
Step 3	switch(config-if)# speed 40000	Sets the speed on the interface to 40 Gbps.

Example

This example shows how to set the speed to 40 Gbps on Ethernet interface 1/2/1:

```
switch# configure terminal
switch(config)# interface ethernet 1/2/1
switch(config-if)# speed 40000
```

Switching Between QSFP and SFP+ Ports

When you break out ports into the 10-GbE mode, you can switch between the first QSFP port and SFP+ ports 1 to 4. Either the first QSFP port or the four SFP+ ports can be active at any time. QSFP is the default port with an interface speed of 40 Gbps.

When the first QSFP port is in the 40-GbE mode, you cannot switch the port to four SFP+ ports and the first QSFP port will be active until you break out the port into the 10-GbE mode. This is because SFP+ ports do not support the 40-GbE mode.

SUMMARY STEPS

1. switch# **configure terminal**
2. switch(config)# **[no] hardware profile front portmode qsfp | sfp-plus**
3. switch(config)# **interface breakout module** *module number* **port port rangemap 10g-4x**
4. (Optional) switch(config)# **copy running-config startup-config**

DETAILED STEPS

	Command or Action	Purpose
Step 1	switch# configure terminal	Enters global configuration mode.
Step 2	switch(config)# [no] hardware profile front portmode qsfp sfp-plus	Activates the specified port mode. <ul style="list-style-type: none"> • qsfp—The front panel QSFP port is active • sfp-plus—The front panel SFP+ ports 1 to 4 are active The no form of this command activates the QSFP port. Note If the first QSFP port speed is 40 Gbps, this command will run, but the SFP+ ports will not become active until after the speed is changed to 10 Gbps.

	Command or Action	Purpose
Step 3	switch(config)# interface breakout module <i>module number</i> port port rangemap 10g-4x	Enables you to configure the module in 10g mode. When you are changing the portmode from QSFP to SFP+, the hardware profile front portmode command takes effect only after breaking out the first QSFP port as displayed in this command.
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 change the portmode from QSFP to SFP+:

```
switch# show int e1/1 transceiver
Ethernet1/1
transceiver is present
type is QSFP-40G-SR
name is CISCO
part number is AFBR-79EIPZ-CS1
revision is 02
serial number is AVP1645S1QT
nominal bitrate is 10300 MBit/sec per channel
Link length supported for 50/125um fiber is 30 m
Link length supported for 50/125um fiber is 100 m
cisco id is --
cisco extended id number is 16

switch# show running-config | inc portmode
hardware profile portmode 32X40G
hardware profile front portmode qsfp

switch# configure terminal
switch(config)# hardware profile front portmode sfp-plus
switch(config)# interface breakout module 1 port 1 map 10g-4x
switch(config)# copy running-config startup-config
```

This example shows how to make the QSFP port active:

```
switch# configure terminal
switch(config)# hardware profile front portmode qsfp
switch(config)# copy running-config startup-config
```

Disabling Link Negotiation

By default, auto-negotiation is enabled on all 1G SFP+ and 40G QSFP ports and it is disabled on 10G SFP+ ports. Auto-negotiation is by default enabled on all 1G and 10G Base-T ports. It cannot be disabled on 1G and 10G Base-T ports.

This command is equivalent to the Cisco IOS **speed non-negotiate** command.



Note The auto-negotiation configuration is not applicable on 10-Gigabit Ethernet ports. When auto-negotiation is configured on a 10-Gigabit port, the following error message is displayed:

```
ERROR: Ethernet1/40: Configuration does not match the port capability
```



Note 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:

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

SUMMARY STEPS

1. switch# **configure terminal**
2. switch(config)# **interface ethernet slot/port**
3. switch(config-if)# **no negotiate auto**
4. (Optional) switch(config-if)# **negotiate auto**

DETAILED STEPS

	Command or Action	Purpose
Step 1	switch# configure terminal	Enters global configuration mode.
Step 2	switch(config)# interface ethernet slot/port	Selects the interface and enters interface mode.
Step 3	switch(config-if)# no negotiate auto	Disables link negotiation on the selected Ethernet interface (1-Gigabit port).
Step 4	(Optional) switch(config-if)# negotiate auto	Enables link negotiation on the selected Ethernet interface. The default for 1-Gigabit Ethernet ports is enabled. Note This command is not applicable for 10GBASE-T ports. It should not be used on 10-GBASE-T ports.

Example

This example shows how to disable auto-negotiation on a specified Ethernet interface (1-Gigabit port):

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

```
switch(config-if) # no negotiate auto
switch(config-if) #
```

This example shows how to enable auto-negotiation on a specified Ethernet interface (1-Gigabit port):

```
switch# configure terminal
switch(config) # interface ethernet 1/5
switch(config-if) # negotiate auto
switch(config-if) #
```

Disabling SVI Autostate

You can configure a SVI to remain active even if no interfaces are up in the corresponding VLAN. This enhancement is called Autostate Disable.

When you enable or disable autostate behavior, it is applied to all the SVIs in the switch unless you configure autostate per SVI .



Note Autostate behavior is enabled by default.

SUMMARY STEPS

1. switch# **configure terminal**
2. switch(config)# **feature interface-vlan**
3. switch(config)# **system default interface-vlan [no] autostate**
4. (Optional) switch(config)# **interface vlan** *interface-vlan-number*
5. (Optional) switch(config-if)# **[no] autostate**
6. (Optional) switch(config)# **show interface-vlan** *interface-vlan*
7. (Optional) switch(config)# **copy running-config startup-config**

DETAILED STEPS

	Command or Action	Purpose
Step 1	switch# configure terminal	Enters global configuration mode.
Step 2	switch(config)# feature interface-vlan	Enables the interface-vlan feature.
Step 3	Required: switch(config)# system default interface-vlan [no] autostate	Configures the system to enable or disable the Autostate default behavior.
Step 4	(Optional) switch(config)# interface vlan <i>interface-vlan-number</i>	Creates a VLAN interface. The number range is from 1 to 4094.
Step 5	(Optional) switch(config-if)# [no] autostate	Enables or disables Autostate behavior per SVI.
Step 6	(Optional) switch(config)# show interface-vlan <i>interface-vlan</i>	Displays the enabled or disabled Autostate behavior of the SVI.

	Command or Action	Purpose
Step 7	(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 disable the systems Autostate default for all the SVIs on the switch:

```
switch# configure terminal
switch(config)# feature interface-vlan
switch(config)# system default interface-vlan no autostate
switch(config)# interface vlan 50
switch(config-if)# no autostate
switch(config)# copy running-config startup-config
```

This example shows how to enable the systems autostate configuration:

```
switch(config)# show interface-vlan 2
Vlan2 is down, line protocol is down, autostate enabled
Hardware is EtherSVI, address is 547f.ee40.a17c
MTU 1500 bytes, BW 1000000 Kbit, DLY 10 usec
```

Configuring a Default Interface

The default interface feature allows you to clear the existing configuration of multiple interfaces such as Ethernet, loopback, management, VLAN, and port-channel interfaces. All user configuration under a specified interface will be deleted. On a Cisco Nexus C3408-S switch, the number of interfaces you can configure using the **default interface ethernet** command, at a time, is limited to a maximum of 64 ports.

SUMMARY STEPS

1. switch# **configure terminal**
2. switch(config)# **default interface** *type interface number*
3. switch(config)# **exit**

DETAILED STEPS

	Command or Action	Purpose
Step 1	switch# configure terminal	Enters global configuration mode.
Step 2	switch(config)# default interface <i>type interface number</i>	Deletes the configuration of the interface and restores the default configuration. The following are the supported interfaces: <ul style="list-style-type: none"> • ethernet • loopback • mgmt • port-channel • vlan

	Command or Action	Purpose
Step 3	switch(config)# exit	Exits global configuration mode.

Example

This example shows how to delete the configuration of an Ethernet interface and revert it to its default configuration:

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

Configuring the CDP Characteristics

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

SUMMARY STEPS

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

DETAILED STEPS

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

	Command or Action	Purpose
Step 5	(Optional) switch(config)# [no] cdp timer <i>seconds</i>	Sets the transmission frequency of CDP updates in seconds. The range is 5 to 254; the default is 60 seconds. Use the no form of the command to return to its default setting.

Example

This example shows how to configure CDP characteristics:

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

Enabling or Disabling CDP

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

SUMMARY STEPS

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

DETAILED STEPS

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

Example

This example shows how to enable CDP for an Ethernet port:

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

```
switch(config-if)# cdp enable
```

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

Enabling the Error-Disabled Detection

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



Note Base ports in Cisco Nexus 5500 never get error disabled due to pause rate-limit like in the Cisco Nexus 5020 or 5010 switch.

SUMMARY STEPS

1. switch# **configure terminal**
2. switch(config)# **errdisable detect cause** {all / link-flap / loopback}
3. switch(config)# **shutdown**
4. switch(config)# **no shutdown**
5. switch(config)# **show interface status err-disabled**
6. (Optional) switch(config)# **copy running-config startup-config**

DETAILED STEPS

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

```

switch# configure terminal
switch(config)# errdisable detect cause all
switch(config)# shutdown
switch(config)# no shutdown
switch(config)# show interface status err-disabled
switch(config)# copy running-config startup-config

```

Enabling the Error-Disabled Recovery

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

SUMMARY STEPS

1. switch# **configure terminal**
2. switch(config)# **errdisable recovery cause** {all | udld | bpduguard | link-flap | failed-port-state | pause-rate-limit | loopback}
3. switch(config)# **show interface status err-disabled**
4. (Optional) switch(config)# **copy running-config startup-config**

DETAILED STEPS

	Command or Action	Purpose
Step 1	switch# configure terminal	Enters global configuration mode.
Step 2	switch(config)# errdisable recovery cause {all udld bpduguard link-flap failed-port-state pause-rate-limit loopback}	Specifies a condition under which the interface automatically recovers from the err-disabled state, and the device retries bringing the interface up. The device waits 300 seconds to retry. The default is disabled.
Step 3	switch(config)# show interface status err-disabled	Displays information about err-disabled interfaces.
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 enable err-disabled recovery under all conditions:

```

switch# configure terminal
switch(config)# errdisable recovery cause loopback
switch(config)# show interface status err-disabled
switch(config)# copy running-config startup-config

```

Configuring the Error-Disabled Recovery Interval

You can use this procedure to configure the err-disabled recovery timer value. The range is from 30 to 65535 seconds. The default is 300 seconds.

SUMMARY STEPS

1. switch# **configure terminal**
2. switch(config)# **errdisable recovery interval interval**
3. switch(config)# **show interface status err-disabled**
4. (Optional) switch(config)# **copy running-config startup-config**

DETAILED STEPS

	Command or Action	Purpose
Step 1	switch# configure terminal	Enters global configuration mode.
Step 2	switch(config)# errdisable recovery interval interval	Specifies the interval for the interface to recover from the err-disabled state. The range is from 30 to 65535 seconds. The default is 300 seconds.
Step 3	switch(config)# show interface status err-disabled	Displays information about err-disabled interfaces.
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 enable err-disabled recovery under all conditions:

```
switch# configure terminal
switch(config)# errdisable recovery interval 32
switch(config)# show interface status err-disabled
switch(config)# copy running-config startup-config
```

Disabling the Error-Disabled Recovery

You can disable recovery of an interface from the err-disabled state.

SUMMARY STEPS

1. switch# **configure terminal**
2. switch(config)# **no errdisable recovery cause {all | udd | bpduguard | link-flap | failed-port-state | pause-rate-limit | loopback}**
3. (Optional) switch(config)# **show interface status err-disabled**
4. (Optional) switch(config)# **copy running-config startup-config**

DETAILED STEPS

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

	Command or Action	Purpose
Step 2	switch(config)# no errdisable recovery cause {all / udd / bpduguard / link-flap / failed-port-state / pause-rate-limit / loopback}	Specifies a condition under which the interface reverts back to the default err-disabled state.
Step 3	(Optional) switch(config)# show interface status err-disabled	Displays information about err-disabled interfaces.
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 disable err-disabled recovery:

```
switch# configure terminal
switch(config)# no errdisable recovery cause loopback
switch(config)# show interface status err-disabled
switch(config)# copy running-config startup-config
```

Configuring the Debounce Timer

You can enable the debounce timer for Ethernet ports by specifying a debounce time, in milliseconds (ms), or disable the timer by specifying a debounce time of 0. By default, the debounce timer is set to 100 ms, which results in the debounce timer not running.



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.

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

SUMMARY STEPS

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

DETAILED STEPS

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

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

Example

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

```
switch# configure terminal
switch(config)# interface ethernet 3/1
switch(config-if)# link debounce time 1000
```

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

```
switch# configure terminal
switch(config)# interface ethernet 3/1
switch(config-if)# link debounce time 0
```

Configuring the Description Parameter

You can provide textual interface descriptions for the Ethernet ports.

SUMMARY STEPS

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

DETAILED STEPS

	Command or Action	Purpose
Step 1	switch# configure terminal	Enters global configuration mode.
Step 2	switch(config)# interface <i>type slot/port</i>	Enters interface configuration mode for the specified interface.
Step 3	switch(config-if)# description <i>test</i>	Specifies the description for the interface.

Example

This example shows how to set the interface description to Server 3 interface:

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

Disabling and Restarting Ethernet Interfaces

You can shut down and restart an Ethernet interface. This action disables all of the interface functions and marks the interface as being down on all monitoring displays.

SUMMARY STEPS

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

DETAILED STEPS

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

Example

This example shows how to disable an Ethernet port:

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

This example shows how to restart an Ethernet interface:

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

Configuring Downlink Delay

You can operationally enable uplink SFP+ ports before downlink RJ-45 ports after a reload on a Cisco Nexus 3048 switch by delaying enabling the RJ-45 ports in the hardware until the SFP+ ports are enabled.

SUMMARY STEPS

1. switch# **configure terminal**
2. switch(config)# **downlink delay enable | disable [timeout time-out]**

DETAILED STEPS

	Command or Action	Purpose
Step 1	switch# configure terminal	Enters global configuration mode.
Step 2	switch(config)# downlink delay enable disable [timeout time-out]	Enables or disables downlink delay and configures the timeout.

Example

This example shows how to enable downlink delay and configure the delay timeout on the switch:

```
switch# configure terminal
switch(config)# downlink delay enable timeout 45
```

Displaying Interface Information

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

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

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

This example shows how to display the physical Ethernet interface:

```
switch# show interface ethernet 1/1
Ethernet1/1 is up
Hardware is 1000/10000 Ethernet, address is 000d.eca3.5f08 (bia 000d.eca3.5f08)
MTU 1500 bytes, BW 10000000 Kbit, DLY 10 usec,
    reliability 255/255, txload 190/255, rxload 192/255
Encapsulation ARPA
Port mode is trunk
full-duplex, 10 Gb/s, media type is 1/10g
Input flow-control is off, output flow-control is off
Auto-mdix is turned on
Rate mode is dedicated
Switchport monitor is off
```

```

Last clearing of "show interface" counters never
5 minute input rate 942201806 bytes/sec, 14721892 packets/sec
5 minute output rate 935840313 bytes/sec, 14622492 packets/sec
Rx
 129141483840 input packets 0 unicast packets 129141483847 multicast packets
 0 broadcast packets 0 jumbo packets 0 storm suppression packets
 8265054965824 bytes
 0 No buffer 0 runt 0 Overrun
 0 crc 0 Ignored 0 Bad etype drop
 0 Bad proto drop
Tx
 119038487241 output packets 119038487245 multicast packets
 0 broadcast packets 0 jumbo packets
 7618463256471 bytes
 0 output CRC 0 ecc
 0 underrun 0 if down drop      0 output error 0 collision 0 deferred
 0 late collision 0 lost carrier 0 no carrier
 0 babble
 0 Rx pause 8031547972 Tx pause 0 reset

```

This example shows how to display the physical Ethernet capabilities:

```

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

  MDIX:                 no
  FEX Fabric:           yes

```

This example shows how to display the physical Ethernet transceiver:

```

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

```

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

```

switch# show interface brief

-----
Ethernet      VLAN  Type Mode  Status Reason          Speed  Port

```

Interface					Ch #	
Eth1/1	200	eth	trunk	up	none	10G(D) --
Eth1/2	1	eth	trunk	up	none	10G(D) --
Eth1/3	300	eth	access	down	SFP not inserted	10G(D) --
Eth1/4	300	eth	access	down	SFP not inserted	10G(D) --
Eth1/5	300	eth	access	down	Link not connected	1000(D) --
Eth1/6	20	eth	access	down	Link not connected	10G(D) --
Eth1/7	300	eth	access	down	SFP not inserted	10G(D) --
...						

This example shows how to display the CDP neighbors:

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

Device ID          Local Intrfce  Hldtme  Capability  Platform  Port ID
dl3-dist-1        mgmt0          148     S I         WS-C2960-24TC  Fas0/9
n5k (FLC12080012) Eth1/5         8       S I s       N5K-C5020P-BA  Eth1/5
```

MIBs for Layer 2 Interfaces

MIB	MIB Link
IF-MIB	To locate and download MIBs, go to the following URL: http://www.cisco.com/public/sw-center/netmgmt/cmtk/mibs.shtml
MAU-MIB Limited support includes only the following MIB Objects: <ul style="list-style-type: none"> • ifMauType (Read-only) GET • ifMauAutoNegSupported (Read-only) GET • ifMauTypeListBits (Read-only) GET • ifMauDefaultType (Read-write) GET-SET • ifMauAutoNegAdminStatus (Read-write) GET-SET • ifMauAutoNegCapabilityBits (Read-only) GET • ifMauAutoNegAdvertisedBits (Read-write) GET-SET 	



CHAPTER 3

Configuring Layer 3 Interfaces

- [Information About Layer 3 Interfaces, on page 39](#)
- [Guidelines and Limitations for Layer 3 Interfaces, on page 43](#)
- [Default Settings for Layer 3 Interfaces, on page 43](#)
- [SVI Autostate Disable, on page 43](#)
- [DHCP Client Discovery, on page 43](#)
- [MAC-Embedded IPv6 Address, on page 44](#)
- [Configuring Layer 3 Interfaces, on page 45](#)
- [Verifying the Layer 3 Interfaces Configuration, on page 56](#)
- [Triggering the Layer 3 Interface Consistency Checker, on page 57](#)
- [Monitoring Layer 3 Interfaces, on page 58](#)
- [Configuration Examples for Layer 3 Interfaces, on page 59](#)
- [Example of Changing VRF Membership for an Interface, on page 60](#)
- [Related Documents for Layer 3 Interfaces, on page 62](#)
- [MIBs for Layer 3 Interfaces, on page 62](#)
- [Standards for Layer 3 Interfaces, on page 62](#)
- [Feature History for Layer 3 Interfaces, on page 62](#)

Information About Layer 3 Interfaces

Layer 3 interfaces forward 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 Layer 2 (switchports) by default. You can change this default behavior using the **no switchport** command from interface configuration mode. To change multiple ports at one time, you can specify a range of interfaces and then apply the **no switchport** command.

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

You can assign a static MAC address to a Layer 3 interface. The default MAC address for a Layer 3 interface is the MAC address of the virtual device context (VDC) that is associated with it. You can change the default MAC address of the Layer 3 interface by using the **mac-address** command from the interface configuration mode. A static MAC address can be configured on SVI, Layer 3 interfaces, port channels, Layer 3 subinterfaces, and tunnel interfaces. You can also configure static MAC addresses on a range of ports and port channels. However, all ports must be in Layer 3. Even if one port in the range of ports is in Layer 2, the command is rejected and an error message appears. For information on configuring MAC addresses, see the Layer 2 Switching Configuration Guide for your device.

You can also create a Layer 3 port channel from routed interfaces.

Routed interfaces and subinterfaces 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 or a port channel.

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 VLAN that is 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 that are carried by the trunking port.

Figure 2: Subinterfaces for VLANs



VLAN Interfaces

A VLAN interface or a 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 configure it. The system automatically takes a checkpoint prior to disabling the feature, and you can roll back to this checkpoint. For information about rollbacks and checkpoints, see the System Management Configuration Guide for your device.




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 on IP addresses and IP routing, see the Unicast Routing Configuration Guide for your device.

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 3: 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 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 that is transmitted over a loopback interface is immediately received by this interface. Loopback interfaces emulate a physical interface.

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.

Tunnel Interfaces

Cisco NX-OS supports tunnel interfaces as IP tunnels. IP tunnels can encapsulate a same- layer or higher layer protocol and transport the result over IP through a tunnel that is created between two routers.

Guidelines and Limitations for Layer 3 Interfaces

Layer 3 interfaces have the following configuration guidelines and limitations:

- The VLAN/SVI is not removed from the Layer 3 interface table, after the configuration is removed. The VLAN itself should be removed from the Layer 3 interface table.
- 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.
- 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.
- 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.
- Cisco Nexus 3000 Series switches punt multicast Layer 2 traffic to the CPU if the Layer 3 MTU is not the same for all Layer 3 interfaces, and if the MTU QoS was changed to jumbo. All Layer 3 interfaces must have the same Layer 3 MTU to avoid this issue.

Default Settings for Layer 3 Interfaces

The default setting for the Layer 3 Admin state is Shut.

SVI Autostate Disable

The SVI Autostate Disable feature enables the Switch Virtual Interface (SVI) to be in the “up” state even if no interface is in the “up” state in the corresponding VLAN.

An SVI is also a virtual routed interface that connects a VLAN on the device to the Layer 3 router engine on the same device. The ports in a VLAN determine the operational state of the corresponding SVI. An SVI interface on a VLAN comes “up” when at least one port in the corresponding VLAN is in the Spanning Tree Protocol (STP) forwarding state. Similarly, the SVI interface goes “down” when the last STP forwarding port goes down or to any other state. This characteristic of SVI is called 'Autostate'.

You can create SVIs to define Layer 2 or Layer 3 boundaries on VLANs, or use the SVI interface to manage devices. In the second scenario, the SVI Autostate Disable feature ensures that the SVI interface is in the “up” state even if no interface is in the “up” state in the corresponding VLAN.

DHCP Client Discovery

Cisco NX-OS Release 6.0(2)U3(1) introduced DHCP client discovery on SVIs. Cisco NX-OS Release 6.0(2)U4(1) adds DHCP client discovery support for IPv6 addresses and 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.

If the switch is reloaded and, at the same time, the router and DNS options are disabled on the server side, after the switch comes up, a new IP address is assigned to the SVI. However, the stale **ip route** command and **ip name-server** command will still exist in the switch configuration. You must manually remove these commands from the configuration.

Limitations for Using DHCP Client Discovery on Interfaces

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

- This feature is supported only on physical Ethernet interfaces, management interfaces, and SVIs.
- Starting with Cisco NX-OS Release 6.0(2)U4(1), this feature is supported on non-default virtual routing and forwarding (VRF) instances as well.
- 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.

MAC-Embedded IPv6 Address

Beginning with Cisco NX-OS Release 6.0(2)U4(1), 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, 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 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.

Configuring Layer 3 Interfaces

Configuring a Routed Interface

SUMMARY STEPS

1. switch# **configure terminal**
2. switch(config)# **interface ethernet slot/port**
3. switch(config-if)# **no switchport**
4. switch(config-if)# [**ip|ipv6**]ip-address/length
5. (Optional) switch(config-if)# **medium {broadcast | p2p}**
6. (Optional) switch(config-if)# **show interfaces**
7. (Optional) switch(config-if)# **copy running-config startup-config**

DETAILED STEPS

	Command or Action	Purpose
Step 1	switch# configure terminal	Enters global configuration mode.
Step 2	switch(config)# interface ethernet slot/port	Enters interface configuration mode.
Step 3	switch(config-if)# no switchport	Configures the interface as a Layer 3 interface and deletes any configuration specific to Layer 2 on this interface. Note To convert a Layer 3 interface back into a Layer 2 interface, use the switchport command.
Step 4	switch(config-if)# [ip ipv6]ip-address/length	Configures an IP address for this interface.
Step 5	(Optional) switch(config-if)# medium {broadcast p2p}	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.
Step 6	(Optional) switch(config-if)# show interfaces	Displays the Layer 3 interface statistics.
Step 7	(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 shows how to configure an IPv4-routed Layer 3 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

```

Configuring a Subinterface

Before you begin

- Configure the parent interface as a routed interface.
- Create the port-channel interface if you want to create a subinterface on that port channel.

