



P Commands

- [peer-gateway](#), on page 2
- [peer-keepalive destination](#), on page 3
- [peer-switch](#), on page 7
- [port-channel limit](#), on page 8
- [port-channel load-balance](#), on page 10
- [port-channel load-balance hash-modulo](#), on page 12
- [port-channel load-defer](#), on page 13
- [port-profile](#), on page 14
- [profile \(EOAM\)](#), on page 17

peer-gateway

To configure the device to send virtual port-channel (vPC) packets to the device's MAC address, use the **peer-gateway** command. To return to the default value, use the **no** form of this command.

```
peer-gateway
no peer-gateway
```

Syntax Description This command has no arguments or keywords.

Command Default None

Command Modes vpc-domain configuration mode

Release	Modification
4.2(1)	This command was introduced.

Usage Guidelines Use the **peer-gateway** command to have a vPC peer device act as the gateway even for packets that are sent to the vPC peer device's MAC address.

This command does not require a license.

Examples

This example shows how to configure the device to use the switch gateway even for the packets that are sent to the vPC:

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

Command	Description
vpc-domain	Configures a vPC domain and enters the vpc-domain configuration mode.

peer-keepalive destination

To configure the virtual port-channel (vPC) peer-keepalive link and message between vPC peer devices, use the **peer-keepalive destination** command.

```
peer-keepalivedestination 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-bytes tos-byte-value] [source ip address] [udp-port number] [vrf {name | management | vpc-keepalive}]
```

Syntax Description

<i>ip address</i>	IP address of the remote vPC peer device. Note You can either enter an IPv4 or IPv6 address.
hold-timeout <i>secs</i>	(Optional) Specifies when the peer-keepalive link goes down, the secondary vPC peer device waits the hold-timeout interval. The range is from 3 to 10. During the hold-timeout, the vPC secondary device does not take any action based on any keepalive messages received, because the keepalive might be received just temporarily, such as if a supervisor fails a few seconds after the peer link goes down.
interval <i>msecs</i>	Specifies the number of milliseconds that you want between sending keepalive messages to the remote vPC peer device. This variable configures the interval between sending peer-keepalive messages to the remote vPC peer device and the maximum period to wait to receive a keepalive message from the remote vPC peer device. The range is from 400 to 10,000.
timeout <i>secs</i>	(Optional) Specifies that the timeout 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 range is from 3 and 20. 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.

precedence <i>prec-value</i>	(Optional) Specifies the precedence value for the peer-keepalive message. Valid values are as follows: <ul style="list-style-type: none"> • 0 to 7 • network (7) • internet (6) • critical (5) • flash-override (4) • flash (3) • immediate (2) • priority (1) • routine (0)
tos tos-value	(Optional) Specifies the precedence or ToS value for the peer-keepalive message. Valid values are as follows: <ul style="list-style-type: none"> • 0, 1, 2, 4, 8 • max-reliability (2) • max-throughput (4) • min-delay (8) • min-monetary-cost (1) • normal (0) <p>Note The only valid values are shown here.</p>
tos-byte	(Optional) Specifies the precedence, or 8-bit ToS value, for the peer-keepalive message. A higher numerical value indicates the higher throughput priority. The range is from 0 to 255.
source	(Optional) Specifies the IP address of the local vPC peer device. Note Must be an IPv4 address.
<i>number</i>	(Optional) Number of the UDP port to send and receive the vPC peer-keepalive messages. The range is from 1024 to 6500.
<i>name</i>	(Optional) Name of the virtual routing and forwarding (VRF) instance that you want to use for the vPC peer-keepalive link and messages.
vrf vrf-name	(Optional) Specifies a VRF instance.
management	(Optional) Specifies the management interface.
vpc-keepalive	(Optional) Specifies a vPC keepalive.

Command Default

Peer-keepalive is disabled.
 Hold-timeout is 3 seconds.
 Interval is 1000 milliseconds.
 Timeout is 5 seconds.
 Precedence is default, with a level of 6 (internet).
 UDP port is 3200.
 VRF is management VRF.

Command Modes

vpc-domain configuration mode

Command History

Release	Modification
4.1(3)	This command was introduced.

Usage Guidelines

You must enable the vPC feature before you can configure the peer-keepalive parameters. The vPC keepalive messages notify the system if one of the vPC peer devices goes down.

You must configure the peer-keepalive messages on each of the vPC peer devices to enable the functionality.

Although the keepalive messages can transmit over any Layer 3 topology, we recommend that you create and configure a separate VRF with Layer 3 ports on each vPC peer device as the source and destination for the vPC keepalive messages. The default ports and VRF for the peer-alive link are the management ports and the management VRF. Do not use the peer link itself for the vPC peer-keepalive messages.

Ensure that both the source and destination IP addresses used for the peer-keepalive messages are unique in your network.

