CONFIGURATION GUIDE

Extreme SLX-OS Layer 2 Configuration Guide, 17r.1.01

Supporting the ExtremeRouting SLX 9850 and ExtremeSwitching SLX 9540 Devices

53-1005157-05
November 2018

Legal Notice

Extreme Networks, Inc. reserves the right to make changes in specifications and other information contained in this document and its website without prior notice. The reader should in all cases consult representatives of Extreme Networks to determine whether any such changes have been made.

The hardware, firmware, software or any specifications described or referred to in this document are subject to change without notice.

Trademarks

Extreme Networks and the Extreme Networks logo are trademarks or registered trademarks of Extreme Networks, Inc. in the United States and/or other countries.

All other names (including any product names) mentioned in this document are the property of their respective owners and may be trademarks or registered trademarks of their respective companies/owners.

For additional information on Extreme Networks trademarks, please see: www.extremenetworks.com/company/legal/trademarks

Software Licensing

Some software files have been licensed under certain open source or third-party licenses. End-user license agreements and open source declarations can be found at: www.extremenetworks.com/support/policies/software-licensing

Support

For product support, phone the Global Technical Assistance Center (GTAC) at 1-800-998-2408 (toll-free in U.S. and Canada) or +1-408-579-2826. For the support phone number in other countries, visit: http://www.extremenetworks.com/support/contact/

For product documentation online, visit: https://www.extremenetworks.com/documentation/

Contents

Preface

This section discusses the conventions used in this guide, ways to provide feedback, additional help, and other Extreme Networks® publications.

Conventions

This section discusses the conventions used in this guide.

Notes, cautions, and warnings

Notes, cautions, and warning statements may be used in this document. They are listed in the order of increasing severity of potential hazards.

NOTE

A Note provides a tip, guidance, or advice, emphasizes important information, or provides a reference to related information.

ATTENTION

An Attention statement indicates a stronger note, for example, to alert you when traffic might be interrupted or the device might reboot.

CAUTION

A Caution statement alerts you to situations that can be potentially hazardous to you or cause damage to hardware, firmware, software, or data.

DANGER

A Danger statement indicates conditions or situations that can be potentially lethal or extremely hazardous to you. Safety labels are also attached directly to products to warn of these conditions or situations.

Text formatting conventions

Text formatting conventions such as boldface, italic, or Courier font may be used to highlight specific words or phrases.

Format	Description
bold text	Identifies command names.
	Identifies keywords and operands.
	Identifies the names of GUI elements.
	Identifies text to enter in the GUI.
italic text	Identifies emphasis.
	Identifies variables.
	Identifies document titles.

Format	Description
`Courier font`	Identifies CLI output. Identifies command syntax examples.

Command syntax conventions

Bold and italic text identify command syntax components. Delimiters and operators define groupings of parameters and their logical relationships.

Convention	Description
bold text	Identifies command names, keywords, and command options.
italic text	Identifies a variable.
[ ]	Syntax components displayed within square brackets are optional. Default responses to system prompts are enclosed in square brackets.
{ **x** | **y** | **z** }	A choice of required parameters is enclosed in curly brackets separated by vertical bars. You must select one of the options.
x | **y**	A vertical bar separates mutually exclusive elements.
< >	Nonprinting characters, for example, passwords, are enclosed in angle brackets.
...	Repeat the previous element, for example, *member*[*member*...].
\	Indicates a "soft" line break in command examples. If a backslash separates two lines of a command input, enter the entire command at the prompt without the backslash.

Documentation and Training

To find Extreme Networks product guides, visit our documentation pages at:

Current Product Documentation	www.extremenetworks.com/documentation/
Archived Documentation (for earlier versions and legacy products)	www.extremenetworks.com/support/documentation-archives/
Release Notes	www.extremenetworks.com/support/release-notes
Hardware/Software Compatibility Matrices	https://www.extremenetworks.com/support/compatibility-matrices/
White papers, data sheets, case studies, and other product resources	https://www.extremenetworks.com/resources/

Open Source Declarations

Some software files have been licensed under certain open source licenses. More information is available at: www.extremenetworks.com/support/policies/open-source-declaration/.

Training

Extreme Networks offers product training courses, both online and in person, as well as specialized certifications. For more information, visit www.extremenetworks.com/education/.

Getting Help

If you require assistance, contact Extreme Networks using one of the following methods:

Extreme Portal	Search the GTAC (Global Technical Assistance Center) knowledge base, manage support cases and service contracts, download software, and obtain product licensing, training, and certifications.
The Hub	A forum for Extreme Networks customers to connect with one another, answer questions, and share ideas and feedback. This community is monitored by Extreme Networks employees, but is not intended to replace specific guidance from GTAC.
Call GTAC	For immediate support: 1-800-998-2408 (toll-free in U.S. and Canada) or +1 408-579-2826. For the support phone number in your country, visit: www.extremenetworks.com/support/contact

Before contacting Extreme Networks for technical support, have the following information ready:

- Your Extreme Networks service contract number and/or serial numbers for all involved Extreme Networks products
- A description of the failure
- A description of any action(s) already taken to resolve the problem
- A description of your network environment (such as layout, cable type, other relevant environmental information)
- Network load at the time of trouble (if known)
- The device history (for example, if you have returned the device before, or if this is a recurring problem)
- Any related RMA (Return Material Authorization) numbers

Subscribing to Service Notifications

You can subscribe to email notifications for product and software release announcements, Vulnerability Notices, and Service Notifications.

1. Go to www.extremenetworks.com/support/service-notification-form.
2. Complete the form with your information (all fields are required).
3. Select the products for which you would like to receive notifications.

   NOTE
   You can modify your product selections or unsubscribe at any time.

4. Click **Submit**.

Providing Feedback to Us

Quality is our first concern at Extreme Networks, and we have made every effort to ensure the accuracy and completeness of this document. We are always striving to improve our documentation and help you work better, so we want to hear from you! We welcome all feedback but especially want to know about:

- Content errors or confusing or conflicting information.
- Ideas for improvements to our documentation so you can find the information you need faster.
- Broken links or usability issues.

If you would like to provide feedback to the Extreme Networks Information Development team, you can do so in two ways:

- Use our short online feedback form at https://www.extremenetworks.com/documentation-feedback/.

- Email us at documentation@extremenetworks.com.

Please provide the publication title, part number, and as much detail as possible, including the topic heading and page number if applicable, as well as your suggestions for improvement.

About This Document

Supported hardware and software

In those instances in which procedures or parts of procedures documented here apply to some devices but not to others, this guide identifies exactly which devices are supported and which are not.

Although many different software and hardware configurations are tested and supported by Extreme Networks, Inc. for SLX-OS Release 17r.1.01, documenting all possible configurations and scenarios is beyond the scope of this document.

The following hardware platforms are supported by this release:

- ExtremeRouting SLX 9850-4
- ExtremeRouting SLX 9850-8
- ExtremeSwitching SLX 9540

Interface module capabilities

The following table lists the supported capabilities for the following SLX 9850 interface modules:

- BR-SLX9850-10Gx72S-M
- BR-SLX9850-100Gx36CQ-M
- BR-SLX9850-10Gx72S-D
- BR-SLX9850-100Gx36CQ-D

TABLE 1 SLX 9850 interface modules capabilities

Capability	Modular interface module
MPLS	Yes
Packet buffer memory per interface module	12GB (BR-SLX9850-10Gx72S-M) 36GB (BR-SLX9850-100Gx36CQ-M) 8GB (BR-SLX9850-10Gx72S-D) 24GB (BR-SLX9850-100Gx36CQ-D)

What's new in this document

The following table includes descriptions of new information added to this guide for the SLX OS 17r.1.01 software release.

TABLE 2 Summary of enhancements in SLX OS release 17r.1.01

Feature	Description	Described in
VPLS and VLL MCT	SLX-OS MCT acts as a data center gateway to connect to another data center through either the VPLS or VLL WAN connection.	VPLS and VLL MCT on page 68

TABLE 2 Summary of enhancements in SLX OS release 17r.1.01 (continued)

Feature	Description	Described in
VXLAN Layer 2 gateway support for bridge domains	The SLX-OS device provides VXLAN Layer 2 gateway support for bridge domains in addition to supporting Layer 2 VLANS.	VXLAN Layer 2 gateway support for bridge domains on page 48

For complete information, refer to the *SLX-OS 17r.1.01 Release Notes*.

Link Aggregation

Link aggregation overview

Link aggregation allows you to bundle multiple physical Ethernet links to form a single logical trunk providing enhanced performance and redundancy. The aggregated trunk is referred to as a Link Aggregation Group (LAG). The LAG is viewed as a single link by connected devices, the Spanning Tree Protocol, IEEE 802.1Q VLANs, and so on. When one physical link in the LAG fails, the other links stay up. There is no disruption to traffic.

To configure links to form a LAG, the physical links must be of the same speed. Link aggregation can be done by manually configuring the LAG, or by dynamically configuring the LAG using the IEEE 802.3ad Link Aggregation Control Protocol (LACP).

When queuing traffic from multiple input sources to the same output port, all input sources are given the same weight, regardless of whether the input source is a single physical link or a trunk with multiple member links.

NOTE
The LAG or LAG interface is also referred to as a *port-channel* in the SLX 9850 platform.

The benefits of link aggregation are summarized as follows:

- Increased bandwidth (The logical bandwidth can be dynamically changed as the demand changes.)
- Increased availability
- Load sharing
- Rapid configuration and reconfiguration

Each LAG consists of the following components:

- A MAC address that is different from the MAC addresses of the LAG's individual member links.
- An interface index for each link to identify the link to the neighboring devices.
- An administrative key for each link. Only the links with the same administrative key value can be aggregated into a LAG. On each link configured to use LACP, LACP automatically configures an administrative key value equal to the port-channel identification number.

The SLX 9850 platform supports the following LAG types:

- Static LAG— In static link aggregation, links are added into a LAG without exchanging any control packets between the partner systems. The distribution and collection of frames on static links is determined by the operational status and administrative state of the link.
- Dynamic, standards-based LAG using LACP—Dynamic link aggregation uses LACP to negotiate with links that can be added and removed from a LAG. Typically, two partner systems sharing multiple physical Ethernet links can aggregate a number of those physical links using LACP. LACP creates a LAG on both partner systems and identifies the LAG by the LAG ID. All links with the same administrative key, and all links that are connected to the same partner switch become members of the LAG. LACP continuously exchanges LACPDUs to monitor the health of each member link.

The SLX 9850 platform supports the following trunk type:

- Static, standards-based LAG

The SLX 9850 platform supports the following LAG scalability configuration:

- The **default** profile supports 256 LAGs (64 ports per LAG)
- The **Lag-profile-1** profile supports 512 LAGs (32 ports per LAG)

NOTE

The following example enables the lag hardware profile for scaling to 512. The user has to save and reload to activate a new profile. Use the same procedure to revert to the default profile.

```
device# configure terminal
Entering configuration mode terminal
device(config)# hardware
device(config-hardware)# profile lag lag-profile-1
%Warning: To activate the new profile config, please run 'copy running-config startup-config'
followed by 'reload system'.
device(config-hardware)#
```

LAG load sharing

Extreme devices can be configured for load sharing over a LAG by using the following:

- Hash-based load sharing

Hash-based load sharing

The Extreme device tries to share the traffic load evenly across the ports in LAG group, while ensuring that packets in the flow are not reordered. Individual flows are assigned a LAG index to identify them. An improved hash-based load sharing algorithm has the following enhancements:

- Better distribution
- Support for 32-port LAGs when the maximum number of LAGs in the system is 512.
- Support for 64-port LAGs when the maximum number of LAGs in the system is 256.
- An increased number of fields in the packet header that can be used for load balancing
- Enhanced load sharing in configurations of ECMP with LAGs.

Configuring LAG hashing

To configure symmetric LAG hashing on an SLX 9850 device, complete the following tasks.

1. Define where to start the picking headers for the key generation using the **lag hash hdr-start <fwd |term>** command.
   - fwd— start from the header that is used for the forwarding of the packet (inner header). This is the default option.
   - term— start from the last terminated header (outer header), that is the header below forwarding header. In the case of switching traffic, there would not be any header below forwarding header, thus hashing will not be visible.
2. Configure the number of headers to be considered for LAG hashing using the **lag hash hdr-count <count>** command. The default value is 1. There can be a maximum of 3 headers based on the first header selected using the command in the previous step.

The following options provide other LAG configurations to achieve specific tasks.

- Configure hash rotate using the **lag hash rotate <rotate-number>** command to provide different options for randomness of hashing. The number can be between 0 and 15. The default value is 3.
- Configure hash normalize by using the **lag hash normalize** command if there is a need to use the same hash in both directions. The normalize option is disabled by default.

- Allow the source port to be included in the hashing configuration using the **lag hash srcport** command. The source port is not used for hashing by default.
- To skip the entire MPLS label stack and pick only the BOS label for hashing, use the **lag hash bos <start | skip>**. The command default is If MPLS header is used for hashing, it will use all labels including BOS label for hashing.
  - start— start from BOS. This is the default option.
  - skip— hash from header next to BOS.
- Enter the **lag hash pwctrlword skip** command to skip password control word in the hashing configuration.
- The following MPLS transit node LSR hashing configuration options are available when using the **lag hash speculate-mpls** command. The default option is using the MPLS labels.
  - enable— Enables Speculative MPLS.
  - inner-eth— Enables inner ethernet header hash for L2VPN.
  - inner-ip-raw— Enables inner IPv4 header hash for L2VPN raw mode.
  - inner-ip-tag— Enables inner IPv4 header hash for L2VPN tag mode.
  - inner-ipv6-raw— Enables inner IPv6 header hash for L2VPN raw mode.
  - inner-ipv6-tag— Enables inner IPv6 header hash for L2VPN tag mode.
- Select the fields to be used for LAG hashing per-header type by entering the **[no] lag hash** *protocol-type packet-fields-to-be-used-for-hashing* command.
- Select the protocol header type using one of the following commands. By default, all the header parameters are enabled as shown here. You can disable or enable a parameter only one at any given instant.

  NOTE
  Using the **no** form of the following commands will mask a certain field in the configuration and that field will not be used for load-balance hashing.

  - Ethernet headers. (By default, all the header parameters are enabled as shown here. You can disable or enable a parameter only one at any given instant.) :
    - `[no] load-balance hash ethernet <sa-mac>`
    - `[no] load-balance hash ethernet <da-mac>`
    - `[no] load-balance hash ethernet <etype>`
    - `[no] load-balance hash ethernet <vlan>`
  - IPv4 and L4 headers: `[no] load-balance hash ip < src-ip > <dst-ip > < protocol > < src-l4-port> < dst-l4-port >`
  - IPv6 and L4 headers: `[no] load-balance hash ipv6 < ipv6-src-ip > < ipv6-dst-ip> < ipv6-next-hdr> <ipv6-src-l4-port> < ipv6-dst-l4-port>`
  - MPLS: `[no] load-balance hash mpls < label1 > <label2> < label3>`

Load balancing mechanism on different traffic types

The following table provides information about load balancing on different traffic types.

TABLE 3 Load balancing on different traffic types

Traffic type	Header field	Description
Layer 2/ Layer 3 packet load balancing	• Ethernet DA, SA, Etype, Vlan-id • IPv4/v6 dst IP, src IP • L4 Src-Port, Dst-Port	• Ethernet destination address, source address, ethernet type, VLAN ID load balancing • IPv4/v6 destination address, source address load balancing • Layer 4 source and destination port-based load balancing

TABLE 3 Load balancing on different traffic types (continued)

Traffic type	Header field	Description
VPLS/ VLL packet load balancing	CE to PE router traffic can use the following fields for load-balancing similar to the Layer 2/ Layer 3 traffic) • Ethernet DA, SA, Etype, Vlan-id • IPv4/v6 dst IP, src IP • L4 Src-Port, Dst-Port PE to CE router traffic can use the following fields for load-balancing • Customer (inner) ethernet DA, SA, Etype, Vlan-id • Customer (inner) IPv4/v6 dst IP, Ipv4/Ipv6 src IP, protocol • Customer (inner) L4 Src-Port, Dst-Port	CE to PE router traffic • Ethernet destination address, source address, ethernet type, VLAN ID load balancing • IPv4/v6 destination address, source address load balancing • Layer 4 source and destination port-based load balancing PE to CE router traffic • Customer ethernet destination and source address, ethernet type, VLAN ID load balancing • Customer IPv4/v6 destination address, source address load balancing • Customer Layer 4 source and destination port-based load balancing
MPLS LSR load balancing • Extreme SLX 9850 device provides multiple options to handle different MPLS transit hashing scenarios • The hashing options are mutually exclusive. If one option is enabled, the other option will be disabled.	IP over MPLS traffic going over transit node	Extreme supports speculate-mpls option as default which speculates the IPv4/IPv6 header after the MPLS labels and use the fields for hashing. This hashing scenario is handled by the **lag hash speculate-mpls enable** command in the global mode.
L2VPN (VPLS/VLL) traffic • The hashing options are mutually exclusive. If one option is enabled, the other option will be disabled.	L2VPN tagged mode with IPv4 inner payload	This scenario is handled using the **lag hash speculate-mpls inner-ip-tag** command in the global mode. Some sections of the IPv4 source and destination address fields are also used for load-balance hashing.
	L2VPN raw mode with IPv4 inner payload	This scenario is handled using the **lag hash speculate-mpls inner-ip-raw** command. Some sections of the IPv4 source and destination address fields are also used for load-balance hashing.
	L2VPN tagged mode with IPv6 inner payload	This scenario is handled using the **lag hash speculate-mpls inner-ipv6-tag** command. Some sections of the IPv6 source and destination address fields are also used for load-balance hashing.
	L2VPN raw mode with IPV6 inner payload	This scenario is handled using the **lag hash speculate-mpls inner-ipv6-raw** command. Some sections of the IPv6 source and destination address fields are also used for load-balance hashing.

Displaying LAG hashing

Use the **show port-channel load-balance** command to display the configured parameters for LAG hashing.

```
device# show port-channel load-balance
Header parameters
        Ethernet Mask: sa-mac da-mac etype vlan
        ip: src-ip dst-ip protocol src-l4-port dst-l4-port
        ipv6: ipv6-src-ip ipv6-dst-ip ipv6-next-hdripv6-src-l4-port ipv6-dst-l4-port
        mpls: label1 label2 label3
```

```
Hash Settings
        hdr-start:FWD, hdr-count:1, bos-start:0, bos-skip:0, skip-cw:0
        normalize:0, rotate:3, include_src_port:0, Disable: L2 0, ipv4 0, ipv6 0, mpls 0

mpls_speculate: Enabled

load-balance-type hash-based
```

Link Aggregation Control Protocol

Link Aggregation Control Protocol (LACP) is an IEEE 802.3ad standards-based protocol that allows two partner systems to dynamically negotiate attributes of physical links between them to form logical trunks. LACP determines whether a link can be aggregated into a LAG. If a link can be aggregated into a LAG, LACP puts the link into the LAG. All links in a LAG inherit the same administrative characteristics.

LACP operates in two modes:

- *Active mode*— LACP initiates the LACPDU exchange regardless of whether the partner system sends LACPDUs.
- *Passive mode* — LACP responds to Link Aggregation Control Protocol Data Units (LACPDUs) initiated by its partner system but does not initiate the LACPDU exchange.

LAG distribution process and conditions

The LAG aggregator is associated with the collection and distribution of Ethernet frames. The collection and distribution process is required to guarantee the following:

- Inserting and capturing control PDUs.
- Restricting the traffic of a given conversation to a specific link.
- Load balancing between individual links.
- Handling dynamic changes in LAG membership.

On each port, link aggregation control does the following:

- Maintains configuration information to control port aggregation.
- Exchanges configuration information with other devices to form LAGs.
- Attaches ports to and detaches ports from the aggregator when they join or leave a LAG.
- Enables or disables an aggregator's frame collection and distribution functions.

LAG configuration guidelines:

- Each link in the SLX 9850 hardware can be associated with a LAG; a link cannot be associated with more than one LAG. The process of adding and removing links to and from a LAG is controlled statically or dynamically (through LACP).
- The maximum number of port members that may be assigned to a LAG depends on the LAG profile configuration. By default, the SLX 9850 platform can have a maximum of 256 LAGs with the maximum of 64 ports in each LAG. With the LAG profile 1 , the platform can have a maximum of 512 LAGs with the maximum of 32 ports in each LAG.
- Use the **show hardware profile current** command to view the current LAG profile. When the LAG profile is changed from "default" profile to "LAG-PROFILE-1", the device enables LAG scaling up to 512 LAGs.
- Interfaces configured as switchport interfaces cannot be aggregated into a LAG. However, a LAG can be configured as a switchport.

Configuring and managing Link Aggregation

The following sections discuss working with Link Aggregation on Extreme devices.

Configuring a new port channel interface

Follow this procedure to create a new port channel interface at the global configuration mode.

1. Enter the **configure terminal** command to access global configuration mode.

```
device# configure terminal
```

2. Enter the **interface port-channel** command to create a new port channel interface at the global configuration level.

```
device(config)# interface port-channel 30
```

NOTE
The port-channel interface ranges from 1 to 512.

The following example creates a new port channel interface of 30.

```
device# configure terminal
device(config)# interface port-channel 30
```

After creating a new port channel, you can do "no shutdown" or "shutdown" to bring up or down the port-channel as follows.

```
device# configuration terminal
device(config)# interface Port-channel 30
2016/10/17-20:31:21, [NSM-1004], 302, M2 | Active | DCE, INFO, SLX,  Port-channel 30 is created.
device(config-Port-channel-30)#
device(config-Port-channel-30)# no shutdown
2016/10/17-20:31:26, [NSM-1019], 303, M2 | Active | DCE, INFO, SLX,  Interface Port-channel 30 is
administratively up.
device(config-Port-channel-30)#
```

Deleting a port channel interface

Follow this procedure to delete a port channel interface and all member interfaces from the specified LAG at the global configuration mode.

1. Enter the **configure terminal** command to access global configuration mode.

```
device# configure terminal
```

2. Enter the **no interface port-channel** command to delete an existing port channel interface at the global configuration level.

```
device(config)# no interface port-channel 30
```

NOTE
The port-channel interface ranges from 1 to 512.

The following example deletes the existing port channel interface 30 from the specified LAG.

```
device# configure terminal
device(config)# no interface port-channel 30
```

Adding a member port to a port channel

Follow this procedure to add a port to a specific port channel interface at the interface configuration level. If the port channel is not created, the **channel-group** command creates the port channel and also adds a port to the port channel.

1. Enter the **configure terminal** command to access global configuration mode.

```
device# configure terminal
device(config)#
```

2. Enter the **interface port-channel** command to add a port channel interface at the global configuration level.

```
device(config)# interface port-channel 30
device(conf-Port-channel-30)#
```

3. Configure the **interface ethernet** command to enable the interface.

```
device(conf-Port-channel-30)# interface ethernet 1/5
device(conf-if-eth-1/5)#
```

4. Add a port to the port channel interface as static.

```
device(conf-if-eth-1/5)# channel-group 10 mode on
```

5. Add a port to the port channel interface as a dynamic (using LACP), active or passive mode.

```
device(conf-if-eth-1/5)# channel-group 10 mode active

device(conf-if-eth-1/5)# channel-group 10 mode passive
```

The following example is for a static LAG configuration with the mode ON.

```
device# configure terminal
device(config)# interface port-channel 30
device(conf-Port-channel-30)# interface ethernet 1/5
device(conf-if-eth-1/5)# channel-group 10 mode on
```

The following example adds a port 1/5 to the existing dynamic port channel interface 30 with the mode active.

```
device# configure terminal
device(config)# interface port-channel 30
device(conf-Port-channel-30)# interface ethernet 1/5
device(conf-if-eth-1/5)# channel-group 30 mode active
```

NOTE

Run the **no shutdown** command to bring the above interface online.

```
device(conf-if-eth-1/5)# no shutdown
2016/10/18-03:47:15, [NSM-1019], 528, M2 | Active | DCE, INFO, SLX,  Interface Ethernet 1/5 is
administratively up.2016/10/18-03:47:15, [NSM-1001], 529, M2 | Active | DCE, INFO, SLX,  Interface
Ethernet 1/5 is online.
```

The following example adds a port 1/5 to the existing dynamic port channel interface 30 with the mode passive.

```
device# configure terminal
device(config)# interface port-channel 30
device(conf-Port-channel-30)# interface ethernet 1/5
device(conf-if-eth-1/5)# channel-group 30 mode passive
```

Deleting a member port from a port channel

Follow this procedure to delete a member port from a port channel interface at the interface configuration level.

Delete a port from the port channel interface.

```
device(conf-if-eth-1/5)# no channel-group
```

The following example deletes a port 1/5 from the existing port channel interface 30 in a LAG.

```
device# configure terminal
device(config)# interface ethernet 1/5
device(conf-if-eth-1/5)# no channel-group
```

Configuring the minimum number of LAG member links

Follow this procedure to configure the minimum number of LAG member links that should be functional so that the port-channel interface is operationally up.

This configuration allows a port-channel to operate at a certain minimum bandwidth at all times. If the bandwidth of the port-channel drops below the minimum number, then the port-channel is declared operationally DOWN even though it has operationally UP members.

1. Enter the **configure terminal** command to access global configuration mode.

```
device# configure terminal
device(config)#
```

2. Enter the **interface port-channel** command at the global configuration level.

```
device(config)# interface port-channel 30
device(conf-Port-channel-30)#
```

3. Configure the minimum number of LAG member links at the port-channel interface configuration mode.

```
device(conf-Port-channel-30)# minimum-links 5
```

NOTE
The number of links ranges from 1 to 64. The default minimum links is 1.

The following example sets min-link 5 to the existing port channel interface 30 in a LAG.

```
device# configure terminal
device(config)# interface port-channel 30
device(conf-Port-channel-30)# minimum-links 5
```

Configuring the LACP system priority

The switch must be in privileged EXEC mode.

You configure the LACP system priority on each switch running LACP. LACP uses the system priority with the switch MAC address to form the system ID and also during negotiation with other switches.

The system priority value must be a number in the range of 1 through 65535. The higher the number, the lower the priority. The default priority is 32768.

To configure the global LACP system priority, perform the following steps:

1. Enter the **configure terminal** command to access global configuration mode.

2. Specify the LACP system priority.

```
device(config)# lacp system-priority 25000
```

3. To reset the system priority to the default value.

```
device(config)# no lacp system-priority
```

Configuring the LACP port priority

Follow this procedure to configure the LACP port priority of a member port of a specific port-channel interface.

1. Enter the **configure terminal** command to access global configuration mode.

```
device# configure terminal
device(config)#
```

2. Enter the **interface port-channel** command to add a port channel interface at the global configuration level.

```
device(config)# interface port-channel 30
device(conf-Port-channel-30)#
```

3. Configure the **interface ethernet** command and add the port to the port-channel interface.

```
device(conf-Port-channel-30)# interface ethernet 1/5
device(conf-if-eth-1/5)#channel-group 30 mode active
```

4. Configure the LACP port priority 12 for the member port.

```
device(conf-if-eth-1/5)# lacp port-priority 12
```

NOTE
The LACP port priority value ranges from 1 to 65535. The default value is 32768.

5. To rest the configured port priority to the default value.

```
device(conf-if-eth-1/5)# no lacp port-priority
```

The example sets the port priority as 12.

```
device# configure terminal
device(config)# interface port-channel 30
device(conf-Port-channel-30)# interface ethernet 1/5
device(conf-if-eth-1/5)# channel-group 30 mode active
device(conf-if-eth-1/5)# lacp port-priority 12
```

Configuring the LACP timeout period

The device must be in privileged EXEC mode.

The LACP timeout period indicates how long LACP waits before timing out the neighboring device. The **short** timeout period is 3 seconds and the **long** timeout period is 90 seconds. The default is **long**.

To configure the LACP timeout period on an interface, perform the following steps:

1. Enter the **configure terminal** command to access global configuration mode.
2. Enter the **interface** command, specifying the interface type and the slot/port.

```
device(config)# interface ethernet 1/1
```

3. Enter the **no shutdown** command to enable the interface.
4. Specify the LACP timeout short period for the interface.

```
device(conf-if-eth 1/1)# lacp timeout short
```

5. Specify the LACP timeout long period for the interface.

```
device(conf-if-eth 1/1)# lacp timeout long
```

Configuring LACP default Up

Follow this procedure to activate an LACP link in the absence of PDUs on the interface mode.

Consider the following when using the **lacp default-up** command:

- The command is available only if the configured interface is a dynamic member of a port-channel interface.
- The command is not supported on static LAGs.
- The command is not supported on port-channel interfaces.

1. Enter the **configure terminal** command to access global configuration mode.
2. Enter the **interface** command, specifying the interface type and the slot/port.

```
device(config)# interface ethernet 1/1
```

3. Specify LACP default-up for the interface.

```
device(conf-if-eth-1/1)# lacp default-up
```

4. Enter the no form of the command to disable the configuration.

```
device(conf-if-eth-1/1)#  no lacp default-up
```

LACP PDU forwarding

By default, LACP PDUs received on an interface where LACP is not configured are discarded. For scenarios in which the interface requires LACP PDU packet forwarding, you can configure the device to forward the LACP PDU on the VLAN on which it is received using the **lacp-pdu-forward enable** command in the interface configuration mode or port channel configuration mode.

Since the destination address of the PDU is a multicast MAC, the frame will be flooded on the VLAN. If the VLAN on which the LACP PDU is received is a regular VLAN, the PDU will be flooded on the VLAN. If the VLAN on which the PDU is received is a service delimiter for a bridge domain, the LACP PDU is flooded on the bridge domain accordingly.

LACP PDU forwarding is supported only on physical interfaces and static port channel interfaces. LACP PDUs cannot be forwarded if they are received on a LACP based dynamic port channel. LACP PDU forwarding enabled on a static port channel applies to all the member ports. If LACP is enabled on a port, it overrides the LACP PDU forwarding configuration and the PDUs are trapped in the CPU.

Configuring LACP PDU forwarding on a port-channel interface

Perform the following steps to configure LACP PDU forwarding on a port-channel interface.

1. Enter the global configuration mode.

```
device# configure terminal
device(config)#
```

2. Enter the **interface port-channel** command to add a port channel interface at the global configuration level.

```
device(config)# interface port-channel 10
device(conf-Port-channel-10)#
```

3. Configure LACP PDU forwarding on the port-channel interface.

```
device(conf-Port-channel-10)# lacp-pdu-forward enable
```

LACP PDU forwarding is supported only on static port channel interfaces.

The following example enables LACP forwarding on a port-channel interface.

```
device# configure terminal
device(config)# interface port-channel 10
device(conf-Port-channel-10)# lacp-pdu-forward enable
```

Configuring LACP PDU forwarding on a physical interface

Perform the following steps to configure LACP PDU forwarding on a physical interface.

1. Enter the global configuration mode.

```
device# configure terminal
device(config)#
```

2. Specify the physical interface on which LACP PDU forwarding needs to enabled.

```
device(config)# interface ethernet 4/1
device(conf-if-eth-4/1)#
```

3. Configure LACP PDU forwarding on the physical interface.

```
device(conf-if-eth-4/1)# lacp-pdu-forward enable
```

The following example enables LACP forwarding on a port-channel interface.

```
device# configure terminal
device(config)# interface ethernet 4/1
device(conf-if-eth-4/1)# lacp-pdu-forward enable
```

Displaying port-channel information

Various show commands are used to display information for a specific port-channel interface.

Before displaying the port-channel information, you should have created a port-channel interface in a LAG to generate details.

1. Use the **show port-channel summary** command to display brief information of all port-channels.

```
device# show port-channel summary
Flags:  D - Down               P - Up in port-channel (members)
        U - Up (port-channel)  * - Primary link in port-channel
        S - Switched
        M - Not in use. Min-links not met
===== ============== ========== ===============
Group Port-channel    Protocol   Member ports
===== ============== ========== ===============
1     Po 1    (D)     None       Eth 2/125 (D)
                                 Eth 4/125 (D)
2     Po 2    (D)     None       Eth 2/126 (D)
                                 Eth 4/126 (D)
10    Po 10   (U)     LACP       Eth 2/4* (P)
                                 Eth 2/18 (P)
100   Po 100  (U)     None       Eth 2/10* (P)
                                          Eth 2/11 (P)
```

2. Use the **show port-channel detail** command to display detailed information of all the port-channels.

```
device# show port-channel detail
 Static Aggregator: Po 1
 Aggregator type: Standard
 Number of Ports: 2
 Member ports:
   Eth 2/125
   Eth 4/125


 Static Aggregator: Po 2
 Aggregator type: Standard
 Number of Ports: 2
 Member ports:
   Eth 2/126
   Eth 4/126


 Static Aggregator: Po 100
 Aggregator type: Standard
 Number of Ports: 2
 Member ports:
   Eth 2/10   *
   Eth 2/11


 LACP Aggregator: Po 10
Aggregator type: Standard
  Actor System ID - 0x8000,76-8e-f8-0a-98-00
  Admin Key: 0010 - Oper Key 0010
  Receive link count: 2 - Transmit link count: 2
  Individual: 0 - Ready: 1
  Partner System ID - 0x8000,76-8e-f8-0a-68-00
  Partner Oper Key 0010
 Number of Ports: 2
 Member ports:
   Link: Eth 2/4 (0x18820016) sync: 1   *
   Link: Eth 2/18 (0x18890084) sync: 1
```

3. Use the **show port-channel** *number* command to display detailed information of a specific port-channel interface

```
device# show port-channel 10
Port-channel 10 is admin down, line protocol is down (admin down)
Hardware is AGGREGATE, address is 00e0.0c70.cc07
    Current address is 00e0.0c70.cc07
Interface index (ifindex) is 671088650
Minimum number of links to bring Port-channel up is 1
MTU 1548 bytes
LineSpeed Actual     : Nil
Allowed Member Speed : 10000 Mbit
Priority Tag disable
Forward LACP PDU: Enable
Last clearing of show interface counters: 00:29:09
Queueing strategy: fifo
Receive Statistics:
    0 packets, 0 bytes
    Unicasts: 0, Multicasts: 0, Broadcasts: 0
    64-byte pkts: 0, Over 64-byte pkts: 0, Over 127-byte pkts: 0
    Over 255-byte pkts: 0, Over 511-byte pkts: 0, Over 1023-byte pkts: 0
    Over 1518-byte pkts(Jumbo): 0
    Runts: 0, Jabbers: 0, CRC: 0, Overruns: 0
    Errors: 0, Discards: 0
Transmit Statistics:
    0 packets, 0 bytes
    Unicasts: 0, Multicasts: 0, Broadcasts: 0
    Underruns: 0
    Errors: 0, Discards: 0
Rate info:
    Input 0.000000 Mbits/sec, 0 packets/sec, 0.00% of line-rate
    Output 0.000000 Mbits/sec, 0 packets/sec, 0.00% of line-rate
Time since last interface status change: 00:29:09
```

Displaying LACP system-id information

Follow this procedure to display LACP system ID and priority information.

Enter the **show lacp sys-id** command to display LACP information for the system ID and priority.

```
device# show lacp sys-id
System ID: 0x8000,76-8e-f8-0a-98-00
```

Displaying LACP statistics

Follow this procedure to display LACP statistics for a port-channel interface or for all port-channel interfaces.

Before displaying the LACP port-channel information, it is recommended that you create a port-channel interface in a LAG to generate details.

Enter the **show lacp counters** command to display LACP statistics for a port-channel.

```
device# show lacp counter
Traffic statistics
Port            LACPDUs            Marker           Pckt err     Sent     Recv     Sent     Recv
Sent     Recv
Aggregator Po 3  Eth 1/6                          110      0           0        0
0        0
```

Clearing LACP counter statistics on a LAG

This topic describes how to clear LACP counter statistics on a single LAG.

To clear LACP counter statistics on a LAG, use the following command:

Enter the **clear lacp** *LAG_group_number* **counters** command to clear the LACP counter statistics for the specified LAG group number.

```
device# clear lacp 42 counters
```

Clearing LACP counter statistics on all LAG groups

This topic describes how to clear the LACP counter statistics for all LAG groups.

