

# Cisco Nexus 3550-T NX-OS Interfaces Configuration Guide, Release 10.6(x)

**First Published:** 2025-08-14

**Last Modified:** 2026-04-27

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

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

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

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

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

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

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

Cisco and the Cisco logo are trademarks or registered trademarks of Cisco and/or its affiliates in the U.S. and other countries. To view a list of Cisco trademarks, go to this URL: https://www.cisco.com/c/en/us/about/legal/trademarks.html. 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)

© 2025–2026 Cisco Systems, Inc. All rights reserved.



# CONTENTS


Trademarks ?

PREFACE

Preface xi

Audience xi

Document Conventions xi

Related Documentation for Cisco Nexus 3550-T Switches xii

Documentation Feedback xii

Communications, Services, and Additional Information xii

CHAPTER 1

New and Changed Information 1

New and Changed Information 1

CHAPTER 2

Interfaces Configuration Guide 3

About Interfaces 3

Physical Interfaces 3

Ethernet Interfaces 3

Access Ports 3

Routed Ports 3

Management Interface 4

Port-Channel Interfaces 4

Loopback Interfaces 4

CHAPTER 3

Configuring Static and Dynamic NAT Translation 5

Network Address Translation Overview 5

Information About Static NAT 5

Dynamic NAT Overview 7

Timeout Mechanisms 7
NAT Inside and Outside Addresses 8
Pool Support for Dynamic NAT 9
Guidelines and Limitations for Static and Dynamic NAT 9
Restrictions for Dynamic NAT 10
Configuring Static NAT 10
Enabling Static NAT 10
Configuring NAT on an Interface 11
Enabling Static NAT for an Inside Source Address 12
Enabling Static NAT for an Outside Source Address 12
Configuring Static PAT for an Inside Source Address 13
Configuring Static PAT for an Outside Source Address 14
Enabling and Disabling no-alias Configuration 14
Configuration Example for Static NAT and PAT 16
Verifying the Static NAT Configuration 17
Configuring Dynamic NAT 18
Configuring Dynamic Translation and Translation Timeouts 18
Configuring Dynamic NAT Pool 20
Configuring Source Lists 21
Clearing Dynamic NAT Translations 22
Verifying Dynamic NAT Configuration 22
Example: Configuring Dynamic Translation and Translation Timeouts 25

CHAPTER 4 **Configuring Layer 2 Interfaces 27**
Information About Access and Trunk Interfaces 27
About Access and Trunk Interfaces 27
IEEE 802.1Q Encapsulation 28
Access VLANs 29
Native VLAN IDs for Trunk Ports 29
Allowed VLANs 30
Default Interfaces 30
Switch Virtual Interface and Autostate Behavior 30
Counter Values 30
Prerequisites for Layer 2 Interfaces 31

Guidelines and Limitations for Layer 2 Interfaces 31
Default Settings for Layer 2 Interfaces 32
Configuring Access and Trunk Interfaces 33
Configuring a Layer 2 Access Port 33
Configuring Access Host Ports 35
Configuring Trunk Ports 36
Configuring the Allowed VLANs for Trunking Ports 37
Configuring a Default Interface 39
Changing the System Default Port Mode to Layer 2 40
Verifying the Interface Configuration 41
Monitoring the Layer 2 Interfaces 42
Configuration Examples for Access and Trunk Ports 43
Related Documents 43

CHAPTER 5 **Configuring Port Channels** 45
About Port Channels 45
Port Channels 46
Port-Channel Interfaces 47
Basic Settings 48
Compatibility Requirements 48
Load Balancing Using Port Channels 50
LACP 51
LACP Overview 51
Port-Channel Modes 51
LACP ID Parameters 53
LACP System Priority 53
LACP Port Priority 53
LACP Administrative Key 53
LACP-Enabled and Static Port Channels Differences 53
LACP Compatibility Enhancements 54
LACP Port-Channel Minimum Links and MaxBundle 54
LACP Fast Timers 55
Prerequisites for Port Channeling 55
Guidelines and Limitations 55

Default Settings 56
Configuring Port Channels 57
Creating a Port Channel 57
Adding a Layer 2 Port to a Port Channel 58
Adding a Layer 3 Port to a Port Channel 60
Configuring the Bandwidth and Delay for Informational Purposes 62
Shutting Down and Restarting the Port-Channel Interface 63
Configuring a Port-Channel Description 65
Enabling LACP 66
Configuring LACP Port-Channel Port Modes 67
Configuring LACP Port-Channel Minimum Links 68
Configuring the LACP Port-Channel MaxBundle 69
Configuring the LACP Fast Timer Rate 70
Configuring the LACP System Priority 71
Configuring the LACP Port Priority 72
Configuring LACP System MAC and Role 73
Disabling LACP Graceful Convergence 74
Reenabling LACP Graceful Convergence 75
Disabling LACP Suspend Individual 76
Reenabling LACP Suspend Individual 78
Configuring Delayed LACP 79
Verifying the Port-Channel Configuration 80
Monitoring the Port-Channel Interface Configuration 81
Example Configurations for Port Channels 81
Related Documents 82

CHAPTER 6 **Configuring vPCs 83**
Information About vPCs 83
vPC Overview 83
Hitless vPC Role Change 85
vPC Terminology 85
vPC Peer-Link Overview 86
Features That You Must Manually Configure on the Primary and Secondary Devices 88
Configuring Layer 3 Backup Routes on a vPC Peer-Link 89

Peer-Keepalive Link and Messages 89
vPC Peer-Gateway 90
vPC Domain 91
vPC Topology 91
Compatibility Parameters for vPC Interfaces 92
Configuration Parameters That Must Be Identical 92
Configuration Parameters That Should Be Identical 94
Consequences of Parameter Mismatches 94
vPC Number 94
Moving Other Port Channels into a vPC 95
vPC Interactions with Other Features 95
vPC and LACP 95
vPC Peer-Links and STP 96
vPC Peer Switch 97
vPC and ARP or ND 98
vPC Multicast— IGMP, and IGMP Snooping 98
vPC Peer-Links and Routing 99
CFSoE 99
vPC and Orphan Ports 100
vPC Recovery After an Outage 100
Autorecovery 100
vPC Peer Roles After a Recovery 100
Guidelines and Limitations 101
Best Practices for Layer 3 and vPC Configuration 103
Layer 3 and vPC Configuration Overview 103
Supported Topologies for Layer 3 and vPC 103
Peering with an External Router Using Layer 3 Links 104
Peering Between vPC Devices for a Backup Routing Path 105
Direct Layer 3 Peering Between Routers 105
Peering Between Two Routers with vPC Devices as Transit Switches 106
Peering with an External Router on Parallel Interconnected Routed Ports 106
Peering between vPC Switch Pairs on Parallel Interconnected Routed Ports 107
Peering Over a PC Interconnection and Dedicated Interswitch Link Using non-vPC VLAN 107
Peering Directly Over a vPC Connection 108

Default Settings 110
Configuring vPCs 110
Enabling vPCs 111
Disabling vPCs 112
Creating a vPC Domain and Entering vpc-domain Mode 112
Configuring a vPC Keepalive Link and Messages 113
Creating a vPC Peer-Link 115
Configuring a vPC Peer-Gateway 116
Configuring Fast Convergence 117
Configuring LACP vPC Convergence 118
Configuring a Graceful Consistency Check 120
Checking the Configuration Compatibility on a vPC Peer-Link 121
Moving Other Port Channels into a vPC 121
Manually Configuring a vPC Domain MAC Address 123
Manually Configuring the System Priority 124
Manually Configuring the vPC Peer Device Role 125
Configuring for Recovery After an Outage 126
Configuring Reload Restore 126
Configuring an Autorecovery 128
Configuring the Suspension of Orphan Ports 130
Configuring Delay Restore on an Orphan Port 131
Configuring the vPC Peer Switch 132
Configuring a Pure vPC Peer Switch Topology 132
Configuring Hitless vPC Role Change 134
Use Case Scenario for vPC Role Change 135
Verifying the vPC Configuration 135
Monitoring vPCs 136
Configuration Examples for vPCs 136

CHAPTER 7 **Configuring Unidirectional Link Detection** 139
Unidirectional Link Detection 139
UDLD Modes 140
Configuring the UDLD Mode 140

CHAPTER 8 Multicast Fairness Tuning 143

Multicast fairness 143

Guidelines and limitations for multicast fairness tuning 143

Configure multicast fairness tuning 144

Verify the Multicast Fairness Tuning Configuration 145

CHAPTER 9 Configuring Layer 3 Interfaces 151

About Layer 3 Interfaces 151

Routed Interfaces 151

VLAN Interfaces 152

Changing VRF Membership for an Interface 152

Notes About Changing VRF Membership for an Interface 153

Loopback Interfaces 153

High Availability 153

DHCP Client 153

Limitations for Using DHCP Client on Interfaces 154

Prerequisites for Layer 3 Interfaces 154

Guidelines and Limitations for Layer 3 Interfaces 154

Default Settings 155

Configuring Layer 3 Interfaces 155

Configuring a Routed Interface 155

Configuring a VLAN Interface 156

Enabling Layer 3 Retention During VRF Membership Change 158

Configuring a Loopback Interface 158

Assigning an Interface to a VRF 159

Configuring a DHCP Client on an Interface 160

Verifying the Layer 3 Interfaces Configuration 160

Monitoring the Layer 3 Interfaces 162

Configuration Examples for Layer 3 Interfaces 162

Example of Changing VRF Membership for an Interface 162

Related Documents 164



# Preface

This preface includes the following sections:


- Audience, on page xi
- Document Conventions, on page xi
- Related Documentation for Cisco Nexus 3550-T Switches, on page xii
- Documentation Feedback, on page xii
- Communications, Services, and Additional Information, on page xii


## Audience

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

## Document Conventions

Command descriptions use the following conventions:

| Convention | Description |
|---|---|
| **bold** | Bold text indicates the commands and keywords that you enter literally as shown. |
| *Italic* | Italic text indicates arguments for which you supply the values. |
| [x] | Square brackets enclose an optional element (keyword or argument). |
| [x \| y] | Square brackets enclosing keywords or arguments that are separated by a vertical bar indicate an optional choice. |
| {x \| y} | Braces enclosing keywords or arguments that are separated by a vertical bar indicate a required choice. |
| [x {y \| z}] | Nested set of square brackets or braces indicate optional or required choices within optional or required elements. Braces and a vertical bar within square brackets indicate a required choice within an optional element. |

| Convention | Description |
|---|---|
| `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 includes 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 that 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. |

# Related Documentation for Cisco Nexus 3550-T Switches

The entire Cisco Nexus 3550-T switch documentation set is available at the following URL:

https://www.cisco.com/c/en/us/support/switches/nexus-3550-series/series.html

# Documentation Feedback

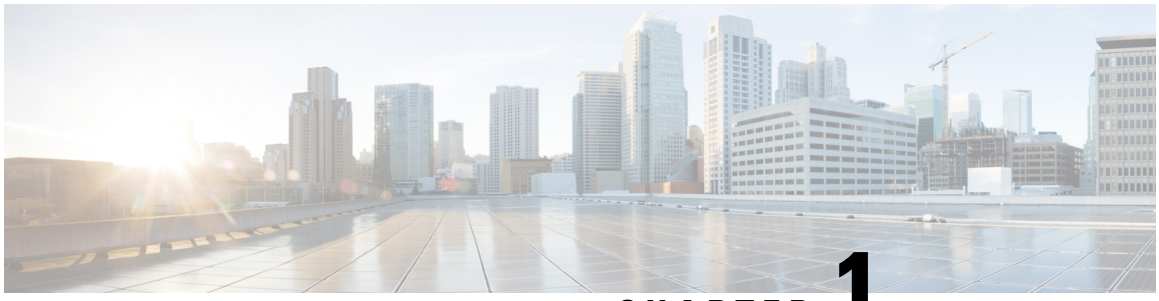
To provide technical feedback on this document, or to report an error or omission, please send your comments to nexus9k-docfeedback@cisco.com. We appreciate your feedback.

# Communications, Services, and Additional Information

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

### Cisco Bug Search Tool

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



CHAPTER 1

# New and Changed Information

This section contains the new and changed information for a release.

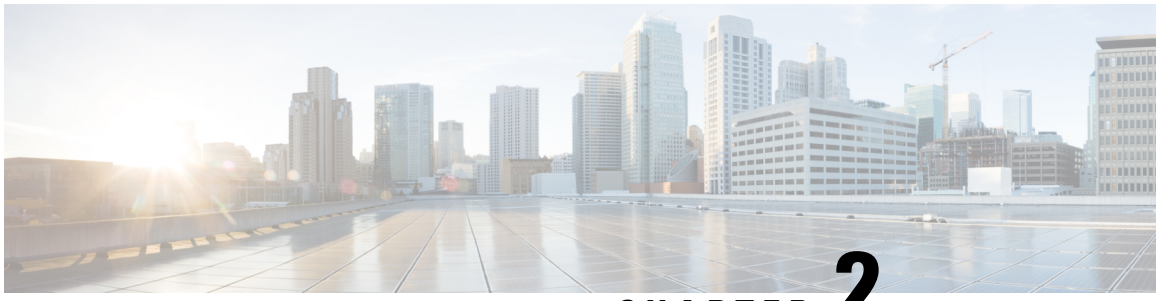

- New and Changed Information, on page 1


# New and Changed Information

*Table 1: New and Changed Information for Cisco Nexus 3550-T NX-OS Release 10.6(x)*

| Feature | Description | Changed in Release | Where Documented |
|---|---|---|---|
| 25G support on 3550-T | • Added support for 25G on first 40 physical ports of 3550-T<br>• Added support for multicast fairness on 25G ports for 3550-T | 10.6(3)F | Physical Interfaces, on page 3<br>Multicast fairness, on page 143<br>Configure multicast fairness tuning, on page 144<br>Verify the Multicast Fairness Tuning Configuration, on page 145 |
| NA | There are no new features supported for this release | 10.6(1)F | NA |

New and Changed Information



CHAPTER 2

# Interfaces Configuration Guide

This preface includes the following sections:


- About Interfaces, on page 3


# About Interfaces

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

## Physical Interfaces

Beginning with NX-OS Release 10.6(3)F, 3550-T platform supports 25G speed on first 40 front panel ports. These interfaces do not support forward error correction (FEC). When connecting to a peer device, disable FEC on the other device to ensure compatibility.

## Ethernet Interfaces

Ethernet interfaces include routed ports.

Cisco Nexus® 3550-T switch has the following guidelines and limitations:

- Mixed speed is not supported within the same quad.

### Access Ports

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

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

### Routed Ports

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

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

# Management Interface

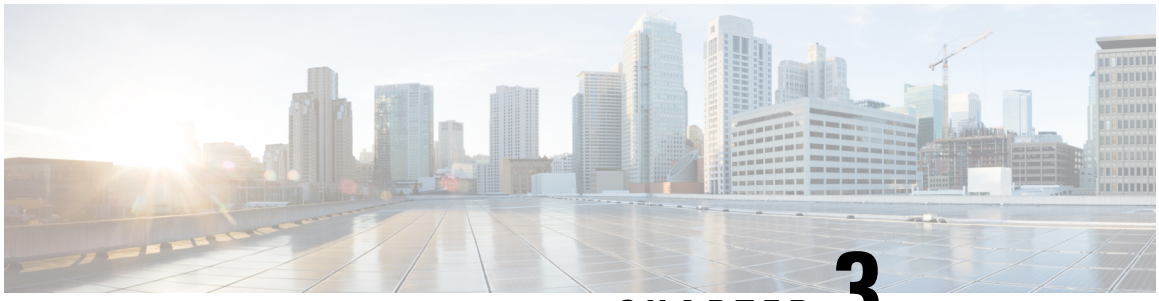
You can use the management Ethernet interface to connect the device to a network for remote management using a Telnet client, the Simple Network Management Protocol (SNMP), or other management agents. The management port (mgmt0) is autosensing and operates in full-duplex mode at a speed of 1000 Mb/s.

# Port-Channel Interfaces

A port channel is a logical interface that is an aggregation of multiple physical interfaces. You can bundle up to 8 individual links to physical ports into a port channel to improve bandwidth and redundancy. You can also use port channeling to load balance traffic across these channeled physical interfaces. For more information about port-channel interfaces, see the *Configuring Port Channels* section.

# Loopback Interfaces

A virtual loopback interface is a virtual interface with a single endpoint that is always up. Any packet that is transmitted over a virtual loopback interface is immediately received by that interface. Loopback interfaces emulate a physical interface.



CHAPTER 3

# Configuring Static and Dynamic NAT Translation


- Network Address Translation Overview, on page 5
- Information About Static NAT, on page 5
- Dynamic NAT Overview, on page 7
- Timeout Mechanisms, on page 7
- NAT Inside and Outside Addresses, on page 8
- Pool Support for Dynamic NAT, on page 9
- Guidelines and Limitations for Static and Dynamic NAT, on page 9
- Restrictions for Dynamic NAT, on page 10
- Configuring Static NAT, on page 10
- Configuring Dynamic NAT, on page 18


## 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 IP addresses to inside global IP 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

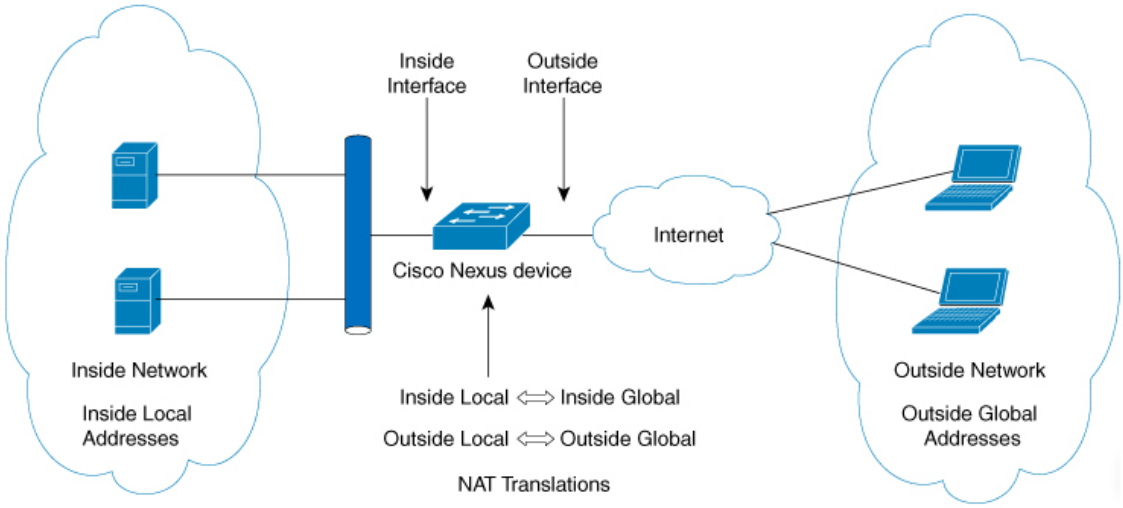
NAT, which means that you can add or remove a NAT translation in the NAT configuration without affecting the existing NAT traffic flows.

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 .

The main difference between static and dynamic NAT is that for 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 based on configured and applicable NAT timeouts. Static entries exist all the time irrespective of the device receiving traffic. However, for both static and dynamic NAT, each host can use different address or port for each subsequent translation based on different configurations, like overload.

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 1: Static NAT***




These are key terms to help you understand static NAT:

- NAT inside interface—The Layer 3 interface that faces the private network.
- NAT outside interface—The Layer 3 interface that faces the public network.
- Local address—Any address that appears on the inside (private) portion of the network.
- Global address—Any address that appears on the outside (public) portion of the network.
- Legitimate IP address—An address that is assigned by the Network Information Center (NIC) or service provider.
- Inside local address—The IP address assigned to a host on the inside network. This address does not need to be a legitimate IP address.

- Outside local address—The IP address of an outside host as it appears to the inside network. It does not have to be a legitimate address, because it is allocated from an address space that can be routed on the inside network.
- Inside global address—A legitimate IP address that represents one or more inside local IP addresses to the outside world.
- Outside global address—The IP address that the host owner assigns to a host on the outside network. The address is a legitimate address that is allocated from an address or network space that can be routed.

# Dynamic NAT Overview

Dynamic Network Address Translation (NAT) translates a group of real IP addresses into mapped IP addresses that are routable on a destination network. Dynamic NAT establishes a one-to-one mapping between unregistered and registered IP addresses; however, the mapping can vary depending on the registered IP address that is available at the time of communication.

A dynamic NAT configuration automatically creates a firewall between your internal network and outside networks or the Internet. Dynamic NAT allows only connections that originate inside the stub domain—a device on an external network cannot connect to devices in your network, unless your device has initiated the contact.

Dynamic NAT translations do not exist in the NAT translation table until a device receives traffic that requires translation. Dynamic translations are cleared or timed out when not in use to make space for new entries based on configured and applicable NAT timeouts. Usually, NAT translation entries are cleared based on timers. The default minimum timeout for dynamic NAT translations is 3600 seconds.



**Note** The **ip nat translation sampling-timeout** command is not supported. Statistics are collected every 60 seconds for the installed NAT policies. These statistics are used to determine if the flow is active or not.

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.

# 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. This release supports syn-timeout and finrst-timeout . The following NAT translation timeout timers are supported on the switch:

The following NAT translation timeout timers are supported on the switch:

- syn-timeout —Timeout value for TCP data packets that send the SYN request, but do not receive a SYN-ACK reply. The timeout value ranges from 1 second to 172800 seconds. The default value is 60 seconds.
- finrst-timeout —Timeout value for the flow entries when a connection is terminated by receiving RST or FIN packets. Use the same keyword to configure the behavior for both RST and FIN packets. The timeout value ranges from 1 second to 172800 seconds. The default value is 60 seconds.

If a FIN packet is received after the connection is established, SYN >SYN-ACK >FIN, the finrst timer starts.

If a FIN-ACK is received from the other side, the translation entry is cleared immediately, else it clears after the timeout value completes.

If an RST packet is received after the connection is established, SYN >SYN-ACK >RST, the translation entry is cleared immediately.

- tcp-timeout —Timeout value for TCP translations for which connections have been established after a three-way handshake (SYN, SYN-ACK, ACK). If no active flow occurs after the connection has been established, the translations expire as per the configured timeout value.

  The timeout value ranges from 60 seconds to 172800 seconds; default is 3600 seconds.

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

  The timeout value ranges from 60 seconds to 172800 seconds; default is 3600 seconds.

- timeout—Timeout value for dynamic NAT translations.

  The timeout value ranges from 60 seconds to 172800 seconds; default is 3600 seconds.

- icmp-timeout— Timeout value for ICMP packets.

  The timeout value ranges from 60 seconds to 172800 seconds; default is 3600 seconds.



**Note** When you create dynamic entries without timeouts configured, they take the default timeout of 3600 seconds. After you change the default timeout values to the new values, the translation entries created after will pick up the latest timeout values.

# NAT Inside and Outside Addresses

NAT inside refers to networks owned by an organization that must be translated. When NAT is configured, hosts within this network will have addresses in one space (known as the local address space) that will appear to those outside the network as being in another space (known as the global address space).

Similarly, NAT outside refers to those networks to which the stub network connects. They are not generally under the control of the organization. Hosts in outside networks can be subject to translation and can have local and global addresses.

NAT uses the following definitions:

- Local address—A local IP address that appears on the inside of a network.
- Global address—A global IP address that appears on the outside of a network.
- Inside local address—The IP address that is assigned to a host on the inside network. The address is probably not a legitimate IP address assigned by the Internet Network Information Center (InterNIC) or a service provider.
- Inside global address—A legitimate IP address (assigned by InterNIC or a service provider) that represents one or more inside local IP addresses to the outside world.

- Outside local address—The IP address of an outside host as it appears to the inside network. The address is not necessarily legitimate; it was allocated from the address space that is routable on the inside.
- Outside global address—The IP address that is assigned to a host on the outside network by the owner of the host. The address was allocated from a globally routable address or a network space.

# Pool Support for Dynamic NAT

Cisco NX-OS provides pool support for dynamic NAT. Dynamic NAT allows the configuration of a pool of global addresses that can be used to dynamically allocate a global address from the pool for every new translation. The addresses are returned to the pool after the session ages out or is closed. This allows for a more efficient use of addresses based on requirements.

Support for PAT includes the use of the global address pool. This further optimizes IP address utilization. PAT exhausts one IP address at a time with the use of port numbers. If no port is available from the appropriate group and more than one IP address is configured, PAT moves to the next IP address and gets the allocation based on the user defined pool (ignoring the source port or attempting to preserve it).

With dynamic NAT and PAT, each host uses a different address or port for each subsequent translation. The main difference between dynamic NAT and static NAT is that static NAT allows a remote host to initiate a connection to a translated host if an access list exists that allows it, while dynamic NAT does not.

# Guidelines and Limitations for Static and Dynamic NAT

Static NAT has the following configuration guidelines and limitations:

- NAT is supported for IPv4 Unicast only.
- If the translated IP is part of the outside interface subnet, then use the **ip proxy-arp** command on the NAT outside interface. If the **add-route** keyword is used, **ip proxy-arp** should be enabled.
- The Cisco Nexus device supports NAT on the following interface types:
  - Switch Virtual Interfaces (SVIs)
  - Physical layer 3 interfaces
  - Port channel layer 3 interfaces
- Non-TCP/ UDP packets are always software translated.
- **show** commands with the **internal** keyword are not supported.
- 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.
- When configuring a large number of translations (more than 100), it is faster to configure the translations before configuring the NAT interfaces.
- Twice NAT is not supported.
- Configuring NAT inside and outside rules together is not supported.
- NAT configurations such as IP NAT inside or IP NAT outside are not supported on loopback interfaces.

# Restrictions for Dynamic NAT

The following restrictions apply to dynamic Network Address Translation (NAT):

- **show** commands with the **internal** keyword are not supported.
- VXLAN routing is not supported on Cisco Nexus devices.
- Fragmented packets are not supported.
- Application layer gateway (ALG) translations are not supported. ALG, also known as application-level gateway, is an application that translates IP address information inside the payload of an application packet.
- Egress ACLs are not applied to translated packets.
- MIBs are not supported.
- Cisco Data Center Network Manager (DCNM) is not supported.
- Multiple global virtual device contexts (VDCs) are not supported on Cisco Nexus devices.
- Dynamic NAT translations are not synchronized with active and standby devices.
- Stateful NAT is not supported. However, NAT and Hot Standby Router Protocol (HSRP) can coexist.
- The timeout value for take up to the configured time-out + 119 seconds.
- For dynamic NAT, pool overload and interface overload are not supported for the outside NAT.
- The Cisco Nexus devices does not support NAT and VLAN Access Control Lists (VACLs) that are configured on an interface at the same time.
- NAT configurations such as ip nat inside or ip nat outside are not supported on loopback interfaces.

# Configuring Static NAT

## Enabling Static NAT

**Procedure**

| | Command or Action | Purpose |
|---|---|---|
| **Step 1** | switch# **configure terminal** | Enters global configuration mode. |
| **Step 2** | switch(config)# **feature nat** | Enables the static NAT feature on the device. |
| **Step 3** | switch(config)# **copy running-config startup-config** | Saves the change persistently through reboots and restarts by copying the running configuration to the startup configuration. |

# Configuring 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.<br>**Note**<br>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

The following two examples show how to configure NAT from the inside:

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

```
switch# configure terminal
switch(config)# interface vlan 100
switch(config-if)# ip nat inside
```

The following two examples show how to configure NAT from the outside:

```
switch# configure terminal
switch(config)# interface ethernet 1/5
switch(config-if)# ip nat outside
```

```
switch# configure terminal
switch(config)# interface vlan 102
switch(config-if)# ip nat ouside
```

# Enabling Static NAT for an Inside Source Address

For inside source translation, the source address of the packet gets translated from the inside to the outside interface. 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* [**vrf** *vrf-name*] [**match-in-vrf**] [**add-route**][**group** *group-id* ] | Configures static NAT to translate the inside local address to the inside global address or to translate the opposite (the inside global traffic to the inside local traffic). |
| **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 destination address gets translated from inside to outside interfaces. On the return traffic, the destination outside global IP address gets translated back to the outside local IP address.

### Procedure

| | Command or Action | Purpose |
|---|---|---|
| **Step 1** | switch# **configure terminal** | Enters global configuration mode. |
| **Step 2** | switch(config)# **ip nat outside source static** *outsideGlobalIP outsideLocalIP* [**vrf** *vrf-name* [**match-in-vrf**] [**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). When an inside |

| | Command or Action | Purpose |
|---|---|---|
| | | translation without ports is configured, an implicit add route is performed. The original add route functionality is an option while configuring an outside translation. |
| **Step 3** | (Optional) switch(config)# **copy running-config startup-config** | Saves the change persistently through reboots and restarts by copying the running configuration to the startup configuration. |

### Example

This example show how to configure static NAT for an outside source address:

```
switch# configure terminal
switch(config)# ip nat outside source static 2.2.2.2 6.6.6.6
switch(config)# copy running-config startup-config
```

# Configuring Static PAT for an Inside Source Address

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

### Procedure

| | Command or Action | Purpose |
|---|---|---|
| **Step 1** | switch# **configure terminal** | Enters global configuration mode. |
| **Step 2** | switch(config)# **ip nat inside source static** {*inside-local-address inside-global-address* \| {**tcp**\| **udp**} *inside-local-address* {*local-tcp-port* \| *local-udp-port*} *inside-global-address* {*global-tcp-port* \| *global-udp-port*}} {**vrf** *vrf-name* {**match-in-vrf**} } | 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*}} {**add-route**} {**vrf** *vrf-name* {**match-in-vrf**}} | 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
```

# Enabling and Disabling no-alias Configuration

NAT devices own Inside Global (IG) and Outside Local (OL) addresses and they are responsible for responding to any ARP requests directed to these addresses. When the IG/OL address subnet matches with the local interface subnet, NAT installs an IP alias and an ARP entry, in this case the device uses local-proxy-arp to respond to ARP requests.

The *no-alias* feature responds to ARP requests of all the translated IPs from a given NAT pool address range if the address range is in same subnet of the outside interface.

If no-alias is enabled on an interface with NAT configuration, the outside interface will not respond to any ARP requests in its subnet. When no-alias is disabled, the ARP requests for IPs in same subnet as of outside interface are served.



**Note** When you downgrade to any older releases that does not support this feature, configurations with *no-alias* option may be deleted.

### Procedure

| | Command or Action | Purpose |
|---|---|---|
| **Step 1** | switch# **configure terminal** | Enters global configuration mode. |
| **Step 2** | switch(config)# **feature nat** | Enables the static NAT feature on the device. |
| **Step 3** | switch(config)# **show run nat** | Displays NAT configuration. |
| **Step 4** | switch(config)# **show ip nat-alias** | Displays the information whether or not the alias is created.<br>**Note**<br>By default, alias is created. To disable the alias, you must append *no-alias* keyword to the command. |
| **Step 5** | switch(config)# **clear ip nat-alias** *ip address*/*all* | Removes entries from alias list. To remove a specific entry you must provide the IP address that you want to remove. To remove all entries, use the all keyword. |

### Example

This example shows the interface information:

```
switch# configure terminal
switch(config)# show ip int b
IP Interface Status for VRF "default"(1)
Interface            IP Address      Interface Status
Lo0                  100.1.1.1       protocol-up/link-up/admin-up
Eth1/1               7.7.7.1         protocol-up/link-up/admin-up
Eth1/3               8.8.8.1         protocol-up/link-up/admin-up
```

This example shows the running configuration:

```
switch# configure terminal
switch(config)# show running-config nat
!Command: show running-config nat
!Running configuration last done at: Thu Aug 23 11:57:01 2018
!Time: Thu Aug 23 11:58:13 2018

version 9.2(2) Bios:version 07.64
feature nat
interface Ethernet1/1
  ip nat inside
interface Ethernet1/3
  ip nat outside
switch(config)#
```

This example shows how to configure alias:

```
switch# configure terminal
switch(config)# ip nat inside source static 1.1.1.2 8.8.8.3
switch(config)# ip nat outside source static 2.2.2.1 7.7.7.3
switch(config)# show ip nat-alias
Alias Information for Context: default
```

```
Address         Interface
7.7.7.2          Ethernet1/1
8.8.8.2          Ethernet1/3
switch(config)#
```

This example shows the output of *show ip nat-alias*. By default, alias is enabled.

```
switch# configure terminal
switch(config)# show ip nat-alias
Alias Information for Context: default
Address         Interface
7.7.7.2          Ethernet1/1
8.8.8.2          Ethernet1/3
switch(config)#
```

This example shows how to disable alias:

```
switch# configure terminal
switch(config)# ip nat inside source static 1.1.1.2 8.8.8.3 no-alias
switch(config)# ip nat outside source static 2.2.2.1 7.7.7.3 no-alias
switch(config)# show ip nat-alias
Alias Information for Context: default
Address         Interface
7.7.7.2          Ethernet1/1
8.8.8.2          Ethernet1/3
switch(config)#

** None of the entry got appended as alias is disabled for above CLIs.
switch(config)#
```

This example shows how to clear alias. Use *clear ip nat-alias* to remove an entry from alias list. You can remove a single entry by specifying the IP address or remove all the alias entries.