The vPC keepalive messages are IP/UDP messages.

This command accepts only IPv4 addresses.

The device assumes that its vPC peer device is down when the device does not receive any messages from the peer during the timeout period. We recommend that you configure the timeout value to be three times the interval value.

You can configure either the **precedence**, **tos**, or **tos-byte** value to ensure throughput for the vPC peer-keepalive message.

**Note**

We recommend that you create a separate VRF and assign a Layer 3 port on each vPC peer device for the peer-keepalive link.

This command does not require a license.

Examples

This example shows how to configure the IP address of the remote vPC peer device for the fault-tolerant link:

```
switch(config-vpc-domain)# peer-keepalive destination 172.28.231.85
```

Related Commands

Command	Description
show running-config vpc all	Displays information on the vPC peer-keepalive status. If the feature is not enabled, the system displays an error when you enter this command.
show vpc peer-keepalive	Displays information on the vPC peer-keepalive status. If the feature is not enabled, the system displays an error when you enter this command.

peer-switch

To enable the virtual port channel (vPC) switch pair to appear as a single Spanning Tree Protocol (STP) root in the Layer 2 topology, use the **peer-switch** command. To disable the peer switch vPC topology, use the **no** form of this command.

peer-switch
no peer-switch

Syntax Description This command has no arguments or keywords.

Command Default Peer switch Layer 2 topology is disabled.

Command Modes vPC domain configuration mode

Command History	Release	Modification
	5.0(2)	This command was introduced.

Usage Guidelines This command does not require a license.

Examples

This example shows how to enable the vPC switch pair to appear as a single STP root in the Layer 2 topology:

```
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.
```

Related Commands	Command	Description
	vpc domain	Creates a virtual port-channel (vPC) domain.

port-channel limit

To configure more than 244 virtual port channels (vPCs), use the **port-channel limit** command. To disable this feature, use the **no** form of this command.

port-channel limit
no port-channel limit

Syntax Description This command has no arguments or keywords.

Command Default Limit to 244 vPCs

Command Modes vPC domain configuration

Command History

Release	Modification
6.1(2)E1	This command was introduced.

Usage Guidelines

To enable this command, first enter the fabricpath multicast load-balance command.

Following guidelines when using the no port-channel limit command:

- Entering this command causes the peer links and vPCs to go up and down and could cause traffic losses.
- Only F2 series modules support this configuration. It cannot be configured on VDCs that do not have an F2 series module.
- Entering this command changes FabricPath MAC addresses that are used by vPC+ port channels. It leads to some transient flooding until the MAC addresses are learned again.
- In-service software upgrades (ISSUs) and In-service software downgrades (ISSDs) are not supported.
- Remove the no port-channel limit configuration before attempting an ISSD to an image that does not support this configuration. To revert to an earlier configuration, the number of vPCs that you must be 244 or less.
- To unconfigure the FabricPath multicast load-balance configuration, you must first remove the no port-channel limit configuration.

This command does not require a license.

Examples

This example shows how to configure the maximum number of supported vPCs:

```
switch# switchto vdc peer1
Cisco Nexus Operating System (NX-OS) Software
TAC support: http://www.cisco.com/tac
Copyright (c) 2002-2012, Cisco Systems, Inc. All rights reserved.
The copyrights to certain works contained in this software are
owned by other third parties and used and distributed under
license. Certain components of this software are licensed under
the GNU General Public License (GPL) version 2.0 or the GNU
Lesser General Public License (LGPL) Version 2.1. A copy of each
such license is available at
http://www.opensource.org/licenses/gpl-2.0.php and
```

```

http://www.opensource.org/licenses/lgpl-2.1.php
switch# config t
Enter configuration commands, one per line. End with CNTL/Z.
switch-peer1(config)# vpc domain 1
switch(config-vpc-domain)# port-channel limit
switch(config-vpc-domain)# no port-channel limit
switch(config-vpc-domain)#

```

This example shows how to configure no port-channel limit:

```

switch# configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
switch(config)# vpc domain 1
switch(config-vpc-domain)# fabricpath multicast load-balance
switch(config-vpc-domain)# no port-channel limit
switch(config-vpc-domain)#

```

This example shows how to enable support of more than 244 vPC+ port channels:

```

switch1# configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
switch1(config)# vpc domain 1
switch1(config-vpc-domain)# fabricpath multicast load-balance
switch1(config-vpc-domain)# no port-channel limit

```

Related Commands

Command	Description
show vpc brief	Displays a brief status of the vPC.

port-channel load-balance

To set the load-balancing method among the interfaces in the channel-group bundle, use the **port-channel load-balance** command. To return the system priority to the default value, use the **no** form of this command.

port-channel load-balance *method* [**module** *slot*]
noport-channel load-balance *method* [**module** *slot*]

Syntax Description

<i>method</i>	Load-balancing method. See the “Usage Guidelines” section for a list of valid values.
module <i>slot</i>	(Optional) Specifies the module slot number.