To clear LACP counter statistics on all LAG groups, use the following command:

Enter the **clear lacp counter** command to clear the LACP counter statistics for all LAG groups.

```
device# clear lacp counter
```

Troubleshooting LACP

To troubleshoot problems with your LACP configuration, use the following troubleshooting tips.

If a standard IEEE 802.1AX-based dynamic trunk is configured on a link and the link is not able to join the LAG, do the following:

- Make sure that both ends of the link are configured as **standard** for the trunk type.
- Make sure that both ends of the link are *not* configured for **passive** mode. They must be configured as **active** /**active**, **active** / **passive**, or **passive** /**active**.
- Make sure that the port-channel interface is in the administrative "up" state by ensuring that the **no shutdown** command was entered on the interface on both ends of the link.
- Make sure that the links that are part of the LAG are connected to the same neighboring switch.
- Make sure that the system ID of the switches connected by the link is unique. You can verify this by entering the **show lacp sys-id** command on both switches.
- Make sure that LACPDUs are being received and transmitted on both ends of the link and that there are no error PDUs. You can verify this by entering the **show lacp counters** *number* command and looking at the receive mode (rx) and transmit mode (tx) statistics. The statistics should be incrementing and should not be at zero or a fixed value. If the PDU rx count is not incrementing, check the interface for possible CRC errors by entering the **show interface** *link-name* command on the neighboring switch. If the PDU tx count is not incrementing, check the operational status of the link by entering the **show interface** *link-name* command and verifying that the interface status is "up."

When a link has problem, the **show port-channel** command displays the following message:

```
Mux machine state: Deskew not OK.
```

If a static trunk is configured on a link and the link is not able to join the LAG, do the following:

- Make sure that both ends of the link are configured as **standard** for trunk type and verify that the mode is "on."
- Make sure that the port-channel interface is in the administrative "up" state by ensuring that the **no shutdown** command was entered on the interface on both ends of the link.

Unidirectional Link Detection

Unidirectional Link Detection overview

Unidirectional Link Detection (UDLD) monitors a link between two Extreme devices and blocks the ports on both ends of the link if there is a unidirectional failure.

UDLD protocol detects and blocks broken unidirectional links in the network. This is done through the exchange of UDLD protocol data units (PDU) between devices on a physical link. Both ends of the link must support the same proprietary UDLD protocol to detect the unidirectional link condition.

A unidirectional link is assumed when the UDLD stops receiving UDLD PDUs from the other end of the link. The device then blocks the physical link. The physical link will still be up but the line protocol will be down. UDLD PDUs continue to be transmitted and received on the link.

UDLD is disabled by default. To use the UDLD protocol, the protocol must first be enabled globally and then on each individual physical port. When enabled globally and on a physical port, the device starts transmitting UDLD PDUs periodically on the port.

How UDLD works

The following shows a simple four-switch network in which two paths connect to each switch. STP blocks traffic on as many ports as necessary so that only one operational path exists from the STP root bridge to all nodes in the network.

FIGURE 1 Four-switch example for UDLD

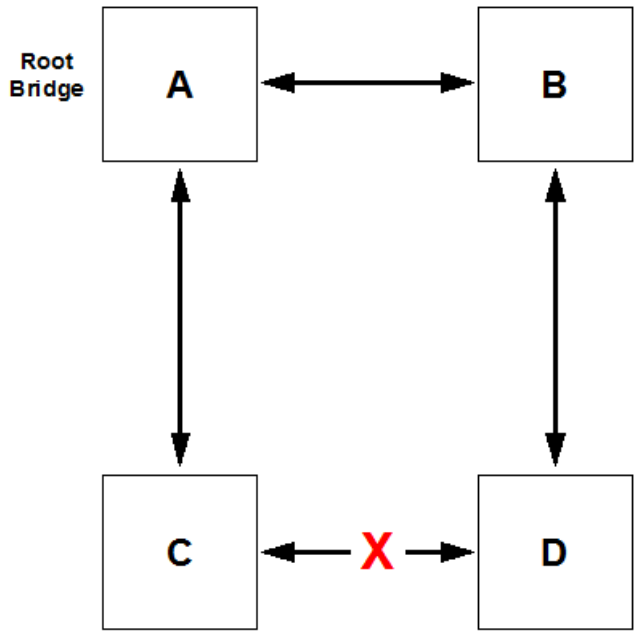

In the previous figure, STP detects that the port on Switch D that is connected to Switch C should be put into a blocked state. Therefore, no data traffic gets transmitted or received on this port. Data traffic remains blocked as long as Switch D receives bridge protocol data units (BPDUs) from both switches C and B.

If the link between Switch C and Switch D becomes unidirectional (for reasons such as hardware failure or incorrect cabling) in the direction from D to C, Switch D ages out the status that it was receiving BPDUs from Switch C. This eventually causes STP to put the port in a forwarding state, thus allowing all data traffic. This creates a loop for all BUM traffic that enters the network. BUM traffic can go from Switch B to Switch D to Switch C to Switch A, and then back to Switch B.

To prevent this loop from forming, UDLD can be used to detect that the link between Switch C and Switch D has become unidirectional.

UDLD considerations and restrictions

Note the following for UDLD:

- UDLD is used in conjunction with the Spanning Tree Protocol.
- UDLD runs only on physical ports assigned to a port channel.
- UDLD is supported on directly connected switches only.
- The protocol must be running on both ends of the link.
- The default timeout is 2.5 (2500 ms) seconds.
- Tagged UDLD is not supported.
- The feature is not compatible with Cisco UDLD protocol.
- Extreme uses a proprietary implementation that interoperates with Multi-Service IronWare devices, FastIron devices, and VDX devices.
- The UDLD interface statistics lose some accuracy after a failover.
- Upon executing the **no protocol udld** command, all of the UDLD global and UDLD interface configuration changes will be removed and the protocol will revert back to its initial disabled state. This means that all interfaces that have been enabled for the protocol stops transmitting and receiving UDLD PDUs.

Configuring UDLD

Complete the following steps to configure basic UDLD on your device.

1. Enter global configuration mode.

   ```
   device# configure terminal
   ```

2. Enable the UDLD protocol and enter protocol UDLD configuration mode.

   ```
   device(config)# protocol udld
   ```

3. Change the hello interval to 2000 milliseconds.

   ```
   device(config-udld)# hello 20
   ```

   This changes the interval at which UDLD PDUs are transmitted. The default interval, in counts of one hundred milliseconds is 500 ms.

4. Change the timeout multiplier from the default of 5.

```
device(config-udld)# multiplier 8
```

   This changes the timeout multiplier value to affect the UDLD PDU timeout interval. The UDLD timeout interval is the product of the hello time interval at the other end of the link and the timeout multiplier value.

   When the remote is Multi-Service IronWare or FastIron devices, the timeout equals local hello interval × multiplier value.

5. Return to global configuration mode.

```
device(config-udld)# exit
```

6. Enter interface configuration mode for an port.

```
device(config)# interface ethernet 5/1
```

7. Enable UDLD on the interface.

```
device(config-if-eth-5/1)# udld enable
```

8. Repeat the preceding step for each port on which you wish to enable UDLD.

   NOTE

   When the UDLD protocol is enabled on one end of a link, the timeout period might elapse before the UDLD protocol is enabled on the other end of the link. In this case, the link becomes temporarily blocked. When the UDLD protocol is enabled at the other end of the link and a UDLD PDU is received, UDLD automatically unblocks the link.

9. Return to privileged exec mode.

```
device(config-if-eth-5/1)# end
```

10. Verify the configuration.
    - Show the UDLD global configuration.

```
device# show udld
UDLD Global Information
  Admin State:       UDLD enabled
  UDLD hello time:   1000 milliseconds
  UDLD timeout:      5000 milliseconds
```

    - Show UDLD status for an interface

```
device# show udld interface ethernet 2/6
Global Admin State:  UDLD enabled

UDLD information for Ethernet 2/6
  UDLD Admin State:                  Enabled
  Interface Operational State:       Bidirectional link
  Remote hello time:                 Unknown
  Remote MAC Addr:                   0024.3890.0d81
  Local system id:  0x9f01fee0       Remote system id: 0x24900c00
  Local port :      2/6              Remote port :     9/2
  Local link id:    0x0              Remote link id:   0x0
  Last Xmt Seq Num: 43849            Last Rcv Seq Num: 43880
```

    - Show UDLD statistics.

```
device# show udld statistics
UDLD Interface statistics for Ethernet 2/7
Frames transmitted:         260
Frames received:            223
Frames discarded:           0
Frames with error:          0
Remote port id changed:     0
Remote MAC address changed: 0
```

      NOTE

      The **show interface** command also indicates whether UDLD is enabled.

11. Save the configuration.

```
device# copy running-config startup-config
```

UDLD configuration example

```
device# configure terminal
device(config)# protocol udld
device(config-udld)# hello 20
device(config-udld)# multiplier 8
device(config-udld)# exit
device(config)# interface ethernet 5/1
device(config-if-eth-5/1)# udld enable
device(config-if-eth-5/1)# end
device# show udld
device# copy running-config startup-config
```

VLANs

802.1Q VLAN overview

IEEE 802.1Q VLANs provide the capability to overlay the physical network with multiple virtual networks. VLANs allow you to isolate network traffic between virtual networks and reduce the size of administrative and broadcast domains.

A VLAN contains end stations that have a common set of requirements that are independent of physical location. You can group end stations in a VLAN even if they are not physically located in the same LAN segment. VLANs are typically associated with IP subnetworks and all the end stations in a particular IP subnet belong to the same VLAN. Traffic between VLANs must be routed. VLAN membership is configurable on a per-interface basis.

Configuring VLANs

The following sections discuss working with VLANs on Extreme devices.

Configuring a VLAN

Follow this procedure to configure a VLAN in the Extreme device at the global configuration level.

1. Enter the **configure terminal** command to access global configuration mode.

```
device# configure terminal
```

2. Enter the **vlan** command to create a topology group at the global configuration level.

```
device(config)# vlan 5
device(config-vlan-5)#
```

NOTE
The **no vlan** command removes the existing VLAN instance from the device.

Configuring a switchport interface

Follow this procedure to configure a switchport interface in the device to send and receive data packets.

1. Enter the **configure terminal** command to access global configuration mode.

```
device# configure terminal
```

2. Enter the **interface ethernet** command to configure the interface mode.

```
device(config)# interface ethernet 0/1
```

3. Enter the **switchport** command to configure a switchport interface.

```
device(conf-if-eth-0/1)# switchport
```

Configuring the switchport interface mode

Do the following to set the switchport interface as access or trunk. This configuration works only when the interface is set as switchport.

1. Enter the **configure terminal** command to access global configuration mode.

```
device# configure terminal
```

2. Enter the **interface ethernet** command to configure the interface mode.

```
device(config)# interface ethernet 0/1
```

3. Enter the **switchport** command to set the interface as switchport.

```
device(conf-if-eth-0/1)# switchport
```

4. Enter the **switchport mode** command to configure the switchport interface in trunk mode.

```
device(conf-if-eth-0/1)# switchport mode trunk
```

NOTE
The default mode is access. Enter the **switchport mode access** command to set the mode as *access*.

NOTE
Before you change the switch port mode from **switchport mode access** with an explicit **switchport access vlan** to **switchport mode trunk-no-default-native**, you must enter the **no switchport** command on the interface level, and then enter the **switchport** command to set the interface as a switchport. Now you can configure the **switchport mode trunk-no-default-native** command.

Configuring the switchport access VLAN type

Do the following to change the switchport access VLAN type. This configuration works only when the interface is set as switchport.

Ensure that reserved VLANs are not used. Use the **no switchport access vlan** command to set the default VLAN as the access VLAN.

1. Enter the **configure terminal** command to access global configuration mode.

```
device# configure terminal
```

2. Enter the **interface ethernet** command to specify an Ethernet interface.

```
device(config)# interface ethernet 0/1
```

3. Enter the **switchport** command to set the interface as switchport.

```
device(conf-if-eth-0/1)# switchport
```

4. Enter the **switchport access vlan** command to set the mode of the interface to *access* and specify a VLAN.

```
device(conf-if-eth-0/1)# switchport access vlan 10
```

This example sets the mode of a specific port-channel interface to *trunk*.

```
device# configure terminal
device(config)# interface port-channel 35
device(config-port-channel-35)# switchport mode trunk
```

Configuring a VLAN in trunk mode

Do the following to add or remove VLANs on a Layer 2 interface in trunk mode. The configuration is also used to configure the VLANs to send and receive data packets.

Ensure that reserved VLANs are not used.

1. Enter the **configure terminal** command to access global configuration mode.

```
device# configure terminal
```

2. Enter the **interface ethernet** command to specify an Ethernet interface.

```
device(config)# interface ethernet 0/1
```

3. Enter the **switchport** command to set the interface as switchport.

```
device(conf-if-eth-0/1)# switchport
```

4. Enter the **switchport trunk allowed vlan** command to set the mode of the interface to *trunk* and add a VLAN.

```
device(conf-if-eth-0/1)# switchport trunk allowed vlan add 5
```

The example sets the mode of the Ethernet interface to *trunk*.

```
device# configure terminal
device(config)# interface ethernet 0/1
device(conf-if-eth-0/1)# switchport mode trunk
```

The example sets the mode of a port-channel interface to *trunk* and allows all VLANs.

```
device# configure terminal
device(config)# interface port-channel 35
device(config-Port-channel-35)# switchport trunk allowed vlan all
```

Configuring a native VLAN on a trunk port

Do the following to set native VLAN characteristics on a trunk port for classifying the untagged traffic data packets.

Ensure that reserved VLANs are not used.

1. Enter the **configure terminal** command to access global configuration mode.

```
device# configure terminal
device(config)#
```

2. Enter the **interface ethernet** command to configure the interface mode.

```
device(config)# interface ethernet 0/1
```

3. Enter the **switchport** command to set the interface as switchport.

```
device(conf-if-eth-0/1)# switchport
```

4. Enter the **switchport trunk native-vlan** command to set native VLAN characteristics to *access* and specify a VLAN.

```
device(conf-if-eth-0/1)# switchport trunk native-vlan 300
```

This example removes the configured native VLAN on the Ethernet interface.

```
device# configure terminal
device(config)# interface ethernet 0/1
device(conf-if-eth-0/1)# no switchport trunk native-vlan 300
```

Enabling VLAN tagging for native traffic

Do the following to enable tagging for native traffic on a specific interface.

Ensure that reserved VLANs are not used.

1. Enter the **configure terminal** command to access global configuration mode.

```
device# configure terminal
```

2. Enter the **interface ethernet** command to configure the interface mode.

```
device(config)# interface ethernet 0/1
```

3. Enter the **switchport** command to set the interface as switchport.

```
device(conf-if-eth-0/1)# switchport
```

4. Enter the **switchport trunk tag native-vlan** command to enable tagging for native traffic data VLAN characteristics on a specific interface.

```
device(conf-if-eth-0/1)# switchport trunk tag native-vlan
```

This example enables tagging for native traffic data on a specific Ethernet interface.

```
device# configure terminal
device(config)# interface ethernet 0/1
device(conf-if-eth-0/1)# switchport trunk tag native-vlan
```

This example disables the native VLAN tagging on a port-channel.

```
device# configure terminal
device(config)# interface port-channel 35
device(config-Port-channel-35)# no switchport trunk tag native
```

Displaying the status of a switchport interface

Do the following to display detailed Layer 2 information for all switchport interfaces.

Enter the **show interface switchport** to display the detailed Layer 2 information for all interfaces.

```
device# show interface switchport
 Interface name          : Eth 0/1
 Switchport mode         : access
 Ingress filter          : enable
 Acceptable frame types : all
 Default Vlan            : 1
 Active Vlans            : 1
 Inactive Vlans          : -
 Interface name          : Port-channel 5
 Switchport mode         : access
 Ingress filter          : enable
 Acceptable frame types : all
 Default Vlan            : 1
 Active Vlans            : 1
```

Displaying the switchport interface type

Do the following to display detailed Layer 2 information for a specific interface.

Enter the **show interface switchport** to display the detailed Layer 2 information for a specific interface.

```
device# show interface ethernet 0/1 switchport
 Interface name           : ethernet 0/1
 Switchport mode          : trunk
 Fcoeport enabled         : no
 Ingress filter           : enable
 Acceptable frame types   : vlan-tagged only
 Native Vlan              : 1
 Active Vlans             : 1,5-10
 Inactive Vlans           : -
```

The example displays the detailed Layer 2 information for a port-channel interface.

```
device# show interface port-channel 5 switchport
 Interface name           : Port-channel 5
 Switchport mode          : access
 Fcoeport enabled         : no
 Ingress filter           : enable
 Acceptable frame types   : vlan-untagged only
 Default Vlan             : 1
 Active Vlans             : 1
 Inactive Vlans           : -
```

Verifying a switchport interface running configuration

Do the following to display the running configuration information for the Layer 2 properties for a specific interface.

Enter the **show running-config interface** to display the running configuration information for a specific interface.

```
device# show running-config interface ethernet 0/1 switchport
interface interface Eth 0/1
 switchport
 switchport mode trunk
 switchport trunk allowed vlan add 5-10
 switchport trunk tag native-vlan
```

This example displays the running configuration information for a port-channel interface.

```
device# show running-config interface port-channel 5 switchport
interface Port-channel 5
 switchport
 switchport mode access
 switchport access vlan 1
```

Displaying VLAN information

Do the following to display information about a specific VLAN.

Enter the **show vlan** to display information about VLAN 1.

```
device# show vlan 1
VLAN Name State Ports
(u)-Untagged, (t)-Tagged
(c)-Converged
================================================================
1 default ACTIVE Eth 0/1(t) Eth 0/4(t) Eth 0/5(t) Eth 0/8(t)
```

Enabling Layer 3 routing for VLANs

Do the following to enable Layer 3 routing on a VLAN.

1. Enter global configuration mode.

```
device# configure terminal
```

2. Create a VLAN.

```
device(config)# vlan 200
```

3. Create a virtual Ethernet (VE), assign an IP address and mask, and enable the interface.

```
device(config)# interface ve 200
device(config-Ve-200)# ip address 10.2.2.1/24
device(config-Ve-200)# no shutdown
```

A VE interface can exist without a VLAN configuration, but it must be provisioned in the VLAN in order to be used.

4. Enter the **router-interface** command and specify the VLAN.

```
device(config-vlan-200)# router-interface ve 200
```

VLAN statistics

Devices gather statistics for all ports and port channels on configured VLANs.

Use the **statistics** command in the VLAN configuration mode to enable statistics on a VLAN.

NOTE

Statistics has to be manually enabled for a specific VLAN, since it is not enabled by default for VLANs.

Please note that:

- The statistics reported are not real-time statistics since they depend upon the load on the system.
- Statistics has to be manually enabled for a specific VLAN. This ensures better utilization of the statistics resources in the hardware.
- Statistics for VLANs with VE interfaces consider only the switched frames. Packets which are routed into or out of the VE interface are not counted.
- Enabling statistics on a VLAN has a heavy impact on the data traffic.

Enabling statistics on a VLAN

Follow this procedure to enable statistics on a VLAN.

1. Enter the global configuration mode.

```
device# configure terminal
```

2. Enter the **vlan** command to specify a VLAN for statistics collection.

```
device(config)# vlan 5
device(config-vlan-5)#
```

3. Enter the **statistics** command to enable statistics for all ports and port channels on configured VLANs.

```
device(config-vlan-5)# statistics
```

NOTE

Use the **no statistics** command to disable statistics on VLANs.

```
device(config-vlan-5)# no statistics
```

Displaying statistics for VLANs

Do the following to display statistics information for VLANs.

Enter the **show statistics vlan** command to view the statistics for all ports and port channels on all configured VLANs.

```
device# show statistics vlan

Vlan 10 Statistics
Interface     RxPkts           RxBytes          TxPkts           TxBytes
eth 0/1       821729           821729           95940360         95940360
eth 0/2       884484           885855           95969584         95484555
po  1        8884             8855              9684             9955

Vlan 20 Statistics
Interface     RxPkts           RxBytes          TxPkts           TxBytes
eth 0/6       821729           821729           95940360         95940360
eth 0/21     8884             8855              9684             9955
po  2      884484           885855           95969584         95484555
```

TABLE 4 Output descriptions of the show statistics vlan command

Field	Description
Interface	The interface whose counter statistics are displayed.
RxPkts	The number of packets received at the specified port.
RxBytes	The number of bytes received at the specified port.

TABLE 4 Output descriptions of the show statistics vlan command (continued)

Field	Description
TxPkts	The number of packets transmitted from the specified port.
TxBytes	The number of bytes transmitted from the specified port.

Displaying VLAN statistics for a specific VLAN

Enter the **show statistics vlan** *vlan ID* command to view the statistics for a specific VLAN. Here *vlan ID* is the specific VLAN ID.

```
device# show statistics vlan 10

Vlan 10 Statistics
Interface     RxPkts          RxBytes         TxPkts          TxBytes
eth 0/1       821729          821729          95940360        95940360
eth 0/2       884484          885855          95969584        95484555
po  1      8884            8855            9684            9955
```

Clearing statistics on VLANs

Follow the procedure to clear statistics' information for VLANs.

Enter the **clear statistics vlan** command to clear the statistics for all ports and port channels on all configured VLANs.

```
device# clear statistics vlan
```

Clearing statistics for a specific VLAN

Enter the **clear statistics vlan** *vlan ID* command to clear the statistics for a specific VLAN. Here *vlan ID* is the specific VLAN ID.

```
device# clear statistics vlan 10
```

VE route-only mode

By default, physical ports and port-channels (LAG ports) support both Layer 2 switching and Layer 3 routing. VE route-only mode enables these ports to act exclusively as a Layer 3 virtual Ethernet (VE) interface, dropping all ingress packets that require Layer 2 switching.

The MAC learning of dropped packets is not affected by this feature. ARP requests (broadcast), LACP, and BPDU packet processing are also not affected. The following table lists the effects of this mode on a variety of features.

TABLE 5 Effects of VE route-only mode on features

Feature	Ingress port as route-only port (incoming frames)	Egress port as route-only port (outgoing frames)
Packets requiring switching	Drop, learn MAC address	Forwarded/switched
Packets requiring routing	Forwarded	Forwarded
ARP requests	Trapped/punted, ARP response generated	Forwarded
LACP packets	Trapped/punted and processed	Forwarded
STP/BPDU packets	Trapped/punted and processed	Forwarded

Note the following considerations and limitations:

- Egress packets through a port configured as route-only are transmitted irrespective of whether they are switched or routed.

- This feature is enabled on the active management module (MM), and is available on the other MM after a failover.
- The number of TCAM entries required for this feature depends on the maximum number of physical ports. On a there are approximately 40 physical ports per PPE , and a maximum of 512 LAG ports for the entire system. Therefore the maximum number of TCAM entries for route-only support is 40. These entries are available on all TCAM profiles.

Configuring VE route-only mode on a physical port

Do the following to configure VE route-only mode on a physical port.

1. Enter global configuration mode.

```
device# configure terminal
```

2. Create a VLAN.

```
device(config)# vlan 100
```

3. Specify an Ethernet interface.

```
device(config)# interface ethernet 1/2
```

4. Enter the **switchport** command to configure Layer 2 characteristics.

```
device(conf-if-eth-1/2)# switchport
```

5. Specify trunk mode.

```
device(conf-if-eth-1/2)# switchport mode trunk
```

6. Tag the port to a VLAN.

```
device(conf-if-eth-1/2)# switchport mode trunk allowed vlan add 100
```

7. Enter the **route-only** command to enable Layer 3 routing exclusively on the port.

```
device(conf-if-eth-1/2)# route-only
```

   Use the **no route-only** command to revert to default Layer 2 and Layer 3 behavior.

8. Enable the interface and exit to global configuration mode.

```
device(conf-if-eth-1/2)# no shutdown
device(conf-if-eth-1/2)# exit
```

9. Verify the Ethernet configuration.

```
device(conf-if-eth-1/2)# do show running-cocnfig interface ethernet 1/2
switchport
switchport mode trunk
switchport trunk allowed vlan add 100
route only
no shutdown
```

10. Verify the port statistics for switching packets dropped.

```
device(conf-if-eth-1/2)# do show interface ethernet 1/2
               Ethernet 1/2 is up, line protocol is up (connected)
               Hardware is Ethernet, address is 768e.f80a.033c
               Current address is 768e.f80a.033c
            ….
          Rate info:
               Input 0.001008 Mbits/sec, 0 packets/sec, 0.00% of line-rate
               Output 0.000252 Mbits/sec, 0 packets/sec, 0.00% of line-rate
               Route-Only Packets Dropped: 17
               Time since last interface status change: 21:35:41
```

11. Enter virtual Ethernet (VE) configuration mode and specify the VLAN.

```
device(config)# interface ve 100
```

12. Assign an IP address and mask and enable the interface.

```
device(config-Ve-100)# ip address 10.2.2.2/24
device(config-Ve-100)# no shutdown
```

13. Confirm the VE configuration.

```
device(config-Ve-100)# do show running-config interface ve 100
interface Ve 100
 ip proxy-arp
 ip address 10.2.2.2/24
 no shutdown
```

Configuring VE route-only mode on a LAG port

Do the following to configure VE route-only mode on a LAG (port-channel) port.

1. Enter global configuration mode.

```
device# configure terminal
```

2. Create a VLAN.

```
device(config)# vlan 100
```

3. Specify a port-channel interface.

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

4. Enter the **switchport** command to configure Layer 2 characteristics.

```
device(config-Port-channel-1)# switchport
```

5. Specify trunk mode.

```
device(config-Port-channel-1)# switchport mode trunk
```

6. Tag the port to a VLAN.

```
device(config-Port-channel-1)# switchport trunk allowed vlan add 100
```

7. Enable tagging on native VLAN traffic.

```
device(config-Port-channel-1)# switchport trunk tag native-vlan
```

8. Enter the **route-only** command top enable Layer 3 routing exclusively on the port.

```
device(config-Port-channel-1)# route-only
```

   Use the **no route-only** command to revert to default Layer 2 and Layer 3 behavior.

9. Enable the interface and exit to global configuration mode.

```
device(config-Port-channel-1)# no shutdown
device(config-Port-channel-1)# exit
```

10. Verify the port-channel configuration.

```
device(config-Port-channel-1)# do show running-cocnfig interface port-channel 1
 interface Port-channel 1
 switchport
 switchport mode trunk
 switchport trunk allowed vlan add 100,200
 switchport trunk tag native-vlan
 route-only
 no shutdown
```

11. Enter virtual Ethernet (VE) configuration mode and specify the VLAN.

```
device(config)# interface ve 100
```

12. Assign an IP address and mask and enable the interface.

```
device(config-Ve-100)# ip address 10.2.2.2/24
device(config-Ve-100)# no shutdown
```

13. Verify the VE configuration.

```
device(config-Ve-100)# )# do show running interface ve 100
interface Ve 100
 ip proxy-arp
 ip address 10.2.2.2/24
 no shutdown
```

Configuring virtual routing interfaces

The Extreme device sends Layer 3 traffic at Layer 2 within a protocol-based VLAN. However, Layer 3 traffic from one protocol-based VLAN to another must be routed. If you want the device to be able to send Layer 3 traffic from one protocol-based VLAN to another on the same device, you must configure a virtual routing interface on each protocol-based VLAN, then configure routing parameters on the virtual routing interfaces.

A *virtual routing interface* is a logical routing interface that the Extreme device uses to route Layer 3 protocol traffic between protocol-based VLANs. It is a logical port on which you can configure Layer 3 routing parameters.

For example, to enable a Extreme device to route IP traffic from one IP protocol VLAN to another, you must configure a virtual routing interface on each IP protocol VLAN, then configure the appropriate IP routing parameters on each of the virtual routing interfaces.

To attach a router interface to a VLAN, using the **router-interface** command:

```
device# configure terminal
device(config)# vlan 2
device(config-vlan-2)# router-interface ve 2
```

NOTE
Only one router VE interface can be mapped to a VLAN. The VLAN ID and the VE ID need not be the same.

Use the **no router interface ve** command to remove the router VE interface.

VXLAN Layer 2 Gateways

VXLAN Layer 2 gateways overview

SLX-OS devices can act as Layer 2 gateways.

The following figure illustrates an example Layer 2 gateway topology.

FIGURE 2 Layer 2 gateway topology

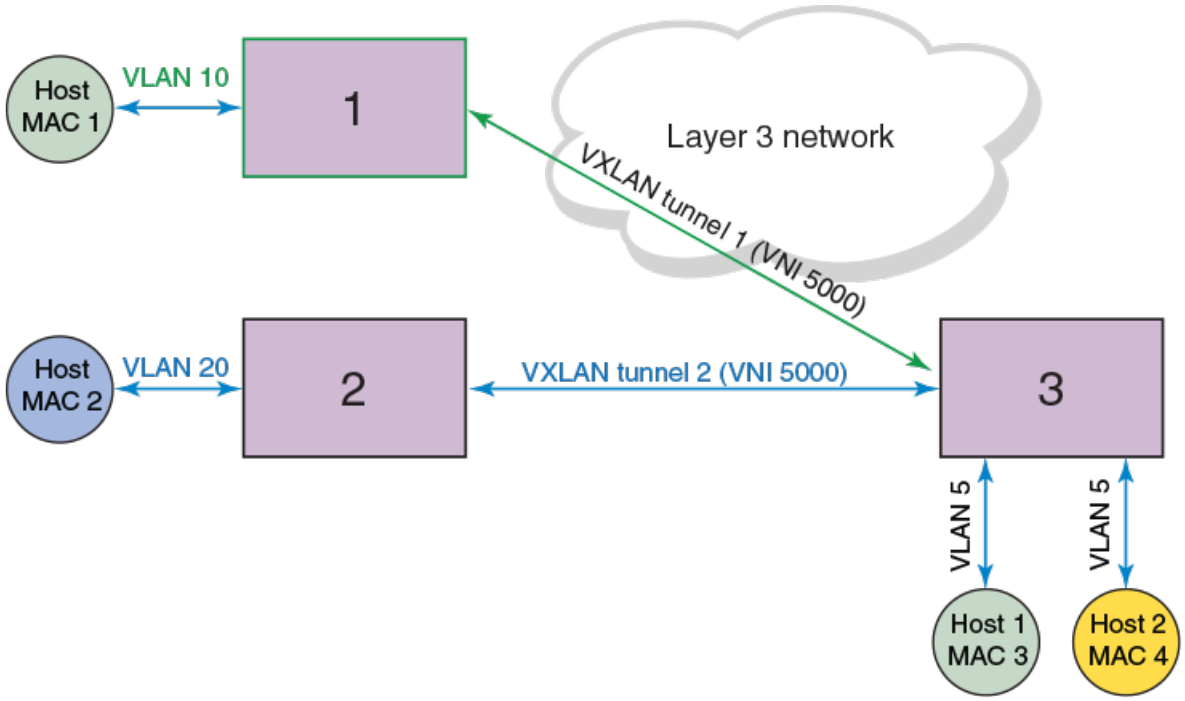

Device	IP address	Mapping
1: SLX 9850 series	1.1.1.1	VNI 5000 < > VLAN 10
2: VDX 6740 series	2.2.2.2	VNI 5000 < > VLAN 20
3: SLX 9540 series	3.3.3.3	VNI 5000 < > VLAN 5

VLANs on each node are extended through common Virtual Network Instance (VNI) 5000. The MAC addresses of local hosts are learned on access points, and MAC addresses are learned on VXLAN tunnels. A split-horizon topology is supported. All nodes participating in the VLAN must be connected through VXLAN tunnels, because there is no tunnel-to-tunnel flooding of broadcast, unknown unicast, and multicast (BUM) traffic.

In this topology, Devices 1, 2, and 3 are VXLAN Layer 2 gateway devices. On Device 3, tunnel 1 and tunnel 2 are mapped to VLAN 5. VLAN 5 has two hosts, MAC 3 and MAC 4. Device 3 is connected to two other hosts, Device 1 and Device 2, which connect to hosts MAC 1 and MAC 2, respectively, through VXLAN tunnels 1 and 2, respectively. If MAC 3 needs to establish traffic to MAC 1, initially there will be BUM flooding and upon a response from MAC 1, MAC 1 is learned through tunnel 1. Subsequent traffic goes directly from MAC 3 to Device 1 on tunnel 1. Traffic in the reverse direction comes from Device 1, is decapsulated, and goes to MAC 3.

VXLAN Layer 2 gateways considerations and limitations

Note the following considerations and limitations for VXLAN Layer 2 gateways.

- Up to 50 tunnels are supported.
- A maximum of 8 ECMP paths are supported.
- Layer 2 snooping is not supported.
- VRRPe source IP addresses and EVPN Multi-Chassis Trunks (MCTs) are not supported.
- QoS, TTL, and MTU values are not configurable. The MTU is based on the IP interface MTU. If a packet is bigger than the IP interace MTU minus the VXLAN header, then the packet is dropped. The default TTL value is 255. The default QoS value is 0, which is applied to DSCP field of the IP header.
- The maximum number of MAC addresses across all VXLAN tunnels in a node is 4000.
- VXLAN tunnels have the standard UDP header encapsulated with the standard defined value of 4789. This value is not configurable. SLX-OS expects VXLAN tunnel packets to be received with this value.
- Only VXLAN extended TCAM profiles are supported.
- VXLAN tunnels are not supported when the counter profile 1 or 4 is configured. These profiles do not allocate hardware resources for TX statistics, which is needed for VXLAN tunnels.
- The tunnel TX bytes statistics do not account for the outer VLAN header size.
- When BUM packets are flooded from one tunnel to another, they are expected to be dropped. However, the TX statistics counter on the outbound tunnel increments.

Configuring VXLAN Layer 2 gateways

Follow these basic steps to configure a VXLAN Layer 2 gateway.

1. Enter global configuration mode.

```
device# configure terminal
```

2. Enter the **overlay-gateway** command, specify the name of a gateway, and enter VXLAN overlay gateway configuration mode.

```
device(config)# overlay-gateway GW1
```

3. Enter the **type** command and specify **l2-extension**.

```
device(config-overlay-gw-GW1)# type l2-extension
```

4. Enter the **map vlan vni** command and specify **l2-extension**.

```
device(config-overlay-gw-GW1)# map vlan 5 vni 5000
```

5. Enter the **map bridge-domain** command and specify a bridge domain and VNI.

```
device(config-overlay-gw-GW1)# map bridge-domain 1 vni 2000
```

6. Enter the **ip interface** command and specify a loopback ID.

```
device(config-overlay-gw-GW1)# ip interface loopback 1
```

7. Enter the **site** command, specify a site name, and enter VXLAN overlay gateway site configuration mode.

```
device(config-overlay-gw-GW1)# site mysite1
```

This mode configures a VXLAN tunnel to the site node.

