Cisco Nexus 3548 Switch NX-OS Interfaces Configuration Guide, Release 6.x

First Published: 2013-04-12

Last Modified: 2018-02-05

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 NONINFRINGEMENT OR ARISING FROM A COURSE OF DEALING, USAGE, OR TRADE PRACTICE.

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

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: www.cisco.com go trademarks. Third-party trademarks mentioned are the property of their respective owners. The use of the word partner does not imply a partnership relationship between Cisco and any other company. (1721R)

© 2017–2018 Cisco Systems, Inc. All rights reserved.

CONTENTS

PREFACE

Preface ix

Audience ix

Document Conventions ix

Documentation Feedback x

CHAPTER 1

New and Changed Information 1

New and Changed Information 1

CHAPTER 2

Configuring Layer 2 Interfaces 3

Information About Ethernet Interfaces 3

Interface Command 3

About 40-Gbps Interface Speed 4

Unidirectional Link Detection Parameter 4

Default UDLD Configuration 5

UDLD Aggressive and Nonaggressive Modes 5

SVI Autostate 5

Cisco Discovery Protocol 6

Default CDP Configuration 6

Error-Disabled State 7

MTU Configuration 7

Debounce Timer Parameters 8

Guidelines and Limitations for Layer 2 Interfaces 8

Configuring Ethernet Interfaces 8

Configuring the UDLD Mode 8

Configuring the Interface Speed 10

Configuring 40-Gigabit Interface Speed 11

Disabling Link Negotiation 12
Disabling SVI Autostate 12
Configuring the CDP Characteristics 13
Enabling or Disabling CDP 14
Enabling the Error-Disabled Detection 15
Enabling the Error-Disabled Recovery 16
Configuring the Error-Disabled Recovery Interval 17
Configuring the Description Parameter 17
Disabling and Restarting Ethernet Interfaces 18
Configuring the Debounce Timer 18
Verifying the Layer 2 Interfaces Configuration 19
Displaying Interface Information 20
Default Physical Ethernet Settings 22
MIBs for Layer 2 Interfaces 22

CHAPTER 3 **Configuring Layer 3 Interfaces 23**
Information About Layer 3 Interfaces 23
Routed Interfaces 23
Subinterfaces 24
VLAN Interfaces 24
Loopback Interfaces 25
Licensing Requirements for Layer 3 Interfaces 25
Guidelines and Limitations for Layer 3 Interfaces 26
Default Settings for Layer 3 Interfaces 26
Configuring Layer 3 Interfaces 26
Configuring a Routed Interface 26
Configuring a Subinterface 27
Configuring the Bandwidth on an Interface 28
Configuring a VLAN Interface 29
Configuring a Loopback Interface 29
Assigning an Interface to a VRF 30
Verifying the Layer 3 Interfaces Configuration 31
Monitoring Layer 3 Interfaces 32
Configuration Examples for Layer 3 Interfaces 33

Related Documents for Layer 3 Interfaces 33
MIBs for Layer 3 Interfaces 34
Standards for Layer 3 Interfaces 34

CHAPTER 4 **Configuring Port Channels** 35
Information About Port Channels 35
Understanding Port Channels 35
Compatibility Requirements 36
Load Balancing Using Port Channels 38
Understanding LACP 39
LACP Overview 39
LACP ID Parameters 40
Channel Modes 40
LACP Marker Responders 41
LACP-Enabled and Static Port Channel Differences 41
LACP Port Channel MinLinks 42
Configuring Port Channels 42
Creating a Port Channel 42
Adding a Port to a Port Channel 43
Configuring Load Balancing Using Port Channels 44
Enabling LACP 46
Configuring the Channel Mode for a Port 46
Configuring LACP Port Channel MinLinks 47
Configuring the LACP Fast Timer Rate 48
Configuring the LACP System Priority and System ID 49
Configuring the LACP Port Priority 50
Verifying Port Channel Configuration 50
Verifying the Load-Balancing Outgoing Port ID 51

CHAPTER 5 **Configuring Virtual Port Channels** 53
Information About vPCs 53
vPC Overview 53
Terminology 54
vPC Terminology 54

vPC Domain 54
Peer-Keepalive Link and Messages 55
Compatibility Parameters for vPC Peer Links 56
Configuration Parameters That Must Be Identical 57
Configuration Parameters That Should Be Identical 58
Viewing Type-1 Inconsistency Check 58
Per-VLAN Consistency Check 59
vPC Auto-Recovery 59
vPC Peer Links 59
vPC Peer Link Overview 60
vPC Number 61
vPC Interactions with Other Features 61
vPC and LACP 61
vPC Peer Links and STP 61
CFSoE 62
vPC Peer Switch 62
Guidelines and Limitations for vPCs 63
Verifying the vPC Configuration 63
Viewing the Graceful Type-1 Check Status 64
Viewing a Global Type-1 Inconsistency 64
Viewing an Interface-Specific Type-1 Inconsistency 66
Viewing a Per-VLAN Consistency Status 67
vPC Default Settings 69
Configuring vPCs 69
Enabling vPCs 69
Disabling vPCs 70
Creating a vPC Domain 70
Configuring a vPC Keepalive Link and Messages 72
Creating a vPC Peer Link 74
Checking the Configuration Compatibility 75
Enabling vPC Auto-Recovery 76
Configuring the Restore Time Delay 77
Excluding VLAN Interfaces from Shutting Down a vPC Peer Link Fails 77
Configuring the VRF Name 78

Moving Other Port Channels into a vPC 79
Manually Configuring a vPC Domain MAC Address 80
Manually Configuring the System Priority 81
Manually Configuring a vPC Peer Switch Role 81
Configuring Layer 3 over vPC 82

CHAPTER 6 **Configuring Static and Dynamic NAT Translation** 85
Network Address Translation Overview 85
Information About Static NAT 86
Dynamic NAT Overview 87
Timeout Mechanisms 88
NAT Inside and Outside Addresses 89
Pool Support for Dynamic NAT 89
Static and Dynamic Twice NAT Overview 90
Licensing Requirements for Static NAT 90
Guidelines and Limitations for Static NAT 91
Restrictions for Dynamic NAT 92
Guidelines and Limitations for Dynamic Twice NAT 92
Configuring Static NAT 93
Enabling Static NAT 93
Configuring Static NAT on an Interface 93
Enabling Static NAT for an Inside Source Address 94
Enabling Static NAT for an Outside Source Address 94
Configuring Static PAT for an Inside Source Address 95
Configuring Static PAT for an Outside Source Address 96
Configuring Static Twice NAT 96
Configuration Example for Static NAT and PAT 98
Example: Configuring Static Twice NAT 98
Verifying the Static NAT Configuration 99
Configuring Dynamic NAT 99
Configuring Dynamic Translation and Translation Timeouts 99
Configuring Dynamic NAT Pool 102
Configuring Source Lists 103
Configuring Dynamic Twice NAT for an Inside Source Address 104

Configuring Dynamic Twice NAT for an Outside Source Address 105
Clearing Dynamic NAT Translations 106
Verifying Dynamic NAT Configuration 106
Verifying NAT Statistics 108
Clearing NAT Statistics 109
Example: Configuring Dynamic Translation and Translation Timeouts 110
Information About VRF Aware NAT 110
Configuring VRF Aware NAT 110

CHAPTER 7 **Configuring IP Event Dampening 113**
IP Event Dampening 113
IP Event Dampening Overview 113
Interface State Change Events 114
Suppress Threshold 114
Half-Life Period 114
Reuse Threshold 114
Maximum Suppress Time 114
Affected Components 115
Route Types 115
Supported Protocols 115
How to Configure IP Event Dampening 115
Enabling IP Event Dampening 115
Verifying IP Event Dampening 116

Preface

The preface contains the following sections:

- Audience, on page ix
- Document Conventions, on page ix
- Documentation Feedback, on page x

Audience

This publication is for network administrators who configure and maintain Cisco Nexus devices.

Document Conventions

Note As part of our constant endeavor to remodel our documents to meet our customers' requirements, we have modified the manner in which we document configuration tasks. As a result of this, you may find a deviation in the style used to describe these tasks, with the newly included sections of the document following the new format.

Command descriptions use the following conventions:

Convention	Description
bold	Bold text indicates the commands and keywords that you enter literally as shown.
Italic	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.

Convention	Description
[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.
`variable`	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
`screen font`	Terminal sessions and information the switch displays are in screen font.
`boldface screen font`	Information you must enter is in boldface screen font.
italic screen font	Arguments for which you supply values are in italic screen font.
< >	Nonprinting characters, such as passwords, are in angle brackets.
[ ]	Default responses to system prompts are in square brackets.
!, #	An exclamation point (!) or a pound sign (#) at the beginning of a line of code indicates a comment line.

This document uses the following conventions:

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

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

Documentation Feedback

To provide technical feedback on this document, or to report an error or omission, please send your comments to: .

We appreciate your feedback.

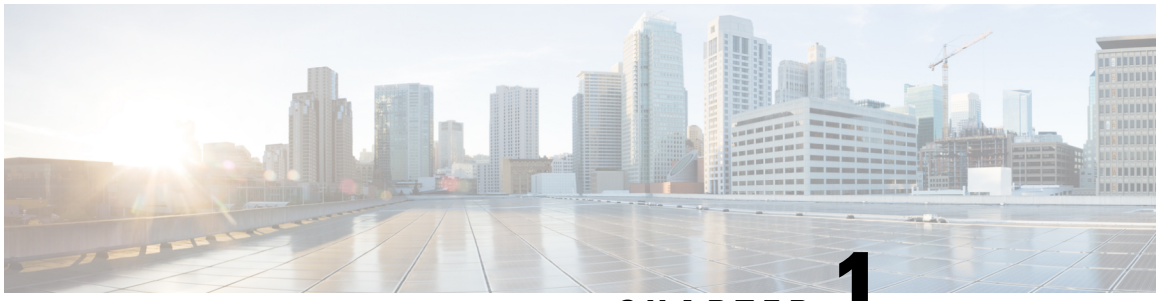

CHAPTER 1

New and Changed Information

This chapter contains the following sections:

- 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 changes made to this guide or all new features in a particular release.