SUMMARY STEPS

1. (Optional) switch(config-if)# **copy running-config startup-config**
2. switch(config)# **interface ethernet slot/port.number**
3. switch(config-if)# **[ip | ipv6] address ip-address/length**
4. switch(config-if)# **encapsulation dot1Q vlan-id**
5. (Optional) switch(config-if)# **show interfaces**
6. (Optional) switch(config-if)# **copy running-config startup-config**

DETAILED STEPS

	Command or Action	Purpose
Step 1	(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.
Step 2	switch(config)# interface ethernet slot/port.number	Enters interface configuration mode. The range for the <i>slot</i> is from 1 to 255. The range for the <i>port</i> is from 1 to 128.
Step 3	switch(config-if)# [ip ipv6] address ip-address/length	Configures an IP address for this interface.
Step 4	switch(config-if)# encapsulation dot1Q vlan-id	Configures IEEE 802.1Q VLAN encapsulation on the subinterface. The range for the <i>vlan-id</i> is from 2 to 4093.
Step 5	(Optional) switch(config-if)# show interfaces	Displays the Layer 3 interface statistics.
Step 6	(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 shows how to create a subinterface:

```
switch# configure terminal
switch(config)# interface ethernet 2/1
switch(config-if)# ip address 192.0.2.1/8
switch(config-if)# encapsulation dot1Q 33
switch(config-if)# copy running-config startup-config
```

Configuring the Bandwidth on an Interface

You can configure the bandwidth for a routed interface, port channel, or subinterface.

SUMMARY STEPS

1. switch# **configure terminal**
2. switch(config)# **interface ethernet** *slot/port*
3. switch(config-if)# **bandwidth** [*value* | **inherit** [*value*]]
4. (Optional) switch(config-if)# **copy running-config startup-config**

DETAILED STEPS

	Command or Action	Purpose
Step 1	switch# configure terminal	Enters global configuration mode.
Step 2	switch(config)# interface ethernet <i>slot/port</i>	Enters interface configuration mode. The range for the <i>slot</i> is from 1 to 255. The range for the <i>port</i> is from 1 to 128.
Step 3	switch(config-if)# bandwidth [<i>value</i> inherit [<i>value</i>]]	Configures the bandwidth parameter for a routed interface, port channel, or subinterface, as follows: <ul style="list-style-type: none"> • value—Size of the bandwidth in kilobytes. The range is from 1 to 10000000. • inherit—Indicates that all subinterfaces of this interface inherit either the bandwidth value (if a value is specified) or the bandwidth of the parent interface (if a value is not specified).
Step 4	(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 shows how to configure Ethernet interface 2/1 with a bandwidth value of 80000:

```
switch# configure terminal
```

```
switch(config)# interface ethernet 2/1
switch(config-if)# bandwidth 80000
switch(config-if)# copy running-config startup-config
```

Configuring a VLAN Interface

SUMMARY STEPS

1. switch# **configure terminal**
2. switch(config)# **feature interface-vlan**
3. switch(config)# **interface vlan** *number*
4. switch(config-if)# [**ip | ipv6**] **address** *ip-address/length*
5. switch(config-if)# **no shutdown**
6. (Optional) switch(config-if)# **show interface vlan** *number*
7. (Optional) switch(config-if)# **copy running-config startup-config**

DETAILED STEPS

	Command or Action	Purpose
Step 1	switch# configure terminal	Enters global configuration mode.
Step 2	switch(config)# feature interface-vlan	Enables VLAN interface mode.
Step 3	switch(config)# interface vlan <i>number</i>	Creates a VLAN interface. The <i>number</i> range is from 1 to 4094.
Step 4	switch(config-if)# [ip ipv6] address <i>ip-address/length</i>	Configures an IP address for this interface.
Step 5	switch(config-if)# no shutdown	Brings the interface up administratively.
Step 6	(Optional) switch(config-if)# show interface vlan <i>number</i>	Displays the VLAN interface statistics. The <i>number</i> range is from 1 to 4094.
Step 7	(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 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.

SUMMARY STEPS

1. **configure terminal**
2. **system vrf-member-change retain-l3-config**

DETAILED STEPS

	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> <p>Warning: Will retain L3 configuration when vrf member change on interface.</p>	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.

Configuring a Loopback Interface

Before you begin

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

SUMMARY STEPS

1. **switch# configure terminal**
2. **switch(config)# interface loopback *instance***
3. **switch(config-if)# [ip | ipv6] address *ip-address/length***
4. (Optional) **switch(config-if)# show interface loopback *instance***
5. (Optional) **switch(config-if)# copy running-config startup-config**

DETAILED STEPS

	Command or Action	Purpose
Step 1	switch# configure terminal	Enters global configuration mode.
Step 2	switch(config)# interface loopback <i>instance</i>	Creates a loopback interface. The <i>instance</i> range is from 0 to 1023.

	Command or Action	Purpose
Step 3	switch(config-if)# [ip ipv6] address <i>ip-address/length</i>	Configures an IP address for this interface.
Step 4	(Optional) switch(config-if)# show interface loopback <i>instance</i>	Displays the loopback interface statistics. The <i>instance</i> range is from 0 to 1023.
Step 5	(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 shows how to create a loopback interface:

```
switch# configure terminal
switch(config)# interface loopback 0
switch(config-if)# ip address 192.0.2.100/8
switch(config-if)# copy running-config startup-config
```

Assigning an Interface to a VRF

Before you begin

Assign the IP address for a tunnel interface after you have configured the interface for a VRF.

SUMMARY STEPS

1. switch# **configure terminal**
2. switch(config)# **interface** *interface-typenumber*
3. switch(config-if)#**vrf member** *vrf-name*
4. switch(config-if)# [ip | ipv6]*ip-address/length*
5. (Optional) switch(config-if)# **show vrf** [*vrf-name*] **interface** *interface-type number*
6. (Optional) switch(config-if)# **show interfaces**
7. (Optional) switch(config-if)# **copy running-config startup-config**

DETAILED STEPS

	Command or Action	Purpose
Step 1	switch# configure terminal	Enters global configuration mode.
Step 2	switch(config)# interface <i>interface-typenumber</i>	Enters interface configuration mode.
Step 3	switch(config-if)# vrf member <i>vrf-name</i>	Adds this interface to a VRF.
Step 4	switch(config-if)# [ip ipv6] <i>ip-address/length</i>	Configures an IP address for this interface. You must do this step after you assign this interface to a VRF.

	Command or Action	Purpose
Step 5	(Optional) switch(config-if)# show vrf [<i>vrf-name</i>] interface <i>interface-type number</i>	Displays VRF information.
Step 6	(Optional) switch(config-if)# show interfaces	Displays the Layer 3 interface statistics.
Step 7	(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 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 an Interface MAC Address

You can configure a static MAC address on SVI, Layer 3 interfaces, port channels, Layer 3 subinterfaces, and tunnel interfaces. You can also configure static MAC addresses on a range of ports and port channels. However, all ports must be in Layer 3. Even if one port in the range of ports is in Layer 2, the command is rejected and an error message appears.

SUMMARY STEPS

1. switch# **configure terminal**
2. switch(config)# **interface ethernet** *slot/port*
3. switch(config-if)# [**no**] **mac-address** *static router MAC address*
4. switch(config-if)# **show interface ethernet** *slot/port*

DETAILED STEPS

	Command or Action	Purpose
Step 1	switch# configure terminal	Enters global configuration mode.
Step 2	switch(config)# interface ethernet <i>slot/port</i>	Enters interface configuration mode.
Step 3	switch(config-if)# [no] mac-address <i>static router MAC address</i>	Configures the interface MAC address. The no form removes the configuration. You can enter the MAC address in any one of the four supported formats: <ul style="list-style-type: none"> • E.E.E • EE-EE-EE-EE-EE-EE • EE:EE:EE:EE:EE:EE • EEEE.EEEE.EEEE Do not enter any of the following invalid MAC addresses:

	Command or Action	Purpose
		<ul style="list-style-type: none"> • Null MAC address—0000.0000.0000 • Broadcast MAC address—FFFF.FFFF.FFFF • Multicast MAC address—0100.DAAA.ADDD
Step 4	switch(config-if)# show interface ethernet slot/port	(Optional) Displays all information for the interface.

Example

This example shows how to configure an interface MAC address:

```
switch# configure terminal
switch(config)# interface ethernet 3/3
switch(config-if)# mac-address aaaa.bbbb.dddd
switch(config-if)# show interface ethernet 3/3
switch(config-if)#
```

Configuring a MAC-Embedded IPv6 Address

SUMMARY STEPS

1. switch# **configure terminal**
2. switch(config)# **interface type slot/port**
3. switch(config-if)# **no switchport**
4. switch(config-if)# **mac-address ipv6-extract**
5. switch(config-if)# **ipv6 address ip-address/length**
6. switch(config-if)# **ipv6 nd mac-extract [exclude nud-phase]**
7. (Optional) switch(config)# **show ipv6 icmp interface type slot/port**

DETAILED STEPS

	Command or Action	Purpose
Step 1	switch# configure terminal	Enters global configuration mode.
Step 2	switch(config)# interface type slot/port	Enters the interface configuration mode for the specified interface.
Step 3	switch(config-if)# no switchport	Configures the interface as a Layer 3 interface and deletes any configuration specific to Layer 2 on this interface. Note To convert a Layer 3 interface back into a Layer 2 interface, use the switchport command.
Step 4	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.

	Command or Action	Purpose
Step 5	switch(config-if)# ipv6 address <i>ip-address/length</i>	Configures an IPv6 address for this interface.
Step 6	switch(config-if)# ipv6 nd mac-extract [exclude nud-phase]	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) switch(config)# show ipv6 icmp interface <i>type slot/port</i>	Displays IPv6 Internet Control Message Protocol version 6 (ICMPv6) interface information.

Example

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

```
switch# configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
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
switch(config)#
```

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

```
switch# configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
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
switch(config-if)#
```

Configuring SVI Autostate Disable

You can configure a SVI to remain active even if no interfaces are up in the corresponding VLAN. This enhancement is called Autostate Disable.

SUMMARY STEPS

1. switch# **configure terminal**
2. switch(config)# **system default interface-vlan autostate**
3. switch(config)# **feature interface-vlan**
4. switch(config)# **interface vlan *vlan id***
5. (config-if)# **[no] autostate**
6. (config-if)# **end**

7. show running-config interface vlan *vlan id*

DETAILED STEPS

	Command or Action	Purpose
Step 1	switch# configure terminal	Enters global configuration mode.
Step 2	switch(config)# system default interface-vlan autostate	Reenables the system default autostate behavior on Switching Virtual Interface (SVI) in a VLAN. Use the no form of the command to disable the autostate behavior on SVI.
Step 3	switch(config)# feature interface-vlan	Enables the creation of VLAN interfaces SVI.
Step 4	switch(config)# interface vlan <i>vlan id</i>	Disables the VLAN interface and enters interface configuration mode.
Step 5	(config-if)# [no] autostate	Disables the default autostate behavior of SVIs on the VLAN interface.
Step 6	(config-if)# end	Returns to privileged EXEC mode.
Step 7	show running-config interface vlan <i>vlan id</i>	(Optional) Displays the running configuration for a specific port channel.

Example

This example shows how to configure the SVI Autostate Disable feature:

```
switch# configure terminal
switch(config)# system default interface-vlan autostate
switch(config)# feature interface-vlan
switch(config)# interface vlan 2
switch(config-if)# no autostate
switch(config-if)# end
```

Configuring a DHCP Client on an Interface

You can configure the IP address of a DHCP client on an SVI, a management interface, or a physical Ethernet interface.

SUMMARY STEPS

1. switch# **configure terminal**
2. switch(config)# **interface ethernet *type slot/port* | mgmt *mgmt-interface-number* | vlan *vlan id***
3. switch(config-if)# **[no] ip | ipv6 address dhcp**
4. (Optional) switch(config)# **copy running-config startup-config**

DETAILED STEPS

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

	Command or Action	Purpose
Step 2	switch(config)# interface ethernet <i>type slot/port</i> mgmt <i>mgmt-interface-number</i> vlan <i>vlan id</i>	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] 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 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 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 dhcp
```

Verifying the Layer 3 Interfaces Configuration

Use one of the following commands to verify the configuration:

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.

Command	Purpose
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> private-vlan mapping	Displays the VLAN interface private VLAN information.
show interface vlan <i>number</i> status	Displays the VLAN interface administrative status and protocol status.

Triggering the Layer 3 Interface Consistency Checker

You can manually trigger the Layer 3 interface consistency checker to compare the hardware and software configuration of all physical interfaces in a module and display the results. To manually trigger the Layer 3 Interface consistency checker and display the results, use the following command in any mode:

SUMMARY STEPS

1. **show consistency-checker l3-interface module** *slot*

DETAILED STEPS

	Command or Action	Purpose
Step 1	show consistency-checker l3-interface module <i>slot</i>	Starts the Layer 3 interface consistency check on all Layer 3 physical interfaces of a module that are up and displays its results.

Example

This example shows how to trigger the Layer 3 interface consistency check and display its results:

```
switch# show consistency-checker l3-interface module 1
L3 LIF Checks: L3 Vlan, CML Flags, IPv4 Enable
Consistency Check: PASSED
No inconsistencies found for:
  Ethernet1/17
  Ethernet1/49
  Ethernet1/50
```

Monitoring Layer 3 Interfaces

Use one of the following commands to display statistics about the feature:

Command	Purpose
load-interval <i>seconds</i> counter { 1 2 3 } <i>seconds</i>	Sets three different sampling intervals to bit-rate and packet-rate statistics. The range is from 5 seconds 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 <i>load-interval-id</i>	Displays the Layer 3 interface input and output counters. The load interval ID specifies a single load interval ID to display the input and output rates. The load interval ID ranges between 1 and 3.
show interface ethernet <i>slot/port</i> counters 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 <i>slot/port</i> counters error	Displays the Layer 3 interface input and output errors.
show interface ethernet <i>slot/port</i> counters snmp	Displays the Layer 3 interface counters reported by SNMP MIBs. You cannot clear these counters.
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).

Command	Purpose
show interface loopback <i>number</i> counters 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).
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>counters</i> snmp	Displays the VLAN interface counters reported by SNMP MIBs. You cannot clear these counters.

Configuration Examples for Layer 3 Interfaces

This example shows how to configure Ethernet subinterfaces:

```
switch# configuration terminal
switch(config)# interface ethernet 2/1.10
switch(config-if)# description Layer 3 for VLAN 10
switch(config-if)# encapsulation dot1q 10
switch(config-if)# ip address 192.0.2.1/8
switch(config-if)# copy running-config startup-config
```

This example shows how to configure a VLAN interface:

```
switch# configuration terminal
switch(config)# interface vlan 100
switch(config-if)# no switchport

switch(config-if)# ipv6 address 33:0DB::2/8
switch(config-if)# copy running-config startup-config
```

This example shows how to configure Switching Virtual Interface (SVI) Autostate Disable:

```
switch# configure terminal
switch(config)# system default interface-vlan autostate
switch(config)# feature interface-vlan
switch(config)# interface vlan 2
switch(config-if)# no autostate
switch(config-if)# end
switch# show running-config interface vlan 2
```

This example shows how to configure a loopback interface:

```
switch# configuration terminal
switch(config)# interface loopback 3
switch(config-if)# no switchport
```

```
switch(config-if)# ip address 192.0.2.2/32
switch(config-if)# copy running-config startup-config
```

This example shows how to configure the three sample load intervals for an Ethernet port:

```
switch# configure terminal
switch(config)# interface ethernet 1/3
switch(config-if)# load-interval counter 1 5
switch(config-if)# load-interval counter 2 135
switch(config-if)# load-interval counter 3 225
switch(config-if)#
```

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 for Layer 3 Interfaces

Related Topics	Document Title
Command syntax	<i>Cisco Nexus 3000 Series Command Reference</i>
IP	“Configuring IP” chapter in the <i>Cisco Nexus 3000 Series NX-OS Unicast Routing Configuration Guide</i>
VLAN	“Configuring VLANs” chapter in the <i>Cisco Nexus 3000 Series NX-OS Layer 2 Switching Configuration Guide</i>

MIBs for Layer 3 Interfaces

MIB	MIB Link
CISCO-IF-EXTENSION-MIB	To locate and download MIBs, go to the following URL: http://www.cisco.com/public/sw-center/netmgmt/cmtk/mibs.shtml
ETHERLIKE-MIB	

Standards for Layer 3 Interfaces

No new or modified standards are supported by this feature, and support for existing standards has not been modified by this feature.

Feature History for Layer 3 Interfaces

Feature Name	Release	Feature Information
show interface vlan <i>vlan-id</i> counters command	5.0(3)U3(1)	The show interface vlan <i>vlan-id</i> counters command has been enhanced to correctly show input and output packet counts.



CHAPTER 4

Configuring Port Channels

- [Information About Port Channels](#), on page 63
- [Configuring Port Channels](#), on page 72
- [Verifying Port Channel Configuration](#), on page 82
- [Triggering the Port Channel Membership Consistency Checker](#), on page 83
- [Verifying the Load-Balancing Outgoing Port ID](#), on page 84
- [Feature History for Port Channels](#), on page 84
- [Port Profiles](#), on page 85
- [Configuring Port Profiles](#), on page 86
- [Creating a Port Profile](#), on page 86
- [Entering Port-Profile Configuration Mode and Modifying a Port Profile](#), on page 88
- [Assigning a Port Profile to a Range of Interfaces](#), on page 88
- [Enabling a Specific Port Profile](#), on page 89
- [Inheriting a Port Profile](#), on page 90
- [Removing a Port Profile from a Range of Interfaces](#), on page 91
- [Removing an Inherited Port Profile](#), on page 92

Information About Port Channels

A port channel bundles individual interfaces into a group 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 create a port channel by bundling compatible interfaces. You can configure and run either static port channels or port channels running the Link Aggregation Control Protocol (LACP).

Any configuration changes that you apply to the port channel are applied to each member interface of that port channel. For example, if you configure Spanning Tree Protocol (STP) parameters on the port channel, Cisco NX-OS applies those parameters to each interface in the port channel.

You can use static port channels, with no associated protocol, for a simplified configuration. For more efficient use of the port channel, you can use the Link Aggregation Control Protocol (LACP), which is defined in IEEE 802.3ad. When you use LACP, the link passes protocol packets.

Related Topics

[LACP Overview](#), on page 69

Understanding Port Channels

Using port channels, Cisco NX-OS provides wider bandwidth, redundancy, and load balancing across the channels.

You can collect ports into a static port channel or you can enable the Link Aggregation Control Protocol (LACP). Configuring port channels with LACP requires slightly different steps than configuring static port channels. For information on port channel configuration limits, see the *Verified Scalability* document for your platform. For more information about load balancing, see [Load Balancing Using Port Channels, on page 66](#).



Note Cisco NX-OS does not support Port Aggregation Protocol (PAgP) for port channels.

A port channel bundles individual links into a channel group to create a single logical link that provides the aggregate bandwidth of several physical links. If a member port within a port channel fails, traffic previously carried over the failed link switches to the remaining member ports within the port channel.

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 operate in full-duplex mode. When you are running static port channels without LACP, the individual links are all in the on channel mode; you cannot change this mode without enabling LACP.



Note You cannot change the mode from ON to Active or from ON to Passive.

You can create a port channel directly by creating the port-channel interface, or you can create a channel group that acts to aggregate individual ports into a bundle. When you associate an interface with a channel group, Cisco NX-OS creates a matching port channel automatically if the port channel does not already exist. You can also create the port channel first. In this instance, Cisco NX-OS creates an empty channel group with the same channel number as the port channel and takes the default configuration.



Note A 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.

Compatibility Requirements

When you add an interface to a port channel group, Cisco NX-OS checks certain interface attributes to ensure that the interface is compatible with the channel group. Cisco NX-OS also checks a number of operational attributes for an interface before allowing that interface to participate in the port-channel aggregation.

The compatibility check includes the following operational attributes:

- Port mode
- Access VLAN
- Trunk native VLAN
- Allowed VLAN list
- Speed

- 802.3x flow control setting
- MTU
- Broadcast/Unicast/Multicast Storm Control setting
- Priority-Flow-Control
- Untagged CoS

Use the **show port-channel compatibility-parameters** command to see the full list of compatibility checks that Cisco NX-OS uses.

You can only add interfaces configured with the channel mode set to on to static port channels. You can also 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.

When the interface joins a port channel, the following individual parameters are replaced with the values on the port channel:

- Bandwidth
- MAC address
- Spanning Tree Protocol

The following interface parameters remain unaffected when the interface joins a port channel:

- Description
- CDP
- LACP port priority
- Debounce

After you enable forcing a port to be added to a channel group by entering the **channel-group force** command, the following two conditions occur:

- When an interface joins a port channel, the following parameters are removed and they are operationally replaced with the values on the port channel; however, this change will not be reflected in the running configuration for the interface:
 - QoS
 - Bandwidth
 - Delay
 - STP
 - Service policy
 - ACLs
- When an interface joins or leaves a port channel, the following parameters remain unaffected:
 - Beacon
 - Description

- CDP
- LACP port priority
- Debounce
- UDLD
- Shutdown
- SNMP traps

Load Balancing Using Port Channels

Cisco NX-OS load balances traffic across all operational interfaces in a port channel by reducing part of the binary pattern formed from the addresses in the frame to a numerical value that selects one of the links in the channel. Port channels provide load balancing by default.

The basic configuration uses the following criteria to select the link:

- For a Layer 2 frame, it uses the source and destination MAC addresses.
- For a Layer 3 frame, it uses the source and destination MAC addresses and the source and destination IP addresses.
- For a Layer 4 frame, it uses the source and destination MAC addresses and the source and destination IP addresses.



Note You have the option to include the source and destination port number for the Layer 4 frame.

You can configure the switch to use one of the following methods (see the following table for more details) 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
- Destination TCP/UDP port number
- Source TCP/UDP port number
- Source and destination TCP/UDP port number

Table 3: Port Channel Load-Balancing Criteria

Configuration	Layer 2 Criteria	Layer 3 Criteria	Layer 4 Criteria
Destination MAC	Destination MAC	Destination MAC	Destination MAC
Source MAC	Source MAC	Source MAC	Source MAC
Source and destination MAC	Source and destination MAC	Source and destination MAC	Source and destination MAC
Destination IP	Destination MAC	Destination MAC, destination IP	Destination MAC, destination IP
Source IP	Source MAC	Source MAC, source IP	Source MAC, source IP
Source and destination IP	Source and destination MAC	Source and destination MAC, source and destination IP	Source and destination MAC, source and destination IP
Destination TCP/UDP port	Destination MAC	Destination MAC, destination IP	Destination MAC, destination IP, destination port
Source TCP/UDP port	Source MAC	Source MAC, source IP	Source MAC, source IP, source port
Source and destination TCP/UDP port	Source and destination MAC	Source and destination MAC, source and destination IP	Source and destination MAC, source and destination IP, source and destination port

Use the option that provides the balance criteria with the greatest variety in your configuration. For example, if the traffic on a port channel is going only to a single MAC address and you use the destination MAC address as the basis of port-channel load balancing, the port channel always chooses the same link in that port channel; using source addresses or IP addresses might result in better load balancing.

Regardless of the load-balancing algorithm configured, 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 This does not apply to Cisco Nexus 3500 Series switches.



Note The hardware multicast hw-hash command is not supported on Cisco Nexus 3000 Series switches and Cisco Nexus 3100 Series switches. It is recommended not to configure this command on these switches. By default, Cisco Nexus 3000 Series switches and Cisco Nexus 3100 Series switches hash multicast traffic.



Note Only the default load-balancing methods are currently supported based on src-dst ip and l4 ports for IP packets and src-dst mac for non-ip packets on the Cisco Nexus 34180YC and 3464C switches

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 supports only unicast traffic.

The resilient hashing system in Cisco Nexus 3100 Series switches maps flows to physical ports. In case a link fails, the flows assigned to the failed link are redistributed uniformly among the working links. The existing flows through the working links are not rehashed and their packets are not delivered out of order.

Resilient hashing is supported only by ECMP groups and on port channel interfaces. When a link is added to the port channel or ECMP group, some of the flows hashed to the existing links are rehashed to the new link, but not across all existing links.

Resilient hashing supports IPv4 and IPv6 unicast traffic, but it does not support IPv4 multicast traffic.

Resilient hashing is not supported on the Cisco Nexus 34180YC and 3464C switches.

Hashing for NVGRE Traffic

You can use Network Virtualization using Generic Routing Encapsulation (NVGRE) to virtualize and extend a network so that Layer 2 and Layer 3 topologies are created across distributed data centers. NVGRE uses encapsulation and tunneling. NVGRE endpoints are network devices that act as interfaces between the physical and virtualized networks.

Data frames are encapsulated or decapsulated at NVGRE endpoints using GRE tunneling. The endpoints obtain the destination address for each data frame from the Tenant Network Identifier (TNI). The Key field in the GRE header holds the 24-bit TNI. Each TNI represents a specific tenant's subnet address.

Cisco NX-OS Release 6.0(2)U2(1) supports hashing for transit NVGRE traffic. You can configure the switch to include the GRE Key field present in the GRE header in hash computations when NVGRE traffic is forwarded over a port channel or an Equal Cost Multipath (ECMP).

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.

Cisco NX-OS Release 6.0(2)U2(3) introduces symmetric hashing. 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.

Symmetric hashing is supported only on Cisco Nexus 3100 Series switches.

Only the following load-balancing algorithms support symmetric hashing:

- source-dest-ip-only
- source-dest-port-only
- source-dest-ip
- source-dest-port
- source-dest-ip-gre

Understanding LACP


LACP Overview



Note You must enable the LACP feature before you can configure and use LACP functions.

The following figure shows how individual links can be combined into LACP port channels and channel groups as well as function as individual links.

Figure 4: Individual Links Combined into a Port Channel



With LACP, just like with static port channels, you can bundle up to 16 interfaces in a channel group.



Note When you delete the port channel, Cisco NX-OS automatically deletes the associated channel group. All member interfaces revert to their previous configuration.

You cannot disable LACP while any LACP configurations are present.

LACP ID Parameters

LACP uses the following parameters:

- LACP system priority—Each system that runs LACP has an LACP system priority value. You can accept the default value of 32768 for this parameter, or you can configure a value between 1 and 65535. LACP uses the system priority with the MAC address to form the system ID and also uses the system priority during negotiation with other devices. A higher system priority value means a lower priority.



Note The LACP system ID is the combination of the LACP system priority value and the MAC address.

- LACP port priority—Each port configured to use LACP has an LACP port priority. You can accept the default value of 32768 for the LACP port priority, or you can configure a value between 1 and 65535. LACP uses the port priority with the port number to form the port identifier. LACP uses the port priority to decide which ports should be put in standby mode when there is a limitation that prevents all compatible ports from aggregating and which ports should be put into active mode. A higher port priority value means a lower priority for LACP. You can configure the port priority so that specified ports have a lower priority for LACP and are most likely to be chosen as active links, rather than hot-standby links.
- LACP administrative key—LACP automatically configures an administrative key value equal to the channel-group number on each port configured to use LACP. The administrative key defines the ability of a port to aggregate with other ports. A port's ability to aggregate with other ports is determined by these factors:
 - Port physical characteristics, such as the data rate, the duplex capability, and the point-to-point or shared medium state
 - Configuration restrictions that you establish

Channel Modes

Individual interfaces in port channels are configured with channel modes. When you run static port channels, with no protocol, the channel mode is always set to on. After you enable LACP globally on the device, you enable LACP for each channel by setting the channel mode for each interface to active or passive. You can configure either channel mode for individual links in the LACP channel group.



Note You must enable LACP globally before you can configure an interface in either the active or passive channel mode.

The following table describes the channel modes.

Table 4: Channel Modes for Individual Links in a Port Channel

Channel Mode	Description
passive	LACP mode that places a port into a passive negotiating state, in which the port responds to LACP packets that it receives but does not initiate LACP negotiation.

Channel Mode	Description
active	LACP mode that places a port into an active negotiating state, in which the port initiates negotiations with other ports by sending LACP packets.
on	<p>All static port channels, that is, that are not running LACP, remain in this mode. If you attempt to change the channel mode to active or passive before enabling LACP, the device returns an error message.</p> <p>You enable LACP on each channel by configuring the interface in that channel for the channel mode as either active or passive. When an LACP attempts to negotiate with an interface in the on state, it does not receive any LACP packets and becomes an individual link with that interface; it does not join the LACP channel group.</p> <p>The no lacp suspend-individual configuration is supported by default on Cisco Nexus 3000 switches.</p>

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.

Ports can form an LACP port channel when they are in different LACP modes as long as the modes are compatible as in the following examples:

- A port in active mode can form a port channel successfully with another port that is in active mode.
- A port in active mode can form a port channel with another port in passive mode.
- A port in passive mode cannot form a port channel with another port that is also in passive mode because neither port will initiate negotiation.
- A port in on mode is not running LACP.

LACP Marker Responders

Using port channels, data traffic may be dynamically redistributed due to either a link failure or load balancing. LACP uses the Marker Protocol to ensure that frames are not duplicated or reordered because of this redistribution. Cisco NX-OS supports only Marker Responders.

LACP-Enabled and Static Port Channel Differences

The following table provides a brief summary of major differences between port channels with LACP enabled and static port channels. For information about the maximum configuration limits, see the *Verified Scalability* document for your device.

Table 5: Port Channels with LACP Enabled and Static Port Channels

Configurations	Port Channels with LACP Enabled	Static Port Channels
Protocol applied	Enable globally.	Not applicable.

Configurations	Port Channels with LACP Enabled	Static Port Channels
Channel mode of links	Can be either: <ul style="list-style-type: none"> • Active • Passive 	Can only be On.

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 MinLinks feature does the following:

- Configures the minimum number of port channel interfaces that must be linked and bundled in the LACP port channel.
- Prevents a low-bandwidth LACP port channel from becoming active.
- Causes the LACP port channel to become inactive if only a few active members ports 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.

Configuring Port Channels

Creating a Port Channel

You can create a port channel before creating a channel group. Cisco NX-OS automatically creates the associated channel group.



Note If you want LACP-based port channels, you need to enable LACP.



Note Channel member ports cannot be a source or destination SPAN port.

SUMMARY STEPS

1. switch# **configure terminal**
2. switch(config)# **interface port-channel** *channel-number*
3. switch(config)# **no interface port-channel** *channel-number*

DETAILED STEPS

	Command or Action	Purpose
Step 1	switch# configure terminal	Enters global configuration mode.
Step 2	switch(config)# interface port-channel <i>channel-number</i>	Specifies the port-channel interface to configure, and enters the interface configuration mode. The range is from 1 to 4096. Cisco NX-OS automatically creates the channel group if it does not already exist.
Step 3	switch(config)# no interface port-channel <i>channel-number</i>	Removes the port channel and deletes the associated channel group.

Example

This example shows how to create a port channel:

```
switch# configure terminal
switch (config)# interface port-channel 1
```

Adding a Port to a Port Channel

You can add a port to a new channel group or to a channel group that already contains ports. Cisco NX-OS creates the port channel associated with this channel group if the port channel does not already exist.



Note If you want LACP-based port channels, you need to enable LACP.

SUMMARY STEPS

1. switch# **configure terminal**
2. switch(config)# **interface** *type slot/port*
3. (Optional) switch(config-if)# **switchport mode trunk**
4. (Optional) switch(config-if)# **switchport trunk** {**allowed vlan** *vlan-id* | **native vlan** *vlan-id*}
5. switch(config-if)# **channel-group** *channel-number*
6. (Optional) switch(config-if)# **no channel-group**

DETAILED STEPS

	Command or Action	Purpose
Step 1	switch# configure terminal	Enters global configuration mode.
Step 2	switch(config)# interface <i>type slot/port</i>	Specifies the interface that you want to add to a channel group and enters the interface configuration mode.
Step 3	(Optional) switch(config-if)# switchport mode trunk	Configures the interface as a trunk port.
Step 4	(Optional) switch(config-if)# switchport trunk {allowed vlan <i>vlan-id</i> native vlan <i>vlan-id</i>}	Configures necessary parameters for a trunk port.
Step 5	switch(config-if)# channel-group <i>channel-number</i>	Configures the port in a channel group and sets the mode. The channel-number range is from 1 to 4096. Cisco NX-OS creates the port channel associated with this channel group if the port channel does not already exist. This is called implicit port channel creation.
Step 6	(Optional) switch(config-if)# no channel-group	Removes the port from the channel group. The port reverts to its original configuration.

Example

This example shows how to add an Ethernet interface 1/4 to channel group 1:

```
switch# configure terminal
switch (config)# interface ethernet 1/4
switch(config-if)# switchport mode trunk
switch(config-if)# channel-group 1
```

Configuring Load Balancing Using Port Channels

You can configure the load-balancing algorithm for port channels that applies to the entire device.



Note If you want LACP-based port channels, you need to enable LACP.

SUMMARY STEPS

1. switch# **configure terminal**
2. switch(config)# **port-channel load-balance ethernet** {[**destination-ip** | **destination-ip-gre** | **destination-mac** | **destination-port** | **source-dest-ip** | **source-dest-ip-gre** | **source-dest-mac** | **source-dest-port** | **source-ip** | **source-ip-gre** | **source-mac** | **source-port**] **symmetric** | **crc-poly**}
3. (Optional) switch(config)# **no port-channel load-balance ethernet**
4. (Optional) switch# **show port-channel load-balance**

DETAILED STEPS

	Command or Action	Purpose
Step 1	switch# configure terminal	Enters global configuration mode.
Step 2	switch(config)# port-channel load-balance ethernet {[destination-ip destination-ip-gre destination-mac destination-port source-dest-ip source-dest-ip-gre source-dest-mac source-dest-port source-ip source-ip-gre source-mac source-port] symmetric crc-poly }	Specifies the load-balancing algorithm and hash for the device. The range depends on the device. The default is source-dest-mac . Note The optional destination-ip-gre , source-dest-ip-gre and source-ip-gre keywords are used to include the NVGRE key in the hash computation. Inclusion of the NVGRE key is not enabled by default in the case of port channels. You must configure it explicitly by using these optional keywords. The optional symmetric keyword is used to enable or disable symmetric hashing. Symmetric hashing forces bi-directional traffic to use the same physical interface. Only the following load-balancing algorithms support symmetric hashing: <ul style="list-style-type: none"> • source-dest-ip-only • source-dest-port-only • source-dest-ip • source-dest-port • source-dest-ip-gre
Step 3	(Optional) switch(config)# no port-channel load-balance ethernet	Restores the default load-balancing algorithm of source-dest-mac.
Step 4	(Optional) switch# show port-channel load-balance	Displays the port-channel load-balancing algorithm.

Example

This example shows how to configure source IP load balancing for port channels:

```
switch# configure terminal
switch (config)# port-channel load-balance ethernet source-ip
```

This example shows how to configure symmetric hashing for port channels:

```
switch# configure terminal
switch (config)# port-channel load-balance ethernet source-dest-ip-only symmetric
```

Enabling LACP

LACP is disabled by default; you must enable LACP before you begin LACP configuration. You cannot disable LACP while any LACP configuration is present.

LACP learns the capabilities of LAN port groups dynamically and informs the other LAN ports. Once LACP identifies correctly matched Ethernet links, it facilitates grouping the links into a port channel. The port channel is then added to the spanning tree as a single bridge port.

SUMMARY STEPS

1. switch# **configure terminal**
2. switch(config)# **feature lacp**
3. (Optional) switch(config)# **show feature**

DETAILED STEPS

	Command or Action	Purpose
Step 1	switch# configure terminal	Enters global configuration mode.
Step 2	switch(config)# feature lacp	Enables LACP on the switch.
Step 3	(Optional) switch(config)# show feature	Displays enabled features.

Example

This example shows how to enable LACP:

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

Configuring the Channel Mode for a Port

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 protocol, all interfaces on both sides of the link remain in the on channel mode.

Before you begin

Ensure that you have enabled the LACP feature.

SUMMARY STEPS

1. switch# **configure terminal**
2. switch(config)# **interface** *type slot/port*
3. switch(config-if)# **channel-group** *channel-number* [**force**] [**mode** {**on** | **active** | **passive**}]
4. switch(config-if)# **no channel-group** *number mode*

DETAILED STEPS

	Command or Action	Purpose
Step 1	switch# configure terminal	Enters global configuration mode.
Step 2	switch(config)# interface <i>type slot/port</i>	Specifies the interface to configure, and enters the interface configuration mode.
Step 3	switch(config-if)# channel-group <i>channel-number</i> [force] [mode { on active passive }]	<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>force—Specifies that the LAN port be forcefully added to the channel group.</p> <p>mode—Specifies the port channel mode of the interface.</p> <p>active—Specifies that when you enable LACP, this command enables LACP on the specified interface. The interface is in an active negotiating state in which the port initiates negotiations with other ports by sending LACP packets.</p> <p>on—(Default mode) Specifies that all port channels that are not running LACP remain in this mode.</p> <p>passive—Enables LACP only if an LACP device is detected. The interface is in a passive negotiation state in which the port responds to LACP packets that it receives but does not initiate LACP negotiation.</p> <p>When you run port channels with no associated protocol, the channel mode is always on.</p>
Step 4	switch(config-if)# no channel-group <i>number mode</i>	Returns the port mode to on for the specified interface.

Example

This example shows how to set the LACP-enabled interface to active port-channel mode for Ethernet interface 1/4 in channel group 5:

```
switch# configure terminal
switch (config)# interface ethernet 1/4
switch(config-if)# channel-group 5 mode active
```

Configuring LACP Port Channel MinLinks

The MinLink feature works only with LACP port channels. The device allows you to configure this feature in non-LACP port channels, but the feature is not operational.



Important We recommend that you configure the LACP MinLink feature on both ends of your LACP port channel, that is, on both the switches. Configuring the **lACP min-links** command on only one end of the port channel might result in link flapping.

SUMMARY STEPS

1. switch# **configure terminal**
2. switch(config)# **interface port-channel** *number*
3. switch(config-if)# [**no**] **lACP min-links** *number*
4. (Optional) switch(config)# **show running-config interface port-channel** *number*

DETAILED STEPS

	Command or Action	Purpose
Step 1	switch# configure terminal	Enters global configuration mode.
Step 2	switch(config)# interface port-channel <i>number</i>	Specifies the interface to configure.
Step 3	switch(config-if)# [no] lACP min-links <i>number</i>	Configures the number of minimum links. The default value for <i>number</i> is 1. The range is from 1 to 16. Note Starting with Release 7.0(3)I2(1), the maximum number of supported LACP min-links is 16. Use the no form of this command to disable this feature.
Step 4	(Optional) switch(config)# show running-config interface port-channel <i>number</i>	Displays the port channel configuration of the interface.

Example

This example shows how to configure the minimum number of links that must be up for the bundle as a whole to be labeled *up*:

```
switch#configure terminal
switch(config)#interface port-channel 3
switch(config-if)#lACP min-links 3
switch(config)#show running-config interface port-channel 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: <pre>switch(config)# no lacp max-bundle</pre>	Restores the default port-channel max-bundle configuration.

Before you begin

Ensure that you are in the correct port-channel interface.

SUMMARY STEPS

1. **configure terminal**
2. **interface port-channel** *number*
3. **lacp max-bundle** *number*
4. **show running-config interface port-channel** *<number>*

DETAILED STEPS

	Command or Action	Purpose
Step 1	configure terminal Example: <pre>switch# configure terminal</pre> <pre>switch(config)#</pre>	Enters global configuration mode.
Step 2	interface port-channel <i>number</i> Example: <pre>switch(config)# interface port-channel 3</pre> <pre>switch(config-if)#</pre>	Specifies an interface to configure.
Step 3	lacp max-bundle <i>number</i> Example: <pre>switch(config-if)# lacp max-bundle <number></pre>	Configures the maximum number of active bundled LACP ports that are allowed in a port channel. 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 <i>pc_max_links_config</i> and <i>pc_max_active_members</i> that is allowed in the port channel.
Step 4	show running-config interface port-channel <i><number></i> Example:	(Optional) Displays the port-channel configuration for the interface.

	Command or Action	Purpose
	switch(config-if)# show running-config interface port-channel 3	

Example

This example shows how to configure the maximum number of active bundled LACP ports:

```
switch# configure terminal
switch# interface port-channel 3
switch (config-if)# lacp max-bundle 3
switch (config-if)# show running-config interface port-channel 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.

Before you begin

Ensure that you have enabled the LACP feature.

SUMMARY STEPS

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

DETAILED STEPS

	Command or Action	Purpose
Step 1	switch# configure terminal	Enters global configuration mode.
Step 2	switch(config)# interface <i>type slot/port</i>	Specifies the interface to configure and enters the interface configuration mode. Note you can set the lacp rate only on the ports that are administratively down.
Step 3	switch(config-if)# lacp rate fast	Configures the fast rate (one second) at which LACP control packets are sent to an LACP-supported interface.

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 and System ID

The LACP system ID is the combination of the LACP system priority value and the MAC address.

Before you begin

Ensure that you have enabled the LACP feature.

SUMMARY STEPS

1. switch# **configure terminal**
2. switch(config)# **lacp system-priority** *priority*
3. (Optional) switch# **show lacp system-identifier**

DETAILED STEPS

	Command or Action	Purpose
Step 1	switch# configure terminal	Enters global configuration mode.
Step 2	switch(config)# lacp system-priority <i>priority</i>	Configures the system priority for use with LACP. Valid values are 1 through 65535, and higher numbers have lower priority. The default value is 32768.
Step 3	(Optional) switch# show lacp system-identifier	Displays the LACP system identifier.

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

You can configure each link in the LACP port channel for the port priority.

Before you begin

Ensure that you have enabled the LACP feature.

SUMMARY STEPS

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

DETAILED STEPS

	Command or Action	Purpose
Step 1	switch# configure terminal	Enters global configuration mode.
Step 2	switch(config)# interface <i>type slot/port</i>	Specifies the interface to configure, and enters the interface configuration mode.
Step 3	switch(config-if)# lACP port-priority <i>priority</i>	Configures the port priority for use with LACP. Valid values are 1 through 65535, and higher numbers have lower priority. The default value is 32768.

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
```

Verifying Port Channel Configuration

Use the following command to verify the port channel configuration information:

Command	Purpose
show interface port channel <i>channel-number</i>	Displays the status of a port channel interface.
show feature	Displays enabled features.
show resource	Displays the number of resources currently available in the system.
show lACP {counters interface <i>type slot/port</i> neighbor port-channel system-identifier}	Displays LACP information.
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 summary	Displays a summary for the port channel interfaces.
show port-channel traffic	Displays the traffic statistics for port channels.

Command	Purpose
show port-channel usage	Displays the range of used and unused channel numbers.
show port-channel database	Displays information on current running of the port channel feature.
show port-channel load-balance	Displays information about load-balancing using port channels.

Triggering the Port Channel Membership Consistency Checker

You can manually trigger the port channel membership consistency checker to compare the hardware and software configuration of all ports in a port channel and display the results. To manually trigger the port channel membership consistency checker and display the results, use the following command in any mode:

SUMMARY STEPS

1. switch# **show consistency-checker membership port-channels**

DETAILED STEPS

	Command or Action	Purpose
Step 1	switch# show consistency-checker membership port-channels	Starts a port channel membership consistency check on the member ports of a port channel and displays its results.

Example

This example shows how to trigger a port channel membership consistency check and display its results:

```
switch# show consistency-checker membership port-channels
Checks: Trunk group and trunk membership table.
Consistency Check: PASSED
No Inconsistencies found for port-channel1111:
  Module:1, Unit:0
    ['Ethernet1/4', 'Ethernet1/5', 'Ethernet1/6']
No Inconsistencies found for port-channel2211:
  Module:1, Unit:0
    ['Ethernet1/7', 'Ethernet1/8', 'Ethernet1/9', 'Ethernet1/10']
No Inconsistencies found for port-channel3311:
  Module:1, Unit:0
    ['Ethernet1/11', 'Ethernet1/12', 'Ethernet1/13', 'Ethernet1/14']
No Inconsistencies found for port-channel4095:
  Module:1, Unit:0
    ['Ethernet1/33', 'Ethernet1/34', 'Ethernet1/35', 'Ethernet1/36', 'Ethernet1
/37', 'Ethernet1/38', 'Ethernet1/39', 'Ethernet1/40', 'Ethernet1/41', 'Ethernet1
/42', 'Ethernet1/43', 'Ethernet1/44', 'Ethernet1/45', 'Ethernet1/46', 'Ethernet1
/47', 'Ethernet1/48', 'Ethernet1/29', 'Ethernet1/30', 'Ethernet1/31', 'Ethernet1
/32']
```

Verifying the Load-Balancing Outgoing Port ID

Command Guidelines

The **show port-channel load-balance** command allows you to verify which ports a given frame is hashed to on a port channel. You need to specify the VLAN and the destination MAC in order to get accurate results.



Note Certain traffic flows are not subject to hashing such as when there is a single port in a port-channel.

The **show port-channel load-balance** command supports only unicast traffic hashing. Multicast traffic hashing is not supported.

To display the load-balancing outgoing port ID, perform one of the tasks:

Command	Purpose
switch# show port-channel load-balance forwarding-path interface port-channel <i>port-channel-id</i> vlan <i>vlan-id</i> dst-ip src-ip dst-mac src-mac l4-src-port <i>port-id</i> l4-dst-port <i>port-id</i> ether-type ether-type ip-proto ip-proto	Displays the outgoing port ID.

Example

This example shows how to display the load balancing outgoing port ID:

```
switch# show port-channel load-balance forwarding-path interface port-channel 10 vlan 1
dst-ip 1.225.225.225 src-ip 1.1.10.10 src-mac aa:bb:cc:dd:ee:ff
l4-src-port 0 l4-dst-port 1
Missing params will be substituted by 0's. Load-balance Algorithm on switch: source-dest-port
  crc8_hash:204 Outgoing port id: Ethernet 1/1 Param(s) used to calculate load balance:
dst-port: 0
src-port: 0
dst-ip: 1.225.225.225
src-ip: 1.1.10.10
dst-mac: 0000.0000.0000
src-mac: aabb.ccdd.eeff
```

Feature History for Port Channels

Feature Name	Release	Feature Information
Minimum Links	5.0(3)U3(1)	Added information about setting up and using the Minimum Links feature.

Port Profiles

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.

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.

Configuring Port Profiles

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:
 - .
 - -
 - _

SUMMARY STEPS

1. **configure terminal**
2. **port-profile** [**type** {**ethernet** | **interface-vlan** | **port-channel**}] *name*
3. **exit**
4. (Optional) **show port-profile**
5. (Optional) **copy running-config startup-config**

DETAILED STEPS

	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	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 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.

SUMMARY STEPS

1. **configure terminal**
2. **port-profile** [type {**ethernet** | **interface-vlan** | **port-channel**}] *name*
3. **exit**
4. (Optional) **show port-profile**
5. (Optional) **copy running-config startup-config**

DETAILED STEPS

	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>	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.

SUMMARY STEPS

1. **configure terminal**
2. **interface** [ethernet *slot/port* | **interface-vlan** *vlan-id* | **port-channel** *number*]
3. **inherit port-profile** *name*
4. **exit**
5. (Optional) **show port-profile**
6. (Optional) **copy running-config startup-config**

DETAILED STEPS

	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.