8. Enter the **ip address** command and specify an IP address.

```
device(config-overlay-gw-GW1-site-mysite1)# ip address 1.1.1.1
```

9. Enter the **extend vlan add** command and specify a VLAN.

```
device(config-overlay-gw-GW1-site-mysite1)# extend vlan add 5
```

10. Enter the **extend bridge-domain add** command and specify a bridge domain.

```
device(config-overlay-gw-GW1-site-mysite1)# extend bridge-domain add 1
```

11. Enter the **activate** command to activate the site.

```
device(config-overlay-gw-GW1-site-mysite1)# activate
```

12. In privileged EXEC mode, enter the **show overlay-gateway** command to confirm the gateway configuration.

```
device# show overlay-gateway
Overlay Gateway "GW1", ID 1,
Admin state up
IP address 3.3.3.3 (loopback 1), Vrfdefault-vrf
Number of tunnels 1
Packet count: RX 17909 TX 1247
Byte count : RX (500125) TX 356626
```

13. In privileged EXEC mode, enter the **show tunnel** command to confirm the tunnel configuration.

```
device# show tunnel
Tunnel 101, mode VXLAN
Ifindex0x7c400065, Admin state up, Operstate up
Source IP 3.3.3.3, Vrf: default-vrf
Destination IP 1.1.1.1
Active next hops on node 1:
IP: 4.4.4.5, Vrf: default-vrf
Egress L3 port: Ve45, Outer SMAC: 609c.9f5a.4415
Outer DMAC: 609c.9f5a.0015, ctag: 0
BUM forwarder: yes
```

14. In privileged EXEC mode, enter the **show vlan** command to confirm the VLAN configuration.

```
device# show vlan 5
VLAN             Name            State                      Ports           Classification
(R)-RSPAN                                                   (u)-Untagged
                                                            (t)-Tagged
================ =============== ========================== =============== ====================
5                VLAN05          ACTIVE                     Eth 2/1(t)
                                                            Eth 2/5(t)
                                                            tu61441  vni5000
```

15. In privileged EXEC mode, enter the **show mac** command to confirm the MAC configuration.

```
device# sow mac
VlanId/BDId    Mac-address       Type      State      Ports/LIF/PW
35 (V)         609c.9f5a.5b15    Dynamic   Active     Po 35
45 (V)         609c.9f5a.4415    Dynamic   Active     Po 45
5 (V)          0000.0400.0011    Dynamic   Active     tu61441
5 (V)          0000.0500.0011    Dynamic   Active     Eth 0/5
5 (V)          0000.0400.0011    Dynamic   Active     tu61441
5 (V)          0000.0500.0011    Dynamic   Active     Eth 0/5
Total MAC addresses : 6
SLX#
```

VXLAN Layer 2 gateway support for bridge domains

The SLX-OS device provides VXLAN Layer 2 gateway support for bridge domains in addition to supporting Layer 2 VLANS. VXLAN gateway support to bridge domains allows a maximum of 4K bridge-domain Virtual Network Interface (VNI) mappings along with a maximum of 4K VLAN VNI mappings, for a a total of 8K mappings.

Since a bridge domain supports different port and VLAN endpoints, all of its traffic can be extended to a remote node using one VNI.

Also, VXLAN gateway support to bridge domains enables VLAN translation of traffic on both sides of the network. The local VLANs can use different VLAN tags on either side of the network and map to the same VNI.

NOTE
Only point-to-multipoint bridge domain types are supported to extend over VXLAN tunnels. Point-to-point bridge-domain type is not supported.

You can extend the bridge domain under a site configuration. You can configure the bridge domain to VNI mapping automatically with auto mode where the bridge domain to the VNI is mapped implicitly. For example,VLAN 1 through 4096 is mapped to VNI 1 through 4096 and the bridge domain is mapped to 4097. You can also configure the bridge domain to a VNI map manually similar to a VLAN.

NOTE
The default tagging mode for a bridge domain is RAW mode.

Configuring VXLAN Layer 2 Gateway support for bridge domains

Before performing this configuration, configure a p2mp bridge domain.

The following steps configure a VXLAN Layer 2 gateway to support bridge domains.

1. Enter global configuration mode.

```
device# configure terminal
```

2. Create a VXLAN overlay gateway and access the overlay gateway configuration mode.

```
device(config)# overlay-gateway gateway1
device(config-overlay-gw-gateway1)#
```

3. Configure the Layer 2 type extension.

```
device(config-overlay-gw-gateway1)# type layer2-extension
```

4. Specify a loopback interface.

```
device(config-overlay-gw-gateway1)# ip interface loopback 1
```

5. Enable the mapping of a bridge domain to a VNI.

```
device(config-overlay-gw-gateway1)# map bridge-domain 1 vni 999
```

6. Create a remote Layer 2 extension site in a VXLAN overlay gateway and access site configuration mode.

```
device(config-overlay-gw-gateway1)# site bd1
```

7. Specify the destination IPv4 address of a tunnel.

```
device(config-site-bd1)# ip address 10.67.67.1
```

8. Configure a bridge domain for the tunnel to the site.

```
device(config-site-bd1)# extend bridge-domain add 1
```

9. Exit site configuration mode.

```
device(config-site-bd1)# exit
```

10. Activate the gateway.

```
device(config-overlay-gateway-gateway1)# activate
```

The previous steps are provided in the following configuration example.

```
device# configure terminal
device(config)# overlay-gateway gateway1
device(config-overlay-gw-gateway1)# type layer2-extension
device(config-overlay-gw-gateway1)# ip interface loopback 1
device(config-overlay-gw-gateway1)# map bridge-domain 1 vni 999
device(config-overlay-gw-gateway1)# site bd1
device(config-site-bd1)# ip address 10.67.67.1
device(config-site-bd1)# extend bridge-domain add 1
device(config-site-bd1)# exit
device(config-overlay-gateway-gateway1)# activate
```

VXLAN Layer 2 gateway payload tag processing

The SLX-OS device provides the following modes for the processing of the payload tag that is received on the attachment-circuit packets.

- VXLAN RFC compliant mode
- Enhanced payload tag transport mode

VXLAN RFC compliant mode

In RFC compliant mode, the VLAN tag in the packet is not carried in the packet and need to be stripped at the ingress device before sending the VXLAN encapsulated packet into the network.

To configure RFC compliant mode, configure the bridge domain in raw mode as shown in the following example.

```
pw-profile test
    vc-mode raw

bridge-domain 10 p2mp
```

```
   pw-profile test
!
```

Then extend the bridge domain in the overlay gateway, as shown in the following example.

```
overlay-gateway gateway1
type layer2-extension
ip interface loopback 1
map bridge-domain 10 vni 999
site vcs1
 ip address 10.67.67.1
 extend bridge-domain add 10
!
```

NOTE

The SLX-OS device supports RFC compliant mode with the bridge-domain based VXLAN service only.

Enhanced payload tag transport mode

In enhanced payload tag transport mode, one VLAN tag from the traffic is carried as part of the VXLAN encapsulated packet as an inner payload tag. This tag can carry the PCP value to include the priority information and also can interoperate with other devices. This tag is removed in the remote device capable of this behavior.

NOTE

This mode is not interoperable with RFC compliant mode.

This mode is supported for VLAN-based VXLAN service and bridge-domain based VXLAN service with tag mode.

To configure enhanced payload tag transport mode, configure the bridge domain in tagging mode as shown in the following example.

```
pw-profile test
    vc-mode tag

bridge-domain 10 p2mp
   pw-profile test
!
```

Then extend the bridge domain in the overlay gateway, as shown in the following example.

```
overlay-gateway gateway1
type layer2-extension
ip interface Loopback 1
map bridge-domain 10 vni 999
site vcs1
 ip address 10.67.67.1
 extend bridge-domain add 10
!
```

Multi-Chassis Trunking (MCT)

MCT Overview

Multi-Chassis Trunking (MCT) is trunking that initiates at a single MCT-unaware server or switch and terminates at two MCT-aware switches. MCT allows the links to the two MCT-aware switches to appear to a downstream device as if they are coming from a single device on a single Link Aggregation (LAG) trunk interface or physical port.

NOTE
The SLX-OS device does not support Layer 2 protocols over MCT. You must provide a loop-free topology.

NOTE
MCT does not support any variant of Spanning Tree Protocol (STP). STP is disabled by default and must not be enabled with MCT.

In a datacenter network environment, LAG trunks provide link level redundancy and increased capacity. However, they do not provide switch-level redundancy. If the switch connected to the LAG trunk fails, the entire trunk loses network connectivity.

With MCT, member links of the LAG trunk are connected to two MCT-aware devices. A configuration between the devices enable data flow and control messages between them to establish a logical Inter-Chassis Link (ICL). In this model, if one MCT device fails, a data path remains through the other device.

SLX-OS Layer 2 MCT is based on RFC 7432 (BGP MPLS-Based Ethernet VPN). The MP-BGP EVPN extension is the control plane on the SLX-OS device to perform both MAC synchronization and cluster management. MAC synchronization between the MCT peers synchronizes the MAC tables between the MCT nodes for node resiliency and faster convergence.

For the data plane, the SLX-OS device supports the MPLS forwarding mechanism to leverage MPLS LER functionality.

SLX-OS MCT provides Layer 3 protocol support for IPv4 and IPv6 BGP, OSPF, and IS-IS through a VE interface. The VE over MCT interface MAC address is synchronized to the MCT peer through BGP using an EVPN MAC route.

SLX-OS supports MCT over VPLS or VLL Layer 2 VPN services, and IPv4 or IPv6 Virtual Routing Redundancy Protocol (VRRP) and VRRP Extended (VRRP-E).

MCT terminology

Before implementing MCT in your network, you must understand some key terms and definitions.

MCT peer devices	A pair of SLX-OS device configured as peers. The LAG interface is spread across two MCT peer devices and acts as the single logical endpoint to the MCT client.

NOTE
MCT is supported across the same chassis type only; for example, SLX-9850 <---> SLX-9850.

MCT client	The MCT client is the device that connects with the MCT peer devices. It can be a switch or an endpoint server host in the single-level MCT topology or another pair of MCT switches in a multi-tier MCT topology.
MP-BGP EVPN extension	The control plane for Layer 2 MCT on the SLX-OS device.
MCT Cluster Client Edge Port (CCEP)	A port on one of the MCT peer devices that is a member of the LAG interface to the MCT client. To have a running MCT instance, at least one Link Aggregation Interface is needed with a member port on each peer device. While there is a LAG on the client device, CCEP on the MCT device can be a LAG or a physical port.
MCT Cluster Edge Port (CEP)	A port on MCT peer device that a member of a MCT VLAN and is not a Cluster Client Edge Port.
MCT VLANs	VLANs on which MCT clients are operating. These VLANs are explicitly configured in the MCT configuration.

SLX-OS MCT control plane

Multiprotocol-BGP (MP-BGP) EVPN extension, as specified in RFC 7432, is used as the SLX-OS MCT control plane.

The control plane consist of the following components:

- EVPN instance—Mapped to a Layer 2 VLAN that the RFC refers to the VLAN-Based service interface.
- Ethernet segment ID (ESI)—10-byte integer that uniquely identifies the set of links connecting MCT PEs to the client CE. SLX-OS MCT supports both dynamic and static LAG between MCT PE and CE and uses ESI type 0 that is encoded as follows:
  - 1-byte ESI type = 0
  - 9 byte ESI value = user-input through the SLX-OS **esi** command
- MP-BGP route distinguisher (RD)—Encoded using RD type 1 as defined in RFC 4364 that consists of the following subfields:
  - 4-byte administrator subfield that is set with the 4-byte router ID
  - 2-byte assigned number subfield that is encoded with the all zeros for the Ethernet segment (ES) route, client ID for the Ethernet auto-discovery route, and EVPN ID (VLAN ID) for MAC and multicast routes.
- MP-BGP EVPN capability—When a BGP session is brought up to a MCT peer, BGP indicates to the peer that it is EVPN capable using BGP capability advertisement with the following information:
  - Capability code = 1 (MP-BGP)
  - AFI = 25 (L2VPN)
  - SAFI = 70 (EVPN)

  If a BGP session already existed to the same peer, the existing BGP session is flapped to allow the advertisement of the EVPN capability.
- MP-BGP EVPN route types—Includes Ethernet Auto-Discovery (A-D), MAC/IP Advertisement, Inclusive Multicast Ethernet Tag, and Ethernet Segment routes.
- MPLS Label Assignment—Statically assigned label ranges for ESI label, EVPN unicast label and EVPN BUM label.
  - The ESI label is generated based on the start ESI label value plus the client ID.
  - EVPN unicast or BUM label is generated based on the start unicast or BUM label plus the VLAN ID.
- Designated forwarder—A PE in a set of multi-homing PEs connected to the same Ethernet segment that is elected for sending BUM traffic to a client for a VLAN ID on an Ethernet segment.

SLX-OS MCT operates in dual-homing mode.

MP-BGP EVPN Routes

RFC 7432 defines EVPN Network Layer Reachability Information (NLRI) with the format as shown in the following figure:

FIGURE 3 EVPN NLRI format

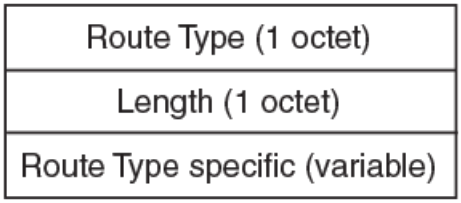

SLX-OS MCT supports the following route types.

TABLE 6 SLX-OS MCT route types

Route type	Route name	SLX-OS usage
1	Ethernet Auto-Discovery (per Ethernet Segment only)	Mass MAC withdraw and Designated forwarder election. The PE advertises one Ethernet A-D route for each client interface. When the client interface goes down, the PE withdraws the Ethernet A-D route which is served as a trigger for the remote PE to remove all MAC addresses learned over the affected client interface instead of withdrawing an individual MAC.
2	MAC/IP Advertisement	MAC synchronization of the MAC addresses between two MCT peers.
3	Inclusive Multicast Ethernet Tag	Advertisement of ingress replication usage and multicast label expected for each EVPN instance when receiving BUM traffic.
4	Ethernet Segment (ES)	Designated forwarder election. The ES route is used to update the remote peer when the MCT client is deployed and undeployed.

Designated forwarder election

Designated forwarder election is triggered in the following scenarios:

- A client is deployed locally or remotely.
- BGP session comes up.
- CCEP goes up or down.

To elect a designated forwarder (DF) for a VLAN ID on an Ethernet segment, each PE exchanges its router IP address with its multi-homing PEs through the Ethernet Segment route. The following algorithm uses the IP address to select the DF.

1. Upon the discovery of a new ES, a PE advertises the ES route and waits a default of three seconds for its peer to advertise the ES route.
2. When the timer expires, the PE builds a sorted list of PE IP addresses including its own address connected to the same ES.
3. The PE with the ordinal number that equals (V mod N) is elected the DF. V is the VLAN ID and N is the number of PEs.
4. When the ES link fails, the PE withdraws its ES route which triggers the selection process to select a new DF. When a PE node failure occurs, DF election is also triggered when the PE is up and down.

   NOTE
   DF election is triggered only after the client is deployed on both MCT devices.

SLX-OS MCT data plane traffic

For the discussion of the MCT data plane traffic, refer to the following topology diagram and configuration.

FIGURE 4 EVPN MPLS forwarding topology

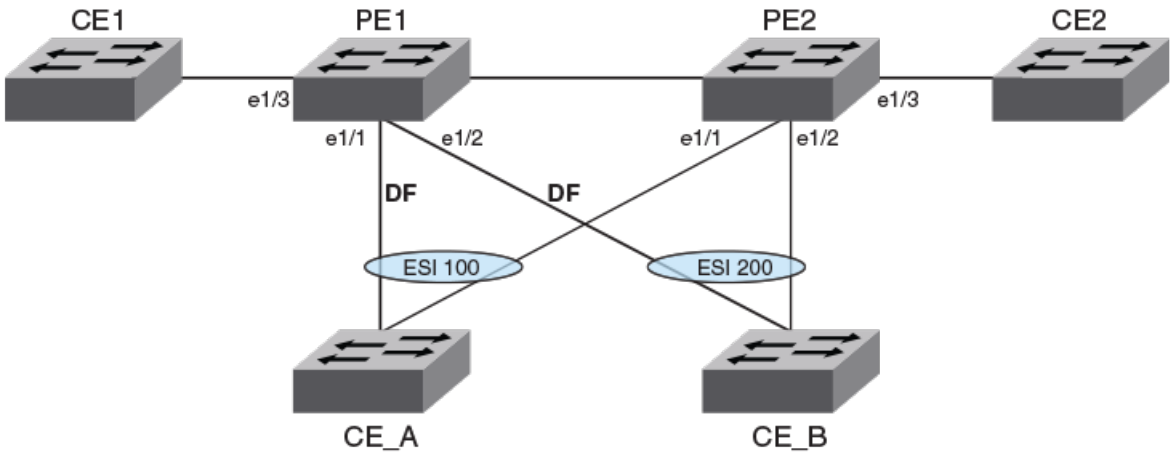

- Cluster VLAN 1000 is configured on both PE1 and PE2.
- Cluster VLAN 1000 has client interface e1/1 and e1/2 that belong to ESI 100 and 200 respectively and a cluster end point on e1/3.
- PE1 is elected as Designated Forwarder (DF) for VLAN 1000 for both ESI 100 and 200.
- VLAN 1000 is mapped to VSI/EVPN instance 1000.

Unicast traffic sent between CE1 and CE_A or CE_B follows normal L2 forwarding.

Forwarding unicast traffic between PEs

To send unicast traffic received over a CEP or CCEPs on PE1 to the remote PE2, PE1 pushes an previously-received PE2 EVPN label. The following figure shows the packet encapsulation.

FIGURE 5 EVPN unicast forwarding between cluster PEs

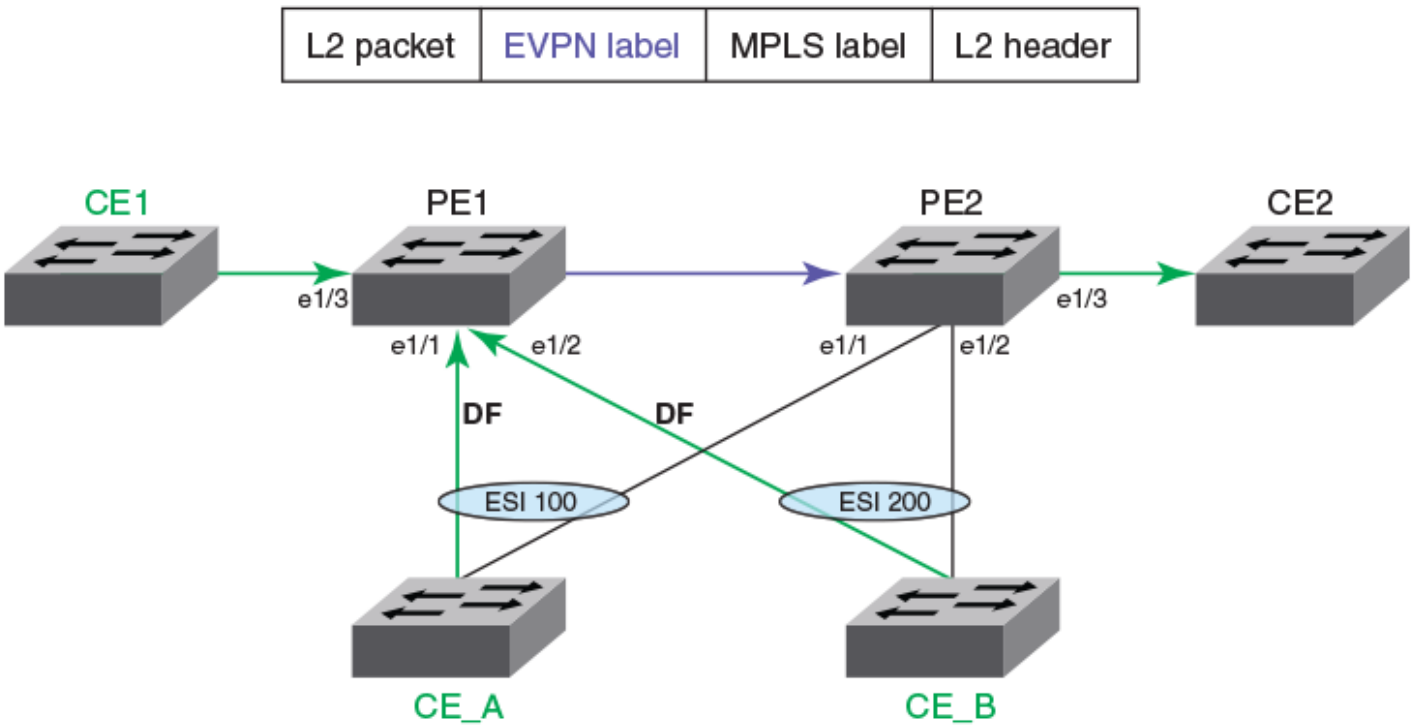

For an SLX-OS device, it advertises one label for all the MACs learned within a bridge domain.

> **NOTE**
> This traffic forwarding is also applicable for MCT VPLS and MCT VLL. For MCT VLL, there are no CEP ports. There is no EVPN label. Only the MPLS label is involved for communication between cluster peers. For each bridge domain, the forwarding rules are defined based on the DF role.

Flooding traffic received from a CEP

For unicast traffic received over a CEP that PE1 needs to flood on VLAN 1000, PE1 sends a copy to e1/1, e1/2, and to PE2. The copy sent on e1/1 and e1/2 has basic L2 encapsulation as in the case of a normal L2 VLAN. The copy sent to PE2 has a multicast label previously received from PE2 encapsulated as shown in the following figure.

FIGURE 6 EVPN BUM forwarding between cluster PEs

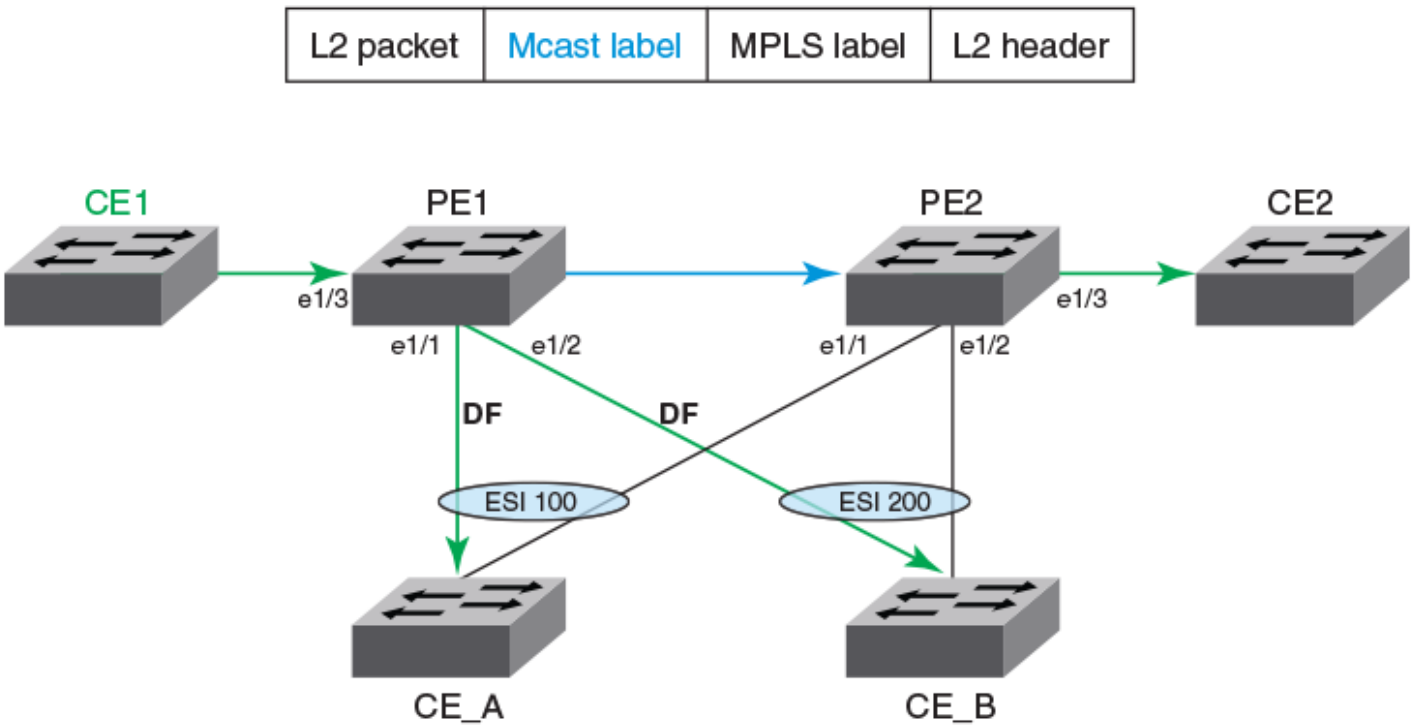

For SLX-OS device, it advertises one multicast label for each EVPN instance.

Flooding traffic received from CCEP

For traffic received over a CCEP from CE_A on PE1 to be flooded to remote PE2, PE1 must push previously received multicast and ESI labels from PE2. The following figure shows the packet encapsulation.

FIGURE 7 EVPN source port suppression based on DF election

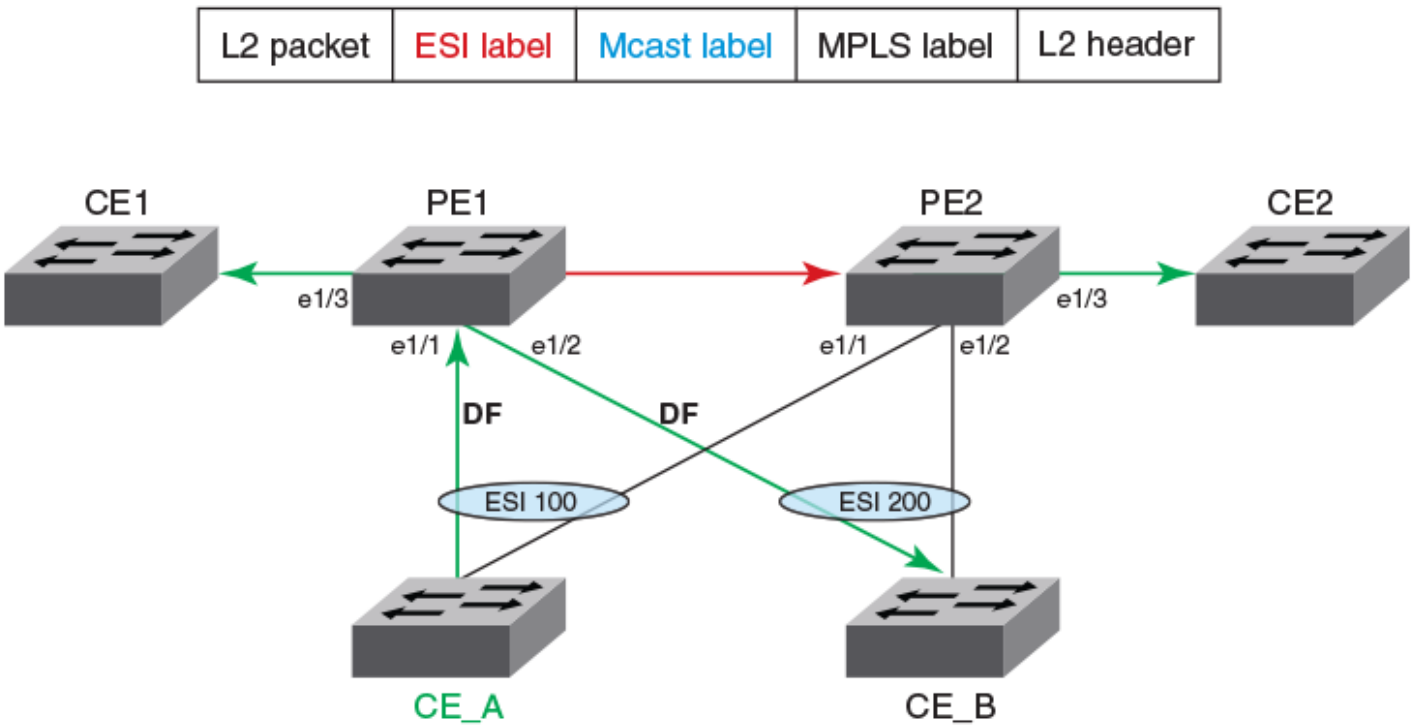

Since source port suppression is based on the DF election result, PE2 suppresses copies to be sent to CE_A and CE_B because e1/1 and e1/2 are not elected as DF for ESI 100 and 200 respectively.

In the following figure, source port suppression is done based on the ESI label carried in the packet itself.

FIGURE 8 EVPN source port suppression based on the ESI label

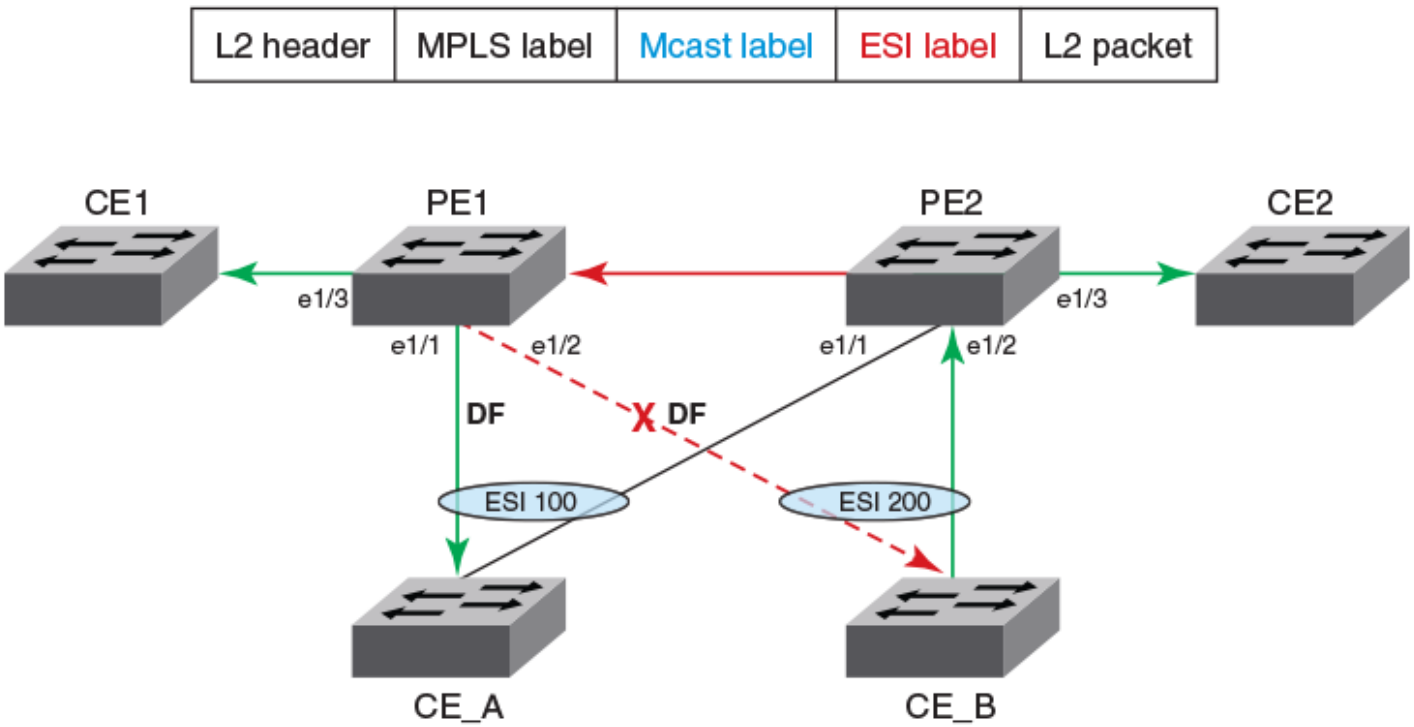

For traffic received from CE_B on PE2 that is flooded to PE1, PE1 sends a copy to CE1 and CE_A on e1/1. However, it suppresses a copy to be sent to CE_B. The suppression is done because the copy received from PE2 carried an ESI label that PE1 previously advertised to PE2 for ESI 200.

MAC management

Each MAC entry maintains information about all the owners of the MAC entry who learned and advertised it. The information about each owner is maintained in the form of the MAC database (MDB).

A MDB entry contains peer and client information with the cost. A local MAC entry has cost 0 and MAC addresses learned from an MCT peer has a cost of 1. The MDB entry with the lowest cost is chosen as the filtering database (FDB).

In MCT topologies, MAC addresses can be learned either on local CCEP or CEP interfaces, or from the remote MCT node. If the same MAC is learned from both MCT nodes, then MAC entries learned locally have higher priority than the one learned from the remote peer.

For remote MAC addresses, aging is disabled, and they can only be deleted when delete notifications are received from the remote node that advertised it before for learning.

The following terminologies are associated with MCT MAC entries.

- Cluster Local (CL)—MAC addresses learned locally on CEP ports
- Cluster Remote (CR)—MAC addresses learned on remote CEP ports
- Cluster Client Local (CCL)—MAC addresses learned locally on a client interface
- Cluster Client Remote (CCR)—MAC addresses learned on a remote client interface

Static MAC handling

Configuration of static MAC entries is allowed over MCT enabled VLANs and CEP and CCEP interfaces.

The MCT static MAC addresses configured on a local node are advertised to remote MCT node for learning. While advertising the MAC using the BGP MAC advertisement route, it uses the MAC mobility extended-community route to identify the MAC as static using the sticky MAC field. On the remote node, when MAC advertisement is received for a static MAC address, the sticky MAC information is saved along with the MAC entry.

When an MCT static MAC address is deleted, a BGP MAC withdrawal route is sent to the remote peer to delete the MAC entry from its database.

When a CEP interface is down and if any static MAC entries are present, MAC Delete messages are sent to the remote node to flush the entries.

When a CCEP interface is down and if any static MAC entries are associated with the client, the MAC addresses are moved to point to the remote MCT peer. The MAC addresses are moved back to the CCEP when the interface comes back up.

On a local MCT node, when a cluster is UP and you configure a static MAC on a CEP or CCEP interface, the node synchronizes the MAC address to the remote MCT node. The remote node processes the MAC address and adds it to the FDB. On the remote MCT node, you can configure the same MAC as the static MAC address for the client 1 CCEP interface since it is configured on the same client CCEP interface. No additional static MAC configurations on the remote node are required since the same MAC are already part of the local MCT node.

When the cluster is down on the local and remote MCT nodes, both nodes are independent as clusters that can be independently configured with the static MAC addresses for the CEP or CCEP interface. However, when the cluster is brought up, the static MAC addresses are synchronized from both nodes and the addresses on the remote node are rejected since the local configuration takes precedence. The misconfiguration remains until you correct it.

MAC learning

MAC learning over CEP interfaces is like basic Layer 2 learning and the EVPN MAC advertisement route is sent to the cluster peer to synchronize the learned MAC address. MAC learning over CCEP interfaces is two-step process in which the MAC entry is added into the MDB first. The best MAC entry is chosen and installed into the FDB and the EVPN MAC advertisement route is sent to the cluster peer for synchronization of the learned MAC address.

The following rules are used for MAC learning:

- If a static MAC address is configured on the CEP port, it is learned as the CL and the EVPN MAC AD route message is sent to the peer. In this case, the ESI is set to NULL. On the peer MCT node, the CR MAC address programmed on the cluster peer is static towards the MCT peer.
- If a static MAC address is configured on the CCEP port, an EVPN MAC advertisement route is sent to the peer. In this case, the MAC entry is associated with the ESI of the MCT client. The peer MCT node programs the MAC address as static over the local CCEP interface.
- Dynamic MAC learning from the CEP is similar to basic Layer 2 MAC learning. An EVPN MAC advertisement route is sent to the peer. In this case, the ESI is set to invalid or NULL. On the peer MCT node, the CR MAC address programmed on the cluster peer is static towards the MCT peer.
- Dynamic MAC learning from the CCEP occurs as a CCL MAC. An EVPN MAC advertisement route is sent to the peer. In this case, the MAC entry is associated with the ESI of the cluster client. The peer MCT node programs the MAC address as static over the local CCEP interface.

MAC aging rules

The following rules are defined for MAC aging:

- The local MAC age over CEP interface is similar to the Layer 2 MAC age. After the local MAC delete, an EVPN MAC withdrawal route is sent to the MCT peer.
- The local MAC age over CCEP interface is considered aged only if all MCT nodes age out the entry. When the MAC that ages on one of the MCT node local MDB is deleted, if the remote MDB present MAC is reprogrammed as the CCR, else the MAC is removed from the local FDB, an EVPN MAC withdrawal route is sent to MCT peer.
- The remote MAC addresses of the CR and CCR that are learned through the EVPN MAC advertisement route does not age out. They can only be removed by the EVPN MAC withdraw messages from the peer.

MAC movement

A MAC address is considered to be moved when the same MAC address is received on a different interface with same VLAN. In MCT, a MAC movement is allowed on both local and remote nodes.