Command Default

Layer 2 packets—**src-dst-mac**

Layer 3 packets—**src-dst-ip**

Command Modes

Global configuration mode

Command History

Release	Modification
6.2(2)	Added the new method “vlan-only” which is applicable at the module level (for F2/F2e modules only).
5.1(3)	The word ethernet was removed from the command name.
4.0	This command was introduced.

Usage Guidelines

When you do not specify a module, you are configuring load balancing for the entire device. When you use the module parameter, you are configuring load balancing for the specified modules.

A load-balance hash will be effective on the egress module only when the desired load-balance hash is configured on the ingress module also. So make sure that you configure the desired port-channel load-balance hash on the module where the traffic ingresses.

Valid *method* values are as follows:

- **dst-ip**—Loads distribution on the destination IP address.
- **dst-mac**—Loads distribution on the destination MAC address.
- **dst-port**—Loads distribution on the destination port.
- **src-dst-ip**—Loads distribution on the source XOR-destination IP address.
- **src-dst-mac**—Loads distribution on the source XOR-destination MAC address.
- **src-dst-port**—Loads distribution on the source XOR-destination port.
- **src-ip**—Loads distribution on the source IP address.
- **src-mac**—Loads distribution on the source MAC address.
- **src-port**—Loads distribution on the source port.

- `vlan-only`—Loads distribution on the vlan modules only.



Note You cannot configure load balancing using port channels per virtual device context (VDC). You must be in the default VDC to configure this feature; if you attempt to configure this feature from another VDC the system returns an error.

Use the **module** keyword to configure the module independently for port-channeling and load-balancing mode. The remaining module uses current load-balancing method configured for the entire device or the default method if you have not configured a method for the entire device. When you enter the **no** form with the **module** keyword, the load-balancing method for the specified module takes the current load-balancing method that is in use for the entire device. If you configured a load-balancing method for the entire device, the specified module uses that configured method rather than the default **src-dst-ip/src-dst-mac**. The per-module configuration takes precedence over the load-balancing method configured for the entire device.

You can configure one load-balancing mode for the entire device, a different mode for specified modules, and another mode for other specified modules. The per-module configuration takes precedence over the load-balancing configuration for the entire device.

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

This command does not require a license.

Examples

This example shows how to set the load-balancing method for the entire device to use the source port:

```
switch(config)# port-channel load-balance src-port
```

This example shows how to set the load-balancing method for the module level (for F2/F2e modules only).

```
switch(config)# port-channel load-balance vlan-only module 1
ERROR: Command is valid for F2/F2E Module only
switch(config)# port-channel load-balance vlan-only module 4
switch(config)#
```

Related Commands

Command	Description
<code>show port-channel load-balance</code>	Displays information about port-channel load balancing.

port-channel load-balance hash-modulo

To enable the modulo hash for Cisco nexus 7000 Series modules, use the **port-channel load-balance hash-modulo** command. To turn off this feature command, use the **no** form of this command.

port-channel load-balance hash-modulo force
no port-channel load-balance hash-modulo force

Syntax Description

force	Specifies the force.
--------------	----------------------

Command Default

Disabled

Command Modes

Global configuration mode

Command History

Release	Modification
6.1(3)	This command was introduced.

Usage Guidelines



Caution

Once you enter the force keyword, the command immediately reinitializes all of the port channels.

By default, when the system comes if the system allows M1 Series Module capability and so, the port-channel load-balance hash-modulo displays an error. Enter the system module-type command to remove “M1 Series Module capability” first and then the command will work.

This command does not require a license.

Examples

This example shows how to enable the modulo hash for the Cisco nexus 7000 Series Module:

```
switch# port-channel load-balance hash-modulo
This command will reinitialize all the port-channels. Do you want to continue(y/n)? [no] y
Warning: This operation may take some time to complete
switch(config)#
```

This example shows how to specify the force:

```
switch# port-channel load-balance hash-modulo force
Warning: This operation may take some time to complete
```

This example shows how to turn on and off this feature:

```
switch(config)# no port-channel load-balance hash-modulo force
Warning: This operation may take some time to complete
switch(config)#
```

Related Commands

Command	Description
show port-channel load-balance	Displays information about port-channel load balancing.

port-channel load-defer

To set the load defer time interval, use the **port-channel load-defer** command. To return the system priority to the default value, use the **no** form of this command.

port-channel load-defer *seconds*
no port-channel load-defer *seconds*

Syntax Description

<i>seconds</i>	Time interval in seconds. The range is from 1 to 1800.
----------------	--

Command Default

120 seconds

Command Modes

Global configuration mode

Command History

Release	Modification
5.1(2)	This command was introduced.