SUMMARY STEPS

1. **configure terminal**

2. **port-profile** [type {ethernet | interface-vlan | port-channel}] *name*
3. **state enabled**
4. **exit**
5. (Optional) **show port-profile**
6. (Optional) **copy running-config startup-config**

DETAILED STEPS

	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.

SUMMARY STEPS

1. **configure terminal**
2. **port-profile** *name*
3. **inherit port-profile** *name*
4. **exit**
5. (Optional) **show port-profile**
6. (Optional) **copy running-config startup-config**

DETAILED STEPS

	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>inherit port-profile name</code>	Inherits another port profile onto the existing one. The original port profile assumes all the configurations of the inherited 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 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.

SUMMARY STEPS

1. `configure terminal`
2. `interface [ethernet slot/port | interface-vlan vlan-id | port-channel number]`
3. `no inherit port-profile name`
4. `exit`
5. (Optional) `show port-profile`
6. (Optional) `copy running-config startup-config`

DETAILED STEPS

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

	Command or Action	Purpose
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.

SUMMARY STEPS

1. **configure terminal**
2. **port-profile** *name*
3. **no inherit port-profile** *name*
4. **exit**
5. (Optional) **show port-profile**
6. (Optional) **copy running-config startup-config**

DETAILED STEPS

	Command or Action	Purpose
Step 1	configure terminal	Enters the global configuration mode.
Step 2	port-profile <i>name</i>	Enters the port-profile configuration mode for the specified port profile.
Step 3	no inherit port-profile <i>name</i>	Removes an inherited port profile from this port profile.
Step 4	exit	Exits the port-profile configuration mode.

	Command or Action	Purpose
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 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)#
```




CHAPTER 5

Configuring IP Tunnels

- [Information About IP Tunnels, on page 95](#)
- [Prerequisites for IP Tunnels, on page 96](#)
- [Guidelines and Limitations for IP Tunnels, on page 96](#)
- [Default Settings for IP Tunneling, on page 100](#)
- [Configuring IP Tunnels, on page 101](#)
- [Verifying the IP Tunnel Configuration, on page 110](#)
- [Configuration Examples for IP Tunneling, on page 110](#)
- [Related Documents for IP Tunnels, on page 111](#)
- [Standards for IP Tunnels, on page 111](#)
- [Feature History for Configuring IP Tunnels, on page 111](#)

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 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 generic routing encapsulation (GRE), and IP-in-IP encapsulation and decapsulation as carrier protocols.
- **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.

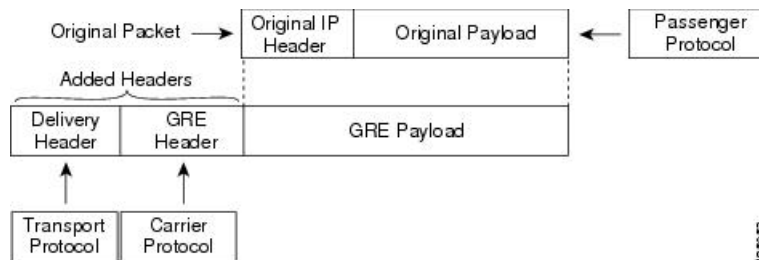
You must enable the tunnel feature before you can configure it.

GRE Tunnels

You can use GRE as the carrier protocol for a variety of passenger protocols. The selection of tunnel interfaces can also be based on the PBR policy.

The 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 5: GRE PDU



Point-to-Point IP-in-IP Tunnel Encapsulation and Decapsulation

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. The selection of these tunnel interfaces can also be based on the PBR policy. This type of tunnel will carry both inbound and outbound traffic.

Multi-Point IP-in-IP Tunnel Decapsulation

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.

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 have installed the Enterprise Services license for Cisco NX-OS.
- You must enable the tunneling feature in a device before you can configure and enable any IP tunnels.

Guidelines and Limitations for IP Tunnels

IP tunnels have the following configuration guidelines and limitations:

- Guidelines for **source-direct** and **ipv6ipv6-decapsulate-any** options for tunnels:

- 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** command. The IP tunnel supports the **tunnel source** command with interface, IPv4 address, IPv6 address, or IPv4 prefix. 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.
 - The **tunnel mode ipip decapsulate-any** is supported for decapsulating IPv4 payload over IPv4 transport. The **tunnel mode ipv6ipv6 decapsulate-any** command supports IPv6 payload over IPv6 transport.
 - The **tunnel source direct** 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** command.
 - The **tunnel source direct** command is supported only on decapsulate-any tunnel modes, for example, **tunnel mode ipip decapsulate-any** and **tunnel mode ipv6ipv6 decapsulate-any**.
 - Auto-recovery is not supported for source-direct.
 - 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 3000 Series switches.
 - To enable IPv6 on ipv6ipv6 decap-any tunnel interface, you must configure a valid IPv6 address or configure the ipv6 address using use-link-local-only CLI command under the interface tunnel interface.
- The hardware limitations on a source direct tunnel are as follows:
 - Source direct tunnel supports Cisco Nexus 3000 Series switches with Network Forwarding Engine (NFE), Application Spine Engine (ASE), and Leaf Spine Engine (LSE). There are limitations in cases 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. In the switches with ASE and LSE, the IPv4 individual scale can be 256 but it is shared with IPv6. If the user plans to configure both v4 and v6 decap any tunnel in the same system, the scale numbers for the switches with NFE for individual IPv4 and IPv6 cannot be guaranteed. However, the scale numbers for the switches with ASE and LSE for individual IPv4 and IPv6 are guaranteed.

There is no CLI command to change these pre-carved scale numbers, for example, allocating X for IPv4 and Y for IPv6.

Whenever the tunnel decap table gets full, the TABLE_FULL error is displayed. If an entry gets deleted after the table is full, the table full error is cleared.

If the tunnel-decap-table is full, the user gets a syslog similar to as follows:

```
2017 Apr 26 10:10:51 switch %$ VDC-1 %$
%IPFIB-2-FIB_HW_IP_TUNNEL_DECAP_TABLE_FULL:
IP TUNNEL decap hardware table full. IP tunnel decapsulation may not work for
some GRE/IPinIP traffic
```

If the table is full and if an entry is deleted from the table because of an interface being operationally down or removal of IP address, the clear syslog for the table is displayed. Deleting of a tunnel removes all the entries that are added as part of that tunnel.

```
2002 Sep 26 10:11:37 switch %$ VDC-1 %$
%IPFIB-2-FIB_HW_IP_TUNNEL_DECAP_TABLE_FULL_CLRDR: IP TUNNEL decap hardware table
full exception cleared
```

• **Table 6: 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.

- The IP-in-IP tunnel decapsulation is supported on IPv6 enabled networks.

```
interface loopback0
  ip address 2001:0:0:4::1/128
!
interface Tunnel 1
  ipv6 address use-link-local-only
  tunnel mode tunnel mode ipv6ip6 decapsulate-any
  tunnel source loopback0
  description IPinIP Decapsulation Interface
  mtu 1476
  no shutdown
```

- Cisco NX-OS software supports the GRE header defined in IETF RFC 2784. Cisco NX-OS software does not support tunnel keys and other options from IETF RFC 1701.
- The Cisco Nexus device supports the following maximum number tunnels:
 - GRE and IP-in-IP regular tunnels-8 tunnels
 - Multipoint IP-in-IP tunnels-32 tunnels
- GRE and IP-in-IP tunnel termination is not supported on SVIs for vPC VLANs.
- Each tunnel will consume one Equal Cost Multipath (ECMP) adjacency.
- The Cisco Nexus device does not support the following features:
 - Path maximum transmission unit (MTU) discovery
 - Tunnel interface statistics
 - Access control lists (ACLs)
 - Unicast reverse path forwarding (URPF)
 - Multicast traffic and associated multicast protocols such as Internet Group Management Protocol (IGMP) and Protocol Independent Multicast (PIM)
- Cisco NX-OS software does not support the Web Cache Control Protocol (WCCP) on tunnel interfaces.
- Cisco NX-OS software supports only Layer-3 traffic.
- Cisco NX-OS software supports ECMP across tunnels and ECMP for tunnel destination.
- IPv6-in-IPv6 tunnels is not supported.
- Limited control protocols, such as Border Gateway Protocol (BGP), Open Shortest Path First (OSPF), and Enhanced Interior Gateway Routing Protocol (EIGRP), are supported for GRE tunnels.
- Starting with Release 6.0(2)U5(1), Cisco Nexus 3000 Series switches drop all the packets when the tunnel is not configured. The packets are also dropped when the tunnel is configured but the tunnel interface is not configured or the tunnel interface is in shut down state.

Point to Point tunnel (Source and Destination) – Cisco Nexus 3000 Series switches decapsulate all IP-in-IP packets destined to it when the command **feature tunnel** is configured and there is an operational tunnel interface configured with the tunnel source and the destination address that matches the incoming packets' outer source and destination addresses. If there is not a source and destination packet match or if the interface is in shutdown state, the packet is dropped.

Decapsulate Tunnel (Source only) - Cisco Nexus 3000 Series switches decapsulate all IP-in-IP packets destined to it when the command **feature tunnel** is configured and there is an operational tunnel interface configured with the tunnel source address that matches the incoming packets' outer destination addresses. If there is not a source packet match or if the interface is in shutdown state, the packet is dropped.

- Starting with Release 6.0(2)U6(1), Cisco Nexus 3000 Series switches support IPv6 in IPv4 with GRE header only. The new control protocols that are supported on the tunnel are:
 - BGP with v6
 - OSPFv3
 - EIGRP over v6
- GRE v4/v6 tunnel configuration is supported only in the default routing mode. It does not support the multicast traffic or multicast protocols, for example, IGMP/PIM. It does not support ACL/QoS policies. It supports a maximum of 8 tunnels in the switch, whether they are all IPinIP or GRE; or any combination of both. The packets that are sent/received over the tunnel and that are destined for the switch, are not counted in the tunnel statistics.
- The Cisco Nexus 3000 Series switches ASIC supports the GRE encapsulation and decapsulation in the hardware.
- On the encapsulation side, the Cisco Nexus 3000 Series switches performs a single lookup in the hardware.
- Since Cisco Nexus 3000 Series switches perform a single lookup in the hardware, the software has to keep the hardware information up-to-date with any changes related to the second lookup, for example, the tunnel destination adjacency.
- On the decapsulation side, the Cisco Nexus 3000 Series switches have a separate table to perform the outer IP header lookup and it does not need an ACL for the same.
- RFC5549 is not supported over tunnels.
- On Cisco Nexus N3K-C34180YC switches, you may not be able to enable tunnel feature or configure tunnels.

Default Settings for IP Tunneling

The following table lists the default settings for IP tunnel parameters.

Table 7: Default IP Tunnel Parameters

Parameters	Default
Tunnel feature	Disabled

Configuring IP Tunnels

Enabling Tunneling

Before you begin

You must enable the tunneling feature before you can configure any IP tunnels.

SUMMARY STEPS

1. switch# **configure terminal**
2. switch(config)# **feature tunnel**
3. switch(config)# **exit**
4. switch(config)# **show feature**
5. (Optional) switch# **copy running-config startup-config**

DETAILED STEPS

	Command or Action	Purpose
Step 1	switch# configure terminal	Enters global configuration mode.
Step 2	switch(config)# feature tunnel	Enables the tunnel feature on the switch.
Step 3	switch(config)# exit	Returns to configuration mode.
Step 4	switch(config)# show feature	Displays the tunnel feature on the switch.
Step 5	(Optional) switch# 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 enable the tunnel feature:

```
switch# configure terminal
switch(config)# feature tunnel
switch(config)# exit
switch(config)# copy running-config startup-config
```

Creating a Tunnel Interface

You can create a tunnel interface and then configure this logical interface for your IP tunnel. GRE mode is the default tunnel mode.

Before you begin

Both the tunnel source and the tunnel destination must exist within the same virtual routing and forwarding (VRF) instance.

Ensure that you have enabled the tunneling feature.

SUMMARY STEPS

1. switch# **configure terminal**
2. switch(config)# [**no**] **interface tunnel** *number*
3. switch(config)# **tunnel mode** {**gre ip** | **ipip** {**ip** | **decapsulate-any**}}
4. switch(config)# **tunnel source** {*ip address* | *interface-name*}
5. switch(config)# **tunnel destination** {*ip address* | *host-name*}
6. (Optional) switch(config)# **tunnel use-vrf** *vrf-name*
7. (Optional) switch(config)# **show interface tunnel** *number*
8. (Optional) switch# **copy running-config startup-config**

DETAILED STEPS

	Command or Action	Purpose
Step 1	switch# configure terminal	Enters global configuration mode.
Step 2	switch(config)# [no] interface tunnel <i>number</i>	Creates a new tunnel interface.
Step 3	switch(config)# 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	switch(config)# tunnel source { <i>ip address</i> <i>interface-name</i> }	Configures the source address for this IP tunnel.
Step 5	switch(config)# tunnel destination { <i>ip address</i> <i>host-name</i> }	Configures the destination address for this IP tunnel.
Step 6	(Optional) switch(config)# tunnel use-vrf <i>vrf-name</i>	Uses the configured VRF instance to look up the tunnel IP destination address.
Step 7	(Optional) switch(config)# show interface tunnel <i>number</i>	Displays the tunnel interface statistics.
Step 8	(Optional) switch# 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 create a tunnel interface:

```

switch# configure terminal
switch(config)# interface tunnel 1
switch(config)# tunnel source ethernet 1/2
switch(config)# tunnel destination 192.0.2.1
switch(config)# copy running-config startup-config

```

Configuring a Tunnel Interface

The **tunnel source direct** and **tunnel mode ipv6ipv6 decapsulate-any** commands are supported on Cisco Nexus 3000 Series switches.

The **tunnel mode ipv6ipv6 decapsulate-any** command supports 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.

SUMMARY STEPS

1. **configure terminal**
2. **interface tunnel** *number*
3. **tunnel mode** {gre ip | ipip | {ip | decapsulate-any}}
4. (Optional) **tunnel mode ipv6ipv6 decapsulate-any**
5. **tunnel source direct**
6. **show interfaces tunnel** *number*
7. **mtu** *value*
8. **copy running-config startup-config**

DETAILED STEPS

	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
		<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	(Optional) tunnel mode ipv6ip6 decapsulate-any	Supports IPv6 payload over IPv6 transport (IPv6inIPv6 packets) This step is applicable for IPv6 networks only.
Step 5	tunnel source direct	Configures IP-in-IP tunnel decapsulation on any directly connected IP addresses. this option is now supported only when the IP-in-IP decapsulation is used to source route the packets through the network.
Step 6	show interfaces tunnel <i>number</i> Example: <pre>switch(config-if)# show interfaces tunnel 1</pre>	(Optional) Displays the tunnel interface statistics.
Step 7	mtu <i>value</i>	<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>
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 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 loopback0
  ip address 2001:0:0:4::1/128
!
interface Tunnel1
  tunnel mode ipip decapsulate-any ipv6
  tunnel source loopback0
  description IPinIP Decapsulation Interface
  mtu 1476
  no shutdown

show running-config interface tunnel 1
interface Tunnel1
  tunnel mode ipv6ipv6 decapsulate-any
  tunnel source direct
  no shutdown

show interface tunnel 1
Tunnel1 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 Tunnel Interface Based on Policy Based Routing

You can create a tunnel interface and then configure this logical interface for your IP tunnel based on the PBR policy.

Before you begin

Ensure that you have enabled the tunneling feature.

SUMMARY STEPS

1. switch# **configure terminal**
2. switch(config)# [**no**] **interface tunnel** *number*
3. switch(config)# **ip address** *ip address*
4. switch(config)# **route-map** *map-name*
5. switch(config-route-map)# **match ip address access-list-name** *name*
6. switch(config-route-map)# **set ip next-hop** *address*

DETAILED STEPS

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

	Command or Action	Purpose
Step 2	switch(config)# [no] interface tunnel <i>number</i>	Creates a new tunnel interface.
Step 3	switch(config)# ip address <i>ip address</i>	Configures an IP address for this interface.
Step 4	switch(config)# route-map <i>map-name</i>	Assigns a route map for IPv4 policy-based routing to the interface
Step 5	switch(config-route-map)# match ip address access-list-name <i>name</i>	Matches an IPv4 address against one or more IP access control lists (ACLs). This command is used for policy-based routing and is ignored by route filtering or redistribution.
Step 6	switch(config-route-map)# set ip next-hop <i>address</i>	Sets the IPv4 next-hop address for policy-based routing. To select tunnel interfaces, you must specify the Tunnel IP addresses as next-hop addresses. This command uses the first valid next-hop address if multiple addresses are configured. Use the load-share option to select ECMP across next-hop entries.

Example

This example shows how to configure a tunnel interface that is based on PBR:

```
switch# configure terminal
switch(config)# interface tunnel 1
switch(config)# ip address 1.1.1.1/24
switch(config)# route-map pbr1
switch(config-route-map)# match ip address access-list-name pbr1
switch(config-route-map)# set ip next-hop 1.1.1.1
```

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.

You can set a tunnel interface to GRE tunnel mode, ipip mode, or ipip decapsulate-only mode. GRE mode is the default tunnel mode. Starting with Release 6.0(2)U6(1), Cisco Nexus 3000 Series switches support IPv6 payload over IPv4 tunnel with GRE header only.

Before you begin

Ensure that you have enabled the tunneling feature.

SUMMARY STEPS

1. switch# **configure terminal**
2. switch(config)# **interface tunnel** *number*
3. switch(config-if)# **tunnel mode** {gre ip | ipip {ip | decapsulate-any}}
4. switch(config-if)# **tunnel use-vrf** *vrf-name*
5. switch(config-if)# **ipv6 address** *IPv6 address*
6. (Optional) switch(config-if)# **show interface tunnel** *number*

7. switch(config-if)# **mtu** *value*
8. (Optional) switch(config-if)# **copy running-config startup-config**

DETAILED STEPS

	Command or Action	Purpose
Step 1	switch# configure terminal	Enters global configuration mode.
Step 2	switch(config)# interface tunnel <i>number</i>	Enters a tunnel interface configuration mode.
Step 3	switch(config-if)# 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	Required: switch(config-if)# tunnel use-vrf <i>vrf-name</i>	Configures tunnel VRF name.
Step 5	Required: switch(config-if)# ipv6 address <i>IPv6 address</i>	Configures the IPv6 address. Note The tunnel source and the destination addresses are still the same (IPv4 address.)
Step 6	(Optional) switch(config-if)# show interface tunnel <i>number</i>	Displays the tunnel interface statistics.
Step 7	switch(config-if)# mtu <i>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
```

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.

SUMMARY STEPS

1. switch# **configure terminal**
2. switch(config)# **interface tunnel** *number*
3. switch(config)# **vrf member** *vrf-name*
4. switch(config)# **ip address** *ip-prefix/length*
5. (Optional) switch(config)# **show vrf** [*vrf-name*] **interface** *interface-type number*
6. (Optional) switch(config-if)# **copy running-config startup-config**

DETAILED STEPS

	Command or Action	Purpose
Step 1	switch# configure terminal	Enters global configuration mode.
Step 2	switch(config)# interface tunnel <i>number</i>	Enters interface configuration mode.
Step 3	switch(config)# vrf member <i>vrf-name</i>	Adds this interface to a VRF.
Step 4	switch(config)# ip address <i>ip-prefix/length</i>	Configures an IP address for this interface. You must do this step after you assign this interface to a VRF.
Step 5	(Optional) switch(config)# show vrf [<i>vrf-name</i>] interface <i>interface-type number</i>	Displays VRF information.
Step 6	(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 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

Use the following commands to verify the configuration:

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> 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

This 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 1/3 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
interface ethernet1/2
 ip address 192.0.2.55/8

router B:
feature tunnel
interface tunnel 0
 ip address 209.165.20.1/8
 tunnel source ethernet 1/3
 tunnel destination 192.0.2.55
 tunnel mode gre ip
interface ethernet 1/3
 ip address 192.0.2.2/8
```

Related Documents for IP Tunnels

Related Topics	Document Title
IP tunnel commands	<i>Cisco Nexus 3000 Series Interfaces Command Reference</i>

Standards for IP Tunnels

No new or modified standards are supported by this feature, and support for existing standards has not been modified by this feature.

Feature History for Configuring IP Tunnels

Table 8: Feature History for Configuring IP Tunnels

Feature Name	Release	Feature Information
Multi-point and Point-to-Point IP-in-IP encapsulation and decapsulation	6.0(2)U2(1)	Support for these tunnel modes was added.
IP tunnels	5.0(3)U4(1)	This feature was introduced.



CHAPTER 6

Configuring Virtual Port Channels

- [Information About vPCs, on page 113](#)
- [Guidelines and Limitations for vPCs, on page 121](#)
- [Verifying the vPC Configuration, on page 122](#)
- [vPC Default Settings, on page 128](#)
- [Configuring vPCs, on page 129](#)
- [Configuring Layer 3 over vPC, on page 145](#)

Information About vPCs

vPC Overview

A virtual port channel (vPC) allows links that are physically connected to two Cisco Nexus devices or Cisco Nexus Fabric Extenders to appear as a single port channel by a third device (see the following figure). The third device can be a switch, server, or any other networking device. You can configure vPCs in topologies that include Cisco Nexus devices connected to Cisco Nexus Fabric Extenders. A vPC can provide multipathing, which allows you to create redundancy by enabling multiple parallel paths between nodes and load balancing traffic where alternative paths exist.

You configure the EtherChannels by using one of the following:

- No protocol
- Link Aggregation Control Protocol (LACP)

When you configure the EtherChannels in a vPC—including the vPC peer link channel—each switch can have up to 16 active links in a single EtherChannel.



Note You must enable the vPC feature before you can configure or run the vPC functionality.

To enable the vPC functionality, you must create a peer-keepalive link and a peer-link under the vPC domain for the two vPC peer switches to provide the vPC functionality.

To create a vPC peer link you configure an EtherChannel on one Cisco Nexus device by using two or more Ethernet ports. On the other switch, you configure another EtherChannel again using two or more Ethernet ports. Connecting these two EtherChannels together creates a vPC peer link.



Note We recommend that you configure the vPC peer-link EtherChannels as trunks.

The vPC domain includes both vPC peer devices, the vPC peer-keepalive link, the vPC peer link, and all of the EtherChannels in the vPC domain connected to the downstream device. You can have only one vPC domain ID on each vPC peer device.



Note Always attach all vPC devices using EtherChannels to both vPC peer devices.

A vPC provides the following benefits:

- Allows a single device to use an EtherChannel across two upstream devices
- Eliminates Spanning Tree Protocol (STP) blocked ports
- Provides a loop-free topology
- Uses all available uplink bandwidth
- Provides fast convergence if either the link or a switch fails
- Provides link-level resiliency
- Assures high availability

Terminology

vPC Terminology

The terminology used in vPCs is as follows:

- vPC—combined EtherChannel between the vPC peer devices and the downstream device.
- vPC peer device—One of a pair of devices that are connected with the special EtherChannel known as the vPC peer link.
- vPC peer link—link used to synchronize states between the vPC peer devices.
- vPC member port—Interfaces that belong to the vPCs.
- vPC domain—domain that includes both vPC peer devices, the vPC peer-keepalive link, and all of the port channels in the vPC connected to the downstream devices. It is also associated to the configuration mode that you must use to assign vPC global parameters. The vPC domain ID must be the same on both switches.
- vPC peer-keepalive link—The peer-keepalive link monitors the vitality of a vPC peer Cisco Nexus device. The peer-keepalive link sends configurable, periodic keepalive messages between vPC peer devices.

No data or synchronization traffic moves over the vPC peer-keepalive link; the only traffic on this link is a message that indicates that the originating switch is operating and running vPCs.

vPC Domain

To create a vPC domain, you must first create a vPC domain ID on each vPC peer switch using a number from 1 to 1000. This ID must be the same on a set of vPC peer devices.

You can configure the EtherChannels and vPC peer links by using LACP or no protocol. When possible, we recommend that you use LACP on the peer-link, because LACP provides configuration checks against a configuration mismatch on the EtherChannel.

The vPC peer switches use the vPC domain ID that you configure to automatically assign a unique vPC system MAC address. Each vPC domain has a unique MAC address that is used as a unique identifier for the specific vPC-related operations, although the switches use the vPC system MAC addresses only for link-scope operations, such as LACP. We recommend that you create each vPC domain within the contiguous network with a unique domain ID. You can also configure a specific MAC address for the vPC domain, rather than having the Cisco NX-OS software assign the address.

The vPC peer switches use the vPC domain ID that you configure to automatically assign a unique vPC system MAC address. The switches use the vPC system MAC addresses only for link-scope operations, such as LACP or BPDUs. You can also configure a specific MAC address for the vPC domain.

We recommend that you configure the same VPC domain ID on both peers and, the domain ID should be unique in the network. For example, if there are two different VPCs (one in access and one in aggregation) then each vPC should have a unique domain ID.

After you create a vPC domain, the Cisco NX-OS software automatically creates a system priority for the vPC domain. You can also manually configure a specific system priority for the vPC domain.



Note If you manually configure the system priority, you must ensure that you assign the same priority value on both vPC peer switches. If the vPC peer switches have different system priority values, the vPC will not come up.

Peer-Keepalive Link and Messages

The Cisco NX-OS software uses a peer-keepalive link between the vPC peers to transmit periodic, configurable keepalive messages. You must have Layer 3 connectivity between the peer switches to transmit these messages; the system cannot bring up the vPC peer link unless a peer-keepalive link is already up and running.

If one of the vPC peer switches fails, the vPC peer switch on the other side of the vPC peer link senses the failure when it does not receive any peer-keepalive messages. The default interval time for the vPC peer-keepalive message is 1 second. You can configure the interval between 400 milliseconds and 10 seconds. You can also configure a timeout value with a range of 3 to 20 seconds; the default timeout value is 5 seconds. The peer-keepalive status is checked only when the peer-link goes down.

The vPC peer-keepalive can be carried either in the management or default VRF on the Cisco Nexus device. When you configure the switches to use the management VRF, the source and destination for the keepalive messages are the mgmt 0 interface IP addresses. When you configure the switches to use the default VRF, an SVI must be created to act as the source and destination addresses for the vPC peer-keepalive messages.

Ensure that both the source and destination IP addresses used for the peer-keepalive messages are unique in your network and these IP addresses are reachable from the VRF associated with the vPC peer-keepalive link.



Note We recommend that you configure the vPC peer-keepalive link on the Cisco Nexus device to run in the management VRF using the mgmt 0 interfaces. If you configure the default VRF, ensure that the vPC peer link is not used to carry the vPC peer-keepalive messages.

Compatibility Parameters for vPC Peer Links

Many configuration and operational parameters must be identical on all interfaces in the vPC. After you enable the vPC feature and configure the peer link on both vPC peer switches, Cisco Fabric Services (CFS) messages provide a copy of the configuration on the local vPC peer switch configuration to the remote vPC peer switch. The system then determines whether any of the crucial configuration parameters differ on the two switches.

Enter the **show vpc consistency-parameters** command to display the configured values on all interfaces in the vPC. The displayed configurations are only those configurations that would limit the vPC peer link and vPC from coming up.

The compatibility check process for vPCs differs from the compatibility check for regular EtherChannels.

New Type 2 Consistency Check on the vPC Port-Channels

A new type 2 consistency check has been added to validate the switchport mac learn settings on the vPC port-channels. The CLI **show vpc consistency-check vPC <vpc no.>** has been enhanced to display the local and peer values of the switchport mac-learn configuration. Because it is a type 2 check, vPC is operationally up even if there is a mismatch between the local and the peer values, but the mismatch can be displayed from the CLI output.