```
switch# configure terminal
switch(config)# clear ip nat-alias address 7.7.7.2
switch(config)# show ip nat-alias
Alias Information for Context: default
Address         Interface
8.8.8.2          Ethernet1/3
switch(config)#
switch(config)# clear ip nat-alias all
switch(config)# show ip nat-alias
switch(config)#
```

# 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
```

# Verifying the Static NAT Configuration

To display the static NAT configuration, perform this task:

### Procedure

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

### Example

This example shows how to display the static NAT configuration:

```
switch# sh ip nat translations
Pro Inside global      Inside local       Outside local      Outside global

--- ---                ---                51.3.1.1           104.1.1.1
--- ---                ---                95.4.1.1           95.3.1.1
--- ---                ---                96.4.1.1           96.3.1.1
--- ---                ---                51.40.1.1          140.1.1.1
--- ---                ---                51.42.1.1          142.1.2.1
--- ---                ---                51.1.2.1           102.1.2.1
--- 11.1.1.1           101.1.1.1          ---                ---
--- 11.3.1.1           103.1.1.1          ---                ---
--- 11.39.1.1          139.1.1.1          ---                ---
--- 11.41.1.1          141.1.1.1          ---                ---
--- 95.1.1.1           149.1.1.1          ---                ---
--- 96.1.1.1           149.2.1.1          ---                ---
    130.1.1.1:590      30.1.1.100:5000    ---                ---
    130.2.1.1:590      30.2.1.100:5000    ---                ---
    130.3.1.1:590      30.3.1.100:5000    ---                ---
    130.4.1.1:590      30.4.1.100:5000    ---                ---
    130.1.1.1:591      30.1.1.101:5000    ---                ---


switch# sh ip nat translations verbose
Pro Inside global      Inside local       Outside local      Outside global
any ---                ---                22.1.1.3           22.1.1.2
    Flags:0x200009 time-left(secs):-1 id:0 state:0x0 grp_id:10
any 11.1.1.130         11.1.1.3           ---                ---
    Flags:0x1 time-left(secs):-1 id:0 state:0x0 grp_id:0
```

```
any 11.1.1.133         11.1.1.33          ---                ---
    Flags:0x1 time-left(secs):-1 id:0 state:0x0 grp_id:10
any 11.1.1.133         11.1.1.33          22.1.1.3           22.1.1.2
    Flags:0x200009 time-left(secs):-1 id:0 state:0x0 grp_id:0
tcp 10.1.1.100:64490   10.1.1.2:0         20.1.1.2:0         20.1.1.2:0
    Flags:0x82 time-left(secs):43192 id:31 state:0x3 grp_id:0 vrf: default
N3550T-1#
```

# Configuring Dynamic NAT

## Configuring Dynamic Translation and Translation Timeouts

### Procedure

| | Command or Action | Purpose |
|---|---|---|
| **Step 1** | **enable**<br>**Example:**<br>`Switch> enable` | Enables privileged EXEC mode.<br>• Enter your password if prompted. |
| **Step 2** | **configure terminal**<br>**Example:**<br>`Switch# configure terminal` | Enters global configuration mode. |
| **Step 3** | **ip access-list** *access-list-name*<br>**Example:**<br>`Switch(config)# ip access-list acl1` | Defines an access list and enters access-list configuration mode. |
| **Step 4** | **permit** *protocol source source-wildcard* **any**<br>**Example:**<br>`Switch(config-acl)# permit ip 10.111.11.0/24 any` | Sets conditions in an IP access list that permit traffic matching the conditions. |
| **Step 5** | **deny** *protocol source source-wildcard* **any**<br>**Example:**<br>`Switch(config-acl)# deny udp 10.111.11.100/32 any` | Sets conditions which disallows NAT translation. |
| **Step 6** | **exit**<br>**Example:**<br>`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* [**vrf** *vrf-name* [**match-in-vrf**] [**add-route**] [**overload**] | Establishes dynamic source translation by specifying the access list defined in Step 3. |

| | Command or Action | Purpose |
|---|---|---|
| | **Example:**<br>`Switch(config)# ip nat inside source list acl1 interface ethernet 1/1 overload` | |
| **Step 8** | **interface** *type number*<br>**Example:**<br>`Switch(config)# interface ethernet 1/4` | Configures an interface and enters interface configuration mode. |
| **Step 9** | **ip address** *ip-address mask*<br>**Example:**<br>`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**<br>**Example:**<br>`Switch(config-if)# ip nat inside` | Connects the interface to an inside network, which is subject to NAT.<br>**Note**<br>Configuration not supported on loopback interface. |
| **Step 11** | **exit**<br>**Example:**<br>`Switch(config-if)# exit` | Exits interface configuration mode and returns to global configuration mode. |
| **Step 12** | **interface** *type number*<br>**Example:**<br>`Switch(config)# interface ethernet 1/1` | Configures an interface and enters interface configuration mode. |
| **Step 13** | **ip address** *ip-address mask*<br>**Example:**<br>`Switch(config-if)# ip address 172.16.232.182 255.255.255.240` | Sets a primary IP address for an interface. |
| **Step 14** | **ip nat outside**<br>**Example:**<br>`Switch(config-if)# ip nat outside` | Connects the interface to an outside network.<br>**Note**<br>Configuration not supported on loopback interface. |
| **Step 15** | **exit**<br>**Example:**<br>`Switch(config-if)# exit` | Exits interface configuration mode and returns to global configuration mode. |
| **Step 16** | **ip nat translation max-entries** *number-of-entries*<br>**Example:** | Specifies the maximum number of dynamic NAT translations. The number of entries can be between 1 and 1023. |

| | Command or Action | Purpose |
|---|---|---|
| | `Switch(config)# ip nat translation max-entries 300` | |
| **Step 17** | **ip nat translation timeout** *seconds*<br>**Example:**<br>`switch(config)# ip nat translation timeout 13000` | Specifies the timeout value for dynamic NAT translations. |
| **Step 18** | **end**<br>**Example:**<br>`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 255.255.255.0
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)#
```

# 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.<br>Displays additional information for each translation table entry, including when an entry was created and used. |
| **show run nat** | Displays NAT configuration. |
| **show ip nat max** | Displays active Network Address Translation (NAT) maximum values. |
| **show ip nat statistics** | Monitor NAT statistics. |

### Example

This example shows how to display IP NAT Max values:

```
switch# show ip nat max

IP NAT Max values
====================
Max Dyn Translations:80
Max all-host:0
No.Static:0
No.Dyn:1
No.Dyn-ICMP:1
====================
Switch(config)#
```

This example shows how to display NAT Statistics:

```
switch# show ip nat statistics

IP NAT Statistics
================================================
Stats Collected since: Mon Feb 24 18:27:34 2020
------------------------------------------------
Total active translations: 1
No.Static: 0
No.Dyn:    1
No.Dyn-ICMP:   1
------------------------------------------------
Total expired Translations: 0
SYN timer expired:          0
FIN-RST timer expired:      0
Inactive timer expired:     0
------------------------------------------------
Total Hits:  2                  Total Misses:  2
In-Out Hits: 0                  In-Out Misses: 2
Out-In Hits: 2                  Out-In Misses: 0
------------------------------------------------
Total SW Translated Packets: 2
In-Out SW Translated: 2
Out-In SW Translated: 0
------------------------------------------------
Total SW Dropped Packets: 0
In-Out SW Dropped: 0
Out-In SW Dropped: 0

Address alloc. failure drop:      0
Port alloc.  failure drop:        0
Dyn. Translation max limit drop:  0
ICMP max limit drop:              0
Allhost max limit drop:           0
------------------------------------------------
Total TCP session established: 0
Total TCP session closed:      0
------------------------------------------------
NAT Inside Interfaces:  1
Ethernet1/34

NAT Outside Interfaces: 1
Ethernet1/32
------------------------------------------------
Inside source list:
++++++++++++++++++++

Access list: T2
RefCount: 1
Pool: T2     Overload
Total addresses: 10
```

```
Allocated: 1    percentage: 10%
Missed: 0


Outside source list:
++++++++++++++++++++
--------------------------------------------------
==================================================
Switch(config)#
Switch(config)#


**No.Dyn-ICMP  field is to display the no of icmp dynamic translations , its a subset of
"No.Dyn" field.
```



**Note** Beginning with Cisco NX-OS Release 10.2(3u), the **No.Dyn-ICMP** field is a subset of **No.Dyn** field and it displays the number of ICMP dynamic translations.

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
```

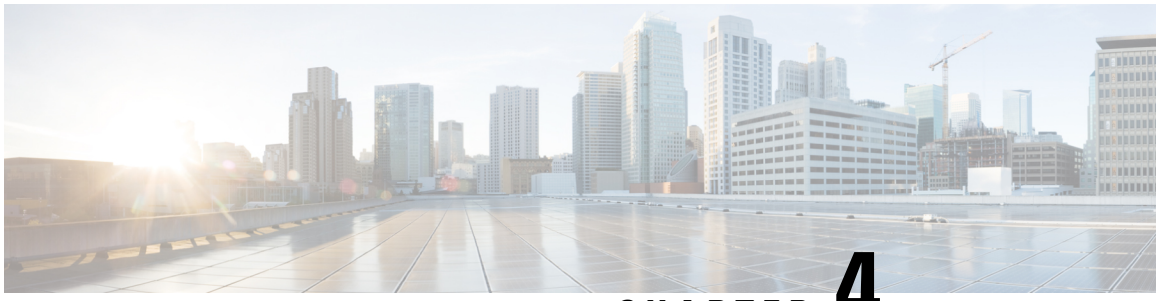
Outside pool without overload

```
switch# show ip nat translation
Pro  Inside global      Inside local       Outside local      Outside global
any  ---                ---                177.7.1.1:0        77.7.1.64:0
any  ---                ---                40.146.1.1:0       40.46.1.64:0
any  ---                ---                10.4.146.1:0       10.4.46.64:0
```

# Example: Configuring Dynamic Translation and Translation Timeouts

The following example shows how to configure dynamic overload Network Address Translation (NAT) by specifying an access list:

```
Switch> enable
Switch# configure terminal
Switch(config)# ip access-list acl1
Switch(config-acl)# permit ip 10.111.11.0/24 any
Switch(config-acl)# deny udp 10.111.11.100/32 any
Switch(config-acl)# exit
Switch(config)# ip nat inside source list acl1 interface ethernet 1/1 overload
Switch(config)# interface ethernet 1/4
Switch(config-if)# ip address 10.111.11.39 255.255.255.0
Switch(config-if)# ip nat inside
Switch(config-if)# exit
Switch(config)# interface ethernet 1/1
Switch(config-if)# ip address 172.16.232.182 255.255.255.240
Switch(config-if)# ip nat outside
Switch(config-if)# exit
Switch(config)# ip nat translation max-entries 300
Switch(config)# ip nat translation timeout 13000
Switch(config)# end
```



CHAPTER 4

# Configuring Layer 2 Interfaces


- Information About Access and Trunk Interfaces, on page 27
- Prerequisites for Layer 2 Interfaces, on page 31
- Guidelines and Limitations for Layer 2 Interfaces, on page 31
- Default Settings for Layer 2 Interfaces, on page 32
- Configuring Access and Trunk Interfaces , on page 33
- Verifying the Interface Configuration, on page 41
- Monitoring the Layer 2 Interfaces, on page 42
- Configuration Examples for Access and Trunk Ports, on page 43
- Related Documents, on page 43


# Information About Access and Trunk Interfaces

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

## About Access and Trunk Interfaces

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

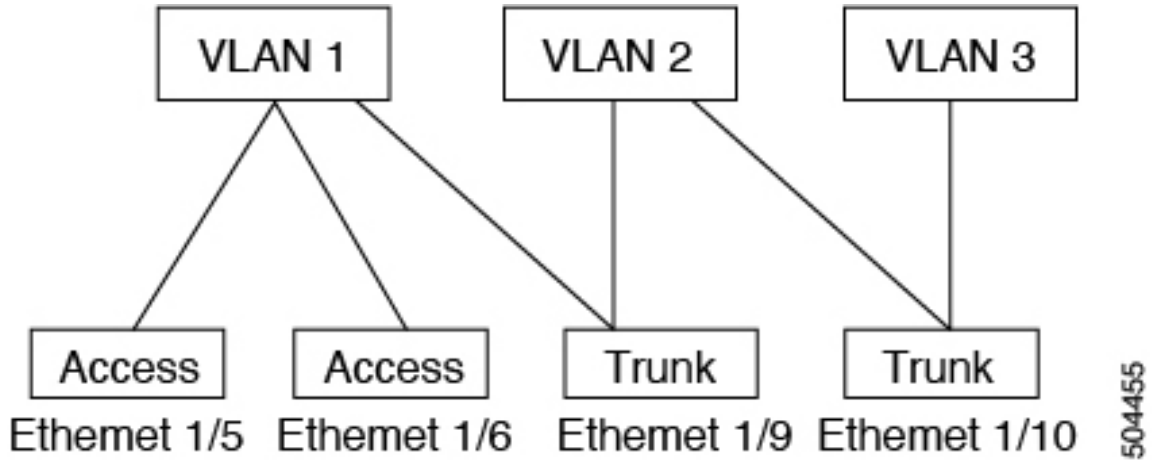
- An access port can have only one VLAN configured on that port; it can carry traffic for only one VLAN.
- A trunk port can have two or more VLANs configured on that port; it can carry traffic for several VLANs simultaneously.

By default, all the ports on the Cisco Nexus® 3550-T switches are Layer 3 ports/Layer 2 ports.

You can make all ports Layer 2 ports using the setup script or by entering the **system default switchport**command. See the *Cisco Nexus® 3550-T Fundamentals Configuration* section for information about using the setup script. To configure the port as a Layer 2 port using the CLI, use the **switchport** command.

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

*Figure 2: Trunk and Access Ports and VLAN Traffic*




**Note** See the *Cisco Nexus® 3550-T Layer 2 Switching Configuration* section for information about VLANs.

In order to correctly deliver the traffic on a trunk port with several VLANs, the device uses the IEEE 802.1Q encapsulation, or tagging, method (see the "IEEE 802.1Q Encapsulation" section for more information).

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

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

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

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

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

# IEEE 802.1Q Encapsulation

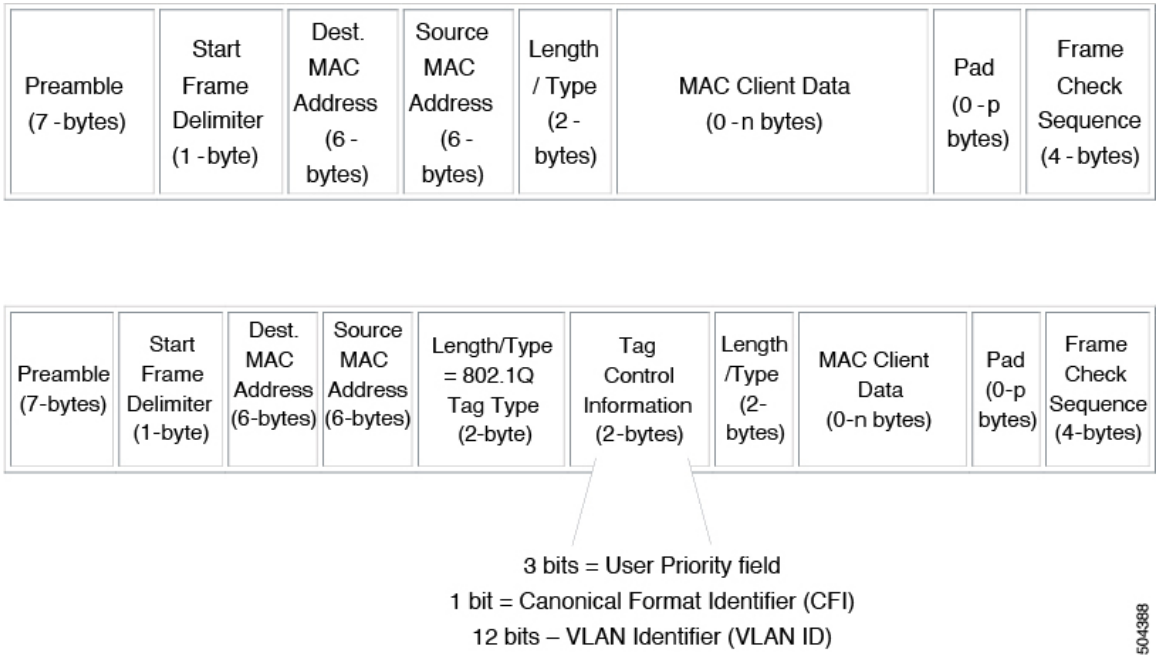
**Note** For information about VLANs, see the *Cisco Nexus® 3550-T Layer 2 Switching Configuration* section.

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

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

the VLANs. Also, the encapsulated VLAN tag allows the trunk to move traffic end-to-end through the network on the same VLAN.

***Figure 3: Header Without and With 802.1Q Tag***




# Access VLANs

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

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

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

# Native VLAN IDs for Trunk Ports

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

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

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

# Allowed VLANs

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

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



**Note** See the *Cisco Nexus® 3550-T Layer 2 Switching Configuration* section for more information about STP.

# Default Interfaces

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



**Note** All 48 ports can be selected for the default interface.

# Switch Virtual Interface and Autostate Behavior

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

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

# Counter Values

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

| Configuration | Packet Size | Incremented Counters | Traffic |
|---|---|---|---|
| L2 port | <1500 | Input error | Dropped |

**Note** Greater than 64 bytes packet with bad CRC–The CRC counter increments.

# Prerequisites for Layer 2 Interfaces

Layer 2 interfaces have the following prerequisites:

- You are logged onto the device.
- By default, Cisco NX-OS configures Layer 3 parameters. If you want to configure Layer 2 parameters, you need to switch the port mode to Layer 2. You can change the port mode by using the **switchport** command.
- You must configure the port as a Layer 2 port before you can use the **switchport mode** command. By default, all ports on the device are Layer 3 ports. By default, all ports on the Cisco Nexus® 3550-T devices are Layer 2 ports.

# Guidelines and Limitations for Layer 2 Interfaces

VLAN trunking has the following configuration guidelines and limitations:

- A port can be either a Layer 2 or a Layer 3 interface; it cannot be both simultaneously.
- When you change a Layer 3 port to a Layer 2 port or a Layer 2 port to a Layer 3 port, all layer-dependent configuration is lost. When you change an access or trunk port to a Layer 3 port, all information about the access VLAN, native VLAN, allowed VLANs, and so forth, is lost.
- Do not connect devices with access links because access links may partition a VLAN.
- When connecting Cisco devices through an 802.1Q trunk, make sure that the native VLAN for an 802.1Q trunk is the same on both ends of the trunk link. If the native VLAN on one end of the trunk is different from the native VLAN on the other end, spanning tree loops might result.
- Disabling spanning tree on the native VLAN of an 802.1Q trunk without disabling spanning tree on every VLAN in the network can cause spanning tree loops. You must leave spanning tree enabled on the native VLAN of an 802.1Q trunk. If you cannot leave spanning tree enabled, you must disable spanning tree on every VLAN in the network. Make sure that your network has no physical loops before you disable spanning tree.
- When you connect two Cisco devices through 802.1Q trunks, the devices exchange spanning tree bridge protocol data units (BPDUs) on each VLAN allowed on the trunks. The BPDUs on the native VLAN of the trunk are sent untagged to the reserved IEEE 802.1D spanning tree multicast MAC address (01-80-C2-00-00-00). The BPDUs on all other VLANs on the trunk are sent tagged to the reserved Cisco Shared Spanning Tree (SSTP) multicast MAC address (01-00-0c-cc-cc-cd).
- Because Cisco devices transmit BPDUs to the SSTP multicast MAC address on VLANs other than the native VLAN of the trunk, non-Cisco devices do not recognize these frames as BPDUs and flood them on all ports in the corresponding VLAN. Other Cisco devices connected to the non-Cisco 802.1Q cloud receive these flooded BPDUs. This BPDU reception allows Cisco switches to maintain a per-VLAN spanning tree topology across a cloud of non-Cisco 802.1Q devices. The non-Cisco 802.1Q cloud that

separates the Cisco devices is treated as a single broadcast segment between all devices connected to the non-Cisco 802.1Q cloud through 802.1Q trunks.

- Make certain that the native VLAN is the same on all of the 802.1Q trunks that connect the Cisco devices to the non-Cisco 802.1Q cloud.
- If you are connecting multiple Cisco devices to a non-Cisco 802.1Q cloud, all of the connections must be through 802.1Q trunks. You cannot connect Cisco devices to a non-Cisco 802.1Q cloud through access ports because doing so places the access port on the Cisco device into the spanning tree "port inconsistent" state and no traffic will pass through the port.
- You can group trunk ports into port-channel groups, but all trunks in the group must have the same configuration. When a group is first created, all ports follow the parameters set for the first port to be added to the group. If you change the configuration of one of these parameters, the device propagates that setting to all ports in the group, such as the allowed VLANs and the trunk status. For example, if one port in a port group ceases to be a trunk, all ports cease to be trunks.
- When MAC addresses are cleared on a VLAN with the clear mac address-table dynamic command, the dynamic ARP (Address Resolution Protocol) entries on that VLAN are refreshed.
- If a static ARP entry exists on the VLAN and no MAC address to port mapping is present, the supervisor may generate an ARP request to learn the MAC address. Upon learning the MAC address, the adjacency entry points to the correct physical port.
- Cisco NX-OS does not support transparent bridging between two VLANs when one of the SVIs is on the Cisco Nexus 3550-T using the BIA MAC (burned-in MAC address). This occurs when the BIA MAC is shared between SVIs/VLANs. A MAC, different from the BIA MAC, can be configured under the SVI for transparent bridging to work properly.
- You may get an error message when you attempt to configure the interface mode to trunk and trunk VLANs simultaneously. On Cisco NX-OS interfaces, the default value of interface mode is access. To implement any trunk related configurations, you must first change the interface mode to trunk and then configure the trunk VLAN ranges.
- Spanning of VLAN tagged packets is not supported on Cisco Nexus® 3550-T switches.

# Default Settings for Layer 2 Interfaces

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

***Table 2: Default Access and Trunk Port Mode Parameters***

| Parameters | Default |
|---|---|
| Switchport mode | Access |
| Allowed VLANs | 1 to 3967<br>**Note**<br>A maximum of 255 VLANs are supported. |
| Access VLAN ID | VLAN1 |

| Parameters | Default |
|---|---|
| Native VLAN ID | VLAN1 |
| Native VLAN ID tagging | Disabled |
| Administrative state | Shut |
| SVI autostate | Enabled |

# Configuring Access and Trunk Interfaces

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

## Configuring a Layer 2 Access Port

You can configure a Layer 2 port as an access port. An access port transmits packets on only one, untagged VLAN. You specify which VLAN traffic that the interface carries, which becomes the access VLAN. If you do not specify a VLAN for an access port, that interface carries traffic only on the default VLAN. The default VLAN is VLAN1.

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

### Before you begin

Ensure that you are configuring a Layer 2 interface.

### Procedure

| | Command or Action | Purpose |
|---|---|---|
| **Step 1** | **configure terminal**<br>**Example:**<br>`switch# configure terminal`<br>`switch(config)#` | Enters global configuration mode. |
| **Step 2** | **interface ethernet** {{*type slot/port*} \| {**port-channel** *number*}}<br>**Example:**<br>`switch(config)# interface ethernet 1/5`<br>`switch(config-if)#` | Specifies an interface to configure, and enters interface configuration mode. |
| **Step 3** | **switchport mode** [**access** \| **trunk**]<br>**Example:** | Sets the interface as a nontrunking nontagged, single-VLAN Layer 2 interface. An access port |

| | Command or Action | Purpose |
|---|---|---|
| | `switch(config-if)# switchport mode access` | can carry traffic in one VLAN only. By default, an access port carries traffic for VLAN1; to set the access port to carry traffic for a different VLAN, use the **switchport access vlan** command. |
| **Step 4** | **switchport access vlan** *vlan-id*<br>**Example:**<br>`switch(config-if)# switchport access vlan 5` | Specifies the VLAN for which this access port will carry traffic. If you do not enter this command, the access port carries traffic on VLAN1 only; use this command to change the VLAN for which the access port carries traffic. |
| **Step 5** | **exit**<br>**Example:**<br>`switch(config-if)# exit`<br>`switch(config)#` | Exits the interface configuration mode. |
| **Step 6** | **show interface**<br>**Example:**<br>`switch# show interface` | (Optional) Displays the interface status and information. |
| **Step 7** | **no shutdown**<br>**Example:**<br>`switch# configure terminal`<br>`switch(config)# int e1/5`<br>`switch(config-if)# no shutdown` | (Optional) Clears the errors on the interfaces and VLANs where policies correspond with hardware policies. This command allows policy programming to continue and the port to come up. If policies do not correspond, the errors are placed in an error-disabled policy state. |
| **Step 8** | **copy running-config startup-config**<br>**Example:**<br>`switch(config)# copy running-config startup-config` | (Optional) Copies the running configuration to the startup configuration. |

### Example

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

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

# Configuring Access Host Ports

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

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

### Before you begin

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

### Procedure

| | Command or Action | Purpose |
|---|---|---|
| **Step 1** | **configure terminal**<br>**Example:**<br>`switch# configure terminal`<br>`switch(config)#` | Enters global configuration mode. |
| **Step 2** | **interface ethernet** *type slot/port*<br>**Example:**<br>`switch(config)# interface ethernet 1/3`<br>`switch(config-if)#` | Specifies an interface to configure, and enters interface configuration mode. |
| **Step 3** | **switchport host**<br>**Example:**<br>`switch(config-if)# switchport host` | Sets the interface to be an access host port, which immediately moves to the spanning tree forwarding state and disables port channeling on this interface.<br>**Note**<br>Apply this command only to end stations. |
| **Step 4** | **exit**<br>**Example:**<br>`switch(config-if-range)# exit`<br>`switch(config)#` | Exits the interface mode. |
| **Step 5** | **show interface**<br>**Example:**<br>`switch# show interface` | (Optional) Displays the interface status and information. |
| **Step 6** | **no shutdown**<br>**Example:** | (Optional) Clears the errors on the interfaces and VLANs where policies correspond with hardware policies. This command allows policy |

| | Command or Action | Purpose |
|---|---|---|
| | `switch# configure terminal`<br>`switch(config)# int e1/3`<br>`switch(config-if)# no shutdown` | programming to continue and the port to come up. If policies do not correspond, the errors are placed in an error-disabled policy state. |
| **Step 7** | **copy running-config startup-config**<br>**Example:**<br>`switch(config)# copy running-config startup-config` | (Optional) Copies the running configuration to the startup configuration. |

### Example

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

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

# Configuring Trunk Ports

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

**Note** The device supports 802.1Q encapsulation only.

### Before you begin

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

### Procedure

| | Command or Action | Purpose |
|---|---|---|
| **Step 1** | **configure terminal**<br>**Example:**<br>`switch# configure terminal`<br>`switch(config)#` | Enters global configuration mode. |
| **Step 2** | **interface** {*type slot/port* \| **port-channel** *number*}<br>**Example:**<br>`switch(config)# interface ethernet 1/4`<br>`switch(config-if)#` | Specifies an interface to configure, and enters interface configuration mode. |

| | Command or Action | Purpose |
|---|---|---|
| **Step 3** | **switchport mode** [**access** \| **trunk**]<br>**Example:**<br>`switch(config-if)# switchport mode trunk` | Sets the interface as a Layer 2 trunk port. A trunk port can carry traffic in one or more VLANs on the same physical link (VLANs are based on the trunk-allowed VLANs list). By default, a trunk interface can carry traffic for all VLANs. To specify that only certain VLANs are allowed on the specified trunk, use the **switchport trunk allowed vlan** command. |
| **Step 4** | **exit**<br>**Example:**<br>`switch(config-if)# exit`<br>`switch(config)#` | Exits the interface mode. |
| **Step 5** | **show interface**<br>**Example:**<br>`switch# show interface` | (Optional) Displays the interface status and information. |
| **Step 6** | **no shutdown**<br>**Example:**<br>`switch# configure terminal`<br>`switch(config)# int e1/4`<br>`switch(config-if)# no shutdown` | (Optional) Clears the errors on the interfaces and VLANs where policies correspond with hardware policies. This command allows policy programming to continue and the port to come up. If policies do not correspond, the errors are placed in an error-disabled policy state. |
| **Step 7** | **copy running-config startup-config**<br>**Example:**<br>`switch(config)# copy running-config startup-config` | (Optional) Copies the running configuration to the startup configuration. |

### Example

This example shows how to set Ethernet 1/4 as a Layer 2 trunk port:

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

# Configuring the Allowed VLANs for Trunking Ports

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

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

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

### Before you begin

Before you configure the allowed VLANs for the specified trunk ports, ensure that you are configuring the correct interfaces and that the interfaces are trunks.

## Procedure

| | Command or Action | Purpose |
|---|---|---|
| **Step 1** | **configure terminal**<br>**Example:**<br>`switch# configure terminal`<br>`switch(config)#` | Enters global configuration mode. |
| **Step 2** | **interface** {**ethernet** *slot/port* \| **port-channel** *number*}<br>**Example:**<br>`switch(config)# interface ethernet 1/3` | Specifies an interface to configure, and enters interface configuration mode. |
| **Step 3** | **switchport trunk allowed vlan** {*vlan-list* **add** *vlan-list* \| **all** \| **except** *vlan-list* \| **none** \| **remove** *vlan-list*}<br>**Example:**<br>`switch(config-if)# switchport trunk allowed vlan add 15-20#` | Sets the allowed VLANs for the trunk interface. The default is to allow all VLANs on the trunk interface: 1 to 3967 and 4048 to 4094. Only 255 VLANs are supported on the Cisco Nexus 3550-T switch.<br>**Note**<br>You cannot add internally allocated VLANs as allowed VLANs on trunk ports. The system returns a message if you attempt to list an internally allocated VLAN as an allowed VLAN. |
| **Step 4** | **exit**<br>**Example:**<br>`switch(config-if)# exit`<br>`switch(config)#` | Exits the interface mode. |
| **Step 5** | **show vlan**<br>**Example:**<br>`switch# show vlan` | (Optional) Displays the status and information for VLANs. |

| | Command or Action | Purpose |
|---|---|---|
| **Step 6** | **no shutdown**<br>**Example:**<br>`switch# configure terminal`<br>`switch(config)# int e1/3`<br>`switch(config-if)# no shutdown` | (Optional) Clears the errors on the interfaces and VLANs where policies correspond with hardware policies. This command allows policy programming to continue and the port to come up. If policies do not correspond, the errors are placed in an error-disabled policy state. |
| **Step 7** | **copy running-config startup-config**<br>**Example:**<br>`switch(config)# copy running-config startup-config` | (Optional) Copies the running configuration to the startup configuration. |

### Example

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

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

# Configuring a Default Interface

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

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

### Procedure

| | Command or Action | Purpose |
|---|---|---|
| **Step 1** | **configure terminal**<br>**Example:**<br>`switch# configure terminal`<br>`switch(config)#` | Enters global configuration mode. |
| **Step 2** | **default interface** *int-if* [**checkpoint** *name*]<br>**Example:**<br>`switch(config)# default interface ethernet 1/3 checkpoint test8` | Deletes the configuration of the interface and restores the default configuration. Use the **?** keyword to display the supported interfaces. |

| | Command or Action | Purpose |
|---|---|---|
| | | Use the **checkpoint** keyword to store a copy of the running configuration of the interface before clearing the configuration. |
| **Step 3** | **exit**<br>**Example:**<br>`switch(config)# exit`<br>`switch(config)#` | Exits global configuration mode. |
| **Step 4** | **show interface**<br>**Example:**<br>`switch# show interface` | (Optional) Displays the interface status and information. |
| **Step 5** | **no shutdown**<br>**Example:**<br>`switch# configure terminal`<br>`switch(config)# int e1/3`<br>`switch(config-if)# no shutdown` | (Optional) Clears the errors on the interfaces and VLANs where policies correspond with hardware policies. This command allows policy programming to continue and the port to come up. If policies do not correspond, the errors are placed in an error-disabled policy state. |

### Example

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

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

# Changing the System Default Port Mode to Layer 2

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

### Procedure

| | Command or Action | Purpose |
|---|---|---|
| **Step 1** | **configure terminal**<br>**Example:**<br>`switch# configure terminal`<br>`switch(config)#` | Enters global configuration mode. |
| **Step 2** | **system default switchport** [**shutdown**]<br>**Example:**<br>switch(config-if)# **system default switchport** | Sets the default port mode for all interfaces on the system to Layer 2 access port mode and enters interface configuration mode. By default, all the interfaces are Layer 3. |

| | Command or Action | Purpose |
|---|---|---|
| | | **Note**<br>When the **system default switchport shutdown** command is issued:<br>• Any Layer 2 port that is not specifically configured with **no shutdown** are shutdown. To avoid the shutdown, configure the Layer 2 port with **no shut** |
| **Step 3** | **exit**<br>**Example:**<br>`switch(config-if)# exit`<br>`switch(config)#` | Exits the interface configuration mode. |
| **Step 4** | **show interface brief**<br>**Example:**<br>`switch# show interface brief` | (Optional) Displays the status and information for interfaces. |
| **Step 5** | **no shutdown**<br>**Example:**<br>`switch# configure terminal`<br>`switch(config)# int e1/3`<br>`switch(config-if)# no shutdown` | (Optional) Clears the errors on the interfaces and VLANs where policies correspond with hardware policies. This command allows policy programming to continue and the port to come up. If policies do not correspond, the errors are placed in an error-disabled policy state. |
| **Step 6** | **copy running-config startup-config**<br>**Example:**<br>`switch(config)# copy running-config startup-config` | (Optional) Copies the running configuration to the startup configuration. |

### Example

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

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

# Verifying the Interface Configuration

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

| Command | Purpose |
|---|---|
| **show interface ethernet** *slot/port* [**brief** \| \| **counters** \| **debounce** \| **description** \| **flowcontrol** \| **mac-address** \| **status** \| **transceiver**] | Displays the interface configuration. |

| Command | Purpose |
|---|---|
| **show interface brief** | Displays interface configuration information, including the mode. |
| **show interface switchport** | Displays information, including access and trunk interface, information for all Layer 2 interfaces. |
| **show interface trunk** [**module** *module-number* \| **vlan** *vlan-id*] | Displays trunk configuration information. |
| **show interface capabilities** | Displays information about the capabilities of the interfaces. |
| **show running-config** [**all**] | Displays information about the current configuration.<br>The **all** command displays the default and current configurations. |
| **show running-config interface ethernet** *slot/port* | Displays configuration information about the specified interface. |
| **show running-config interface port-channel** *slot/port* | Displays configuration information about the specified port-channel interface. |
| **show running-config interface vlan** *vlan-id* | Displays configuration information about the specified VLAN interface. |

# Monitoring the Layer 2 Interfaces

Use the following commands to display Layer 2 interfaces:

| Command | Purpose |
|---|---|
| **clear counters interface** [**interface**] | Clears the counters. |
| **load- interval** {**counter** {**1** \| **2** \| **3**}} *seconds* | Cisco Nexus 3550-T devices set three different sampling intervals to bit-rate and packet-rate statistics. |
| **show interface counters** [**module** *module*] | Displays input and output octets unicast packets, multicast packets, and broadcast packets. |
| **show interface counters detailed** [**all**] | Displays input packets, bytes, and multicast as well as output packets and bytes.<br>**Note**<br>Ignore *Output Dropped Errors* as it represents the cumulative ingress drops in the traffic that is directed to the port. The ingress drops on any port are displayed as part of *Input Discard Errors.* |

| Command | Purpose |
|---|---|
| **show interface counters errors** [**module** *module*] | Displays information on the number of error packets.<br>**Note**<br>Ignore *OutDiscards* as it represents the cumulative ingress drops in the traffic that is directed to the port. The ingress drops on any port are displayed as part of *InDiscards*. |

# Configuration Examples for Access and Trunk Ports

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

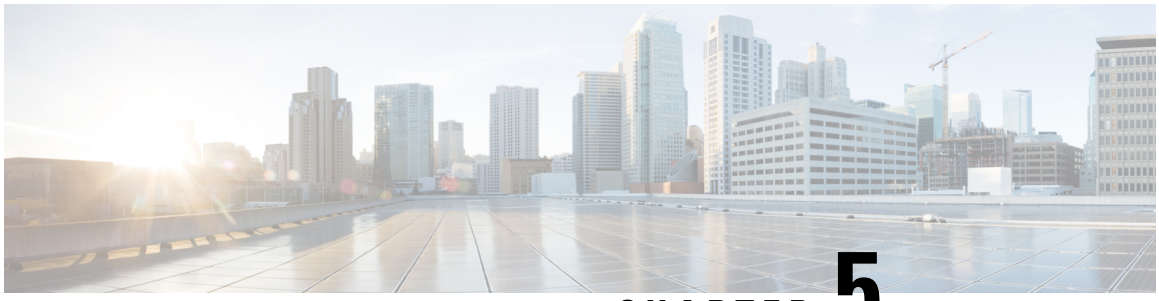
```
switch# configure terminal
switch(config)# interface ethernet 1/30
switch(config-if)# switchport
switch(config-if)# switchport mode access
switch(config-if)# switchport access vlan 5
switch(config-if)#
```

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

```
switch# configure terminal
switch(config)# interface ethernet 1/35
switch(config-if)# switchport
switch(config-if)# switchport mode trunk
switch(config-if)# switchport trunk native vlan 10
switch(config-if)# switchport trunk allowed vlan 5, 10
switch(config-if)# exit
```

# Related Documents

| Related Documents | Document Title |
|---|---|
| Configuring Layer 3 interfaces | *Configuring Layer 2 Interfaces* section |
| Port Channels | *Configuring Port Channels* section |
| System management | *Cisco Nexus® 3550-T System Management Configuration* chapter |
| High availability | *Cisco Nexus® Series High Availability and Redundancy Guide* |
| Licensing | *Cisco NX-OS Licensing Guide* |
| Release Notes | *Cisco Nexus® Series NX-OS Release Notes* |



CHAPTER 5

# Configuring Port Channels


- About Port Channels, on page 45
- Port Channels, on page 46
- Port-Channel Interfaces, on page 47
- Basic Settings, on page 48
- Compatibility Requirements, on page 48
- Load Balancing Using Port Channels, on page 50
- LACP, on page 51
- Prerequisites for Port Channeling, on page 55
- Guidelines and Limitations, on page 55
- Default Settings, on page 56
- Configuring Port Channels, on page 57


## About Port Channels

A port channel is an aggregation of multiple physical interfaces that creates a logical interface. You can bundle up to 8 individual active links into a port channel to provide increased bandwidth and redundancy. Port channeling also load balances traffic across these physical interfaces. The port channel stays operational as long as at least one physical interface within the port channel is operational.