Feature	Description	Added or Changed in Release	Where Documented
NAT Statistics	Added support for NAT Statistics	6.0(2)A8(9)	Guidelines and Limitations for Dynamic Twice NAT, on page 92 Verifying NAT Statistics, on page 108 Clearing NAT Statistics, on page 109
NAT with ECMP	Added ECMP support for NAT	6.0(2)A8(3)	Guidelines and Limitations for Dynamic Twice NAT, on page 92
Layer 3 over vPC	Added support for configuring Layer 3 over vPC.	6.0(2)A8(1)	Configuring Layer 3 over vPC, on page 82
VPC	Added support for VPC in warp mode.	6.0(2)A8(1)	Guidelines and Limitations for vPCs, on page 63 Viewing Type-1 Inconsistency Check, on page 58

Feature	Description	Added or Changed in Release	Where Documented
VRF Aware NAT	This feature is supported.	6.0(2)A7(1)	Information About VRF Aware NAT, on page 110
IP Event Dampening	This feature is supported.	6.0(2)A7(1)	IP Event Dampening Overview, on page 113
Link Debounce Timer	This feature is supported.	6.0(2)A7(1)	Debounce Timer Parameters, on page 8
Twice NAT	This feature was introduced.	6.0(2)A4(1)	Static and Dynamic Twice NAT Overview, on page 90
TCP State Enhancements	**syn-timeout** and **finrst-timeout** timers were introduced.	6.0(2)A4(1)	Timeout Mechanisms, on page 88
Virtual Port Channels	Added information to configure virtual port channels (vPC).	6.0(2)A3(1)	Configuring Virtual Port Channels, on page 53
NAT Enhancements	Introduced pool support for dynamic NAT. The support for PAT now includes the use of the global address pool.	6.0(2)A3(1)	Pool Support for Dynamic NAT, on page 89
40-Gigabit Interface Support	You can enable 40-Gigabit interface speed on the first port of a group of four adjacent ports. Up to 12 interfaces in 40-Gigabit speed are supported.	6.0(2)A1(1)	Configuring Layer 2 Interfaces, on page 3
Static and Dynamic NAT Translation	Dynamic Network Address Translation (NAT) translates a group of real IP addresses into mapped IP addresses that are routable on a destination network.	6.0(2)A1(1)	Static and Dynamic Twice NAT Overview, on page 90

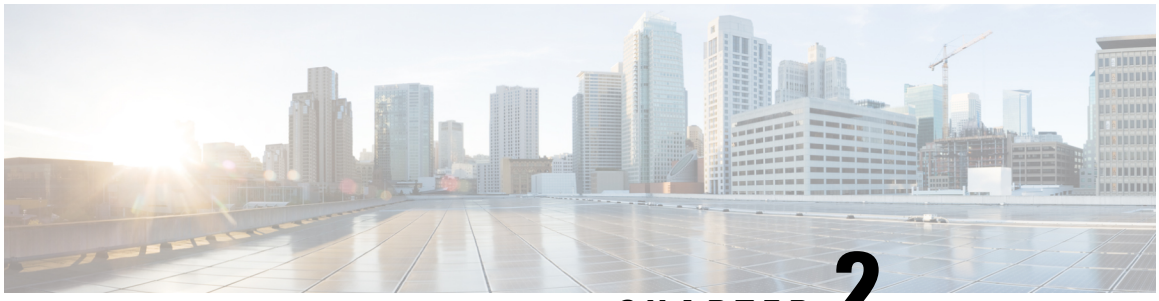

CHAPTER 2

Configuring Layer 2 Interfaces

This chapter contains the following sections:

- Information About Ethernet Interfaces, on page 3
- Guidelines and Limitations for Layer 2 Interfaces, on page 8
- Configuring Ethernet Interfaces, on page 8
- Verifying the Layer 2 Interfaces Configuration, on page 19
- Displaying Interface Information, on page 20
- Default Physical Ethernet Settings , on page 22
- MIBs for Layer 2 Interfaces, on page 22

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.

About 40-Gbps Interface Speed

You can enable 40-Gigabits per second (Gbps) speed on up to 12 interfaces. You enable 40-Gbps speed on the first port of a group of four adjacent ports. For example, you enable 40-Gbps speed on port 1 of port group 1-4, port 5 of port group 5-8, and port 9 of port group 9-12, and so on. The 40-Gbps port numbering is Ethernet interface 1/1, 1/5, 1/9, 1/13, 1/17, and so on.

The configuration is applied to the first port, not on the remaining three ports in the group. The remaining ports act like the ports without an enhanced small form-factor pluggable (SFP+) transceiver inserted. The configuration takes effect immediately. You do not need to reload the switch.

An SFP+ transceiver security check is performed only on the first port of the group.

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

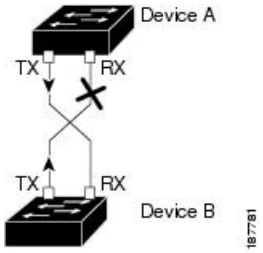

Default UDLD Configuration

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	Enabled

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.

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 inteface 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 reenable 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.

MTU Configuration

The Cisco Nexus device 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.

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.

Guidelines and Limitations for Layer 2 Interfaces

- 40-Gbps Ethernet interfaces do not support the following features:
  - Switched Port Analyzer (SPAN)
  - Encapsulated Remote Switched Port Analyzer (ERSPAN)
  - Warp SPAN
  - Private Virtual Local Area Network (PVLAN)
  - Active buffer monitoring
  - Latency monitoring
  - Link level flow control
  - Precession Time Protocol (PTP)
  - Image downgrade after 40-Gbps interface configuration
  - Configuration rollback
- If you set the 40-Gbps interface speed on an interface and the link is up, the CLI shows the first port as up and the remaining three ports as down. If any of the four links are down, the CLI shows all of the links as down.

Configuring Ethernet Interfaces

Configuring the UDLD Mode

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

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

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

Procedure

	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.
Step 5	switch(config)# **interface** *type slot*/*port*	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** *interface*	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
```

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.

Procedure

	Command or Action	Purpose
Step 1	switch# **configure terminal**	Enters global configuration mode.
Step 2	switch(config)# **interface** *type slot*/*port*	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** *speed*	Sets the speed on the interface. This command can only be applied to a physical Ethernet interface. The *speed* argument can be set to one of the following: • 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 40-Gigabit Interface Speed

Before you begin

To achieve 40-Gbps port speed, each of the four ports in an adjacent port group must have a 10-Gbps SFP installed. All four SFP+ must be capable of 10-Gbps speed and must be the same type of port. By default, all ports are 10-Gbps ports.

Procedure

	Command or Action	Purpose
Step 1	switch# **configure terminal**	Enters global configuration mode.
Step 2	switch(config)# **interface** *type slot*/*port-range*	Enters interface configuration mode for the specified range of interfaces.
Step 3	switch(config-if-rang)# **shut**	Shuts down the range of interfaces that you specified.
Step 4	switch(config-if-rang)# **exit**	Exits the current configuration mode.
Step 5	switch(config-if)# **interface** *type slot*/*port*	Enters interface configuration mode for the interface. You specify the first port in the four adjacent port group to configure that port with 40-Gbps speed. For example, you specify interface 1/1, which is the first port in port group 1/1 through 1/4, to configure that port with 40-Gbps speed. **Note** All four adjacent ports must have 10-Gbps Ethernet SFP transceivers installed.
Step 6	switch(config-if)# **speed 40000**	Sets the speed on the interface for 40 Gbps.
Step 7	switch(config-if)# **no shut**	Brings up the range of interfaces.

Example

This example shows how to set the speed to 40 Gbps on Ethernet interface 1/33:

```
switch# configure terminal
switch(config)# interface ethernet 1/33-36
switch(config-if-rang)# shut
switch(config-if-rang)# exit
switch(config)# interface ethernet 1/33
switch(config-if)# speed 40000
switch(config-if)# no shut
```

Disabling Link Negotiation

You can disable link negotiation using the **no negotiate auto** command. By default, auto-negotiation is enabled on 1-Gigabit ports and disabled on 10-Gigabit ports. The **no negotiate auto** command is supported on 100M port with full duplex setting.

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

Note Auto negotiation configuration is not applicable on 10-Gigabit 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
```

Procedure

	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 ports is enabled. **Note** This command is not applicable for 10GBase-T ports. It should not be used on 10GBase-T ports.

Example

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.

Procedure

	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** *interface-vlan-number*	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** *interface-vlan*	Displays the enabled or disabled Autostate behavior of the SVI.
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 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.

Procedure

	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.
Step 5	(Optional) switch(config)# [**no**] **cdp timer** *seconds*	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.

Procedure

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

	Command or Action	Purpose
Step 2	switch(config)# **interface** *type slot*/*port*	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.

Procedure

	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.

	Command or Action	Purpose
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).

Procedure

	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.

Procedure

	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
```

Configuring the Description Parameter

You can provide textual interface descriptions for the Ethernet ports.

Procedure

	Command or Action	Purpose
Step 1	switch# **configure terminal**	Enters global configuration mode.
Step 2	switch(config)# **interface** *type slot*/*port*	Enters interface configuration mode for the specified interface.
Step 3	switch(config-if)# **description** *test*	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. This information is communicated to other network servers through all dynamic routing protocols. When shut down, the interface is not included in any routing updates.

Procedure

	Command or Action	Purpose
Step 1	switch# **configure terminal**	Enters global configuration mode.
Step 2	switch(config)# **interface** *type slot*/*port*	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 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 debounce feature is available for 10G and 40G interfaces only.

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

Procedure

	Command or Action	Purpose
Step 1	switch# **configure terminal**	Enters global configuration mode.
Step 2	switch(config)# **interface** *type slot/port*	Enters interface configuration mode for the specified interface.
Step 3	switch(config-if)# **link debounce time** *milliseconds*	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
```

Verifying the Layer 2 Interfaces Configuration

Use one of the following commands to verify the configuration:

Command	Purpose
show interface ethernet *slot/port* **brief**	Displays the Layer 2 interface operational status. **Note** If you have 40-Gbps interface speed set on an interface and the link is up, the CLI shows the first port as up and the remaining three ports as down. If any one of the four links are down, the CLI shows all of the links as down.