```
switch# sh vpc consistency-parameters vpc 1112
```

Legend:

Type 1 : vPC will be suspended in case of mismatch

Name	Type	Local Value	Peer Value
Shut Lan	1	No	No
STP Port Type	1	Default	Default
STP Port Guard	1	None	None
STP MST Simulate PVST	1	Default	Default
nve configuration	1	nve	nve
lag-id	1	[(fa0, 0-23-4-ee-be-64, 8458, 0, 0), (8000, f4-4e-5-84-5e-3c, 457, 0, 0)]	[(fa0, 0, 0), (8000, 0, 0)]
mode	1	active	active
Speed	1	10 Gb/s	10 Gb/s
Duplex	1	full	full
Port Mode	1	trunk	trunk
Native Vlan	1	1	1
MTU	1	1500	1500
Admin port mode	1		
Switchport MAC Learn	2	Enable	Disable>
Newly added consistency parameter			
vPC card type	1	Empty	Empty
Allowed VLANs	-	311-400	311-400
Local suspended VLANs	-	-	-

Configuration Parameters That Must Be Identical

The configuration parameters in this section must be configured identically on both switches at either end of the vPC peer link.



Note You must ensure that all interfaces in the vPC have the identical operational and configuration parameters listed in this section.

Enter the **show vpc consistency-parameters** command to display the configured values on all interfaces in the vPC. The displayed configurations are only those configurations that would limit the vPC peer link and vPC from coming up.

The switch automatically checks for compatibility of these parameters on the vPC interfaces. The per-interface parameters must be consistent per interface, and the global parameters must be consistent globally.

- Port-channel mode: on, off, or active
- Link speed per channel
- Duplex mode per channel
- Trunk mode per channel:
 - Native VLAN
 - VLANs allowed on trunk
 - Tagging of native VLAN traffic
- Spanning Tree Protocol (STP) mode
- STP region configuration for Multiple Spanning Tree (MST)
- Enable or disable state per VLAN
- STP global settings:
 - Bridge Assurance setting
 - Port type setting—We recommend that you set all vPC interfaces as normal ports
 - Loop Guard settings
- STP interface settings:
 - Port type setting
 - Loop Guard
 - Root Guard

If any of these parameters are not enabled or defined on either switch, the vPC consistency check ignores those parameters.



Note To ensure that none of the vPC interfaces are in the suspend mode, enter the **show vpc brief** and **show vpc consistency-parameters** commands and check the syslog messages.

Configuration Parameters That Should Be Identical

When any of the following parameters are not configured identically on both vPC peer switches, a misconfiguration might cause undesirable behavior in the traffic flow:

- MAC aging timers
- Static MAC entries
- VLAN interface—Each switch on the end of the vPC peer link must have a VLAN interface configured for the same VLAN on both ends and they must be in the same administrative and operational mode.

Those VLANs configured on only one switch of the peer link do not pass traffic using the vPC or peer link. You must create all VLANs on both the primary and secondary vPC switches, or the VLAN will be suspended.

- Private VLAN configuration
- All ACL configurations and parameters
- Quality of service (QoS) configuration and parameters—Local parameters; global parameters must be identical
- STP interface settings:
 - BPDU Filter
 - BPDU Guard
 - Cost
 - Link type
 - Priority
 - VLANs (Rapid PVST+)

To ensure that all the configuration parameters are compatible, we recommend that you display the configurations for each vPC peer switch once you configure the vPC.

Per-VLAN Consistency Check

Type-1 consistency checks are performed on a per-VLAN basis when spanning tree is enabled or disabled on a VLAN. VLANs that do not pass the consistency check are brought down on both the primary and secondary switches while other VLANs are not affected.

vPC Auto-Recovery

When both vPC peer switches reload and only one switch reboots, auto-recovery allows that switch to assume the role of the primary switch and the vPC links will be allowed to come up after a predetermined period of time. The reload delay period in this scenario can range from 240 to 3600 seconds.

When vPCs are disabled on a secondary vPC switch due to a peer-link failure and then the primary vPC switch fails or is unable to forward traffic, the secondary switch reenables the vPCs. In this scenario, the vPC waits for three consecutive keepalive failures to recover the vPC links.

The vPC auto-recovery feature is disabled by default.

vPC Peer Links

A vPC peer link is the link that is used to synchronize the states between the vPC peer devices.



Note You must configure the peer-keepalive link before you configure the vPC peer link or the peer link will not come up.

vPC Peer Link Overview

You can have only two switches as vPC peers; each switch can serve as a vPC peer to only one other vPC peer. The vPC peer switches can also have non-vPC links to other switches.

To make a valid configuration, you configure an EtherChannel on each switch and then configure the vPC domain. You assign the EtherChannel on each switch as a peer link. For redundancy, we recommend that you should configure at least two dedicated ports into the EtherChannel; if one of the interfaces in the vPC peer link fails, the switch automatically falls back to use another interface in the peer link.



Note We recommend that you configure the EtherChannels in trunk mode.

Many operational parameters and configuration parameters must be the same in each switch connected by a vPC peer link. Because each switch is completely independent on the management plane, you must ensure that the switches are compatible on the critical parameters. vPC peer switches have separate control planes. After configuring the vPC peer link, you should display the configuration on each vPC peer switch to ensure that the configurations are compatible.



Note You must ensure that the two switches connected by the vPC peer link have certain identical operational and configuration parameters.

When you configure the vPC peer link, the vPC peer switches negotiate that one of the connected switches is the primary switch and the other connected switch is the secondary switch. By default, the Cisco NX-OS software uses the lowest MAC address to elect the primary switch. The software takes different actions on each switch—that is, the primary and secondary—only in certain failover conditions. If the primary switch fails, the secondary switch becomes the operational primary switch when the system recovers, and the previously primary switch is now the secondary switch.

You can also configure which of the vPC switches is the primary switch. If you want to configure the role priority again to make one vPC switch the primary switch, configure the role priority on both the primary and secondary vPC switches with the appropriate values, shut down the EtherChannel that is the vPC peer link on both switches by entering the **shutdown** command, and reenab the EtherChannel on both switches by entering the **no shutdown** command.

MAC addresses that are learned over vPC links are also synchronized between the peers.

Configuration information flows across the vPC peer links using the Cisco Fabric Services over Ethernet (CFSOE) protocol. All MAC addresses for those VLANs configured on both switches are synchronized between vPC peer switches. The software uses CFSOE for this synchronization.

If the vPC peer link fails, the software checks the status of the remote vPC peer switch using the peer-keepalive link, which is a link between vPC peer switches, to ensure that both switches are up. If the vPC peer switch is up, the secondary vPC switch disables all vPC ports on its switch. The data then forwards down the remaining active links of the EtherChannel.

The software learns of a vPC peer switch failure when the keepalive messages are not returned over the peer-keepalive link.

Use a separate link (vPC peer-keepalive link) to send configurable keepalive messages between the vPC peer switches. The keepalive messages on the vPC peer-keepalive link determines whether a failure is on the vPC peer link only or on the vPC peer switch. The keepalive messages are used only when all the links in the peer link fail.

vPC Number

Once you have created the vPC domain ID and the vPC peer link, you can create EtherChannels to attach the downstream switch to each vPC peer switch. That is, you create one single EtherChannel on the downstream switch with half of the ports to the primary vPC peer switch and the other half of the ports to the secondary peer switch.

On each vPC peer switch, you assign the same vPC number to the EtherChannel that connects to the downstream switch. You will experience minimal traffic disruption when you are creating vPCs. To simplify the configuration, you can assign the vPC ID number for each EtherChannel to be the same as the EtherChannel itself (that is, vPC ID 10 for EtherChannel 10).



Note The vPC number that you assign to the EtherChannel that connects to the downstream switch from the vPC peer switch must be identical on both vPC peer switches.

vPC Interactions with Other Features

vPC and LACP

The Link Aggregation Control Protocol (LACP) uses the system MAC address of the vPC domain to form the LACP Aggregation Group (LAG) ID for the vPC.

You can use LACP on all the vPC EtherChannels, including those channels from the downstream switch. We recommend that you configure LACP with active mode on the interfaces on each EtherChannel on the vPC peer switches. This configuration allows you to more easily detect compatibility between switches, unidirectional links, and multihop connections, and provides dynamic reaction to run-time changes and link failures.

The vPC peer link supports 16 EtherChannel interfaces.



Note When you manually configure the system priority, you must ensure that you assign the same priority value on both vPC peer switches. If the vPC peer switches have different system priority values, vPC does not come up.

vPC Peer Links and STP

When you first bring up the vPC functionality, STP reconverges. STP treats the vPC peer link as a special link and always includes the vPC peer link in the STP active topology.

We recommend that you set all the vPC peer link interfaces to the STP network port type so that Bridge Assurance is automatically enabled on all vPC peer links. We also recommend that you do not enable any of the STP enhancement features on vPC peer links.

You must configure a list of parameters to be identical on the vPC peer switches on both sides of the vPC peer link.

STP is distributed; that is, the protocol continues running on both vPC peer switches. However, the configuration on the vPC peer switch elected as the primary switch controls the STP process for the vPC interfaces on the secondary vPC peer switch.

The primary vPC switch synchronizes the STP state on the vPC secondary peer switch using Cisco Fabric Services over Ethernet (CFSOE).

The vPC manager performs a proposal/handshake agreement between the vPC peer switches that sets the primary and secondary switches and coordinates the two switches for STP. The primary vPC peer switch then controls the STP protocol for vPC interfaces on both the primary and secondary switches.

The Bridge Protocol Data Units (BPDUs) use the MAC address set for the vPC for the STP bridge ID in the designated bridge ID field. The vPC primary switch sends these BPDUs on the vPC interfaces.



Note Display the configuration on both sides of the vPC peer link to ensure that the settings are identical. Use the **show spanning-tree** command to display information about the vPC.

CFSOE

The Cisco Fabric Services over Ethernet (CFSOE) is a reliable state transport mechanism that you can use to synchronize the actions of the vPC peer devices. CFSOE carries messages and packets for many features linked with vPC, such as STP and IGMP. Information is carried in CFS/CFSOE protocol data units (PDUs).

When you enable the vPC feature, the device automatically enables CFSOE, and you do not have to configure anything. CFSOE distributions for vPCs do not need the capabilities to distribute over IP or the CFS regions. You do not need to configure anything for the CFSOE feature to work correctly on vPCs.

You can use the **show mac address-table** command to display the MAC addresses that CFSOE synchronizes for the vPC peer link.



Note Do not enter the **no cfs eth distribute** or the **no cfs distribute** command. CFSOE must be enabled for vPC functionality. If you do enter either of these commands when vPC is enabled, the system displays an error message.

When you enter the **show cfs application** command, the output displays "Physical-eth," which shows the applications that are using CFSOE.

Guidelines and Limitations for vPCs

vPCs have the following configuration guidelines and limitations:

- vPC is not supported between different types of Cisco Nexus 3000 Series switches.
- VPC peers should have same reserved VLANs for VXLAN. Different reserved VLANs on the peers may lead to undesired behavior with VXLAN.
- The output of the **sh vpc brief** CLI command displays two additional fields, Delay-restore status and Delay-restore SVI status.
- vPC is not qualified with IPv6.
- You must enable the vPC feature before you can configure vPC peer-link and vPC interfaces.
- You must configure the peer-keepalive link before the system can form the vPC peer link.

- The vPC peer-link needs to be formed using a minimum of two 10-Gigabit Ethernet interfaces.
- We recommend that you configure the same vPC domain ID on both peers and the domain ID should be unique in the network. For example, if there are two different vPCs (one in access and one in aggregation) then each vPC should have a unique domain ID.
- Only port channels can be in vPCs. A vPC can be configured on a normal port channel (switch-to-switch vPC topology) and on a port channel host interface (host interface vPC topology).
- You must configure both vPC peer switches; the configuration is not automatically synchronized between the vPC peer devices.
- Check that the necessary configuration parameters are compatible on both sides of the vPC peer link.
- You might experience minimal traffic disruption while configuring vPCs.
- You should configure all port channels in the vPC using LACP with the interfaces in active mode.
- You might experience traffic disruption when the first member of a vPC is brought up.
- OSPF over vPC and BFD with OSPF are supported on Cisco Nexus 3000 and 3100 Series switches.

SVI limitation: 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 using **no bfd echo** at the SVI configuration level.

- When a Layer 3 link is used for peer-keepalive instead of the mgmt interface, and the CPU queues are congested with control plane traffic, vPC peer-keepalive packets could be dropped. The CPU traffic includes routing protocol, ARP, Glean, and IPMC miss packets. When the peer-keepalive interface is a Layer 3 link instead of a mgmt interface, the vPC peer-keepalive packets are sent to the CPU on a low-priority queue.

If a Layer 3 link is used for vPC peer-keepalives, configure the following ACL to prioritize the vPC peer-keepalive:

```
ip access-list copp-system-acl-routingproto2
30 permit udp any any eq 3200
```

Here, 3200 is the default UDP port for keepalive packets. This ACL must match the configured UDP port in case the default port is changed.

- vPC is not supported on the Cisco Nexus 34180YC platform.

Verifying the vPC Configuration

Use the following commands to display vPC configuration information:

Command	Purpose
switch# show feature	Displays whether vPC is enabled or not.
switch# show port-channel capacity	Displays how many EtherChannels are configured and how many are still available on the switch.
switch# show running-config vpc	Displays running configuration information for vPCs.

Command	Purpose
switch# show vpc brief	Displays brief information on the vPCs.
switch# show vpc consistency-parameters	Displays the status of those parameters that must be consistent across all vPC interfaces.
switch# show vpc peer-keepalive	Displays information on the peer-keepalive messages.
switch# show vpc role	Displays the peer status, the role of the local switch, the vPC system MAC address and system priority, and the MAC address and priority for the local vPC switch.
switch# show vpc statistics	Displays statistics on the vPCs. Note This command displays the vPC statistics only for the vPC peer device that you are working on.

For information about the switch output, see the Command Reference for your Cisco Nexus Series switch.

Viewing the Graceful Type-1 Check Status

This example shows how to display the current status of the graceful Type-1 consistency check:

```
switch# show vpc brief
Legend:
          (*) - local vPC is down, forwarding via vPC peer-link

vPC domain id           : 10
Peer status             : peer adjacency formed ok
vPC keep-alive status   : peer is alive
Configuration consistency status: success
Per-vlan consistency status : success
Type-2 consistency status : success
vPC role                : secondary
Number of vPCs configured : 34
Peer Gateway            : Disabled
Dual-active excluded VLANs : -
Graceful Consistency Check : Enabled
Auto-recovery status    : Disabled
Delay-restore status     : Timer is off.(timeout = 30s)
Delay-restore SVI status : Timer is off.(timeout = 10s)

vPC Peer-link status
-----
id   Port   Status Active vlans
--   -
1    Po1    up      1
```

Viewing a Global Type-1 Inconsistency

When a global Type-1 inconsistency occurs, the vPCs on the secondary switch are brought down. The following example shows this type of inconsistency when there is a spanning-tree mode mismatch.

The example shows how to display the status of the suspended vPC VLANs on the secondary switch:

```
switch(config)# show vpc
Legend:
          (*) - local vPC is down, forwarding via vPC peer-link

vPC domain id          : 10
Peer status            : peer adjacency formed ok
vPC keep-alive status  : peer is alive
Configuration consistency status: failed
Per-vlan consistency status : success
Configuration consistency reason: vPC type-1 configuration incompatible - STP
                                Mode inconsistent

Type-2 consistency status : success
vPC role                 : secondary
Number of vPCs configured : 2
Peer Gateway             : Disabled
Dual-active excluded VLANs : -
Graceful Consistency Check : Enabled
```

vPC Peer-link status

```
-----
id  Port  Status Active vlans
--  ---  -----
1   Po1   up    1-10
```

vPC status

```
-----
id  Port      Status Consistency Reason          Active vlans
-----
20  Po20      down*  failed    Global compat check failed -
30  Po30      down*  failed    Global compat check failed -
```

The example shows how to display the inconsistent status (the VLANs on the primary vPC are not suspended) on the primary switch:

```
switch(config)# show vpc
Legend:
          (*) - local vPC is down, forwarding via vPC peer-link

vPC domain id          : 10
Peer status            : peer adjacency formed ok
vPC keep-alive status  : peer is alive
Configuration consistency status: failed
Per-vlan consistency status : success
Configuration consistency reason: vPC type-1 configuration incompatible - STP Mo
de inconsistent

Type-2 consistency status : success
vPC role                 : primary
Number of vPCs configured : 2
Peer Gateway             : Disabled
Dual-active excluded VLANs : -
Graceful Consistency Check : Enabled
```

vPC Peer-link status

```
-----
id  Port  Status Active vlans
--  ---  -----
1   Po1   up    1-10
```

vPC status

```
-----
id  Port      Status Consistency Reason          Active vlans
-----
20  Po20      up    failed    Global compat check failed 1-10
```

```
30      Po30      up      failed      Global compat check failed 1-10
```

Viewing an Interface-Specific Type-1 Inconsistency

When an interface-specific Type-1 inconsistency occurs, the vPC port on the secondary switch is brought down while the primary switch vPC ports remain up. The following example shows this type of inconsistency when there is a switchport mode mismatch.

This example shows how to display the status of the suspended vPC VLAN on the secondary switch:

```
switch(config-if)# show vpc brief
Legend:
      (*) - local vPC is down, forwarding via vPC peer-link

vPC domain id          : 10
Peer status            : peer adjacency formed ok
vPC keep-alive status  : peer is alive
Configuration consistency status: success
Per-vlan consistency status : success
Type-2 consistency status : success
vPC role               : secondary
Number of vPCs configured : 2
Peer Gateway           : Disabled
Dual-active excluded VLANs : -
Graceful Consistency Check : Enabled
Auto-recovery status   : Disabled
Delay-restore status   : Timer is off.(timeout = 30s)
Delay-restore SVI status : Timer is off.(timeout = 10s)

vPC Peer-link status
-----
id  Port  Status Active vlans
--  ---  -----
1   Po1   up     1

vPC status
-----
id  Port  Status Consistency Reason Active vlans
-----
20  Po20  up     success success 1
30  Po30  down*  failed  Compatibility check failed -
                        for port mode
```

This example shows how to display the inconsistent status (the VLANs on the primary vPC are not suspended) on the primary switch:

```
switch(config-if)# show vpc brief
Legend:
      (*) - local vPC is down, forwarding via vPC peer-link

vPC domain id          : 10
Peer status            : peer adjacency formed ok
vPC keep-alive status  : peer is alive
Configuration consistency status: success
Per-vlan consistency status : success
Type-2 consistency status : success
vPC role               : primary
Number of vPCs configured : 2
Peer Gateway           : Disabled
Dual-active excluded VLANs : -
```

```

Graceful Consistency Check      : Enabled
Auto-recovery status           : Disabled
Delay-restore status           : Timer is off.(timeout = 30s)
Delay-restore SVI status       : Timer is off.(timeout = 10s)

vPC Peer-link status
-----
id   Port   Status Active vlans
--   -
1    Po1    up     1

vPC status
-----
id   Port      Status Consistency Reason              Active vlans
-----
20   Po20       up     success      success                          1
30   Po30       up     failed       Compatibility check failed 1
                                         for port mode

```

Viewing a Per-VLAN Consistency Status

To view the per-VLAN consistency or inconsistency status, enter the **show vpc consistency-parameters vlans** command.

Example

This example shows how to display the consistent status of the VLANs on the primary and the secondary switches.