The allowed system wide limit for maximum port channel is 48. There are 6 port groups with each port group containing 8 ports. You can create a maximum of 8 port channels per port group. You can have a maximum of 8 physical ports, that belong to the same port group, bundled into a given port channel. You cannot create port channels from ports that belong to different port groups.

***Table 3: Port group name and port group range***

| Port group name | Port group range |
|---|---|
| 1 | 1/1 to 1/8 |
| 2 | 1/9 to 1/16 |
| 3 | 1/17 to 1/24 |
| 4 | 1/25 to 1/32 |
| 5 | 1/33 to 1/40 |

| Port group name | Port group range |
|---|---|
| 6 | 1/41 to 1/48 |

You can create a Layer 2 port channel by bundling compatible Layer 2 interfaces, or you can create Layer 3 port channels by bundling compatible Layer 3 interfaces. You cannot combine Layer 2 and Layer 3 interfaces in the same port channel.

You can also change the port channel from Layer 3 to Layer 2. See the *Configuring Layer 2 Interfaces* chapter for information about creating Layer 2 interfaces.

A Layer 2 port channel interface and it's member ports can have different STP parameters. Changing the STP parameters of the port channel does not impact the STP parameters of the member ports because a port channel interface takes precedence if the member ports are bundled.

**Note** After a Layer 2 port becomes part of a port channel, all switchport configurations must be done on the port channel; you can no longer apply switchport configurations to individual port-channel members. You cannot apply Layer 3 configurations to an individual port-channel member either; you must apply the configuration to the entire port channel.

You can use static port channels, with no associated aggregation protocol, for a simplified configuration.

For more flexibility, you can use the Link Aggregation Control Protocol (LACP), which is defined in IEEE 802.3ad. When you use LACP, the link passes protocol packets. You cannot configure LACP on shared interfaces.

See the *LACP Overview* section for information about LACP.

# Port Channels

A port channel bundles physical links into a channel group to create a single logical link that provides the aggregate bandwidth of up to 8 physical links. If a member port within a port channel fails, the traffic previously carried over the failed link switches to the remaining member ports within the port channel.

However, you can enable the LACP to use port channels more flexibly. Configuring port channels with LACP and static port channels require a slightly different procedure (see the *Configuring Port Channels* section).

**Note** The device does not support Port Aggregation Protocol (PAgP) for port channels.

Each port can be in only one port channel. All the ports in a port channel must be compatible; they must use the same speed and duplex mode (see the *Compatibility Requirements* section). When you run static port channels with no aggregation protocol, the physical links are all in the on channel mode; you cannot change this mode without enabling LACP (see the *Port-Channel Modes* section).

You can create port channels directly by creating the port-channel interface, or you can create a channel group that acts to aggregate individual ports into a bundle. When you associate an interface with a channel group, the software creates a matching port channel automatically if the port channel does not already exist. In this instance, the port channel assumes the Layer 2 or Layer 3 configuration of the first interface. You can also create the port channel first. In this instance, the Cisco NX-OS software creates an empty channel group with

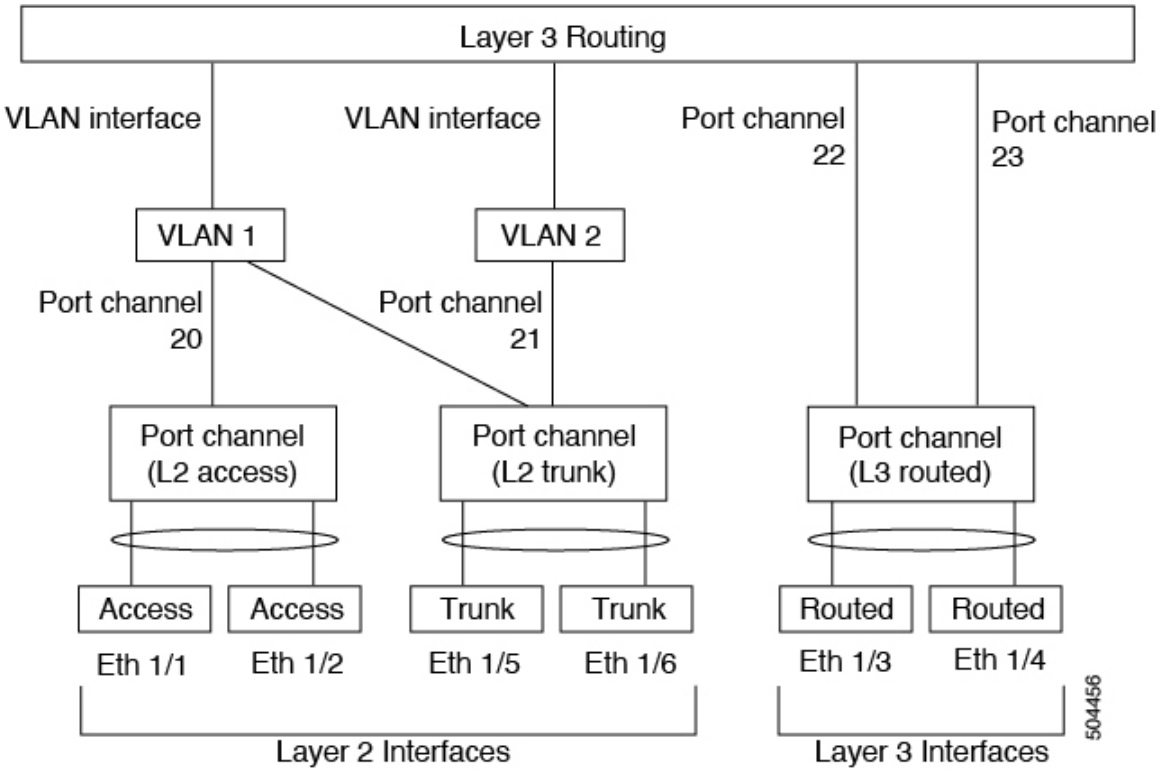
the same channel number as the port channel and takes the default Layer 2 or Layer 3 configuration, as well as the compatibility configuration (see the *Compatibility Requirements* section).

**Note** The port channel is operationally up when at least one of the member ports is up and that port's status is channeling. The port channel is operationally down when all member ports are operationally down.

# Port-Channel Interfaces

The following shows port-channel interfaces.

***Figure 4: Port-Channel Interfaces***




You can classify port-channel interfaces as Layer 2 or Layer 3 interfaces. In addition, you can configure Layer 2 port channels in either access or trunk mode. Layer 3 port-channel interfaces have routed ports as channel members.

You can configure a Layer 3 port channel with a static MAC address. If you do not configure this value, the Layer 3 port channel uses the router MAC of the first channel member to come up. See the *Cisco Nexus® 3550-T Layer 2 Switching Configuration* section for information about configuring static MAC addresses on Layer 3 port channels.

See the *Configuring Layer 2 Interfaces* chapter for information about configuring Layer 2 ports in access or trunk mode and the *Configuring Layer 3 Interfaces* chapter for information about configuring Layer 3 interfaces.

# Basic Settings

You can configure the following basic settings for the port-channel interface:

- Bandwidth—Use this setting for informational purposes only; this setting is to be used by higher-level protocols.
- Delay—Use this setting for informational purposes only; this setting is to be used by higher-level protocols.
- Description
- Duplex
- IP addresses
- Shutdown
- Speed

# Compatibility Requirements

When you add an interface to a channel group, the software checks certain interface attributes to ensure that the interface is compatible with the channel group. For example, you cannot add a Layer 3 interface to a Layer 2 channel group. The Cisco NX-OS software also checks a number of operational attributes for an interface before allowing that interface to participate in the port-channel aggregation.

The compatibility check includes the following operational attributes:

- Network layer
- (Link) speed capability
- Speed configuration
- Duplex capability
- Duplex configuration
- Port mode
- Access VLAN
- Trunk native VLAN
- Tagged or untagged
- Allowed VLAN list
- Flow-control capability
- Flow-control configuration
- Media type, either copper or fiber

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

You can only add interfaces configured with the channel mode set to on to static port channels, and you can only add interfaces configured with the channel mode as active or passive to port channels that are running LACP. You can configure these attributes on an individual member port. If you configure a member port with an incompatible attribute, the software suspends that port in the port channel.

Alternatively, you can force ports with incompatible parameters to join the port channel if the following parameters are the same:

- (Link) speed capability
- Speed configuration
- Duplex capability
- Duplex configuration
- Flow-control capability
- Flow-control configuration

When the interface joins a port channel, some of its individual parameters are removed and replaced with the values on the port channel as follows:

- Bandwidth
- Delay
- Extended Authentication Protocol over UDP
- VRF
- IP address
- MAC address
- Spanning Tree Protocol

Many interface parameters remain unaffected when the interface joins or leaves a port channel as follows:

- Beacon
- Description
- CDP
- LACP port priority
- Debounce
- UDLD
- MDIX
- Rate mode
- Shutdown
- SNMP trap

**Note** When you delete the port channel, the software sets all member interfaces as if they were removed from the port channel.

# Load Balancing Using Port Channels

The Cisco NX-OS software load balances traffic across all operational interfaces in a port channel by hashing the addresses in the frame to a numerical value that selects one of the links in the channel. Port channels provide load balancing by default. Port-channel load balancing uses MAC addresses, or IP addresses to select the link. Port-channel load balancing uses either source or destination addresses, or both source and destination addresses.

You can configure the load- balancing mode to apply to all port channels that are configured on the entire device. You can configure one load-balancing mode for the entire device. You cannot configure the load-balancing method per port channel.

You can configure the type of load-balancing algorithm used. You can choose the load-balancing algorithm that determines which member port to select for egress traffic by looking at the fields in the frame.

The default load-balancing mode for Layer 3 interfaces is the source and destination IP L4 ports, and the default load-balancing mode for non-IP traffic is the source and destination MAC address. Use the **port-channel load-balance** command to set the load-balancing method among the interfaces in the channel-group bundle. The default method for Layer 2 packets is src-dst-mac. The default method for Layer 3 packets is src-dst IP.

You can configure the device to use one of the following methods to load balance across the port channel:

- Destination IP address
- Source IP address
- Source and destination IP address
- Source MAC address
- Destination MAC address
- Source and destination MAC address

Non-IP and Layer 3 port channels both follow the configured load-balancing method, using the source, destination, or source and destination parameters. For example, when you configure load balancing to use the source IP address, all non-IP traffic uses the source MAC address to load balance the traffic while the Layer 3 traffic load balances the traffic using the source IP address. Similarly, when you configure the destination MAC address as the load-balancing method, all Layer 3 traffic uses the destination IP address while the non-IP traffic load balances using the destination MAC address.

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

The multicast traffic uses the following methods for load balancing with port channels:

- Multicast traffic with Layer 4 information—Source IP address, source port, destination IP address, destination port
- Multicast traffic without Layer 4 information—Source IP address, destination IP address

- Non-IP multicast traffic—Source MAC address, destination MAC address

**Note** Devices that run Cisco IOS can optimize the behavior of the member ports ASICs if a failure of a single member occurred by running the port-channel hash-distribution command. The Cisco Nexus 3550-T device performs this optimization by default and does not require or support this command. Cisco NX-OS does support the customization of the load-balancing criteria on port channels through the port-channel load-balance command for the entire device.

# LACP

LACP allows you to configure up to 4 interfaces into a port channel.

## LACP Overview

The Link Aggregation Control Protocol (LACP) for Ethernet is defined in IEEE 802.1AX and IEEE 802.3ad. This protocol controls how physical ports are bundled together to form one logical channel.

**Note** You must enable LACP before you can use LACP. By default, LACP is disabled. See the *Enabling LACP* section for information about enabling LACP.

The system automatically takes a checkpoint before disabling the feature, and you can roll back to this checkpoint. See the *Cisco Nexus® 3550-T System Management Configuration* section for information about rollbacks and checkpoints.

Individual links can be combined into LACP port channels and channel groups as well as function as individual links.

With LACP, you can bundle up to 4 interfaces in a channel group.

**Note** When you delete the port channel, the software automatically deletes the associated channel group. All member interfaces revert to their original configuration.

You cannot disable LACP while any LACP configurations are present.

## Port-Channel Modes

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

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

The following table describes the channel modes.

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

| Channel Mode | Description |
|---|---|
| **passive** | The LACP is enabled on this port channel and the ports are in a passive negotiating state. Ports responds to LACP packets that it receives but does not initiate LACP negotiation. |
| **active** | The LACP is enabled on this port channel and the ports are in an active negotiating state. Ports initiate negotiations with other ports by sending LACP packets. |
| **on** | The LACP is disabled on this port channel and the ports are in a non-negotiating state. The **on** state of the port channel represents the static mode.<br>The port will not verify or negotiate port channel memberships. If you attempt to change the channel mode to active or passive before enabling LACP, the device displays an error message. When an LACP attempts to negotiate with an interface in the **on** state, it does not receive any LACP packets and becomes an individual link with that interface, it does not join the LACP channel group. The **on** state is the default port-channel mode |

Both the passive and active modes allow LACP to negotiate between ports to determine if they can form a port channel based on criteria such as the port speed and the trunking state.The passive mode is useful when you do not know whether the remote system, or partner, supports LACP.

Two devices can form an LACP port channel when their ports are in different LACP modes if the modes are compatible as in the following example:

*Table 5: Channel Modes Compatibility*

| Device 1 > Port-1 | Device 2 > Port-2 | Result |
|---|---|---|
| Active | Active | Can form a port channel. |
| Active | Passive | Can form a port channel. |
| Passive | Passive | Cannot form a port channel because no ports can initiate negotiation. |

| Device 1 > Port-1 | Device 2 > Port-2 | Result |
|---|---|---|
| On | Active | Cannot form a port channel because LACP is enabled only on one side. |
| On | Passive | Cannot form a port channel because LACP is not enabled. |

# LACP ID Parameters

This section describes the LACP parameters.

## LACP System Priority

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

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

## LACP Port Priority

Each port that is configured to use LACP has an LACP port priority. You can accept the default value of 32768 for the LACP port priority, or you can configure a value between 1 and 65535. LACP uses the port priority with the port number to form the port identifier.

LACP uses the port priority to decide which ports should be put in standby mode when there is a limitation that prevents all compatible ports from aggregating and which ports should be put into active mode. A higher port priority value means a lower priority for LACP. You can configure the port priority so that specified ports have a lower priority for LACP and are most likely to be chosen as active links, rather than hot-standby links.

## LACP Administrative Key

LACP automatically configures an administrative key value equal to the channel-group number on each port configured to use LACP. The administrative key defines the ability of a port to aggregate with other ports. A port's ability to aggregate with other ports is determined by these factors:

- Port physical characteristics, such as the data rate and the duplex capability
- Configuration restrictions that you establish

# LACP-Enabled and Static Port Channels Differences

The following table summarizes the major differences between port channels with LACP enabled and static port channels.

*Table 6: Port Channels with LACP Enabled and Static Port Channels*

| Configurations | Port Channels with LACP Enabled | Static Port Channels |
|---|---|---|
| Protocol applied | Enable globally | Not applicable |
| Channel mode of links | Can be either:<br>• Active<br>• Passive | Can only be On |
| Maximum number of links in channel | 4 | 4 |

# LACP Compatibility Enhancements

When a Cisco Nexus 3550-T device is connected to a non-Nexus peer, its graceful failover defaults may delay the time that is taken to bring down a disabled port or cause traffic from the peer to be lost. To address these conditions, the **lacp graceful-convergence** command was added.

By default, LACP sets a port to suspended state if it does not receive an LACP PDU from the peer. **lacp suspend-individual** is a default configuration on Cisco Nexus® 3550-T switches. This command puts the port in suspended state if it does not receive any LACP PDUs. In some cases, although this feature helps in preventing loops created due to misconfigurations, it can cause servers fail to boot up because they require LACP to logically bring up the port. You can put a port into an individual state by using the **no lacp suspend-individual**. Port in individual sate takes attributes of the individual port based on the port configuration.

LACP port-channels exchange LACP PDUs for quick bundling of links when connecting a server and a switch. However, the links go into suspended state when the PDUs are not received.

The **delayed LACP** feature enables one port-channel member, the delayed-LACP port, to come up first as a member of a regular port-channel before LACP PDUs are received. After it is connected in LACP mode, other members, the auxiliary LACP ports, are brought up. This avoids having the links becoming suspended when PDUs are not received.

Which port in the port-channel comes up first depends on the port-priority value of the ports. A member link in a port channel with lowest priority value, will come come up first as a LACP delayed port. Regardless of the operational status of the links, the configured priority of a LACP port is used to select the delayed-lacp port

This feature supports Layer 2 port channels and trunk mode spanning tree and has the following limitations:

- Using **no lacp suspend-individual** and **lacp mode delay** on a same port channel is not recommended because it can put non-lacp delayed ports in individual state. As a best practice, you must avoid combining these two configurations.
- Not supported on Layer 3 port channels.

# LACP Port-Channel Minimum Links and MaxBundle

A port channel aggregates similar ports to provide increased bandwidth in a single manageable interface.

The introduction of the minimum links and maxbundle feature further refines LACP port-channel operation and provides increased bandwidth in one manageable interface.

The LACP port-channel minimum links feature does the following:

- Configures the minimum number of ports that must be linked up and bundled in the LACP port channel.
- Prevents the low-bandwidth LACP port channel from becoming active.
- Causes the LACP port channel to become inactive if there are few active members ports to supply the required minimum bandwidth.

The LACP MaxBundle defines the maximum number of bundled ports allowed in a LACP port channel.

The LACP MaxBundle feature does the following:

- Defines an upper limit on the number of bundled ports in an LACP port channel.
- Allows hot-standby ports with fewer bundled ports. (For example, in an LACP port channel with five ports, you can designate two of those ports as hot-standby ports.)



**Note** The minimum links and maxbundle feature works only with LACP port channels. However, the device allows you to configure this feature in non-LACP port channels, but the feature is not operational.

## LACP Fast Timers

You can change the LACP timer rate to modify the duration of the LACP timeout. Use the lacp rate command to set the rate at which LACP control packets are sent to an LACP-supported interface. You can change the timeout rate from the default rate (30 seconds) to the fast rate (1 second). This command is supported only on LACP-enabled interfaces. To configure the LACP fast time rate, see the *Configuring the LACP Fast Timer Rate* section.

# Prerequisites for Port Channeling

Port channeling has the following prerequisites:

- You must be logged onto the device.
- All ports for a single port channel must be either Layer 2 or Layer 3 ports.
- All ports for a single port channel must meet the compatibility requirements. See the Compatibility Requirements, on page 48 section for more information about the compatibility requirements.

# Guidelines and Limitations

Port channeling has the following configuration guidelines and limitations:

- **show** commands with the **internal** keyword are not supported.

- The LACP port-channel minimum links and maxbundle feature is not supported for host interface port channels.
- Enable LACP before you can use that feature.
- You can configure multiple port channels on a device.
- Do not put shared and dedicated ports into the same port channel. (See the *Configuring Basic Interface Parameters* chapter for information about shared and dedicated ports.)
- For Layer 2 port channels, ports with different STP port path costs can form a port channel if they are compatibly configured with each other. See the Compatibility Requirements, on page 48 section for more information about the compatibility requirements.
- In STP, the port-channel cost is based on the aggregated bandwidth of the port members.
- After you configure a port channel, the configuration that you apply to the port channel interface affects the port channel member ports. The configuration that you apply to the member ports affects only the member port where you apply the configuration.
- LACP does not support half-duplex mode. Half-duplex ports in LACP port channels are put in the suspended state.
- A maximum of 48 port channels can be supported by Cisco Nexus 3550-T switches system-wide.
- On a Cisco Nexus 3550-T switch, 8 ports are part of the same port group. All the ports in the same port group must have same speed. Maximum of eight member ports per port-channel. All the member ports must be in the same port group.

# Default Settings

The following table lists the default settings for port-channel parameters.

***Table 7: Default Port-Channel Parameters***

| Parameters | Default |
|---|---|
| Port channel | Admin up |
| Load balancing method for Layer 3 interfaces | Source and destination IP address |
| Load balancing method for Layer 2 interfaces | Source and destination MAC address |
| Load balancing per module | Disabled |
| LACP | Disabled |
| Channel mode | on |
| LACP system priority | 32768 |
| LACP port priority | 32768 |
| Minimum links for LACP | 1 |

| Parameters | Default |
|---|---|
| Maxbundle | 8 |
| Maxbundle | 4 |

# Configuring Port Channels

**Note** See the *Configuring Layer 3 Interfaces* chapter for information about configuring IPv4 addresses on the port-channel interface.

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

## Creating a Port Channel

You can create a port channel before you create a channel group. The software automatically creates the associated channel group.

**Note** When the port channel is created before the channel group, the port channel should be configured with all of the interface attributes that the member interfaces are configured with. Use the **switchport mode trunk** {*allowed vlan vlan-id* | *native vlan-id*} command to configure the members.

This is required only when the channel group members are Layer 2 ports (switchport) and trunks (switchport mode trunk).

**Note** Use the **no interface port-channel** command to remove the port channel and delete the associated channel group.

| Command | Purpose |
|---|---|
| **no interface port-channel** *channel-number*<br>**Example:**<br>`switch(config)# no interface port-channel 1` | Removes the port channel and deletes the associated channel group. |

### Before you begin

Enable LACP if you want LACP-based port channels.

### Procedure

| | Command or Action | Purpose |
|---|---|---|
| **Step 1** | **configure terminal**<br>**Example:**<br>`switch# configure terminal`<br>`switch(config)#` | Enters global configuration mode. |
| **Step 2** | **interface port-channel** *channel-number*<br>**Example:**<br>`switch(config)# interface port-channel 1`<br>`switch(config-if)` | Specifies the port-channel interface to configure, and enters the interface configuration mode. The range is from 1 to 4096. The Cisco NX-OS software automatically creates the channel group if it does not already exist. |
| **Step 3** | **show port-channel summary**<br>**Example:**<br>`switch(config-router)# show port-channel summary` | (Optional) Displays information about the port channel. |
| **Step 4** | **no shutdown**<br>**Example:**<br>`switch# configure terminal`<br>`switch(config)# int e1/1`<br>`switch(config-if)# no shutdown` | (Optional) Clears the errors on the interfaces and VLANs where policies correspond with hardware policies. This command allows policy programming to continue and the port to come up. If policies do not correspond, the errors are placed in an error-disabled policy state. |
| **Step 5** | **copy running-config startup-config**<br>**Example:**<br>`switch(config)# copy running-config startup-config` | (Optional) Copies the running configuration to the startup configuration. |

### Example

This example shows how to create a port channel:

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

See the Compatibility Requirements, on page 48 section for details on how the interface configuration changes when you delete the port channel.

# Adding a Layer 2 Port to a Port Channel

You can add a Layer 2 port to a new channel group or to a channel group that already contains Layer 2 ports. The software creates the port channel associated with this channel group if the port channel does not already exist.

**Note** Use the **no channel-group** command to remove the port from the channel group.

| Command | Purpose |
|---|---|
| **no channel-group**<br>**Example:**<br>`switch(config)# no channel-group` | Removes the port from the channel group. |

### Before you begin

Enable LACP if you want LACP-based port channels.

All Layer 2 member ports must run in full-duplex mode and at the same speed

## Procedure

| | Command or Action | Purpose |
|---|---|---|
| **Step 1** | **configure terminal**<br>**Example:**<br>`switch# configure terminal`<br>`switch(config)#` | Enters global configuration mode. |
| **Step 2** | **interface** *type slot/port*<br>**Example:**<br>`switch(config)# interface ethernet 1/4`<br>`switch(config-if)#` | Specifies the interface that you want to add to a channel group, and enters the interface configuration mode. |
| **Step 3** | **switchport**<br>**Example:**<br>`switch(config)# switchport` | Configures the interface as a Layer 2 access port. |
| **Step 4** | **switchport mode trunk**<br>**Example:**<br>`switch(config)# switchport mode trunk` | (Optional) Configures the interface as a Layer 2 trunk port. |
| **Step 5** | **switchport trunk** {**allowed vlan** *vlan-id* \| **native** *vlan-id*}<br>**Example:**<br>`switch(config)# switchport trunk native 3`<br>`switch(config-if)#` | (Optional) Configures necessary parameters for a Layer 2 trunk port. |
| **Step 6** | **channel-group** *channel-number*[**force**] [**mode** {**on** \| **active** \| **passive**}]<br>**Example:** | Configures the port in a channel group and sets the mode. The channel-number range is from 1 to 4096. This command creates the port channel associated with this channel group if the port |

| | Command or Action | Purpose |
|---|---|---|
| | • `switch(config-if)# channel-group 5`<br>• `switch(config-if)# channel-group 5 force` | channel does not already exist. All static port-channel interfaces are set to mode **on**. You must set all LACP-enabled port-channel interfaces to **active** or **passive**. The default mode is **on**.<br>(Optional) Forces an interface with some incompatible configurations to join the channel. The forced interface must have the same speed, duplex, and flow control settings as the channel group. |
| **Step 7** | **show interface** *type slot/port*<br>**Example:**<br>`switch# show interface port channel 5` | (Optional) Displays interface information. |
| **Step 8** | **no shutdown**<br>**Example:**<br>`switch# configure terminal`<br>`switch(config)# int e1/4`<br>`switch(config-if)# no shutdown` | (Optional) Clears the errors on the interfaces and VLANs where policies correspond with hardware policies. This command allows policy programming to continue and the port to come up. If policies do not correspond, the errors are placed in an error-disabled policy state. |
| **Step 9** | **copy running-config startup-config**<br>**Example:**<br>`switch(config)# copy running-config startup-config` | (Optional) Copies the running configuration to the startup configuration. |

### Example

This example shows how to add a Layer 2 Ethernet interface 1/4 to channel group 5:

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

# Adding a Layer 3 Port to a Port Channel

You can add a Layer 3 port to a new channel group or to a channel group that is already configured with Layer 3 ports. The software creates the port channel associated with this channel group if the port channel does not already exist.

If the Layer 3 port that you are adding has a configured IP address, the system removes that IP address before adding the port to the port channel. After you create a Layer 3 port channel, you can assign an IP address to the port-channel interface.

**Note** Use the **no channel-group** command to remove the port from the channel group. The port reverts to its original configuration. You must reconfigure the IP addresses for this port.

| Command | Purpose |
|---|---|
| **no channel-group**<br>**Example:**<br>`switch(config)# no channel-group` | Removes the port from the channel group. |

### Before you begin

Enable LACP if you want LACP-based port channels.

Remove any IP addresses configured on the Layer 3 interface.

## Procedure

| | Command or Action | Purpose |
|---|---|---|
| **Step 1** | **configure terminal**<br>**Example:**<br>`switch# configure terminal`<br>`switch(config)#` | Enters global configuration mode. |
| **Step 2** | **interface** *type slot/port*<br>**Example:**<br>`switch(config)# interface ethernet 1/4`<br>`switch(config-if)#` | Specifies the interface that you want to add to a channel group, and enters the interface configuration mode. |
| **Step 3** | **no switchport**<br>**Example:**<br>`switch(config-if)# no switchport` | Configures the interface as a Layer 3 port. |
| **Step 4** | **channel-group** *channel-number* [**force**] [**mode** {**on** \| **active** \| **passive**}]<br>**Example:**<br>• `switch(config-if)# channel-group 5`<br>• `switch(config-if)# channel-group 5 force` | Configures the port in a channel group and sets the mode. The channel-number range is from 1 to 4096. The Cisco NX-OS software creates the port channel associated with this channel group if the port channel does not already exist.<br>(Optional) Forces an interface with some incompatible configurations to join the channel. The forced interface must have the same speed, duplex, and flow control settings as the channel group. |
| **Step 5** | **show interface** *type slot/port*<br>**Example:** | (Optional) Displays interface information. |

| | Command or Action | Purpose |
|---|---|---|
| | `switch# show interface ethernet 1/4` | |
| **Step 6** | **no shutdown**<br>**Example:**<br>`switch# configure terminal`<br>`switch(config)# int e1/1`<br>`switch(config-if)# no shutdown` | (Optional) Clears the errors on the interfaces and VLANs where policies correspond with hardware policies. This command allows policy programming to continue and the port to come up. If policies do not correspond, the errors are placed in an error-disabled policy state. |
| **Step 7** | **copy running-config startup-config**<br>**Example:**<br>`switch(config)# copy running-config startup-config` | (Optional) Copies the running configuration to the startup configuration. |

### Example

This example shows how to add a Layer 3 Ethernet interface 1/5 to channel group 6 in on mode:

```
switch# configure terminal
switch (config)# interface ethernet 1/5
switch(config-if)# switchport
switch(config-if)# channel-group 6
```

This example shows how to create a Layer 3 port-channel interface and assign the IP address:

```
switch# configure terminal
switch (config)# interface port-channel 4
switch(config-if)# ip address 192.0.2.1/8
```

# Configuring the Bandwidth and Delay for Informational Purposes

The bandwidth of the port channel is determined by the number of total active links in the channel.