Displaying Interface Information

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

Command	Purpose
switch# **show interface** *type slot/port*	Displays the detailed configuration of the specified interface.
switch# **show interface** *type slot/port* **capabilities**	Displays detailed information about the capabilities of the specified interface. This option is available only for physical interfaces.
switch# **show interface** *type slot/port* **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
d13-dist-1               mgmt0           148     S I      WS-C2960-24TC  Fas0/9
n5k(FLC12080012)         Eth1/5          8       S I s    N5K-C5020P-BA  Eth1/5
```

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.

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: • 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	

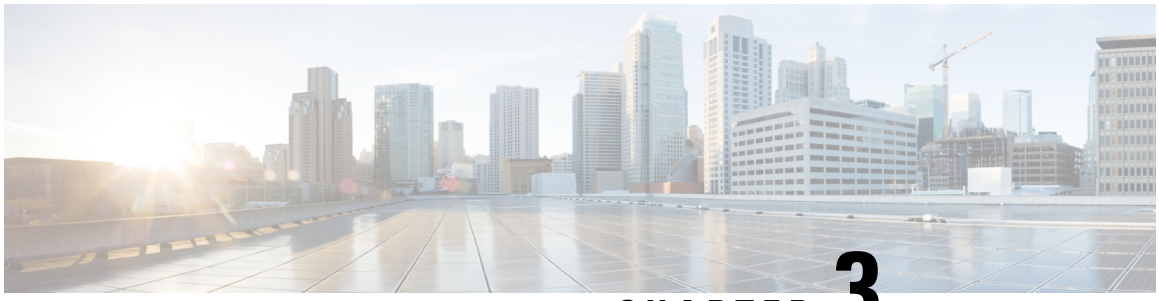

CHAPTER 3

Configuring Layer 3 Interfaces

This chapter contains the following sections:

- Information About Layer 3 Interfaces, on page 23
- Licensing Requirements for Layer 3 Interfaces, on page 25
- Guidelines and Limitations for Layer 3 Interfaces, on page 26
- Default Settings for Layer 3 Interfaces, on page 26
- Configuring Layer 3 Interfaces, on page 26
- Verifying the Layer 3 Interfaces Configuration, on page 31
- Monitoring Layer 3 Interfaces, on page 32
- Configuration Examples for Layer 3 Interfaces, on page 33
- Related Documents for Layer 3 Interfaces, on page 33
- MIBs for Layer 3 Interfaces, on page 34
- Standards for Layer 3 Interfaces, on page 34

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 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

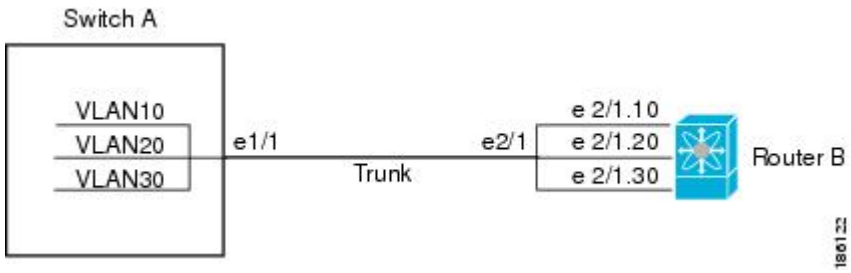

VLAN Interfaces

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 1and VLAN 10 communicates at Layer 3 over VLAN interface 10.

Figure 3: Connecting Two VLANs with VLAN Interfaces

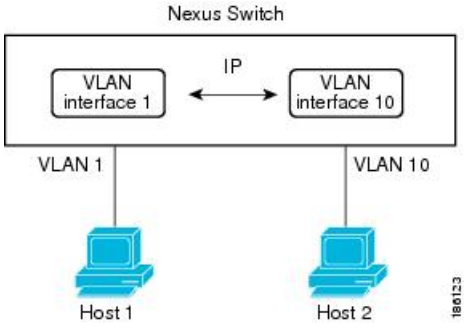

Loopback Interfaces

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.

Licensing Requirements for Layer 3 Interfaces

This feature does not require a license. Any feature not included in a license package is bundled with the Cisco NX-OS system images and is provided at no extra charge to you. For a complete explanation of the Cisco NX-OS licensing scheme, see the *Cisco NX-OS Licensing Guide*.

Guidelines and Limitations for Layer 3 Interfaces

Layer 3 interfaces have the following configuration guidelines and limitations:

- 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.
-
-

Default Settings for Layer 3 Interfaces

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

Configuring Layer 3 Interfaces

Configuring a Routed Interface

Procedure

	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(conifg-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***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.

	Command or Action	Purpose
		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.

Procedure

	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 *slot* is from 1 to 255. The range for the *port* is from 1 to 128.
Step 3	switch(config-if)# **ip 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 *vlan-id* is from 2 to 4093.

	Command or Action	Purpose
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.

Procedure

	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. The range for the *slot* is from 1 to 255. The range for the *port* is from 1 to 128.
Step 3	switch(conifg-if)# **bandwidth** [*value* | **inherit** [*value*]]	Configures the bandwidth parameter for a routed interface, port channel, or subinterface, as follows: • *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

Procedure

	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** *number*	Creates a VLAN interface. The *number* range is from 1 to 4094.
Step 4	switch(config-if)# **ip address** *ip-address/length*	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** *number*	Displays the VLAN interface statistics. The *number* 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
```

Configuring a Loopback Interface

Before you begin

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

Procedure

	Command or Action	Purpose
Step 1	switch# **configure terminal**	Enters global configuration mode.
Step 2	switch(config)# **interface loopback** *instance*	Creates a loopback interface. The *instance* range is from 0 to 1023.
Step 3	switch(config-if)# **ip address** *ip-address/length*	Configures an IP address for this interface.
Step 4	(Optional) switch(config-if)# **show interface loopback** *instance*	Displays the loopback interface statistics. The *instance* 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.

Procedure

	Command or Action	Purpose
Step 1	switch# **configure terminal**	Enters global configuration mode.
Step 2	switch(config)# **interface** *interface-typenumber*	Enters interface configuration mode.
Step 3	switch(conifg-if)#**vrf member** *vrf-name*	Adds this interface to a VRF.
Step 4	switch(config-if)# **ip***ip-address/length*	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-if)# **show vrf** [*vrf-name*] **interface** *interface-type number*	Displays VRF information.
Step 6	(Optional) switch(config-if)# **show interfaces**	Displays the Layer 3 interface statistics.

	Command or Action	Purpose
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
```

Verifying the Layer 3 Interfaces Configuration

Use one of the following commands to verify the configuration:

Command	Purpose
show interface ethernet *slot/port*	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 *slot/port* **brief**	Displays the Layer 3 interface operational status.
show interface ethernet *slot/port* **capabilities**	Displays the Layer 3 interface capabilities, including port type, speed, and duplex.
show interface ethernet *slot/port* **description**	Displays the Layer 3 interface description.
show interface ethernet *slot/port* **status**	Displays the Layer 3 interface administrative status, port mode, speed, and duplex.
show interface ethernet *slot/port.number*	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 *channel-id.number*	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 *number*	Displays the loopback interface configuration, status, and counters.
show interface loopback *number* **brief**	Displays the loopback interface operational status.

Command	Purpose
show interface loopback *number* **description**	Displays the loopback interface description.
show interface loopback *number* **status**	Displays the loopback interface administrative status and protocol status.
show interface vlan *number*	Displays the VLAN interface configuration, status, and counters.
show interface vlan *number* **brief**	Displays the VLAN interface operational status.
show interface vlan *number* **description**	Displays the VLAN interface description.
show interface vlan *number* **status**	Displays the VLAN interface administrative status and protocol status.

Monitoring Layer 3 Interfaces

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

Command	Purpose
show interface ethernet *slot*/*port* **counters**	Displays the Layer 3 interface statistics (unicast, multicast, and broadcast).
show interface ethernet *slot*/*port* **counters brief**	Displays the Layer 3 interface input and output counters.
show interface ethernet *slot*/*port* **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 *slot*/*port* **counters error**	Displays the Layer 3 interface input and output errors.
show interface ethernet *slot*/*port* **counters snmp**	Displays the Layer 3 interface counters reported by SNMP MIBs. You cannot clear these counters.
show interface ethernet *slot*/*port***.***number* **counters**	Displays the subinterface statistics (unicast, multicast, and broadcast).
show interface port-channel *channel-id***.***number* **counters**	Displays the port-channel subinterface statistics (unicast, multicast, and broadcast).
show interface loopback *number* **counters**	Displays the loopback interface input and output counters (unicast, multicast, and broadcast).
show interface loopback *number* **counters detailed** [**all**]	Displays the loopback interface statistics. You can optionally include all 32-bit and 64-bit packet and byte counters (including errors).

Command	Purpose
show interface loopback *number* **counters errors**	Displays the loopback interface input and output errors.
show interface vlan *number* **counters**	Displays the VLAN interface input and output counters (unicast, multicast, and broadcast).
show interface vlan *number* **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 *counters* **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)# ip address 192.0.2.1/8

switch(config-if)# copy running-config startup-config
```

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
```

Related Documents for Layer 3 Interfaces

Related Topics	Document Title
Command syntax	Cisco Nexus 3548 Switch NX-OS Interfaces Command Reference
IP	"Configuring IP" chapter in the *Cisco Nexus 3548 Switch NX-OS Unicast Routing Configuration Guide*

Related Topics	Document Title
VLAN	"Configuring VLANs" chapter in the *Cisco Nexus 3548 Switch NX-OS Layer 2 Switching Configuration Guide*

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.

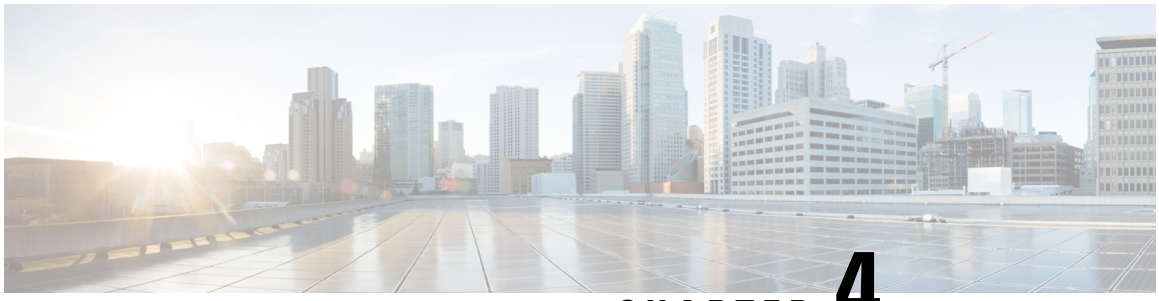

CHAPTER 4

Configuring Port Channels

This chapter contains the following sections:

- Information About Port Channels, on page 35
- Configuring Port Channels, on page 42
- Verifying Port Channel Configuration, on page 50
- Verifying the Load-Balancing Outgoing Port ID , on page 51

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 39

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 38.

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 default port-channel load balance parameter for all Layer 2, Layer 3 and Layer 4 frames is the source and destination IP addresses only. This criteria can be changed using the **port-channel load-balance ethernet** command. Load balancing based only on MAC addresses occurs only when the Ethertype is not set to 0800 in the Layer 2 packet header. When the Ethertype is 0800, then load balancing takes place based on IP addresses in the IP packet header irrespective of the port-channel load balancing parameters defined on the command line. In addition, if the packet has Ethertype 0800, and it does not have a valid IP header, the packet will be flagged for a parsing error and subsequently dropped.

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

Configuration	Layer 2 Criteria	Layer 3 Criteria	Layer 4 Criteria
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.

The unicast and multicast traffic is load-balanced across port-channel links based on the configured load-balancing algorithm shown in the **show port-channel load-balancing** command output.

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

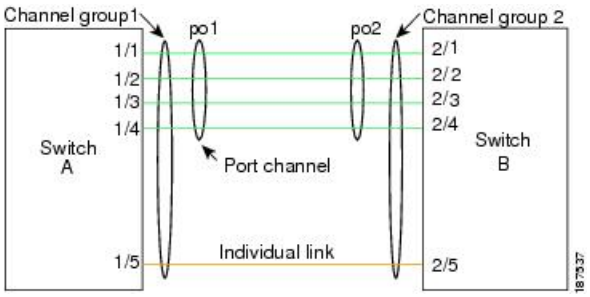

With LACP, just like with static port channels, you can bundle up to 16 interfaces in a channel group.

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.
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	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. 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.

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: • Active • Passive	Can only be On.

LACP Port Channel MinLinks

A port channel aggregates similar ports to provide increased bandwidth in a single manageable interface. The MinLinks feature allows you to define the minimum number of interfaces from a LACP bundle that must fail before the port channel goes down.

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.

Note The MinLinks 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.

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.

Procedure

	Command or Action	Purpose
Step 1	switch# **configure terminal**	Enters global configuration mode.
Step 2	switch(config)# **interface port-channel** *channel-number*	Specifies the port-channel interface to configure, and enters the interface configuration mode. The range is from 1 to 4096. Cisco

	Command or Action	Purpose
		NX-OS automatically creates the channel group if it does not already exist.
Step 3	switch(config)# **no interface port-channel** *channel-number*	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.

Procedure

	Command or Action	Purpose
Step 1	switch# **configure terminal**	Enters global configuration mode.
Step 2	switch(config)# **interface** *type slot*/*port*	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** *vlan-id* | **native vlan** *vlan-id*}	Configures necessary parameters for a trunk port.
Step 5	switch(config-if)# **channel-group** *channel-number*	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.

Note For load-balancing FC traffic across SAN PO members in Nexus 5672UP-16G switch, the **port-channel load-balance ethernet** command is not needed. The load-balancing happens by default.

Procedure

	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-mac** | **destination-port** | **source-dest-ip** | **source-dest-mac** | **source-dest-port** | **source-ip** | **source-mac** | **source-port**] | **crc-poly**}	Specifies the load-balancing algorithm for the device. The range depends on the device. The default is **source-dest-ip**.

	Command or Action	Purpose
		Note • source-dest-ip-only • source-dest-port-only • source-dest-ip • source-dest-port • source-dest-ip-gre The Cisco Nexus 5500 Platform switches support 8 hash polynomials that can be used for compression on the hash-parameters. Depending on variations in the hash parameters for egress traffic flows from a port channel, different polynomials could provide different load distribution results. The default hash polynomial is CRC8a. The variable can be configured as follows: • CRC8a • CRC8b • CRC8c • CRC8d • CRC8e • CRC8f • CRC8g
Step 3	(Optional) switch(config)# **no port-channel load-balance ethernet**	Restores the default load-balancing algorithm of source-dest-ip.
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
```