```

switch(config-if)# show vpc brief
Legend:
          (*) - local vPC is down, forwarding via vPC peer-link

vPC domain id                : 10
Peer status                   : peer adjacency formed ok
vPC keep-alive status         : peer is alive
Configuration consistency status: success
Per-vlan consistency status   : success
Type-2 consistency status     : success
vPC role                      : secondary
Number of vPCs configured     : 2
Peer Gateway                  : Disabled
Dual-active excluded VLANs    : -
Graceful Consistency Check    : Enabled
Auto-recovery status         : Disabled
Delay-restore status         : Timer is off.(timeout = 30s)
Delay-restore SVI status     : Timer is off.(timeout = 10s)

vPC Peer-link status
-----
id   Port   Status Active vlans
--   -
1    Po1    up     1-10

vPC status
-----
id   Port      Status Consistency Reason              Active vlans
-----
20   Po20       up     success      success                          1-10

```

```
30    Po30    up    success    success    1-10
```

Entering **no spanning-tree vlan 5** command triggers the inconsistency on the primary and secondary VLANs:

```
switch(config)# no spanning-tree vlan 5
```

This example shows how to display the per-VLAN consistency status as Failed on the secondary switch:

```
switch(config)# show vpc brief
```

Legend:

(*) - local vPC is down, forwarding via vPC peer-link

```
vPC domain id          : 10
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           : Disabled
Dual-active excluded VLANs : -
Graceful Consistency Check : Enabled
Auto-recovery status   : Disabled
Delay-restore status    : Timer is off.(timeout = 30s)
Delay-restore SVI status : Timer is off.(timeout = 10s)
```

vPC Peer-link status

```
-----
id   Port   Status Active vlans
--   ---   -
1    Po1    up    1-4,6-10
-----
```

vPC status

```
-----
id   Port   Status Consistency Reason           Active vlans
-----
20   Po20    up    success    success    1-4,6-10
30   Po30    up    success    success    1-4,6-10
-----
```

This example shows how to display the per-VLAN consistency status as Failed on the primary switch:

```
switch(config)# show vpc brief
```

Legend:

(*) - local vPC is down, forwarding via vPC peer-link

```
vPC domain id          : 10
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               : primary
Number of vPCs configured : 2
Peer Gateway           : Disabled
Dual-active excluded VLANs : -
Graceful Consistency Check : Enabled
Auto-recovery status   : Disabled
Delay-restore status    : Timer is off.(timeout = 30s)
Delay-restore SVI status : Timer is off.(timeout = 10s)
```

```
vPC Peer-link status
```

```
-----
id   Port   Status Active vlans
--   -
1    Po1    up     1-4,6-10
-----
```

```
vPC status
```

```
-----
id   Port   Status Consistency Reason           Active vlans
--   -
20   Po20   up     success  success           1-4,6-10
30   Po30   up     success  success           1-4,6-10
-----
```

This example shows the inconsistency as STP Disabled:

```
switch(config)# show vpc consistency-parameters vlans
```

```
-----
Name                               Type Reason Code           Pass Vlans
-----
STP Mode                            1    success                0-4095
STP Disabled                       1    vPC type-1           0-4,6-4095
                                     configuration
                                     incompatible - STP is
                                     enabled or disabled on
                                     some or all vlans
STP MST Region Name                 1    success                0-4095
STP MST Region Revision              1    success                0-4095
STP MST Region Instance to          1    success                0-4095
  VLAN Mapping
STP Loopguard                       1    success                0-4095
STP Bridge Assurance                 1    success                0-4095
STP Port Type, Edge                  1    success                0-4095
BPDUFilter, Edge BPDUGuard
STP MST Simulate PVST                1    success                0-4095
Pass Vlans                           -
                                     0-4,6-4095
-----
```

vPC Default Settings

The following table lists the default settings for vPC parameters.

Table 9: 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

Enabling vPCs

You must enable the vPC feature before you can configure and use vPCs.

SUMMARY STEPS

1. switch# **configure terminal**
2. switch(config)# **feature vpc**
3. (Optional) switch# **show feature**
4. (Optional) switch# **copy running-config startup-config**

DETAILED STEPS

	Command or Action	Purpose
Step 1	switch# configure terminal	Enters global configuration mode.
Step 2	switch(config)# feature vpc	Enables vPCs on the switch.
Step 3	(Optional) switch# show feature	Displays which features are enabled on the switch.
Step 4	(Optional) switch# copy running-config startup-config	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
```

Disabling vPCs

You can disable the vPC feature.



Note When you disable the vPC feature, the Cisco Nexus device clears all the vPC configurations.

SUMMARY STEPS

1. switch# **configure terminal**
2. switch(config)# **no feature vpc**
3. (Optional) switch# **show feature**
4. (Optional) switch# **copy running-config startup-config**

DETAILED STEPS

	Command or Action	Purpose
Step 1	switch# configure terminal	Enters global configuration mode.
Step 2	switch(config)# no feature vpc	Disables vPCs on the switch.
Step 3	(Optional) switch# show feature	Displays which features are enabled on the switch.
Step 4	(Optional) switch# copy running-config startup-config	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
```

Creating a vPC Domain

You must create identical vPC domain IDs on both the vPC peer devices. This domain ID is used to automatically form the vPC system MAC address.

Before you begin

Ensure that you have enabled the vPC feature.

You must configure both switches on either side of the vPC peer link.

SUMMARY STEPS

1. switch# **configure terminal**
2. switch(config)# **vpc domain** *domain-id*
3. (Optional) switch# **show vpc brief**
4. (Optional) switch# **copy running-config startup-config**

DETAILED STEPS

	Command or Action	Purpose
Step 1	switch# configure terminal	Enters global configuration mode.
Step 2	switch(config)# vpc domain <i>domain-id</i>	Creates a vPC domain on the switch, and enters the vpc-domain configuration mode. There is no default <i>domain-id</i> ; the range is from 1 to 1000. Note You can also use the vpc domain command to enter the vpc-domain configuration mode for an existing vPC domain.

	Command or Action	Purpose
Step 3	(Optional) switch# show vpc brief	Displays brief information about each vPC domain.
Step 4	(Optional) switch# copy running-config startup-config	Copies the running configuration to the startup configuration.

Example

This example shows how to create a vPC domain:

```
switch# configure terminal
switch(config)# vpc domain 5
```

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.

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.

Ensure that both the source and destination IP addresses used for the peer-keepalive message are unique in your network and these IP addresses are reachable from the Virtual Routing and Forwarding (VRF) instance associated with the vPC peer-keepalive link.



Note We recommend that you configure a separate VRF instance and put a Layer 3 port from each vPC peer switch into that VRF instance for the vPC peer-keepalive link. Do not use the peer link itself to send vPC peer-keepalive messages.

Before you begin

Ensure that you have enabled the vPC feature.

You must configure the vPC peer-keepalive link before the system can form the vPC peer link.

You must configure both switches on either side of the vPC peer link.

SUMMARY STEPS

1. switch# **configure terminal**
2. switch(config)# **vpc domain** *domain-id*
3. switch(config-vpc-domain)# **peer-keepalive destination** *ipaddress* [**hold-timeout** *secs* | **interval** *msecs* {**timeout** *secs*} | **precedence** {*prec-value*} | **network** | **internet** | **critical** | **flash-override** | **flash** | **immediate** | **priority** | **routine**} | **tos** {*tos-value*} | **max-reliability** | **max-throughput** | **min-delay** | **min-monetary-cost** | **normal**} | **tos-byte** *tos-byte-value*} | **source** *ipaddress* | **vrf** {*name*} | **management vpc-keepalive**}]
4. (Optional) switch(config-vpc-domain)# **vpc peer-keepalive destination** *ipaddress* **source** *ipaddress*
5. (Optional) switch# **show vpc peer-keepalive**

6. (Optional) switch# copy running-config startup-config

DETAILED STEPS

	Command or Action	Purpose
Step 1	switch# configure terminal	Enters global configuration mode.
Step 2	switch(config)# vpc domain <i>domain-id</i>	Creates a vPC domain on the switch if it does not already exist, and enters the vpc-domain configuration mode.
Step 3	switch(config-vpc-domain)# peer-keepalive destination <i>ipaddress</i> [hold-timeout <i>secs</i> interval <i>msecs</i> { timeout <i>secs</i> } precedence { <i>prec-value</i> network internet critical flash-override flash immediate priority routine } tos { <i>tos-value</i> max-reliability max-throughput min-delay min-monetary-cost normal } tos-byte <i>tos-byte-value</i> } source <i>ipaddress</i> vrf { <i>name</i> management vpc-keepalive }]	Configures the IPv4 address for the remote end of the vPC peer-keepalive link. Note The system does not form the vPC peer link until you configure a vPC peer-keepalive link. The management ports and VRF are the defaults.
Step 4	(Optional) switch(config-vpc-domain)# vpc peer-keepalive destination <i>ipaddress source ipaddress</i>	Configures a separate VRF instance and puts a Layer 3 port from each vPC peer device into that VRF for the vPC peer-keepalive link.
Step 5	(Optional) switch# show vpc peer-keepalive	Displays information about the configuration for the keepalive messages.
Step 6	(Optional) switch# copy running-config startup-config	Copies the running configuration to the startup configuration.

Example

This example shows how to configure the destination IP address for the vPC-peer-keepalive link:

```
switch# configure terminal
switch(config)# vpc domain 5
switch(config-vpc-domain)# peer-keepalive destination 10.10.10.42
```

This example shows how to set up the peer keepalive link connection between the primary and secondary vPC device:

```
switch(config)# vpc domain 100
switch(config-vpc-domain)# peer-keepalive destination 192.168.2.2 source 192.168.2.1
Note:-----: Management VRF will be used as the default VRF ::-----
switch(config-vpc-domain)#
```

This example shows how to create a separate VRF named vpc_keepalive for the vPC keepalive link and how to verify the new VRF:

```
vrf context vpc_keepalive
interface Ethernet1/31
  switchport access vlan 123
interface Vlan123
  vrf member vpc_keepalive
```

```

ip address 123.1.1.2/30
no shutdown
vpc domain 1
peer-keepalive destination 123.1.1.1 source 123.1.1.2 vrf
vpc_keepalive

L3-NEXUS-2# show vpc peer-keepalive

vPC keep-alive status           : peer is alive
--Peer is alive for             : (154477) seconds, (908) msec
--Send status                   : Success
--Last send at                  : 2011.01.14 19:02:50 100 ms
--Sent on interface             : Vlan123
--Receive status                : Success
--Last receive at               : 2011.01.14 19:02:50 103 ms
--Received on interface         : Vlan123
--Last update from peer        : (0) seconds, (524) msec

vPC Keep-alive parameters
--Destination                   : 123.1.1.1
--Keepalive interval            : 1000 msec
--Keepalive timeout             : 5 seconds
--Keepalive hold timeout        : 3 seconds
--Keepalive vrf                 : vpc_keepalive
--Keepalive udp port            : 3200
--Keepalive tos                 : 192

```

The services provided by the switch , such as ping, ssh, telnet, radius, are VRF aware. The VRF name need to be configured or specified in order for the correct routing table to be used.

```

L3-NEXUS-2# ping 123.1.1.1 vrf vpc_keepalive
PING 123.1.1.1 (123.1.1.1): 56 data bytes
64 bytes from 123.1.1.1: icmp_seq=0 ttl=254 time=3.234 ms
64 bytes from 123.1.1.1: icmp_seq=1 ttl=254 time=4.931 ms
64 bytes from 123.1.1.1: icmp_seq=2 ttl=254 time=4.965 ms
64 bytes from 123.1.1.1: icmp_seq=3 ttl=254 time=4.971 ms
64 bytes from 123.1.1.1: icmp_seq=4 ttl=254 time=4.915 ms

--- 123.1.1.1 ping statistics ---
5 packets transmitted, 5 packets received, 0.00% packet loss
round-trip min/avg/max = 3.234/4.603/4.971 ms

```

Creating a vPC Peer Link

You can create a vPC peer link by designating the EtherChannel that you want on each switch as the peer link for the specified vPC domain. We recommend that you configure the EtherChannels that you are designating as the vPC peer link in trunk mode and that you use two ports on separate modules on each vPC peer switch for redundancy.

Before you begin

Ensure that you have enabled the vPC feature.

You must configure both switches on either side of the vPC peer link

SUMMARY STEPS

1. switch# **configure terminal**
2. switch(config)# **interface port-channel** *channel-number*

3. switch(config-if)# **vpc peer-link**
4. (Optional) switch# **show vpc brief**
5. (Optional) switch# **copy running-config startup-config**

DETAILED STEPS

	Command or Action	Purpose
Step 1	switch# configure terminal	Enters global configuration mode.
Step 2	switch(config)# interface port-channel <i>channel-number</i>	Selects the EtherChannel that you want to use as the vPC peer link for this switch, and enters the interface configuration mode.
Step 3	switch(config-if)# vpc peer-link	Configures the selected EtherChannel as the vPC peer link, and enters the vpc-domain configuration mode.
Step 4	(Optional) switch# show vpc brief	Displays information about each vPC, including information about the vPC peer link.
Step 5	(Optional) switch# copy running-config startup-config	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)# vpc peer-link
```

Checking the Configuration Compatibility

After you have configured the vPC peer link on both vPC peer switches, check that the configurations are consistent on all vPC interfaces.

The following QoS parameters support Type 2 consistency checks

- Network QoS—MTU and Pause
- Input Queuing —Bandwidth and Absolute Priority
- Output Queuing—Bandwidth and Absolute Priority

In the case of a Type 2 mismatch, the vPC is not suspended. Type 1 mismatches suspend the vPC.

SUMMARY STEPS

1. switch# **show vpc consistency-parameters**{global|interface port-channel*channel-number*}

DETAILED STEPS

	Command or Action	Purpose
Step 1	<code>switch# show vpc consistency-parameters {global interface port-channel channel-number}</code>	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# show vpc consistency-parameters global
Legend:
      Type 1 : vPC will be suspended in case of mismatch
Name                               Type Local Value                               Peer Value
-----
QoS                                 2      ([], [], [], [], [], ([], [], [], [], [],
      []
Network QoS (MTU)                   2      (1538, 0, 0, 0, 0, 0) (1538, 0, 0, 0, 0, 0)
Network Qos (Pause)                 2      (F, F, F, F, F, F)   (1538, 0, 0, 0, 0, 0)
Input Queuing (Bandwidth)           2      (100, 0, 0, 0, 0, 0) (100, 0, 0, 0, 0, 0)
Input Queuing (Absolute             2      (F, F, F, F, F, F)   (100, 0, 0, 0, 0, 0)
Priority)
Output Queuing (Bandwidth)           2      (100, 0, 0, 0, 0, 0) (100, 0, 0, 0, 0, 0)
Output Queuing (Absolute           2      (F, F, F, F, F, F)   (100, 0, 0, 0, 0, 0)
Priority)
STP Mode                             1      Rapid-PVST          Rapid-PVST
STP Disabled                         1      None                None
STP MST Region Name                  1      ""                  ""
STP MST Region Revision              1      0                   0
STP MST Region Instance to          1
VLAN Mapping
STP Loopguard                       1      Disabled            Disabled
STP Bridge Assurance                 1      Enabled              Enabled
STP Port Type, Edge                  1      Normal, Disabled,   Normal, Disabled,
BPDUFilter, Edge BPDUGuard          Disabled            Disabled
STP MST Simulate PVST                1      Enabled              Enabled
Allowed VLANs                        -      1,624                1
Local suspended VLANs                -      624                  -
switch#
```

Enabling vPC Auto-Recovery

SUMMARY STEPS

1. switch# **configure terminal**
2. switch(config)# **vpc domain domain-id**
3. switch(config-vpc-domain)# **auto-recovery reload-delay delay**

DETAILED STEPS

	Command or Action	Purpose
Step 1	switch# configure terminal	Enters global configuration mode.
Step 2	switch(config)# vpc domain <i>domain-id</i>	Enters vpc-domain configuration mode for an existing vPC domain.
Step 3	switch(config-vpc-domain)# auto-recovery reload-delay <i>delay</i>	Enables the auto-recovery feature and sets the reload delay period. The default is disabled.

Example

This example shows how to enable the auto-recovery feature in vPC domain 10 and set the delay period for 240 seconds:

```
switch(config)# vpc domain 10
switch(config-vpc-domain)# auto-recovery reload-delay 240
```

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

This example shows how to view the status of the auto-recovery feature in vPC domain 10:

```
switch(config-vpc-domain)# show running-config vpc
!Command: show running-config vpc
!Time: Tue Dec 7 02:38:44 2010
```

```
version 5.0(3)U2(1)
feature vpc
vpc domain 10
 peer-keepalive destination 10.193.51.170
 auto-recovery
```

Configuring the Restore Time Delay

You can configure a restore timer that delays the vPC from coming back up until after the peer adjacency forms and the VLAN interfaces are back up. This feature avoids packet drops if the routing tables fail to converge before the vPC is once again passing traffic.

Before you begin

Ensure that you have enabled the vPC feature.

You must configure both switches on either side of the vPC peer link with the following procedures.

SUMMARY STEPS

1. switch# **configure terminal**
2. switch(config)# **vpc domain** *domain-id*
3. switch(config-vpc-domain)# **delay restore** *time*

4. (Optional) switch# copy running-config startup-config

DETAILED STEPS

	Command or Action	Purpose
Step 1	switch# configure terminal	Enters global configuration mode.
Step 2	switch(config)# vpc domain <i>domain-id</i>	Creates a vPC domain on the switch if it does not already exist, and enters vpc-domain configuration mode.
Step 3	switch(config-vpc-domain)# delay restore <i>time</i>	Configures the time delay before the vPC is restored. The restore time is the number of seconds to delay bringing up the restored vPC peer device. The range is from 1 to 3600. The default is 30 seconds.
Step 4	(Optional) switch# copy running-config startup-config	Copies the running configuration to the startup configuration.

Example

This example shows how to configure the delay reload time for a vPC link:

```
switch(config)# vpc domain 1
switch(config-vpc-domain)# delay restore 10
switch(config-vpc-domain)#
```

Configuring Delay Restore on an Orphan Port

You can configure **delay restore orphan-port** command on Cisco Nexus 3000 Series switches to configure a restore timer that delays the bringing up of restored device's orphan port.

SUMMARY STEPS

1. **configure terminal**
2. **switch(config) # vpc domain <domain>**
3. **switch(config) # peer-switch**
4. **switch(config) # show vpc peer-keepalive**
5. **switch(config) # delay restore { time }**
6. **switch(config) # peer-gateway**
7. **switch(config) # delay restore orphan-port**

DETAILED STEPS

	Command or Action	Purpose
Step 1	configure terminal	Enters global configuration mode.
Step 2	switch(config) # vpc domain <domain>	Configure the VPC domain number.
Step 3	switch(config) # peer-switch	Define the peer switch.

	Command or Action	Purpose
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 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 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, not on port channel member ports.

Before you begin

Ensure that you have enabled the vPC feature.

Ensure that you are in the correct VDC (or use the `switchto vdc` command).

SUMMARY STEPS

1. `configure terminal`
2. `show vpc orphan-ports`
3. `interface type slot/port`
4. `vpc orphan-ports suspend`
5. `exit`
6. `copy running-config startup-config`

DETAILED STEPS

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

	Command or Action	Purpose
Step 2	show vpc orphan-ports Example: <pre>switch(config)# show vpc orphan-ports switch(config-vpc-domain)#</pre>	(Optional) Displays a list of the orphan ports.
Step 3	interface type slot/port Example: <pre>switch(config)# interface ethernet 3/1 switch(config-if)#</pre>	Specifies an interface to configure, and enters interface configuration mode.
Step 4	vpc orphan-ports suspend Example: <pre>switch(config-if)# vpc orphan-ports suspend</pre>	Configures the selected interface as a vPC orphan port to be suspended by the secondary peer in case of vPC failure.
Step 5	exit Example: <pre>switch(config-if)# exit</pre>	Exits the interface configuration mode.
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 configure an interface as a vPC orphan port to be suspended by the secondary peer in case of vPC failure:

```
switch# configure terminal
switch(config)# interface ethernet 3/1
switch(config-if)# vpc orphan-ports suspend
switch(config-if)#
```

Excluding VLAN Interfaces from Shutting Down a vPC Peer Link Fails

When a vPC peer-link is lost, the vPC secondary switch suspends its vPC member ports and its switch virtual interface (SVI) interfaces. All Layer 3 forwarding is disabled for all VLANs on the vPC secondary switch. You can exclude specific SVI interfaces so that they are not suspended.

Before you begin

Ensure that the VLAN interfaces have been configured.

SUMMARY STEPS

1. switch# **configure terminal**
2. switch(config)# **vpc domain** *domain-id*
3. switch(config-vpc-domain)# **dual-active exclude interface-vlan** *range*

DETAILED STEPS

	Command or Action	Purpose
Step 1	switch# configure terminal	Enters global configuration mode.
Step 2	switch(config)# vpc domain <i>domain-id</i>	Creates a vPC domain on the switch if it does not already exist, and enters vpc-domain configuration mode.
Step 3	switch(config-vpc-domain)# dual-active exclude interface-vlan <i>range</i>	Specifies the VLAN interfaces that should remain up when a vPC peer-link is lost. <i>range</i> —Range of VLAN interfaces that you want to exclude from shutting down. The range is from 1 to 4094.

Example

This example shows how to keep the interfaces on VLAN 10 up on the vPC peer switch if a peer link fails:

```
switch# configure terminal
switch(config)# vpc domain 5
switch(config-vpc-domain)# dual-active exclude interface-vlan 10
switch(config-vpc-domain)#
```

Configuring the VRF Name

The switch services, such as ping, ssh, telnet, radius, are VRF aware. You must configure the VRF name in order for the correct routing table to be used.

You can specify the VRF name.

SUMMARY STEPS

1. switch# **ping ipaddress vrf** *vrf-name*

DETAILED STEPS

	Command or Action	Purpose
Step 1	switch# ping ipaddress vrf <i>vrf-name</i>	Specifies the virtual routing and forwarding (VRF) name to use. The VRF name is case sensitive and can be a maximum of 32 characters..

Example

This example shows how to specify the VRF named `vpc_keepalive`:

```

switch# ping 123.1.1.1 vrf vpc_keepalive
PING 123.1.1.1 (123.1.1.1): 56 data bytes
64 bytes from 123.1.1.1: icmp_seq=0 ttl=254 time=3.234 ms
64 bytes from 123.1.1.1: icmp_seq=1 ttl=254 time=4.931 ms
64 bytes from 123.1.1.1: icmp_seq=2 ttl=254 time=4.965 ms
64 bytes from 123.1.1.1: icmp_seq=3 ttl=254 time=4.971 ms
64 bytes from 123.1.1.1: icmp_seq=4 ttl=254 time=4.915 ms

--- 123.1.1.1 ping statistics ---
5 packets transmitted, 5 packets received, 0.00% packet loss
round-trip min/avg/max = 3.234/4.603/4.971 ms

```

Moving Other Port Channels into a vPC

Before you begin

Ensure that you have enabled the vPC feature.

You must configure both switches on either side of the vPC peer link with the following procedure.

SUMMARY STEPS

1. switch# **configure terminal**
2. switch(config)# **interface port-channel** *channel-number*
3. switch(config-if)# **vpc** *number*
4. (Optional) switch# **show vpc brief**
5. (Optional) switch# **copy running-config startup-config**

DETAILED STEPS

	Command or Action	Purpose
Step 1	switch# configure terminal	Enters global configuration mode.
Step 2	switch(config)# interface port-channel <i>channel-number</i>	Selects the port channel that you want to put into the vPC to connect to the downstream switch, and enters interface configuration mode. Note A vPC can be configured on a normal port channel (physical vPC topology) and on a port channel host interface (host interface vPC topology)
Step 3	switch(config-if)# vpc <i>number</i>	Configures the selected port channel into the vPC to connect to the downstream switch. The range is from 1 to 4096. The vPC <i>number</i> that you assign to the port channel that connects to the downstream switch from the vPC peer switch must be identical on both vPC peer switches.

	Command or Action	Purpose
Step 4	(Optional) switch# show vpc brief	Displays information about each vPC.
Step 5	(Optional) switch# copy running-config startup-config	Copies the running configuration to the startup configuration.

Example

This example shows how to configure a port channel that will connect to the downstream device:

```
switch# configure terminal
switch(config)# interface port-channel 20
switch(config-if)# vpc 5
```

Manually Configuring a vPC Domain MAC Address



Note Configuring the system address is an optional configuration step.

Before you begin

Ensure that you have enabled the vPC feature.

You must configure both switches on either side of the vPC peer link.

SUMMARY STEPS

1. switch# **configure terminal**
2. switch(config)# **vpc domain** *domain-id*
3. switch(config-vpc-domain)# **system-mac** *mac-address*
4. (Optional) switch# **show vpc role**
5. (Optional) switch# **copy running-config startup-config**

DETAILED STEPS

	Command or Action	Purpose
Step 1	switch# configure terminal	Enters global configuration mode.
Step 2	switch(config)# vpc domain <i>domain-id</i>	Selects an existing vPC domain on the switch, or creates a new vPC domain, and enters the vpc-domain configuration mode. There is no default <i>domain-id</i> ; the range is from 1 to 1000.
Step 3	switch(config-vpc-domain)# system-mac <i>mac-address</i>	Enters the MAC address that you want for the specified vPC domain in the following format: <i>aaaa.bbbb.cccc</i> .

	Command or Action	Purpose
Step 4	(Optional) switch# show vpc role	Displays the vPC system MAC address.
Step 5	(Optional) switch# copy running-config startup-config	Copies the running configuration to the startup configuration.

Example

This example shows how to configure a vPC domain MAC address:

```
switch# configure terminal
switch(config)# vpc domain 5
switch(config-if)# system-mac 23fb.4ab5.4c4e
```

Manually Configuring the System Priority

When you create a vPC domain, the system automatically creates a vPC system priority. However, you can also manually configure a system priority for the vPC domain.

Before you begin

Ensure that you have enabled the vPC feature.

You must configure both switches on either side of the vPC peer link.

SUMMARY STEPS

1. switch# **configure terminal**
2. switch(config)# **vpc domain** *domain-id*
3. switch(config-vpc-domain)# **system-priority** *priority*
4. (Optional) switch# **show vpc brief**
5. (Optional) switch# **copy running-config startup-config**

DETAILED STEPS

	Command or Action	Purpose
Step 1	switch# configure terminal	Enters global configuration mode.
Step 2	switch(config)# vpc domain <i>domain-id</i>	Selects an existing vPC domain on the switch, or creates a new vPC domain, and enters the vpc-domain configuration mode. There is no default <i>domain-id</i> ; the range is from 1 to 1000.
Step 3	switch(config-vpc-domain)# system-priority <i>priority</i>	Enters the system priority that you want for the specified vPC domain. The range of values is from 1 to 65535. The default value is 32667.

	Command or Action	Purpose
Step 4	(Optional) switch# show vpc brief	Displays information about each vPC, including information about the vPC peer link.
Step 5	(Optional) switch# copy running-config startup-config	Copies the running configuration to the startup configuration.

Example

This example shows how to configure a vPC peer link:

```
switch# configure terminal
switch(config)# vpc domain 5
switch(config-if)# system-priority 4000
```

Manually Configuring a vPC Peer Switch Role

By default, the Cisco NX-OS software elects a primary and secondary vPC peer switch after you configure the vPC domain and both sides of the vPC peer link. However, you may want to elect a specific vPC peer switch as the primary switch for the vPC. Then, you would manually configure the role value for the vPC peer switch that you want as the primary switch to be lower than the other vPC peer switch.

vPC does not support role preemption. If the primary vPC peer switch fails, the secondary vPC peer switch takes over to become operationally the vPC primary switch. However, the original operational roles are not restored when the formerly primary vPC comes up again.

Before you begin

Ensure that you have enabled the vPC feature.

You must configure both switches on either side of the vPC peer link.

SUMMARY STEPS

1. switch# **configure terminal**
2. switch(config)# **vpc domain** *domain-id*
3. switch(config-vpc-domain)# **role priority** *priority*
4. (Optional) switch# **show vpc brief**
5. (Optional) switch# **copy running-config startup-config**

DETAILED STEPS

	Command or Action	Purpose
Step 1	switch# configure terminal	Enters global configuration mode.
Step 2	switch(config)# vpc domain <i>domain-id</i>	Selects an existing vPC domain on the switch, or creates a new vPC domain, and enters the vpc-domain configuration

	Command or Action	Purpose
		mode. There is no default <i>domain-id</i> ; the range is from 1 to 1000.
Step 3	switch(config-vpc-domain)# role priority <i>priority</i>	Enters the role priority that you want for the vPC system priority. The range of values is from 1 to 65535. The default value is 32667.
Step 4	(Optional) switch# show vpc brief	Displays information about each vPC, including information about the vPC peer link.
Step 5	(Optional) switch# copy running-config startup-config	Copies the running configuration to the startup configuration.

Example

This example shows how to configure a vPC peer link:

```
switch# configure terminal
switch(config)# vpc domain 5
switch(config-if)# role priority 4000
```

Configuring Layer 3 over vPC

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.

SUMMARY STEPS

1. switch# **configure terminal**
2. switch(config)# **vpc domain** *domain-id*
3. switch(config-vpc-domain)#**layer3 peer-router**
4. switch(config-vpc-domain)# **exit**
5. (Optional) switch#**show vpc brief**
6. (Optional) switch#**copy running-config startup-config**

DETAILED STEPS

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

	Command or Action	Purpose
	switch# configure terminal switch(config)#	
Step 2	switch(config)# vpc domain <i>domain-id</i> Example: switch(config)# vpc domain 5 switch(config-vpc-domain)#	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>.
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
```




CHAPTER 7

Configuring Static and Dynamic NAT Translation

- [Network Address Translation Overview, on page 149](#)
- [Information About Static NAT, on page 150](#)
- [Dynamic NAT Overview, on page 151](#)
- [Timeout Mechanisms, on page 151](#)
- [NAT Inside and Outside Addresses, on page 153](#)
- [Pool Support for Dynamic NAT, on page 153](#)
- [Static and Dynamic Twice NAT Overview, on page 154](#)
- [Guidelines and Limitations for Static NAT, on page 154](#)
- [Restrictions for Dynamic NAT, on page 155](#)
- [Guidelines and Limitations for Dynamic Twice NAT, on page 156](#)
- [Configuring Static NAT, on page 156](#)
- [Configuring Dynamic NAT, on page 164](#)
- [Information About VRF Aware NAT, on page 174](#)
- [Configuring VRF Aware NAT, on page 174](#)

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.

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 6: 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. The minimum value of the sampling-timeout in the **ip nat translation sampling-timeout** command is 120 seconds. However it is recommended to configure the **ip nat translation sampling-timeout** value to 15 minutes or higher.

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

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. **syn-timeout** and **finrst-timeout** is supported only on Cisco Nexus 3500 Series switches. 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 **syn-timeout** is not supported on Cisco 3100 Series 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.
 - If an RST packet is received after the connection is established, SYN-->SYN-ACK-->RST, the flows are expired after the configured timeout value.
 - 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.



Note If dynamic pool-based configuration is used and a FIN-ACK is received, the translation entry is not cleared.

finrst-timeout is not supported on Cisco 3100 Series switches.

The timeout value ranges from 1 second to 172800 seconds. The default value is 60 seconds.

- **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. This timeout value starts after the sampling timeout value completes.

The timeout value ranges from 60 seconds to 172800 seconds, including the sampling-timeout.

- **udp-timeout**—Timeout value for all NAT UDP packets.

The timeout value ranges from 60 seconds to 172800 seconds, including the sampling-timeout.

- **timeout**—Timeout value for dynamic NAT translations.

The timeout value ranges from 60 seconds to 172800 seconds, including the sampling-timeout.

- **sampling-timeout**—Time after which the device checks for dynamic translation activity.

The timeout value ranges from 120 seconds to 172800 seconds.



Note **sampling-timeout** is not supported on Cisco 3100 Series switches.

The **tcp-timeout**, **udp-timeout**, and the **timeout** value timers are triggered after the timeout configured for the **ip nat translation sampling-timeout** command expires.

The SYN, FIN and RST timers are not used for dynamic pool-based NAT.



Note All the above timers will take additional time (01 to 30 seconds) to expire. This additional time is to randomize the timer expiry events for performance and optimization.

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

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 tries to allocate the original source port again. This process continues until PAT runs out of available ports and IP addresses.

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.

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.

Guidelines and Limitations for Static NAT

Static NAT has the following configuration guidelines and limitations:

- NAT is supported on Cisco Nexus 31108PC-V, Cisco Nexus 31108TC-V, Cisco Nexus 3132Q-V, Cisco Nexus 3132Q/3132Q-X, Cisco Nexus 3164Q, Cisco Nexus 3172PQ, Cisco Nexus 3172TQ, Cisco Nexus 31128PQ switches. However NAT is not supported on Cisco Nexus 3048 and Cisco Nexus 3064 switches.
- NAT supports up to 1024 translations which include both static and dynamic NAT.
- The Cisco Nexus device supports NAT on the following interface types:
 - Switch Virtual Interfaces (SVIs)
 - Routed ports
 - Layer 3 port channels
- NAT is supported for IPv4 Unicast only.

- The Cisco Nexus device does not support the following:
 - 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.
- Egress ACLs are applied to the original packets and not the NAT translated packets.
- HSRP and VRRP are not supported on a NAT interface.
- Warp mode latency performance is not supported on packets coming from the outside to the inside domain.
- 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.
- If the translated IP is part of the outside interface subnet, then use the **ip local-proxy-arp** command on the NAT outside interface.
- Twice NAT is not supported. (Twice NAT is a variation of NAT in that both the source and destination addresses are modified by NAT as a datagram crosses address domains (inside to outside or outside to inside.)
- 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.
- 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):

- 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.
- NAT and VLAN Access Control Lists (VACLs) are not supported together on an interface. You can configure either NAT or VACLs on an interface.
- Egress ACLs are not applied to translated packets.
- MIBs are not supported.
- Cisco Data Center Network Manager (DCNM) is not supported.

- Dynamic NAT on traffic coming from outside domains is not supported.
- 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.
- When creating a new translation on a Cisco Nexus 3548 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.
- ICMP NAT translation is supported only on Cisco Nexus 3500 Series switches.
- NAT configurations such as ip nat inside or ip nat outside are not supported on loopback interfaces.

Guidelines and Limitations for Dynamic Twice NAT

See the following guidelines for configuring dynamic twice 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.
- The maximum number of supported ICMP translations or flow entries is 176 for an optimal utilization of the TCAM space.
- NAT is ECMP aware and it supports a maximum of 24 ECMP paths.

Configuring Static NAT

Enabling Static NAT

SUMMARY STEPS

1. switch# **configure terminal**
2. switch(config)# **feature nat**
3. (Optional) switch(config)# **copy running-config startup-config**

DETAILED STEPS

	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.

	Command or Action	Purpose
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.

Configuring Static NAT on an Interface

SUMMARY STEPS

1. switch# **configure terminal**
2. switch(config)# **interface** *type slot/port*
3. switch(config-if)# **ip nat** {**inside** | **outside**}
4. (Optional) switch(config)# **copy running-config startup-config**

DETAILED STEPS

	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).

SUMMARY STEPS

1. switch# **configure terminal**
2. switch(config)# **ip nat inside source static** *local-ip-address global-ip-address* [**group group-id**]
3. (Optional) switch(config)# **copy running-config startup-config**

DETAILED STEPS

	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</i> [group group-id]	Configures static NAT to translate the inside global address to the inside local address or to translate the opposite (the inside local traffic to the inside global traffic).
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.

SUMMARY STEPS

1. switch# **configure terminal**
2. switch(config)# **ip nat outside source static** *global-ip-address local-ip-address* [**group group-id**] [**add-route**]
3. (Optional) switch(config)# **copy running-config startup-config**

DETAILED STEPS

	Command or Action	Purpose
Step 1	switch# configure terminal	Enters global configuration mode.
Step 2	switch(config)# ip nat outside source static <i>global-ip-address local-ip-address [group group-id]</i> [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).
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).

SUMMARY STEPS

1. switch# **configure terminal**
2. switch(config)# **ip nat inside source static** *{inside-local-address outside-local-address | {tcp|udp} inside-local-address {local-tcp-port | local-udp-port} inside-global-address {global-tcp-port | global-udp-port}}*
3. (Optional) switch(config)# **copy running-config startup-config**

DETAILED STEPS

	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-address outside-local-address {tcp udp}</i> <i>inside-local-address {local-tcp-port local-udp-port}</i> <i>inside-global-address {global-tcp-port global-udp-port}}</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).

SUMMARY STEPS

1. switch# **configure terminal**
2. switch(config)# **ip nat outside source static** *{outside-global-address outside-local-address | {tcp | udp} outside-global-address {global-tcp-port | global-udp-port} outside-local-address {global-tcp-port | global-udp-port}}*
3. (Optional) switch(config)# **copy running-config startup-config**

DETAILED STEPS

	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 outside-local-address {tcp udp} outside-global-address {global-tcp-port global-udp-port} outside-local-address {global-tcp-port global-udp-port}}</i>	Maps static NAT to an outside global port to an outside local 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 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.

SUMMARY STEPS

1. **enable**
2. **configure terminal**
3. **ip nat inside source static** *inside-local-ip-address inside-global-ip-address* [**group** *group-id*]
4. **ip nat outside source static** *outside-global-ip-address outside-local-ip-address* [**group** *group-id*] [**add-route**]
5. **interface** *type number*
6. **ip address** *ip-address mask*
7. **ip nat inside**
8. **exit**
9. **interface** *type number*
10. **ip address** *ip-address mask*
11. **ip nat outside**
12. **end**

DETAILED STEPS

	Command or Action	Purpose
Step 1	enable Example: <pre>switch> enable</pre>	Enables privileged EXEC mode. • 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 inside-global-ip-address</i> [group <i>group-id</i>] 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. • The group keyword determines the group to which a translation belongs.
Step 4	ip nat outside source static <i>outside-global-ip-address 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. • 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.

	Command or Action	Purpose
Step 7	ip nat inside Example: switch(config-if)# ip nat inside	Connects the interface to an inside network, which is subject to NAT. Note This configuration is not supported on loopback interface.
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 This configuration is 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
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:

SUMMARY STEPS

1. switch# show ip nat translations

DETAILED STEPS

	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

```

```

--- ---
--- 11.1.1.1      101.1.1.1      51.1.2.1      102.1.2.1
--- 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 ---           ---

```

Configuring Dynamic NAT

Configuring Dynamic Translation and Translation Timeouts

SUMMARY STEPS

1. **enable**
2. **configure terminal**
3. **ip access-list** *access-list-name*
4. **permit** *protocol source source-wildcard any*
5. **deny** *protocol source source-wildcard any*
6. **exit**
7. **ip nat inside source list** *access-list-name interface type number overload*
8. **interface** *type number*
9. **ip address** *ip-address mask*
10. **ip nat inside**
11. **exit**
12. **interface** *type number*
13. **ip address** *ip-address mask*
14. **ip nat outside**
15. **exit**
16. **ip nat translation tcp-timeout** *seconds*
17. **ip nat translation max-entries** *number-of-entries*
18. **ip nat translation udp-timeout** *seconds*
19. **ip nat translation timeout** *seconds*
20. **end**

DETAILED STEPS

	Command or Action	Purpose
Step 1	enable Example:	Enables privileged EXEC mode. <ul style="list-style-type: none"> • Enter your password if prompted.

	Command or Action	Purpose
	Switch> enable	
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. The deny rule is treated as a permit rule, and the packets matching the criteria mentioned in the deny rule are forwarded without NAT translation.
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> 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.
Step 11	exit Example:	Exits interface configuration mode and returns to global configuration mode.

	Command or Action	Purpose
	<code>Switch(config-if)# exit</code>	Note This configuration is not supported on loopback interface.
Step 12	interface <i>type number</i> Example: <code>Switch(config)# interface ethernet 1/1</code>	Configures an interface and enters interface configuration mode.
Step 13	ip address <i>ip-address mask</i> Example: <code>Switch(config-if)# ip address 172.16.232.182 255.255.255.240</code>	Sets a primary IP address for an interface.
Step 14	ip nat outside Example: <code>Switch(config-if)# ip nat outside</code>	Connects the interface to an outside network. Note This configuration is not supported on loopback interface.
Step 15	exit Example: <code>Switch(config-if)# exit</code>	Exits interface configuration mode and returns to global configuration mode.
Step 16	ip nat translation tcp-timeout <i>seconds</i> Example: <code>Switch(config)# ip nat translation tcp-timeout 50000</code>	Specifies the timeout value for TCP-based dynamic NAT entries. <ul style="list-style-type: none">• Dynamically created NAT translations are cleared when the configured timeout limit is reached. All configured timeouts are triggered after the timeout configured for the ip nat translation sampling-timeout command expires.
Step 17	ip nat translation max-entries <i>number-of-entries</i> Example: <code>Switch(config)# ip nat translation max-entries 300</code>	Specifies the maximum number of dynamic NAT translations. The number of entries can be between 1 and 1023.
Step 18	ip nat translation udp-timeout <i>seconds</i> Example: <code>Switch(config)# ip nat translation udp-timeout 45000</code>	Specifies the timeout value for UDP-based dynamic NAT entries. <ul style="list-style-type: none">• Dynamically created NAT translations are cleared when the configured timeout limit is reached. All configured timeouts are triggered after the timeout configured for the ip nat translation sampling-timeout command expires.
Step 19	ip nat translation timeout <i>seconds</i> Example: <code>switch(config)# ip nat translation timeout 13000</code>	Specifies the timeout value for dynamic NAT translations.

	Command or Action	Purpose
Step 20	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

SUMMARY STEPS

1. switch# **configure terminal**
2. switch(config)# **feature nat**
3. switch(config)# **ip nat pool** *pool-name* [*startip endip*] {**prefix** *prefix-length* | **netmask** *network-mask*}
4. (Optional) switch(config-ipnat-pool)# **address** *startip endip*
5. (Optional) switch(config)# **no ip nat pool** *pool-name*

DETAILED STEPS

	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.

SUMMARY STEPS

1. switch# **configure terminal**
2. (Optional) switch# **ip nat inside source list** *list-name* **pool** *pool-name* [**overload**]
3. (Optional) switch# **ip nat outside source list** *list-name* **pool** *pool-name* [**add-route**]

DETAILED STEPS

	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</i> pool <i>pool-name</i> [overload]	Creates a NAT inside source list with pool with or without overloading.
Step 3	(Optional) switch# ip nat outside source list <i>list-name</i> pool <i>pool-name</i> [add-route]	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.

SUMMARY STEPS

1. switch# **configure terminal**
2. switch(config)# **ip nat outside source static** *outside-global-ip-address outside-local-ip-address* | [**tcp** | **udp**] *outside-global-ip-address outside-global-port outside-local-ip-address outside-local-port* [**group group-id**] [**add-route**] [**dynamic**]
3. switch(config)# **ip nat inside source list** *access-list-name* [**interface type slot/port overload** | **pool pool-name**] [**group group-id**] [**dynamic**]
4. switch(config)# **ip nat pool** *pool-name* [*startip endip*] {**prefix prefix-length** | **netmask network-mask**}
5. switch(config)# **interface type slot/port**
6. switch(config-if)# **ip nat outside**
7. switch(config-if)# **exit**
8. switch(config)# **interface type slot/port**
9. switch(config-if)# **ip nat inside**

DETAILED STEPS

	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 outside-local-ip-address</i> [tcp udp] <i>outside-global-ip-address outside-global-port outside-local-ip-address outside-local-port</i> [group group-id] [add-route] [dynamic]	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] [group group-id] [dynamic]	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.

	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 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.

SUMMARY STEPS

1. switch# **configure terminal**
2. switch(config)# **ip nat inside source static** *inside-local-ip-address* *inside-global-ip-address* | [**tcp** | **udp**] *inside-local-ip-address* *local-port* *inside-global-ip-address* *global-port* [**group** *group-id*] [**dynamic**]
3. switch(config)# **ip nat outside source list** *access-list-name* [**interface** *type slot/port* **pool** *pool-name*] [**group** *group-id*] [**add-route**] [**dynamic**]
4. switch(config)# **ip nat pool** *pool-name* [*startip endip*] {**prefix** *prefix-length* | **netmask** *network-mask*}

5. switch(config)# **interface** *type slot/port*
6. switch(config-if)# **ip nat outside**
7. switch(config-if)# **exit**
8. switch(config)# **interface** *type slot/port*
9. switch(config-if)# **ip nat inside**

DETAILED STEPS

	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</i> <i>inside-global-ip-address global-port</i> [group group-id] [dynamic]	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</i> [interface <i>type slot/port pool pool-name</i>] [group group-id] [add-route] [dynamic]	Establishes dynamic source translation by creating a NAT outside source list with pool.
Step 4	switch(config)# ip nat pool <i>pool-name</i> [<i>startip endip</i>] { prefix prefix-length netmask network-mask }	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_2 pool pool_2 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 ip nat translations verbose	Displays active Network Address Translation (NAT) translations including dynamic translations in a more readable format.
show run nat	Displays NAT configuration.

Example

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
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

switch# sh ip nat translations verbose
Pro Inside global      Inside local      Outside local     Outside global
any 1.1.1.1            10.1.1.2         ---              ---
    Flags:0x1 Entry-id:0 State:0x0 Group_id:0 Format(H:M:S) Time-left:0:0:-1
icmp 101.1.0.1:65351  101.0.0.1:0      102.1.0.1:231    102.1.0.1:231
    Flags:0x82 Entry-id:101 State:0x3 Group_id:0 VRF: red Format(H:M:S) Time-left:12:0:9
udp  101.1.0.1:65383   101.0.0.1:63     102.1.0.1:63     102.1.0.1:63
    Flags:0x82 Entry-id:103 State:0x3 Group_id:0 VRF: red Format(H:M:S) Time-left:12:0:9
tcp  101.1.0.1:64549   101.0.0.1:8809   102.1.0.1:9087   102.1.0.1:9087
    Flags:0x82 Entry-id:102 State:0x1 Group_id:0 VRF: red Format(H:M:S) Time-left:12:0:9
    syn:0:1:9 fin-rst:12:0:9

```

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

switch# show ip nat translations verbose
Pro Inside global      Inside local      Outside local     Outside global
any 1.1.1.1            10.1.1.2         ---              ---
    Flags:0x1 Entry-id:0 State:0x0 Group_id:0 Format(H:M:S) Time-left:0:0:-1
any 101.1.0.1         101.0.0.1        ---              ---
    Flags:0x0 Entry-id:92 State:0x3 Group_id:0 VRF: red Format(H:M:S) Time-left:12:0:11

```

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 tcp-timeout 50000
Switch(config)# ip nat translation max-entries 300
Switch(config)# ip nat translation udp-timeout 45000
Switch(config)# ip nat translation timeout 13000
Switch(config)# end

```

Information About VRF Aware NAT

VRF aware NAT is supported by static and dynamic NAT configurations. When the 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.

When the 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. A NAT outside configuration is not supported on a non-default VRF interface when the NAT inside is configured on a default VRF interface.

When overlapping addresses are configured across different VRFs for a NAT inside interface, a NAT outside interface should not be the default VRF interface. For example, vrfA and vrfB are configured as NAT inside interfaces with same source subnets and a NAT outside interface is configured as the default VRF. NAT is not supported in a configuration like this because of the ambiguity in routing packets from a NAT outside interface to NAT inside interface.

Configuring VRF Aware NAT

Before you begin

Ensure that you enable NAT on the switch.