The following table describes the allowed MAC movements in MCT.

TABLE 7 MCT MAC movement

MAC movement scenario	Behavior
Local dynamic MAC move from CEP1 to the CEP2 edge interface on MCT1.	On local node MCT1, the MAC address is updated to point to the new CEP2 interface. There is no MAC route update required to the remote MCT node. As on the remote node, the MAC always point towards the MCT peer for all CR MAC addresses.
Local dynamic MAC move from CEP1 edge interface to the CCEP1 client interface on MCT1.	On local node MCT1, the MAC address is updated to point to the new client interface CCEP1. A MAC update route is sent with the new ESI of client 1. The remote node updates the MAC address to point to the CCEP of client 1.
CCEP1 interface (client 1) to CCEP2 interface (client 2) on MCT1.	On local node MCT1, the MAC address is updated to point to the new client interface CCEP2. A MAC update is sent with the new ESI of client 2 to the remote node. The remote node updates the MAC address to point to the CCEP of client 2.
Local dynamic MAC move from CCEP1 interface (client 1) to CEP1 edge interface on MCT1.	On local node MCT1, the MAC address is updated to point to the new edge interface CEP1. A MAC update is sent with the new ESI 0 to the remote node. The remote node MCT2 updates the MAC address pointing to the MCT1 node.
For a MAC learned on a CEP port locally (MCT1). Dynamic MAC move to a CEP port on the remote node (MCT2)	On the MCT2 node for the CR MAC learned from MCT1, it is considered as a MAC move when it is learned on a CEP port. The MAC is updated as local on MCT2 and now points to the CL on the CEP port on MCT2 instead of pointing to MCT1 node MCT2 sends an updated MAC to MCT1. MCT1 updates the MAC as remote and points to the MCT2 .
For a MAC learned on a CEP port locally (MCT1). Dynamic MAC move to CCEP1 on MCT2.	On the MCT2 node for the CR MAC learned from MCT1, it is considered as a MAC move when the same MAC is learned on a CCEP1 port. The MAC is updated as CCL on MCT2 and now points to the local CCEP1 port on MCT2 instead of pointing to the MCT1 node. MCT2 sends a CCL MAC updated to MCT1. MCT1 updates the MAC as CCR and point to the CCEP1 port.
For a MAC CCL learned on a CCEP1 port locally (MCT1). Dynamic MAC move to CCEP2 on remote MCT2 node.	On the MCT2 node for the CCR MAC learned from MCT1 for client 1, it is considered as a MAC move when the same MAC is learned on client 2 over the CCEP2 port. The MAC is updated as CCL on MCT2 and now points to the local CCEP2 port on MCT2 instead of pointing to CCEP1. From MCT2, it sends a CCL MAC updated to MCT1. MCT1 updates the MAC as CCR and points to the CCEP2 port.

TABLE 7 MCT MAC movement (continued)

MAC movement scenario	Behavior
For a MAC CCL learned on a CCEP1 port locally (MCT1). Dynamic MAC move to CEP port on remote MCT2 node.	On the MCT2 node for the CCR MAC learned from MCT1 for client1, it is considered as a MAC move when the same MAC is learned on the CEP port. The MAC is updated as CL on MCT2 and points to the local CEP port on MCT2 instead of pointing to CCEP1. From MCT2, it sends a CL MAC updated to MCT1. MCT1 updates the MAC as CR and points to the MCT2.

MAC address deletion

The following rules are defined for MAC address deletion. Note that every deletion triggers the MAC resolution algorithm and reprograms the MAC entry if required.

- If the CEP interface is down, MAC addresses are deleted locally and individual MAC deletion messages are sent to the peer.
- If the CCEP local port is down and the remote CCEP is down, MAC addresses are deleted locally and the ESI withdraw message is sent to the MCT peer instead of sending individual MAC delete messages.
- If the CCEP local port is down and the remote CCEP is up, all local MAC addresses are moved to point to the remote MCT peer including the static MAC addresses associated with the CCEP.
- When the client entry is undeployed, all MAC addresses are deleted locally, and the ESI withdraw message is sent to the MCT peer to delete all associated client MAC addresses.

Automatic RSVP LSP bring up for MCT

The automatic bring up of an RSVP LSP for MCT functionality allows MPLS to automatically create and delete an RSVP LSP when you deploy or undeploy the MCT cluster with a peer-interface configuration. This functionality handles MM failover and the MPLS process restart for automatically-created LSPs.

NOTE
This LSP is a non-CSPF LSP.

Automatic LSP creation

When you deploy the MCT cluster, MPLS receives the cluster peer interface and a peer address for the LSP creation. The peer interface can be any Layer 3 interface (physical or VE). The peer address is the peer router ID.

Upon receiving the cluster deploy update, MPLS is enabled on the peer interface and the RSVP LSP is created with a default template with a unique auto-generated name, *MCT_peer-ipaddr_peer-interface-index*. A random number suffix is generated if this name is already in use in the user configurations. You can see the LSP by using the **show mpls lsp** command.

Automatic LSP deletion

When the cluster is undeployed, MCT updates the MPLS daemon to delete the automatically-created RSVP LSP. Then the MPLS daemon deletes the corresponding LSP. MPLS is disabled on the peer interface if it is not MPLS enabled from the CLI configuration.

Configuration considerations

- MCT does not support any variant of Spanning Tree Protocol (STP). STP is disabled by default and must not be enabled with MCT.
- The SLX-OS device does not assume that the MCT peers are directly connected.

- The cluster peer address is required to be the peer router ID and must be configured on both MCT nodes.
- On both MCT nodes, you must configure the same client ID.
- Since the The SLX-OS device supports both dynamic and static LAG between the MCT PE and CE, it uses Ethernet segment identifier (ESI) type 0 regardless of the LAG type.
  - You configure the 9-byte ESI value that is used to form a 10-byte integer globally unique ESI.
  - You must configure the same 9-byte ESI value for each client on both MCT devices.
- You must configure and activate a BGP EVPN neighbor as the peer interface.
- Since EVPN-based L2 MCT uses the MPLS forwarding plane, you must either configure a peer interface or configure MPLS to bring up either an LDP or RSVP LSP.
- The SLX-OS device uses the peer interface to setup the MPLS data path automatically. The LSP is internally created or removed on the cluster deployment or undeployment respectively. The LSP is the outgoing peer interface and has an MCT prefix when it is displayed by the **show mpls lsp** command.

  The configured peer interface must be the best nexthop to reach the MCT peer. If not, the automatically-created LSP fails to come UP.
- The automatically-created LSP can co-exist with other explicit MPLS configurations. You can configure additional LDP configurations and RSVP LSPs.

  When co-existence occurs on the device, consider the following:
  - As part of the MCT cluster configuration, the MPLS daemon automatically creates and signals the RSVP LSP with the default template and a unique auto-generated name.

    If you create an LSP with same name as an MCT LSP name *before* the MCT LSP is created, then MCT LSP is created with a similar name but with a different number suffix to make it unique.

    Example: MCT_20.21.22.23_9876543_123456
  - You cannot create an LSP with same name as a pre-existing MCT LSP. The attempted configuration is rejected and an error message is displayed.
  - When the router port or the VE interface is MPLS enabled from the MCT configuration, you cannot disable the router port or the VE interface. The configuration to disable the router port or the VE interface from the CLI is rejected with an appropriate error message.
  - You can disable MPLS by using the **no router mpls** command. Upon command execution, the automatically-created LSP along with the MPLS-enabled peer-interface configuration are disabled and re-enabled.
- In the case of LACP, to ensure two MCT peers send the same system ID and key but a different port ID to each client, you must configure the same cluster ID and client ID on both nodes. The LACP fields will be set as follows:
  - Key = MCT_LACP_KEY_BASE (3000) + client_ID
  - Port ID (16-bit unique value) = 5-bit (slot value) + 8-bit (port value + 3-bit (MCT position offset)

Configuring the BGP EVPN peer

Create a BGP EVPN address family to configure and activate the cluster EVPN peer. This configuration is associated with the MCT cluster peer configuration.

Before configuring a BGP EVPN peer, ensure you configure a loopback interface.

1. In privileged EXEC mode, enter global configuration mode.

```
device# configure terminal
```

2. Enable BGP routing.

```
device(config-terminal)# router bgp
```

3. Configure the EVPN peer with the autonomous system number (ASN).

```
device(config-bgp-router)# neighbor 10.1.1.1 remote-as 100
```

4. Configure the EVPN peer through the loopback interface.

```
device(config-bgp-router)# neighbor 10.1.1.1 update-source loopback 1
```

5. Enter EVPN address family configuration mode

```
device(config-bgp-router)# address-family l2vpn evpn
```

6. Activate the EVPN peer.

```
device(config-bgp-evpn)# neighbor 10.1.1.1 activate
```

The following example are the steps in the previous configuration.

```
device# configure terminal
device(config-terminal)# router bgp
device(config-bgp-router)# neighbor 10.1.1.1 remote-as 100
device(config-bgp-router)# neighbor 10.1.1.1 update-source loopback 1
device(config-bgp-router)# address-family l2vpn evpn
device(config-bgp-evpn)# neighbor 10.1.1.1 activate
```

Configuring MCT

Configure the local and remote MCT cluster and cluster clients.

Before configuring MCT, ensure that the following configurations exist:

- Layer 3 interface for the cluster peer interface
- VLANs for the cluster members
- Port channel for Link Aggregation or an Ethernet interface as a client interface

Perform the following steps.

1. In privileged EXEC mode, enter global configuration mode.

```
device# configure terminal
```

2. Create a cluster on the device.

```
device (config)# cluster MCT1 1
```

3. Add the VLANs as members to the cluster.

```
device(config-cluster-1)# member vlan add 2,4-170
```

4. Configure the peer IP address.

```
device(config-cluster-1)# peer 10.1.1.1
```

   The peer IP address must be the remote peer's router ID. This address corresponds with the neighbor in BGP EVPN address family configuration for the peer.

5. Configure the peer interface.

```
device(config-cluster-1)# peer-interface Ve 10
```

   The peer interface should be a valid Layer 3 interface. You should configure the peer interface before deploying the configuration.

6. Deploy the cluster.

```
device(config-cluster-1)# deploy
```

   The MPLS LSP is created when the cluster is deployed and is removed when the cluster in undeployed.

7. Create the client for the cluster and access cluster client configuration mode.

```
device(config-cluster-1)# client MCT1-client 200
```

   On both MCT nodes, you must configure the same client ID.

8. Configure the interface to the cluster client instance.

```
device(config-cluster-client-200)# client-interface port-channel 3
```

   The port channel specifies the LAG ID.

   The client interface can also be a physical interface, for example:

```
device(config-cluster-client-200)# client-interface Ethernet 2/5
```

   The client interface cannot be added under multiple client entries.

9. Set the 9-octet Ethernet Segment ID (ESI) value which is used to uniquely identify the cluster client.

```
device(config-cluster-client-200)# esi 00.a1.b2.c3.d4.e5.f6.89.00
```

   You must configure the same value on both MCT nodes to create the MCT client LAG.

   The ESI cannot be added under multiple client entries.

10. Deploy the cluster client.

```
device(config-cluster-client-200)# deploy
```

11. After configuring the local MCT cluster and client, configure the remote MCT cluster and client.

The following example is the steps in the previous configuration.

```
device# configure terminal
device (config)# cluster MCT1 1
device(config-cluster-1)# member vlan add 2,4-170
device(config-cluster-1)# peer 10.1.1.1
device(config-cluster-1)# peer-interface Ve 10
device(config-cluster-1)# deploy
device(config-cluster-1)# client MCT1-client 200
device(config-cluster-client-200)# client-interface port-channel 3
device(config-cluster-client-200)# esi 00.a1.b2.c3.d4.e5.f6.89.00
device(config-cluster-client-200)# deploy
```

Taking the MCT node offline for maintenance

If you need to take an MCT device offline for maintenance or an upgrade, perform the following steps to minimize traffic loss.

1. Verify that the MCT node that is peer to the node being taken offline is in loose client-isolation mode.

```
device# show cluster 1
Cluster MCT1 1
==============
Cluster State: Deploy
Client Interfaces Shutdown: FALSE
Client Isolation Mode: Strict
Configured Member Vlan Range: 2, 4-7
Active Member Vlan Range: 2, 4-7
...
```

2. If the peer node is in strict client-isolation mode, configure it to loose mode.

```
device(config)# cluster MCT1 1
device(config-cluster-1)# client-isolation loose
```

3. Disable the MCT clients from the MCT node that you will take offline, as shown in the following example.

```
device(config)# cluster MCT1 1
device(config-cluster-1)# client-interfaces-shutdown
```

4. Isolate the MCT node that you will take offline from the core of the network by shutting down all uplink interfaces.

   NOTE
   Do not write the configuration changes made in the previous steps to the startup-configuration file.

To bring the MCT node back online, perform one of the following actions.

- If you upgraded or downgraded the device, select the **coldboot** option under the firmware download menu.
- For any other reason, execute the **reload system** command. Since the changed configuration was not saved, the reload reverts the configuration changes that had taken the MCT node offline.

Configuring additional MCT cluster parameters

The SLX-OS device has additional cluster commands with default values. You can change the parameters for these commands in cluster configuration mode.

Changing the client-isolation mode

Isolation mode defines the action to be taken when the BGP control session goes down between the MCT nodes while the cluster is in deployed state. When the client-isolation mode is strict, the client interface will be shutdown.

By default, client-isolation mode is based on the peer IP address. The node with the lower peer IP address is set to the client-isolation mode of loose, while the node with the higher peer IP address is set to the client-isolation mode of strict. You can override this behavior by configuring strict client-isolation mode on both nodes.

Use the **client-isolation strict** command to configure the strict mode on both nodes, as shown in the following example.

```
device(config-cluster-1)# client-isolation strict
```

Use the **client-isolation loose** command to configure the loose mode on both nodes, as shown in the following example.

```
device(config-cluster-1)# client-isolation strict
```

Changing the designated-forwarder hold timer value

Upon expiration of the designated-forwarder hold timer, the reelection of the designated forwarder is considered.

By default, the hold time is three seconds. Use the **designated-forwarder-hold-time** command to change the time in seconds from 1 to 60 seconds, as shown in the following example.

```
device(config-cluster-1)# designated-forwarder-hold-time 35
```

Moving the traffic from an MCT node to the remote node

Use the **client-interfaces-shutdown** command to move all the traffic on the node to the remote MCT node by disabling the local client interfaces administratively, as shown in the following example.

```
device(config-cluster-1)# client-interfaces-shutdown
```

Displaying MCT information

You can display detailed MCT information and related MCT MAC addresses.

To display the EVPN neighbor information, use **show ip bgp neighbors** command. This information includes the peer configured for the EVPN address family, the undeplyed MCT cluster, and the negotiation of the EVPN address family.

Displaying the cluster information

The following example shows the information of the cluster on the SLX-OS device.

```
device# show cluster 1

Cluster MCT1 1
==============
Cluster State: Deploy
Client Interfaces Shutdown: FALSE
Client Isolation Mode: Loose
Configured Member Vlan Range: 2, 4-7
Active Member Vlan Range: 2, 4-7

 Peer Info:
 ----------
 Peer IP: 10.1.1.1, State: Up
 Peer Interface: Ve 20

 PW client Info:
 ---------------

 Client Info:
 ------------
 Name     Id   ESI                           Interface          Local/Remote State
 access1  100  00.a1.b2.c3.d4.e5.f6.89.00     Ethernet 1/8         Up/UP
 access2  200  00.11.22.33.44.55.66.77.88     Port-Channel 3       Dep/UnDep
 access3  300  00.24.46.0e.cd.ab.66.16.00     Ethernet 2/3         Up/Down
```

NOTE
When you delete an IP router ID that is used as the neighbor ID and IP address on an MCT peer, the **show cluster** command on the MCT peer devices displays inconsistent cluster states.

Displaying the cluster client information

The following example displays all client information for cluster 1.

```
device# show cluster 1 client
Client Info:
------------
Name      Id    Label(Lo/Re)    Interface         Local/Remote State
access1   100   NA/ 798721      Ethernet 1/8      UnDep/Dep
access2   200   798722/798722   Port-Channel 3    Up/UP
access3   300   798723/798723   Ethernet 2/3      Down/Up
```

The following example displays client 100 information for cluster 10.

```
device# show cluster 10 client 100
Client Info:
------------
Client: access1, client-id: 100, Deployed, State: Up
Interface: Ethernet 1/8
Vlans : 1-10, 100
Elected DF for vlans: 2, 4, 8, 10, 100
```

Displaying member VLAN information

The following example displays the member VLAN information for the cluster.

```
device# show cluster member vlan
VLAN-ID    Mcast-label(Lo/Re)    Unicast-label(Lo/Re)    Forwarding State
   101     817253 / 817253       800869 / 800869              Up
   102     817254 / 817254       800870 / 800870              Up
   103     817255 / 817255       800871 / 800871              Up
   104     817256 / 817256       800872 / 800872              Up
   105     817257 / 817257       800873 / 800873              Up
   106     817258 / 817258       800874 / 800874              Up
```

Displaying and clearing the MAC address table cluster information

The following example displays the MCT cluster information in the MAC address table.

```
device# show mac-address-table cluster 1
VlanId/BDId Mac-address     Type         State   Ports/LIF/PW
102 (V)     0001.0001.0001  Dynamic-CCL  Active  Eth 1/2
102 (V)     0001.0001.0111  Static-CCL   Active  po 100
102 (V)     0001.0001.0111  CCR          Active  Eth 3/66
104 (V)     0024.387c.8f00  Dynamic-CCL  Active  Po 3
107 (V)     768e.f80b.2801  CR           Active  16.16.16.16
```

You can also view the MAC entries for a specific client.

Clearing the MCT cluster MAC table entries

You can clear all cluster entries from the MAC address table or the entries for a specified client. The following example clears the MAC entries for client 3 of cluster 1.

```
device# clear mac-address-table cluster 1 client 3
```

Only the local MAC entries are deleted from the current node. Individual MAC withdrawal flush messages are sent through the EVPN. However, BGP still batches multiple routes to the remote node.

When the remote MCT peer receives the MAC withdrawal message, it only deletes the remote MAC entry. To clear MAC addresses on both nodes, you must issue **clear mac-address-table** commands on both MCT nodes.

VPLS and VLL MCT

VPLS and VLL MCT are used for data center interconnection in which SLX-OS MCT acts as a data center gateway to connect to another data center through either the VPLS or VLL WAN connection.

> NOTE
> For more information on VPLS and VLL, refer to the "VPLS and VLL Layer 2 VPN services" chapter.

For VPLS MCT, a point-to-multipoint (p2mp) bridge domain is added to the MCT cluster. For VLL MCT, a point-to-point (p2p) bridge domain is added to the MCT cluster. The VPLS or VLL horizon is added as a pseudowire (PW) client.

VLL MCT supports PW redundancy. At any point of time, one active-active PW path exists to reach the destination. The node on which the PW is active is called the active node. The endpoint traffic coming from the standby node traverses through the MCT PW session to the active node for that instance and the active MCT node takes care of the forwarding to the remote VPLS or VLL peer.

> NOTE
> For SLX-OS, the MCT cluster requires both nodes to be on SLX-OS devices. However, an SLX-OS MCT cluster that connects to a Extreme MLX cluster through VPLS or VLL is supported.

Control plane for VPLS or VLL MCT

As with Layer 2 MCT, VPLS or VLL MCT uses MP-BGP EVPN for the control plane. However, an Ethernet segment ID (ESI) controls all pseudowires (PWs) and is encoded the same as the Layer 2 CCEP ESI. This ESI is called the PW horizon ESI.

The bridge domain is mapped to an EVPN instance. For each BD, the default EVPN ID is the BD ID plus 4,096. A user configured EVPN ID is not supported. The EVPN ID for the VLAN is the VLAN ID.

For VPLS or VLL MCT, the physical and LAG CCEP operate in active/active multi-homing mode. However, the PW operates in active/standby mode.

The designated forwarder (DF) state of the PW ESI represents the active PW node state for a VPLS or VLL instance. The DF election process for the PW ESI is the same as the Layer 2 ESI process. However, for VLL MCT in the following dynamic cases, VLL does not change its role and is driven through the PW horizon client.

- If the local endpoint is down, the remote endpoint is up.
- If the active-active PW is down on the active node, the active-active PW is up on the standby node.

PW redundancy for VLL MCT

PW redundancy for VLL MCT allows the quick failover of traffic to the backup PWs. To support active and standby PWs to remote PEs, the Preferential status bit in the PW status TLV is exchanged to indicate whether the PW forwarding is active or standby.

Status TLV support is enabled through a VLL instance if one of the following is true.

- VLL is configured with two remote peers.
- The VLL endpoint is a MCT client CCEP port.

To support PW redundancy, configure two VLL peers under one VLL instance. One PW is for each VLL peer. Among these PWs, an active-active PW is selected and used for traffic flow to the remote side. An active-active PW is selected based on the local and remote PW redundancy preferences. A remote PW redundancy preference is received by the PW status TLV. When the bit is set, it indicates PW forwarding standby. When the bit is cleared, it indicates PW forwarding active.

PW state in VPLS or VLL MCT

The PW state in VPLS or VLL MCT is controlled by two entities. The MCT module controls its MCT state. The PW remote peer provides its PW redundancy state. Together, they decide the operational (forwarding) state of the PW. The following table shows the PW state decisions.

MCT state	PW remote state	Operational state	PW signaling state
DF	Active	Active	Active
DF	Standby	Standby	Active
Non-DF	Any	Standby	Standby

When the PW is in active operational state, the data plane objects (such as LIF or MGID, or cross-connect for VLL) is created and be programmed into the hardware. When the PW is in standby operational state, the data plane is programmed as if this PW is down.

NOTE
The SLX-OS PW state table is the same as the Extreme NetIron VPLS-MCT or VLL-MCT PW state table to ensure compatibility when facing an MLX MCT cluster over a VPLS or VLL connection.

VLL-MCT data plane

A topology for VLL MCT is provided in the following figure. Two MCT clusters face each other and four PWs connect the clusters.

NOTE
This topology is for only one BD. Since the DF state is selected per BD, another BD can use a different PE as the active node.

FIGURE 9 VLL MCT topology

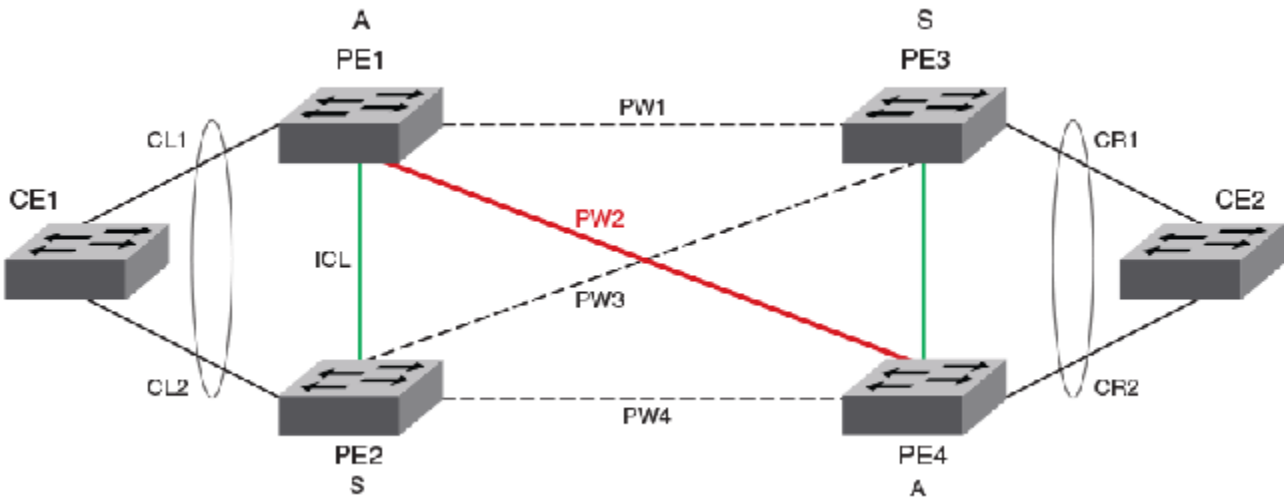

When VLL MCT is activated, only one PW is operationally active between the MCT clusters, as represented by the solid line. The standby PWs are represented by the dotted lines.

The ICL link between the MCT nodes is a BGP EVPN connection. It is not a spoke PW.

> NOTE
> VLL MCT does not use MAC learning. BUM traffic handling is not required. It uses cross-connect instead of VSI. VLL MCT does not use the EVPN label.

Steady state traffic

Based on the previous figure, the steady state traffic is as follows:

CE1->PE1->**PW2**->PE4->CE2
CE1->PE2->ICL[Split Horizon PW or ICL]->PE1->**PW2**->PE4->CE2

Client Link down protection

When the client link (CL1) is down, the device does not change the MCT status for this VLL. Traffic from the client will be received on CL2 to PE2 and forwarded using Spoke PW from PE2 to PE1. The traffic flow from the client is as follows:

CE1 -> PE2 -> [Split Horizon PW or ICL] -> PE1 -> **PW2** -> PE4 -> CE2

Active MCT Node protection

VLL MCT provides protection when one PE node has a failure including a software or hardware failure, or a power down. In the case when the active MCT node (PE1) is down, the standby MCT node acts as active and uses corresponding PWs for the traffic flow from the client. The traffic flow from the client is as follows:

CE1 -> PE2 -> **PW4** -> PE4 -> CE2

NOTE

For active-active PW link protection when the PW redundancy status changes, the device relies on the MPLS configuration. There should not be a case where PW2 is down and PW4 is up. The MPLS configuration ensure that both PW2 and PW4 are UP or DOWN. When PW2 is not active-active due to a role change on PE4, PW1 will become active-active.

VPLS-MCT data plane

The main case topology for VPLS MCT is provided in the following figure. Two MCT clusters face each other and four PWs connect the clusters.

NOTE

This topology is for only one BD. Since the DF state is selected per BD, another BD can use a different PE as the active node.

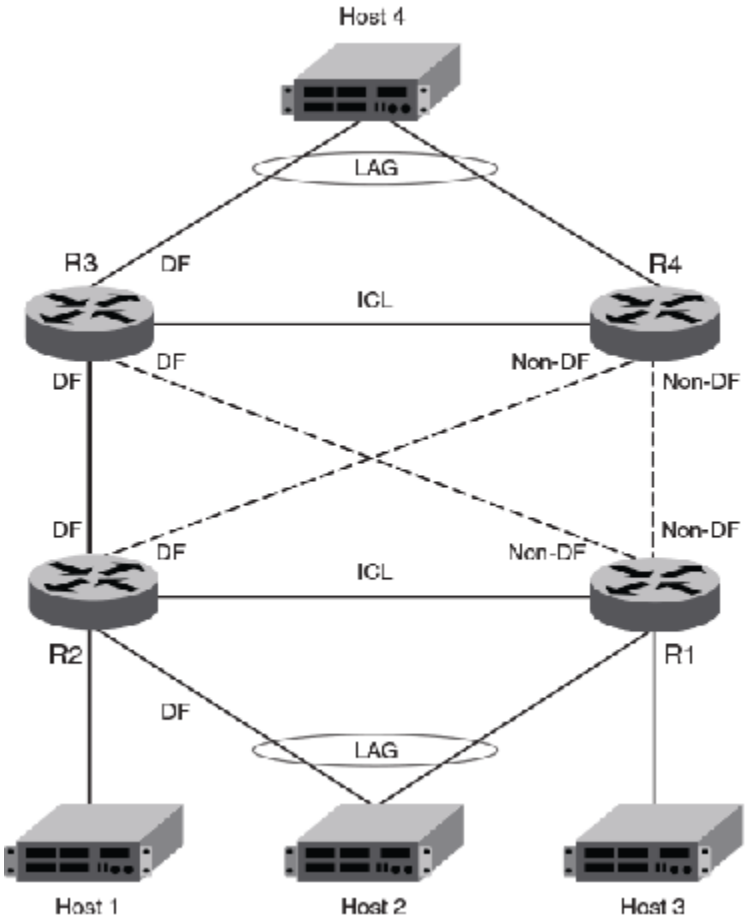

When VPLS MCT is activated, only one PW is operationally active between the MCT clusters, as represented by the solid line. The standby PWs are represented by the dotted lines.

The ICL link between the MCT nodes is a BGP EVPN connection. It is not a spoke PW.

VPLS MCT BUM traffic

The following figure illustrates how a BUM packet that starts from host 4 travels through the VPLS-MCT network to reach hosts 1, 2, and 3.

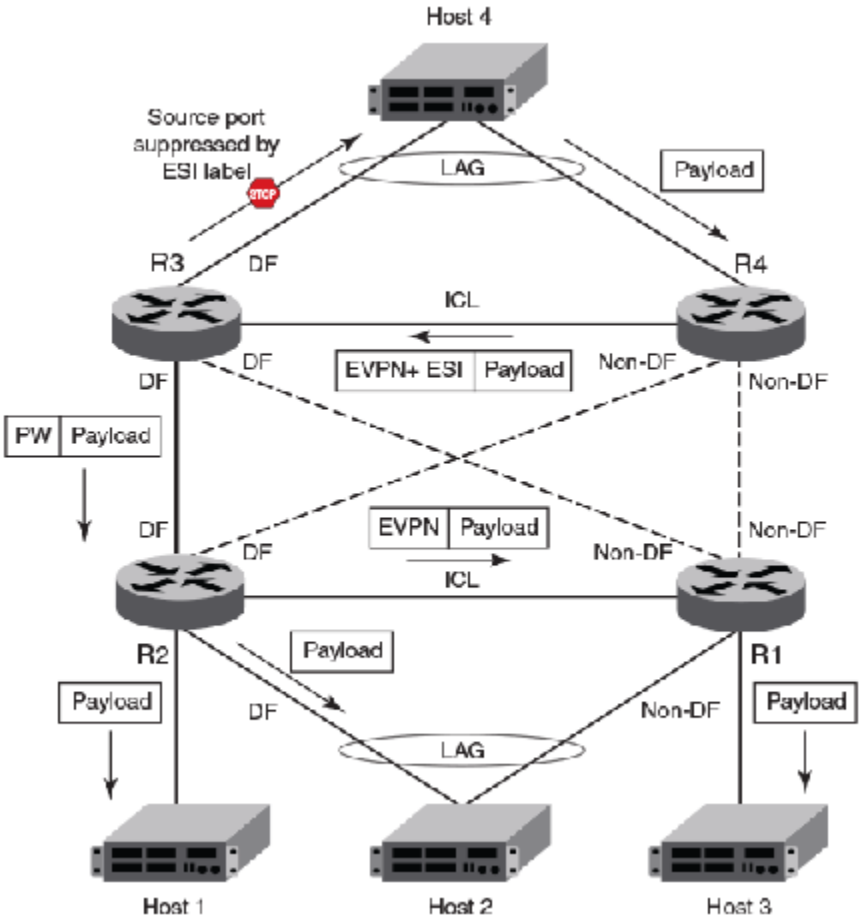

VPLS-MCT PE node protection

VPLS MCT provides protection when one PE node has a failure, including a software or hardware failure, or a power-down event. When the active PE fails in MCT, the standby PE becomes the active PE and all PWs on this node transits into MCT active state.

The following figure shows the BUM packet flow after an MCT PE switch-over event.

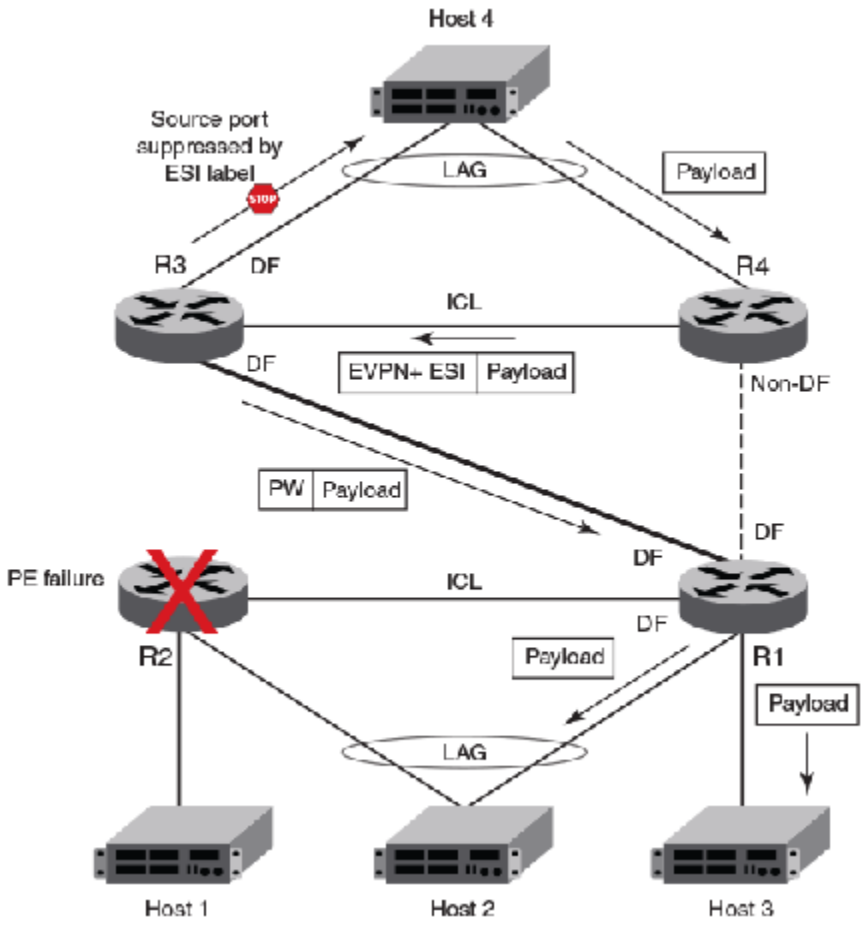

NOTE

VPLS MCT does not support PW link protection.

VPLS MAC management

In VPLS-MCT topologies, MAC addresses can be learned either on local PW or from the remote MCT node. The same MAC can be learned from either locally or remotely but not from both MCT nodes.

For MCT remote MAC addresses, aging is disabled, and they can only be deleted when delete notifications are received from the remote MCT node that advertised it before for learning.

The following terminologies are associated with MCT MAC entries.

- Cluster Local (CL)—MAC addresses learned locally on CEP ports
- Cluster Remote (CR)—MAC addresses learned on remote CEP ports
- Cluster Client Local (CCL)—MAC addresses learned locally on a client interface
- Cluster Client Remote (CCR)—MAC addresses learned on a remote client interface

VPLS MAC addresses that are learned locally are classified as CL, and the corresponding VPLS MAC addresses that are learned remotely are classified as CR.

Static MAC handling

Static MAC configuration over the local VPLS endpoints is supported. Static MAC pointing to the PW that is established with the remote VPLS PE is not supported.

MAC learning

MAC addresses learned from the PW on the active PE triggers EVPN MAC synchronization message that are sent to the peer. The PW ESI is used in this MAC route. VPLS CR MAC addresses point to the active MCT node since no local forwarding path on the standby PE traffic is expected to be switched by the active MCT node.

MAC aging

When the VPLS MAC ages on the active node, the MAC address is locally flushed and the EVPN MAC withdrawal route is sent to remote MCT node to flush the MAC.

VPLS MAC movement

A MAC address is considered to be moved when the same MAC address is received on a different interface with same VLAN. In MCT, a MAC movement is allowed on both local and remote nodes.

The following table describes the allowed VPLS MAC movements in MCT.