Enabling LACP

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

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

Procedure

	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.

Procedure

	Command or Action	Purpose
Step 1	switch# **configure terminal**	Enters global configuration mode.
Step 2	switch(config)# **interface** *type slot*/*port*	Specifies the interface to configure, and enters the interface configuration mode.
Step 3	switch(config-if)# **channel-group** *channel-number* [**force**] [**mode** {**on** | **active** | **passive**}]	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.

	Command or Action	Purpose
		force—Specifies that the LAN port be forcefully added to the channel group. **mode**—Specifies the port channel mode of the interface. **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. **on**—(Default mode) Specifies that all port channels that are not running LACP remain in this mode. **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. When you run port channels with no associated protocol, the channel mode is always on.
Step 4	switch(config-if)# **no channel-group** *number* **mode**	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
```

This example shows how to forcefully add an interface to the channel group 5:

```
switch(config)# interface ethernet 1/1
switch(config-if)# channel-group 5 force
switch(config-if)#
```

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.

Procedure

	Command or Action	Purpose
Step 1	switch# **configure terminal**	Enters global configuration mode.
Step 2	switch(config)# **interface port-channel** *number*	Specifies the interface to configure and enters interface configuration mode.
Step 3	switch(config-if)# [**no**] **lacp min-links** *number*	Specifies the port channel interface to configure the number of minimum links and enters the interface configuration mode. The default value for *number* is 1. The range is from 1 to 16. Use the **no** form of this command to disable this feature.
Step 4	(Optional) switch(config)# **show running-config interface port-channel** *number*	Displays the port channel MinLinks configuration.

Example

This example shows how to configure the minimum number of port channel interfaces on module 3:

```
switch# configure terminal
switch(config) # interface port-channel 3
switch(config-if) # lacp min-links 3
switch(config-if) #
```

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.

Procedure

	Command or Action	Purpose
Step 1	switch# **configure terminal**	Enters global configuration mode.
Step 2	switch(config)# **interface** *type slot/port*	Specifies the interface to configure and enters the interface configuration mode.
Step 3	switch(config-if)# **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.

Procedure

	Command or Action	Purpose
Step 1	switch# **configure terminal**	Enters global configuration mode.
Step 2	switch(config)# **lacp system-priority** *priority*	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.

Procedure

	Command or Action	Purpose
Step 1	switch# **configure terminal**	Enters global configuration mode.
Step 2	switch(config)# **interface** *type slot*/*port*	Specifies the interface to configure, and enters the interface configuration mode.
Step 3	switch(config-if)# **lacp port-priority** *priority*	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*channel-number*	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.

Command	Purpose
show lacp {**counters** | **interface** *type slot*/*port* | **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** *channel-number*]	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.
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. Based on the system—Source and destination IP, and MAC addresses are used for port channel load balance calculation. By default the **source-destination-ip** address is used. Based on protocol—For non-IP traffic, only the source and destination MAC address is used, and for IP traffic, both source and destination IP and MAC addresses are used for the port channel load balance calculation.

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.

Note In warp mode, the output contains two destination ports: one when there is no match in the warp table and one when there is a match in the warp table. A Layer 2 port match means that the source and destination MAC addresses are learned in the MAC table whereas a Layer 3 port match means the IP address is resolved.

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** *port-channel-id* **src-interface** *source-interface* **vlan** vlan-id **dst-ip src-ip dst-mac src-mac l4-src-port** *port-id* **l4-dst-port** *port-id* **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
```

Example

This example shows the output of the **port-channel load-balance** command while the device is in warp mode:

```
switch# show port-channel load-balance forwarding-path interface port-channel 1 src-interface
 ethernet 1/6 vlan 1 src-ip 1.1.1.1 dst-ip 2.2.2.2
Missing params will be substituted by 0's.
Load-balance Algorithm on switch: source-dest-ip
        Outgoing port id (no cache hit): Ethernet1/29
        Outgoing port id (cache hit): Ethernet1/32
Param(s) used to calculate load-balance:
        dst-ip:   2.2.2.2
        src-ip:   1.1.1.1
        dst-mac:  0000.0000.0000
        src-mac:  0000.0000.0000
        VLAN: 1
```

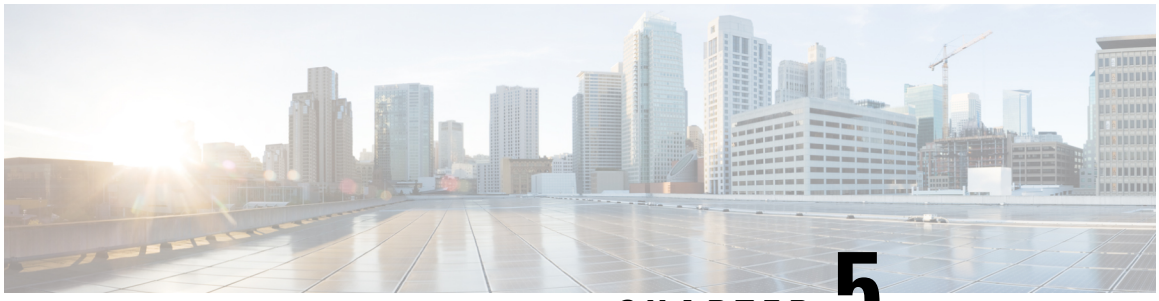

CHAPTER 5

Configuring Virtual Port Channels

This chapter contains the following sections:

- Information About vPCs, on page 53
- Guidelines and Limitations for vPCs, on page 63
- Verifying the vPC Configuration, on page 63
- vPC Default Settings, on page 69
- Configuring vPCs, on page 69

Information About vPCs

vPC Overview

A virtual port channel (vPC) allows links that are physically connected to two different 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.

You can configure a hold-timeout and a timeout value simultaneously.

Hold-timeout value—The hold-timeout value range is between 3 to 10 seconds, with a default value of 3 seconds. This timer starts when the vPC peer link goes down. The purpose of the hold-timeout period is to prevent false-positive cases.

If you configure a hold-timeout value that is lower than the timeout value, then the vPC system ignores vPC peer-keepalive messages for the hold-timeout period and considers messages for the reminder of the timeout period. If no keepalive message is received for this period, the vPC secondary device takes over the role of the primary device. For example, if the hold-timeout value is 3 seconds and the timeout value is 5 seconds, for the first 3 seconds vPC keepalive messages are ignored (such as, when accommodating a supervisor failure for a few seconds after peer link failure) and keepalive messages are considered for the remaining timeout period of 2 seconds. After this period, the vPC secondary device takes over as the primary device, in case there is no keep alive message.

Timeout value—The timeout value range is between 3 to 20 seconds, with a default value of 5 seconds. This timer starts at the end of the hold-timeout interval. If you configure a timeout value that is lower than or equal to the hold-timeout value, then the timeout duration is initiated after the hold-timeout period. For example, if the timeout value is 3 seconds and the hold-timeout value is 5 seconds, the timeout period starts after 5 seconds.

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,                       [(fa0,
                                                     0-23-4-ee-be-64, 8458,
0-23-4-ee-be-64, 8458,
                                                     0, 0), (8000,               0, 0),
(8000,
                                                     f4-4e-5-84-5e-3c, 457,
f4-4e-5-84-5e-3c, 457,
                                                     0, 0)]                      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.
- 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.

Viewing Type-1 Inconsistency Check

After you have configured the Virtual Port Channels (vPC) peer link on both vPC peer switches, check that the configurations are consistent on all vPC interfaces. Starting with Cisco NX-OS Release 6.0(2)A8(1), vPC in warp mode is supported.

Note You must ensure that both the vPC peers are in the same forwarding mode. In case of a forwarding mode mismatch, vPCs are suspended.