Examples

This example shows how to set the load defer time interval:

```
switch(config)# port-channel load-defer 100
switch(config)#
```

Related Commands

Command	Description
show port-channel load-balance	Displays information about port-channel load balancing.

port-profile

To create a port profile and enter the port-profile configuration mode or to enter into the port-profile configuration mode of a previously created port profile, use the **port-profile** command. To remove the port profile, use the **no** form of this command.

port-profile [**type** {**ethernet** | **interface-vlan** | **port-channel**}] *name*
no port-profile [**type** {**ethernet** | **interface-vlan** | **port-channel**}] *name*

Syntax Description

type	(Optional) Specifies the type of interfaces.
ethernet	Specifies Layer 2 or Layer 3 interfaces.
interface-vlan	Specifies VLAN network interfaces.
port-channel	Specifies port-channel interfaces.
<i>name</i>	Name of the port profile.

Command Default

None

Command Modes

Interface configuration mode
 Port-profile configuration mode

Command History

Release	Modification
4.2(1)	This command was introduced.

Usage Guidelines

Use the **port-profile** command to group configuration commands and apply them to several interfaces simultaneously. All interfaces in the range must be the same type. The maximum number of interfaces that can inherit a single port profile is 512.

The port-profile name must be globally unique across types and networks.

Each port profile can be applied only to a specific type of interface; the choices are as follows:

- Ethernet
- VLAN network interface
- Port channel



Note

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

A subset of commands are available under the port-profile configuration mode, depending on which interface type you specify. Layer 3 and CTS commands are not supported by port profiles.

You can configure the following port-profile operations:

- Create port profiles
- Delete port profiles
- Add commands to and delete commands from port profiles
- Inherit port profiles at interfaces
- Enable and disable port profiles
- Inherit between port profiles
- Configure maximum number of ports that a profile can inherit

You inherit the port profile when you attach the port profile to an interface or range of interfaces. The maximum number of interfaces that can inherit a single profile is 512. When you attach, or inherit, a port profile to an interface or range of interfaces, the system applies all the commands in that port profile to the interfaces.

Additionally, you can have one port profile inherit another port profile, which allows the initial port profile to assume all of the commands of the second, inherited port profile that do not conflict with the initial port profile. Four levels of inheritance are supported except for the **switchport private-vlan mapping** and **private-vlan mapping** commands, which support only one level of inheritance. See the **inherit port-profile** command for information about inheriting an additional port profile and assigning port profiles to specified interfaces.

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

- Commands that you enter under the interface mode take precedence over the port profile's commands if there is a conflict. However, the port profile retains that command in the port profile.
- The port profile's commands take precedence over default commands on the interface, unless it is explicitly overridden by the default command.
- When a range of interfaces inherits a second port profile, the commands of the initial port profile override those commands of the second port profile if there is a conflict.
- After you inherit a port profile onto an interface or range of interfaces, you can override individual configuration values by entering the new value at the interface configuration level. If you then remove the individual configuration values at the interface configuration level, the interface again uses the values in the port profile again.
- There are no default configurations associated with a port profile.



Note You cannot use port profiles with Session Manager.

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

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

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

You must enable each specific port profile using the **state-enabled** command.

This command does not require a license.

Examples

This example shows how to configure, name a port profile, and enter the port-profile configuration mode:

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

Related Commands

Command	Description
state-enable	Enables a specified port profile.
show port-profile	Displays information about port profiles.

profile (EOAM)

To attach an Ethernet OAM profile to an interface, use the **profile** command in interface Ethernet OAM configuration mode. To remove the profile from the interface, use the **no** form of this command.

profile name
noprofile name

Syntax Description

name Text name of the Ethernet OAM profile to attach to the interface.

Command Default

None

Command Modes

Interface Ethernet OAM configuration (config-if-eoam)

Command History

Release	Modification
7.3(0)D1(1)	This command was introduced.

Usage Guidelines

When an Ethernet OAM profile is attached to an interface using this command, all of the parameters configured for the profile are applied to the interface.

Individual parameters that are set by the profile configuration can be overridden by configuring them directly on the interface.

This command does not require a license.

The following example shows how to attach an Ethernet OAM profile to a Ethernet interface.

```
switch# configure terminal
switch(config)# interface ethernet 2/19
switch(config-if)# ethernet oam
switch(config-if-eoam)# profile Profile_1
```

Related Commands

Command	Description
ethernet oam profile	Creates an EOAM profile and enters EOAM configuration mode.
ethernet oam	Enables Ethernet Link OAM, with default values, on an interface and enter interface Ethernet OAM configuration mode.
show ethernet oam configuration	Displays the current active Ethernet OAM configuration on an interface.
show ethernet oam interfaces	Displays the current state of Ethernet OAM interfaces.