MAC movement scenario	Behavior
Local dynamic MAC move from PW A to PW B on MCT1.	On local node MCT1, the MAC address is updated to point to the new PW interface. There is no MAC route update required to the remote MCT node. As on the remote node, the MAC always point towards the MCT peer for all VPLS addresses.
Local dynamic MAC move from PW to the Layer 2 CCEP1 client interface on MCT1.	On local node MCT1, the MAC address is updated to point to the new client interface CCEP1. A MAC update route is sent with the new ESI of client 1. The remote node updates the MAC address to point to the CCEP of client 1.
For a MAC learned on a PW locally (MCT1). Dynamic MAC move to CCEP1 on MCT2.	On the MCT2 node for the CR MAC learned from MCT1, it is considered as a MAC move when the same mac is learned on a CCEP1 port. The MAC is updated as CCL on MCT2 and now points to the local CCEP1 port on MCT2 instead of pointing to the MCT1 (PW) node. MCT2 sends a CCL MAC update to MCT1. MCT1 updates the MAC as CCR and points to the CCEP1 port.
For a MAC CCL learned on a CCEP1 port locally (MCT1). Dynamic MAC move to the PW on the remote MCT2 node.	On the MCT2 node for the CCR MAC learned from MCT1 for client 1, it is considered as a MAC move when the same MAC is learned on the PW. The MAC is updated as the CL on MCT2 and now points to the PW on MCT2 instead of pointing to CCEP1. From MCT2, it sends a CL MAC update to MCT1. MCT1 updates the MAC to point to MCT2.
Local dynamic MAC move from PW (MCT1) to CEP1 client interface on MCT1.	On local node MCT1, the MAC address is updated to point to the new interface CEP1. A MAC advertise route is sent with ESI 0 to the remote MCT node. The remote node MCT2 updates the MAC address to point to the MCT1 node.
For a MAC CL learned on a PW locally (MCT1). Dynamic MAC move to CEP on MCT2.	On the MCT2 node for the CR MAC learned from MCT1, it is considered as a MAC move when the same mac is learned on the CEP port. The MAC is updated as CL on MCT2 and points to the local CEP1 port. From MCT2, it sends a CL MAC updated to MCT1. MCT1 updates the MAC as CR and points to the MCT2 node.
For a MAC CL learned on a CEP1 port locally (MCT1). Dynamic MAC move to PW on remote MCT2 node.	On MCT2 node for the CR MAC learned from MCT1 for CEP, it will be considered as a MAC move when the same mac is learned on the PW. The MAC is updated as CL on MCT2 and now points to the PW on MCT2.

MAC movement scenario	Behavior
	MCT2 sends a CL MAC updated to MCT1 with the ESI of the PW client, MCT1 now should updated the MAC point to the MCT2.

MAC address deletion

The following rules are defined for MAC address deletion. Every MAC deletion triggers the MAC resolution algorithm and reprograms the MAC entry if required.

- If a PW is down, MAC addresses flushed locally and individual MAC deletion messages are sent to the MCT Peer. This is similar to the Layer 2 CEP port-down handling.
- If PW client is undeploy on MCT 1, only one MAC withdraw message is send to MCT 2.

  All MAC addresses tagged to the PW client are flushed.
- If MCT 2 detects that MCT 1 is down or if the EVPN session is down, all VPLS MAC addresses that are learned from MCT 1 are flushed.

Configuration considerations and limitations for VPLS and VLL MCT

- When the same primary interface is configured as MCT cluster client interface, multiple logical interfaces belongings to same primary physical (eth/port-channel) interface cannot be the same MCT BD.
- Hitless ISSU is not supported. Before starting ISSU, issue the **client-interface shutdown** command on the PE where the ISSU is planned to gracefully move the traffic to the MCT peer. Similarly, use the **force-standby** or **no deploy** command for the PW CCEP before starting ISSU.
- Routing over EVPN is not supported.
- Statistics are not supported.
- For VPLS MCT, consider the following:
  - Configuring a cluster peer as a BD peer impacts data traffic.
  - You can use the **client-interfaces-shutdown** command to shutdown traffic on one node before a software upgrade. After you issue this command, all PWs are put into standby mode.
  - Client-interface shutdown brings down all CCEP interfaces and the BGP session. Other nodes attempt client-isolation logic after the BGP session is down, and you may see the Strict behavior.
  - Logical-interface shutdown brings down the admin state of the CCEP LIF interface if the parent port is an MCT client interface and does not trigger a DF re-election.
  - Protection for a PW link failure is not supported. Active forwarding paths does not occur between the nodes.
- For VLL MCT, consider the following:
  - Cross-connect is used instead of VSI.
  - MAC learning is not required.
  - BUM traffic handling is not required.
  - The MCT role does not depends upon endpoint as well as the PW redundancy status.

Configuring MCT for VPLS or VLL

Configuration of VPLS or VLL for MCT requires the adding of member bridge domains to the MCT cluster and a PW client.

- Before configuring VPLS MCT, configure a point-to-multipoint (p2mp) bridge domain.
- Before configuring VLL MCT, configure a point-to-point (p2p) bridge domain.

For information on configuring VLL or VPLS bridge domains, refer to the "VPLS and VLL Layer 2 VPN services" chapter.

For information on configuring the MCT cluster and client, refer to the Configuring MCT on page 63. Their full configuration is provided in the example after the steps.

Perform the following steps to configure MCT for VPLS or VLL.

1. In privileged EXEC mode, enter global configuration mode.

```
device# configure terminal
```

2. Access the cluster on the device.

```
device (config)# cluster MCT1 1
```

3. If the cluster is deployed, undeploy it.

```
device(config-cluster-1)# undeploy
```

4. Add the bridge-domains as members to the cluster.

```
device(config-cluster-1)# member bridge-domain add 1-5
```

   For VPLS, the bridge domain is p2mp. For VLL, the bridge domain is p2p.

5. Deploy the cluster.

```
device(config-cluster-1)# deploy
```

6. Create the PW client for the cluster and access PW cluster client configuration mode.

```
device(config-cluster-1)# client-pw
```

   Only one instance of the PW client represents all VPLS or VLL PWs over all bridge domains.

7. Set the 9-octet Ethernet Segment ID (ESI) value which is used to uniquely identify the PW client.

```
device(config-cluster-client-pw)# esi 01:02:03:04:05:06:07:08:0a
```

   You must configure the same value on both MCT nodes.

   The ESI cannot be added under multiple client entries.

8. Deploy the PW client.

```
device(config-cluster-client-pw)# deploy
```

9. After configuring the local MCT cluster and PW client, configure the remote MCT cluster and PW client.

The following example is the steps in the previous configuration with the additional configuration of the MCT cluster and client.

```
device# configure terminal
device (config)# cluster MCT1 1
device(config-cluster-1)# member vlan add 2,4-170
device(config-cluster-1)# member bridge-domain add 1-5
device(config-cluster-1)# peer 10.1.1.1
device(config-cluster-1)# peer-interface Ve 10
device(config-cluster-1)# deploy
device(config-cluster-1)# client MCT1-client 200
device(config-cluster-client-200)# client-interface port-channel 3
device(config-cluster-client-200)# esi 00.a1.b2.c3.d4.e5.f6.89.00
device(config-cluster-client-200)# deploy
device(config-cluster-1)# client-pw
device(config-cluster-client-pw)# esi 01:02:03:04:05:06:07:08:0a
device(config-cluster-client-pw)# deploy
```

Displaying information related to VPLS and VLL MCT

The following examples display PW client and bridge domain information for VPLS and VLL MCT.

Displaying PW client information on an MCT cluster

The following example displays the configuration and state information of the PW client on the MCT cluster.

```
device# show cluster 1
Cluster c1 1
================
Cluster State: Deployed
Client Isolation Mode: Loose
Configured Member Vlan Range: 100-101
Active Member Vlan Range: 100-101
Configured Member BD Range: 1000-1001
Active Member BD Range: 1000-1001
No. of Peers: 1
No. of Clients: 2


 Peer Info:
 ==========
 Peer IP: 10.38.38.38, State: Up
  Peer Interface: Ethernet 3/1


 Client Info:
 ============
 Name           Id          ESI                Interface                Local/Remote State
 ----          ----       -----------         ---------                ------------------
 c3             3         a:b:1:2:3:0:0:0:0     Port-channel 100          Up / Up
 Client-PW      34816     a:b:c:d:0:0:0:0:0     PW                        Up / Up
```

The following example displays only PW client and its bridge-domain information on the MCT cluster.

```
device# show cluster 1 client-pw
Client Info:
============
Client: Client-pw, client-id: 34816, Deployed, State: Up
Interface: PW
Bridge-domains: 8100-8101
Elected DF for Bridge-domains:
8100
```

The following example displays the multicast and unicast labels, and forwarding state for the cluster member bridge domain.

```
device# show cluster member bridge-domain
 BD-ID     Mcast-label(Lo/Re)     Unicast-label(Lo/Re)     Forwarding state
-------    ------------------     --------------------     ----------------
  1000       822248/ -1                 805864/ 0                Down
  1001       822249/ -1                 805865/ 0                Down
```

Displaying the MCT state on a bridge domain

In the **show bridge-domain** output, the MCT Enabled field displays whether the bridge domain is configured under a cluster configuration.

```
device# show bridge-domain 501
Bridge-domain Type: MP , VC-ID: 501, MCT Enabled: TRUE
Number of configured end-points: 3 , Number of Active end-points: 3
VE if-indx: 0, Local switching: TRUE, bpdu-drop-enable: TRUE
PW-profile: default, mac-limit: 0
VLAN: 501, Tagged ports: 1(1 up), Un-tagged ports: 0 (0 up)
Tagged Ports: eth1/6.501
Un-tagged Ports:
```

```
Total VPLS peers: 2 (2 Operational):
VC id: 501, Peer address: 10.9.9.9 , State: Operational , uptime: 1
hr 40 min 51 sec
Load-balance: True , Cos Enabled: False,
Tunnel cnt: 4
rsvp p101 (cos_enable:False cos_value:0)
rsvp p102 (cos_enable:False cos_value:0)
rsvp p103 (cos_enable:False cos_value:0)
rsvp p104 (cos_enable:False cos_value:0)
Assigned LSPs count:4 Assigned LSPs:p101 p102 p103 p104
Local VC lbl: 988042, Remote VC lbl: 985332,
Local VC MTU: 1500, Remote VC MTU: 1500,
Local VC-Type: 5, Remote VC-Type: 5
VC id: 501, Peer address: 56.56.56.56 , State: Operational , uptime: 1
hr 40 min 13 sec
Load-balance: True , Cos Enabled: False,
Tunnel cnt: 1
rsvp q101 (cos_enable:False cos_value:0)
Assigned LSPs count:0 Assigned LSPs:
Local VC lbl: 988043, Remote VC lbl: 986039,
Local VC MTU: 1500, Remote VC MTU: 1500,
Local VC-Type: 5, Remote VC-Type: 5
```

Displaying MAC address information for a VPLS bridge domain on an MCT cluster

The following example displays the MAC address table for bridge domain of an MCT cluster.

```
device# show mac-address-table bridge-domain
Vlan/BDId    Mac-address       Type           State      Ports/LIF/PeerIp
100(B)       001b.ed0b.ae00    Dynamic-CL     Active      2.2.2.2
100(B)       0230.0774.0aac    Dynamic-CCL     Active      eth 1/2
100(B)       0230.0774.0aac    Dynamic-CCR     Active      po100
200(B)       003d.ed0b.ae00    Dynamic-CCL     Active      3.3.3.3
200(B)       0230.0774.0aac    Dynamic-CCL     Active      3.3.3.3
```

The following example displays all the MAC addresses learned from other VPLS PE nodes over MCT bridge domains.

```
device# show mac-address-table bridge-domain
Vlan/BDId    Mac-address       Type           State      Ports/LIF/PeerIp
100(B)       001b.ed0b.ae00    Dynamic-CCL     Active      2.2.2.2
200(B)       003d.ed0b.ae00    Dynamic-CCL     Active      3.3.3.3
200(B)       0230.0774.0aac    Dynamic-CCL     Active      3.3.3.3
```

Enabling Layer3 routing for an MCT VLAN

Layer 3 routing is supported for IPv4 and IPv6 BGP, OSPF, and IS-IS routing protocols on an MCT VLAN.

The enabling of the Layer 3 protocols on the MCT VLAN are the same as enabling them on a VE interface. You must first create the VE interface for an MCT VLAN.

> **NOTE**
> For VE over an MCT VLAN interface, you cannot enable MPLS on it.

The following configuration example enables OSPFv2 and OSPFv3 protocols on VE 200 for the MCT member VLAN 2.

```
router ospf
  area 0

ipv6 router ospf
  area 0

vlan 2
 router-interface Ve 200
```

```
interface Ve 200
 ipv6 address 2001::1/64
 ip address 10.2.2.1/24

 ip ospf area 0
 ipv6 ospf area 0
 !
 no shutdown
!
```

Using MCT with VRRP and VRRP-E

Standard VRRP and VRRP-E configuration commands, VMAC generation login, and scaling numbers apply when used with MCT. VRRP advertisement packets are exchanged through the MCT ICL.

The MCT device that acts as the Virtual Routing Redundancy Protocol (VRRP) and VRRP Extended (VRRP-E) backup router performs as a Layer 2 switch to pass the packets to the VRRP or VRRP-E master router for forwarding. Through MAC synchronization, the VRRP or VRRP-E backup router learns the virtual MAC (VMAC) on the Inter-Chassis Link (ICL) represented by the MPLS cloud in the diagram. The data traffic and control traffic both pass through the ICL MPLS cloud link from the backup router. If VRRP-E short path forwarding is enabled, the backup router can forward the packets directly, instead of sending them to the master.

> NOTE
> Short path forwarding is only supported on VRRP-E.

In the diagram below, when an ARP request from the S1 switch device is sent through the direct link to the VRRP or VRRP-E backup router (PE2), a broadcast packet is sent to the VRRP/E master router (PE1) for processing through the ICL (MPLS cloud). When the ARP request is received by the PE1 device, PE1 sends a reply through the direct link to S1. If the ARP reply was received before the MAC address for the MCT on S1 is learned, the reply packet may be flooded to both the Customer Client Edge Port (CCEP) ports and ICL ports.

Using VRRP or VRRP-E, data traffic received from a client device on a backup router is Layer 2 switched to the master device. If VRRP-E short path forwarding is enabled, traffic received on the backup device may be forwarded by the backup if the route to the destination device is shorted than through the master device.

MCT short path forwarding configuration using VRRP-E example

In this example configuration, we are assuming that MCT is using the VRRP-E short path forwarding. When short path forwarding is enabled, packets from either the E1 or E2 devices with a destination of the E4 device can be routed through the PE 2 device which is a VRRP-E backup device. Short path forwarding is designed for load-balancing and allows packets to use the shortest path, and in this case, PE2 is directly connected to E4 so the packets will travel through PE2.

FIGURE 10 MCT short path forwarding

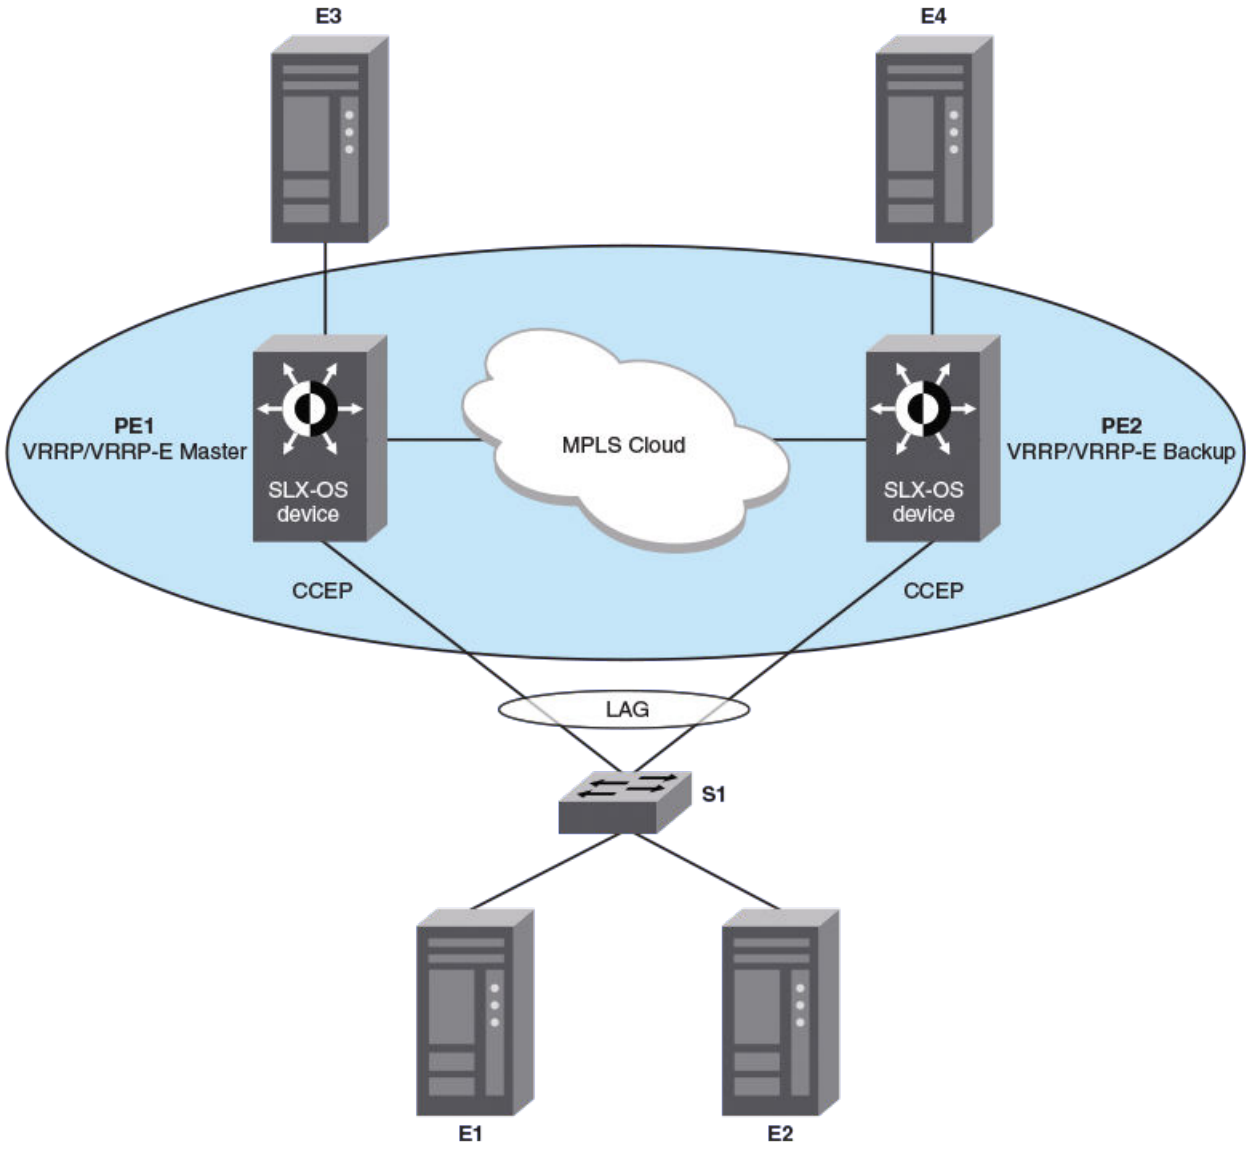

PE1 configuration

The following example configures the OSPF, BGP, and MPLS protocols with cluster configuration for MCT for the PE1 router in the diagram. A VRRP-E priority value of 110 (higher than the device at PE2) allows the PE1 device to assume the role of VRRP-E master.

```
ip router-id 10.19.19.19
 router ospf
   area 0

 interface Loopback 200
  no shutdown
  ip ospf area 0
  ip address 10.19.19.19/32

 router bgp
 local-as 100
```

```
 neighbor 10.32.32.32 remote-as 100
 neighbor 10.32.32.32 update-source loopback 200
 address-family ipv4 unicast
!
 address-family ipv6 unicast
 address-family l2vpn evpn
 neighbor 10.32.32.32 activate
 exit

!
 interface Ethernet 2/3
  ip address 10.1.8.19/24
  ip ospf area 0
  no shutdown
!
router mpls
 mpls-int Ethernet 2/3
 lsp to32
  to 10.32.32.32
  enable
!
vlan 100
!
interface Ethernet 2/5
 switchport
 switchport mode trunk-no-default-native
 switchport trunk allow vlan add 100
 no shutdown

cluster c1 1
 peer 10.32.32.32
 member vlan add 100
 deploy
  client c1 1
   client-interface Ethernet 2/5
    esi 01:02:03:04:05:06:07:08:09
    deploy
!
vlan 100
 router-interface Ve 100
!
protocol vrrp-extended
interface Ve 100
 ip proxy-arp
 ip address 10.2.3.6/24
 vrrp-extended-group 1
  priority 110
  short-path-forwarding
  virtual-ip 10.2.3.4
 no shutdown
!
interface Ve 100
 ipv6 address fe80::1:2 link-local
 ipv6 address 3313::2/64
 ipv6 vrrp-extended-group 1
 virtual-ip 3313::1
```

PE2 configuration

The following example configures the OSPF, BGP, and MPLS protocols with cluster configuration for MCT for the PE2 router in the diagram. A VRRP-E priority value of 80 (lower than the device at PE1) allows the PE2 device to assume the role of a VRRP-E backup device.

```
ip router-id 10.32.32.32
 router ospf
  area 0

 interface Loopback 100
```

```
  no shutdown
  ip ospf area 0
  ip address 10.32.32.32/32

 router bgp
 local-as 100
 neighbor 10.19.19.19 remote-as 100
 neighbor 10.19.19.19 update-source loopback 100
 address-family ipv4 unicast
!
 address-family ipv6 unicast
!
 address-family l2vpn evpn
 neighbor 10.19.19.19 activate
!
!
interface Ethernet 2/3
 ip address 10.1.8.32/24
 no shutdown
  ip ospf area 0
!
router mpls
 mpls-int Ethernet 2/3
  lsp to19
   to 10.19.19.19
   enable
!
vlan 100
!
interface Ethernet 2/7
 switchport
 switchport mode trunk-no-default-native
 switchport trunk allow vlan add 100
 no shutdown
!
cluster c1 1
 peer 10.19.19.19
 member vlan add 100
 deploy
  client c1 1
    esi 01:02:03:04:05:06:07:08:09
    client-interface Ethernet 2/7
    deploy
!
vlan 100
 router-interface Ve 100
!
protocol vrrp-extended
interface Ve 100
 ip proxy-arp
 ip address 10.2.3.5/24
 vrrp-extended-group 1
  priority 80
  short-path-forwarding
  virtual-ip 10.2.3.4
 no shutdown
!
interface Ve 100
 ipv6 address fe80::1:1 link-local
 ipv6 address 3313::3/64
 ipv6 vrrp-extended-group 1
 virtual-ip 3313::1
```

MCT use cases

An L2 MCT solution can be deployed at the access, aggregation, and the core of the data center. However, SLX-OS device is targeted for the data center core.

L2 MCT in the data center core

The following diagram shows a typical 3-tier data center where access and aggregation layers are running Layer 2 and the core is running Layer 2 and Layer 3. The access and aggregation can be standalone Extreme switches or any other third party switches.

FIGURE 11 Typical 3-tier data center

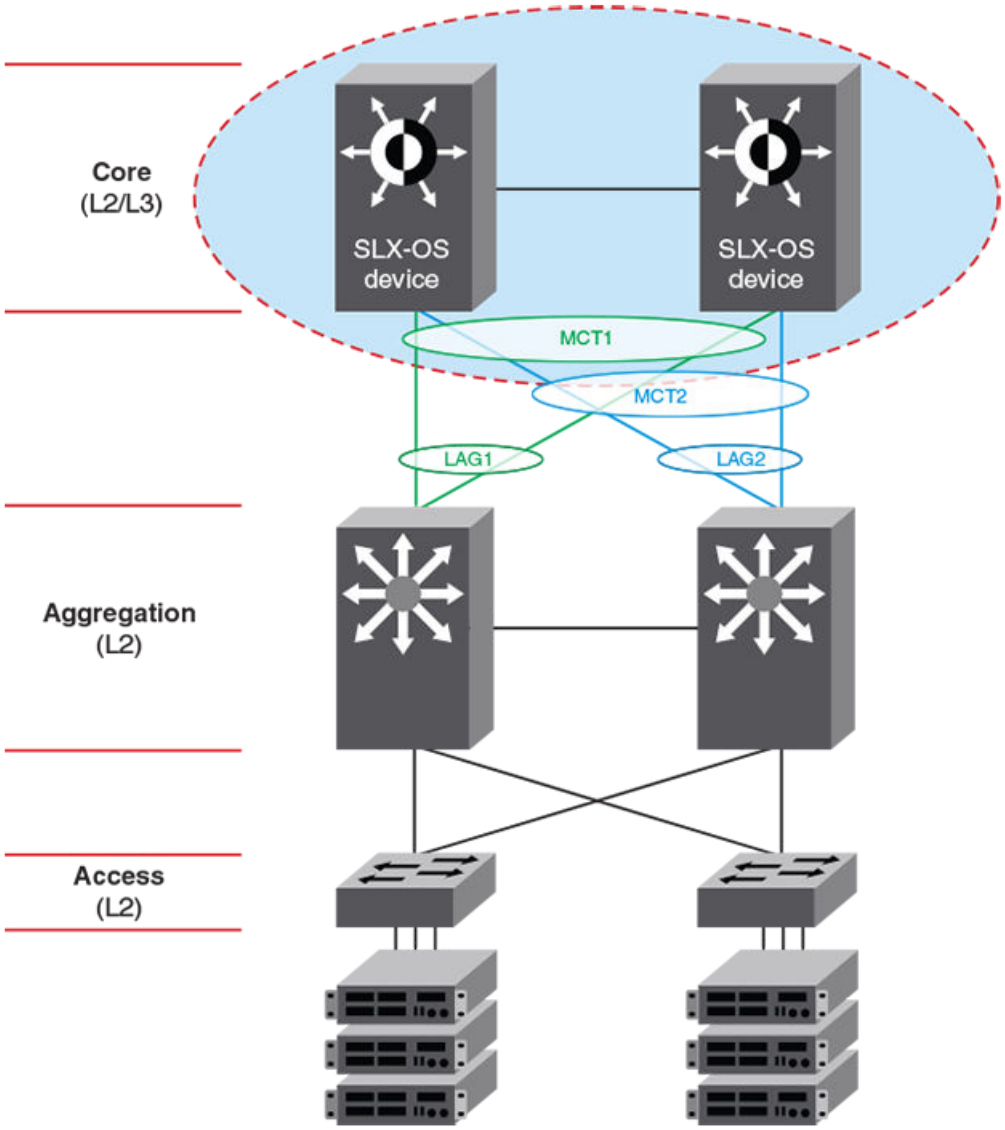

Another variation of this use case is when the aggregation layer is a virtual cluster of switches which is transparent to SLX-OS devices in the core layer.

FIGURE 12 L2 MCT in the data center core connecting to a VCS

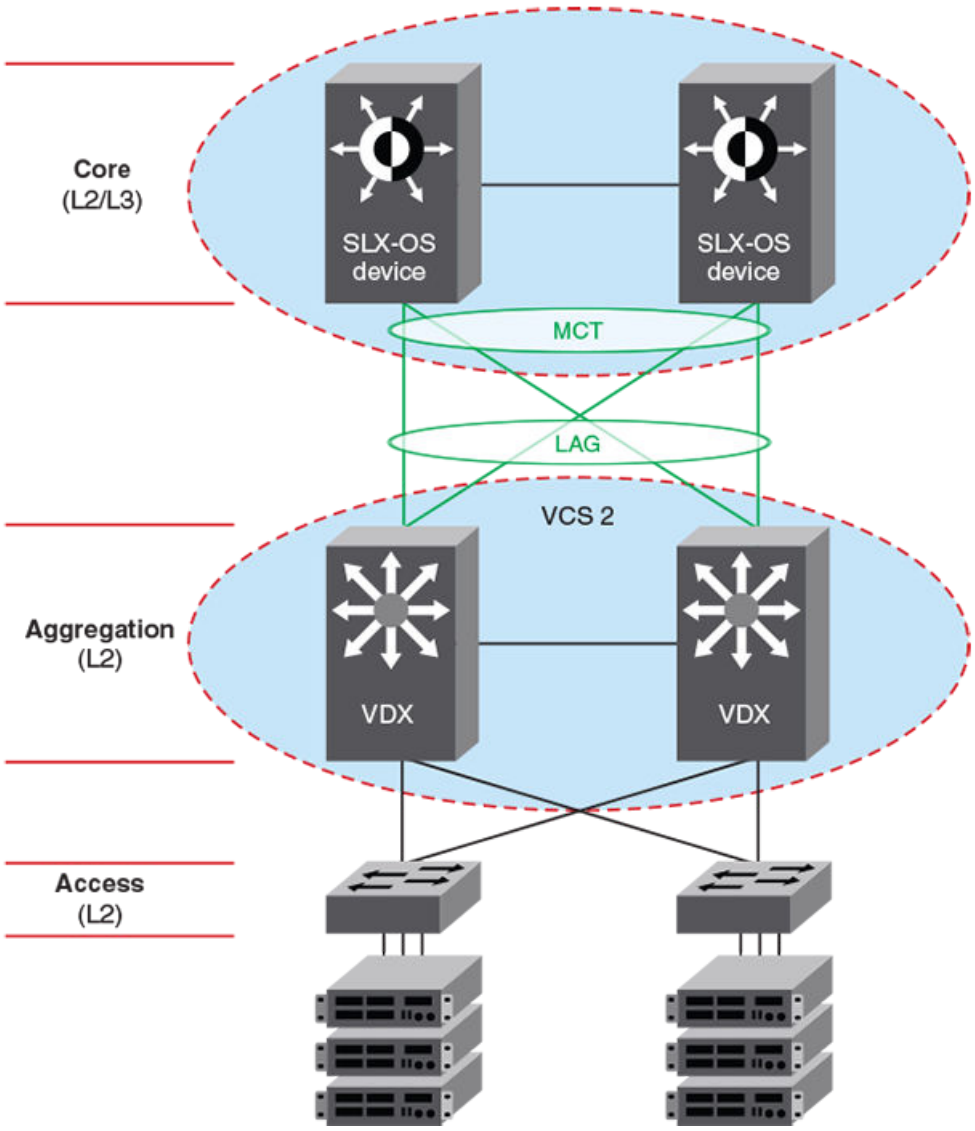

L2 MCT in a data center with a collapsed core and aggregation

The following diagram describes a scenario where VCS fabric of VDX 8870 switches is deployed at the access layer. With the availability of 10G and 40G interface, access switches can connect directly to the core without the need to have a separate aggregation layer.

FIGURE 13 L2 MCT with collapsed core and aggregation

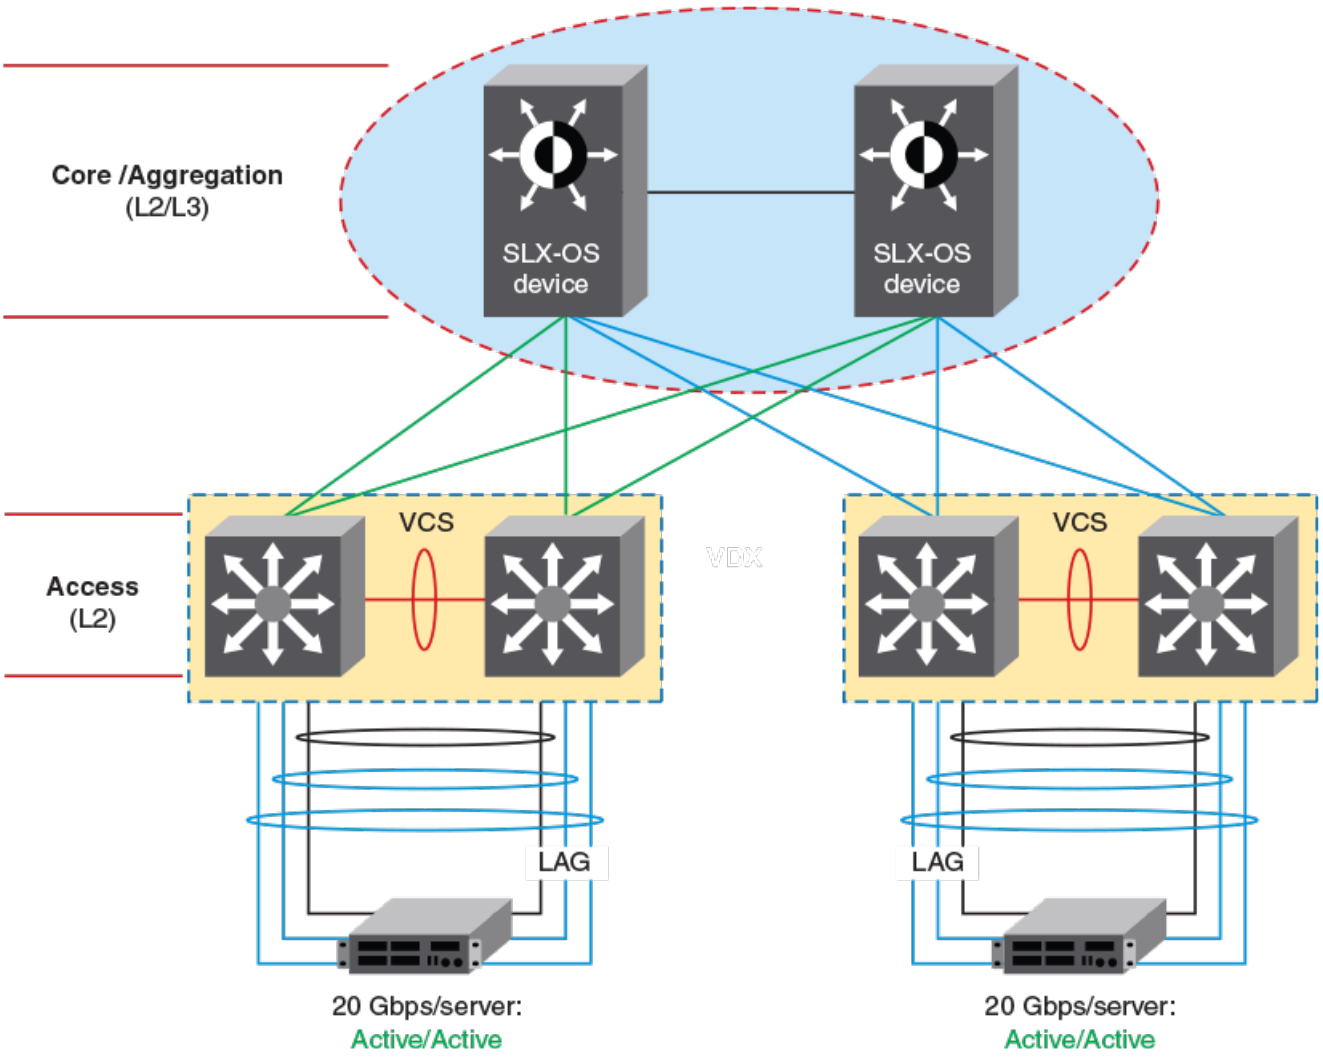

VPLS and VLL Layer 2 VPN services

VPLS overview

Virtual Private LAN Service (VPLS) is a Layer 2 Virtual Private Network (L2 VPN) architecture that provides multipoint Ethernet LAN services.

VPLS provides transparent LAN services across provider edge (PE) devices using Internet Protocol (IP) or Multiprotocol Label Switching (MPLS) as the transport technology.

Because it emulates LAN switching, VPLS is considered to be a L2 service that operates over Layer 3 (L3) clouds.

VPLS provides point-to-multipoint (p2mp) functionality.

The following figure shows a VPLS topology in which switched packets traverse a network.