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,
                                  0, 0)                  0, 0)
Network Qos (Pause)         2     (F, F, F, F, F, F, F,  (F, F, F, F, F, F,F,
                                  F)                     F)
Network Qos (WRED)          2     (F, F, F, F, F, F, F,  (F, F, F, F, F, F,F,
                                  F)                     F)
```

```
Network Qos (ECN)           2     (F, F, F, F, F, F, F,   (F, F, F, F, F, F,F,
                                  F)                      F)
Output Queuing (Bandwidth)  2     (100, 0, 0, 0, 0, 0,    (100, 0, 0, 0, 0, 0,
                                  0, 0)                   0, 0)
Output Queuing (Absolute    2     (F, F, F, F, F, F, F,   (F, F, F, F, F, F,F,
Priority)                         F)                      F)
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
HW profile Forwarding Mode  1     warp                    warp

<<<<<<<<<< Both Local and remote VPC have same forwarding mode.
IGMP Snooping Group-Limit   2     8190                    8190
Interface-vlan admin up     2     10                      10
Interface-vlan routing      2     10                      10
capability
Allowed VLANs               -     10                      10
Local suspended VLANs       -     -                       -
```

-

Per-VLAN Consistency Check

Some 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

The vPC auto-recovery feature re-enables vPC links in the following scenarios:

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.

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 reenable 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.

vPC Peer Switch

The vPC peer switch feature was added to address performance concerns around STP convergence. This feature allows a pair of Cisco Nexus 3500 Series switches to appear as a single STP root in the Layer 2 topology. This eliminates the need to pin the STP root to the vPC primary switch and improves vPC convergence if the vPC primary switch fails.

To avoid loops, the vPC peer link is excluded from the STP computation. In vPC peer switch mode, STP BPDUs are sent from both vPC peer devices to avoid issues related to STP BPDU timeout on the downstream switches, which can cause traffic disruption.

vPC peer switch can be used with the pure peer switch topology in which the devices all belong to the vPC.

Note Peer-switch is supported on networks that use vPC, and STP-based redundancy is not supported. If the vPC peer-link fails in a hybrid peer-switch configuration, you can lose traffic. In this scenario, the vPC peers use the same STP root ID as well as the same bridge ID. The access switch traffic is split in two with half going to the first vPC peer and the other half to the second vPC peer. With peer link failure, there is no impact to the north/south traffic but the east/west traffic is lost.

Guidelines and Limitations for vPCs

vPCs have the following configuration guidelines and limitations:

- vPC is not qualified with IPv6.
- VPC is now supported in Warp mode on the Cisco Nexus 3500 Series platform.
- 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.
- 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.

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.

Command	Purpose
switch# **show port-channel capacity**	Displays how many EtherChannels are configured and how many are still available on the switch.
switch# **show running-config vpc**	Displays running configuration information for vPCs.
switch# **show vpc brief**	Displays brief information on the vPCs.
switch# **show vpc consistency-parameters**	Displays the status of those parameters that must be consistent across all vPC interfaces.
switch# **show vpc peer-keepalive**	Displays information on the peer-keepalive messages.
switch# **show vpc role**	Displays the peer status, the role of the local switch, the vPC system MAC address and system priority, and the MAC address and priority for the local vPC switch.
switch# **show vpc statistics**	Displays statistics on the vPCs. **Note** This command displays the vPC statistics only for the vPC peer device that you are working on.

For 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

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

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
```

```
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

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

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

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 6: 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.

Procedure

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

	Command or Action	Purpose
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.

Procedure

	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.

Procedure

	Command or Action	Purpose
Step 1	switch# **configure terminal**	Enters global configuration mode.
Step 2	switch(config)# **vpc domain** *domain-id*	Creates a vPC domain on the switch, and enters the vpc-domain configuration mode. There is no default *domain-id* ; the range is from 1 to 1000. **Note** You can also use the **vpc domain** command to enter the vpc-domain configuration mode for an existing vPC domain.
Step 3	switch(config-vpc-domain)# **fast-convergence**	Enables the vPC optimizations feature. Use the **[no] fast-convergence** command to disable the vPC optimizations feature. The CLI should be enabled on both the vPC peers to achieve fast-convergence.
Step 4	(Optional) switch# **show vpc brief**	Displays brief information about each vPC domain.
Step 5	(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
```

This example shows how to enforce the global level type-2 consistency check for the fast-convergence configuration.

```
switch# show vpc consistency-parameters global

    Legend:
        Type 1 : vPC will be suspended in case of mismatch

Name                        Type  Local Value            Peer Value
-------------               ----  ---------------------- -----------------------
Vlan to Vn-segment Map      1     No Relevant Maps       No Relevant Maps
QoS                         2     ([], [], [], [], [],   ([], [], [], [], [],
                                  [], [], [])            [], [], [])
Network QoS (MTU)           2     (1538, 0, 0, 0, 0, 0,  (1538, 0, 0, 0, 0, 0,
```

```
                                    0, 0)                 0, 0)

 ------------------------------------------------------------------------------
 ------------------------------------------------------------------------------
VTP pruning status          2     Disabled              Disabled
IGMP Snooping Group-Limit   2     8000                  8000
Fast Convergence            2     Enable                Enable
Interface-vlan admin up     2     101-120
Interface-vlan routing      2     1,101-120             1
capability
Allowed VLANs               -     -                     -
Local suspended VLANs       -     -                     -
```

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.

Procedure

	Command or Action	Purpose
Step 1	switch# **configure terminal**	Enters global configuration mode.
Step 2	switch(config)# **vpc domain** *domain-id*	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** *ipaddress* [**hold-timeout** *secs* | **interval** *msecs* {**timeout** *secs*} | **precedence**	Configures the IPv4 address for the remote end of the vPC peer-keepalive link.

	Command or Action	Purpose
	{*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**}]	**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** *ipaddress* **source** *ipaddress*	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

Procedure

	Command or Action	Purpose
Step 1	switch# **configure terminal**	Enters global configuration mode.
Step 2	switch(config)# **interface port-channel** *channel-number*	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.

	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)# 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.

Procedure

	Command or Action	Purpose
Step 1	switch# **show vpc consistency-parameters** {**global**|**interface port-channel***channel-number*}	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

Procedure

	Command or Action	Purpose
Step 1	switch# **configure terminal**	Enters global configuration mode.
Step 2	switch(config)# **vpc domain** *domain-id*	Enters vpc-domain configuration mode for an existing vPC domain.
Step 3	switch(config-vpc-domain)# **auto-recovery reload-delay** *delay*	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

YouYou 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.

Procedure

	Command or Action	Purpose
Step 1	switch# **configure terminal**	Enters global configuration mode.
Step 2	switch(config)# **vpc domain** *domain-id*	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** *time*	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)#
```

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.

Procedure

	Command or Action	Purpose
Step 1	switch# **configure terminal**	Enters global configuration mode.
Step 2	switch(config)# **vpc domain** *domain-id*	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** *range*	Specifies the VLAN interfaces that should remain up when a vPC peer-link is lost. range—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.

Procedure

	Command or Action	Purpose
Step 1	switch# **ping** *ipaddress* **vrf** *vrf-name*	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 specifiy 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.

Procedure

	Command or Action	Purpose
Step 1	switch# **configure terminal**	Enters global configuration mode.
Step 2	switch(config)# **interface port-channel** *channel-number*	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** *number*	Configures the selected port channel into the vPC to connect to the downstream switch. The range is from 1 to 4096. The vPC *number* 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.
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.

Procedure

	Command or Action	Purpose
Step 1	switch# **configure terminal**	Enters global configuration mode.
Step 2	switch(config)# **vpc domain** *domain-id*	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 *domain-id*; the range is from 1 to 1000.
Step 3	switch(config-vpc-domain)# **system-mac** *mac-address*	Enters the MAC address that you want for the specified vPC domain in the following format: aaaa.bbbb.cccc.
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.

Procedure

	Command or Action	Purpose
Step 1	switch# **configure terminal**	Enters global configuration mode.
Step 2	switch(config)# **vpc domain** *domain-id*	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 *domain-id*; the range is from 1 to 1000.
Step 3	switch(config-vpc-domain)# **system-priority** *priority*	Enters the system priority that you want for the specified vPC domain. The range of values is from 1 to 65535. The default value is 32667.
Step 4	(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.

Procedure

	Command or Action	Purpose
Step 1	switch# **configure terminal**	Enters global configuration mode.
Step 2	switch(config)# **vpc domain** *domain-id*	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 *domain-id*; the range is from 1 to 1000.
Step 3	switch(config-vpc-domain)# **role priority** *priority*	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

Beginning with Cisco NX-OS Release 6.0(2)A8(1) and later releases, you can configure vPC peer devices to act as the gateway for packets that are destined to the vPC peer device's MAC address.

When you attach a Layer 3 device to a vPC domain, the peering of routing protocols using a VLAN carried out on the vPC peer-link is supported with the following requirements:

- Ensure that you have enabled the vPC feature.
- Ensure that you are in the correct VDC (or use the switchto vdc command).
- Ensure that you enable peer-gateway and peer-routing on Layer 3 over vPC on both the peers.
- Ensure that the peer link is up

If routing protocol adjacencies are needed between vPC peer devices and a generic Layer 3 device, you must use physical routed interfaces for the interconnection. Use of the vPC peer-gateway feature does not change this requirement.

- Ensure that you have enabled the vPC feature.
- Ensure that you are in the correct VDC (or use the switchto vdc command).
- Peer-gateway and peer-routing on Layer 3 over vPC are enabled on both the peers.
- Ensure that the peer link is up

Procedure

	Command or Action	Purpose
Step 1	switch# **configure terminal**	Enters global configuration mode.
Step 2	switch(config)#**vpc domain** *domain-id*	Creates a vPC domain on the device and enters the vpc-domain configuration mode for configuration purposes. There is no default; the range is from 1 to 1000.
Step 3	switch(config-vpc-domain)# **peer-gateway**	Enables Layer 3 forwarding for packets destined to the peer's gateway MAC address.
Step 4	switch(config-vpc-domain)# **layer3 peer-router**	Enables the Layer 3 device to form peering adjacency with both peers. **Note** Configure this command in both the peers.
Step 5	switch(config-vpc-domain)#**exit**	Exits vpc-domain configuration mode.
Step 6	(Optional) switch# **show vpc brief**	Displays brief information about each vPC domain. **Note** 'Operational Layer3 Peer-router' field will be shown as enabled only when **layer3 peer-router** is configured on the both the vPC nodes.
Step 7	(Optional) switch# **copy running-config startup-config**	Copies the running configuration to the startup configuration.

Example

The following example shows how to configure a Layer 3 over vPC:

```
switch# configure terminal
switch(config)# vpc domain 2
switch(config-vpc-domain)# peer-gateway
```

```
switch(config-vpc-domain)# layer3 peer-router
switch(config-vpc-domain)# exit
switch(config)#
```

The following example shows how to verify if the Layer 3 over vPC is configured:

```
switch(config)# show vpc brief
vPC domain id : 2
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 : 7
Peer Gateway : Enabled
Peer gateway excluded VLANs : -
Dual-active excluded VLANs : 502
Graceful Consistency Check : Enabled
Operational Layer3 Peer-router : Enabled
Auto-recovery status : Disabled

vPC Peer-link status
---------------------------------------------
id Port Status Active vlans
---------------------------------------------
 1       Po300        up 1,300,400-403,500-503

vPC Status
----------------------------------------------------------
id  Port   Status  Consistency  Reason  Active vlans
----------------------------------------------------------
1   Po400   up     success     success   400
2   Po500   up     success     success   500
3   Po401   up     success     success   401
4   Po402   up     success     success   402
5   Po403   up     success     success   1
6   Po501   up     success     success   501
7   Po502   up     success     success   502

switch(config)#
```

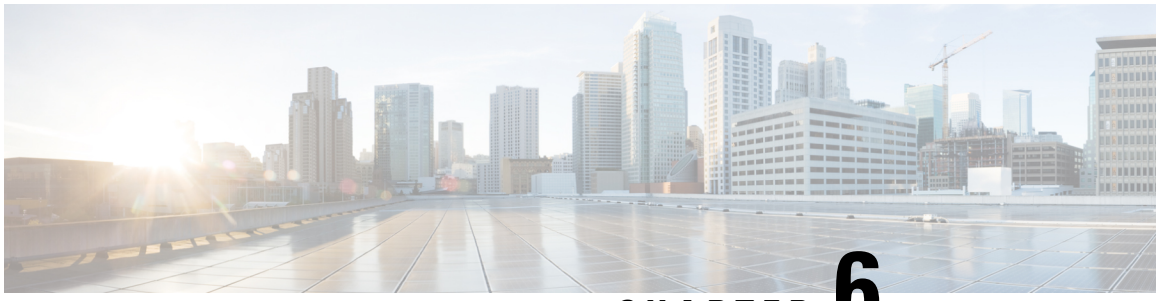

CHAPTER 6

Configuring Static and Dynamic NAT Translation

This chapter contains the following sections:

- Network Address Translation Overview, on page 85
- Information About Static NAT, on page 86
- Dynamic NAT Overview, on page 87
- Timeout Mechanisms, on page 88
- NAT Inside and Outside Addresses, on page 89
- Pool Support for Dynamic NAT, on page 89
- Static and Dynamic Twice NAT Overview, on page 90
- Licensing Requirements for Static NAT, on page 90
- Guidelines and Limitations for Static NAT, on page 91
- Restrictions for Dynamic NAT, on page 92
- Guidelines and Limitations for Dynamic Twice NAT, on page 92
- Configuring Static NAT, on page 93
- Configuring Dynamic NAT, on page 99
- Information About VRF Aware NAT, on page 110
- Configuring VRF Aware NAT, on page 110

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 5: Static NAT

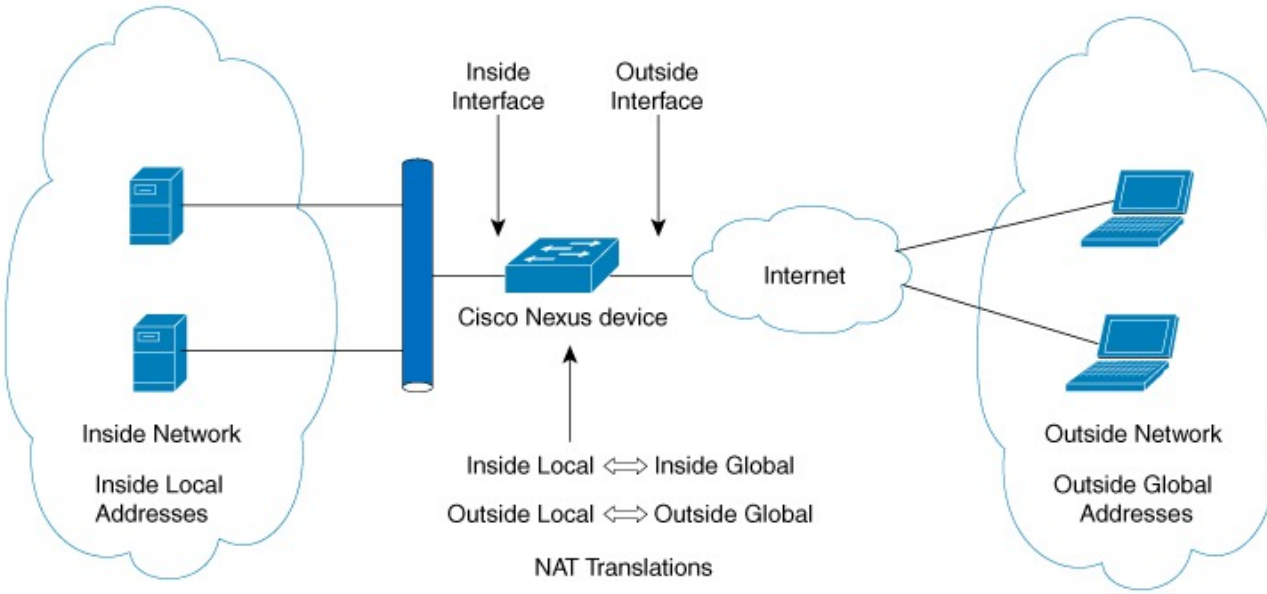

These are key terms to help you understand static NAT:

- 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 avkailable 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. In Cisco NX-OS Release 6.0(2)A1(1e), the minimum value of the sampling-timeout in the **ip nat translation sampling-timeout** command was reduced from 30 minutes to 15 minutes.

Timeout of a dynamic NAT translation involves both the sampling-timeout value and the TCP or UDP timeout value. The sampling-timeout specifies the time after which the device checks for dynamic translation activity. It has a default value of 12 hours. All the other timeouts start only after the sample-timeout times out. After the sampling-timeout, the device inspects the packets that are hitting this translation. The checking happens for the TCP or UDP timeout period. If there are no packets for the TCP or UDP timeout period, the translation is cleared. If activity is detected on the translation, then the checking is stopped immediately and a sampling-timeout period begins.

After waiting for this new sampling-timeout period, the device checks for dynamic translation activity again. During an activity check the TCAM sends a copy of the packet that matches the dynamic NAT translation to the CPU. If the Control Plane Policing (CoPP) is configured at a low threshold, the TCP or UDP packets might not reach the CPU, and the CPU considers this as inactivity of the NAT translation.

Dynamic NAT supports Port Address Translation (PAT) and access control lists (ACLs). PAT, also known as overloading, is a form of dynamic NAT that maps multiple unregistered IP addresses to a single registered IP address by using different ports. Your NAT configuration can have multiple dynamic NAT translations with same or different ACLs. However, for a given ACL, only one interface can be specified.

Timeout Mechanisms

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. Cisco NX-OS Release 6.0(2)A4(1) adds support for **syn-timeout** and **finrst-timeout**. 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.

- **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.

  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.

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

Cisco NX-OS Release 6.0(2)A3(1) introduces 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.

Licensing Requirements for Static NAT

This table shows the licensing requirements for static NAT.

Product	License Requirement
Cisco NX-OS	Static NAT requires a LAN Base license and an Algo Boost license. For a complete explanation of the Cisco NX-OS licensing scheme and how to obtain and apply licenses, see the Cisco NX-OS Licensing Guide. **Note** Make sure the LAN Base Services license is installed on the switch to enable Layer 3 interfaces.

Guidelines and Limitations for Static NAT

Static NAT has the following configuration guidelines and limitations:

- NAT supports up to 1024 translations which include both static and dynamic NAT.
- Cisco Nexus 3500 Series switches do not support static and dynamic NAT on vPC topology.
- 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.
- By default, NAT can go up to 127 translations with 256 TCAM entries. If you need more NAT translations, you need to reduce the TCAM region allocation in other areas and then increase the NAT TCAM region using the **hardware profile tcam region nat** command.
- HSRP and VRRP are supported on NAT inside address and not on NAT outside addresses.
- 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.
- 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.
- Only one of the following features can be enabled on an interface at a time. If more than one of these features is enabled on an interface, only the feature that is enabled last will work:

- NAT
- DHCP Relay
- VACL

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 translations are not synchronized with active and standby devices.
- Stateful NAT is not supported. However, NAT and Hot Standby Router Protocol (HSRP) can coexist.
- Normally, ICMP NAT flows time out after the expiration of the configured sampling-timeout and translation-timeout. However, when ICMP NAT flows present in the switch become idle, they time out immediately after the expiration of the sampling-timeout configured.
- If the translated IP is part of the outside interface subnet, then use the **ip local-proxy-arp** command on the NAT outside interface.
- 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.

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.

- Starting from Cisco NXOS Release 6.0(2)A8(3), NAT is ECMP aware and it supports a maximum of 24 ECMP paths.
- Starting from Cisco NX-OS Release 6.0(2)A8(9), Network Address Translation (NAT) statistics on Cisco Nexus 3548 switches.
- Traceroute is not supported on static and dynamic NAT

Configuring Static NAT

Enabling Static NAT

Procedure

	Command or Action	Purpose
Step 1	switch# **configure terminal**	Enters global configuration mode.
Step 2	switch(config)# **feature nat**	Enables the static NAT feature on the device.
Step 3	(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

Procedure

	Command or Action	Purpose
Step 1	switch# **configure terminal**	Enters global configuration mode.
Step 2	switch(config)# **interface** *type slot*/*port*	Specifies an interface to configure, and enters interface configuration mode.
Step 3	switch(config-if)# **ip nat** {**inside** | **outside**}	Specifies the interface as inside or outside. **Note** Only packets that arrive on a marked interface can be translated.
Step 4	(Optional) switch(config)# **copy running-config startup-config**	Saves the change persistently through reboots and restarts by copying the running configuration to the startup configuration.

Example

This example shows how to configure an interface with static NAT from the inside:

```
switch# configure terminal
switch(config)# interface ethernet 1/4
switch(config-if)# ip nat inside
```

Enabling Static NAT for an Inside Source Address

For inside source translation, the traffic flows from inside interface to the outside interface. NAT translates the inside local IP address to the inside global IP address. On the return traffic, the destination inside global IP address gets translated back to the inside local IP address.

Note When the Cisco Nexus device is configured to translate an inside source IP address (Src:ip1) to an outside source IP address (newSrc:ip2), the Cisco Nexus device implicitly adds a translation for an outside destination IP address (Dst: ip2) to an inside destination IP address (newDst: ip1).

Procedure

	Command or Action	Purpose
Step 1	switch# **configure terminal**	Enters global configuration mode.
Step 2	switch(config)# **ip nat inside source static** *local-ip-address global-ip-address* [**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.

Procedure

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

	Command or Action	Purpose
Step 2	switch(config)# **ip nat outside source static** *global-ip-address local-ip-address* [**group** *group-id*] [**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).

Procedure

	Command or Action	Purpose
Step 1	switch# **configure terminal**	Enters global configuration mode.
Step 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*}} **group** *group-id*	Maps static NAT to an inside local port to an inside global port.
Step 3	(Optional) switch(config)# **copy running-config startup-config**	Saves the change persistently through reboots and restarts by copying the running configuration to the startup configuration.

Example

This example shows how to map UDP services to a specific inside source address and UDP port:

```
switch# configure terminal
switch(config)#  ip nat inside source static udp 20.1.9.2 63 35.48.35.48 130
switch(config)# copy running-config startup-config
```

Configuring Static PAT for an Outside Source Address

You can map services to specific outside hosts using Port Address Translation (PAT).

Procedure

	Command or Action	Purpose
Step 1	switch# **configure terminal**	Enters global configuration mode.
Step 2	switch(config)# **ip nat outside source static** {*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*}}**group** *group-id* **add-route**	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.

Procedure

	Command or Action	Purpose
Step 1	**enable** **Example:** `switch> enable`	Enables privileged EXEC mode. • Enter your password if prompted.
Step 2	**configure terminal** **Example:** `switch# configure terminal`	Enters privileged EXEC mode.

	Command or Action	Purpose
Step 3	**ip nat inside source static** *inside-local-ip-address inside-global-ip-address* [**group** *group-id*] **Example:** `switch(config)# ip nat inside source static 10.1.1.1 192.168.34.4 group 4`	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** *outside-global-ip-address outside-local-ip-address* [**group** *group-id*] [**add-route**] **Example:** `switch(config)# ip nat outside source static 209.165.201.1 10.3.2.42 group 4 add-route`	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** *type number* **Example:** `switch(config)# interface ethernet 1/2`	Configures an interface and enters interface configuration mode.
Step 6	**ip address** *ip-address mask* **Example:** `switch(config-if)# ip address 10.2.4.1 255.255.255.0`	Sets a primary IP address for an interface.
Step 7	**ip nat inside** **Example:** `switch(config-if)# ip nat inside`	Connects the interface to an inside network, which is subject to NAT.
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.

	Command or Action	Purpose
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:

Procedure

	Command or Action	Purpose
Step 1	switch# **show ip nat translations**	Shows the translations for the inside global, inside local, outside local, and outside global IP addresses.

Example

This example shows how to display the static NAT configuration:

```
switch# show ip nat translations
Pro Inside global      Inside local       Outside local      Outside global
any ---                ---                20.4.4.40          220.2.2.20
tcp ---                ---                23.1.1.133:333     210.3.3.33:555
any 160.200.1.140      10.1.1.40          ---                ---
any 160.200.1.140      10.1.1.40          20.4.4.40          220.2.2.20
tcp 172.9.9.142:777    12.2.2.42:444      ---                ---
tcp 172.9.9.142:777    12.2.2.42:444      23.1.1.133:333     210.3.3.33:555
```

Configuring Dynamic NAT

Configuring Dynamic Translation and Translation Timeouts

Procedure

	Command or Action	Purpose
Step 1	**enable** **Example:** `Switch> enable`	Enables privileged EXEC mode. • Enter your password if prompted.
Step 2	**configure terminal** **Example:** `Switch# configure terminal`	Enters global configuration mode.
Step 3	**ip access-list** *access-list-name* **Example:** `Switch(config)# ip access-list acl1`	Defines an access list and enters access-list configuration mode.
Step 4	**permit** *protocol source source-wildcard* **any** **Example:**	Sets conditions in an IP access list that permit traffic matching the conditions.

	Command or Action	Purpose
	`Switch(config-acl)# permit ip 10.111.11.0/24 any`	
Step 5	**deny** *protocol source source-wildcard* **any** **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** *access-list-name* **interface** *type number* **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** *type number* **Example:** `Switch(config)# interface ethernet 1/4`	Configures an interface and enters interface configuration mode.
Step 9	**ip address** *ip-address mask* **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:** `Switch(config-if)# exit`	Exits interface configuration mode and returns to global configuration mode.
Step 12	**interface** *type number* **Example:** `Switch(config)# interface ethernet 1/1`	Configures an interface and enters interface configuration mode.
Step 13	**ip address** *ip-address mask* **Example:** `Switch(config-if)# ip address 172.16.232.182 255.255.255.240`	Sets a primary IP address for an interface.

	Command or Action	Purpose
Step 14	**ip nat outside** **Example:** `Switch(config-if)# ip nat outside`	Connects the interface to an outside network.
Step 15	**exit** **Example:** `Switch(config-if)# exit`	Exits interface configuration mode and returns to global configuration mode.
Step 16	**ip nat translation tcp-timeout** *seconds* **Example:** `Switch(config)# ip nat translation tcp-timeout 50000`	Specifies the timeout value for TCP-based dynamic NAT entries. • 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 [all-host]** *number-of-entries* **Example:** `Switch(config)# ip nat translation max-entries 300`	Specifies the maximum number of dynamic NAT translations. The number of entries can be between 1 and 1023.
Step 18	**ip nat translation udp-timeout** *seconds* **Example:** `Switch(config)# ip nat translation udp-timeout 45000`	Specifies the timeout value for UDP-based dynamic NAT entries. • 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** *seconds* **Example:** `switch(config)# ip nat translation timeout 13000`	Specifies the timeout value for dynamic NAT translations.
Step 20	**ip nat translation syn-timeout** {*seconds* | **never**} **Example:** `switch(config)# ip nat translation syn-timeout 20`	Specifies the 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. The **never** keyword specifies that the SYN timer will not be run.

	Command or Action	Purpose
Step 21	**ip nat translation finrst-timeout** {*seconds* | **never**} **Example:** `switch(config)# ip nat translation finrst-timeout 30`	Specifies the timeout value for the flow entries when a connection is terminated by receiving finish (FIN) or reset (RST) packets. Use the same keyword to configure the behavior for both RST and FIN packets. The timeout value ranges from 1 second to 172800 seconds. The default value is 60 seconds. The **never** keyword specifies that the FIN or RST timer will not be run.
Step 22	**end** **Example:** `Switch(config)# end`	Exits global configuration mode and returns to privileged EXEC mode.

Configuring Dynamic NAT Pool

You can create a NAT pool by either defining the range of IP addresses in a single **ip nat pool** command or by using the **ip nat pool** and **address** commands

Procedure

	Command or Action	Purpose
Step 1	switch# **configure terminal**	Enters global configuration mode.
Step 2	switch(config)# **feature nat**	Enables the NAT feature on the device.
Step 3	switch(config)# **ip nat pool** *pool-name* [*startip endip*] {**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 4	(Optional) switch(config-ipnat-pool)# **address** *startip endip*	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** *pool-name*	Deletes the specified NAT pool.

Example

This example shows how to create a NAT pool with a prefix length:

```
switch# configure terminal
switch(config)# ip nat pool pool1 30.1.1.1 30.1.1.2 prefix-length 24
switch(config)#
```

This example shows how to create a NAT pool with a network mask:

```
switch# configure terminal
switch(config)# ip nat pool pool5 20.1.1.1 20.1.1.5 netmask 255.0.255.0
switch(config)#
```

This example shows how to create a NAT pool and define the range of global IP addresses using the **ip nat pool** and **address** commands:

```
switch# configure terminal
switch(config)# ip nat pool pool7 netmask 255.255.0.0
switch(config-ipnat-pool)# address 40.1.1.1 40.1.1.5
switch(config-ipnat-pool)#
```

This example shows how to delete a NAT pool:

```
switch# configure terminal
switch(config)# no ip nat pool pool4
switch(config)#
```

Configuring Source Lists

You can configure a source list of IP addresses for the inside interface and the outside interface.

Before you begin

Ensure that you configure a pool before configuring the source list for the pool.

Procedure

	Command or Action	Purpose
Step 1	switch# **configure terminal**	Enters global configuration mode.
Step 2	(Optional) switch# **ip nat inside source list** *list-name* **pool** *pool-name* [**overload**]	Creates a NAT inside source list with pool with or without overloading.
Step 3	(Optional) switch# **ip nat outside source list** *list-name* **pool** *pool-name* [**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.

Procedure

	Command or Action	Purpose
Step 1	switch# **configure terminal**	Enters global configuration mode.
Step 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**]	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** *access-list-name* [**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.
Step 4	switch(config)# **ip nat pool** *pool-name* [*startip endip*] {**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** *type slot/port*	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** *type slot/port*	Configures an interface and enters interface configuration mode.

	Command or Action	Purpose
Step 9	switch(config-if)# **ip nat inside**	Connects the interface to an inside network, which is subject to NAT.

Example

This example shows how to configure dynamic twice NAT for an inside source address:

```
switch# configure terminal
Enter configuration commands, one per line.  End with CNTL/Z.
switch(config)# ip nat outside source static 2.2.2.2 4.4.4.4 group 20 dynamic
switch(config)# ip nat inside source list acl_1 pool pool_1 overload group 20 dynamic
switch(config)# ip nat pool pool_1 3.3.3.3 3.3.3.10 prefix-length 24
switch(config)# interface Ethernet1/8
switch(config-if)# ip nat outside
switch(config-if)# exit
switch(config)# interface Ethernet1/15
switch(config-if)# ip nat inside
```

Configuring Dynamic Twice NAT for an Outside Source Address

For an outside source address translation, the traffic flows from the outside interface to the inside interface. You can configure dynamic twice NAT for an outside source address.

Before you begin

Ensure that you enable NAT on the switch.

Procedure

	Command or Action	Purpose
Step 1	switch# **configure terminal**	Enters global configuration mode.
Step 2	switch(config)# **ip nat inside source static** *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**]	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** *access-list-name* [**interface** *type slot/port* **pool** *pool-name*] [**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** *pool-name* [*startip endip*] {**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** *type slot/port*	Configures an interface and enters interface configuration mode.

	Command or Action	Purpose
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** *type slot/port*	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** *global-ip-address local-ip-address* [**outside** *local-ip-address global-ip-address*] | **outside** *local-ip-address global-ip-address*]	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 including dynamic 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 ip nat translations vrf all	Displays active Network Address Translation (NAT) translations including dynamic translations along with its VRF name.
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# show ip nat translations verbose
Pro Inside global Inside local Outside local Outside global
any 1.1.1.1 10.1.1.2 --- ---
    Vrf: default    Group_id:0
icmp 101.1.0.1:65351 101.0.0.1:0 102.1.0.1:231 102.1.0.1:231
    VRF: red   Group_id:0
```

```
 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
    VRF: red   Group_id:0
 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
    VRF: red   Group_id:0
 Format(H:M:S) Time-left:12:0:9
```

Verifying NAT Statistics

To display Network Address Translation (NAT) statistics, perform the following task:

Command	Purpose
show ip nat statistics	Display Network Address Translation (NAT) statistics.