You configure the bandwidth and delay on port-channel interfaces for informational purposes.

### Procedure

| | Command or Action | Purpose |
|---|---|---|
| **Step 1** | **configure terminal**<br>**Example:**<br>`switch# configure terminal`<br>`switch(config)#` | Enters global configuration mode. |

| | Command or Action | Purpose |
|---|---|---|
| **Step 2** | **interface port-channel** *channel-number*<br>**Example:**<br>`switch(config)# interface port-channel 2`<br>`switch(config-if)#` | Specifies the port-channel interface that you want to configure, and enters the interface mode. |
| **Step 3** | **bandwidth** *value*<br>**Example:**<br>`switch(config-if)# bandwidth 60000000`<br>`switch(config-if)#` | Specifies the bandwidth, which is used for informational purposes. The range is from 1 to 3,200,000,000 kbs. The default value depends on the total active interfaces in the channel group. |
| **Step 4** | **delay** *value*<br>**Example:**<br>`switch(config-if)# delay 10000`<br>`switch(config-if)#` | Specifies the throughput delay, which is used for informational purposes. The range is from 1 to 16,777,215 tens of microseconds. The default value is 10 microseconds. |
| **Step 5** | **exit**<br>**Example:**<br>`switch(config-if)# exit`<br>`switch(config)#` | Exits the interface mode and returns to the configuration mode. |
| **Step 6** | **show interface port-channel** *channel-number*<br>**Example:**<br>`switch# show interface port-channel 2` | (Optional) Displays interface information for the specified port channel. |
| **Step 7** | **copy running-config startup-config**<br>**Example:**<br>`switch(config)# copy running-config startup-config` | (Optional) Copies the running configuration to the startup configuration. |

### Example

This example shows how to configure the informational parameters of the bandwidth and delay for port channel 5:

```
switch# configure terminal
switch (config)# interface port-channel 5
switch(config-if)# bandwidth 60000000
switch(config-if)# delay 10000
switch(config-if)#
```

# Shutting Down and Restarting the Port-Channel Interface

You can shut down and restart the port-channel interface. When you shut down a port-channel interface, no traffic passes and the interface is administratively down.

### Procedure

| | Command or Action | Purpose |
|---|---|---|
| **Step 1** | **configure terminal**<br>**Example:**<br>`switch# configure terminal`<br>`switch(config)#` | Enters global configuration mode. |
| **Step 2** | **interface port-channel** *channel-number*<br>**Example:**<br>`switch(config)# interface port-channel 2`<br>`switch(config-if)#` | Specifies the port-channel interface that you want to configure, and enters the interface mode. |
| **Step 3** | **shutdown**<br>**Example:**<br>`switch(config-if)# shutdown`<br>`switch(config-if)#` | Shuts down the interface. No traffic passes and the interface displays as administratively down. The default is no shutdown.<br>**Note**<br>Use the **no shutdown** command to open the interface.<br>The interface displays as administratively up. If there are no operational problems, traffic passes. The default is no shutdown. |
| **Step 4** | **exit**<br>**Example:**<br>`switch(config-if)# exit`<br>`switch(config)#` | Exits the interface mode and returns to the configuration mode. |
| **Step 5** | **show interface port-channel** *channel-number*<br>**Example:**<br>`switch(config-router)# show interface port-channel 2` | (Optional) Displays interface information for the specified port channel. |
| **Step 6** | **no shutdown**<br>**Example:**<br>`switch# configure terminal`<br>`switch(config)# int e1/4`<br>`switch(config-if)# no shutdown` | (Optional) Clears the errors on the interfaces and VLANs where policies correspond with hardware policies. This command allows policy programming to continue and the port to come up. If policies do not correspond, the errors are placed in an error-disabled policy state. |
| **Step 7** | **copy running-config startup-config**<br>**Example:**<br>`switch(config)# copy running-config startup-config` | (Optional) Copies the running configuration to the startup configuration. |

### Example

This example shows how to bring up the interface for port channel 2:

```
switch# configure terminal
switch (config)# interface port-channel 2
switch(config-if)# no shutdown
```

# Configuring a Port-Channel Description

You can configure a description for a port channel.

### Procedure

| | Command or Action | Purpose |
|---|---|---|
| **Step 1** | **configure terminal**<br>**Example:**<br>`switch# configure terminal`<br>`switch(config)#` | Enters global configuration mode. |
| **Step 2** | **interface port-channel** *channel-number*<br>**Example:**<br>`switch(config)# interface port-channel 2`<br>`switch(config-if)#` | Specifies the port-channel interface that you want to configure, and enters the interface mode. |
| **Step 3** | **description**<br>**Example:**<br>`switch(config-if)# description engineering`<br>`switch(config-if)#` | Allows you to add a description to the port-channel interface. You can use up to 80 characters in the description. By default, the description does not display; you must configure this parameter before the description displays in the output. |
| **Step 4** | **exit**<br>**Example:**<br>`switch(config-if)# exit`<br>`switch(config)#` | Exits the interface mode and returns to the configuration mode. |
| **Step 5** | **show interface port-channel** *channel-number*<br>**Example:**<br>`switch# show interface port-channel 2` | (Optional) Displays interface information for the specified port channel. |
| **Step 6** | **copy running-config startup-config**<br>**Example:**<br>`switch(config)# copy running-config startup-config` | (Optional) Copies the running configuration to the startup configuration. |

### Example

This example shows how to add a description to port channel 2:

```
switch# configure terminal
switch (config)# interface port-channel 2
switch(config-if)# description engineering
```

# Enabling LACP

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

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

To configure LACP, you must do the following:

- Enable LACP globally by using the **feature lacp** command.
- You can use different modes for different interfaces within the same LACP-enabled port channel. You can change the mode between **active** and **passive** for an interface only if it is the only interface that is designated to the specified channel group.

### Procedure

| | Command or Action | Purpose |
|---|---|---|
| **Step 1** | **configure terminal**<br>**Example:**<br>`switch# configure terminal`<br>`switch(config)#` | Enters global configuration mode. |
| **Step 2** | **feature lacp**<br>**Example:**<br>`switch(config)# feature lacp` | Enables LACP on the device. |
| **Step 3** | **copy running-config startup-config**<br>**Example:**<br>`switch(config)# copy running-config startup-config` | (Optional) Copies the running configuration to the startup configuration. |

### Example

This example shows how to enable LACP:

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

# Configuring LACP Port-Channel Port Modes

After you enable LACP, you can configure the channel mode for each individual link in the LACP port channel as **active** or **passive**. This channel configuration mode allows the link to operate with LACP.

When you configure port channels with no associated aggregation protocol, all interfaces on both sides of the link remain in the **on** channel mode.

### Procedure

| | Command or Action | Purpose |
|---|---|---|
| **Step 1** | **configure terminal**<br>**Example:**<br>`switch# configure terminal`<br>`switch(config)#` | Enters global configuration mode. |
| **Step 2** | **interface** *type slot/port*<br>**Example:**<br>`switch(config)# interface ethernet 1/4`<br>`switch(config-if)#` | Specifies the interface that you want to add to a channel group, and enters the interface configuration mode. |
| **Step 3** | **channel-group** *number* **mode** {**active** \| **on** \| **passive**}<br>**Example:**<br>`switch(config-if)# channel-group 5 mode active` | 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.<br>When you run port channels with no associated aggregation protocol, the port-channel mode is always on.<br>The default port-channel mode is **on**. |
| **Step 4** | **show port-channel summary**<br>**Example:**<br>`switch(config-if)# show port-channel summary` | (Optional) Displays summary information about the port channels. |
| **Step 5** | **copy running-config startup-config**<br>**Example:**<br>`switch(config)# copy running-config startup-config` | (Optional) Copies the running configuration to the startup configuration. |

### Example

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

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

# Configuring LACP Port-Channel Minimum Links

You can configure the LACP minimum links feature. Although minimum links and maxbundles work only in LACP, you can enter the CLI commands for these features for non-LACP port channels, but these commands are nonoperational.



**Note** Use the **no lacp min-links** command to restore the default port-channel minimum links configuration.

| Command | Purpose |
|---|---|
| **no lacp min-links**<br>**Example:**<br>`switch(config)# no lacp min-links` | Restores the default port-channel minimum links configuration. |

### Before you begin

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

### Procedure

| | Command or Action | Purpose |
|---|---|---|
| **Step 1** | **configure terminal**<br>**Example:**<br>`switch# configure terminal`<br>`switch(config)#` | Enters global configuration mode. |
| **Step 2** | **interface port-channel** *number*<br>**Example:**<br>`switch(config)# interface port-channel 3`<br>`switch(config-if)#` | Specifies the interface to configure, and enters the interface configuration mode. |
| **Step 3** | **lacp min-links** *number*<br>**Example:**<br>`switch(config-if)# lacp min-links 3` | Specifies the port-channel interface to configure the number of minimum links. The range is from 1 to 4. |
| **Step 4** | **show running-config interface port-channel** *number*<br>**Example:**<br>`switch(config-if)# show running-config interface port-channel 3` | (Optional) Displays the port-channel minimum links configuration. |

### Example

This example shows how to configure the minimum number of port-channel member interfaces to be up/active for the port-channel to be up/active:

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

# Configuring the LACP Port-Channel MaxBundle

You can configure the LACP maxbundle feature. Although minimum links and maxbundles work only in LACP, you can enter the CLI commands for these features for non-LACP port channels, but these commands are nonoperational.



**Note** Use the **no lacp max-bundle** command to restore the default port-channel max-bundle configuration.

| Command | Purpose |
|---|---|
| **no lacp max-bundle**<br>**Example:**<br>`switch(config)# no lacp max-bundle` | Restores the default port-channel max-bundle configuration. |

### Before you begin

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

### Procedure

| | Command or Action | Purpose |
|---|---|---|
| **Step 1** | **configure terminal**<br>**Example:**<br>`switch# configure terminal`<br>`switch(config)#` | Enters global configuration mode. |
| **Step 2** | **interface port-channel** *number*<br>**Example:**<br>`switch(config)# interface port-channel 3`<br>`switch(config-if)#` | Specifies the interface to configure, and enters the interface configuration mode. |
| **Step 3** | **lacp max-bundle** *number*<br>**Example:** | Specifies the port-channel interface to configure max-bundle. |

| | Command or Action | Purpose |
|---|---|---|
| | `switch(config-if)# lacp max-bundle` | The default value for the port-channel max-bundle is 8. The allowed range is from 1 to 8.<br>The default value for the port-channel max-bundle is 4. The allowed range is from 1 to 4.<br>**Note**<br>Even if the default value is 8, the number of active members in a port channel is the minimum of the pc_max_links_config and pc_max_active_members that is allowed in the port channel.<br>Even if the default value is 4, the number of active members in a port channel is the minimum of the pc_max_links_config and pc_max_active_members that is allowed in the port channel. |
| **Step 4** | **show running-config interface port-channel** *number*<br>**Example:**<br>`switch(config-if)# show running-config interface port-channel 3` | (Optional) Displays the port-channel max-bundle configuration. |

### Example

This example shows how to configure the port channel interface max-bundle:

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

# Configuring the LACP Fast Timer Rate

You can change the LACP timer rate to modify the duration of the LACP timeout. Use the **lacp rate** command to set the rate at which LACP control packets are sent to an LACP-supported interface. You can change the timeout rate from the default rate (30 seconds) to the fast rate (1 second). This command is supported only on LACP-enabled interfaces.



**Note** We do not recommend changing the LACP timer rate. HA and SSO are not supported when the LACP fast rate timer is configured.

### Before you begin

Ensure that you have enabled the LACP feature.

### Procedure

| | Command or Action | Purpose |
|---|---|---|
| **Step 1** | **configure terminal**<br>**Example:**<br>`switch# configure terminal`<br>`switch(config)#` | Enters global configuration mode. |
| **Step 2** | **interface** *type slot/port*<br>**Example:**<br>`switch(config)# interface ethernet 1/4`<br>`switch(config-if)#` | Specifies the interface to configure and enters the interface configuration mode. |
| **Step 3** | **lacp rate fast**<br>**Example:**<br>`switch(config-if)# lacp rate fast` | Configures the fast rate (one second) at which LACP control packets are sent to an LACP-supported interface.<br>To reset the timeout rate to its default, use the **no** form of the command. |

### Example

This example shows how to configure the LACP fast rate on Ethernet interface 1/4:

```
switch# configure terminal
switch (config)# interface ethernet 1/4
switch(config-if)# lacp rate fast
```

This example shows how to restore the LACP default rate (30 seconds) on Ethernet interface 1/4.

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

# Configuring the LACP System Priority

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

### Before you begin

Enable LACP.

### Procedure

| | Command or Action | Purpose |
|---|---|---|
| **Step 1** | **configure terminal**<br>**Example:**<br>`switch# configure terminal`<br>`switch(config)#` | Enters global configuration mode. |
| **Step 2** | **lacp system-priority** *priority*<br>**Example:**<br>`switch(config)# lacp system-priority 40000` | Configures the system priority for use with LACP. Valid values are from 1 through 65535, and higher numbers have a lower priority. The default value is 32768. |
| **Step 3** | **show lacp system-identifier**<br>**Example:**<br>`switch(config-if)# show lacp system-identifier` | (Optional) Displays the LACP system identifier. |
| **Step 4** | **copy running-config startup-config**<br>**Example:**<br>`switch(config)# copy running-config startup-config` | (Optional) Copies the running configuration to the startup configuration. |

### Example

This example shows how to set the LACP system priority to 2500:

```
switch# configure terminal
switch(config)# lacp system-priority 2500
```

# Configuring the LACP Port Priority

When you enable LACP, you can configure each link in the LACP port channel for the port priority.

### Before you begin

Enable LACP.

### Procedure

| | Command or Action | Purpose |
|---|---|---|
| **Step 1** | **configure terminal**<br>**Example:**<br>`switch# configure terminal`<br>`switch(config)#` | Enters global configuration mode. |

| | Command or Action | Purpose |
|---|---|---|
| **Step 2** | **interface** *type slot/port*<br>**Example:**<br>`switch(config)# interface ethernet 1/4`<br>`switch(config-if)#` | Specifies the interface that you want to add to a channel group, and enters the interface configuration mode. |
| **Step 3** | **lacp port-priority** *priority*<br>**Example:**<br>`switch(config-if)# lacp port-priority 40000` | Configures the port priority for use with LACP. Valid values are from 1 through 65535, and higher numbers have a lower priority. The default value is 32768. |
| **Step 4** | **copy running-config startup-config**<br>**Example:**<br>`switch(config-if)# copy running-config startup-config` | (Optional) Copies the running configuration to the startup configuration. |

### Example

This example shows how to set the LACP port priority for Ethernet interface 1/4 to 40000:

```
switch# configure terminal
switch (config)# interface ethernet 1/4
switch(config-if)# lacp port-priority 40000
```

# Configuring LACP System MAC and Role

You can configure the MAC address used by the LACP for protocol exchanges and the optional role. By default, the role is primary.

This procedure is supported on the Cisco Nexus 3550-T switches.

### Before you begin

LACP must be enabled.

### Procedure

| | Command or Action | Purpose |
|---|---|---|
| **Step 1** | **configure terminal**<br>**Example:**<br>`switch# configure terminal` | Enter global configuration mode. |
| **Step 2** | **lacp system-mac** *mac-address* **role** *role-value*<br>**Example:**<br>`switch(config)# lacp system-mac 000a.000b.000c role primary` | Specifies the MAC address to use in the LACP protocol exchanges. The role is optional. Primary is the default. |

| | Command or Action | Purpose |
|---|---|---|
| | `switch(config)# lacp system-mac 000a.000b.000c role secondary` | |
| **Step 3** | (Optional) **show lacp system-identifier**<br>**Example:**<br>`switch(config)# show lacp system-identifier` | Displays the configured MAC address. |
| **Step 4** | **copy running-config startup-config**<br>**Example:**<br>`switch(config)# copy running-config startup-config` | Copies the running configuration to the startup configuration. |

### Example

The following example shows how to configure the role of a switch as primary.

```
Switch1# sh lacp system-identifier
32768,0-b-0-b-0-b
Switch1# sh run  | grep lacp
feature lacp
lacp system-mac 000b.000b.000b role primary
```

The following example shows how to configure the role of a switch as secondary.

```
Switch2# sh lacp system-identifier
32768,0-b-0-b-0-b
Switch2# sh run | grep lacp
feature lacp
lacp system-mac 000b.000b.000b role secondary
```

# Disabling LACP Graceful Convergence

By default, LACP graceful convergence is enabled. In situations where you need to support LACP interoperability with devices where the graceful failover defaults may delay the time taken for a disabled port to be brought down or cause traffic from the peer to be lost, you can disable convergence. If the downstream access switch is not a Cisco Nexus device, disable the LACP graceful convergence option.



**Note** The port channel has to be in the administratively down state before the command can be run.

### Before you begin

Enable LACP.

### Procedure

| | Command or Action | Purpose |
|---|---|---|
| **Step 1** | **configure terminal**<br>**Example:**<br>`switch# configure terminal`<br>`switch(config)#` | Enters global configuration mode. |
| **Step 2** | **interface port-channel** *number*<br>**Example:**<br>`switch(config)# interface port-channel 1`<br>`switch(config-if)#` | Specifies the port channel interface to configure and enters the interface configuration mode. |
| **Step 3** | **shutdown**<br>**Example:**<br>`switch(config-if) shutdown` | Administratively shuts down the port channel. |
| **Step 4** | **no lacp graceful-convergence**<br>**Example:**<br>`switch(config-if)# no lacp graceful-convergence` | Disables LACP graceful convergence on the port channel. |
| **Step 5** | **no shutdown**<br>**Example:**<br>`switch(config-if) no shutdown` | Brings the port channel administratively up. |
| **Step 6** | **copy running-config startup-config**<br>**Example:**<br>`switch(config)# copy running-config startup-config` | (Optional) Copies the running configuration to the startup configuration. |

#### Example

This example shows how to disable LACP graceful convergence on a port channel:

```
switch# configure terminal
switch (config)# interface port-channel 1
switch(config-if)# shutdown
switch(config-if)# no lacp graceful-convergence
switch(config-if)# no shutdown
```

## Reenabling LACP Graceful Convergence

If the default LACP graceful convergence is once again required, you can reenable convergence.

### Procedure

| | Command or Action | Purpose |
|---|---|---|
| **Step 1** | **configure terminal**<br>**Example:**<br>`switch# configure terminal`<br>`switch(config)#` | Enters global configuration mode. |
| **Step 2** | **interface port-channel** *number*<br>**Example:**<br>`switch(config)# interface port-channel 1`<br>`switch(config-if)#` | Specifies the port channel interface to configure and enters the interface configuration mode. |
| **Step 3** | **shutdown**<br>**Example:**<br>`switch(config-if) shutdown` | Administratively shuts down the port channel. |
| **Step 4** | **lacp graceful-convergence**<br>**Example:**<br>`switch(config-if)# lacp graceful-convergence` | Enables LACP graceful convergence on the port channel. |
| **Step 5** | **no shutdown**<br>**Example:**<br>`switch(config-if) no shutdown` | Brings the port channel administratively up. |
| **Step 6** | **copy running-config startup-config**<br>**Example:**<br>`switch(config)# copy running-config startup-config` | (Optional) Copies the running configuration to the startup configuration. |

### Example

This example shows how to enable LACP graceful convergence on a port channel:

```
switch# configure terminal
switch (config)# interface port-channel 1
switch(config-if)# shutdown
switch(config-if)# lacp graceful-convergence
switch(config-if)# no shutdown
```

# Disabling LACP Suspend Individual

LACP sets a port to the suspended state if it does not receive an LACP PDU from the peer. This process can cause some servers to fail to boot up as they require LACP to logically bring up the port.

**Note** You should only enter the **lacp suspend-individual** command on edge ports. The port channel has to be in the administratively down state before you can use this command.

### Before you begin

Enable LACP.

## Procedure

| | Command or Action | Purpose |
|---|---|---|
| **Step 1** | **configure terminal**<br>**Example:**<br>`switch# configure terminal`<br>`switch(config)#` | Enters global configuration mode. |
| **Step 2** | **interface port-channel** *number*<br>**Example:**<br>`switch(config)# interface port-channel 1`<br>`switch(config-if)#` | Specifies the port channel interface to configure and enters the interface configuration mode. |
| **Step 3** | **shutdown**<br>**Example:**<br>`switch(config-if) shutdown` | Administratively shuts down the port channel. |
| **Step 4** | **no lacp suspend-individual**<br>**Example:**<br>`switch(config-if)# no lacp suspend-individual` | Disables LACP individual port suspension behavior on the port channel. |
| **Step 5** | **no shutdown**<br>**Example:**<br>`switch(config-if) no shutdown` | Brings the port channel administratively up. |
| **Step 6** | **copy running-config startup-config**<br>**Example:**<br>`switch(config)# copy running-config startup-config` | (Optional) Copies the running configuration to the startup configuration. |

### Example

This example shows how to disable LACP individual port suspension on a port channel:

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

```
switch(config-if)# shutdown
switch(config-if)# no lacp suspend-individual
switch(config-if)# no shutdown
```

# Reenabling LACP Suspend Individual

You can reenable the default LACP individual port suspension.

### Procedure

| | Command or Action | Purpose |
|---|---|---|
| **Step 1** | **configure terminal**<br>**Example:**<br>`switch# configure terminal`<br>`switch(config)#` | Enters global configuration mode. |
| **Step 2** | **interface port-channel** *number*<br>**Example:**<br>`switch(config)# interface port-channel 1`<br>`switch(config-if)#` | Specifies the port channel interface to configure and enters the interface configuration mode. |
| **Step 3** | **shutdown**<br>**Example:**<br>`switch(config-if) shutdown` | Administratively shuts down the port channel. |
| **Step 4** | **lacp suspend-individual**<br>**Example:**<br>`switch(config-if)# lacp suspend-individual` | Enables LACP individual port suspension behavior on the port channel. |
| **Step 5** | **no shutdown**<br>**Example:**<br>`switch(config-if) no shutdown` | Brings the port channel administratively up. |
| **Step 6** | **copy running-config startup-config**<br>**Example:**<br>`switch(config)# copy running-config startup-config` | (Optional) Copies the running configuration to the startup configuration. |

### Example

This example shows how to reenable the LACP individual port suspension on a port channel:

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

```
switch(config-if)# lacp suspend-individual
switch(config-if)# no shutdown
```

# Configuring Delayed LACP

The delayed LACP feature enables one port channel member, the delayed LACP port, to come up first as a member of a regular port channel before LACP PDUs are received. You configure the delayed LACP feature using the **lacp mode delay**command on a port channel followed by configuring the LACP port priority on a one member port of the port channel.

### Procedure

| | Command or Action | Purpose |
|---|---|---|
| **Step 1** | **configure terminal** | Enters global configuration mode. |
| **Step 2** | **interface port-channel** *number* | Specifies the port channel interface to configure and enters the interface configuration mode. |
| **Step 3** | **lacp mode delay** | Enables delayed LACP.<br>**Note**<br>To disable delayed LACP, use the **no lacp mode delay** command.<br><br>Complete the configuration of the delayed LACP by configuring the LACP port priority. See the "Configuring the LACP Port Priority" section for details.<br><br>The priority of a LACP port determines the election of the delayed LACP port. The port with the lowest numerical priority is elected.<br><br>When the delayed LACP feature is configured and made effective with a port channel flap, the delayed LACP port operates as a member of a regular port channel, allowing data to be exchanged between the server and switch. After receiving the first LACP PDU, the delayed LACP port transitions from a regular port member to a LACP port member.<br>**Note**<br>The election of the delayed LACP port is not complete or effective until the port channel flaps on the switch or at a remote server. |

### Example

The following example configures delayed LACP.

```
switch# config terminal
switch(config)# interface po 1
switch(config-if)# lacp mode delay
```

```
switch# config terminal
switch(config)# interface ethernet 1/1
switch(config-if)# lacp port-priority 1
switch(config-if)# channel-group 1 mode active
```

The following example disables delayed LACP.

```
switch# config terminal
switch(config)# interface po 1
switch(config-if)# no lacp mode delay
```

# Verifying the Port-Channel Configuration

To display port-channel configuration information, perform one of the following tasks:

| Command | Purpose |
|---|---|
| **show interface port-channel** *channel-number* | Displays the status of a port-channel interface. |
| **show feature** | Displays enabled features. |
| **load- interval** {**interval** *seconds* {**1** \| **2** \| **3**}} | Sets three different sampling intervals to bit-rate and packet-rate statistics. |
| **show port-channel compatibility-parameters** | Displays the parameters that must be the same among the member ports in order to join a port channel. |
| **show port-channel database** [**interface port-channel** *channel-number*] | Displays the aggregation state for one or more port-channel interfaces. |
| **show port-channel load-balance** | Displays the type of load balancing in use for port channels. |
| **show port-channel summary** | Displays a summary for the port-channel interfaces. |
| **show port-channel traffic** | Displays the traffic statistics for port channels. |
| **show port-channel usage** | Displays the range of used and unused channel numbers. |
| **show lacp** {**counters** [**interface port-channel** *channel-number*] \| [**interface** *type/slot*] \| **neighbor** [**interface port-channel** *channel-number*] \| **port-channel** [**interface port-channel** *channel-number*] \| **system-identifier**]]} | Displays information about LACP. |
| **show running-config interface port-channel** *channel-number* | Displays information about the running configuration of the port-channel. |

# Monitoring the Port-Channel Interface Configuration

Use the following commands to display port-channel interface configuration information.

| Command | Purpose |
|---|---|
| **clear counters interface port-channel** *channel-number* | Clears the counters. |
| **clear lacp counters** [**interface port-channel** *channel-number*] | Clears the LACP counters. |
| **load- interval** {**interval** *seconds* {**1** \| **2** \| **3**}} | Sets three different sampling intervals to bit-rate and packet-rate statistics. |
| **show interface counters** [**module** *module*] | Displays input and output octets unicast packets, multicast packets, and broadcast packets. |
| **show interface counters detailed** [**all**] | Displays input packets, bytes, and multicast and output packets and bytes.<br>**Note**<br>Ignore *Output Dropped Errors* as it represents the cumulative ingress drops in the traffic that is directed to the port. The ingress drops on any port are displayed as part of *Input Discard Errors.* |
| **show interface counters errors** [**module** *module*] | Displays information about the number of error packets.<br>**Note**<br>Ignore *OutDiscards* as it represents the cumulative ingress drops in the traffic that is directed to the port. The ingress drops on any port are displayed as part of *InDiscards.* |
| **show lacp counters** | Displays statistics for LACP. |

# Example Configurations for Port Channels

This example shows how to create an LACP port channel and add two Layer 2 interfaces to that port channel:

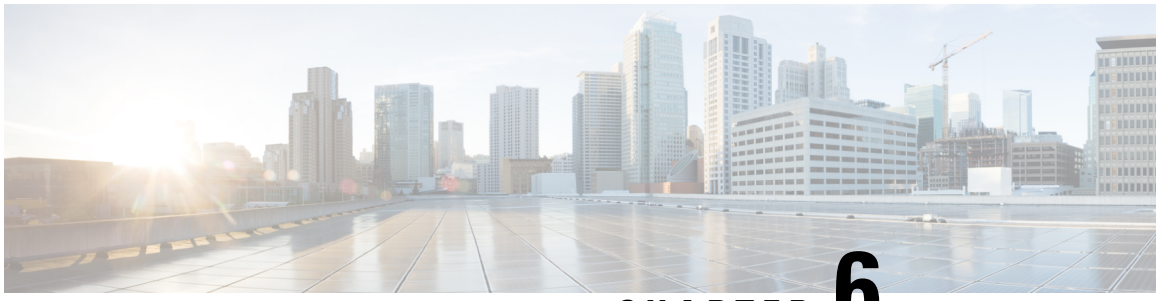
```
switch# configure terminal
switch (config)# feature lacp
switch (config)# interface port-channel 5
switch (config-if)# interface ethernet 1/4
switch(config-if)# switchport
switch(config-if)# channel-group 5 mode active
switch(config-if)# lacp port priority 40000
switch(config-if)# interface ethernet 1/7
switch(config-if)# switchport
switch(config-if)# channel-group 5 mode
```

This example shows how to add two Layer 3 interfaces to a channel group. The Cisco NX-OS software automatically creates the port channel:

```
switch# configure terminal
switch (config)# interface ethernet 1/5
switch(config-if)# no switchport
switch(config-if)# no ip address
switch(config-if)# channel-group 6 mode active
switch (config)# interface ethernet 1/6
switch(config-if)# no switchport
switch(config-if)# no ip address
switch(config-if)# channel-group 6 mode active
switch (config)# interface port-channel 6
switch(config-if)# ip address 192.0.2.1/8
```

# Related Documents

| Related Topic | Document Title |
|---|---|
| System management | *Cisco Nexus 3550-T NX-OS System Management Configuration* section |
| Licensing | *Cisco NX-OS Licensing Guide* |



CHAPTER 6

# Configuring vPCs

This chapter describes how to configure virtual port channels (vPCs) on Cisco NX-OS devices.

You can use any of the interfaces of the Nexus 3550-T device for the vPC peer link.

The port channel compatibility parameters must be the same for all the port channel members on the physical switch.

You cannot configure shared interfaces to be part of a vPC.

**Note** The port channel compatibility parameters must also be the same for all vPC member ports on both peers and therefore you must use the same type of module in each chassis.

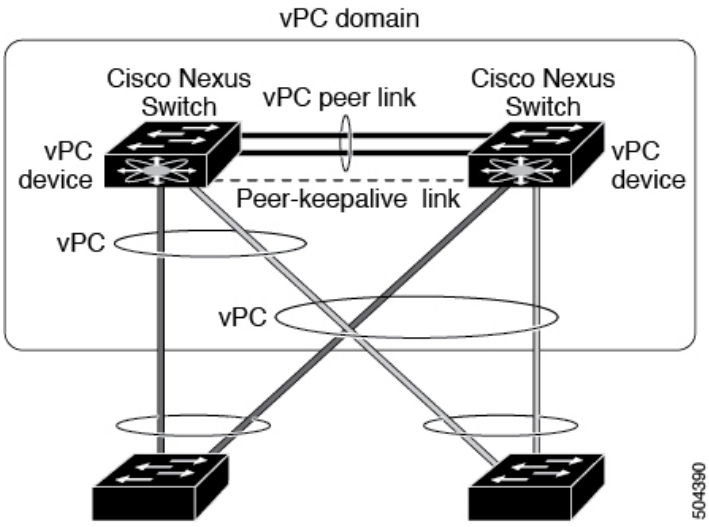

- Information About vPCs, on page 83
- Guidelines and Limitations, on page 101
- Best Practices for Layer 3 and vPC Configuration, on page 103
- Default Settings, on page 110
- Configuring vPCs, on page 110
- Verifying the vPC Configuration, on page 135
- Monitoring vPCs, on page 136
- Configuration Examples for vPCs, on page 136


# Information About vPCs

## vPC Overview

A virtual port channel (vPC) allows links that are physically connected to two Cisco Nexus 3550-T devices to appear as a single port channel by a third device (see figure). The third device can be a switch, server, or any other networking device that supports port channels. A vPC can provide Layer 2 multipathing, which allows you to create redundancy and increase the bisectional bandwidth by enabling multiple parallel paths between nodes and allowing load balancing traffic.

*Figure 5: vPC Architecture*




You can use only Layer 2 port channels in the vPC. You configure the port channels by using one of the following:

- No protocol
- Link Aggregation Control Protocol (LACP)

When you configure the port channels in a vPC—including the vPC Peer-Link channel—without using LACP, each device can have up to four links in a single port channel and all the four members must belong to the same quad. From a given quad, only one port channel is possible.

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

After you enable the vPC functionality, you create the peer-keepalive link, which sends heartbeat messages between the two vPC peer devices.

You can create a vPC Peer-Link by configuring a port channel on one Cisco Nexus 3550-T Series chassis by using two or more Ethernet ports higher speed than 1-Gigabit Ethernet. To ensure that you have the correct hardware to enable and run a vPC, enter the **show hardware feature-capability** command. If you see an X across from the vPC in your command output, your hardware cannot enable the vPC feature.

We recommend that you configure the vPC Peer-Link Layer 2 port channels as trunks. On another Cisco Nexus 3550-T Series chassis, you configure another port channel again using two or more Ethernet ports with spped higher than 1-Gigabit in the dedicated port mode. Connecting these two port channels creates a vPC Peer-Link in which the two linked Cisco Nexus devices appear as one device to a third device. The third device, or downstream device, can be a switch, server, or any other networking device that uses a regular port channel to connect to the vPC.

You can use any of the interfaces of the Nexus 3550-T device for the vPC Peer-Link.

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

In this version, you can connect each downstream device to a single vPC domain ID using a single port channel.

**Note** Devices attached to a vPC domain using port channels should be connected to both of vPC peers.

A vPC (see figure) provides the following benefits:

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

# Hitless vPC Role Change

A virtual port channel (vPC) allows links that are physically connected to two different Cisco Nexus 3550-T switches to appear as a single port channel. The vPC role change feature enables you switch vPC roles between vPC peers without impacting traffic flow. The vPC role switching is done based on the role priority value of the device under the vPC domain. A vPC peer device with lower role priority is selected as the primary vPC device during the vPC Role switch. You can use the vpc role preempt command to switch vPC role between peers.