FIGURE 14 VPLS topology with switching between attachment circuits (ACs) and network core

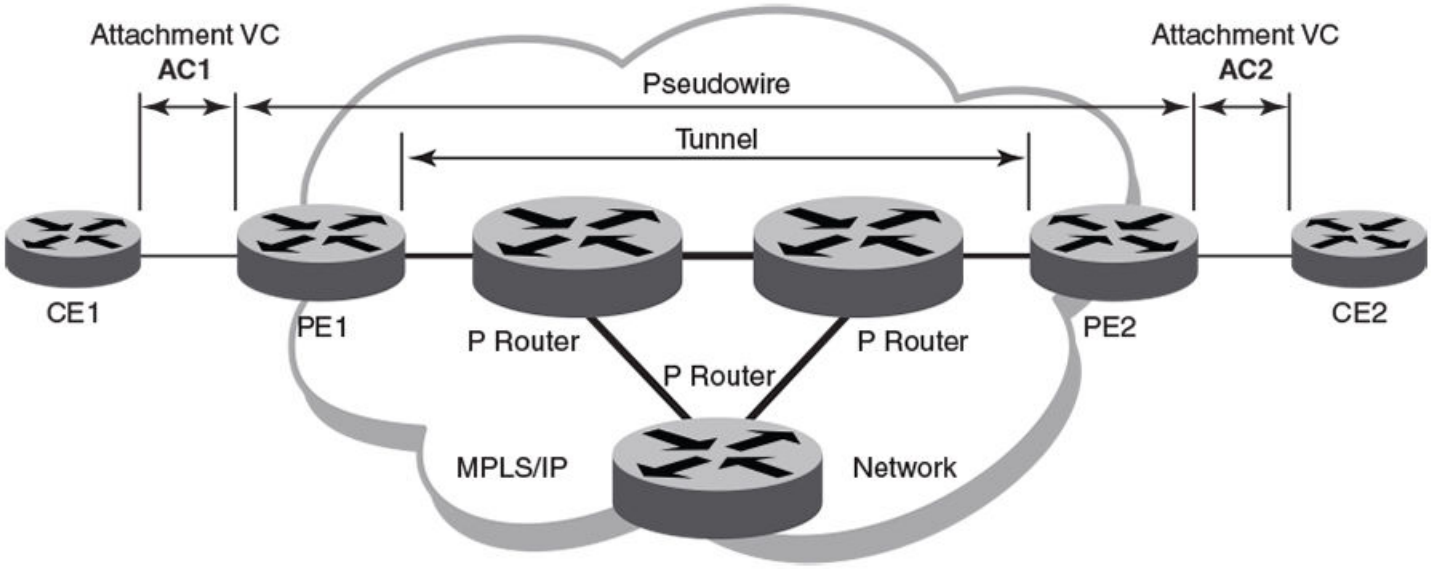

AC1 and AC2 represent L2 connectivity between customer edge (CE) and provider edge (PE) devices.

Pseudowire is a circuit emulation infrastructure that extends L2 connectivity from CE1 to CE2 by way of PE1 and PE2. The tunnel is typically a L3 tunnel on which a L2 circuit is emulated.

In the case of a packet flowing from CE1 to CE2, the packet enters PE1 from CE1 after the forwarding database (FDB) is used to determine the destination MAC address. Then, a virtual connection (VC) label is imposed prior to encapsulation with the tunnel forwarding information, and the packet is sent out onto the wire towards the network core.

Essentially, the topology in the preceding figure shows a L2 VPN enabling the transport of L2 traffic between two or more native Ethernet networks through an underlying Multiprotocol Label Switching (MPLS) provider network. Customer edge (CE) is the last mile and provider edge (PE) is the first mile node for packets transported towards the provider network. The provider intermediary network is an emulated switch (LAN) or wire (LINE) to the CE. The attachment circuit (AC) represents the logical link between the CE and PE.

An AC may be a port, IEEE 802.1q or IEEE 802.1ad (QinQ)) for Ethernet VPNs. A pseudowire (PW) or emulated wire is used as a transport mechanism to tunnel frames between PEs. A PW is characterized by a circuit identifier, which identifies the destination PE.

MPLS tunnels and paths are established by using routing protocols. PW circuits are established by using signaling.

The following figure shows a VPLS topology where switching occurs between two local AC endpoints. This implementation of VPLS does not use VC labels or a pseudowire.

FIGURE 15 VPLS topology with local switching

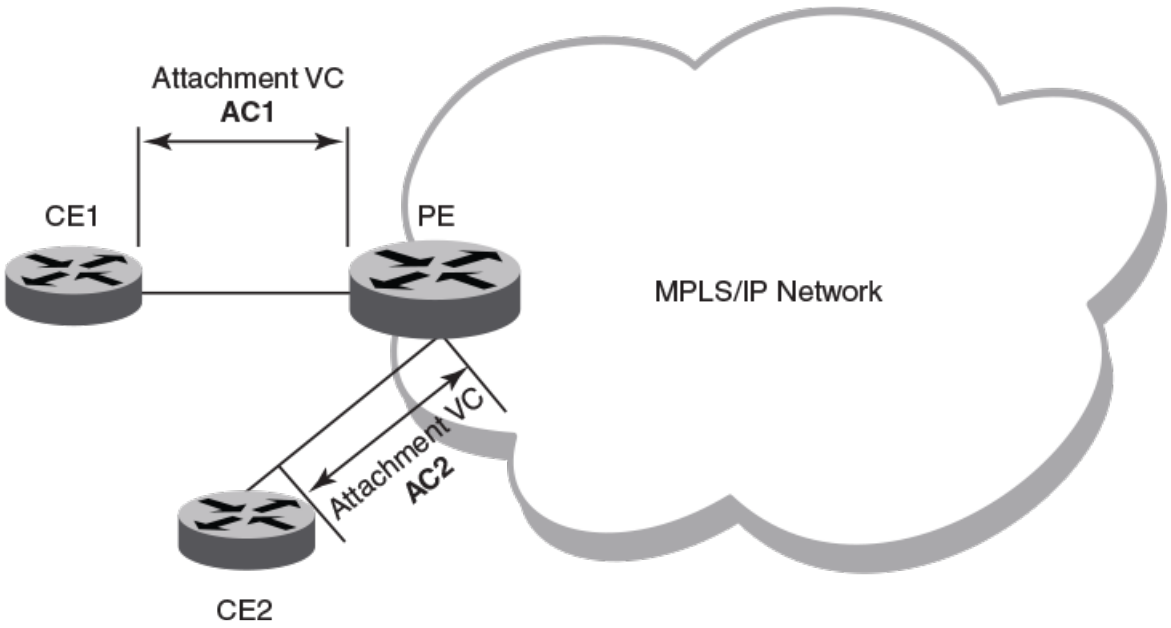

The following figure shows a common VPLS deployment; an enterprise LAN service. The CE devices represent customer edge devices while the PE devices represent provider edge devices.

FIGURE 16 Enterprise LAN service (VPLS)

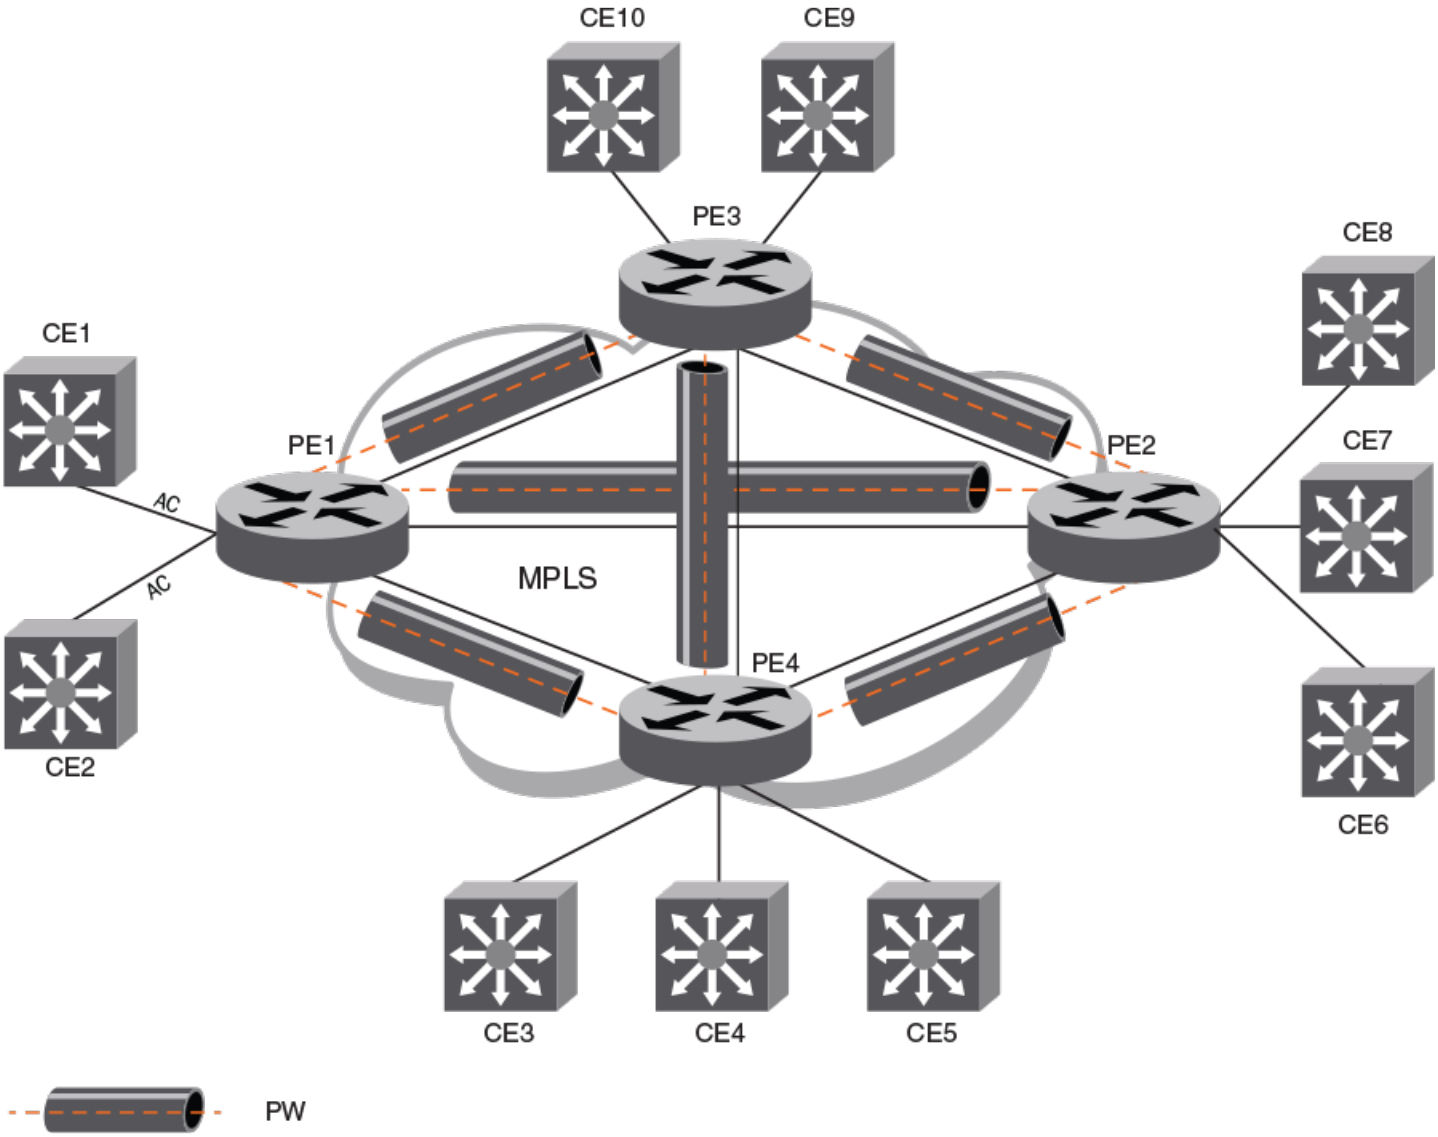

VLL

Virtual Leased Line (VLL) is a Layer 2 Virtual Private Network (L2 VPN) architecture that provides point-to-point Ethernet line or Virtual Private Wire Services (VPWS).

A VLL instance is a special type of VPLS deployment.

VLL provides point-to-point (p2p) connectivity between two access networks or endpoints. Typically, a VLL is used to connect two sites that are geographically apart.

The following figure depicts an enterprise VLL service.

FIGURE 17 Enterprise leased line service (VLL)

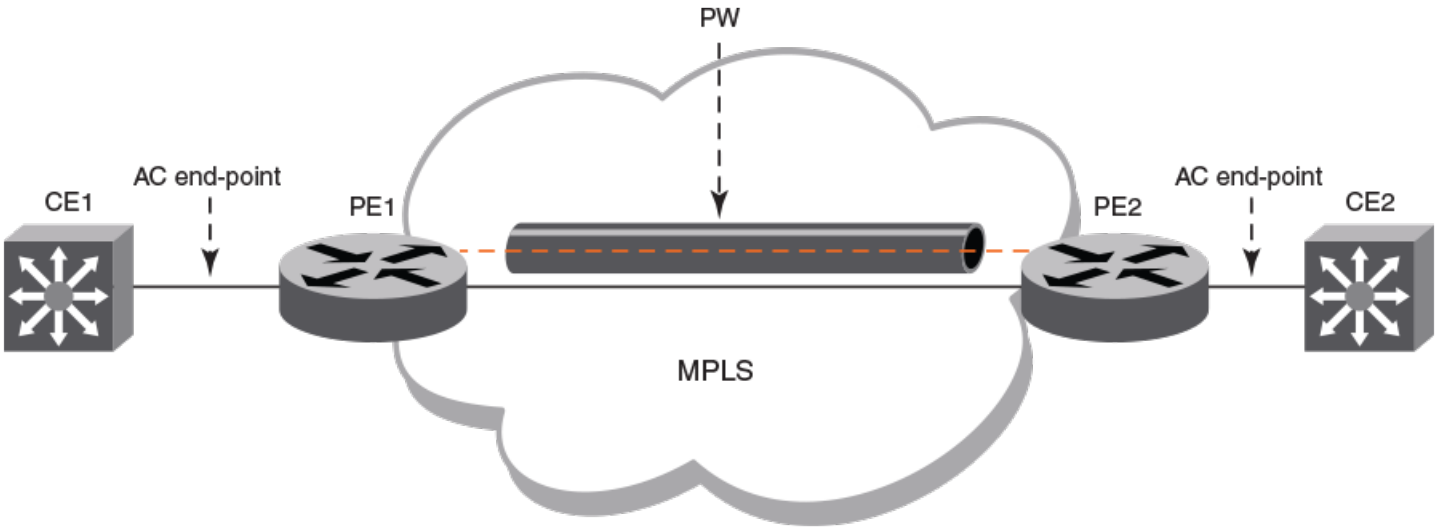

CE1 and CE2 are the customer edge devices in geographically separate sites.

Pseudowire (PW) is a circuit emulation infrastructure that extends L2 connectivity from CE1 to CE2 by way of PE1 and PE2. The tunnel is typically a L3 tunnel on which a L2 circuit is emulated.

VPLS service endpoints

VPLS supports two types of service-endpoints for VPLS and VLL.

Service endpoints can be categorized as:

- AC endpoints
- PW endpoints

An AC endpoint is a L2 link between a PE device and a CE device. The AC endpoint can be an untagged port, or a tagged port with one or more VLANs. AC endpoints with different VLAN tags can be configured in a single VPLS instance.

A VLL instance interconnects two AC endpoints through a pseudowire, while a VPLS instance forms a full mesh interconnection between multiple AC endpoints through multiple PWs.

The following endpoints are supported:

- port-vlan
- port-vlan-vlan
- PW

Both regular port and port-channel interfaces can be used to form port-vlan, untagged port-vlan, and port-vlan-vlan endpoints.

VPLS service endpoints are represented by logical interfaces (LIFs). By using LIFs, features that apply to regular interfaces, such as QoS, can be applied to VPLS service endpoints.

Local switching

The forwarding behavior of AC endpoints in a VPLS instance is controlled by the switching-mode configuration for local endpoints.

When local switching is enabled, traffic is switched and flooded among AC endpoints in addition to between ACs and PWs. When local switching is disabled, the forwarding between AC endpoints is suppressed.

When an unknown unicast packet is received on an AC endpoint and local switching is enabled, the packet is flooded to all other AC endpoints and PW endpoints in the VPLS instance. When local switching is disabled, the unknown packet is only flooded to the PW endpoints in the domain.

Regardless of the local switching configuration, an unknown unicast packet that is received on a PW endpoint is flooded to all AC endpoints.

By default, local switching is enabled.

In a VPLS instance that does not have a PW peer and where all endpoints are AC endpoints (Local VPLS), local switching must be enabled.

To avoid receipt of traffic with different VLAN tags on local endpoints, it is recommended that local switching is disabled in a bridge domain where the PW profile is configured with the VC mode option of **raw-passthrough**. Raw passthorugh mode is designed to forward packets between two VPLS peer devices and is not intended for use with local switching.

Bridge domains

Bridge domain is an infrastructure that supports the implementation of different switching technologies.

A bridge domain is a generic broadcast domain that is not tied to a specific transport technology. Bridge domains support a wide range of service endpoints including regular L2 endpoints and L2 endpoints over L3 technologies.

Bridge domains switch packets between a range of different endpoint types; for example, attachment circuit (AC) endpoints, Virtual Private LAN Service (VPLS) endpoints, Virtual Leased Line (VLL) endpoints, and tunnel endpoints.

VPLS performs multipoint switching, while VLL performs one-to-one switching.

A bridge domain that is created for a VPLS application is also referred to as a VPLS instance.

The following services are bridge-domain capable:

- VPLS—with multiple AC endpoints and pseudowire (PW) logical interfaces (LIFs)
- Local VPLS—with multiple AC endpoints
- VLL—with one AC endpoint and one PW endpoint

Pseudowires

A pseudowire (PW) is a virtual circuit (VC) formed between two PE devices that connect two attachment circuits (ACs).

An Ethernet pseudowire is logically viewed as an L2 nexthop (VC label) that is reachable through an L3 nexthop (LDP label).

The frames from an AC endpoint packet are sent through an ingress pseudowire interface (which abstracts the transport path and packet encapsulations) towards the remote PE. An egress pseudowire interface then abstracts the packet received from a remote PE and hands it over to the corresponding AC end-point.

A pseudowire interface is unidirectional.

PWs support the following underlying MPLS tunnels:

- LDP – Single Path LSP

- RSVP – Single Path LSP
- RSVP – Pri/Sec (Act LSP)
- RSVP – Pri/Sec (Pas LSP)
- FRR: Adaptive LSP (Make Before Break)
- FRR: Protected & detour (1:1)

PWs do not support the following underlying MPLS tunnels:

- FRR: Protected & Bypass (N:1)
- LDP – Multipath LSP (ECMP)
- LDP over RSVP

Pseudowire operation

The pseudowire setup process establishes the reachability of VPLS bridge domain endpoints across an IP or MPLS cloud.

A pseudowire is operational when the following conditions are met:

- VC signaling is established.
- The L3 reachability of the PW peer is resolved.
- At least one AC endpoint within the bridge domain is up.

A pseudowire is non-operational when the following conditions are met:

- No logical interface is configured for the VPLS instance.
- All AC endpoints are non-operational.

Supported pseudowire features

Pseudowires (PWs) support the following features for each configured PW:

- LSP Load Balancing—Load balancing across a maximum of 16 underlying MPLS tunnels.
- Assigned LSP—A maximum of 32 LSPs can be assigned.
- Specific COS—The underlying MPLS tunnel with the closest CoS value is selected for the transport
- Raw, raw-passthrough, or tagged mode—Can be configured by way of the PW profile that is associated with the bridge domain.
- MTU and MTU check—Can be configured by way of the PW profile that is associated with the bridge domain.
- Uniform and pipe mode for QoS
- Statistics—Egress and ingress statistics are supported but must be enabled in the bridge-domain configuration by using the **statistics** command

Unsupported pseudowire features

Pseudowires (PWs) do not support the following features:

- Auto-discovery of peers
- PW redundancy
- Static PW peers
- VC MAC withdraw
- Status TLV update
- VEoVPLS

- OAM
- Multicast snooping
- Extended counters
- High availability—Process restart
- High availability—ISSU

PW VLAN tag manipulation (vc-mode)

The virtual connection (VC) mode configuration for a pseudowire (PW) profile determines how VLAN tags are manipulated when a packet is received or transmitted on the PW.

The following table describes VC modes that are supported on PWs.

TABLE 8 VC modes supported on pseudowires

VC mode	Description
Raw	At VC label imposition, when a tagged packet is received on a tagged AC endpoint, the VLAN tag is removed before it is sent out on the wire. When an untagged packet is received on an untagged AC endpoint it is encapsulated as is and sent out on the wire.
Raw-passthrough	Enables interoperation with third-party devices. When a packet that is destined for a remote peer is received on either a tagged or untagged AC endpoint, it is encapsulated in an MPLS header and sent on to the MPLS cloud without adding or removing VLAN tags. When a packet that is destined for a local endpoint is received on either a tagged or untagged AC endpoint, the MPLS header is removed before sending it on to the local endpoint; VLAN tags in the original packet are not changed in any way.
Tagged	At VC label imposition, when a tagged packet is received on a tagged AC endpoint, the packet is encapsulated as is and sent out on the wire. This applies to dual-tagged and triple-tagged (or more) endpoints also; that is, tags are neither altered or removed but are sent to the remote peer. When an untagged packet is received on an untagged AC endpoint, a dummy tag is added and it is sent out on the wire.

The VC mode is agreed by PE peer devices during the pseudowire signaling process.

A single VPLS instance can have a mixture of tagged and untagged endpoints.

When the VC mode is changed on a device, the PWs are torn down are re-established except in the cases of a change from raw to raw-passthrough or a change from raw-passthrough to raw. The traffic impact is minimal (because PWs are not torn down and re-established) when the VC mode is changed from raw to raw-passthrough (or vice versa).

VC mode is configured by specifying the **vc-mode** option for the **pw-profile** command.

PW statistics

Statistics are supported for both the Ingress PE (packet going over the PW) and egress PE (packet received on the VC-Label of the PW).

The statistics are enabled or disabled per bridge domain and applies to all the PWs which are part of the bridge domain. The logical interface for inbound and outbound statistics are shared resources. Hence, the corresponding PW is operational only if these hardware resources are available.

PW statistics is configured using the **statistics** command.

```
device (config)# bridge-domain 501 p2mp
device (config-bridge-domain-501)# vc-id 501
device (config-bridge-domain-501)# peer 10.9.9.9
device (config-bridge-domain-501)# peer 56.56.56.56
device (config-bridge-domain-501)# statistics
```

The **show statistics** command displays the statistics.

```
device# show statistics bridge-domain 501
Bridge-Domain Statistics
BD Index:501
Interface               Rx Pkts             Rx Bytes              Tx Pkts            Tx
Bytes

10.9.9.9          7183311757358142120     8237025092157046784     824253729
7806901152

56.56.56.56       7183311757358142120     8237025092157046784     82425372
97806901152

device#
```

For more information on the **statistics** command, please refer the *SLX OS CLI reference.*

Supported VPLS features

The following VPLS features are supported:

- Local VPLS (without PWs)
- VPLS—untagged, single-tagged, and dual-tagged endpoints
- Flooding of L2 BPDUs in VPLS
- VPLS tagged, raw, and raw-passthrough modes for virtual circuits
- Dynamic LAG support for VPLS endpoints
- VPLS MTU enforcement
- VPLS static MAC address support for AC endpoints
- VPLS over Multi-Chassis Trunking (MCT)

Unsupported VPLS feature

The Virtual Ethernet (VE) over VPLS feature is not supported.

Configuration of VPLS and VLL

Configuration of a VPLS or VLL instance includes configuring a bridge-domain, configuring a virtual connection (VC) identifier, configuring logical interfaces for attachment circuit (AC) endpoints, and configuring peer IP addresses.

To configure a VPLS or VLL instance, you must complete the following tasks:

- Configure a bridge domain.
- Configure a VC identifier.
- Configure logical interfaces for AC endpoints.
- (Optional) Configure a pseudowire (PW) profile.
- Configure peer IP addresses. Configuring peer IP addresses creates PW endpoints.

  NOTE

  VPLS (or VLL) configuration is separate from the underlying IP or MPLS configuration. MPLS tunnels need to be brought up separately. For further information about the configuration of MPLS tunnels, refer to the *Extreme SLX-OS MPLS Configuration Guide* for the SLX 9850 Router.

QoS treatment in VPLS packet flow

There are default behaviors for Quality of Service (QoS) propagation in VPLS forwarding on PE routers.

On the ingress label-edge router (LER), the final EXP value for the VC label is not dependant on the CoS value in the VC-peer configuration.

By default, for traffic flowing from a CE device to a PE device, 3 bits of the PCP field from the incoming Ethernet frame header are extracted and mapped to an internal CoS value by way of an ingress CoS map. This internal value is then mapped to an outgoing CoS value by way of an egress CoS map. The outgoing CoS value is then inserted into the EXP field in the outgoing VC label. When incoming traffic does not have VLAN tag, the default PCP value that is configured on a port is used.

In the case of traffic received from the network core side, by default the EXP field from the incoming VC label is mapped to an internal CoS value by way of an ingress CoS map. This internal value is then mapped to an outgoing CoS value by way of an egress CoS map. The outgoing CoS value is then inserted into the PCP field in the Ethernet frame header going out to the CE device.

On the egress LER, the CoS value for the VC-peer configuration is not dependant on the final EXP value for the VC label.

The following table shows ingress and egress behavior for different global, tunnel and PW configuration combinations.

TABLE 9 Ingress and Egress LER behavior

Global	Tunnel	PW	AC to PW (per path)	PW to AC (per PW)
Uniform	No CoS	No CoS	Uniform	Uniform
Uniform	CoS	No CoS	Pipe	Uniform
Uniform	No CoS	CoS	Uniform	Pipe
Uniform	CoS	CoS	Pipe	Pipe
Pipe	No CoS	No CoS	Pipe	Pipe
Pipe	CoS	No CoS	Pipe	Pipe
Pipe	No CoS	CoS	Pipe	Pipe
Pipe	CoS	CoS	Pipe	Pipe

Bridge domain statistics

Devices gather statistics for all the logical interfaces and peers in bridge domains.

Use the **statistics** command in the bridge domain configuration mode to enable statistics on a bridge domain.

> NOTE
> Statistics has to be manually enabled for a specific domain, since it is not enabled by default for bridge domains.

Please note that:

- The statistics reported are not real-time statistics since they depend upon the load on the system.
- Statistics has to be manually enabled for a bridge domain. This ensures better utilization of the statistics resources in the hardware.
- Enabling statistics on a bridge domain has a heavy impact on the data traffic.

Configuring a PW profile

A pseudowire (PW) emulates a point-to-point connection over a packet-switching network. PW profile configuration defines PW attributes. After configuration, a PW profile must be attached to a bridge domain.

A pseudowire profile can be shared across multiple bridge domains. Complete the following task to configure a PW profile.

1. From privileged EXEC mode, enter global configuration mode.

```
device# configure terminal
```

2. Create a PW profile and enter configuration mode for the profile.

```
device(config)# pw-profile pw_example
```

3. Configure the virtual connection mode for the profile.

```
device(config-pw-pw_example)# vc-mode tag
```

   In this example, tag mode is configured for the PW profile, pw_example.

4. Configure a maximum transmission unit (MTU) of 1300 for the PW profile.

```
device(config-pw-pw_example)# mtu 1300
```

5. Enforce an MTU check during PW signaling.

```
device(config-pw-pw_example)# mtu-enforce true
```

The following example creates a PW profile named pw_example and configures attributes for the profile.

```
device# configure terminal
device(config)# pw-profile pw_example
device(config-pw-pw_example)# vc-mode tag
device(config-pw-pw_example)# mtu 1300
device(config-pw-pw_example)# mtu-enforce true
```

Configuring a static MAC address over an endpoint in a VPLS instance

A static MAC address can be associated with the logical interface for an attachment circuit (AC) endpoint in a bridge domain.

You can configure a MAC address for a logical interface for an endpoint in a VPLS instance by completing the following task.

> NOTE
> Pre-configuration for the static mac is supported. Pre-configured static mac is shown as inactive mac.

1. From privileged EXEC mode, enter global configuration mode.

```
device# configure terminal
```

2. Create the logical-interface (LIF) entry associated to a physical or port-channel interface, associate the LIF entry to the VPLS instance and configure the static mac associated with theLIF entry.

```
device(config)# mac-address-table static 0011.2345.6789 forward logical-interface ethernet 1/2.200
```

3. Enter privileged EXEC mode.

```
device(config)# exit
```

4. (Optional) Verify the configuration using any of the following commands.

```
device# show mac-address-table
device# show mac-address-table static
device# show mac-address-table bridge-domain [bd-id]
```

Configuring a VPLS instance

A virtual private LAN Service (VPLS) instance provides multipoint LAN services.

Prior to completing the following task, the underlying L3 configuration of MPLS tunnels must be completed. There is a configuration example at the end of this task that shows all the steps in order.

You can configure a VPLS instance by completing the following task.

1. From privileged EXEC mode, enter global configuration mode.

```
device# configure terminal
```

2. Create a multipoint bridge domain.

```
device(config)# bridge-domain 5
```

   By default, the bridge-domain service type is multipoint. In this example, bridge domain 5 is configured as a multipoint service.

3. Configure a virtual connection identifier for the bridge domain.

```
device(config-bridge-domain-5)# vc-id 8
```

4.
   NOTE
   Logical interfaces representing bridge-domain endpoints must be created before they can be bound to a bridge domain.

   Bind the logical interfaces for attachment circuit endpoints to the bridge domain.

```
device(config-bridge-domain-5)# logical-interface ethernet 1/6.400
```

   In this example, Ethernet logical interface 1/6.400 is bound to bridge domain 5.

5. Repeat Step 4 to bind other logical interfaces for attachment circuit endpoints to the bridge domain.

```
device(config-bridge-domain-5)# logical-interface port-channel 2.200
```

   In this example, port channel logical interface 2.200 is bound to bridge domain 5.

6. Configure peer IP addresses to create pseudowire (PW) endpoints.

```
device(config-bridge-domain-5)# peer 10.15.15.15 load-balance
```

   In this example, a peer IP address of 10.15.15.15 is configured under bridge domain 5 and specifies load balancing.

7. Repeat Step 6 to configure more peer IP addresses to create PW endpoints.

```
device(config-bridge-domain-5)# peer 10.12.12.12 lsp lsp1 lsp2
```

   In this example, a peer IP address of 10.12.12.12 under bridge domain 5 and specifies two label-switched paths (lsp1 and lsp2).

8. (Optional) Configure local switching for bridge domain 5.

```
device(config-bridge-domain-5)# local-switching
```

9. (Optional) Enable dropping L2 bridge protocol data units (BPDUs) for bridge domain 5.

```
device(config-bridge-domain-5)# bpdu-drop-enable
```

10. (Optional) Configure a PW profile under the bridge domain 5.

```
device(config-bridge-domain-5)# pw-profile 2
```

The following example creates bridge domain 5 and configures virtual connection identifier 8 for the bridge domain. It binds ethernet and port-channel logical interfaces to the bridge domain and configures peer IP addresses under the domain. It configures local switching, enables dropping of L2 BPDUs, and configures a PW profile for the domain.

```
device# configure terminal
device(config)# bridge-domain 5
device(config-bridge-domain-5)# vc-id 8
device(config-bridge-domain-5)# logical-interface ethernet 1/6.400
device(config-bridge-domain-5)# logical-interface port-channel 2.200
device(config-bridge-domain-5)# peer 10.15.15.15 load-balance
device(config-bridge-domain-5)# peer 10.12.12.12 lsp lsp1 lsp2
device(config-bridge-domain-5)# local-switching
device(config-bridge-domain-5)# bpdu-drop-enable
device(config-bridge-domain-5)# pw-profile 2
```

Configuring a VLL instance

A virtual leased line (VLL) instance provides point-to-point (peer) LAN services.

Prior to completing the following task, the Ethernet logical interface and pseudowire profiles must be created. There is an example at the end of this task that shows all the steps in order.

You can configure a VLL instance by completing the following task.

1. From privileged EXEC mode, enter global configuration mode.

```
device# configure terminal
```

2. Create a point-to-point bridge domain to use VLL services.

```
device(config)# bridge-domain 3 p2p
```

   In this example, bridge domain 3 is created as a point-to-point service. By default, the bridge-domain service type is multipoint.

3. Configure a virtual connection identifier for the bridge domain.

```
device(config-bridge-domain-3)# vc-id 500
```

4. Bind the logical interfaces for attachment circuit endpoints to the bridge domain.

```
device(config-bridge-domain-3)# logical-interface ethernet 1/5.15
```

   In this example, the Ethernet logical interface 1/5.15 is bound to bridge domain 3.

5. Configure peer IP addresses to create pseudowire (PW) endpoints.

```
device(config-bridge-domain-3)# peer 10.10.10.10
```

6. Configure a PW profile under the bridge domain.

```
device(config-bridge-domain-3)# pw-profile to-mpls-nw
```

   The following example configures a PW profile 2 under bridge domain 3.

7. Repeat this configuration on the other peer device with appropriate parameters.
8. Enter Privileged EXEC mode.

```
device(config-bridge-domain-3)# end
```

9. (Optional) Display information about the configured VLL instance.

```
device# show bridge-domain 3

Bridge-domain Type: P2P , VC-ID: 3
Number of configured end-points:2 , Number of Active end-points: 2
VE if-indx: 0, Local switching: FALSE, bpdu-drop-enable: FALSE
PW-profile: default, mac-limit: 0
VLAN: 3, Tagged ports: 1(1 up), Un-tagged ports: 0 (0 up)
Tagged Ports: eth1/5.15
Un-tagged Ports:
Total VLL peers: 1 (1 Operational):
VC id: 3, Peer address: 10.10.10.10, State: Operational, uptime: 18 sec
Load-balance: True , Cos Enabled: False,
Tunnel cnt: 4
rsvp p105 (cos_enable:Falsecos_value:0)
rsvp p106 (cos_enable:Falsecos_value:0)
rsvp p107 (cos_enable:Falsecos_value:0)
rsvp p108 (cos_enable:Falsecos_value:0)
Assigned LSPs count:4 Assigned LSPs:p105 p106 p107 p108
Local VC lbl: 851968, Remote VC lbl: 985331,
Local VC MTU: 1600, Remote VC MTU: 1500,
Local VC-Type: 5, Remote VC-Type: 5
```

The following example shows the creation of a logical interface and a pseudowire profile in addition to the bridge domain and VLL instance configuration.

```
device# configure terminal
device(config)# interface ethernet 1/5
device(conf-if-eth-1/5)# switchport
device(conf-if-eth-1/5)# switchport mode trunk
device(conf-if-eth-1/5)# switchport trunk tag native-vlan
device(conf-if-eth-1/5)# shutdown
device(conf-if-eth-1/5)# logical-interface ethernet 1/5.15
device(conf-if-eth-lif-1/5.15)# vlan 200
device(conf-if-eth-lif-1/5.15)# exit
device(conf-if-eth-1/5)# exit
device(config)# pw-profile to-mpls-nw
device(config-pw-profile-to-mpls-nw)# mtu 1600
device(config-pw-profile-to-mpls-nw)# mtu-enforce true vc-mode tag
device(config-pw-profile-to-mpls-nw)# exit
device(config)# bridge-domain 3 p2p
device(config-bridge-domain-3)# vc-id 500
device(config-bridge-domain-3)# logical-interface ethernet 1/5.15
device(config-bridge-domain-3)# peer 10.10.10.10
device(config-bridge-domain-5)# pw-profile to-mpls-nw
```

Displaying bridge-domain configuration information

Various show commands can be used to display bridge-domain configuration information.