Example

This example shows the sample output from the **show ip nat statistics** command:

```
switch# show ip nat statistics

IP NAT Statistics
==================================================
Total active translations: 2
No.Static: 2
No.Dyn:    0
No.ICMP:   0
--------------------------------------------------
Total expired Translations: 0
SYN timer expired:          0
FIN-RST timer expired:      0
Inactive timer expired:     0
--------------------------------------------------
Total Hits:  0                  Total Misses:  3
In-Out Hits: 0                  In-Out Misses: 0
Out-In Hits: 0                  Out-In Misses: 3
--------------------------------------------------
Total SW Translated Packets: 0
In-Out SW Translated: 0
Out-In SW Translated: 0
--------------------------------------------------
Total SW Dropped Packets: 0
In-Out SW Dropped: 0
Out-In SW Dropped: 0

Address alloc. failure drop:      0
Port alloc.  failure drop:        0
Dyn. Translation max limit drop:  0
ICMP max limit drop:              0
Allhost max limit drop:           0
--------------------------------------------------
NAT Inside Interfaces:  1
Ethernet1/1

NAT Outside Interfaces: 1
Ethernet1/3
--------------------------------------------------
```

```
Inside source list:
++++++++++++++++++++

Access list: ACL1
RefCount: 0
Pool: pool1     Overload
Total addresses: 200
Allocated: 0    percentage: 0%
Missed: 0
```

Clearing NAT Statistics

To clear Network Address Translation (NAT) statistics, perform the following task:

Command	Purpose
clear ip nat Statistics	Clear Network Address Translation (NAT) statistics entries. This command clears only the following parameters: • Total Hits • In-Out Hits • Out-In Hits • Total Misses • In-Out Misses • Out-In Misses • Total SW Translated Packets • In-Out SW Translated • Out-In SW Translated • Total SW Dropped Packets • In-Out SW Dropped • Address alloc. failure drop • Port alloc. failure drop • Dyn. Translation max limit drop • ICMP max limit drop • Allhost max limit drop • Inside / Outside source list • Missed

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.

Procedure

	Command or Action	Purpose
Step 1	switch# **configure terminal**	Enters global configuration mode.
Step 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**]]	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** *LOCAL IP GLOBAL IP* | [*tcp* | *udp LOCAL IP LOCAL PORT GLOBAL IP GLOBAL PORT*] [**group** *group-id*] [**dynamic**] [**vrf** <*vrf-name*> [**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** *type slot/port* [**vrf** <*vrf-name* **ip nat** *inside* | *outside*	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
```

See the output of **show ip nat translation verbose** command.

```
switch# show ip nat translation verbose
Pro Inside global Inside local Outside local Outside global
any 1.1.1.1 10.1.1.2 --- ---
    Vrf: default    Group_id:0
icmp 101.1.0.1:65351 101.0.0.1:0 102.1.0.1:231 102.1.0.1:231
    VRF: red   Group_id:0
 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
    VRF: red   Group_id:0
 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
    VRF: red   Group_id:0
 Format(H:M:S) Time-left:12:0:9
```

See the output of **show ip nat translations vrf all** command.

```
switch# show ip nat translations vrf all
Pro Inside global      Inside local       Outside local      Outside global       Vrf
udp ---                ---                3.3.3.30:100       3.3.3.3:30         default
udp 2.1.1.10:200       2.1.1.1:100        3.3.3.30:100       3.3.3.3:30         default
```

```
icmp 70.1.1.1:65276    10.1.1.2:0         20.1.1.2:0         20.1.1.2:0         default
udp  101.1.0.1:65383   101.0.0.1:63       102.1.0.1:63       102.1.0.1:63       red
```

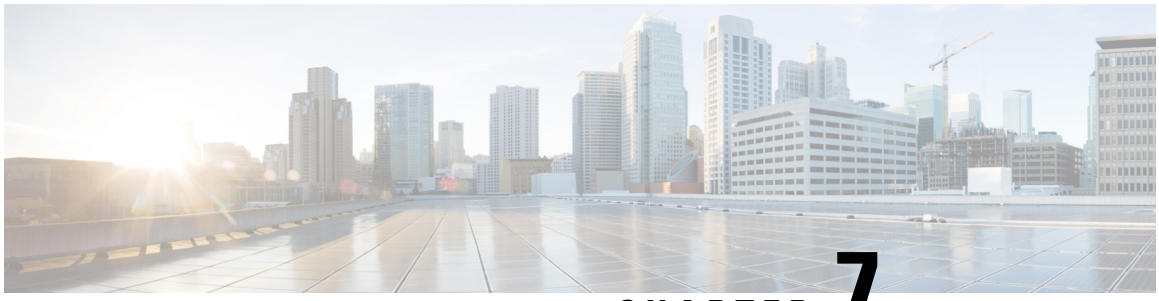

CHAPTER 7

Configuring IP Event Dampening

This chapter includes the following sections:

- IP Event Dampening, on page 113

IP Event Dampening

The IP Event Dampening feature introduces a configurable exponential decay mechanism to suppress the effects of excessive interface flapping events on routing protocols and routing tables in the network. This feature allows the network operator to configure a router to automatically identify and selectively dampen a local interface that is flapping.

Guidelines and Limitations

See the following guidelines and limitations before configuring IP Event Dampening feature:

- Due to changes in the netstack-IP component, all the IP clients observe the impact of dampening or interface.
- For each flap of the interface, a certain penalty is added. The penalty decays exponentially whose parameters are configured.
- When penalty exceeds a certain high level, the interface is dampened. It is unsuppressed when the penalty decays below a low level.
- When an interface is dampened, the IP address and the static routes are removed from the interface. All the clients of IP get an IP delete notification.
- When an interface is unsuppressed, the IP address and the relevant routes are added back. All the clients of IP get an IP address add notification for all the IP addresses of the interface.
- All Layer 3 interfaces that are configured on the Ethernet interface, port changes, and SVI support this feature.

IP Event Dampening Overview

Interface state changes occur when interfaces are administratively brought up or down or if an interface changes state. When an interface changes state or flaps, routing protocols are notified of the status of the routes that are affected by the change in state. Every interface state change requires all affected devices in the

network to recalculate best paths, install or remove routes from the routing tables, and then advertise valid routes to peer routers. An unstable interface that flaps excessively can cause other devices in the network to consume substantial amounts of system processing resources and cause routing protocols to lose synchronization with the state of the flapping interface.

The IP Event Dampening feature introduces a configurable exponential decay mechanism to suppress the effects of excessive interface flapping events on routing protocols and routing tables in the network. This feature allows the network operator to configure a router to automatically identify and selectively dampen a local interface that is flapping. Dampening an interface removes the interface from the network until the interface stops flapping and becomes stable. Configuring the IP Event Dampening feature improves convergence times and stability throughout the network by isolating failures so that disturbances are not propagated. This, in turn, reduces the utilization of system processing resources by other devices in the network and improves overall network stability.

Interface State Change Events

This section describes the interface state change events of the IP Event Dampening feature. This feature employs a configurable exponential decay mechanism that is used to suppress the effects of excessive interface flapping or state changes. When the IP Event Dampening feature is enabled, flapping interfaces are dampened from the perspective of the routing protocol by filtering excessive route updates. Flapping interfaces are identified, assigned penalties, suppressed if necessary, and made available to the network when the interface stabilizes. Figure 1 displays interface state events as they are perceived by routing protocols.

Suppress Threshold

The suppress threshold is the value of the accumulated penalty that triggers the router to dampen a flapping interface. The flapping interface is identified by the router and assigned a penalty for each up and down state change, but the interface is not automatically dampened. The router tracks the penalties that a flapping interface accumulates. When the accumulated penalty reaches the default or preconfigured suppress threshold, the interface is placed in a dampened state.

Half-Life Period

The half-life period determines how fast the accumulated penalty can decay exponentially. When an interface is placed in a dampened state, the router monitors the interface for additional up and down state changes. If the interface continues to accumulate penalties and the interface remains in the suppress threshold range, the interface will remain dampened. If the interface stabilizes and stops flapping, the penalty is reduced by half after each half-life period expires. The accumulated penalty will be reduced until the penalty drops to the reuse threshold. The configurable range of the half-life period timer is from 1 to 30 seconds. The default half-life period timer is 5 seconds.

Reuse Threshold

When the accumulated penalty decreases until the penalty drops to the reuse threshold, the route is unsuppressed and made available to other devices in the network. The range of the reuse value is from 1 to 20000 penalties. The default value is 1000 penalties.

Maximum Suppress Time

The maximum suppress time represents the maximum time an interface can remain dampened when a penalty is assigned to an interface. The maximum suppress time can be configured from 1 to 255 seconds. The

maximum penalty is truncated to maximum 20000 unit. The maximum value of the accumulated penalty is calculated based on the maximum suppress time, reuse threshold, and half-life period.

Affected Components

When an interface is not configured with dampening, or when an interface is configured with dampening but is not suppressed, the routing protocol behavior as a result of interface state transitions is not changed by the IP Event Dampening feature. However, if an interface is suppressed, the routing protocols and routing tables are immune to any further state transitions of the interface until it is unsuppressed.

Route Types

The following interfaces are affected by the configuration of this feature:

- Connected routes:
  - The connected routes of dampened interfaces are not installed into the routing table.
  - When a dampened interface is unsuppressed, the connected routes will be installed into the routing table if the interface is up.
- Static routes:
  - Static routes assigned to a dampened interface are not installed into the routing table.
  - When a dampened interface is unsuppressed, the static route will be installed into the routing table if the interface is up.

Note Only the primary interface can be configured with this feature, and all subinterfaces are subject to the same dampening configuration as the primary interface. IP Event Dampening does not track the flapping of individual subinterfaces on an interface.

Supported Protocols

All the protocols that are used are impacted by the IP Event Dampening feature. The IP Event Dampening feature supports Border Gateway Protocol (BGP), Enhanced Interior Gateway Routing Protocol (EIGRP), Hot Standby Routing Protocol (HSRP), Open Shortest Path First (OSPF), Routing Information Protocol (RIP), and VRRP. Ping and SSH to the concerned interface IP address does not work.

Note The IP Event Dampening feature has no effect on any routing protocols if it is not enabled or an interface is not dampened.

How to Configure IP Event Dampening

Enabling IP Event Dampening

The **dampening** command is entered in interface configuration mode to enable the IP Event Dampening feature. If this command is applied to an interface that already has dampening configured, all dampening

states are reset and the accumulated penalty will be set to 0. If the interface has been dampened, the accumulated penalty will fall into the reuse threshold range, and the dampened interface will be made available to the network. The flap counts, however, are retained.

Procedure

	Command or Action	Purpose
Step 1	**configure terminal**	Enters global configuration mode.
Step 2	**interface** *type number*	Enters interface configuration mode and configures the specified interface.
Step 3	**dampening** [*half-life-period reuse-threshold*] [*suppress-threshold max-suppress* [*restart-penalty*]]	Enables interface dampening. • Entering the **dampening** command without any arguments enables interface dampening with default configuration parameters. • When manually configuring the timer for the *restart-penalty* argument, the values must be manually entered for all arguments.
Step 4	**end**	Exits interface configuration mode.

Verifying IP Event Dampening

Use the **show dampening interface** or **show interface dampening** commands to verify the configuration of the IP Event Dampening feature.

Procedure

	Command or Action	Purpose
Step 1	**show dampening interface**	Displays dampened interfaces.
Step 2	**show interface dampening**	Displays dampened interfaces on the local router.