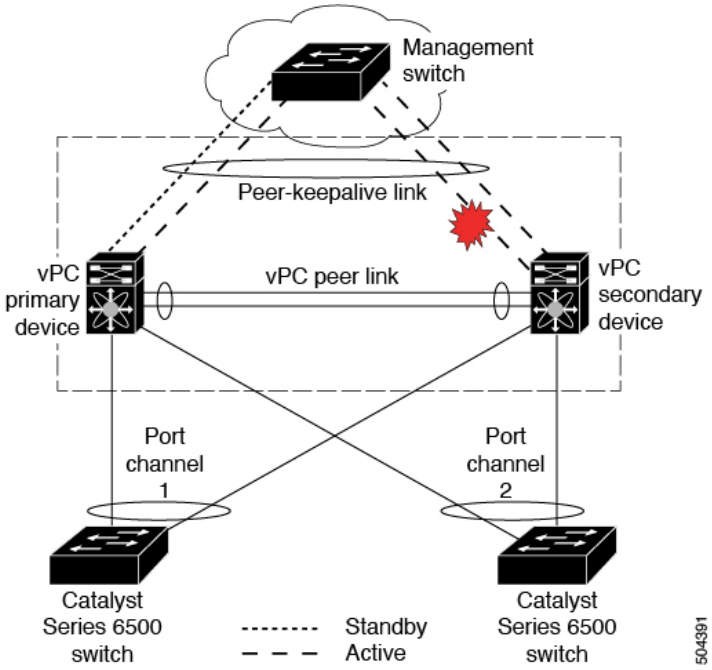
# vPC Terminology

The terminology used in vPCs is as follows:

- vPC—The combined port channel between the vPC peer devices and the downstream device.
- vPC peer device—One of a pair of devices that are connected with the special port channel known as the vPC Peer-Link.
- vPC Peer-Link—The link used to synchronize state between the vPC peer devices. This link must use a 10-Gigabit Ethernet interface.
- vPC member port—An interface that belongs to a vPC.
- vPC domain—This domain includes both vPC peer devices, the vPC peer-keepalive link, and all of the port channels in the vPC connected to the downstream devices. It is also associated to the configuration mode that you must use to assign vPC global parameters.
- vPC peer-keepalive link—The peer-keepalive link monitors the vitality of a vPC peer Cisco Nexus 3550-T Series device. The peer-keepalive link sends configurable, periodic keepalive messages between vPC peer devices.

  We recommend that you associate a peer-keepalive link to a default virtual routing and forwarding (VRF) instance that is mapped to a Layer 3 interface in each vPC peer device.

*Figure 6: Separate Switch Required to Connect Management Ports for vPC Peer-Keepalive Link*




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

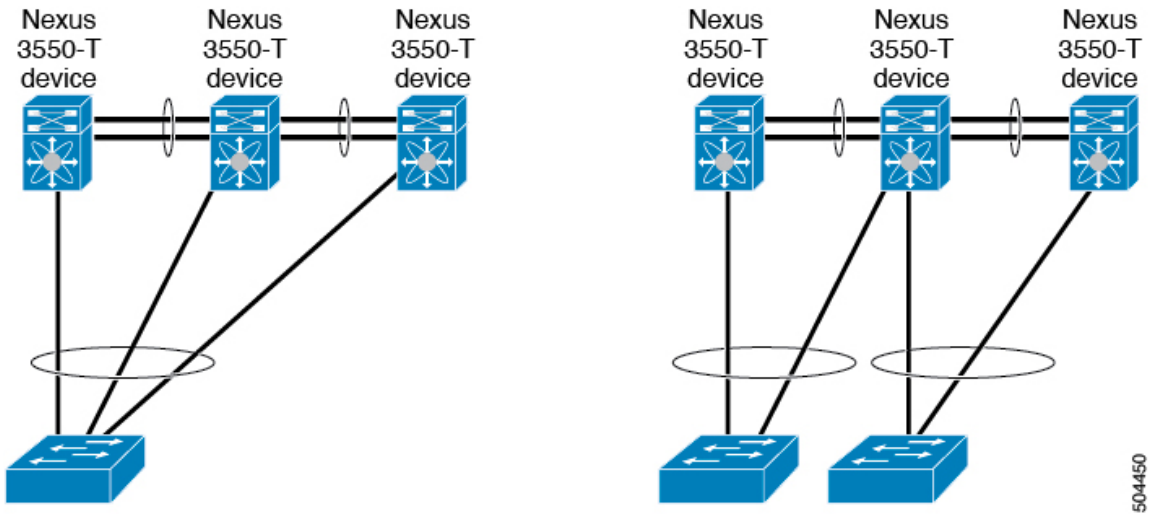
- vPC member port—Interfaces that belong to the vPCs.
- Dual-active— Both vPC peers act as primary. This situation occurs when the peer-keepalive and vPC Peer-Link go down when both the peers are still active. In this case, the secondary vPC assumes that the primary vPC is inactive and acts as the primary vPC.
- Recovery—When the peer-keepalive and the vPC Peer-Link come up, one switch becomes the secondary vPC. On the switch that becomes the secondary vPC, the vPC links go down and come back up.

## vPC Peer-Link Overview

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

See the following figure for invalid vPC peer configurations.

*Figure 7: vPC Peer Configurations That Are Not Allowed*




To make a valid configuration, you first configure a port channel on each device and then configure the vPC domain. You assign the port channel on each device as a vPC Peer-Link, using the same vPC domain ID. For redundancy, we recommend that you should configure at least two of the dedicated ports into the port channel because if one of the interfaces in the vPC Peer-Link fails, the device automatically falls back to use another interface in the vPC Peer-Link.

**Note** We recommend that you configure the Layer 2 port channels in trunk mode.

Many operational parameters and configuration parameters must be the same in each device connected by a vPC Peer-Link (see the Compatibility Parameters for vPC Interfaces section). Because each device is completely independent on the management plane, you must ensure that the devices are compatible on the critical parameters. vPC peer devices have separate control planes. After configuring the vPC Peer-Link, you should display the configuration on each vPC peer device to ensure that the configurations are compatible.

**Note** You must ensure that the two devices connected by the vPC Peer-Link have certain identical operational and configuration parameters. For more information on required configuration consistency, see the Compatibility Parameters for vPC Interfaces section.

When you configure the vPC Peer-Link, the vPC peer devices negotiate that one of the connected devices is the primary device and the other connected device is the secondary device (see the “Configuring vPCs” section). The Cisco NX-OS software uses the lowest MAC address to elect the primary device. The software takes different actions on each device—that is, the primary and secondary—only in certain failover conditions. If the primary device fails, the secondary device becomes the new primary device when the system recovers, and the previously primary device is now the secondary device.

You can also configure which of the vPC devices is the primary device. Changing the priority of the vPC peer devices can cause the interfaces in your network to go up and down. If you want to configure the role priority again to make one vPC device the primary device, configure the role priority on both the primary vPC device with a lower priority value and the secondary vPC device with the higher value. Then, shut down the port

channel that is the vPC Peer-Link on both devices by entering the **shutdown** command, and finally reenable the port channel on both devices by entering the **no shutdown** command.

The software keeps all traffic that forwards across the vPC peer devices as local traffic. A packet that ingresses the port channel uses one of the local links rather than moving across the vPC Peer-Link. Unknown unicast, and broadcast traffic (including STP BPDUs) are flooded across the vPC Peer-Link. The software keeps the multicast forwarding state synchronized on both of the vPC peer devices.

You can configure any of the standard load-balancing schemes on both the vPC Peer-Link devices and the downstream device (see the *Configuring Port Channels* chapter for information about load balancing).

Configuration information flows across the vPC Peer-Links using the Cisco Fabric Services over Ethernet (CFSoE) protocol. (See the CFSoE, on page 99 section for more information about CFSoE.)

All MAC addresses for those VLANs configured on both devices are synchronized between vPC peer devices. The software uses CFSoE for this synchronization. (See the CFSoE, on page 99 section for information about CFSoE.)

If the vPC Peer-Link fails, the software checks the status of the remote vPC peer device using the peer-keepalive link, which is a link between vPC peer devices that ensures that both devices are up. If the vPC peer device is up, the secondary vPC device disables all vPC ports on its device, to prevent loops and disappearing or flooding traffic. The data then forwards down the remaining active links of the port channel.

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

Use a separate link (vPC peer-keepalive link) to send configurable keepalive messages between the vPC peer devices. The keepalive messages on the vPC peer-keepalive link determines whether a failure is on the vPC Peer-Link only or on the vPC peer device. The keepalive messages are used only when all the links in the vPC Peer-Link fail. See the “Peer-Keepalive Link and Messages” section for information about the keepalive message.

## Features That You Must Manually Configure on the Primary and Secondary Devices

You must manually configure the following features to conform to the primary/secondary mapping of each of the vPC peer devices:

- STP root—Configure the primary vPC peer device as the STP primary root device and configure the vPC secondary device to be the STP secondary root device. See the “vPC Peer-Links and STP” section for more information about vPCs and STP.
  - We recommend that you configure the vPC Peer-Link interfaces as STP network ports so that Bridge Assurance is enabled on all vPC Peer-Links.
  - We recommend that you configure Rapid per VLAN Spanning Tree plus (PVST+) so that the primary device is the root for all VLANs and configure Multiple Spanning Tree (MST) so that the primary device is the root for all instances.
- Layer 3 VLAN network interface—Configure Layer 3 connectivity from each vPC peer device by configuring a VLAN network interface for the same VLAN from both devices.
- VRRP active—If you want to use Virtual Router Redundancy Protocol (VRRP) and VLAN interfaces on the vPC peer devices, configure the primary vPC peer device with the VRRP master as highest priority. Configure the backup device to be theVRRP standby and ensure that you have VLAN interfaces on each vPC device that are in the same administrative and operational mode.

While you configure Unidirectional Link Detection (UDLD), note the following recommendations:

- If LACP is used as port-channel aggregation protocol, UDLD is not required in a vPC domain.
- If LACP is not used as the port-channel aggregation protocol (static port-channel), use UDLD in normal mode on vPC member ports.
- If STP is used without Bridge Assurance and if LACP is not used, use UDLD in normal mode on vPC orphan ports.

## Configuring Layer 3 Backup Routes on a vPC Peer-Link

You can use VLAN network interfaces on the vPC peer devices to link to Layer 3 of the network for such applications as VRRP. Ensure that you have a VLAN network interface configured on each peer device and that the interface is connected to the same VLAN on each device. Also, each VLAN interface must be in the same administrative and operational mode. For more information about configuring VLAN network interfaces, see the "Configuring Layer 3 Interfaces" chapter.

If a failover occurs on the vPC Peer-Link, the VLAN interfaces on the vPC peer devices are also affected. If a vPC Peer-Link fails, the system brings down associated VLAN interfaces on the secondary vPC peer device.

You can ensure that specified VLAN interfaces do not go down on the vPC secondary device when the vPC Peer-Link fails.

# Peer-Keepalive Link and Messages

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



**Note** We recommend that you associate the vPC peer-keepalive link to a default VRF mapped to a Layer 3 interface in each vPC peer device. If you do not configure a management VRF, the system uses the management VRF and management ports by default. Do not use the vPC Peer-Link itself to send and receive vPC peer-keepalive messages.

If one of the vPC peer devices fails, the vPC peer device on the other side of the vPC Peer-Link senses the failure by not receiving any peer-keepalive messages. The default interval time for the vPC peer-keepalive message is 1 second, and you can configure the interval between 400 milliseconds and 10 seconds.

You can configure a hold-timeout value with a range of 3 to 10 seconds; the default hold-timeout value is 3 seconds. This timer starts when the vPC Peer-Link goes down. During this hold-timeout period, the secondary vPC peer device ignores vPC peer-keepalive messages, which ensures that network convergence occurs before a vPC action takes place. The purpose of the hold-timeout period is to prevent false-positive cases.

You can also configure a timeout value with a range of 3 to 20 seconds; the default timeout value is 5 seconds. This timer starts at the end of the hold-timeout interval. During the timeout period, the secondary vPC peer device checks for vPC peer-keepalive hello messages from the primary vPC peer device. If the secondary vPC peer device receives a single hello message, that device disables all vPC interfaces on the secondary vPC peer device.

The difference between the hold-timeout and the timeout parameters is as follows:

- During the hold-timeout, the vPC secondary device does not take any action based on any keepalive messages received, which prevents the system taking action when the keepalive might be received just temporarily, such as if a supervisor fails a few seconds after the vPC Peer-Link goes down.
- During the timeout, the vPC secondary device takes action to become the vPC primary device if no keepalive message is received by the end of the configured interval.

See the "Configuring vPC Keepalive Link and Messages" section for information about configuring the timer for the keepalive messages.



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

Peer-keepalive IP addresses must be global unicast addresses. Link-local addresses are not supported.

Use the command-line interface (CLI) to configure the interfaces you are using the vPC peer-keepalive messages as trusted ports. Leave the precedence at the default (6) or configure it higher.

# vPC Peer-Gateway

You can configure vPC peer devices to act as the gateway even for packets that are destined to the vPC peer device's MAC address.

Use the **peer-gateway** command to configure this feature.



**Note** The **peer-gateway exclude-vlan** command that is used when configuring a VLAN interface for Layer 3 backup routing on vPC peer devices is not supported.

Some network-attached storage (NAS) devices or load balancers might have features that help to optimize the performances of particular applications. These features enable the device to avoid a routing-table lookup when responding to a request that originated from a host that is not locally attached to the same subnet. Such devices might reply to traffic using the MAC address of the sender Cisco Nexus 3550-T device rather than the common VRRP gateway. This behavior is noncomplaint with some basic Ethernet RFC standards. Packets that reach a vPC device for the nonlocal router MAC address are sent across the vPC Peer-Link and could be dropped by the built in vPC loop avoidance mechanism if the final destination is behind another vPC.

The vPC peer-gateway capability allows a vPC switch to act as the active gateway for packets that are addressed to the router MAC address of the vPC peer. This feature enables local forwarding of packets without the need to cross the vPC Peer-Link. In this scenario, the feature optimizes use of the vPC Peer-Link and avoids potential traffic loss.

Configuring the peer-gateway feature must be done on both primary and secondary vPC peers and is nondisruptive to the operations of the device or to the vPC traffic. The vPC peer-gateway feature can be configured globally under the vPC domain submode.

When you enable this feature, Cisco NX-OS automatically disables IP redirects on all interface VLANs mapped over a vPC VLAN to avoid generation of IP redirect messages for packets switched through the peer gateway router.

# vPC Domain

You can use the vPC domain ID to identify the vPC Peer-Links and the ports that are connected to the vPC downstream devices.

The vPC domain is also a configuration mode that you use to configure the keepalive messages and other vPC Peer-Link parameters rather than accept the default values. See the “Configuring vPCs” section for more information about configuring these parameters.

To create a vPC domain, you must first create a vPC domain ID on each vPC peer device using a number from 1 to 1000. You can have only one vPC domain per vPC peer.

You must explicitly configure the port channel that you want to act as the vPC Peer-Link on each device. You associate the port channel that you made a vPC Peer-Link on each device with the same vPC domain ID to form a single vPC domain. Within this domain, the system provides a loop-free topology and Layer 2 multipathing.

You can only configure these port channels and vPC Peer-Links statically. You can configure the port channels and vPC Peer-Links either using LACP or no protocol. We recommend that you use LACP with the interfaces in active mode to configure port channels in each vPC, which ensures an optimized, graceful recovery in a port-channel failover scenario and provides configuration checks against configuration mismatches among the port channels themselves.

The vPC peer devices use the vPC domain ID that you configure to automatically assign a unique vPC system MAC address. Each vPC domain has a unique MAC address that is used as a unique identifier for the specific vPC-related operations, although the devices use the vPC system MAC addresses only for link-scope operations, such as LACP. We recommend that you create each vPC domain within the contiguous Layer 2 network with a unique domain ID. You can also configure a specific MAC address for the vPC domain, rather than having the Cisco NX-OS software assign the address.

See the “vPC and Orphan Ports” section for more information about displaying the vPC MAC table.

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

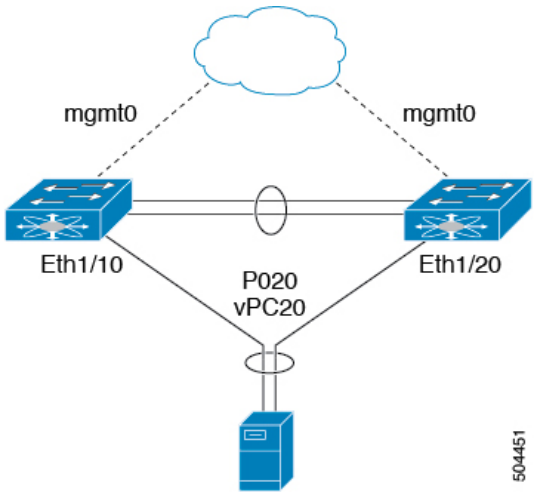


**Note** When manually configuring the system priority, you must ensure that you assign the same priority value on both vPC peer devices. If the vPC peer devices have different system priority values, vPC does not come up.

# vPC Topology

The following figure shows a basic configuration in which the Cisco Nexus 3550-T device ports are directly connected to another switch or host and are configured as part of a port channel that becomes part of a vPC.

*Figure 8: Switch vPC Topology*




In the figure, vPC 20 is configured on port channel 20, which has Eth1/10 on the first device and Eth1/20 on the second as member ports.

# Compatibility Parameters for vPC Interfaces

Many configuration and operational parameters must be identical on all interfaces in the vPC. We recommend that you configure the Layer 2 port channels that you use for the vPC Peer-Link in trunk mode.

After you enable the vPC feature and configure the vPC Peer-Link on both vPC peer devices, Cisco Fabric Services (CFS) messages provide a copy of the configuration on the local vPC peer device configuration to the remote vPC peer device. The system then determines whether any of the crucial configuration parameters differ on the two devices. (See the "vPC and Orphan Ports" section for more information about CFS.)

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

**Note** The port channel compatibility parameters must be the same for all the port channel members on the physical switch. You cannot configure shared interfaces to be part of a vPC.

The compatibility check process for vPCs differs from the compatibility check for regular port channels.

See the "Configuring Port Channels" chapter for information about regular port channels.

## Configuration Parameters That Must Be Identical

The configuration parameters in this section must be configured identically on both devices of the vPC Peer-Link; otherwise, the vPC moves fully or partially into a suspended mode.

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

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

The devices automatically check for compatibility for some of these parameters on the vPC interfaces. The per-interface parameters must be consistent per interface, and the global parameters must be consistent globally:

- Port-channel mode: on, off, or active (port-channel mode can, however, be active/passive on each side of the vPC peer)
- Trunk mode per channel:
  - Native VLAN
  - VLANs allowed on trunk
  - Tagging of native VLAN traffic
- Spanning Tree Protocol (STP) mode
- STP region configuration for Multiple Spanning Tree
- Enable/disable state per VLAN
- STP global settings:
  - Bridge Assurance setting
  - Port type setting
  - Loop Guard settings
- STP interface settings:
  - Port type setting
  - Loop Guard
  - Root Guard

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

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

## Configuration Parameters That Should Be Identical

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

- MAC aging timers
- Static MAC entries
- VLAN interface—Each device on the end of the vPC Peer-Link must have a VLAN interface configured for the same VLAN on both ends and they must be in the same administrative and operational mode. Those VLANs configured on only one device of the vPC Peer-Link do not pass traffic using the vPC or vPC Peer-Link. You must create all VLANs on both the primary and secondary vPC devices, or the VLAN will be suspended.
- All ACL configurations and parameters
- STP interface settings:
  - BPDU Filter
  - BPDU Guard
  - Cost
  - Link type
  - Priority
  - VLANs (Rapid PVST+)
- Dynamic Host Configuration Protocol (DHCP) snooping
- Internet Group Management Protocol (IGMP) snooping
- All routing protocol configurations

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

## Consequences of Parameter Mismatches

You can configure the graceful consistency check feature, which suspends only the links on the secondary peer device when a mismatch is introduced in a working vPC. This feature is configurable only in the CLI and is enabled by default.

The graceful consistency-check command is configured by default.

As part of the consistency check of all parameters from the list of parameters that must be identical, the system checks the consistency of all VLANs.

The vPC remains operational, and only the inconsistent VLANs are brought down. This per-VLAN consistency check feature cannot be disabled and does not apply to Multiple Spanning Tree (MST) VLANs.

# vPC Number

Once you have created the vPC domain ID and the vPC Peer-Link, you create port channels to attach the downstream device to each vPC peer device. That is, you create one port channel to the downstream device

from the primary vPC peer device and you create another port channel to the downstream device from the secondary peer device.

**Note** We recommend that you configure the ports on the downstream devices that connect to a host or a network device that is not functioning as a switch or a bridge as STP edge ports.

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

**Note** The vPC number that you assign to the port channel that connects to the downstream device from the vPC peer device must be identical on both vPC peer devices.

# Moving Other Port Channels into a vPC

**Note** You must attach a downstream device using a port channel to both vPC peer devices.

To connect to the downstream device, you create a port channel to the downstream device from the primary vPC peer device and you create another port channel to the downstream device from the secondary peer device. On each vPC peer device, you assign a vPC number to the port channel that connects to the downstream device. You will experience minimal traffic disruption when you are creating vPCs.

# vPC Interactions with Other Features

## vPC and LACP

LACP uses the system MAC address of the vPC domain to form the LACP Aggregation Group (LAG) ID for the vPC. (See the “Configuring Port Channels” chapter for information about LAG-ID and LACP.)

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

We recommend that you manually configure the system priority on the vPC Peer-Link devices to ensure that the vPC Peer-Link devices have a higher LACP priority than the downstream connected devices. A lower numerical value system priority means a higher LACP priority.

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

## vPC Peer-Links and STP

Although vPCs provide a loop-free Layer 2 topology, STP is still required to provide a fail-safe mechanism to protect against any incorrect or defective cabling or possible misconfiguration. When you first bring up a vPC, STP reconverges. STP treats the vPC Peer-Link as a special link and always includes the vPC Peer-Link in the STP active topology.

We recommend that you set all the vPC Peer-Link interfaces to the STP network port type so that Bridge Assurance is automatically enabled on all vPC Peer-Links. We also recommend that you do not enable any of the STP enhancement features on vPC Peer-Links. If the STP enhancements are already configured, they do not cause any problems for the vPC Peer-Links..

When you are running both MST and Rapid PVST+, ensure that the PVST simulation feature is correctly configured.



**Note** You must configure a list of parameters to be identical on the vPC peer devices on both sides of the vPC Peer-Link. See the “Compatibility Parameters for vPC Interfaces” section for information about these required matched settings.

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

The primary vPC device synchronizes the STP state on the vPC secondary peer device using Cisco Fabric Services over Ethernet (CFSoE). See the “vPC and Orphan Ports” section for information about CFSoE.

The STP process for vPC also relies on the periodic keepalive messages to determine when one of the connected devices on the vPC Peer-Link fails. See the “Peer-Keepalive Link and Messages” section for information about these messages.

The vPC manager performs a proposal/handshake agreement between the vPC peer devices that set the primary and secondary devices and coordinates the two devices for STP. The primary vPC peer device then controls the STP protocol on both the primary and secondary devices. We recommend that you configure the primary vPC peer device as the STP primary root device and configure the secondary VPC device to be the STP secondary root device.

If the primary vPC peer device fails over to the secondary vPC peer device, there is no change in the STP topology.

The BPDUs uses the MAC address set for the vPC for the STP bridge ID in the designated bridge ID field. The vPC primary device sends these BPDUs on the vPC interfaces.

You must configure both ends of vPC Peer-Link with the identical STP configuration for the following parameters:

- STP global settings:
  - STP mode
  - STP region configuration for MST
  - Enable/disable state per VLAN
  - Bridge Assurance setting
  - Port type setting

- Loop Guard settings

- STP interface settings:
  - Port type setting
  - Loop Guard
  - Root Guard



**Note** If any of these parameters are misconfigured, the Cisco NX-OS software suspends all interfaces in the vPC. Check the syslog and enter the **show vpc brief** command to see if the vPC interfaces are suspended.

Ensure that the following STP interface configurations are identical on both sides of the vPC Peer-Links or you may see unpredictable behavior in the traffic flow:

- BPDU Filter
- BPDU Guard
- Cost
- Link type
- Priority
- VLANs (PVRST+)



**Note** Display the configuration on both sides of the vPC Peer-Link to ensure that the settings are identical.

You can use the **show spanning-tree** command to display information about the vPC when that feature is enabled.

We recommend that you configure the ports on the downstream devices as STP edge ports. You should configure all host ports connected to a switch as STP edge ports.

## vPC Peer Switch

The vPC peer switch feature was added to Cisco NX-OS to address performance concerns around STP convergence. This feature allows a pair of Cisco Nexus 3550-T devices to appear as a single STP root in the Layer 2 topology. This feature eliminates the need to pin the STP root to the vPC primary switch and improves vPC convergence if the vPC primary switch fails.

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

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

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

## vPC and ARP or ND

A feature was added to Cisco NX-OS to address table synchronization across vPC peers using the reliable transport mechanism of the Cisco Fabric Service over Ethernet (CFSoE) protocol. You must enable the **ip arp synchronize** command to support faster convergence of address tables between the vPC peers. This convergence overcomes the delay that occurs in ARP table restoration for IPv4 or ND table restoration when the vPC Peer-Link port channel flaps or when a vPC peer comes back online.

## vPC Multicast— IGMP, and IGMP Snooping

The software keeps the multicast forwarding state synchronized on both of the vPC peer devices. The IGMP snooping process on a vPC peer device shares the learned group information with the other vPC peer device through the vPC Peer-Link; the multicast states are always synchronized on both vPC peer devices.

Each vPC peer is a Layer 2 device. Multicast traffic flows from only one of the vPC peer devices. You might see duplicate packets in the following scenarios:

- Orphan hosts
- When the source and receivers are in the Layer 2 vPC cloud in different VLANs with multicast routing enabled and a vPC member link goes down.

You might see negligible traffic loss:

- When you reload the vPC peer device that is forwarding the traffic.

Overall multicast convergence times are scale and vPC role change duration dependent.

The following outlines vPC IGMP/IGMP snooping:

- vPC IGMP/IGMP snooping—The IGMP process in vPC mode synchronizes the designated router (DR) information on both vPC peer devices. Dual DRs are available for IGMP when you are in vPC mode. Dual DRs are not available when you are not in vPC mode, because both vPC peer devices maintain the multicast group information between the peers.

You should enable or disable IGMP snooping identically on both vPC peer devices, and all the feature configurations should be identical. IGMP snooping is on by default.

**Note** Cisco Nexus 3550-T does not support PIM on a vPC VLAN.

## vPC Peer-Links and Routing

The First Hop Redundancy Protocols (FHRPs) interoperate with vPCs. The Virtual Router Redundancy Protocol (VRRP) interoperates with vPCs. We recommend that you dual-attach all Layer 3 devices to both vPC peer devices.

The primary FHRP device responds to ARP requests, even though the secondary vPC device forwards the data traffic.

To simplify initial configuration verification and vPC troubleshooting, you can configure the primary vPC peer device with the FHRP active router highest priority.

When the primary vPC peer device fails over to the secondary vPC peer device, the FHRP traffic continues to flow seamlessly.

We recommend that you configure routing adjacency between the two vPC peer devices to act as a backup routing path. If one vPC peer device loses Layer 3 uplinks, the vPC can redirect the routed traffic to the other vPC peer device and leverage its active Layer 3 uplinks.

You can configure the inter-switch link for a backup routing path in the following ways:

- Create a Layer 3 link between the two vPC peer devices.
- Use the non-VPC VLAN trunk with a dedicated VLAN interface.
- Use a vPC Peer-Link with a dedicated VLAN interface.

We do not recommend that you configure the burnt-in MAC address option (use-bia) for VRRP or manually configure virtual MAC addresses for any FHRP protocol in a vPC environment because these configurations can adversely affect vPC load balancing. The VRRP use-bia option is not supported on vPCs. When you are configuring custom MAC addresses, you must configure the same MAC address on both vPC peer devices.

You can use the **delay restore** command to configure a restore timer that delays the vPC coming back up until after the peer adjacency forms and the VLAN interfaces are back up. This feature enables you to avoid packet drops when the routing tables might not be converged before the vPC is once again passing traffic. Use the **delay restore** command to configure this feature.

To delay the VLAN interfaces on the restored vPC peer device from coming up, use the **interfaces-vlan** option of the **delay restore** command.

## CFSoE

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

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

The CFSoE transport is local to each VDC.

You can use the **show mac address-table** command to display the MAC addresses that CFSoE synchronizes for the vPC Peer-Link.

**Note** Do not enter the **no cfs eth distribute** or the **no cfs distribute** command. You must enable CFSoE for vPC functionality. If you do enter either of these commands with vPC enabled, the system displays an error message.

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

**Note** The software does not support CFS regions.

## vPC and Orphan Ports

When a device that is not vPC-capable connects to each peer, the connected ports are known as orphan ports because they are not members of a vPC. The device's link to one peer will be active (forwarding) and the other link will be standby (blocking) due to STP.

If a vPC Peer-Link failure or restoration occurs, an orphan port's connectivity might be bound to the vPC failure or restoration process. For example, if a device's active orphan port connects to the secondary vPC peer, the device loses any connections through the primary peer if a vPC Peer-Link failure occurs and the vPC ports are suspended by the secondary peer. If the secondary peer were to also suspend the active orphan port, the device's standby port becomes active, provides a connection to the primary peer, and restores connectivity. You can configure in the CLI that specific orphan ports are suspended by the secondary peer when it suspends its vPC ports and are restored when the vPC is restored.

# vPC Recovery After an Outage

In a data center outage, both the vPC peer in vPC domain get reloaded. Occasionally only one peer can be restored. With no functioning peer-keepalive or vPC Peer-Link, the vPC cannot function normally, a method might be available to allow vPC services to use only the local ports of the functional peer.

## Autorecovery

You can configure the Cisco Nexus 3550-T device to restore vPC services when its peer fails to come online by using the **auto-recovery** command. You must save this setting in the startup configuration. On reload, if the vPC Peer-Link is down and three consecutive peer-keepalive messages are lost, the secondary device assumes the primary STP role and the primary LACP role. The software reinitializes the vPCs, bringing up its local ports. Because there are no peers, the consistency check is bypassed for the local vPC ports. The device elects itself to be the STP primary regardless of its role priority and also acts as the primary device for LACP port roles.

## vPC Peer Roles After a Recovery

When the other peer device completes its reload and adjacency forms, the following process occurs:

1. The first vPC peer maintains its current role to avoid any transition reset to other protocols. The peer accepts the other available role.
2. When an adjacency forms, consistency checks are performed and appropriate actions are taken.

# Guidelines and Limitations

vPCs have the following configuration guidelines and limitations:

- Make sure that both vPC peers are in the same mode (regular mode or enhanced mode) before performing a non-disruptive upgrade.

**Note** vPC peering between an enhanced ISSU mode (boot mode lxc) configured switch and a non-enhanced ISSU mode switch is not supported.

- **show** commands with the **internal** keyword are not supported.
- Cisco Nexus 3550-T switches do not support NAT on vPC topology.
- vPC peers must run the same Cisco NX-OS release. During a software upgrade, make sure to upgrade the primary vPC peer first.
- All ports for a given vPC must be in the same VDC.
- You must enable vPCs before you can configure them.
- You must configure the peer-keepalive link and messages before the system can form the vPC Peer-Link.
- Only Layer 2 port channels can be in vPCs.
- Layer 3 multicast over vPC is not supported.
- You must configure both vPC peer devices; the configuration is not sent from one device to the other.
- To configure multilayer (back-to-back) vPCs, you must assign unique vPC domain ID for each respective vPC.
- Check that the necessary configuration parameters are compatible on both sides of the vPC Peer-Link. See the “Compatibility Parameters for vPC Interfaces” section for information about compatibility recommendations.
- You might experience minimal traffic disruption while configuring vPCs.
- The software does not support BIDR PIM on vPCs.
- The software does not support DHCP snooping, DAI, or IPSG in a vPC environment.
- The software does not support CFS regions.
- Port security is not supported.
- When **peer-switch** features are configured under **vpc domain** configuration mode on two Cisco Nexus 3550-T switches, the spanning-tree root changes even for VLANs that are not enabled on the vPC Peer-Link. Both the switches act as one system with one MAC address as the bridge address. This is true even for non-vPC mst-instance or VLANs. Therefore, a non vPC Peer-Link between the two switches gets blocked as a backup link. This is an expected behavior.
- We recommend that you configure all the port channels in the vPC using LACP with the interfaces in active mode.

- Back-to-back, multilayer vPC topologies require unique domain IDs on each respective vPC.
- Having the same Virtual Router Redundancy Protocol (VRRP) group on all nodes on a double sided vPC is supported.
- 
- When using vPCs, we recommend that you use default timers for VRRP. There is no advantage in convergence times when using aggressive timers in vPC configurations.
- If you configure open shortest path first (OSPF) in a vPC environment, use the following timer commands in router configuration mode on the core switch to ensure fast OSPF convergence when a vPC Peer-Link is shut down:

```
switch (config-router)# timers throttle spf 1 50 50
switch (config-router)# timers lsa-arrival 10
```

- The STP port cost is fixed to 200 in a vPC environment.
- To accommodate increased traffic when the vPC goes down and traffic needs to cross the vPC Peer-Link, it is a best practice to use multiple high bandwidth interfaces across linecards for the vPC Peer-Link.
- The **vpc orphan-ports suspend** command also applies to ports in non-vPC VLANs and Layer 3 ports. However, it is recommended to be used with ports in VPC VLANs.
- vPC STP hitless role change feature is supported.
- vPC role change can be performed from either of the peer devices.
- When forming a vPC domain between two Cisco Nexus 3550-T Series switches, both switches must be the exact same model to form a supported vPC domain.
- If the original secondary device has higher role priority value than the original primary device, role swapping cannot be performed. Change the role priority on either vPC device so that the value of the original secondary device is lower than the original primary one. To view the existing role of a device, use the show vpc role command on local and peer switch.
- Always check the existing configured role priority before configuring vPC hitless role change feature
- In a vPC domain, enable the peer-switch command, where both vPC peers have same STP priorities, and ensure it is operational before issuing a role change. If you do not enable the peer-switch command, it can lead to convergence issues. Use **show spanning-tree summary | grep peer** command to verify whether the peer vPC switch is operational or not.
- All the devices that are attached to a vPC domain must be dual homed.
- Layer 3 over vPC is supported on Cisco Nexus 3550-T Series switches for Layer 3 unicast communication only. Layer 3 over vPC is not supported for Layer 3 multicast traffic. For more information see the *Best Practices for Layer 3 and vPC Configuration* section
- The default behavior with Layer 3 peer-router and TTL=1 packet destined to IP of vPC peer is to punt packet to CPU and then forward the software to vPC peer.