- Enter the **show bridge-domain** command to display information about all configured bridge domains.

```
device# show bridge-domain

Total Number of bridge-domains: 3
Number of bridge-domains: 3

Bridge-domain 1
-------------------------------
Bridge-domain Type: mp , VC-ID: 5
Number of configured end-points:  5 , Number of Active end-points: 4
VE if-indx: 1207959555, Local switching: TRUE, bpdu-drop-enable:TRUE
PW-profile: 1, mac-limit: 128000
Number of Mac's learned:90000,     Static-mac count: 10,
VLAN: 100, Tagged ports: 2(2 up), Un-tagged ports: 0 (0 up)
Tagged ports: Eth 0/2/6, eth 0/2/8
Un-tagged ports:

Total PW peers: 2 (2 Operational)
Peer address: 12.12.12.12, State: Operational, Uptime: 2 hr 55 min
      Load-balance: True , Cos enabled:False,
      Assigned LSP;s:
      Tnnl in use: tnl2[RSVP]
      Local VC lbl: 983040, Remote VC lbl: 983040
      Local VC MTU: 1500, Remote VC MTU: 1500,
      Local VC-Type: Ethernet(0x05), Remote VC-Type: Ethernet(0x05)
Peer address: 15.15.15.15, State: Operational, Uptime: 2 hr 55 min
         Load-balance: False , Cos enabled:False,
      Assigned LSP's: lsp1, lsp2
      Tnnl in use: tnl1[MPLS]
      Local VC lbl: 983041, Remote VC lbl: 983043
      Local VC MTU: 1500, Remote VC MTU: 1500 ,
      Local VC-Type: Ethernet(0x05), Remote VC-Type: Ethernet(0x05)

Bridge-domain 2
-------------------------------
Bridge-domain Type: mp , VC-ID: 100
Number of configured end-points:  5 , Number of Active end-points: 4
VE if-indx:  NA, Local switching: FALSE, bpdu-drop-enable:FALSE
PW-profile: profile_1, mac-limit: 262144
Number of Mac's learned:90000,     Static-mac count: 10,
VLAN: 100, Tagged ports: 2(1 up), Un-tagged ports: 0 (0 up)
      Tagged ports: eth 0/2/10, eth 0/1/10
      Un-tagged ports:
VLAN: 150, Tagged ports: 1(1 up), Un-tagged ports: 0 (0 up)
      Tagged ports: eth 0/1/5
      Un-tagged ports:

Bridge-domain 3
-------------------------------
Bridge-domain Type: mp , VC-ID: 200
Number of configured end-points:  5 , Number of Active end-points: 4
VE if-indx: 120793855, Local switching: FALSE, bpdu-drop-enable:FALSE
PW-profile: 2, mac-limit: 262144
Number of Mac's learned:90000,     Static-mac count: 10,
Local switching: TRUE,
VLAN: 500, Tagged ports: 2(2 up), Un-tagged ports: 2 (1 up)
Tagged ports:     eth 0/11/6, eth 0/4/3
Un-tagged ports:

Total VPLS peers: 3 (2 Operational)
Peer address: 5.5.5.5, State: Operational, Uptime: 2 hr 35 min
         Load-balance: False , Cos enabled:False,
      Assigned LSP;s:
      Tnnl in use: tnl2[RSVP]
      Local VC lbl: 983050, Remote VC lbl: 983050
```

```
      Local VC MTU: 1500,Remote VC MTU: 1500,
      Local VC-Type: Ethernet(0x05), Remote VC-Type: Ethernet(0x05)
Peer address: 20.20.20.20, State: Operational, Uptime: 0 hr 18 min
         Load-balance: False , Cos enabled:True,
      Assigned LSP's:
      Tnnl in use: NA,
      Local VC lbl: NA, Remote VC lbl: NA
      Local VC MTU: 1500,Remote VC MTU: 1500,
      Local VC-Type: Ethernet(0x05), Remote VC-Type: Ethernet(0x05)
Peer address: 10.10.10.10, State: Not-Operational (Tunnel Not Available),
         Load-balance: True , Cos enabled:False,
      Assigned LSP's: lsp10, lsp15
      Tnnl in use: NA,
      Peer Index:2
      Local VC lbl: NA, Remote VC lbl: NA
      Local VC MTU: 1500,Remote VC MTU: NA ,
      Local VC-Type: Ethernet(0x05), Remote VC-Type: NA
```

- Enter the **show bridge-domain** command specifying the bridge-domain ID to display information about a specific bridge domain. The following example displays information about bridge domain 501.

```
device# show bridge-domain 501

Bridge-domain 501
------------------------------
Bridge-domain Type: MP , VC-ID: 501
Number of configured end-points:  2 , Number of Active end-points: 2
VE if-indx: 0, Local switching: TRUE, bpdu-drop-enable: TRUE
PW-profile: default, mac-limit: 0
VLAN: 501, Tagged ports: 1(1 up), Un-tagged ports: 0 (0 up)
Tagged Ports: eth1/6.501
Un-tagged Ports:
Total VPLS peers: 1 (1 Operational):

VC id: 501, Peer address: 10.9.9.9, State: Operational, uptime: 2 sec
      Load-balance: False, Cos Enabled: False,
      Tunnel cnt: 1
      rsvp  p101(cos_enable:False cos_value:0)
      Assigned LSPs count:0 Assigned LSPs:
      Local VC lbl: 989042, Remote VC lbl: 983040,
      Local VC MTU: 1500, Remote VC MTU: 1500,
      Local VC-Type: 5, Remote VC-Type: 5
```

The following example shows information about a bridge domain (501) in which the **load-balance** option is configured for the peer device 10.9.9.9.

```
show bridge-domain 501

Bridge-domain 501
------------------------------
Bridge-domain Type: MP , VC-ID: 501
Number of configured end-points:  2 , Number of Active end-points: 2
VE if-indx: 0, Local switching: TRUE, bpdu-drop-enable: TRUE
PW-profile: default, mac-limit: 0
VLAN: 501, Tagged ports: 1(1 up), Un-tagged ports: 0 (0 up)
Tagged Ports: eth1/6.501
Un-tagged Ports:
Total VPLS peers: 1 (1 Operational):

VC id: 501, Peer address: 10.9.9.9, State: Operational, uptime: 48 sec
      Load-balance: True , Cos Enabled: False,
      Tunnel cnt: 16
      rsvp  p101(cos_enable:False cos_value:0)
      rsvp  p102(cos_enable:False cos_value:0)
      rsvp  p103(cos_enable:False cos_value:0)
      rsvp  p104(cos_enable:False cos_value:0)
      rsvp  p105(cos_enable:False cos_value:0)
      rsvp  p106(cos_enable:False cos_value:0)
      rsvp  p107(cos_enable:False cos_value:0)
      rsvp  p108(cos_enable:False cos_value:0)
```

```
      rsvp  p109(cos_enable:False cos_value:0)
      rsvp  p110(cos_enable:False cos_value:0)
      rsvp  p111(cos_enable:False cos_value:0)
      rsvp  p112(cos_enable:False cos_value:0)
      rsvp  p113(cos_enable:False cos_value:0)
      rsvp  p114(cos_enable:False cos_value:0)
      rsvp  p115(cos_enable:False cos_value:0)
      rsvp  p116(cos_enable:False cos_value:0)
      Assigned LSPs count:0 Assigned LSPs:
      Local VC lbl: 989040, Remote VC lbl: 983040,
      Local VC MTU: 1500, Remote VC MTU: 1500,
      Local VC-Type: 5, Remote VC-Type: 5
```

The following example shows information about bridge domain 501 in which the **load-balance** option and four assigned label-switched paths (p101, p102, p103, and p104) are configured for the peer device 10.9.9.9.

```
device# show bridge-domain 501

Bridge-domain 501
------------------------------
Bridge-domain Type: MP , VC-ID: 501
Number of configured end-points:  2 , Number of Active end-points: 2
VE if-indx: 0, Local switching: TRUE, bpdu-drop-enable: TRUE
PW-profile: default, mac-limit: 0
VLAN: 501, Tagged ports: 1(1 up), Un-tagged ports: 0 (0 up)
Tagged Ports: eth1/6.501
Un-tagged Ports:
Total VPLS peers: 1 (1 Operational):

VC id: 501, Peer address: 10.9.9.9, State: Operational, uptime: 4 sec
      Load-balance: True , Cos Enabled: False,
      Tunnel cnt: 4
      rsvp  p101(cos_enable:False cos_value:0)
      rsvp  p102(cos_enable:False cos_value:0)
      rsvp  p103(cos_enable:False cos_value:0)
      rsvp  p104(cos_enable:False cos_value:0)
      Assigned LSPs count:4 Assigned LSPs:p101 p102 p103 p104
      Local VC lbl: 989041, Remote VC lbl: 983040,
      Local VC MTU: 1500, Remote VC MTU: 1500,
      Local VC-Type: 5, Remote VC-Type: 5
```

- Enter the **show bridge-domain brief** command to display summary information about all configured bridge domains.

```
device# show bridge-domain brief

Total Number of bridge-domains configured: 10
Number of VPLS bridge-domains: 5
Macs Dynamically learned: 50360, Macs statically configured: 0

BDID(VC-ID)   TYPE     Intf(up)     PWs(up)  macs
501(501)      P2MP     5(3)         2(2)     50000
502(502)      P2MP     1(1)         1(1)     10
503(503)      P2MP     10(6)        3(1)     0
504(504)      P2MP     1(1)         1(1)     350
505(505)      P2MP     1(1)         1(1)     0
506(506)      P2P      1(1)         1(1)     0
507(507)      P2P      1(1)         1(1)     0
508(508)      P2P      1(1)         1(1)     0
509(509)      P2P      1(1)         1(1)     0
510(510)      P2P      1(1)         1(1)     0
```

Displaying MAC address information for VPLS bridge domains

Various show commands can be used to display MAC address information for bridge domains.

- Enter the **show mac-address-table bridge-domain** command to display information about MAC addresses in VPLS bridge domains. The following example shows details of all MAC addresses learned on all bridge domains.

```
device# show mac-address-table bridge-domain all

VlanId/BD-Id   Mac-address           Type     State     Ports/LIF/peer-ip
629(B)            0011.2222.5555     Dynamic  Active    eth 1/3.100
629(B)            0011.2222.6666     Dynamic  Inactive  eth 1/1.500
629(B)            0011.2222.1122     Dynamic  Active    10.12.12.12
629(B)            0011.2222.3333     static   Inactive  po 5.700
629(B)            0011.0101.5555     Dynamic  Active    eth 1/2.400

Total MAC addresses : 5
```

- Enter the **show mac-address-table bridge-domain** command specifying a bridge domain to display information about MAC addresses for a specific bridge domain.

```
device# show mac-address-table bridge-domain 1

BD-Name   Mac-address     Type     State     Ports/LIF/Peer-IP
1         0011.2222.5555  Dynamic  Active    eth 1/3.200
1         0011.2222.6666  Dynamic  Inactive  eth 2/2.500

Total MAC addresses : 2
```

VPLS MAC withdrawal

The VPLS MAC address withdrawal feature removes MAC addresses that have been dynamically learned, thus providing faster convergence.

A MAC address withdrawal message is send with a MAC list Type Length Value (TLV). 200 MAC addresses are bulked and sent in one Mac TLV message.

> NOTE
> The MAC withdrawal support is only for explicit MAC addresses in MAC withdrawal TLV. Empty MAC list as well as sending MAC withdrawal TLV to specific subset of peers will not be supported.

The maximum number of MAC addresses supported is 5000 in a 5 second interval. The remaining MAC in the AC LIF are not sent. After the 5 second interval, another LIF down event triggers MAC withdrawal message for a new 5 second interval. MAC withdrawal is supported for both VPLS and MCT-VPLS. MPLS signals the MAC withdraw TLV to all the peers.

The **mac-address withdrawal** command enables MAC withdrawal on the bridge domain. The **no** form of the command disables MAC withdrawal.

```
device(config)# bridge-domain 1
device(config-bridge-domain-1)# mac-address withdrawal
device(config-bridge-domain-1)# no mac-address withdrawal
```

For more information about commands, please refer SLX OS Command Reference Guide.

Disabling MAC withdrawal on a bridge domain stops sending of MAC withdraw messages. MAC withdraw messages is received at the receiver end and MAC flush happens even if the feature is not enabled.

The **show bridge-domain** command output is enhanced to displays the MAC withdrawal status.

```
device# show bridge-domain
Bridge-domain 1
------------------------------
Bridge-domain Type: MP , VC-ID: 0
Number of configured end-points:  0 , Number of Active end-points: 0
VE if-indx: 0, Local switching: TRUE, bpdu-drop-enable: TRUE
MAC Withdrawal: Enabled
PW-profile: default, mac-limit: 0
Total VPLS peers: 0 (0 Operational):
device#
```

Enabling statistics on a bridge domain

Follow this procedure to enable statistics on a bridge domain.

1. Enter the global configuration mode.

   ```
   device# configure terminal
   ```

2. Enter the **bridge-domain** command to create a bridge domain at the global configuration level.

   ```
   device(config)# bridge-domain 3
   ```

3. Enter the **statistics** command to enable statistics for all the logical interfaces and peers in bridge domains.

   ```
   device(config-bridge-domain-3)# statistics
   ```

   NOTE

   Use the **no statistics** command to disable statistics on bridge domains.

   ```
   device(config-bridge-domain-3)# no statistics
   ```

Displaying statistics for logical interfaces in bridge domains

Follow the procedure to display statistics' information for logical interfaces in bridge domains.

Enter the **show statistics bridge-domain** command to view the statistics for all logical interfaces and peers on all configured bridge domains.

```
device# show statistics bridge-domain

Bridge Domain 1 Statistics
Interface                    RxPkts          RxBytes          TxPkts          TxBytes
eth 1/1.100       821729          821729          95940360        95940360
eth 1/21.200      884484          885855          95969584        95484555
po  1.300           8884            8855              9684            9955

Bridge Domain 20 Statistics
Interface                    RxPkts          RxBytes          TxPkts          TxBytes
eth 1/6.400       821729          821729          95940360        95940360
eth 1/21.100        8884            8855              9684            9955
po  2.40          884484          885855          95969584        95484555
```

TABLE 10 Output descriptions of the show statistics bridge-domain command

Field	Description
Interface	The interface whose counter statistics are displayed.
RxPkts	The number of packets received at the specified interface.
RxBytes	The number of bytes received at the specified interface.
TxPkts	The number of packets transmitted from the specified interface.
TxBytes	The number of bytes transmitted from the specified interface.

Displaying statistics for a specific bridge domain

Enter the **show statistics bridge-domain** *bd- ID* command to view the statistics for a specific bridge domain. Here *bd-ID* is the specific bridge domain ID.

```
device# show statistics bridge-domain 1

Bridge Domain 1 Statistics
Interface                    RxPkts          RxBytes          TxPkts          TxBytes
eth 1/1.100       821729          821729          95940360         95940360
eth 1/21.200      884484          885855          95969584         95484555
po  1.300           8884            8855              9684             9955
```

Clearing statistics on bridge domains

Follow the procedure to clear statistics' information for logical interfaces in bridge domains.

Enter the **clear statistics bridge-domain** command to clear the statistics for all logical interfaces and peers on all configured bridge domains.

```
device# clear statistics bridge-domain
```

Clearing statistics for a specific bridge domain

Enter the **clear statistics bridge-domain** *bd- ID* command to clear the statistics for a specific bridge domain. Here *bd-ID* is the specific bridge domain ID.

```
device# clear statistics bridge-domain 1
```

Configuration example for VPLS with switching between ACs and network core

VPLS can be configured with switching between attachment circuits (ACs) and the network core. Because VPLS emulates LAN switching, it is considered to be a Layer 2 (L2) service that operates over Layer 3 (L3) clouds.

FIGURE 18 VPLS configuration with switching between attachment circuits (ACs) and network core

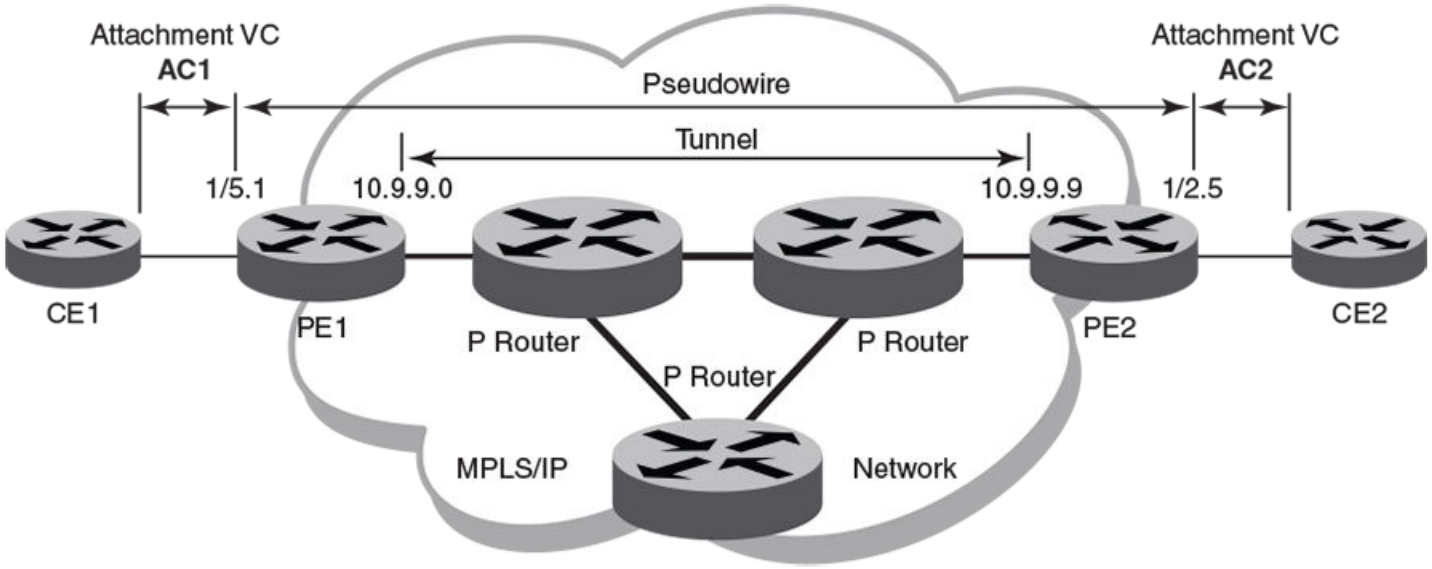

The topology in the preceding figure shows a L2 VPN that enables transport of L2 traffic between two or more native Ethernet networks through an underlying Multiprotocol Label Switching (MPLS) provider network. Customer edge (CE) is the last mile and provider edge (PE) is the first mile node for packets transported towards the provider network. The provider intermediary network is an emulated switch (LAN) or wire (LINE) to the CE. The AC represents the logical link between the CE and PE.

Pseudowire is a circuit emulation infrastructure that extends L2 connectivity from CE1 to CE2 by way of PE1 and PE2. The tunnel is typically a L3 tunnel on which a L2 circuit is emulated.

The following examples show how to configure the provider edge devices (PE1 and PE2) shown in this topology.

PE1

```
device# configure terminal
device(config)# bridge-domain 500 p2mp
device(config-bridge-domain-500)# vc-id 501
device(config-bridge-domain-500)# peer 10.9.9.9 load-balance
device(config-bridge-domain-500)# logical-interface ethernet 1/5.1     ! AC1
device(config-bridge-domain-500)# exit

device(config)# pw-profile default
```

PE2

```
device# configure terminal
device(config)# bridge-domain 300 p2mp
device(config-bridge-domain-300)# vc-id 501
```

```
device(config-bridge-domain-300)# peer 10.9.9.0 load-balance
device(config-bridge-domain-500)# logical-interface ethernet 1/2.5     ! AC2
device(config-bridge-domain-500)# exit

device(config)# pw-profile default
```

802.1d Spanning Tree Protocol

Spanning Tree Protocol overview

The Spanning Tree Protocol (STP) prevents Layer 2 loops in a network by providing redundant links. If a primary link fails, the backup link is activated and network traffic is not affected. STP also ensures that the least cost path is taken when multiple paths exist between ports or VLANs.

The IEEE 802.1d Spanning Tree Protocol (STP) runs on bridges and switches that are 802.1d-compliant.

These variants are Rapid STP (RSTP), Multiple STP (MSTP), Per-VLAN Spanning Tree Plus (PVST+), and Rapid-PVST+ (R-PVST+)

When the spanning tree algorithm is run, the network switches transform the real network topology into a spanning tree topology. In an STP topology any LAN in the network can be reached from any other LAN through a unique path. The network switches recalculate a new spanning tree topology whenever there is a change to the network topology.

For each LAN, the switches that attach to the LAN select a designated switch that is the closest to the root switch. The designated switch forwards all traffic to and from the LAN. The port on the designated switch that connects to the LAN is called the designated port. The switches decide which of their ports is part of the spanning tree. A port is included in the spanning tree if it is a root port or a designated port.

STP runs one spanning tree instance (unaware of VLANs) and relies on long duration forward-delay timers for port state transition between disabled, blocking, listening, learning and forwarding states.

Spanning Tree Protocol configuration notes

Enabling the Spanning Tree Protocol (STP) creates a loop-free topology of Ethernet LANs connected by bridge devices.

The Extreme device supports STP as described in the IEEE 802.1d-1998 specification.

The STP is disabled by default on the Extreme device. Thus, any new VLANs you configure on the Extreme device have STP disabled by default.

Optional features

The following STP configuration features are optional:

- Root guard
- BPDU guard
- PortFast

STP states

Each Layer 2 interface participating in a spanning tree is in one of five states.

A network topology of bridges typically contains redundant connections to provide alternate paths in case of link failures. The redundant connections create a potential for loops in the system. As there is no concept of time to live (TTL) in Ethernet frames, a situation may arise where there is a permanent circulation of frames when the network contains loops. To prevent this, a spanning tree connecting all the bridges is formed in real time.

Every Layer 2 interface running the STP is in one of these states:

State	Action or inaction
Blocking	The interface does not forward frames. Redundant ports are put in a blocking state and enabled when required. This is a transitional state after initialization.
Listening	The interface is identified by the spanning tree as one that should participate in frame forwarding. This is a transitional state after the blocking state for a legacy STP.
Learning	The interface prepares to participate in frame forwarding. This is a transitional state after the blocking state for a legacy STP.
Forwarding	The interface forwards frames. This is a transitional state after the learning state.
Disabled	The interface is not participating in a spanning tree because of shutdown of a port or the port is not operationally up. Any of the other states may transition into this state.

BPDUs

To build a spanning tree for the bridge topology, the bridges must exchange control frames called Bridge Protocol data units (BPDUs).

To construct a spanning tree requires knowledge of the all the participants. The bridges must determine the root bridge and compute the port roles (root, designated, or blocked) with only the information that they have. To ensure that each bridge has enough information, the bridges use BPDUs to exchange information about bridge IDs and root path costs.

A bridge sends a BPDU frame using the unique MAC address of the port itself as a source address, and a destination address of the STP multicast address 01:80:C2:00:00:00.

BPDUs are exchanged regularly (every 2 seconds by default) and enable switches to keep track of network changes and to start and stop forwarding through ports as required.

When a device is first attached to a switch port, it does not immediately forward data. It instead goes through a number of states while it processes inbound BPDUs and determines the topology of the network. When a host is attached, after a listening and learning delay of about 30 seconds, the port always goes into the forwarding state. The time spent in the listening and learning states is determined by the forward delay. However, if instead another switch is connected, the port may remain in blocking mode if it would cause a loop in the network.

There are four types of BPDUs in the original STP specification:

- Configuration BPDU (CBPDU) is used for spanning tree computation.
- Topology Change Notification (TCN) BPDU is used to announce changes in the network topology.
- RSTP BPDU is used for RSTP
- MSTP BPDU is used for MSTP

TCN BPDUs

TCN BPDUs are used to inform other switches of port changes.

TCNs are injected into the network by a non-root switch and propagated to the root. Upon receipt of the TCN, the root switch will set a Topology Change flag in its normal BPDUs. This flag is propagated to all other switches to instruct them to rapidly age out their forwarding table entries.

Consider these configuration rules:

- TCN BPDUs are sent per VLAN.
- TCN BPDUs are sent only in those VLANs in which a topology change is detected.
- TCN BPDUs are sent only in those VLANs for which the bridge is not the root bridge.
- If a topology change is detected on a VLAN for which the bridge is the root bridge, the topology change flag is set in the configuration BPDU that is sent out.

For a given link, in conjunction with the configuration rules, a TCN BPDU is sent out as follows:

- On an access port, only a standard IEEE TCN BPDU is sent out. This TCN BPDU corresponds to a topology change in the access VLAN.
- On a trunk port, if VLAN 1 is allowed (either untagged or untagged), a standard IEEE TCN BPDU is sent for VLAN 1.
- On a trunk port, if the native VLAN is not 1, an untagged TCN BPDU is sent to Cisco or Extreme proprietary MAC address for that VLAN.
- On a trunk port, a tagged TCN BPDU is sent to Cisco or Extreme proprietary MAC address for a tagged VLAN.

As part of the response to TCN BPDUs, the Topology Change and Topology Change Acknowledgment flags are set in all configuration BPDUs corresponding to the VLAN for which the TCN was received.

When a topology change is detected on a trunk port, it is similar to detecting topology changes in each VLAN that is allowed on that trunk port. TCN BPDUs are sent for each VLAN as per the rules.

STP configuration guidelines and restrictions

Follow these configuration guidelines and restrictions when configuring STP and STP variants:

- Only one form of a standing tree protocol, such as STP or RSTP, can be enabled at a time. You must disable one form of xSTP before enabling another.
- When any form of STP is enabled globally, that form of STP is enabled by default on all switch ports.
- LAGs are treated as normal links for any form of STP.
- The STP is disabled by default on the SLX device. Thus, any new VLANs you configure on the SLX device have STP disabled by default.
- PVST/RPVST BPDUs are flooded only if PVST/RPVST is not enabled. STP/RSTP (IEEE) BPDUs are never flooded if STP/RSTP is not enabled.

Understanding the default STP configuration

You should be familiar with STP defaults before you make configuration changes.

TABLE 11 Default STP configuration

Parameter	Default setting
Spanning-tree mode	By default, STP, RSTP, and MSTP are disabled

TABLE 11 Default STP configuration (continued)

Parameter	Default setting
Bridge priority	32768
Bridge forward delay	15 seconds
Bridge maximum aging time	20 seconds
Error disable timeout timer	Disabled
Error disable timeout interval	300 seconds
Port-channel path cost	Standard
Bridge hello time	2 seconds

The following table lists the switch defaults for the interface-specific configuration.

TABLE 12 Default interface specific configuration

Parameter	Default setting
Spanning tree	Enabled on the interface
Automatic edge detection	Disabled
Path cost	2000
Edge port	Disabled
Guard root	Disabled
Hello time	2 seconds
Link type	Point-to-point
Portfast	Disabled
Port priority	128
BPDU restriction	Restriction is disabled.

STP features

The following sections discuss root guard, BPDU guard, and PortFast.

Root guard

Root guard can be used to predetermine a root bridge location and prevent rogue or unwanted switches from becoming the root bridge.

At times it is necessary to protect the root bridge from malicious attack or even unintentional misconfigurations where a bridge device that is not intended to be the root bridge becomes the root bridge, causing severe bottlenecks in the data path. These types of mistakes or attacks can be avoided by configuring root guard on ports of the root bridge.

The root guard feature provides a way to enforce the root bridge placement in the network and allows STP and its variants to interoperate with user network bridges while still maintaining the bridged network topology that the administrator requires. Errors are triggered if any change from the root bridge placement is detected.

When root guard is enabled on a port, it keeps the port in designated FORWARDING state. If the port receives a superior BPDU, which is a root guard violation, it sets the port into a DISCARDING state and triggers a Syslog message and an SNMP trap. No further traffic will be forwarded on this port. This allows the bridge to prevent traffic from being forwarded on ports connected to rogue or wrongly configured STP or RSTP bridges.

Root guard should be configured on all ports where the root bridge should not appear. In this way, the core bridged network can be cut off from the user network by establishing a protective perimeter around it.

Once the port stops receiving superior BPDUs, root guard automatically sets the port back to a FORWARDING state after the timeout period has expired.

BPDU guard

In an STP environment, switches, end stations, and other Layer 2 devices use BPDUs to exchange information that STP will use to determine the best path for data flow.

In a valid configuration, edge port-configured interfaces do not receive BPDUs. If an edge port-configured interface receives a BPDU, an invalid configuration exists, such as the connection of an unauthorized device. The BPDU Guard provides a secure response to invalid configurations because the administrator must manually put the interface back in service.

BPDU guard removes a node that reflects BPDUs back in the network. It enforces the STP domain borders and keeps the active topology predictable by not allowing any network devices behind a BPDU guard-enabled port to participate in STP.

In some instances, it is unnecessary for a connected device, such as an end station, to initiate or participate in an STP topology change. In this case, you can enable the STP BPDU guard feature on the Extreme device port to which the end station is connected. The STP BPDU guard shuts down the port and puts it into an "error disabled" state. This disables the connected device's ability to initiate or participate in an STP topology. A log message is then generated for a BPDU guard violation, and a message is displayed to warn the network administrator of an invalid configuration.

The BPDU Guard provides a secure response to invalid configurations because the administrator must manually put the interface back in service with the **no shutdown** command if error disable recovery is not enabled by enabling the **errdisable-timeout** command. The interface can also be automatically configured to be enabled after a timeout. However, if the offending BPDUs are still being received, the port is disabled again.

Expected behavior in an interface context

When BPDU Guard is enabled on an interface, the device is expected to put the interface in Error Disabled state when BPDU is received on the port when edge-port and BPDU guard is enabled on the switch interface. When the port ceases to receive the BPDUs, it does not automatically switch to edge port mode, you must configure **error disable timeout** or **no shutdown** on the port to move the port back into edge port mode.

Error disable recovery

A port is placed into an error-disabled state when:

- A BPDU guard violation or loop detection violation occurs
- The number of inError packets exceeds the configured threshold
- An EFM-OAM enabled interface receives a critical event from the remote device (functionally equivalent to a disable state)

Once in an error disable state, the port remains in that state until it is re-enabled automatically or manually.

In STP, RSTP, MSTP, PVST+, or R-PVST+ mode, you can specify the time in seconds it takes for an interface to time out. The range is from 10 through 1000000 seconds. The default is 300 seconds. By default, the timeout feature is disabled.

PortFast

PortFast allows an interface to transition quickly to the forwarding state.

Consider the following when configuring PortFast:

- Do not enable PortFast on ports that connect to other devices.
- PortFast only needs to be enabled on ports that connect to workstations or PCs. Repeat this configuration for every port connected to workstations or PCs.
- Enabling PortFast on ports can cause temporary bridging loops, in both trunking and nontrunking mode.
- If BPDUs are received on a PortFast- enabled interface, the interface loses the edge port status unless it receives a **shutdown**/**no shutdown** command.
- PortFast immediately puts the interface into the forwarding state without having to wait for the standard forward time.

STP parameters

The following section discusses bridge parameters.

Bridge parameters

These parameters are set in STP, RSTP, MSTP, PVST+, and R-PVST+.

Bridge priority

Use this parameter to specify the priority of a device and to determine the root bridge.

Each device has a unique bridge identifier called the bridge ID. The bridge ID is an 8 byte value that is composed of two fields: a 2 B bridge priority field and the 6 B MAC address field. The value for the bridge priority ranges from 0 to 61440 in increments of 4096. The default value for the bridge priority is 32768. You use the **bridge-priority** command to set the appropriate values to designate a device as the root bridge or root device. A default bridge ID may appear as 32768.768e.f805.5800. If the bridge priorities are equal, the device with the lowest MAC address is elected the root.

After you decide what device to designate as the root, you set the appropriate device bridge priorities. The device with the lowest bridge priority becomes the root device. When a device has a bridge priority that is lower than that of all the other devices, it is automatically selected as the root.

The root device should be centrally located and not in a "disruptive" location. Backbone devices typically serve as the root because they usually do not connect to end stations. All other decisions in the network, such as which port to block and which port to put in forwarding mode, are made from the perspective of the root device.

You may also specify the bridge priority for a specific VLAN. If the VLAN parameter is not provided, the priority value is applied globally for all per-VLAN instances. However, for the VLANs that have been configured explicitly, the per-VLAN configuration takes precedence over the global configuration.

Bridge Protocol data units (BPDUs) carry information between devices. All the devices in the Layer 2 network, participating in any variety of STP, gather information on other devices in the network through an exchange of BPDUs. As the result of exchange of the BPDUs, the device with the lowest bridge ID is elected as the root bridge

When setting the bridge forward delay, bridge maximum aging time, and the hello time parameters keep in mind that the following relationship should be kept:

```
(2 × (forward-delay − 1)) ≥ max-age ≥ (2 ×(hello-time + 1))
```

Bridge forward delay

The bridge forward delay parameter specifies how long an interface remains in the listening and learning states before the interface begins forwarding all spanning tree instances. The valid range is from 4 through 30 seconds. The default is 15 seconds.

Additionally, you may specify the forward delay for a specific VLAN. If the VLAN parameter is not provided, the bridge forward delay value is applied globally for all per-VLAN instances. However, for the VLANs that have been configured explicitly, the per-VLAN configuration takes precedence over the global configuration.

Bridge maximum aging time

You can use this setting to configure the maximum length of time that passes before an interface saves its BPDU configuration information.

Keeping with the inequality shown above, when configuring the maximum aging time, you must set the value greater than the hello time. The range of values is 6 through 40 seconds while he default is 20 seconds.

You may specify the maximum aging for a specific VLAN. If the VLAN parameter is not provided, the priority value is applied globally for all per-VLAN instances. However, for the VLANs that have been configured explicitly, the per-VLAN configuration takes precedence over the global configuration.

Bridge hello time

You can use this parameter to set how often the device interface broadcasts hello BPDUs to other devices.

Use the **hello-time** command to configure the bridge hello time. The range is from 1 through 10 seconds. The default is 2 seconds.

You may also specify the hello time for a specific VLAN. If the VLAN parameter is not provided, the priority value is applied globally for all per-VLAN instances. However, for the VLANs that have been configured explicitly, the per-VLAN configuration takes precedence over the global configuration.

Error disable timeout parameter

Configure this parameter to enable a timer that brings a port out of the disabled state.

These parameters are be set in STP, RSTP, MSTP, PVST+, and R-PVST+.

When the STP BPDU guard disables a port, the port remains in the disabled state unless the port is enabled manually. The parameter specifies the time in seconds it takes for an interface to time out. The range is from 10 through 1000000 seconds. The default is 300 seconds.

By default, the timeout feature is disabled.

Port-channel path cost parameter

You configure this parameter to specify the port channel path cost.

This parameter can be set in STP, RSTP, MSTP, PVST+, and R-PVST+ mode.

There are two path cost options:

- Custom - Specifies that the path cost changes according to the port channel bandwidth.
- Standard - Specifies that the path cost does not change according to the port channel bandwidth.

The default port cost is standard.

Configuring STP

The following section discusses configuring STP.

Enabling and configuring STP globally

Follow these steps to enable or disable STP and configure STP parameters.

You can enable STP or STP with one or more parameters enabled.

The parameters can be configured individually by:

1. Entering the commands in steps 1 and 2
2. Running the relevant parameter command
3. Verifying the result
4. Saving the configuration

For detailed descriptions of the parameters and features, see the sections STP parameters and STP features.

1. Enter global configuration mode.

   ```
   device# configure terminal
   ```

2. Enable STP globally.

   ```
   device(config)# protocol spanning-tree stp
   ```

   A spanning tree can be disabled by entering the **no protocol spanning-tree stp** command.

3. Describe or name the STP.

   ```
   device(config-stp)# description stp1
   ```

   A description is not required.

4. Specify the bridge priority.

   ```
   device(config-stp)# bridge-priority 4096
   ```

   The bridge with the lowest priority number (highest priority) is designated the root bridge . The range of values is 0 through 61440; values can be set only in increments of 4096. The default priority is 32678.

5. Specify the bridge forward delay.

   ```
   device(config-stp)# forward-delay 20
   ```

   The forward delay specifies how long an interface remains in the listening and learning states before it begins forwarding all spanning tree instances. The valid range is from 4 through 30 seconds. The default is 15 seconds.

6. Configure the maximum aging time.

   ```
   device(config-stp)# max-age 25
   ```

   This parameter controls the maximum length of time that passes before an interface saves its BPDU configuration information. You must set the maximum age to be greater than the hello time. The range is 6 through 40 seconds. The default is 20 seconds.