SUMMARY STEPS

1. switch# **configure terminal**
2. switch(config)# **[no] ip nat** *inside | outside* **source list** *ACL_NAME* [*interface INTERFACE NAME* **overload**] [*pool POOL NAME* **overload**] [**group** *group-id*] [**dynamic**] [**vrf** *<vrf-name>*] [**match-in-vrf**]
3. switch(config)# **[no] ip nat** *inside | outside* **source static** *LOCAL IP GLOBAL IP* [*tcp | udp LOCAL IP LOCAL PORT GLOBAL IP GLOBAL PORT*] [**group** *group-id*] [**dynamic**] [**vrf** *<vrf-name>*] [**match-in-vrf**]
4. switch(config)# **interface** *type slot/port* [**vrf** *<vrf-name>* **ip nat** *inside | outside*]

DETAILED STEPS

	Command or Action	Purpose
Step 1	switch# configure terminal	Enters global configuration mode.
Step 2	switch(config)# [no] ip nat inside outside source list <i>ACL_NAME</i> [<i>interface INTERFACE NAME</i> overload]][<i>pool POOL NAME</i> overload] [group <i>group-id</i>] [dynamic] [vrf < <i>vrf-name</i> > [match-in-vrf]]	Creates or deletes dynamic NAT with VRF specific. The group keyword determines the group to which a translation belongs.
Step 3	switch(config)# [no] ip nat inside outside source static <i>LOCAL IP GLOBAL IP</i> [<i>tcp udp LOCAL IP LOCAL</i> <i>PORT GLOBAL IP GLOBAL PORT</i>] [group <i>group-id</i>] [dynamic] [vrf < <i>vrf-name</i> > [match-in-vrf]]	Creates or deletes a VRF specific static NAT. The group keyword determines the group to which a translation belongs.
Step 4	switch(config)# interface <i>type slot/port</i> [vrf < <i>vrf-name</i> > ip nat <i>inside outside</i>	Enables NAT on a VRF-aware interface.

See the output of **show run nat** command.

```
#show run nat
...
feature nat
ip nat inside source static 1.1.1.1 1.1.1.100 vrf red match-in-vrf
ip nat outside source static 2.2.2.200 2.2.2.2 vrf red match-in-vrf add-route
ip nat inside source list nat-acl-in1 pool pool-in1 vrf red match-in-vrf overload
ip nat outside source list nat-acl-out1 pool pool-out1 vrf red match-in-vrf add-route
interface Ethernet1/3
    ip nat outside
interface Ethernet1/5
    ip nat inside

N3548#show ip nat translation verbose
Pro Inside global      Inside local      Outside local      Outside global
any 1.1.1.1            10.1.1.2          ---                ---
    Flags:0x1  Entry-id:0  State:0x0  Group_id:0  Format(H:M:S)  Time-left:0:0:-1
icmp 101.1.0.1:65351  101.0.0.1:0      102.1.0.1:231    102.1.0.1:231
    Flags:0x82  Entry-id:101  State:0x3  Group_id:0  VRF: red  Format(H:M:S)  Time-left:12:0:9
udp 101.1.0.1:65383  101.0.0.1:63     102.1.0.1:63     102.1.0.1:63
    Flags:0x82  Entry-id:103  State:0x3  Group_id:0  VRF: red  Format(H:M:S)  Time-left:12:0:9
tcp 101.1.0.1:64549  101.0.0.1:8809   102.1.0.1:9087   102.1.0.1:9087
    Flags:0x82  Entry-id:102  State:0x1  Group_id:0  VRF: red  Format(H:M:S)  Time-left:12:0:9

syn:0:1:9  fin-rst:12:0:9
```




CHAPTER 8

Information About Q-in-Q Tunnels

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.

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 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 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 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 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.



Note Selective Q-in-Q tunneling is not supported. All frames entering the tunnel port are subjected to Q-in-Q tagging.

Figure 7: 802.1Q-in-Q Tunnel Ports

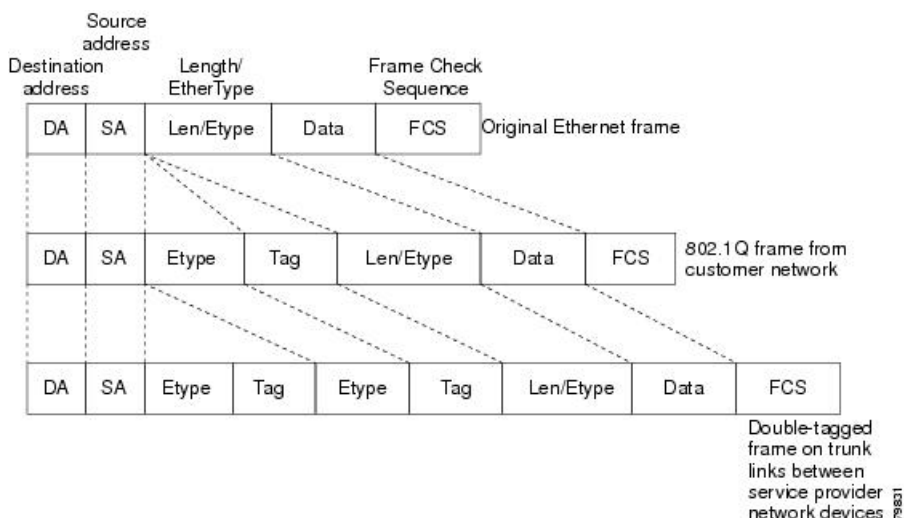


Packets that enter the tunnel port on the service-provider edge switch, which are already 802.1Q-tagged with the appropriate VLAN IDs, are encapsulated with another layer of an 802.1Q tag that contains a VLAN ID that is unique to the customer. The original 802.1Q tag from the customer is preserved in 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.

The following figure shows the differences between the untagged, tagged and double-tagged ethernet frames.

Figure 8: 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, that is multi-level dot1q tagging Q-in-Q, is not supported.

- [Native VLAN Hazard, on page 179](#)
- [Information About Layer 2 Protocol Tunneling, on page 180](#)
- [Guidelines and Limitations for Q-in-Q Tunneling, on page 182](#)
- [Configuring Q-in-Q Tunnels and Layer 2 Protocol Tunneling, on page 183](#)
- [Verifying the Q-in-Q Configuration, on page 187](#)
- [Configuration Example for Q-in-Q and Layer 2 Protocol Tunneling, on page 188](#)
- [Feature History for Q-in-Q Tunnels and Layer 2 Protocol Tunneling, on page 188](#)


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.

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 the 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.

The following figure shows the native VLAN hazard.

Figure 9: Native VLAN Hazard



A couple of ways to solve the native VLAN problem, are as follows:

- 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. 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.

You can configure the switch to allow multi-tagged BPDUs on a tunnel port. If you enable *l2 protocol tunnel allow-double-tag*, 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 comes into the supervisor module cause the CPU load to go up. The load is controlled by Control Plane Policing CoPP configured for packets marked as BPDU.

For example, the following figure shows Customer X has four switches in the same VLAN that are connected through the service-provider network. If the network does not tunnel BPDUs, the switches on the far ends of the network cannot properly run the STP, CDP, 802.1X, and VTP protocols.


Figure 10: 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 following figure shows the resulting topology on the customer's network when BPDU tunneling is not enabled.

Figure 11: Virtual Network Topology Without BPDU Tunneling



Guidelines and Limitations for Q-in-Q Tunneling

Q-in-Q tunnels and Layer 2 tunneling have the following configuration guidelines and limitations:

- Switches in the service-provider network must be configured to handle the increase in MTU size due to Q-in-Q tagging.
- Cisco Nexus 3500 Series switches do not support configuring Q-in-Q Tunneling on Cisco NX-OS Release 7.0(3)I7(2) and the previous releases.
- Cisco Nexus 3500 Series switches do not support Q-in-Q tunneling. However they forward Q-in-Q traffic.
- Selective Q-in-Q tunneling is not supported. All frames that enter the tunnel port will be subject 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 may 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.
- You should use MAC address-based frame distribution.
- 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.
- CDP must be explicitly disabled, as needed, on the dot1Q tunnel port.
- You must disable IGMP snooping on the tunnel VLANs.
- You should run the **vlan dot1Q tag native** command to maintain the tagging on the native VLAN and drop untagged traffic to prevent native VLAN misconfigurations.
- You must manually configure the 802.1Q interfaces to be edge ports.
- Dot1x tunneling is not supported.
- Cisco Nexus 3400 platform switches do not support Q-in-Q.

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 VLAN membership of the port is changed when you enter 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.

Before you begin

You must first configure the interface as a switchport.

SUMMARY STEPS

1. switch# **configure terminal**
2. switch(config)# **interface ethernet slot/port**
3. switch(config-if)# **switchport**
4. switch(config-if)# **[no] switchport mode dot1q-tunnel**
5. switch(config-if)# **[no] l2protocol tunnel allow-double-tag**
6. switch(config-if)# **exit**
7. (Optional) switch(config)# **show dot1q-tunnel [interface if-range]**
8. (Optional) switch(config)# **copy running-config startup-config**

DETAILED STEPS

	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)# [no] switchport mode dot1q-tunnel	Creates an 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)# [no] l2protocol tunnel allow-double-tag	Enable or disable multi tagged BPDU support.
Step 6	switch(config-if)# exit	Exits interface configuration mode.

	Command or Action	Purpose
Step 7	(Optional) switch(config)# show dot1q-tunnel [<i>interface if-range</i>]	Displays all ports that are in dot1q-tunnel mode. Optionally you can specify an interface or range of interfaces to display.
Step 8	(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 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)# exit
switch(config)# exit
switch# show dot1q-tunnel
```

Enabling the Layer 2 Protocol Tunnel

You can enable protocol tunneling on the 802.1Q tunnel port.

SUMMARY STEPS

1. switch# **configure terminal**
2. switch(config)# **interface ethernet slot/port**
3. switch(config)# **switchport**
4. switch(config-if)# **switchport mode dot1q-tunnel**
5. switch(config-if)# **[no] l2protocol tunnel [cdp | stp | vtp]**
6. switch(config-if)# **exit**
7. (Optional) switch(config)# **copy running-config startup-config**

DETAILED STEPS

	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)# switchport	Sets the interface as a Layer 2 switching port.
Step 4	switch(config-if)# switchport mode dot1q-tunnel	Creates an 802.1Q tunnel on the port.
Step 5	switch(config-if)# [no] l2protocol tunnel [cdp stp vtp]	Enables Layer 2 protocol tunneling. Optionally, you can enable CDP, STP, or VTP tunneling.
Step 6	switch(config-if)# exit	Exits interface configuration mode.

	Command or Action	Purpose
Step 7	(Optional) <code>switch(config)# copy running-config startup-config</code>	Saves the change persistently through reboots and restarts by copying 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 Thresholds for Layer 2 Protocol Tunnel Ports

You can specify the port drop and shutdown value for a Layer 2 protocol tunneling port.

SUMMARY STEPS

1. `switch# configure terminal`
2. `switch(config)# interface ethernet slot/port`
3. `switch(config-if)# switchport`
4. `switch(config-if)# switchport mode dot1q-tunnel`
5. `switch(config-if)# [no] l2protocol tunnel drop-threshold [cdp | stp | vtp]`
6. `switch(config-if)# [no] l2protocol tunnel shutdown-threshold [cdp | stp | vtp]`
7. `switch(config-if)# exit`
8. (Optional) `switch(config)# copy running-config startup-config`

DETAILED STEPS

	Command or Action	Purpose
Step 1	<code>switch# configure terminal</code>	Enters global configuration mode.
Step 2	<code>switch(config)# interface ethernet slot/port</code>	Specifies an interface to configure, and enters interface configuration mode.
Step 3	<code>switch(config-if)# switchport</code>	Sets the interface as a Layer 2 switching port.
Step 4	<code>switch(config-if)# switchport mode dot1q-tunnel</code>	Creates an 802.1Q tunnel on the port.
Step 5	<code>switch(config-if)# [no] l2protocol tunnel drop-threshold [cdp stp vtp]</code>	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.

	Command or Action	Purpose
		The no form of this command resets the threshold values to 0 and disables the drop threshold.
Step 6	switch(config-if)# [no] l2protocol tunnel shutdown-threshold [cdp stp vtp]	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 specify the Cisco Discovery Protocol (CDP), Spanning Tree Protocol (STP), or VLAN Trunking Protocol (VTP). Valid values for the packets is from 1 to 4096.
Step 7	switch(config-if)# exit	Exits interface configuration mode.
Step 8	(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 a threshold for a Layer 2 protocol tunnel port:

```
switch# configure terminal
switch(config)# interface ethernet 7/1
switch(config-if)# switchport
switch(config-if)# switchport mode dot1q-tunnel
switch(config)# l2protocol tunnel drop-threshold 3000
switch(config)# l2protocol tunnel shutdown-threshold 3000
switch(config)# exit
switch# copy running-config startup-config
```

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.

SUMMARY STEPS

1. switch# **configure terminal**
2. switch(config)# **interface interface-id**
3. switch(config-if)# **switchport mode dot1q-tunnel**
4. switch(config-if)# **switchport vlan mapping vlan-id-range dot1q-tunnel outer vlan-id**
5. switch(config-if)# **exit**
6. switch# **show interfaces interface-id vlan mapping**
7. switch# **copy running-config startup-config**

DETAILED STEPS

	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	Configure the interface as a tunnel port.
Step 4	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 5	switch(config-if)# exit	Exits the configuration mode.
Step 6	switch# show interfaces <i>interface-id</i> vlan mapping	Verifies the configuration.
Step 7	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 drop all VLANs other than the configured mapping and allowed VLANs.

```
switch(config)# interface port-channel201
switch(config-if)# switchport vlan mapping dot1q-tunnel allowed-vlan 201-204
switch(config-if)# switchport vlan mapping 300-400 dot1q-tunnel 500
switch(config-if)# spanning-tree port type edge trunk
switch(config-if)# spanning-tree bpdupfilter enable
switch(config-if)# vpc 201
```

Verifying the Q-in-Q Configuration

Use the following command to verify the Q-in-Q tunnel and Layer 2 protocol tunneling configuration information:

Command	Purpose
clear l2protocol tunnel counters [<i>interface 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 [<i>interface if-range</i>]	Displays a range of interfaces or all interfaces that are in dot1q-tunnel mode.

Command	Purpose
show l2protocol tunnel [<i>interface if-range</i> vlan <i>vlan-id</i>]	Displays Layer 2 protocol tunnel information for a range of interfaces or 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 Example 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)# vlan configuration 8
switch(config-vlan-config)# no ip igmp snooping
switch(config-vlan-config)# exit
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#
```

Feature History for Q-in-Q Tunnels and Layer 2 Protocol Tunneling

Table 10: Feature History for Q-in-Q Tunnels and Layer 2 Protocol Tunneling

Feature Name	Release	Feature Information
Q-in-Q VLAN Tunnels	6.0(2)U1(1)	This feature was introduced.
L2 Protocol Tunneling	6.0(2)U1(1)	This feature was introduced.



INDEX

40-Gigabit Ethernet interface speed **6**
40-Gigabit Ethernet mode **6**
802.1q tunnel port, creating **183**
 interfaces **183**

A

adding ports **73**
 port channels **73**

B

bandwidth **47**
 configuring **47**

C

channel mode **76**
 port channels **76**
channel modes **70**
 port channels **70**
configuration **56**
 Layer 3 interfaces **56**
 verifying **56**
configuration example **162**
 static nat **162**
configuration examples **59, 110**
 ip tunneling **110**
 Layer 3 interfaces **59**
configuring **30, 33, 45–49, 80–81**
 description parameter **33**
 error-disabled recovery interval **30**
 interface bandwidth **47**
 LACP fast timer rate **80**
 LACP port priority **81**
 loopback interfaces **49**
 routed interfaces **45**
 subinterfaces **46**
 VLAN interfaces **48**
configuring 10 GbE interface speed **20**
configuring 40 GbE interface speed **21**
Configuring a DHCP client on an interface **55**
configuring dynamic pool **167**
configuring LACP **76**

D

debounce timer **12**
 parameters **12**
debounce timer, configuring **32**
 Ethernet interfaces **32**
default interface **12**
default settings **43, 100**
 ip tunnels **100**
 Layer 3 interfaces **43**
DHCP client configuration **43**
DHCP client configuration limitations **44**
DHCP client discovery **43**
disabling **23, 28, 31, 34, 129**
 CDP **28**
 error-disabled recovery **31**
 ethernet interfaces **34**
 link negotiation **23**
 vPCs **129**
downlink delay **14**

E

enabling **28–30**
 CDP **28**
 error-disabled detection **29**
 error-disabled recovery **30**
Ethernet interfaces **5, 32**
 debounce timer, configuring **32**
 interface speed **5**

F

feature history **62, 84, 111, 188**
 ip tunnels **111**
 Layer 3 interfaces **62**
 port channels **84**
 q-in-q tunnels, layer 2 protocol tunneling **188**

G

gre tunnel, configuring **106**
 interfaces **106**

gre tunnels [96](#)
 interfaces [96](#)

I

inside source address [157](#)
 static NAT, configuring [157](#)
 interface information, displaying [35](#)
 layer 2 [35](#)
 interface MAC address, configuring [51](#)
 interface port-channel [79](#)
 interface speed [5, 19](#)
 configuring [19](#)
 Ethernet interfaces [5](#)
 interface tunnel [103](#)
 interface, configuring [157](#)
 static NAT [157](#)
 interfaces [3–4, 39, 41–42, 47–50, 58–59, 95–96, 101, 106, 110, 177, 180, 183–185, 187](#)
 802.1q tunnel port, creating [183](#)
 assigning to a VRF [50](#)
 chassis ID [3](#)
 configuring bandwidth [47](#)
 gre tunnel, configuring [106](#)
 gre tunnels [96](#)
 ip tunnel configuration, verifying [110](#)
 ip tunnels [95](#)
 ipip tunnel decapsulation-only, configuring [106](#)
 ipip tunnel, configuring [106](#)
 layer 2 protocol tunnel [184](#)
 layer 2 protocol tunnel ports, thresholds configuring [185](#)
 layer 2 protocol tunneling [180](#)
 Layer 3 [39, 58–59](#)
 configuration examples [59](#)
 monitoring [58](#)
 loopback [42, 49](#)
 options [3](#)
 q-in-q configuration, verifying [187](#)
 q-in-q tunnels [177](#)
 routed [39](#)
 tunnel [42](#)
 tunnel interface, creating [101](#)
 UDLD [4](#)
 VLAN [41, 48](#)
 configuring [48](#)
 ip tunnel configuration, verifying [110](#)
 interfaces [110](#)
 ip tunneling [110](#)
 configuration examples [110](#)
 ip tunnels [95–96, 100, 111](#)
 default settings [100](#)
 feature history [111](#)
 interfaces [95](#)
 prerequisites [96](#)
 standards [111](#)

ipip decapsulate-only [96](#)

L

LACP [64, 69–72, 76–77](#)
 configuring [76](#)
 marker responders [71](#)
 port channel, minlinks [72, 77](#)
 port channels [69](#)
 system ID [70](#)
 LACP fast timer rate [80](#)
 configuring [80](#)
 lacp max-bundle [79](#)
 LACP port priority [81](#)
 configuring [81](#)
 LACP-enabled vs static [71](#)
 port channels [71](#)
 layer 2 [10, 25, 35](#)
 interface information, displaying [35](#)
 svi autostate [10](#)
 svi autostate, disabling [25](#)
 layer 2 protocol tunnel [184](#)
 interfaces [184](#)
 layer 2 protocol tunneling [180](#)
 interfaces [180](#)
 Layer 3 interfaces [39, 43, 45, 56, 58–59, 62](#)
 configuration examples [59](#)
 configuring routed interfaces [45](#)
 default settings [43](#)
 feature history [62](#)
 interfaces [62](#)
 Layer 3 [62](#)
 feature history [62](#)
 MIBs [62](#)
 related documents [62](#)
 standards [62](#)
 MIBs [62](#)
 monitoring [58](#)
 related documents [62](#)
 standards [62](#)
 verifying [56](#)
 limitations [44](#)
 Link Aggregation Control Protocol [64](#)
 load balancing [74](#)
 port channels [74](#)
 configuring [74](#)
 loopback interfaces [42, 49](#)
 configuring [49](#)

M

MIBs [37, 62](#)
 Layer 2 interfaces [37](#)
 Layer 3 interfaces [62](#)

- monitoring [58](#)
 - Layer 3 interfaces [58](#)
- mtu [103–104](#)
- Multi-point IP-in-IP decapsulation [96](#)

N

- NAT [163](#)
 - verifying [163](#)
- NVGRE traffic [68](#)

O

- outside address [158](#)
 - static NAT, configuring [158](#)

P

- parameters, about [12](#)
 - debounce timer [12](#)
- PAT [162](#)
 - configuration example [162](#)
- physical Ethernet settings [14](#)
- point-to-point IP-in-IP encapsulation and decapsulation [96](#)
- port [159–160](#)
 - static PAT, configuring [159–160](#)
- port channel [82](#)
 - verifying configuration [82](#)
- port channel, minlinks [72, 77](#)
 - LACP [72, 77](#)
- port channeling [64](#)
- port channels [47, 63–64, 66, 69, 71–74, 76, 84, 141](#)
 - adding ports [73](#)
 - channel mode [76](#)
 - compatibility requirements [64](#)
 - configuring bandwidth [47](#)
 - creating [72](#)
 - feature history [84](#)
 - LACP [69](#)
 - LACP-enabled vs static [71](#)
 - load balancing [66, 74](#)
 - port channels [66](#)
 - moving into a vPC [141](#)
 - STP [63](#)
- port mode [17](#)
 - interface [17](#)
- port modes [7](#)
- prerequisites [96](#)
 - ip tunnels [96](#)

Q

- q-in-q configuration, verifying [187](#)
 - interfaces [187](#)

- q-in-q tunnels [177](#)
 - interfaces [177](#)
- q-in-q tunnels, layer 2 protocol [188](#)
 - feature history [188](#)

R

- related documents [62](#)
 - Layer 3 interfaces [62](#)
- resilient hashing [68](#)
- restarting [34](#)
 - ethernet interfaces [34](#)
- routed interfaces [39, 45, 47](#)
 - configuring [45](#)
 - configuring bandwidth [47](#)

S

- security [150](#)
 - static NAT [150](#)
- SFP+ transceiver [5](#)
- show interfaces tunnel [103–104](#)
- show running-config interface port-channel [79](#)
- show vpc brief [138–139](#)
- Small form-factor pluggable (plus) transceiver [5](#)
- standards [62, 111](#)
 - ip tunnels [111](#)
 - Layer 3 interfaces [62](#)
- static nat [162](#)
 - configuration example [162](#)
- static NAT [150, 156–157, 163](#)
 - enabling [156](#)
 - interface, configuring [157](#)
 - security [150](#)
 - verifying [163](#)
- static NAT, configuring [157–158](#)
 - inside source address [157](#)
 - outside address [158](#)
- static PAT [162](#)
 - configuration example [162](#)
- static PAT, configuring [159–160](#)
 - port [159–160](#)
- STP [63](#)
 - port channel [63](#)
- subinterfaces [40, 46–47](#)
 - configuring [46](#)
 - configuring bandwidth [47](#)
- svi autostate [10](#)
 - layer 2 [10](#)
- SVI autostate disable [43](#)
- SVI autostate disable, configuring [54](#)
- svi autostate, disabling [25](#)
 - layer 2 [25](#)
- symmetric hashing [68](#)

T

- tunnel interface [109](#)
 - vrf membership, assigning [109](#)
- tunnel interfaces [42, 105](#)
 - configuring based on PBR [105](#)
- tunnel interfaces, creating [101](#)
 - interfaces [101](#)
- tunnel mode [103–104](#)
- tunnel mode ipip [103](#)
- tunnel mode ipv6ipv6 decapsulate-any [103–104](#)

U

- UDLD [4–5](#)
 - aggressive mode [5](#)
 - defined [4](#)
 - nonaggressive mode [5](#)
- UDLD modeA [15](#)
 - configuring [15](#)

Unidirectional Link Detection [4](#)

V

- verifying [56](#)
 - Layer 3 interface configuration [56](#)
- verifying dynamic NAT configuration [172](#)
- VLAN [41](#)
 - interfaces [41](#)
- VLAN interfaces [48](#)
 - configuring [48](#)
- vPC terminology [114](#)
- vPCs [141](#)
 - moving port channels into [141](#)
- VRF [50](#)
 - assigning an interface to [50](#)
- vrf membership, assigning [109](#)
 - tunnel interface [109](#)