# Best Practices for Layer 3 and vPC Configuration

This section describes best practices for using and configuring Layer 3 with vPC.

## Layer 3 and vPC Configuration Overview

When a Layer 3 device is connected to a vPC domain through a vPC, it has the following views:

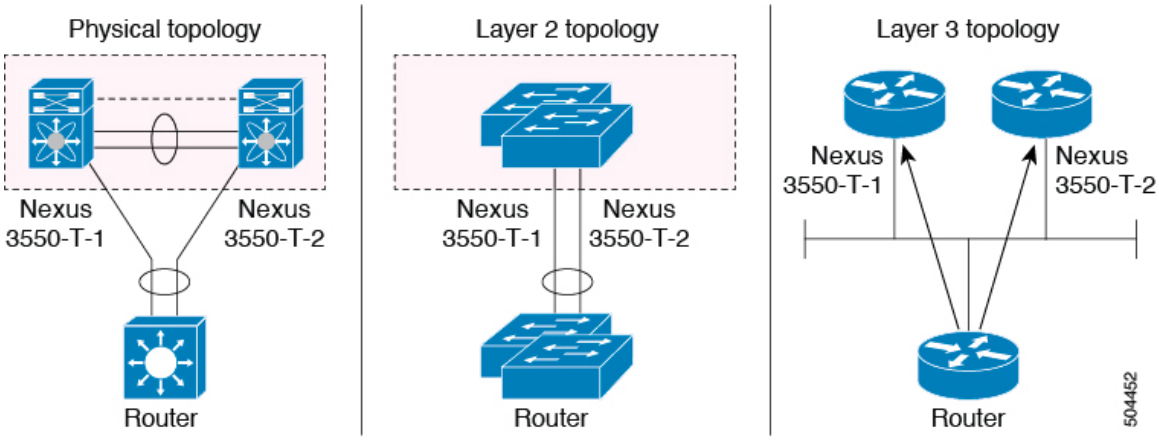
- At Layer 2, the Layer 3 device sees a unique Layer 2 switch presented by the vPC peer devices.
- At Layer 3, the Layer 3 device sees two distinct Layer 3 devices (one for each vPC peer device).

vPC is a Layer 2 virtualization technology, so at Layer 2, both vPC peer devices present themselves as a unique logical device to the rest of the network.

There is no virtualization technology at Layer 3, so each vPC peer device is seen as a distinct Layer 3 device by the rest of the network.

The following figure illustrates the two different Layer 2 and Layer 3 views with vPC.

***Figure 9: Different Views for vPC Peer Devices***

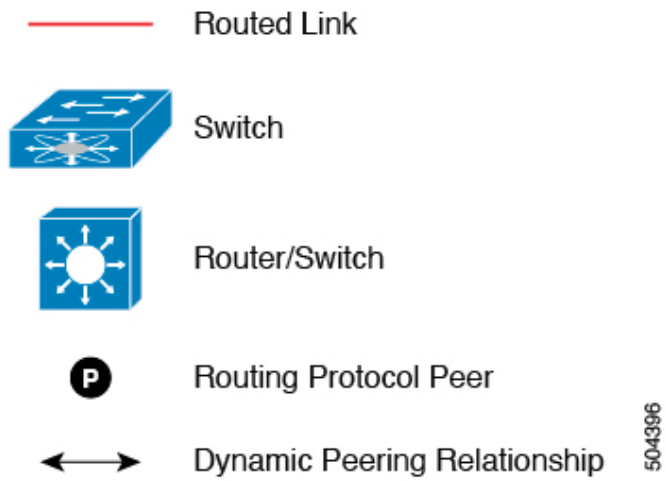



## Supported Topologies for Layer 3 and vPC

This section contains examples of Layer 3 and vPC network topologies.

There are two approaches for Layer 3 and vPC interactions. The first one is by using dedicated Layer 3 links to connect the Layer 3 devices to each vPC peer device. The second one is by allowing the Layer 3 devices to peer with the SVIs defined on each of the vPC peer device, on a dedicated VLAN that is carried on the vPC connection. The following sections describe all the supported topologies leveraging the elements that are described in the legends in the following figure.

*Figure 10: Legend*

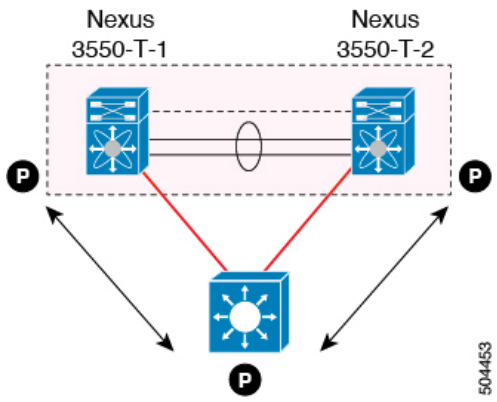



## Peering with an External Router Using Layer 3 Links

This example shows a topology that uses Layer 3 links to connect a Layer 3 device to the Cisco Nexus 3550-T switches that are part of the a vPC domain

**Note** Interconnecting the two entities together in this way allows to support Layer 3 unicast and multicast communication.

*Figure 11: Peering with an External Router Using Layer 3 Links*




Layer 3 devices can initiate Layer 3 routing protocol adjacencies with both vPC peer devices.

One or multiple Layer 3 links can be used to connect a Layer 3 device to each vPC peer device.

Follow these guidelines when connecting a Layer 3 device to the vPC domain using Layer 3 links:

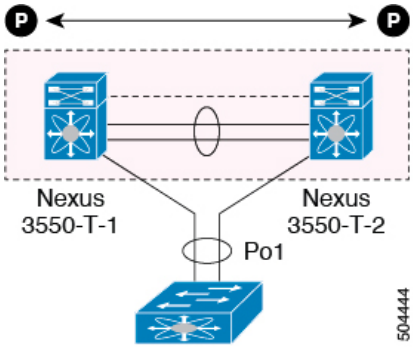
- Use separate Layer 3 links to connect Layer 3 devices to the vPCdomain. Each link represents a point-to-point Layer 3 connection and should get assigned an IP address taken from a small IP subnet (/30 or /31).

## Peering Between vPC Devices for a Backup Routing Path

This example shows peering between the two vPC peer devices with a Layer 3 backup routed path. If the Layer 3 uplinks on vPC peer device 1 or vPC peer device 2 fail, the path between the two peer devices is used to redirect traffic to the switch that has the Layer 3 uplinks in the up state.

The Layer 3 backup routing path can be implemented using a dedicated interface VLAN (such as SVI) over the vPC Peer-Link or by using dedicated Layer 2 or Layer 3 links across the two vPC peer devices.

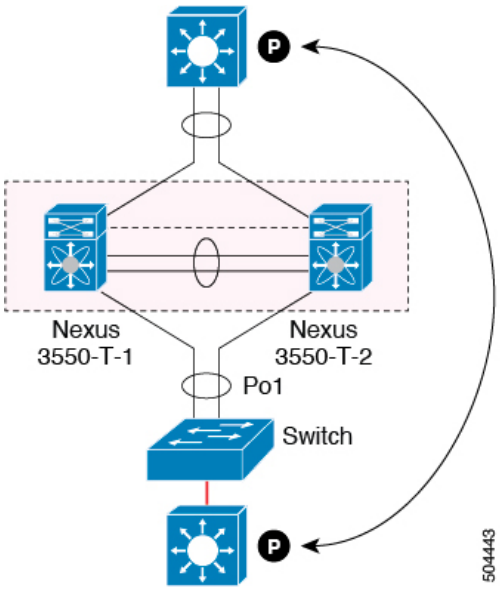
***Figure 12: Peering Between vPC Devices for a Backup Routing Path***




## Direct Layer 3 Peering Between Routers

In this scenario, the Nexus 3550-T devices part of the vPC domain are simply used as a Layer 2 transit path to allow the routers connected to them to establish Layer 3 peering and communication.

***Figure 13: Peering Between Routers***




The Layer 3 devices can peer with each other in following two methods. Peering also depends on the specific device deployed for this role.

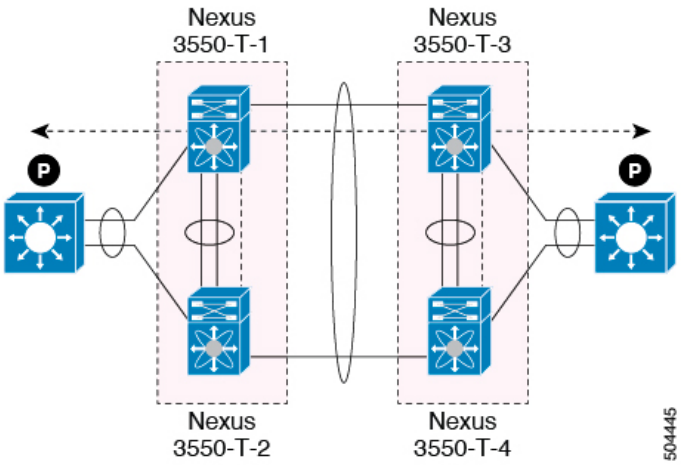
- Defining a VLAN network interface (SVI) for a VLAN that is extended between the Layer 3 devices through the intermediate Cisco Nexus 3550-T vPC peer switches.

- Defining a Layer 3 port-channel interface on each Layer 3 device and establishing a point-to-point Layer 3 peering.

## Peering Between Two Routers with vPC Devices as Transit Switches

This example is similar to the peering between routers topology. In this case also, the Cisco Nexus 3550-T devices that are part of the same vPC domain are only used as Layer 2 transit paths. The difference here is that there are two pairs of Cisco Nexus 3550-T switches. Each switch that is connected with a Layer 3 device using a vPC connection, also establishes a back-to-back vPC connection between them. The difference is that the vPC domains are only used as Layer 2 transit paths.

***Figure 14: Peering Between Two Routers with vPC Devices as Transit Switches***




This topology is commonly used when you want to establish connectivity between separate data centers that are interconnected with direct links (dark fibers or DWDM circuits). The two pairs of Cisco Nexus 3550-T switches, in this case, provide only Layer 2 extension services, allowing the Layer 3 devices to peer with each other at Layer 3.

## Peering with an External Router on Parallel Interconnected Routed Ports

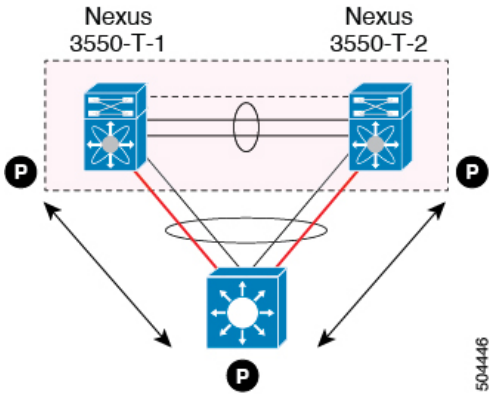
When you require both routed and bridged traffic, use individual Layer 3 links for routed traffic and a separate Layer 2 port-channel for bridged traffic, as shown in following figure.

The Layer 2 links are used for bridged traffic (traffic staying in the same VLAN) or inter-VLAN traffic (assuming vPC domain hosts the interface VLAN and associated VRRP configuration).

The Layer 3 links are used for routing protocol peering adjacency with each vPC peer device.

The purpose of this topology is to attract specific traffic to go through the Layer 3 device. Layer 3 links are also used to carry routed traffic from a Layer 3 device to the vPC domain.

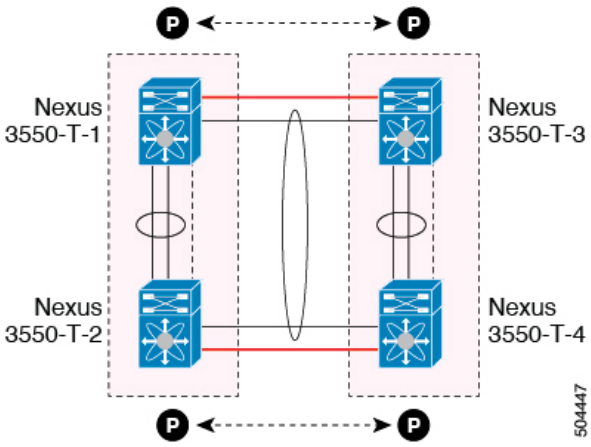
*Figure 15: Peering with an External Router on Parallel Interconnected Routed Ports*




## Peering between vPC Switch Pairs on Parallel Interconnected Routed Ports

An alternative design to what is shown in the previous section (Peering Between Two Routers with vPC Devices as Transit Switches), uses two pairs of Cisco Nexus 3550-T switches that are deployed in each data center for providing both Layer 2 and Layer 3 extension services. When routing protocol peering adjacency is required to be established between the two pairs of Cisco Nexus 3550-T devices, the best practice is to add dedicated Layer 3 links between the two sites as shown in the following example.

*Figure 16: Peering Over a vPC Interconnection on Parallel Interconnected Routed Ports*



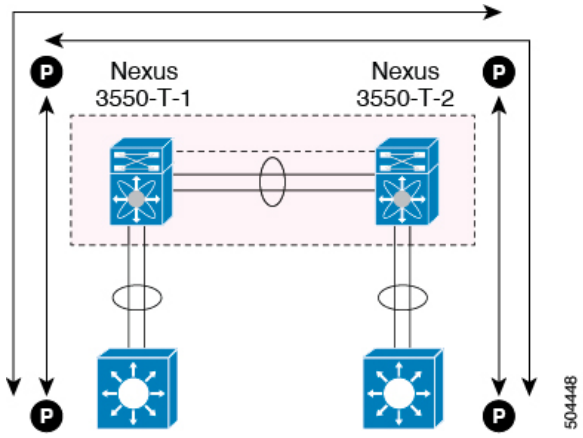

The back-to-back vPC connection between the two data centers carry bridged traffic or inter-VLAN traffic while the dedicated Layer 3 links carry the routed traffic across the two sites.

## Peering Over a PC Interconnection and Dedicated Interswitch Link Using non-vPC VLAN

This example shows when the Layer 3 device is single-attached to the vPC domain, you can use a non-vPC VLAN with a dedicated inter-switch link to establish the routing protocol peering adjacency between the Layer 3 device and each vPC peer device. However, the non-vPC VLAN must be configured to use a static MAC that is different than the vPC VLAN.

**Note** Configuring the vPC VLAN (and vPC Peer-Link) for this purpose is not supported.

*Figure 17: Peering Over a PC Interconnection and Dedicated Interswitch Link Using non-vPC VLAN*

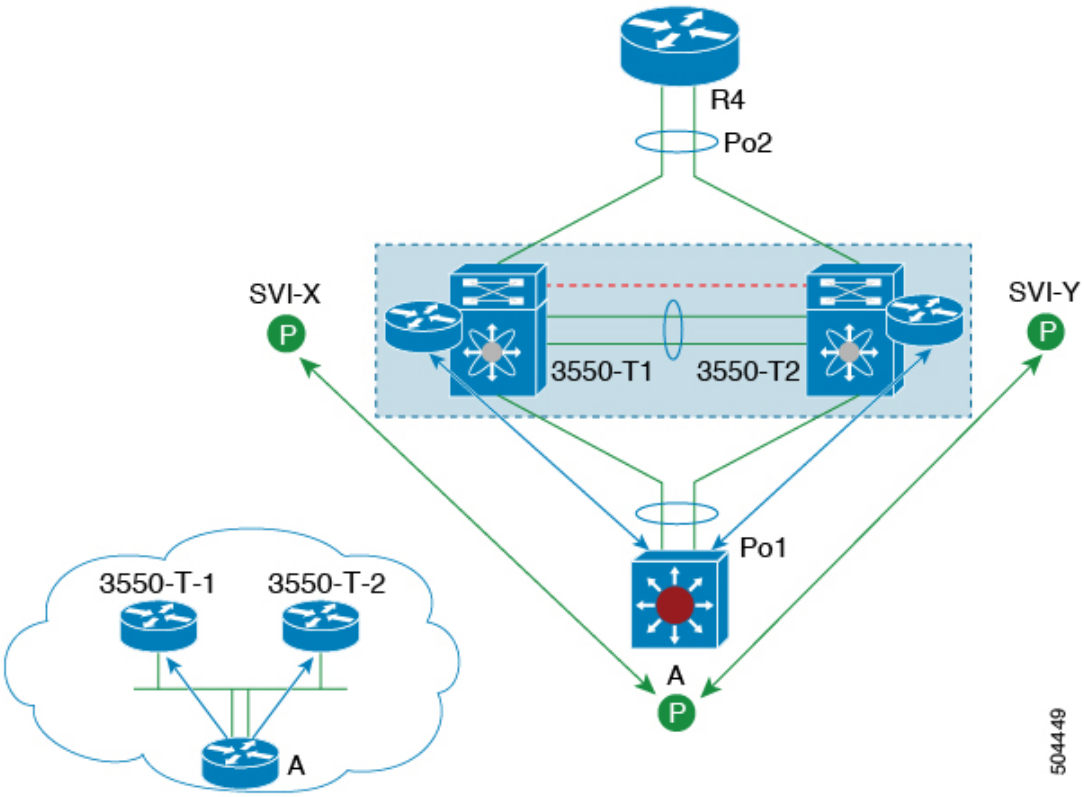



## Peering Directly Over a vPC Connection

An alternate method to establish Layer 3 peering between a Layer 3 router and a pair of Cisco Nexus 3550-T vPC switches.

**Note** Peering directly over a vPC connection is supported only for Layer 3 unicast communication but not for Layer 3 multicast traffic. If you require Layer 3 multicast, you must establish peering over dedicated Layer 3 links

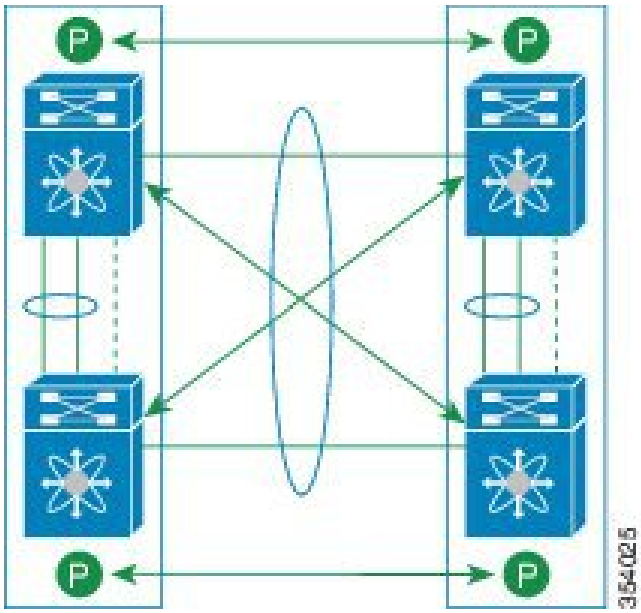
*Figure 18: Supported: Peering Over a vPC Interconnection Where the Router Peers with Both the vPC Peers.*




In this scenario, the Layer 3 peering between the external router and the Cisco Nexus 3550-T switches that are part of a same vPC domain is established directly on a VLAN carried on the vPC connection. The external router in this case peers with SVI interfaces defined on each vPC device. As for the scenario shown in previous figure 12, the external router could use an SVI or a Layer 3 Port-Channel to peer with the vPC devices (multiple SVIs or Port-Channel subinterfaces could be used for a multi-VRF deployment).

This deployment model requires configuring **layer3 peer-router** command as part of the vPC domain. You can adopt the same approach for establishing Layer 2 and Layer 3 connectivity on a vPC back-to-back connection established between two separate pairs of vPC switches.

*Figure 19: Supported: Peering Over a vPC Interconnection Where Each Nexus Device Peers with Two vPC Peers.*




In this deployment model, SVI interfaces in the same VLAN is configured on all the four Cisco Nexus 3550-T switches to establish routing peering and connectivity between them.

# Default Settings

The following table lists the default settings for vPC parameters.

*Table 8: Default vPC Parameters*

| Parameters | Default |
|---|---|
| vPC system priority | 32667 |
| vPC peer-keepalive message | Disabled |
| vPC peer-keepalive interval | 1 second |
| vPC peer-keepalive timeout | 5 seconds |
| vPC peer-keepalive UDP port | 3200 |

# Configuring vPCs

**Note** You must use these procedures on both devices on both sides of the vPC Peer-Link. You configure both of the vPC peer devices using these procedures.

This section describes how to configure vPCs using the command-line interface (CLI).

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

# Enabling vPCs

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

### Procedure

| | Command or Action | Purpose |
|---|---|---|
| **Step 1** | **configure terminal**<br>**Example:**<br>`switch# configure terminal`<br>`switch(config)#` | Enters global configuration mode. |
| **Step 2** | **feature vpc**<br>**Example:**<br>`switch(config)# feature vpc` | Enables vPCs on the device. |
| **Step 3** | **exit**<br>**Example:**<br>`switch(config)# exit`<br>`switch#` | Exits global configuration mode. |
| **Step 4** | **show feature**<br>**Example:**<br>`switch# show feature` | (Optional) Displays which features are enabled on the device. |
| **Step 5** | **copy running-config startup-config**<br>**Example:**<br>`switch# copy running-config startup-config` | (Optional) Copies the running configuration to the startup configuration. |

### Example

This example shows how to enable the vPC feature:

```
switch# configure terminal
switch(config)# feature vpc
switch(config)# exit
switch(config)#
```

# Disabling vPCs



**Note** When you disable the vPC functionality, the device clears all the vPC configurations.

### Procedure

| | Command or Action | Purpose |
|---|---|---|
| **Step 1** | **configure terminal**<br>**Example:**<br>`switch# configure terminal`<br>`switch(config)#` | Enters global configuration mode. |
| **Step 2** | **no feature vpc**<br>**Example:**<br>`switch(config)# no feature vpc` | Disables vPCs on the device. |
| **Step 3** | **exit**<br>**Example:**<br>`switch(config)# exit`<br>`switch#` | Exits global configuration mode. |
| **Step 4** | **show feature**<br>**Example:**<br>`switch# show feature` | (Optional) Displays which features are enabled on the device. |
| **Step 5** | **copy running-config startup-config**<br>**Example:**<br>`switch# copy running-config startup-config` | (Optional) Copies the running configuration to the startup configuration. |

### Example

This example shows how to disable the vPC feature:

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

# Creating a vPC Domain and Entering vpc-domain Mode

You can create a vPC domain and put the vPC Peer-Link port channels into the identical vPC domain on both vPC peer devices. Use a unique vPC domain number throughout a single vPC domain . This domain ID is used to automatically to form the vPC system MAC address.

You can also use this command to enter vpc-domain command mode.

### Procedure

| | Command or Action | Purpose |
|---|---|---|
| **Step 1** | **configure terminal**<br>**Example:**<br>`switch# configure terminal`<br>`switch(config)#` | Enters global configuration mode. |
| **Step 2** | **vpc domain** *domain-id* [**shut** \| **no shut**]<br>**Example:**<br>`switch(config)# vpc domain 5`<br>`switch(config-vpc-domain)#` | Creates a vPC domain on the device, and enters vpc-domain configuration mode for configuration purposes. There is no default; the range is from 1 to 1000. |
| **Step 3** | **exit**<br>**Example:**<br>`switch(config)# exit`<br>`switch#` | Exits vpc-domain configuration mode. |
| **Step 4** | **show vpc brief**<br>**Example:**<br>`switch# show vpc brief` | (Optional) Displays brief information about each vPC domain. |
| **Step 5** | **copy running-config startup-config**<br>**Example:**<br>`switch# copy running-config startup-config` | (Optional) Copies the running configuration to the startup configuration. |

### Example

This example shows how to enter the vpc-domain command mode to configure an existing vPC domain:

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

# Configuring a vPC Keepalive Link and Messages

You can configure the destination IP for the peer-keepalive link that carries the keepalive messages. Optionally, you can configure other parameters for the keepalive messages.

**Note** You must configure the vPC peer-keepalive link before the system can form the vPC Peer-Link.

**Note** We recommend that you configure a default VRF instance and put a Layer 3 port from each vPC peer device into that VRF for the vPC peer-keepalive link. Do not use the vPC Peer-Link itself to send vPC peer-keepalive messages. For information about creating and configuring VRFs, see the *Cisco 3550-T Unicast Configuration Guide*. Ensure that both the source and destination IP addresses use for the peer-keepalive message are unique in your network. The management port and management VRF are the defaults for these keepalive messages.

### Before you begin

Ensure that you have enabled the vPC feature.

## Procedure

| | Command or Action | Purpose |
|---|---|---|
| **Step 1** | **configure terminal**<br>**Example:**<br>`switch# configure terminal`<br>`switch(config)#` | Enters global configuration mode. |
| **Step 2** | **vpc domain** *domain-id* [**shut** \| **no shut**]<br>**Example:**<br>`switch(config)# vpc domain 5`<br>`switch(config-vpc-domain)#` | Creates a vPC domain on the device, and enters vpc-domain configuration mode. |
| **Step 3** | **peer-keepalive destination** *ipaddress* [**hold-timeout** *secs* \| **interval** *msecs* {**timeout** *secs*} \| {**precedence** {*prec-value* \| **network** \| **internet** \| **critical** \| **flash-override** \| **flash** \| **immediate priority** \| **routine**}} \| **tos** {*tos-value* \| **max-reliability** \| **max-throughput** \| **min-delay** \| **min-monetary-cost** \| **normal**}} \|**tos-byte** *tos-byte-value*} \| **source** *ipaddress* \| **vrf** {*name* \| **management vpc-keepalive**}]<br>**Example:**<br>`switch(config-vpc-domain)# peer-keepalive destination 172.28.230.85`<br>`switch(config-vpc-domain)#` | Configures the IPv4 addresses for the remote end of the vPC peer-keepalive link.<br>**Note**<br>The system does not form the vPC Peer-Link until you configure a vPC peer-keepalive link.<br>The management ports and VRF are the defaults.<br>**Note**<br>We recommend that you configure a default VRF and use a Layer 3 port from each vPC peer device in that VRF for the vPC peer-keepalive link. |
| **Step 4** | **exit**<br>**Example:**<br>`switch(config)# exit`<br>`switch#` | Exits global configuration mode. |

| | Command or Action | Purpose |
|---|---|---|
| **Step 5** | **show vpc statistics**<br>**Example:**<br>`switch# show vpc statistics` | (Optional) Displays information about the configuration for the keepalive messages. |
| **Step 6** | **copy running-config startup-config**<br>**Example:**<br>`switch# copy running-config startup-config` | (Optional) Copies the running configuration to the startup configuration. |

### Example

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

```
switch# configure terminal
switch(config)# vpc domain 100
switch(config-vpc-domain)# peer-keepalive destination 172.168.1.2 source 172.168.1.1 vrf
vpc-keepalive
switch(config-vpc-domain)# exit
switch#
```

# Creating a vPC Peer-Link

You create the vPC Peer-Link by designating the port channel that you want on each device as the vPC Peer-Link for the specified vPC domain. We recommend that you configure the Layer 2 port channels that you are designating as the vPC Peer-Link in trunk mode and that you use two ports on separate modules on each vPC peer device for redundancy.

### Before you begin

Ensure that you have enabled the vPC feature.

### Procedure

| | Command or Action | Purpose |
|---|---|---|
| **Step 1** | **configure terminal**<br>**Example:**<br>`switch# configure terminal`<br>`switch(config)#` | Enters global configuration mode. |
| **Step 2** | **interface port-channel** *channel-number*<br>**Example:**<br>`switch(config)# interface port-channel 20`<br>`switch(config-if)#` | Selects the port channel that you want to use as the vPC Peer-Link for this device, and enters interface configuration mode. |

| | Command or Action | Purpose |
|---|---|---|
| **Step 3** | **switchport mode trunk**<br>**Example:**<br>`switch(config-if)# switchport mode trunk` | (Optional) Configures this interface in trunk mode. |
| **Step 4** | **switchport trunk allowed vlan** *vlan-list*<br>**Example:**<br>`switch(config-if)# switchport trunk allowed vlan 1-120,201-3967` | (Optional) Configures the permitted VLAN list. |
| **Step 5** | **vpc peer-link**<br>**Example:**<br>`switch(config-if)# vpc peer-link`<br>`switch(config-vpc-domain)#` | Configures the selected port channel as the vPC Peer-Link, and enters vpc-domain configuration mode. |
| **Step 6** | **exit**<br>**Example:**<br>`switch(config)# exit`<br>`switch#` | Exits vpc-domain configuration mode. |
| **Step 7** | **show vpc brief**<br>**Example:**<br>`switch# show vpc brief` | (Optional) Displays information about each vPC, including information about the vPC Peer-Link. |
| **Step 8** | **copy running-config startup-config**<br>**Example:**<br>`switch# copy running-config startup-config` | (Optional) Copies the running configuration to the startup configuration. |

### Example

This example shows how to configure a vPC Peer-Link:

```
switch# configure terminal
switch(config)# interface port-channel 20
switch(config-if)# switchport mode
switch(config-if)# switchport mode trunk
switch(config-if)# switchport trunk allowed vlan 1-255
switch(config-if)# vpc peer-link
switch(config-vpc-domain)# exit
switch(config)#
```

# Configuring a vPC Peer-Gateway

You can configure vPC peer devices to act as the gateway for packets that are destined to the vPC peer device's MAC address.

### Before you begin

Ensure that you have enabled the vPC feature.

### Procedure

| | Command or Action | Purpose |
|---|---|---|
| **Step 1** | **configure terminal**<br>**Example:**<br>`switch# configure terminal`<br>`switch(config)#` | Enters global configuration mode. |
| **Step 2** | **vpc domain** *domain-id* [**shut** \| **no shut**]<br>**Example:**<br>`switch(config-if)# vpc domain 5`<br>`switch(config-vpc-domain)#` | Creates a vPC domain if it does not already exist, and enters vpc-domain configuration mode. |
| **Step 3** | **peer-gateway**<br>**Example:**<br>`switch(config-vpc-domain)# peer-gateway`<br>**Note**<br>Disable IP redirects on all interface-vlans of this vPC domain for correct operation of this feature. | Enables Layer 3 forwarding for packets destined to the peer's gateway MAC address. |
| **Step 4** | **exit**<br>**Example:**<br>`switch(config)# exit`<br>`switch#` | Exits vpc-domain configuration mode. |
| **Step 5** | **show vpc brief**<br>**Example:**<br>`switch# show vpc brief` | (Optional) Displays information about each vPC, including information about the vPC Peer-Link. |
| **Step 6** | **copy running-config startup-config**<br>**Example:**<br>`switch# copy running-config startup-config` | (Optional) Copies the running configuration to the startup configuration. |

# Configuring Fast Convergence

Fast convergence feature is supported on Cisco Nexus 9000 Series platforms. You can enable or disable the vPC optimizations using this command. To achieve faster convergence, you must enable **[no] fast-convergence** on both vPC peers to achieve fast-convergence. The optimization is archived on secondary switch, vPC member ports and orphan ports with **vpc orphan-ports suspend** command is configured. In case the vPC

Peer-Link fails, these ports will be suspended immediately and traffic will be forwarded to primary vPC peer only to improve convergence.

Beginning with Cisco NX-OS Release 7.0(3)I7(1), fast convergence feature is supported on Cisco Nexus 9000 Series platforms. You can enable or disable the vPC optimizations using this command. To achieve faster convergence, you must enable **[no] fast-convergence** on both vPC peers to achieve fast-convergence. The optimization is archived on secondary switch, vPC member ports and orphan ports with **vpc orphan-ports suspend** command is configured. In case the vPC Peer-Link fails, these ports will be suspended immediately and traffic will be forwarded to primary vPC peer only to improve convergence.

### Procedure

| | Command or Action | Purpose |
|---|---|---|
| **Step 1** | **configure terminal** | Enters global configuration mode. |
| **Step 2** | **switch(config) # vpc domain <domain>** | Configure the VPC domain number. |
| **Step 3** | **switch(config) # peer-switch** | Define the peer switch. |
| **Step 4** | **switch(config) # show vpc peer-keepalive** | Displays information about the peer keepalive messages |
| **Step 5** | **switch(config) # delay restore { time }** | Number of seconds to delay bringing up the restored vPC peer device. The range is from 1 to 3600. |
| **Step 6** | **switch(config) # peer-gateway** | To enable Layer 3 forwarding for packets destined to the gateway MAC address of the virtual Port Channel (vPC), use the peer-gateway command. To disable Layer 3 forwarding packets, use the no form of this command.<br>. |
| **Step 7** | **switch(config) # delay restore orphan-port** | Number of seconds to delay bringing up the restored device's orphan port |
| **Step 8** | **switch(config-vpc-domain)# fast-convergence** | Configure vPC fast convergence. |

# Configuring LACP vPC Convergence

Beginning with Cisco NX-OS Release 7.0(3)I7(1), Link Aggregation Control Protocol (LACP) vPC convergence feature is supported on Cisco Nexus 9200 and 9300 Series Switches. You can configure LACP vPC convergence feature for more efficient use of port channels by reducing convergence time of vPC port channel for member link going down and first member bring up.

Beginning with Cisco NX-OS Release 7.0(3)I7(5), Link Aggregation Control Protocol (LACP) vPC convergence feature is supported on Cisco Nexus 9500 Series Switches with 9700-EX and 9700-FX line cards. This feature is not supported on Nexus 9500 with 9400, 9500, and 9600 and 9600-R line cards.

Link Aggregation Control Protocol (LACP) vPC convergence feature is supported on Cisco Nexus 9500 Series Switches with 9700-EX and 9700-FX line cards. This feature is not supported on Nexus 9500 with 9400, 9500, and 9600 and 9600-R line cards.

Link Aggregation Control Protocol (LACP) vPC convergence feature is supported on Cisco Nexus 9200 and 9300 Series Switches. You can configure LACP vPC convergence feature for more efficient use of port channels by reducing convergence time of vPC port channel for member link going down and first member bring up.

When you configure LACP vPC convergence on a Cisco Nexus 9000 switch, it waits until all the VLANs are initialized and programmed and then send LACP sync PDU, which will start sending traffic to the VPC domain without drops. You may configure the **lacp vpc-convergence** command in a VXLAN and non-VXLAN environments that have vPC port-channels to hosts that support LACP.

## Procedure

| | Command or Action | Purpose |
|---|---|---|
| **Step 1** | **configure terminal** | Enters global configuration mode. |
| **Step 2** | **switch(config) # interface** {*type/slot* \| **portchannel** *number*} | Specifies an interface to configure, and enters interface configuration mode. |
| **Step 3** | **switch(config-if) # lacp vpc-convergence** | Configure LACP convergence. Reduce the convergence time of the vPC port channel for member link going down and first member bring up.<br>**Note**<br>You must enable this command on both the vPC peer switches. This command must be configured only on PortFast ports (vPC port channels on which the spanning-tree port type edge [trunk] is enabled).<br>**Note**<br>In a vPC environment, when this command is not configured on vPC port-channel interfaces to devices that support LACP, and if one of the vPC peers is reloaded or if one of the links is brought up, the link(s) connected to the vPC peer switch in the 'up' state will remain active and forward traffic. The other link(s) may go down and will transition to the 'up' state. The links transitioning to the 'up' state starts initializing VLANs. When the VLANs are initialized, LACP sync PDUs are sent per initialized VLAN, which will bring port-channel into 'up' state and that leads to traffic black-hole for non-idealized VLANs. |

# Configuring a Graceful Consistency Check

You can configure the graceful consistency check feature, which is enabled by default. Unless this feature is enabled, the vPC is completely suspended when a mismatch in a mandatory compatibility parameter is introduced in a working vPC. When this feature is enabled, only the links on the secondary peer device are suspended. See the "Compatibility Parameters for vPC Interfaces" section for information about consistent configurations on the vPCs.

### Procedure

| | Command or Action | Purpose |
|---|---|---|
| **Step 1** | **configure terminal**<br>**Example:**<br>`switch# configure terminal`<br>`switch(config)#` | Enters global configuration mode. |
| **Step 2** | **vpc domain** *domain-id* [**shut** \| **no shut**]<br>**Example:**<br>`switch(config-if)# vpc domain 5`<br>`switch(config-vpc-domain)#` | Creates a vPC domain if it does not already exist, and enters vpc-domain configuration mode. |
| **Step 3** | **graceful consistency-check**<br>**Example:**<br>`switch(config-vpc-domain)# graceful consistency-check` | Specifies that only the links on the secondary peer device are suspended when a mismatch is detected in a mandatory compatibility parameter.<br>Use the **no** form of this command to disable the feature. |
| **Step 4** | **exit**<br>**Example:**<br>`switch(config)# exit`<br>`switch#` | Exits vpc-domain configuration mode. |
| **Step 5** | **show vpc brief**<br>**Example:**<br>`switch# show vpc brief` | (Optional) Displays information on the vPCs. |

### Example

This example shows how to enable the graceful consistency check feature:

```
switch# configure terminal
switch(config)# vpc domain 5
switch(config-vpc-domain)# graceful consistency-check
switch(config-vpc-domain)# exit
switch(config)#
```

# Checking the Configuration Compatibility on a vPC Peer-Link

After you have configured the vPC Peer-Link on both vPC peer devices, check that the configurations are consistent on all vPC interfaces. See the "Compatibility Parameters for vPC Interfaces" section for information about consistent configurations on the vPCs.

### Procedure

| | Command or Action | Purpose |
|---|---|---|
| **Step 1** | **configure terminal**<br>**Example:**<br>`switch# configure terminal`<br>`switch(config)#` | Enters global configuration mode. |
| **Step 2** | **show vpc consistency-parameters** {**global** \| **interface port-channel** *channel-number*}<br>**Example:**<br>`switch(config)# show vpc consistency-parameters global`<br>`switch(config)#` | (Optional) Displays the status of those parameters that must be consistent across all vPC interfaces. |

### Example

This example shows how to check that the required configurations are compatible across all the vPC interfaces:

```
switch# configure terminal
switch(config)# show vpc consistency-parameters global
switch(config)#
```



**Note** Messages regarding the vPC interface configuration compatibility are also logged to the syslog.

# Moving Other Port Channels into a vPC

We recommend that you attach the vPC domain downstream port channel to two devices for redundancy.

To connect to the downstream device, you create a port channel from the downstream device to the primary vPC peer device and you create another port channel from the downstream device to the secondary peer device. On each vPC peer device, you assign a vPC number to the port channel that connects to the downstream device. You will experience minimal traffic disruption when you are creating vPCs.

### Before you begin

Ensure that you have enabled the vPC feature.

Ensure that you are using a Layer 2 port channel.

### Procedure

| | Command or Action | Purpose |
|---|---|---|
| **Step 1** | **configure terminal**<br>**Example:**<br>`switch# configure terminal`<br>`switch(config)#` | Enters global configuration mode. |
| **Step 2** | **interface port-channel** *channel-number*<br>**Example:**<br>`switch(config)# interface port-channel 20`<br>`switch(config-if)#` | Selects the port channel that you want to put into the vPC to connect to the downstream device, and enters interface configuration mode. |
| **Step 3** | **vpc** *number*<br>**Example:**<br>`switch(config-if)# vpc 5`<br>`switch(config-vpc-domain)#` | Configures the selected port channel into the vPC to connect to the downstream device. You can use any module in the device for these port channels. The range is from 1 and 4096.<br>**Note**<br>The vPC number that you assign to the port channel connecting to the downstream device from the vPC peer device must be identical on both vPC peer devices. |
| **Step 4** | **exit**<br>**Example:**<br>`switch(config)# exit`<br>`switch#` | Exits vpc-domain configuration mode. |
| **Step 5** | **show vpc brief**<br>**Example:**<br>`switch# show vpc brief` | (Optional) Displays information on the vPCs. |
| **Step 6** | **copy running-config startup-config**<br>**Example:**<br>`switch# copy running-config startup-config` | (Optional) Copies the running configuration to the startup configuration. |

### Example

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

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

# Manually Configuring a vPC Domain MAC Address

When you create a vPC domain, the Cisco NX-OS software automatically creates a vPC system MAC address, which is used for operations that are confined to the link-scope, such as LACP. However, you might choose to configure the vPC domain MAC address manually.

### Before you begin

Ensure that you have enabled the vPC feature.

### Procedure

| | Command or Action | Purpose |
|---|---|---|
| **Step 1** | **configure terminal**<br>**Example:**<br>`switch# configure terminal`<br>`switch(config)#` | Enters global configuration mode. |
| **Step 2** | **vpc domain** *domain-id* [**shut** \| **no shut**]<br>**Example:**<br>`switch(config)# vpc domain 5`<br>`switch(config-vpc-domain)#` | Enters the vPC domain number that you want to configure. The system enters vpc-domain configuration mode. |
| **Step 3** | **system-mac** *mac-address*<br>**Example:**<br>`switch(config-vpc-domain)# system-mac 23fb.4ab5.4c4e`<br>`switch(config-vpc-domain)#` | Enters the MAC address that you want for the specified vPC domain in the following format: aaaa.bbbb.cccc. |
| **Step 4** | **exit**<br>**Example:**<br>`switch(config-vpc-domain)# exit`<br>`switch#` | Exits vpc-domain configuration mode. |
| **Step 5** | **show vpc role**<br>**Example:**<br>`switch# show vpc brief` | (Optional) Displays the vPC system MAC address. |
| **Step 6** | **copy running-config startup-config**<br>**Example:**<br>`switch# copy running-config startup-config` | (Optional) Copies the running configuration to the startup configuration. |

### Example

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

```
switch# configure terminal
switch(config)# vpc domain 5
switch(config-vpc-domain)# system-mac 13gb.4ab5.4c4e
switch(config-vpc-domain)# exit
switch(config)#
```

# Manually Configuring the System Priority

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



**Note** We recommend that you manually configure the vPC system priority when you are running LACP to ensure that the vPC peer devices are the primary devices on LACP. When you manually configure the system priority, ensure that you configure the same priority value on both vPC peer devices. If these values do not match, vPC does not come up.

### Before you begin

Ensure that you have enabled the vPC feature.

### Procedure

| | Command or Action | Purpose |
|---|---|---|
| **Step 1** | **configure terminal**<br>**Example:**<br>`switch# configure terminal`<br>`switch(config)#` | Enters global configuration mode. |
| **Step 2** | **vpc domain** *domain-id* [**shut** \| **no shut**]<br>**Example:**<br>`switch(config)# vpc domain 5`<br>`switch(config-vpc-domain)#` | Enters the vPC domain number that you want to configure. The system enters vpc-domain configuration mode. |
| **Step 3** | **system-priority** *priority*<br>**Example:**<br>`switch(config-vpc-domain)#`<br>`system-priority 4000`<br>`switch(config-vpc-domain)#` | Enters the system priority that you want for the specified vPC domain. The range of values is from 1 to 65535. The default value is 32667. |
| **Step 4** | **exit**<br>**Example:**<br>`switch(config-vpc-domain)# exit`<br>`switch#` | Exits vpc-domain configuration mode. |
| **Step 5** | **show vpc role**<br>**Example:** | (Optional) Displays the vPC system priority. |

| | Command or Action | Purpose |
|---|---|---|
| | `switch# show vpc role` | |
| **Step 6** | **copy running-config startup-config**<br>**Example:**<br>`switch# copy running-config startup-config` | (Optional) Copies the running configuration to the startup configuration. |

### Example

This example shows how to manually configure the vPC domain system priority:

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

# Manually Configuring the vPC Peer Device Role

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

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

### Before you begin

Ensure that you have enabled the vPC feature.

### Procedure

| | Command or Action | Purpose |
|---|---|---|
| **Step 1** | **configure terminal**<br>**Example:**<br>`switch# configure terminal`<br>`switch(config)#` | Enters global configuration mode. |
| **Step 2** | **vpc domain** *domain-id* [**shut** \| **no shut**]<br>**Example:**<br>`switch(config)# vpc domain 5`<br>`switch(config-vpc-domain)#` | Enters the vPC domain number that you want to configure. The system enters vpc-domain configuration mode. |
| **Step 3** | **role priority** *priority*<br>**Example:** | Enters the role priority that you want for the vPC system priority. The range of values is from 1 to 65636, and the default value is 32667. |

| | Command or Action | Purpose |
|---|---|---|
| | `switch(config-vpc-domain)# role priority 4`<br>`switch(config-vpc-domain)#` | A lower value means that this switch has a better chance of being the primary vPC. |
| **Step 4** | **exit**<br>**Example:**<br>`switch(config)# exit`<br>`switch#` | Exits vpc-domain configuration mode. |
| **Step 5** | **show vpc role**<br>**Example:**<br>`switch# show vpc role` | (Optional) Displays the vPC system priority. |
| **Step 6** | **copy running-config startup-config**<br>**Example:**<br>`switch# copy running-config startup-config` | (Optional) Copies the running configuration to the startup configuration. |

### Example

This example shows how to manually configure the role priority of the vPC peer device:

```
switch# configure terminal
switch(config)# vpc domain 5
switch(config-vpc-domain)# role priority 4
switch(config-vpc-domain)# exit
switch(config)#
```

# Configuring for Recovery After an Outage

If an outage occurs, the vPC waits for a peer adjacency to form on a switch reload. This situation can result in an unacceptably long service disruption. You can configure the Cisco Nexus 3550-T Series device to restore vPC services when its peer fails to come on line.

## Configuring Reload Restore

The **reload restore** command and procedure described in this section is deprecated. We recommend that you use the **auto-recovery** command and procedure described in the “Configuring an Autorecovery” section.

You can configure the Cisco Nexus 3550-T device to restore vPC services when its peer fails to come online by using the **reload restore** command.

### Before you begin

Ensure that you have enabled the vPC feature.

### Procedure

| | Command or Action | Purpose |
|---|---|---|
| **Step 1** | **configure terminal**<br>**Example:**<br>`switch# configure terminal`<br>`switch(config)#` | Enters global configuration mode. |
| **Step 2** | **vpc domain** *domain-id* [**shut** \| **no shut**]<br>**Example:**<br>`switch(config)# vpc domain 5`<br>`switch(config-vpc-domain)#` | Enters the vPC domain number that you want to configure, and enters vpc-domain configuration mode. |
| **Step 3** | **reload restore** [**delay** *time-out*]<br>**Example:**<br>`switch(config-vpc-domain)# reload restore` | Configures the vPC to assume its peer is not functional and to bring up the vPC. The default delay is 240 seconds. You can configure a time-out delay from 240 to 3600 seconds.<br>Use the **no** form of the command to reset the vPC to its default settings. |
| **Step 4** | **exit**<br>**Example:**<br>`switch(config-vpc-domain)# exit`<br>`switch#` | Exits vpc-domain configuration mode. |
| **Step 5** | **show running-config vpc**<br>**Example:**<br>`switch# show running-config vpc` | (Optional) Displays information about the vPC, specifically the reload status. |
| **Step 6** | **show vpc consistency-parameters interface port-channel** *number*<br>**Example:**<br>`switch# show vpc consistency-parameters interface port-channel 1` | (Optional) Displays information about the vPC consistency parameters for the specified interface. |
| **Step 7** | **copy running-config startup-config**<br>**Example:**<br>`switch# copy running-config startup-config` | (Optional) Copies the running configuration to the startup configuration.<br>**Note**<br>To ensure the reload feature is enabled, you should perform this step. |

### Example

This example shows how to set the vPC reload restore feature and save it in the switch startup configuration:

```
switch# configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
switch(config)# vpc domain 5
switch(config-vpc-domain)# reload restore

Warning:
Enables restoring of vPCs in a peer-detached state after reload, will wait for 240
seconds (by default) to determine if peer is un-reachable

switch(config-vpc-domain)# exit
switch(config)# exit
switch# copy running-config startup-config
switch# show running-config vpc

!Command: show running-config vpc
!Time: Wed Mar 24 18:43:54 2010
version 5.0(2)
feature vpc
logging level vpc 6
vpc domain 5
reload restore
```

This example shows how to examine the consistency parameters:

```
switch# show vpc consistency-parameters interface port-channel 1

Legend:
Type 1 : vPC will be suspended in case of mismatch
Name Type Local Value Peer Value
------------- ---- ----------- ---------------
STP Port Type 1 Default -
STP Port Guard 1 None -
STP MST Simulate PVST 1 Default -
mode 1 on -
Speed 1 1000 Mb/s -
Duplex 1 full -
Port Mode 1 trunk -
Native Vlan 1 1 -
MTU 1 1500 -
Allowed VLANs - 1-3967,4048-4093
Local suspended VLANs
```

## Configuring an Autorecovery

You can configure the Cisco Nexus 3550-T Series device to restore vPC services when its peer fails to come online by using the auto-recovery command.

You can configure the Cisco Nexus 3550-T Series device to restore vPC services on the secondary vPC peer when its vPC primary peer fails and bringing down peer-keepalive and vPC Peer-Link, by using the **auto-recovery** command. In case of failure of primary switch where both peer-keepalive and vPC Peer-Links are down secondary switch will suspend vPC member. However, after 3 missed keepalive heartbeats secondary switch resumes the role of a primary switch and bring up vPC member ports. The **auto-recovery reload restore** command can be used in scenarios when vPC primary switch reloads, where secondary switch resumes the role of the vPC primary and bring ip VPC member ports.

### Before you begin

Ensure that you have enabled the vPC feature.

### Procedure

| | Command or Action | Purpose |
|---|---|---|
| **Step 1** | **configure terminal**<br>**Example:**<br>`switch# configure terminal`<br>`switch(config)#` | Enters global configuration mode. |
| **Step 2** | **vpc domain** *domain-id* [**shut** \| **no shut**]<br>**Example:**<br>`switch(config)# vpc domain 5`<br>`switch(config-vpc-domain)#` | Enters the vPC domain number that you want to configure, and enters vpc-domain configuration mode. |
| **Step 3** | **auto-recovery** [**reload-delay** *time*]<br>**Example:**<br>`switch(config-vpc-domain)# auto-recovery` | Configures the vPC to assume its peer is not functional and to bring up the vPC, and specifies the time to wait after a reload to restore the vPC. The default delay is 240 seconds. You can configure a delay from 240 to 3600 seconds.<br>Use the **no** form of the command to reset the vPC to its default settings. |
| **Step 4** | **exit**<br>**Example:**<br>`switch(config-vpc-domain)# exit`<br>`switch#` | Exits vpc-domain configuration mode. |
| **Step 5** | **show running-config vpc**<br>**Example:**<br>`switch# show running-config vpc` | (Optional) Displays information about the vPC, specifically the reload status. |
| **Step 6** | **show vpc consistency-parameters interface port-channel** *number*<br>**Example:**<br>`switch# show vpc consistency-parameters interface port-channel 1` | (Optional) Displays information about the vPC consistency parameters for the specified interface. |
| **Step 7** | **copy running-config startup-config**<br>**Example:**<br>`switch# copy running-config startup-config` | (Optional) Copies the running configuration to the startup configuration.<br>**Note**<br>To ensure the autorecovery feature is enabled, you should perform this step. |

### Example

This example shows how to set the vPC autorecovery feature and save it in the switch startup configuration:

```
switch# configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
switch(config)# vpc domain 5
switch(config-vpc-domain)# auto-recovery
switch(config-vpc-domain)# auto-recovery auto-recovery reload-delay 100

Warning:
Enables restoring of vPCs in a peer-detached state after reload, will wait for 240
seconds to determine if peer is un-reachable

switch(config-vpc-domain)# exit
switch(config)# exit
switch# copy running-config startup-config
```

# Configuring the Suspension of Orphan Ports

When a device that is not vPC-capable connects to each peer, the connected ports are known as orphan ports because they are not members of a vPC. You can explicitly declare physical interfaces as orphan ports to be suspended (shut down) by the secondary peer when it suspends its vPC ports in response to a vPC Peer-Link or peer-keepalive failure. The orphan ports are restored when the vPC is restored.



**Note** You can configure vPC orphan port suspension only on physical ports, portchannels. However, you cannot configure the same on individual port channel member ports.

### Before you begin

Ensure that you have enabled the vPC feature.

### Procedure

| | Command or Action | Purpose |
|---|---|---|
| **Step 1** | **configure terminal**<br>**Example:**<br>`switch# configure terminal`<br>`switch(config)#` | Enters global configuration mode. |
| **Step 2** | **show vpc orphan-ports**<br>**Example:**<br>`switch# show vpc orphan-ports` | (Optional) Displays a list of the orphan ports. |
| **Step 3** | **interface** *type slot/port*<br>**Example:**<br>`switch(config)# interface ethernet 1/3`<br>`switch(config-if)#` | Specifies an interface to configure, and enters interface configuration mode. |

| | Command or Action | Purpose |
|---|---|---|
| **Step 4** | **vpc orphan-port suspend**<br>**Example:**<br>`switch(config-if)# vpc orphan-ports suspend` | Configures the selected interface as a vPC orphan port to be suspended by the secondary peer in the case of a vPC failure. |
| **Step 5** | **exit**<br>**Example:**<br>`switch(config-if)# exit`<br>`switch#` | Exits interface configuration mode. |
| **Step 6** | **copy running-config startup-config**<br>**Example:**<br>`switch# copy running-config startup-config` | (Optional) Copies the running configuration to the startup configuration. |

### Example

This example shows how to configure an interface as a vPC orphan port to be suspended by the secondary peer in the case of a vPC failure:

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

# Configuring Delay Restore on an Orphan Port

You can configure **delay restore orphan-port** command on Cisco Nexus 9000 Series switches to configure a restore timer that delays the bringing up of restored device's orphan port starting from Cisco NX-OS Release 7.0(3)I7(1).

You can configure **delay restore orphan-port** command on Cisco Nexus 9000 Series switches to configure a restore timer that delays the bringing up of restored device's orphan port.



**Note** The delay restore orphan-port command applies only to interfaces that has vpc orphan-port suspend command configured. Other orphan ports may not delay bringing up devices.

### Procedure

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

| | Command or Action | Purpose |
|---|---|---|
| **Step 3** | **switch(config) # peer-switch** | Define the peer switch. |
| **Step 4** | **switch(config) # show vpc peer-keepalive** | Displays information about the peer keepalive messages |
| **Step 5** | **switch(config) # delay restore { time }** | Number of seconds to delay bringing up the restored vPC peer device. The range is from 1 to 3600. |
| **Step 6** | **switch(config) # peer-gateway** | To enable Layer 3 forwarding for packets destined to the gateway MAC address of the virtual Port Channel (vPC), use the peer-gateway command. To disable Layer 3 forwarding packets, use the no form of this command.<br>. |
| **Step 7** | **switch(config) # delay restore orphan-port** | Number of seconds to delay bringing up the restored device's orphan port |

# Configuring the vPC Peer Switch

You can configure the Cisco Nexus 3550-T Series device to make a pair of vPC devices appear as a single STP root in the Layer 2 topology.

## Configuring a Pure vPC Peer Switch Topology

You can configure a pure vPC peer switch topology by using the peer-switch command and then setting the best possible (lowest) spanning tree bridge priority value.

### Before you begin

Ensure that you have enabled the vPC feature.

**Note** When using a non-VPC dedicated trunk link between the VPC peers, the non-VPC VLANs should have a different global priority on the peers to prevent STP from blocking the VLANs.

### Procedure

| | Command or Action | Purpose |
|---|---|---|
| **Step 1** | **configure terminal**<br>**Example:**<br>`switch# configure terminal`<br>`switch(config)#` | Enters global configuration mode. |

| | Command or Action | Purpose |
|---|---|---|
| **Step 2** | **vpc domain** *domain-id* [**shut** \| **no shut**]<br>**Example:**<br>`switch(config)# vpc domain 5`<br>`switch(config-vpc-domain)#` | Enters the vPC domain number that you want to configure, and enters vpc-domain configuration mode. |
| **Step 3** | **peer-switch**<br>**Example:**<br>`switch(config-vpc-domain)# peer-switch` | Enables the vPC switch pair to appear as a single STP root in the Layer 2 topology.<br>Use the **no** form of the command to disable the peer switch vPC topology. |
| **Step 4** | **spanning-tree vlan** *vlan-range* **priority** *value*<br>**Example:**<br>`switch(config)# spanning-tree vlan 1 priority 8192` | Configures the bridge priority of the VLAN. Valid values are multiples of 4096. The default value is 32768. |
| **Step 5** | **exit**<br>**Example:**<br>`switch(config-vpc-domain)# exit`<br>`switch#` | Exits vpc-domain configuration mode. |
| **Step 6** | **show spanning-tree summary**<br>**Example:**<br>`switch# show spanning-tree summary` | (Optional) Displays a summary of the spanning tree port states including the vPC peer switch. |
| **Step 7** | **copy running-config startup-config**<br>**Example:**<br>`switch# copy running-config startup-config` | (Optional) Copies the running configuration to the startup configuration. |

### Example

This example shows how to configure a pure vPC peer switch topology:

```
switch# configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
switch(config)# vpc domain 5
switch(config-vpc-domain)# peer-switch

2010 Apr 28 14:44:44 switch %STP-2-VPC_PEERSWITCH_CONFIG_ENABLED: vPC peer-switch
configuration is enabled. Please make sure to configure spanning tree "bridge" priority as
per recommended guidelines to make vPC peer-switch operational.

switch(config-vpc-domain)# spanning-tree vlan 1 priority 8192
switch(config-vpc-domain)# exit
switch(config)#
```

# Configuring Hitless vPC Role Change

Complete these steps to enable hitless vPC role change.

### Before you begin

- • Ensure that the vPC feature is enabled.
- • Ensure that the vPC Peer-Link is up
- • Verify the role priority of devices

### Procedure

| | Command or Action | Purpose |
|---|---|---|
| **Step 1** | **vpc role preempt**<br>**Example:**<br>`switch# vpc role preempt`<br>`switch(config)#` | Enable hitless vPC role change. |
| **Step 2** | **show vpc role**<br>**Example:**<br>`switch(config)# show vpc role` | (Optional) Verify hitless vPC role change feature. |

### Example

This example on how to configure hitless vPC role change:

```
switch# show vpc rolevPC Role status
----------------------------------------------------
vPC role                        : secondary
vPC system-mac                  : 00:23:04:ee:be:01
vPC system-priority             : 32667
vPC local system-mac            : 8c:60:4f:03:84:41
vPC local role-priority         : 32668
vPC peer system-mac             : 8c:60:4f:03:84:43
vPC peer role-priority          : 32667

! Configure vPC hitless role change on the device!

switch(config)# vpc role preempt
! The following is an output from the show vpc role command after the
vPC hitless feature is configured
switch(config)# show vpc role
vPC Role status
----------------------------------------------------
vPC role                        : primary
vPC system-mac                  : 00:00:00:00:00:00
vPC system-priority             : 32667
vPC local system-mac            : 8c:60:4f:03:84:41
vPC local role-priority         : 32666
vPC peer system-mac             : 8c:60:4f:03:84:43
vPC peer role-priority          : 32667
```

```
switch(config)#
```

## Use Case Scenario for vPC Role Change

The hitless vPC role change feature can be used in the following scenarios:

- Role change request—When you want to change the roles of the peer devices in a vPC domain.
- Primary switch reload—When the devices comes up after a reload and roles are defined, you can use the hitless vPC role change feature to restore the roles. For example, after a reload if the primary device takes the role of operational secondary and the secondary device takes the role of primary operational, you can change the vPC peer roles to their original defined roles using the **vpc role preempt** command.



**Note** Always check the existing device role priority before switching vPC role.

- Dual-active recovery—In a dual-active recovery scenario, the vPC primary switch continues to be (operational) primary, but the vPC secondary switch becomes the targeted primary switch and keeps its vPC member ports up. You can use the vPC hitless feature and restore the device roles. After the Dual-active recovery, if one side is operational primary and the other side operational secondary, then you can use the **vpc role preempt** command to restore the device roles to be primary and secondary

# Verifying the vPC Configuration

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

| Command | Purpose |
|---|---|
| **show feature** | Displays whether the vPC is enabled or not. |
| **show vpc brief** | Displays brief information about the vPCs. |
| **show vpc consistency-parameters** | Displays the status of those parameters that must be consistent across all vPC interfaces. |
| **show running-config vpc** | Displays running configuration information for vPCs. |
| **show port-channel capacity** | Displays how many port channels are configured and how many are still available on the device. |
| **show vpc statistics** | Displays statistics about the vPCs. |
| **show vpc peer-keepalive** | Displays information about the peer-keepalive messages. |
| **show vpc role** | Displays the peer status, the role of the local device, the vPC system MAC address and system priority, and the MAC address and priority for the local vPC device. |

# Monitoring vPCs

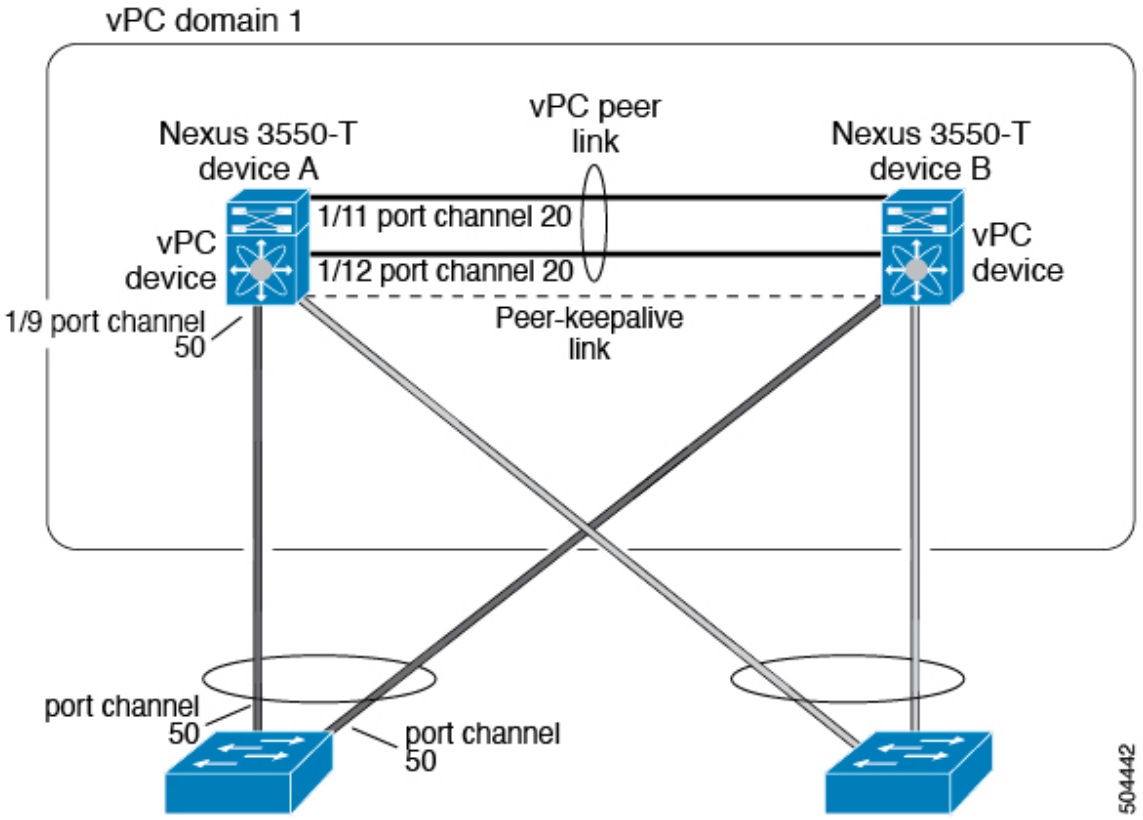
Use the **show vpc statistics**command to display vPC statistics.

**Note** This command displays the vPC statistics only for the vPC peer device that you are working on.

# Configuration Examples for vPCs

The following example shows how to configure vPC on device A as shown in the figure:

***Figure 20: vPC Configuration Example***




1. Enable vPC and LACP.

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

2. (Optional) Configure one of the interfaces that you want to be a vPC Peer-Link in the dedicated port mode.

```
switch(config)# interface ethernet 1/7,
ethernet 1/3, ethernet 1/5. ethernet 1/7
```

```
switch(config-if)# shutdown
switch(config-if)# exit
switch(config)# interface ethernet 1/7

switch(config-if)# no shutdown
switch(config-if)# exit
switch(config)#
```

3. (Optional) Configure the second, redundant interface that you want to be a vPC Peer-Link in the dedicated port mode.

```
switch(config)# interface ethernet 1/2, ethernet 1/4,
ethernet 1/6. ethernet 1/8
switch(config-if)# shutdown
switch(config-if)# exit
switch(config)# interface ethernet 1/2
switch(config-if)# no shutdown
switch(config-if)# exit
switch(config)#
```

4. Configure the two interfaces (for redundancy) that you want to be in the vPC Peer-Link to be an active Layer 2 LACP port channel.

```
switch(config)# interface ethernet 1/1-2
switch(config-if)# switchport
switch(config-if)# switchport mode trunk
switch(config-if)# switchport trunk allowed vlan 1-50
switch(config-if)# switchport trunk native vlan 20
switch(config-if)# channel-group 20 mode active
switch(config-if)# exit
```

5. Create and enable the VLANs.

```
switch(config)# vlan 1-50
switch(config-vlan)# no shutdown
switch(config-vlan)# exit
```

6. Create a separate VRF for the vPC peer-keepalive link and add a Layer 3 interface to that VRF.

```
switch(config)# vrf context pkal
switch(config-vrf)# exit
switch(config)# interface ethernet 1/20
switch(config-if)# vrf member pkal
switch(config-if)# ip address 172.23.145.218/24
switch(config-if)# no shutdown
switch(config-if)# exit
```

7. Create the vPC domain and add the vPC peer-keepalive link.

```
switch(config)# vpc domain 1
switch(config-vpc-domain)# peer-keepalive
 destination 172.23.145.217 source 172.23.145.218 vrf pkal
switch(config-vpc-domain)# exit
```

8. Configure the vPC vPC Peer-Link.

```
switch(config)# interface port-channel 20
switch(config-if)# switchport mode trunk
switch(config-if)# switchport trunk allowed vlan 1-50
switch(config-if)# vpc peer-link
switch(config-if)# exit
switch(config)#
```

9. Configure the interface for the port channel to the downstream device of the vPC.

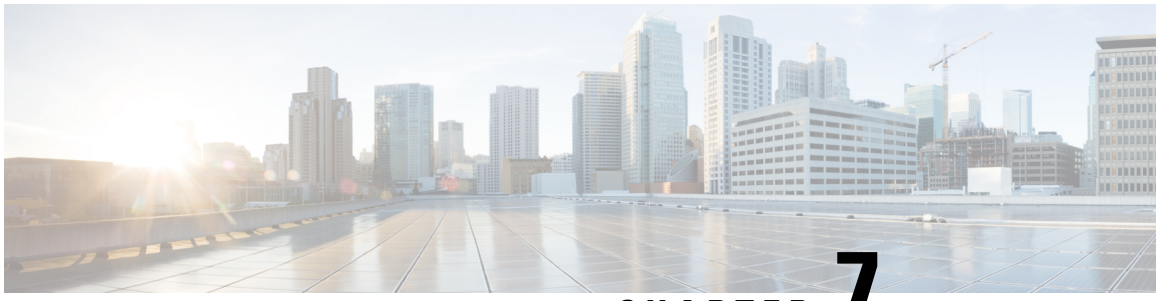
```
switch(config)# interface ethernet 1/9
switch(config-if)# switchport mode trunk
switch(config-if)# allowed vlan 1-50
switch(config-if)# native vlan 20
switch(config-if)# channel-group 50 mode active
switch(config-if)# exit
switch(config)# interface port-channel 50
switch(config-if)# vpc 50
switch(config-if)# exit
switch(config)#
```

**10.** Save the configuration.

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



**Note** If you configure the port channel first, ensure that it is a Layer 2 port channel.



CHAPTER 7

# Configuring Unidirectional Link Detection

This chapter contains the following sections:


- Unidirectional Link Detection, on page 139
- Configuring the UDLD Mode, on page 140


## Unidirectional Link Detection

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

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.

The Cisco Nexus 3550-T device periodically transmits UDLD frames to neighbor devices on LAN ports with UDLD enabled. If the frames are echoed back within a specific time frame and they lack a specific acknowledgment (echo), the link is flagged as unidirectional and the LAN port is shut down. Devices on both ends of the link must support UDLD in order for the protocol to successfully identify and disable unidirectional links. You can configure the transmission interval for the UDLD frames, either globally or for the specified interfaces.

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

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

***Figure 21: Unidirectional Link***



The following table shows the default UDLD configuration.

*Table 9: UDLD Default Configuration*

| Feature | Default Value |
|---|---|
| UDLD global enable state | Globally disabled |
| UDLD per-port enable state for fiber-optic media | Enabled on all Ethernet fiber-optic LAN ports |
| UDLD per-port enable state for twisted-pair (copper) media | Disabled on all 10G Ethernet ports |
| UDLD aggressive mode | Disabled |
| UDLD message interval | 15 seconds |

# UDLD Modes

UDLD can operate in two modes, namely, aggressive mode and non-aggressive mode.

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 frame, UDLD tries to reestablish the connection with the neighbor. After eight failed retries, the port is disabled.

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.

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

# Configuring the UDLD Mode

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

Before you can enable 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.

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

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

*Table 10: CLI Details to Enable or Disable UDLD on Different Interfaces*

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

### Before you begin

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

## Procedure

| | Command or Action | Purpose |
|---|---|---|
| **Step 1** | **configure terminal**<br>**Example:**<br>`switch# configure terminal`<br>`switch(config)#` | Enters global configuration mode. |
| **Step 2** | [**no**] **feature udld**<br>**Example:**<br>`switch(config)# feature udld`<br>`switch(config)#`<br>`switch(config)# no feature udld`<br>`switch(config)#` | Enables/Disables UDLD for the device. |
| **Step 3** | **udld message-time** *seconds*<br>**Example:**<br>`switch(config)# udld message-time 30`<br>`switch(config)#` | (Optional) Specifies the interval between sending UDLD messages. The range is from 7 to 90 seconds, and the default is 15 seconds. |
| **Step 4** | **udld aggressive**<br>**Example:**<br>`switch(config)# udld aggressive`<br>`switch(config)#` | Optional) Specifies UDLD mode to be aggressive.<br>**Note**<br>For copper interfaces, you enter the interface command mode for those interfaces you want to configure for UDLD aggressive mode and issue this command in interface command model. |
| **Step 5** | **interface ethernet** *slot/port*<br>**Example:**<br>`switch(config)# interface ethernet 1/1`<br>`switch(config-if)#` | (Optional) Specifies an interface to configure, and enters interface configuration mode. |

| | Command or Action | Purpose |
|---|---|---|
| **Step 6** | **udld** [**enable** \| **disable**]<br>**Example:**<br>`switch(config-if)# udld enable`<br>`switch(config-if)#` | (Optional) Enables UDLD on the specified copper port or disables UDLD on the specified fiber port.<br>To enable UDLD on copper ports, enter the **udld enable** command. To enable UDLD on fiber ports, enter the **no udld disable** command. |
| **Step 7** | **show udld** [**ethernet** *slot/port* \| **global** \| **neighbors**]<br>**Example:**<br>`switch(config)# show udld`<br>`switch(config)#` | (Optional) Displays the UDLD status. |
| **Step 8** | **exit**<br>**Example:**<br>`switch(config-if-range)# exit`<br>`switch(config)#` | Exits the interface mode. |
| **Step 9** | **copy running-config startup-config**<br>**Example:**<br>`switch(config)# copy running-config startup-config` | (Optional) Copies the running configuration to the startup configuration. |

### Example

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

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

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

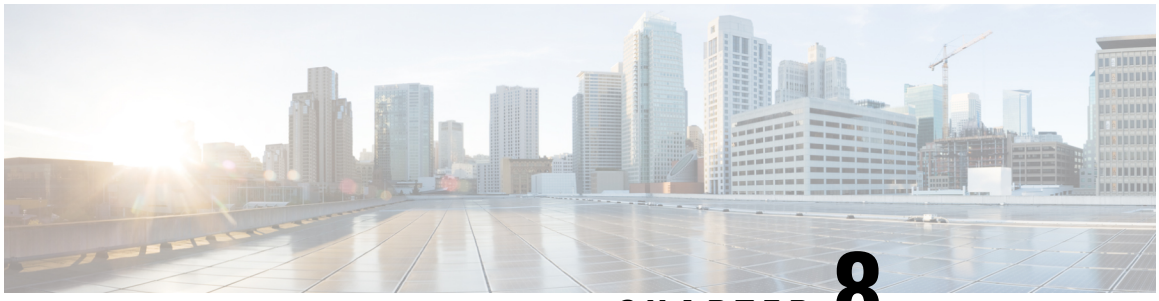
```
switch# configure terminal
switch(config)# feature udld
switch(config)# udld message-time 30
switch(config)#
```

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

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

This example shows how to disable UDLD for the device:

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



CHAPTER 8

# Multicast Fairness Tuning


- Multicast fairness, on page 143
- Guidelines and limitations for multicast fairness tuning, on page 143
- Configure multicast fairness tuning, on page 144
- Verify the Multicast Fairness Tuning Configuration, on page 145


## Multicast fairness

Multicast traffic involves sending data from one source to multiple destinations simultaneously, which can sometimes lead to difference in latency. The multicast fairness tuning feature aims to minimize the latency difference for a multicast stream across different ports.

Beginning with NX-OS Release 10.5(2)F, the multicast fairness tuning feature for Cisco Nexus 3550-T switches allows you to tune the egressing multicast traffic by configuring equalization delay for specific ports. Thus, this feature ensures that the egressing traffic reaches its destinations at almost the same time.

You can tune the latency difference for a multicast stream among the ports by adding delays to the faster ports. However, you need to measure the delay or latency for each port beforehand and be aware of the default latency, only then can you equalize the delay on the faster ports. The deviation is reduced to a negligible difference of less than 250 pico-seconds.

For example, consider a multicast stream is being sent through the interfaces Ethernet 1/2, Ethernet 1/3, and Ethernet 1/4. The timestamp of the multicast stream reveals that it leaves N3550-T from Ethernet 1/2 at 6.85 nano-seconds, from Ethernet 1/3 at 5.70 nano-seconds, and from Ethernet 1/4 at 6.20 nano-seconds. The tuning feature allows you to add around 1000 pico-seconds delay at Ethernet 1/3 and around 600 pico-seconds delay at Ethernet 1/4 to allow each of these streams coming out from these ports within the 250 pico-seconds range.

Beginning with NX-OS Release 10.6(3)F, 25G is supported on 3550-T and multicast fairness tuning feature is supported on 25G 3550-T platform.

## Guidelines and limitations for multicast fairness tuning

Follow these guidelines and limitations while configuring the multicast fairness tuning feature.

- If you send more traffic than the line rate, it leads to congestion and the interface can no longer maintain the fairness. However, as soon as the traffic rate reduces, the fairness is restored.

- If traffic from multiple sources compete to go out of the same interface, then this impacts the latency fairness for that interface.
- If you try to configure latency on a specific port with running traffic, then this results in a short disruption in the traffic in the form of either a drop or corruption.

# Configure multicast fairness tuning

The multicast fairness tuning feature is an interface specific feature, so go to the required interface to configure the feature. Follow these steps to configure the multicast fairness tuning feature.

### Procedure

| | Command or Action | Purpose |
|---|---|---|
| **Step 1** | **configure terminal**<br>**Example:**<br>`switch# configure terminal`<br>`switch(config)#` | Enters global configuration mode. |
| **Step 2** | **interface** *type slot/port*<br>**Example:**<br>`switch(config)# interface ethernet 1/10`<br>`switch(config-if)#` | Specifies an interface to configure and enters the interface configuration mode. |
| **Step 3** | [ **no**] **equalization-delay** *value*<br>**Example:**<br>`switch(config-if)# equalization-delay 10`<br>`switch(config-if)#` | Specify the equalization delay value that you want to configure on the specified interface. The default value is zero (0).<br>The range for equalization delay values is 0 to 15 for both 10G and 25G ports, where 1=100 pico-seconds. Thus, the maximum delay that can be configured is 1500 pico-seconds.<br>If you configure the delay as 4, it is equal to 400 pico-seconds.<br>The **no** form of this command disables equalization delay. |
| **Step 4** | **shut**<br>**Example:**<br>`switch(config-if)# shut`<br>`switch(config)#` | Disables the specified interface. |
| **Step 5** | **no shut**<br>**Example:**<br>`switch(config-if)# no shut`<br>`switch(config)#` | Enables the specified interface. |

| | Command or Action | Purpose |
|---|---|---|
| **Step 6** | **exit**<br>**Example:**<br>`switch(config-if)# exit`<br>`switch(config)#` | Exits the interface configuration mode. |

### Example

This example shows how to configure multicast fairness tuning on a particular interface.

```
switch# configure terminal
switch(config)# interface ethernet 1/10
switch(config-if)# equalization-delay 10
switch(config-if)# shut
switch(config-if)# no shut
switch(config-if)# exit
```

# Verify the Multicast Fairness Tuning Configuration

Perform the relevant show commands listed in the table to display the required information about the multicast fairness tuning configuration.

| Command | Purpose |
|---|---|
| **show interface** *type slot/port* | Displays the interface status and information for the specified interface along with the configured equalization delay in pico-seconds. |
| **show interface** *type slot/port* **equalization-delay** | Displays only the values for equalization delay in pico-seconds for the specified interface. |
| **show interface** *type range of slots/ports* **equalization-delay** | Displays all the values for equalization delay in pico-seconds for the specified range of interfaces. |
| **show running-config** [ **all** ] | Displays information about the current configuration. The **all** option displays the default and current configurations.<br>This command also displays the equalization delay configured for each interface. |

### Sample Outputs for Show Commands

This is a sample output of the **show run interface** *type slot/port* command that displays the equalization delay for the specified interface.

```
show run interface ethernet 1/10
  interface Ethernet1/10
  equalization-delay 10
```

This is a sample output of the **show interface** *type slot/port* command that displays the interface status and information for the specified interface including the information regarding equalization delay for the interface.

### Sample output for 25G

```
N3550T-TOR2# show interface eth1/17

Ethernet1/17 is up

admin state is up, Dedicated Interface

 Hardware: 25000 Ethernet, address: 643f.5f84.489e (bia 643f.5f84.48af)

 MTU 1500 bytes, BW 25000000 Kbit , DLY 10 usec

 reliability 255/255, txload 1/255, rxload 1/255

 Encapsulation ARPA, medium is broadcast

 full-duplex, 25 Gb/s, media type is 25G

 Beacon is turned off

 Auto-Negotiation is turned off FEC mode is Auto

 Input flow-control is off, output flow-control is off

 Auto-mdix is turned off

 Rate mode is dedicated

 Switchport monitor is off

 EtherType is 0x8100

 EEE (efficient-ethernet) : n/a

 admin fec state is auto, oper fec state is auto

Equalization delay 1000 picosec

 Last link flapped 00:03:13

 Last clearing of "show interface" counters never

 2 interface resets

 Load-Interval #1: 30 seconds

 30 seconds input rate 0 bits/sec, 0 packets/sec

 30 seconds output rate 0 bits/sec, 0 packets/sec

 input rate 0 bps, 0 pps; output rate 0 bps, 0 pps

 Load-Interval #2: 5 minute (300 seconds)

 300 seconds input rate 0 bits/sec, 0 packets/sec

 300 seconds output rate 0 bits/sec, 0 packets/sec

 input rate 0 bps, 0 pps; output rate 0 bps, 0 pps
```

```
 RX
 0 unicast packets 0 multicast packets 0 broadcast packets
 0 input packets 0 bytes
 0 jumbo packets 0 storm suppression packets
 0 runts 0 giants 0 CRC 0 no buffer
 0 input error 0 short frame 0 overrun 0 underrun 0 ignored
 0 watchdog 0 bad etype drop 0 bad proto drop 0 if down drop
 0 input with dribble 0 input discard
 0 Rx pause
 0 Stomped CRC
 TX
 0 unicast packets 54 multicast packets 10 broadcast packets
 64 output packets 0 bytes
 0 jumbo packets
 0 output error 0 collision 0 deferred 0 late collision
 0 lost carrier 0 no carrier 0 babble 0 output discard
 0 Tx pause
N3550T-TOR2#
```

### Sample output for 10G

```
N3550T-TOR1# show interface eth1/19
Ethernet1/19 is up
admin state is up, Dedicated Interface
 Hardware: 1000/10000 Ethernet, address: 643f.5f84.3c5a (bia 643f.5f84.3c6d)
 MTU 1500 bytes, BW 10000000 Kbit , DLY 10 usec
 reliability 255/255, txload 1/255, rxload 1/255
 Encapsulation ARPA, medium is broadcast
 full-duplex, 10 Gb/s, media type is 10G
 Beacon is turned off
 Auto-Negotiation is turned on FEC mode is Auto
 Input flow-control is off, output flow-control is off
 Auto-mdix is turned off
 Rate mode is dedicated
```

```
 Switchport monitor is off
 EtherType is 0x8100
 EEE (efficient-ethernet) : n/a
 admin fec state is auto, oper fec state is auto
Equalization delay 1000 picosec
 Last link flapped 17:02:49
 Last clearing of "show interface" counters 02:57:48
 0 interface resets
 Load-Interval #1: 30 seconds
 30 seconds input rate 0 bits/sec, 0 packets/sec
 30 seconds output rate 0 bits/sec, 0 packets/sec
 input rate 0 bps, 0 pps; output rate 0 bps, 0 pps
 Load-Interval #2: 5 minute (300 seconds)
 300 seconds input rate 0 bits/sec, 0 packets/sec
 300 seconds output rate 0 bits/sec, 0 packets/sec
 input rate 0 bps, 0 pps; output rate 0 bps, 0 pps
 RX
 0 unicast packets 0 multicast packets 0 broadcast packets
 0 input packets 0 bytes
 0 jumbo packets 0 storm suppression packets
 0 runts 0 giants 0 CRC 0 no buffer
 0 input error 0 short frame 0 overrun 0 underrun 0 ignored
 0 watchdog 0 bad etype drop 0 bad proto drop 0 if down drop
 0 input with dribble 0 input discard
 0 Rx pause
 0 Stomped CRC
 TX
 0 unicast packets 2133 multicast packets 0 broadcast packets
 2133 output packets 0 bytes
 0 jumbo packets
 0 output error 0 collision 0 deferred 0 late collision
 0 lost carrier 0 no carrier 0 babble 0 output discard
```

```
 0 Tx pause


N3550T-TOR1#
```

### Sample output for 1G

```
N3550T-TOR1# show interface eth1/1

Ethernet1/1 is up

admin state is up, Dedicated Interface

 Hardware: 1000/10000 Ethernet, address: 643f.5f84.3c5b (bia 643f.5f84.3c5b)

 MTU 1500 bytes, BW 1000000 Kbit , DLY 10 usec

 reliability 255/255, txload 1/255, rxload 1/255

 Encapsulation ARPA, medium is broadcast

 Port mode is access

 full-duplex, 1000 Mb/s, media type is 1G

 Beacon is turned off

 Auto-Negotiation is turned on FEC mode is Auto

 Input flow-control is off, output flow-control is off

 Auto-mdix is turned off

 Rate mode is dedicated

 Switchport monitor is off

 EtherType is 0x8100

 EEE (efficient-ethernet) : n/a

 admin fec state is auto, oper fec state is auto

Equalization delay 1000 picosec

 Last link flapped 03:37:28

 Last clearing of "show interface" counters 02:58:31

 0 interface resets

 Load-Interval #1: 30 seconds

 30 seconds input rate 0 bits/sec, 0 packets/sec

 30 seconds output rate 0 bits/sec, 0 packets/sec

 input rate 0 bps, 0 pps; output rate 0 bps, 0 pps

 Load-Interval #2: 5 minute (300 seconds)

 300 seconds input rate 0 bits/sec, 0 packets/sec
```

```
 300 seconds output rate 0 bits/sec, 0 packets/sec

 input rate 0 bps, 0 pps; output rate 0 bps, 0 pps

 RX

 9 unicast packets 2141 multicast packets 0 broadcast packets

 2150 input packets 0 bytes

 0 jumbo packets 0 storm suppression packets

 0 runts 0 giants 0 CRC 0 no buffer

 0 input error 0 short frame 0 overrun 0 underrun 0 ignored

 0 watchdog 0 bad etype drop 0 bad proto drop 0 if down drop

 0 input with dribble 0 input discard

 0 Rx pause

 0 Stomped CRC

 TX

 9 unicast packets 2141 multicast packets 0 broadcast packets

 2150 output packets 0 bytes

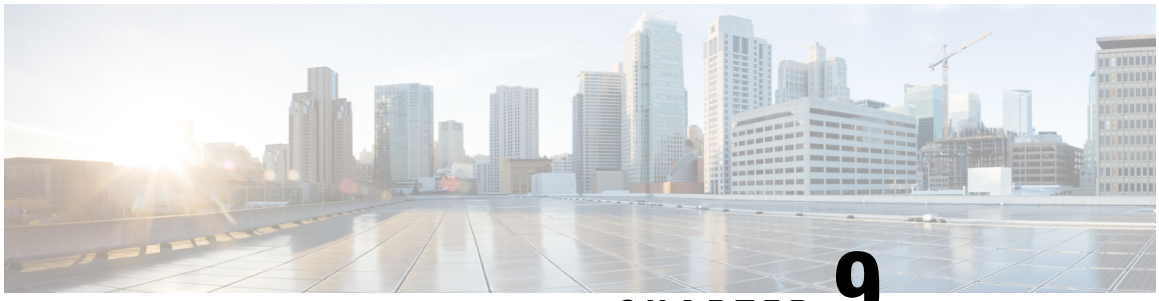
 0 jumbo packets

 0 output error 0 collision 0 deferred 0 late collision

 0 lost carrier 0 no carrier 0 babble 0 output discard

 0 Tx pause


N3550T-TOR1#
```



CHAPTER 9

# Configuring Layer 3 Interfaces


- About Layer 3 Interfaces, on page 151
- Prerequisites for Layer 3 Interfaces, on page 154
- Guidelines and Limitations for Layer 3 Interfaces, on page 154
- Default Settings, on page 155
- Configuring Layer 3 Interfaces, on page 155
- Verifying the Layer 3 Interfaces Configuration, on page 160
- Monitoring the Layer 3 Interfaces, on page 162
- Configuration Examples for Layer 3 Interfaces, on page 162
- Related Documents, on page 164


## About Layer 3 Interfaces

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

### Routed Interfaces

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

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

**Note** The default mode for the Cisco Nexus® 3550-T switch interface is Layer 3.

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

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

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

- Input packets/sec
- Output packets/sec

# VLAN Interfaces

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

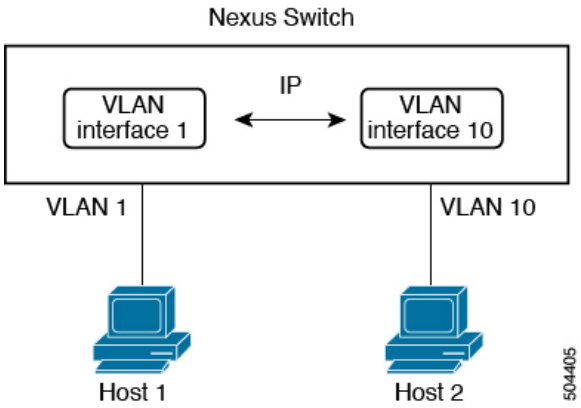
You must enable the VLAN network interface feature before you can see configure it. The system automatically takes a checkpoint prior to disabling the feature, and you can roll back to this checkpoint. See the *Cisco Nexus® 3550-T System Management Configuration* section for information on rollbacks and checkpoints.

**Note** You cannot delete the VLAN interface for VLAN 1.

You can route across VLAN interfaces to provide Layer 3 inter-VLAN routing by configuring a VLAN interface for each VLAN that you want to route traffic to and assigning an IP address on the VLAN interface. For more information about IP addresses and IP routing, see the *Cisco Nexus® 3550-T Unicast Routing Configuration* section.

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

***Figure 22: Connecting Two VLANs with VLAN interfaces***




# Changing VRF Membership for an Interface

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

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



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

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

### Notes About Changing VRF Membership for an Interface

- Momentary traffic loss may occur when changing the VRF name.
- Only the configurations under the interface level are processed when the **system vrf-member-change retain-l3-config** command is enabled. You must manually process any configurations at the router level to accommodate routing protocols after a VRF change.
- The **system vrf-member-change retain-l3-config** command supports interface level configurations with:
  - Layer 3 configurations maintained by the CLI Server, such as **ip address** and all OSPF/ISIS/EIGRP CLIs available under the interface configuration.

## Loopback Interfaces

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

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

## High Availability

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

See the *Cisco Nexus® 3550-T Unicast Routing Configuration* section for complete information about high availability.

## DHCP Client

Cisco NX-OS supports DHCP client for IPv4 and IPv6 addresses on SVIs, physical Ethernet, and management interfaces. You can configure the IP address of a DHCP client by using the **ip address dhcp** or **ipv6 address**

**dhcp** command. These commands send a request from the DHCP client to the DHCP server soliciting an IPv4 or IPv6 address from the DHCP server. The DHCP client on the Cisco Nexus switch identifies itself to the DHCP server. The DHCP server uses this identifier to send the IP address back to the DHCP client.

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

## Limitations for Using DHCP Client on Interfaces

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

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

  If the number of DNS servers configured on the switch is more than six, and if you get a DHCP offer for an SVI with DNS option set, the IP address is not assigned to the SVI.
- A Cisco Nexus 3550-T switch supports a maximum of 10 IPv4 DHCP clients.
- DHCP relay and DHCP client configurations are incompatible and are not supported on the same switch. You must remove the DHCP relay configuration before configuring the DHCP Client on an interface.
- When DHCP snooping is enabled on the VLAN whose SVI is configured with the DHCP client, the DHCP snooping is not enforced on the SVI DHCP client.
- When configuring the IPv4 DHCP client, you must configure with the **ipv4 address use-link-local-only** command before the **ipv4 address dhcp** command.

# Prerequisites for Layer 3 Interfaces

Layer 3 interfaces have the following prerequisites:

- You are familiar with IP addressing and basic configuration. See the *Cisco Nexus® 3550-T Unicast Routing Configuration* section for more information about IP addressing.

# 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.

- IP unnumbered interfaces are not supported.
- Multicast and/or broadcast counters for SVI are not supported.
- Control plane SVI/SI traffic for SVI counters are not supported.
- **show** commands with the **internal** keyword are not supported.



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

# Default Settings

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

***Table 11: Default Layer 3 Interface Parameters***

| Parameters | Default |
|---|---|
| Admin state | Shut |

# Configuring Layer 3 Interfaces

## Configuring a Routed Interface

You can configure any Ethernet port as a routed interface.

### Procedure

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

| | Command or Action | Purpose |
|---|---|---|
| **Step 4** | [**ip address**]<br>**Example:**<br>`switch(config-if)# ip address 192.0.2.1/8` | • Configures an IP address for this interface. See the *Cisco Nexus® 3550-T Unicast Routing Configuration* section for more information about IP addresses. |
| **Step 5** | **show interfaces**<br>**Example:**<br>`switch(config-if)# show interfaces ethernet 1/2` | (Optional) Displays the Layer 3 interface statistics. |
| **Step 6** | **no shutdown**<br>**Example:**<br>`switch#`<br>`switch(config-if)# int e1/2`<br>`switch(config-if)# no shutdown` | (Optional) Clears the errors on the interfaces where policies correspond with hardware policies. This command allows policy programming to continue and the port to come up. If policies do not correspond, the errors are placed in an error-disabled policy state. |
| **Step 7** | **copy running-config startup-config**<br>**Example:**<br>`switch(config)# copy running-config startup-config` | (Optional) Saves the configuration change. |

### Example

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

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

- This example shows how to configure a routed interface:

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

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

# Configuring a VLAN Interface

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

### Procedure

| | Command or Action | Purpose |
|---|---|---|
| **Step 1** | **configure terminal**<br>**Example:**<br>`switch# configure terminal`<br>`switch(config)#` | Enters configuration mode. |
| **Step 2** | **feature interface-vlan**<br>**Example:**<br>`switch(config)# feature interface-vlan` | Enables VLAN interface mode. |
| **Step 3** | **interface vlan** *number*<br>**Example:**<br>`switch(config)# interface vlan 10`<br>`switch(config-if)#` | Creates a VLAN interface. The number range is from 1 to 4094. |
| **Step 4** | [**ip address** *ip-address/length*]<br>**Example:**<br>`switch(config-if)# ip address 192.0.2.1/8` | • Configures an IP address for this VLAN interface. See the *Cisco Nexus® 3550-T Unicast Routing Configuration* section for more information on IP addresses. |
| **Step 5** | **show interface vlan** *number*<br>**Example:**<br>`switch(config-if)# show interface vlan 10` | (Optional) Displays the Layer 3 interface statistics. |
| **Step 6** | **no shutdown**<br>**Example:**<br>`switch(config)# int e1/3`<br>`switch(config)# no shutdown` | (Optional) Clears the errors on the interfaces where policies correspond with hardware policies. This command allows policy programming to continue and the port to come up. If policies do not correspond, the errors are placed in an error-disabled policy state. |
| **Step 7** | **copy running-config startup-config**<br>**Example:**<br>`switch(config-if)# copy running-config startup-config` | (Optional) Saves the configuration change. |

### Example

This example shows how to create a VLAN interface:

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

# Enabling Layer 3 Retention During VRF Membership Change

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

### Procedure

| | Command or Action | Purpose |
|---|---|---|
| **Step 1** | **configure terminal**<br>**Example:**<br>`switch# configure terminal`<br>`switch(config)#` | Enters configuration mode. |
| **Step 2** | **system vrf-member-change retain-l3-config**<br>**Example:**<br>`switch(config)# system vrf-member-change retain-l3-config`<br>`Warning: Will retain L3 configuration when vrf member change on interface.` | Enables Layer 3 configuration retention during VRF membership change.<br>**Note**<br>To disable the retention of the Layer 3 configuration, use the **no system vrf-member-change retain-l3-config** command. |

# Configuring a Loopback Interface

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

### Before you begin

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

### Procedure

| | Command or Action | Purpose |
|---|---|---|
| **Step 1** | **configure terminal**<br>**Example:**<br>`switch# configure terminal`<br>`switch(config)#` | Enters configuration mode. |
| **Step 2** | **interface loopback** *instance*<br>**Example:**<br>`switch(config)# interface loopback 0`<br>`switch(config-if)#` | Creates a loopback interface. The range is from 0 to 1023. |
| **Step 3** | [**ip address** *ip-address/length*]<br>**Example:**<br>`switch(config-if)# ip address 192.0.2.1/8` | • Configures an IP address for this interface. See the *Cisco Nexus® 3550-T Unicast Routing Configuration* section for more information about IP addresses. |

| | Command or Action | Purpose |
|---|---|---|
| **Step 4** | **show interface loopback** *instance*<br>**Example:**<br>`switch(config-if)# show interface loopback 0` | (Optional) Displays the loopback interface statistics. |
| **Step 5** | **copy running-config startup-config**<br>**Example:**<br>`switch(config-if)# copy running-config startup-config` | (Optional) Saves the configuration change. |

### Example

This example shows how to create a loopback interface:

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

# Assigning an Interface to a VRF

You can add a Layer 3 interface to a VRF.

### Procedure

| | Command or Action | Purpose |
|---|---|---|
| **Step 1** | **configure terminal**<br>**Example:**<br>`switch# configure terminal`<br>`switch(config)#` | Enters configuration mode. |
| **Step 2** | **interface** *interface-type number*<br>**Example:**<br>`switch(config)# interface loopback 0`<br>`switch(config-if)#` | Enters interface configuration mode. |
| **Step 3** | **vrf member** *vrf-name*<br>**Example:**<br>`switch(config-if)# vrf member RemoteOfficeVRF` | Adds this interface to a VRF. |
| **Step 4** | **ip address** *ip-prefix/length*<br>**Example:**<br>`switch(config-if)# ip address 192.0.2.1/16` | Configures an IP address for this interface. You must do this step after you assign this interface to a VRF. |

| | Command or Action | Purpose |
|---|---|---|
| **Step 5** | **show vrf** [*vrf-name*] **interface** *interface-type number*<br>**Example:**<br>`switch(config-vrf)# show vrf Enterprise interface loopback 0` | (Optional) Displays VRF information. |
| **Step 6** | **copy running-config startup-config**<br>**Example:**<br>`switch(config-if)# copy running-config startup-config` | (Optional) Saves the configuration change. |

### Example

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

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

# Configuring a DHCP Client on an Interface

You can configure the DHCP client on a management interface, for IPv4 address.

### Procedure

| | Command or Action | Purpose |
|---|---|---|
| **Step 1** | switch# **configure terminal** | Enters global configuration mode. |
| **Step 2** | switch(config)# **interface ethernet** *type slot/port* \| **mgmt** *mgmt-interface-number* \| | Creates a physical Ethernet interface, a management interface. |
| **Step 3** | switch(config-if)# [**no**] [**ip** \| **ipv4] address dhcp** | Requests the DHCP server for an IPv4 address.<br>The **no** form of this command removes any address that was acquired. |
| **Step 4** | Save the configuration. | |

# Verifying the Layer 3 Interfaces Configuration

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

| 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. |
| **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. |
| **show ip interface brief** | Displays interface address and interface status (numbered/unnumbered). |
| **show ip route** | Displays routes learned via OSPF or ISIS. (Includes addresses for best unicast and multicast next-hop.) |

# Monitoring the Layer 3 Interfaces

Use the following commands to display Layer 3 statistics:

| Command | Purpose |
|---|---|
| **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 errors** *slot/port* **detailed** [**all**] | Displays the Layer 3 interface statistics. You can optionally include all 32-bit and 64-bit packet and byte counters (including errors). |
| **show interface ethernet errors** *slot/port* **counters errors** | Displays the Layer 3 interface input and output errors. |
| **show interface ethernet errors** *slot/port* **counters snmp** | Displays the Layer 3 interface counters reported by SNMP MIBs. |
| **show interface loopback** *number* **detailed** [**all**] | Displays the loopback interface statistics. You can optionally include all 32-bit and 64-bit packet and byte counters (including errors). |
| **show interface 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** *number* **counters snmp** | Displays the VLAN interface counters reported by SNMP MIBs. |

# Configuration Examples for Layer 3 Interfaces

This example shows how to configure a loopback interface:

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

## Example of Changing VRF Membership for an Interface

- Enable Layer 3 configuration retention when changing VRF membership.

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

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

- Verify Layer 3 retention.

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

system vrf-member-change retain-l3-config
```

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

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

interface Vlan2002
description TESTSVI
no shutdown
vrf member blue
no ip redirects
ip address 192.168.211.2/27
ip router ospf 1 area 0.0.0.0
preempt delay minimum 300 reload 600
priority 110 forwarding-threshold lower 1 upper 110
ip 192.168.211.1
preempt delay minimum 300 reload 600
priority 110 forwarding-threshold lower 1 upper 110
```

- Verify SVI interface after VRF change.

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

interface Vlan2002
description TESTSVI
no shutdown
vrf member red
no ip redirects
ip address 192.168.211.2/27
ip router ospf 1 area 0.0.0.0
preempt delay minimum 300 reload 600
priority 110 forwarding-threshold lower 1 upper 110
ip 192.168.211.1
preempt delay minimum 300 reload 600
priority 110 forwarding-threshold lower 1 upper 110
```



**Note**

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

# Related Documents

| Related Documents | Document Title |
|---|---|
| IP | *Cisco Nexus® 3550-T Unicast Routing Configuration* section |
| VLANs | *Cisco Nexus® 3550-T Layer 2 Switching Configuration* section |