7. Configure the maximum hello time.

```
device(config-stp)# hello-time 8
```

   The hello time determines how often the switch interface broadcasts hello BPDUs to other devices. The default is 2 seconds while the range is from 1 through 10 seconds.

8. Enable the error disable timeout timer.

```
device(config-stp)# error-disable-timeout enable
```

   This parameter enables a timer that brings a port out of the disabled state. By default, the timeout feature is disabled.

9. Set the error disable timeout timer.

```
device(config-stp)# error-disable-timeout interval 60
```

   When enabled the default is 300 seconds and the range is from 10 through 1000000 seconds.

10. Configure the port channel path cost.

```
device(config-stp)# port-channel path-cost custom
```

   Specifying **custom** means the path cost changes according to the port channel's bandwidth.

11. Return to privileged EXEC mode.

```
device(config-stp)# end
```

12. Verify the configuration.

```
device# show spanning-tree brief

 Spanning-tree Mode: Spanning Tree Protocol

      Root ID      Priority 4096
                   Address 768e.f805.5800
                   Hello Time 8, Max Age 25, Forward Delay 20

      Bridge ID    Priority 4096
                   Address 768e.f805.5800
                   Hello Time 8, Max Age 25, Forward Delay 20

 Interface    Role  Sts  Cost       Prio  Link-type     Edge
 ------------------------------------------------------------
 Eth 0/2      DES   FWD  2000       128   P2P            No
 Eth 0/20     DIS   DIS  20000000   128   P2P            No
 Eth 0/25     DIS   DIS  20000000   128   P2P            No
 Eth 0/30     DIS   DIS  20000000   128   P2P            No
 Eth 0/31     DIS   DIS  2000000    128   P2P            No
```

   Observe that the settings comply with the formula set out in the STP parameter configuration section, as:

   `(2 × (forward delay - 1)) ≥ maximum age ≥ (2 × (hello time + 1))`

   Or in this case `38 ≥ 25 ≥ 18`.

13. Save the configuration.

```
device# copy running-config startup-config
```

STP configuration example

```
device# configure terminal
device(config)# protocol spanning-tree stp
device(config-stp)# description stpForInterface
device(config-stp)# bridge-priority 4096
device(config-stp)# forward-delay 20
device(config-stp)# max-age 25
device(config-stp)# hello-time 8
device(config-stp)# error-disable-timeout enable
device(config-stp)# error-disable-timeout interval 60
device(config-stp)# port-channel path-cost custom
device(config-stp)# end
device# show spanning-tree brief
device# copy running-config startup-config
```

Enabling and configuring STP on an interface

Follow these steps to enable STP and STP features on an interface.

Globally enable STP and STP parameters.

The parameters can be configured individually by:

1. Entering the commands in steps 1-3
2. Running the relevant parameter command
3. Verifying the result
4. Saving the configuration

For detailed descriptions of the parameters and features, see the sections STP parameters and STP features.

1. Enter global configuration mode.

   ```
   device# configure terminal
   ```

2. Enter interface configuration mode.

   ```
   device(config)# interface ethernet 0/20
   ```

3. Enable the interface.

   ```
   device(conf-if-eth-0/20)# no shutdown
   ```

4. Configure the path cost for spanning tree calculations on the interface.

   ```
   device(conf-if-eth-0/20)# spanning-tree cost 10000
   ```

   The lower the path cost means a greater chance that the interface becomes the root port. The range is 1 through 200000000. The default path cost is assigned as per the port speed.

5. Enable BPDU guard on the interface.

   ```
   device(conf-if-eth-0/20)# spanning-tree port-fast bpdu-guard
   ```

   BPDU guard removes a node that reflects BPDUs back in the network. It enforces the STP domain borders and keeps the active topology predictable by not allowing any network devices behind a BPDU guard-enabled port to participate in STP.

6. Configure Root Guard on the interface.

```
device(conf-if-eth-0/20)# spanning-tree guard root
```

   Root Guard protects the root bridge from malicious attacks and unintentional misconfigurations where a bridge device that is not intended to be the root bridge becomes the root bridge.

7. Specify an interface link-type.

```
device(conf-if-eth-0/20)# spanning-tree link-type point-to-point
```

   Specifying a point-to-point link enables rapid spanning tree transitions to the forwarding state. Specifying a shared link disables spanning tree rapid transitions. The default setting is point-to-point.

8. Specify port priority to influence the selection of root or designated ports.

```
device(conf-if-eth-0/20)# spanning-tree priority 64
```

   The range is from 0 through 240 in increments of 16. The default value is 128.

9. Verify the configuration.

```
device# show spanning-tree brief

 Spanning-tree Mode: Spanning Tree Protocol

      Root ID      Priority 4096
                   Address 768e.f805.5800
                   Hello Time 8, Max Age 25, Forward Delay 20

      Bridge ID    Priority 4096
                   Address 768e.f805.5800
                   Hello Time 8, Max Age 25, Forward Delay 20

 Interface    Role  Sts  Cost       Prio  Link-type     Edge
 -------------------------------------------------------------
 Eth 0/2      DES   FWD  2000       128   P2P            No
 Eth 0/20     DES   FWD  1000        64   P2P            No
 Eth 0/25     DIS   DIS  20000000   128   P2P            No
 Eth 0/30     DIS   DIS  20000000   128   P2P            No
 Eth 0/31     DIS   DIS  2000000    128   P2P            No
```

   NOTE

   Observe that the settings comply with the formula set out in the STP parameters section, as:
   `(2 × (forward delay - 1)) ≥ maximum age ≥ (2 × (hello time + 1))`
   or in this case: `38 ≥ 25 ≥ 18`

```
device# show running-config interface ethernet 0/20
interface ethernet 0/20
switchport
switchport mode access
switchport access val 1
spanning-tree cost 1000
spanning-tree guard root
spanning-tree link-type point-to-point
spanning-tree portfast bpdu-guard
spanning-tree priority 64
```

10. Save the settings by copying the running configuration to the startup configuration.

```
device# copy running-config startup-config
```

STP on an interface configuration example

```
device# configure terminal
device(config)# interface ethernet 0/20
device(conf-if-eth-0/20)# no shutdown
device(conf-if-eth-0/20)# spanning-tree cost 10000
device(conf-if-eth-0/20)# spanning-tree port-fast bpdu-guard
device(conf-if-eth-0/20)# spanning-tree guard root
device(conf-if-eth-0/20)# spanning-tree link-type point-to-point
device(conf-if-eth-0/20)# spanning-tree priority 64
device(conf-if-eth-0/20)# end
device# show spanning-tree brief
device# copy running-config startup-config
```

Configuring basic STP parameters

Follow this example to configure basic STP behavior.

1. Enter global configuration mode.

   ```
   device# configure terminal
   ```

2. Enable STP globally

   ```
   device(config)# protocol spanning-tree stp
   ```

3. Name the STP.

   ```
   device(config-stp)# description stp1
   ```

4. Designate the root switch.

   ```
   device(conf-stp)# bridge-priority 28672
   ```

   The priority values can be set only in increments of 4096. The range is 0 through 61440.

5. Specify the bridge forward delay.

   ```
   device(config-stp)# forward-delay 20
   ```

6. Configure the maximum aging time.

   ```
   device(config-stp)# max-age 25
   ```

7. Configure the maximum hello time.

   ```
   device(config-stp)# hello-time 8
   ```

8. Enable the error disable timeout timer.

   ```
   device(config-stp)# error-disable-timeout enable
   ```

9. Set the error disable timeout timer interval.

   ```
   device(config-stp)# error-disable-timeout interval 60
   ```

10. Enable port fast on switch ports.
    a) Configure port fast on Ethernet port 0/1.

```
device(config)# interface ethernet 0/1
device(conf-if-eth-0/1)# spanning-tree portfast
device(conf-if-eth-0/1)# exit
```

    Spanning trees are automatically enabled on switch ports.

    b) Configure port fast on Ethernet port 0/2.

```
device(config)# interface ethernet 0/2
device(conf-if-eth-0/2)# spanning-tree portfast
device(conf-if-eth-0/2)# exit
```

    c) Repeat these commands for every port connected to workstations or PCs.

```
device(config)# interface ethernet ...
```

11. Specify port priorities to influence the selection of the root and designated ports.

```
device(config)# interface ethernet 0/1
device(conf-if-eth-0/1)# spanning-tree priority 1
device(conf-if-eth-0/1)# exit
```

12. Enable the guard root feature.

```
device(config)# interface ethernet 0/12
device(conf-if-eth-0/12)# no shutdown
device(conf-if-eth-0/12)# spanning-tree guard root
```

    Root guard lets the device top participate in the STP but only when the device does not attempt to become the root.

13. Return to privileged exec mode.

```
device(conf-if-eth-0/12)# end
```

14. Verify the configuration.

```
device# show spanning-tree brief
Spanning-tree Mode: Spanning Tree Protocol
Root ID Priority 4096
Address 768e.f805.5800
Hello Time 8, Max Age 25, Forward Delay 20
Bridge ID Priority 4096
Address 768e.f805.5800
Hello Time 8, Max Age 25, Forward Delay 20
Interface     Role     Sts     Cost     Prio     Link-type     Edge
-------------------------------------------------------------------
Eth 0/1       DES      FWD     2000     128      P2P           No
Eth 0/2       DES      FWD     2000     128      P2P           No
Eth 0/12      DES      FWD     2000     128      P2P           No
```

    Observe that the settings comply with the formula set out in the STP parameter configuration section, as:

    `(2 × (forward delay - 1)) ≥ maximum age ≥ (2 × (hello time + 1))`

    or in this case `38 ≥ 25 ≥ 18`.

15. Save the configuration.

```
device# copy running-config startup-config
```

Basic STP configuration example

```
device# configure terminal
device(config)# protocol spanning-tree stp
device(config-stp)# description stp1
device(conf-stp)# bridge-priority 28672
device(config-stp)# forward-delay 20
device(config-stp)# max-age 25
device(config-stp)# hello-time 8
device(config-stp)# error-disable-timeout enable
device(config-stp)# error-disable-timeout interval 60
device(config)# interface ethernet 0/1
device(conf-if-eth-0/1)# spanning-tree portfast
device(conf-if-eth-0/1)# exit
device(config)# interface ethernet 0/2
device(conf-if-eth-0/2)# spanning-tree portfast
device(conf-if-eth-0/2)# exit
device(config)# interface ethernet 0/1
device(conf-if-eth-0/1)# spanning-tree priority 1
device(conf-if-eth-0/1)# exit
device(config)# interface ethernet 0/12
device(conf-if-eth-0/12)# no shutdown
device(conf-if-eth-0/12)# spanning-tree guard root
device(conf-if-eth-0/12)# end
device# show spanning-tree brief
device# copy running-config startup-config
```

Re-enabling an error-disabled port automatically

Enable a port to automatically recover from the error-disabled state after the expiration of an error recovery timer.

1. Enter global configuration mode.

   ```
   device# configure terminal
   ```

2. Enter STP configuration mode.

   ```
   device(config)# protocol spanning-tree stp
   ```

3. Enable the error-disable-timeout timer.

   ```
   device(config-stp)# error-disable-timeout enable
   ```

4. Set an interval after which port shall be enabled.

   ```
   device(config-stp)# error-disable-timeout interval 60
   ```

   The interval range is from 0 to 1000000 seconds, the default is 300 seconds.

5. Return to privileged EXEC mode.

   ```
   device(config-stp)# end
   ```

6. Verify the configuration.

```
device# show spanning-tree
 Spanning-tree Mode: Spanning Tree Protocol

 Root Id: 8000.768e.f805.5800 (self)
 Bridge Id: 8000.768e.f805.5800

 Root Bridge Forward Delay: 15; Hello Time: 2; Max Age: 20; Max-hops: 20
 Configured Forward Delay: 15; Hello Time: 2; Max Age: 20; Max-hops: 20
 Number of topology change(s): 0

 Bpdu-guard errdisable timeout: enabled
 Bpdu-guard errdisable timeout interval: 60 sec
```

Automatically re-enable an error-disabled port configuration example

```
device# configure terminal
device(config)# protocol spanning-tree stp
device(config-stp)# error-disable-timeout enable
device(config-stp)# error-disable-timeout interval 60
device(config-stp)# end
device# show spanning-tree
```

Clearing spanning tree counters

Follow these steps to clear spanning tree counters on all interfaces or on the specified interface.

1. Clear spanning tree counters on all interfaces.

```
device# clear spanning-tree counter
```

2. Clear spanning tree counters on a specified Ethernet interface.

```
device# clear spanning-tree counter interface ethernet 0/3
```

3. Clear spanning tree counters on a specified port channel interface.

```
device# clear spanning-tree counter interface port-channel 12
```

   Port channel interface numbers range from 1 through 64.

Clearing spanning tree-detected protocols

Follow these steps to restart the protocol migration process.

These commands force a spanning tree renegotiation with neighboring devices on either all interfaces or on a specified interface.

1. Restart the spanning tree migration process on all interfaces.

```
device# clear spanning-tree detected-protocols
```

2. Restart the spanning tree migration process on a specific Ethernet interface.

```
device# clear spanning-tree detected-protocols interface ethernet 0/3
```

3. Restart the spanning tree migration process on a specific port channel interface.

```
device# clear spanning-tree detected-protocols port-channel 12
```

   Port channel interface numbers range from 1 through 64.

Shutting down STP

Follow these steps to o shut down STP either globally, on a specific interface, or a specific VLAN.

1. Enter global configuration mode.

```
device# configure terminal
```

2. Shut down STP.
   - Shut down STP globally and return to privileged EXEC mode.

```
device(config)# protocol spanning-tree stp
device(config-stp)# shutdown
device(config-stp)# end
```

   - Shut down STP on a specific interface and return to privileged EXEC mode.

```
device(config)# interface ethernet 1/2
device(conf-if-eth-1/2)# spanning-tree shutdown
device(conf-if-eth-1/2)# end
```

   - Shut down STP on a specific VLAN and return to privileged EXEC mode.

```
device(config)# vlan 10
device(config-vlan-10)# spanning-tree shutdown
device(config-vlan-10)# end
```

3. Verify the configuration.

```
device# show spanning-tree
device#
```

4. Save the running configuration to the startup configuration.

```
device# copy running-config startup-config
```

Shut down STP configuration example

```
device# configure terminal
device(config)# vlan 10
device(config-vlan-10)# spanning-tree shutdown
device(config-stp)# end
device# show spanning-tree
device# copy running-config startup-config
```

NOTE

Shutting down STP on a VLAN is used in this example.

802.1w Rapid Spanning Tree Protocol

Rapid Spanning Tree Protocol overview

The RSTP is a way to provide rapid traffic reconvergence for point-to-point links within a few milliseconds (< 500 milliseconds), following the failure of a bridge or bridge port.

The STP (802.1d) standard was designed at a time when recovering connectivity after an outage within a minute or so was considered adequate performance. With the advent of Layer 3 switching in LAN environments, bridging competes with routed solutions where protocols such as OSPF are able to provide an alternate path in less time.

The RSTP can be seen as evolution of STP standard. It provides rapid convergence of connectivity following the failure of bridge, a bridge port or a LAN. It provides rapid convergence of edge ports, new root ports and port connected through point-to-point links. The port, which qualifies for fast convergence, is derived from the duplex mode of a port. A port operating in full-duplex will be assumed to be point-to-point, while a half-duplex port will be considered as a shared port by default. This automatic setting can be overridden by explicit configuration.

RSTP is designed to be compatible and interoperate with the STP. However, the benefit of the RSTP fast convergence is lost when interacting with legacy STP (802.1d) bridges since the RSTP downgrades itself to the STP when it detects a connection to a legacy bridge.

The states for every Layer 2 interface running the RSTP are as follows:

State	Action
Learning	The interface prepares to participate in frame forwarding.
Forwarding	The interface forwards frames.
Discarding	The interface discards frames. Ports in the discarding state do not take part in the active topology and do not learn MAC addresses.

NOTE
The STP disabled, blocking, and listening states are merged into the RSTP discarding state.

The RSTP port roles for the interface are also different. The RSTP differentiates explicitly between the state of the port and the role it plays in the topology. The RSTP uses the root port and designated port roles defined im the STP, but splits the blocked port role into backup port and alternate port roles:

Backup port	Provides a backup for the designated port and can only exist where two or more ports of the switch are connected to the same LAN; the LAN where the bridge serves as a designated switch.
Alternate port	Serves as an alternate port for the root port providing a redundant path towards the root bridge.

Only the root port and the designated ports are part of the active topology; the alternate and backup ports do not participate in it. When the network is stable, the root and the designated ports are in the forwarding state, while the alternate and backup ports are in the discarding state. When there is a topology change, the new RSTP port roles allow a faster transition of an alternate port into the forwarding state.

For more information about spanning trees, see the introductory sections in the 802.1d Spanning Tree Protocol chapter.

RSTP parameters

The parameters you would normally set when you configure STP are applicable to RSTP. Before you configure RSTP see the STP parameters sections for descriptions of the bridge parameters. the error disable timeout parameter and the port channel path cost parameter.

There is one parameter that can be configured in RSTP that is not available in STP; the transmit hold count. This parameter configures the BPDU burst size by specifying the maximum number of BPDUs transmitted per second for before pausing for 1 second. The range is 1 through 10 while the default is 6. See the section Enabling RSTP and configuring RSTP parameters for the procedure to configure this parameter.

The edge port and auto edge features can be enabled in RSTP as well. See the section Edge port and automatic edge detection and the section Configuring RSTP on an interface for descriptions of these features and how they are configured.

Edge port and automatic edge detection

Configuring the edge port feature makes a port transition directly from initialization to the forwarding state, skipping the listening and learning states.

From an interface, you can configure a device to automatically identify the edge port. The port can become an edge port if no BPDU is received. By default, automatic edge detection is disabled.

Follow these guidelines to configure a port as an edge port:

- When edge port is enabled, the port still participates in a spanning tree.
- A port can become an edge port if no BPDU is received.
- When an edge port receives a BPDU, it becomes a normal spanning tree port and is no longer an edge port.
- Because ports that are directly connected to end stations cannot create bridging loops in the network, edge ports transition directly to the forwarding state and skip the listening and learning states.

  NOTE
  If BPDUs are received on a port fast enabled interface, the interface loses the edge port status unless it receives a **shutdown** or **no shutdown** command.

Configuring RSTP

Enabling and configuring RSTP globally

Follow these steps to enable and configure RSTP.

See the section STP parameters for parameters applicable to all STP variants.

You can enable RSTP or RSTP with one or more parameters enabled. The parameters can be enabled or changed individually by entering the commands in steps 1 and 2, running the parameter command, verifying the result, and then saving the configuration.

1. Enter global configuration mode.

```
device# configure terminal
```

2. Enable RSTP.

```
device(config)# protocol spanning-tree rstp
```

   You can shut down RSTP by entering the **shutdown** command.

3. Designate the root device.

```
device(conf-rstp)# bridge-priority 28582
```

   The range is 0 through 61440 and the priority values can be set only in increments of 4096.

   You can shut down RSTP by entering the **shutdown** command when in RSTP configuration mode.

4. Configure the bridge forward delay value.

```
device(conf-rstp)# forward-delay 15
```

5. Configure the bridge maximum aging time value.

```
device(conf-rstp)# max-age 20
```

6. Enable the error disable timeout timer.
   a) Enable the timer.

```
device(conf-rstp)# error-disable-timeout enable
```

   b) Configure the error disable timeout interval value.

```
device(conf-rstp)# error-disable-timeout interval 60
```

7. Configure the port-channel path cost.

```
device(conf-rstp)# port-channel path-cost custom
```

8. Configure the bridge hello-time value.

```
device(conf-rstp)# hello-time 2
```

9. Specify the transmit hold count.

```
device(config-rstp)# transmit-holdcount 5
```

   This command configures the maximum number of BPDUs transmitted per second.

10. Return to privileged exec mode.

```
device(conf-rstp)# end
```

11. Verify the configuration

```
device# show spanning-tree

 Spanning-tree Mode: Rapid Spanning Tree Protocol

 Root Id: 8000.01e0.5200.0180 (self)
 Bridge Id: 8000.01e0.5200.0180

 Root Bridge Forward Delay: 15; Hello Time: 2; Max Age: 20; Max-hops: 20
 Configured Forward Delay: 15; Hello Time: 2; Max Age: 20; Max-hops: 20
 Number of topology change(s): 0

 Bpdu-guard errdisable timeout: enabled
 Bpdu-guard errdisable timeout interval: 60 sec
```

    Observe that the settings comply with the formula set out in the STP parameters section, as:

    (2 × (forward delay - 1)) ≥ maximum age ≥ (2 × (hello time + 1))

    or in this case: 28 ≥ 20 ≥ 6.

12. Save the configuration.

```
device# copy running-config startup-config
```

Enabling RSTP and configuring RSTP parameters example

```
device# configure terminal
device(config)# protocol spanning-tree rstp
device(conf-rstp)# bridge-priority 28582
device(conf-rstp)# forward-delay 20
device(conf-rstp)# max-age 25
device(conf-rstp)# error-disable-timeout enable
device(conf-rstp)# error-disable-timeout interval 60
device(conf-rstp)# port-channel path-cost custom
device(conf-rstp)# hello-time 5
device(conf-rstp)# transmit-holdcount 5
device(conf-rstp# end
device# show spanning-tree
device# copy running-config startup-config
```

Enabling and configuring RSTP on an interface

Follow these steps to configure RSTP on an Ethernet interface.

You can configure the parameters individually on an interface by doing the following:

1. Entering the commands in Steps 1 through 3.
2. Specifying additional parameters, as appropriate.
3. Verifying the result.
4. Saving the configuration.

For detailed descriptions of the parameters and features, see the sections STP parameters and STP features.

1. Enter global configuration mode.

```
device# configure terminal
```

2. Enter interface subtype configuration mode.

```
device(config)# interface ethernet 0/10
```

3. Enable the interface.

```
device(conf-if-eth-0/10)# no shutdown
```

To disable the spanning tree on the interface you use the **spanning-tree shutdown** command.

4. Specify the port priority on the interface.

```
device(conf-if-eth-0/10)# spanning-tree priority 128
```

The range is from 0 through 240 in increments of 16. The default value is 128.

5. Specify the path cost on the interface.

```
device(conf-if-eth-0/10)# spanning-tree cost 20000000
```

The lower the path cost means a greater chance that the interface becomes the root port. The range is 1 through 200000000. The default path cost is assigned as per the port speed.

6. Enable edge port.

```
device(conf-if-eth-0/10)# spanning-tree edgeport
```

If BPDUs are received on a port fast enabled interface, the interface loses the edge port status unless it receives a **shutdown** or **no shutdown** command.

7. Enable BPDU guard on the interface.

```
device(conf-if-eth-0/10)# spanning-tree edgeport bpdu-guard
```

BPDU guard removes a node that reflects BPDUs back in the network. It enforces the STP domain borders and keeps the active topology predictable by not allowing any network devices behind a BPDU guard-enabled port to participate in STP.

8. Enable automatic edge detection on the interface.

```
device(conf-if-eth-0/10)# spanning-tree autoedge
```

You use this command to automatically identify the edge port. A port becomes an edge port if it receives no BPDUs. By default, automatic edge detection is disabled.

9. Enable root guard on the interface.

```
device(conf-if-eth-0/10)# spanning-tree guard root
```

Root guard protects the root bridge from malicious attacks and unintentional misconfigurations where a bridge device that is not intended to be the root bridge becomes the root bridge.

10. Specify a link type on the interface.

```
device(conf-if-eth-0/10)# spanning-tree link-type point-to-point
```

> NOTE
> The link type is explicitly configured as **point-to-point** rather than **shared**.

11. Return to privileged EXEC mode.

```
device(conf-if-eth-0/10)# end
```

12. Verify the configuration.

```
device# show spanning-tree

 Spanning-tree Mode: Rapid Spanning Tree Protocol

 Root Id: 8000.01e0.5200.0180 (self)
 Bridge Id: 8000.01e0.5200.0180

 Root Bridge Forward Delay: 15; Hello Time: 2; Max Age: 20; Max-hops: 20
 Configured Forward Delay: 15; Hello Time: 2; Max Age: 20; Max-hops: 20
 Number of topology change(s): 0

 Bpdu-guard errdisable timeout: enabled
 Bpdu-guard errdisable timeout interval: 60 sec


 Port Eth 0/10 enabled
    Ifindex: 201359360; Id: 8002; Role: Disabled; State: Disabled
    Designated Path Cost: 0
    Configured Path Cost: 20000000
    Designated Port Id: 0; Port Priority: 128
    Designated Bridge: 0000.0000.0000.0000
    Number of forward-transitions: 0
    Version: Spanning Tree Protocol - Received None - Sent STP
    Edgeport: on; AutoEdge: yes; AdminEdge: no; EdgeDelay: 3 sec
    Configured Root guard: on; Operational Root guard: on
    Bpdu-guard: on
    Link-type: point-to-point
    Received BPDUs: 0; Sent BPDUs: 0

Interface   Role     Sts     Cost       Prio     Link-type     Edge
--------------------------------------------------------------------
Eth 0/10    DES      FWD     20000000   128      P2P           No
```

The **forward-delay**, **hello-time**, and **max-age** parameters are set globally, not on the interface.

Observe that the settings comply with the formula set out in the STP parameters section, as:

(2 × (forward delay - 1)) ≥ maximum age ≥ (2 × (hello time + 1))

or in this case: 28 ≥ 20 ≥ 6.

13. Save the configuration.

```
device# copy running-config startup-config
```

RSTP on an interface configuration example

```
device# configure terminal
device(config)# interface ethernet 0/10
device(conf-if-eth-0/10)# no spanning-tree shutdown
device(conf-if-eth-0/10)# spanning-tree priority 128
device(conf-if-eth-0/10)# spanning-tree cost 20000000
device(conf-if-eth-0/10)# spanning-tree edgeport
device(conf-if-eth-0/10)# spanning-tree edgeport bpdu-guard
device(conf-if-eth-0/10)# spanning-tree autoedge
device(conf-if-eth-0/10)# spanning-tree guard root
device(conf-if-eth-0/10)# spanning-tree link-type point-to-point
device(conf-if-eth-0/10)# end
device# show spanning-tree
device# copy running-config startup-config
```

Configuring a basic RSTP

Follow these steps to configure a basic RSTP.

1. Enter global configuration mode.

```
device# configure terminal
```

2. Enable RSTP.

```
device(config)# protocol spanning-tree rstp
```

3. Designate the root device.

```
device(conf-rstp)# bridge-priority 28582
```

4. Enable the error disable timeout timer value.

```
device(conf-rstp)# error-disable-timeout enable
```

5. Configure the error-disable-timeout interval value.

```
device(conf-rstp)# error-disable-timeout interval 60
```

6. Enable edge port on switch ports.
   a) Enter interface configuration mode for the switchport.

```
device(conf-rstp)# interface ethernet 1/10
```

   b) Enable edge port.

```
device(conf-if-eth-1/10)# spanning-tree edge-port
```

   c) Return to global configuration mode.

```
device(conf-if-eth-1/10)# exit
```

   d) Repeat these steps for all ports that connect to a workstation or PC.
7. Specify port priorities.
   a) Enter interface configuration mode.

```
device(config)# interface ethernet 1/11
```

   b) Configure the port priority.

```
device(conf-if-eth-1/11)# spanning-tree priority 1
```

   c) Return to global configuration mode.

```
device(conf-if-eth-1/11)# exit
```

8. Enable the guard root feature.
   a) Enter interface configuration mode.

```
device(config)# interface ethernet 1/1
```

   b) Configure the port priority.

```
device(conf-if-eth-1/1)# spanning-tree guard root
```

   c) Return to privileged EXEC mode.

```
device(conf-if-eth-1/1)# exit
```

9. Verify the configuration.

```
device# show spanning-tree

Spanning-tree Mode: Rapid Spanning Tree Protocol

Root Id: 4096.01e0.5200.0180 (self)
Bridge Id: 4096.01e0.5200.0180
Root Bridge Forward Delay: 15; Hello Time: 2; Max Age: 20; Max-hops: 20
Configured Forward Delay: 15; Hello Time: 2; Max Age: 20; Max-hops: 20
Number of topology change(s): 0
Bpdu-guard errdisable timeout: disabled
Bpdu-guard errdisable timeout interval: 300 sec
switch# show spanning-tree brief
Spanning-tree Mode: Rapid Spanning Tree Protocol
Root ID Priority 4096
Address 768e.f805.5800
Hello Time 2, Max Age 20, Forward Delay 15
Bridge ID Priority 4096
Address 768e.f805.5800

Interface   Role     Sts     Cost     Prio     Link-type     Edge
-------------------------------------------------------------------
Eth 1/1     DES      FWD     2000     128      P2P           No
Eth 1/10    DES      FWD     2000     128      P2P           No
Eth 1/11    DES      FWD     2000     128      P2P           No
```

   Observe that the settings comply with the formula set out in the STP parameters section, as:

   (2 × (forward delay - 1)) ≥ maximum age ≥ (2 × (hello time + 1))

   or in this case: 28 ≥ 20 ≥ 6.

10. Save the configuration.

```
device# copy running-config startup-config
```

Basic RSTP configuration example

```
device# configure terminal
device(config)# protocol spanning-tree rstp
device(conf-rstp)# bridge-priority 28582
device(conf-rstp)# error-disable-timeout enable
device(conf-rstp)# error-disable-timeout interval 60
device(conf-rstp)# interface etherne 1/10
device(conf-if-eth-1/10)# spanning-tree edge-port
device(conf-if-eth-1/10)# exit
device(config)# interface ethernet 1/11
device(conf-if-eth-1/11)# spanning-tree priority 1
device(conf-if-eth-1/11)# exit
device(config)# interface ethernet 1/1
device(conf-if-eth-1/1)# spanning-tree guard root
device(conf-if-eth-1/1)# exit
device# show spanning-tree
device# copy running-config startup-config
```

Clearing spanning tree counters

Follow these steps to clear spanning tree counters on all interfaces or on the specified interface.

1. Clear spanning tree counters on all interfaces.

   ```
   device# clear spanning-tree counter
   ```

2. Clear spanning tree counters on a specified Ethernet interface.

   ```
   device# clear spanning-tree counter interface ethernet 0/3
   ```

3. Clear spanning tree counters on a specified port channel interface.

   ```
   device# clear spanning-tree counter interface port-channel 12
   ```

   Port channel interface numbers range from 1 through 64.

Clearing spanning tree-detected protocols

Follow these steps to restart the protocol migration process.

These commands force a spanning tree renegotiation with neighboring devices on either all interfaces or on a specified interface.

1. Restart the spanning tree migration process on all interfaces.

   ```
   device# clear spanning-tree detected-protocols
   ```

2. Restart the spanning tree migration process on a specific Ethernet interface.

   ```
   device# clear spanning-tree detected-protocols interface ethernet 0/3
   ```

3. Restart the spanning tree migration process on a specific port channel interface.

   ```
   device# clear spanning-tree detected-protocols port-channel 12
   ```

   Port channel interface numbers range from 1 through 64.

Shutting down RSTP

Follow these steps to o shut down RSTP either globally, on a specific interface, or a specific VLAN.

1. Enter global configuration mode.

```
device# configure terminal
```

2. Shut down RSTP.
   - Shut down STP globally and return to privileged EXEC mode.

```
device(config)# protocol spanning-tree rstp
device(conf-rstp)# shutdown
device(conf-rstp)# end
```

   - Shut down RSTP on a specific interface and return to privileged EXEC mode.

```
device(config)# interface ethernet 1/2
device(conf-if-eth-1/2)# spanning-tree shutdown
device(conf-if-eth-1/2)# end
```

   - Shut down RSTP on a specific VLAN and return to privileged EXEC mode.

```
device(config)# vlan 10
device(config-vlan-10)# spanning-tree shutdown
device(config-vlan-10)# end
```

3. Verify the configuration.

```
device# show spanning-tree
device#
```

4. Save the configuration.

```
device# copy running-config startup-config
```

Shut down RSTP configuration example

```
device# configure terminal
device(config)# vlan 10
device(config-vlan-10)# spanning-tree shutdown
device(config-vlan-10)# end
device# show spanning-tree
device# copy running-config startup-config
```

NOTE

Shutting down RSTP on a VLAN is used in this example.

Per-VLAN Spanning Tree+ and Rapid Per-VLAN Spanning Tree+

PVST+ and R-PVST+ overview

The Per-VLAN Spanning Tree Plus (PVST+) protocol runs a spanning tree instance for each VLAN in the network. The version of PVST+ that uses the RSTP state machine is called Rapid-PVST Plus (R-PVST+). R-PVST+ has one instance of spanning tree for each VLAN on the device.

Both the STP and the RSTP build a single logical topology. A typical network has multiple VLANs. A single logical topology does not efficiently utilize the availability of redundant paths for multiple VLANs. A single logical topology does not efficiently utilize the availability of redundant paths for multiple VLANs. If a port is set to the blocked state or the discarding state for one VLAN (under the STP or the RSTP), it is the same for all other VLANs. PVST+ builds on the STP on each VLAN, and R-PVST+ builds on the RSTP on each VLAN.

PVST+ R-PVST+ provide interoperability with Cisco PVST and R-PVST and other vendor switches which implement Cisco PVST or R-PVST. the PVST+ and R-PVST+ implementations are extensions to PVST and R-PVST, which can interoperate with an STP topology, including MSTP (CIST), on Extreme and other vendor devices sending untagged IEEE BPDUs.

PVST+ and R-PVST+ guidelines and restrictions

Consider the following when configuring PVST+ and R-PVST+:

- Extreme supports PVST+ and R-PVST+ only. The PVST and R-PVST protocols are proprietary to Cisco and are not supported.
- A port native VLAN is the native VLAN ID associated with a trunk port on an Extreme switch. This VLAN ID is associated with all untagged packets on the port. The default native VLAN ID for a trunk port is 1.
- IEEE compliant switches run just one instance of STP protocol shared by all VLANs, creating a Mono Spanning Tree (MST). A group of such switches running a single spanning tree forms an MST region.
- You can configure up to 128 PVST+ or R-PVST+ instances. If you have more than 128 VLANs configured on the switch and enable PVST then the first 128 VLANs are PVST/+ or R-PVST+ enabled.
- In PVST/+ or R-PVST+ mode, when you are connected to a Cisco or MLX switch, the Cisco proprietary MAC address to which the BPDUs are sent/processed must be explicitly configured on a per-port basis.
- In PVST/+ or R-PVST+ mode, when you connect to a Cisco switch using a trunk port, the Extreme switch must have a native VLAN configured on the trunk port (same configuration as on the other side).
- A Common Spanning Tree (CST) is the single spanning tree instance used by Extreme switches to interoperate with 802.1q bridges. This spanning tree instance stretches across the entire network domain (including PVST, PVST+ and 802.1q regions). It is associated with VLAN 1 on the Extreme switch.
- In order to interact with STP and IEEE 802.1q trunk, PVST evolved to PVST+ to interoperate with STP topology by STP BPDU on the native or default VLAN.
- A group of switches running PVST+ is called a PVST+ region.

For more information about spanning trees, see the introductory sections in the Spanning Tree Protocol chapter.

PVST+ and R-PVST+ parameters

The parameters you would normally set when you configure STP are applicable to PVST+ and R-PVST+. Before you configure PVST+ or R-PVST+ parameters see the sections in the Standing Tree Protocol chapter explaining bridge parameters, the error disable timeout parameter and the port channel path cost parameter.

There is one parameter that can be configured in R-PVST+ that is not available in STP or PVST+; the transmit hold count. This parameter configures the BPDU burst size by specifying the maximum number of BPDUs transmitted per second for before pausing for 1 second. The range is 1 through 10 while the default is 6. See the section Configuring R-PVST+ for the procedure to configure this parameter.

Bridge protocol data units in different VLANs

PVST+ uses the spanning tree instance for VLAN 1 to join the CST in the network to build the CST, PVST+ processes and sends standard IEEE Bridge protocol data units (BPDUs) on all the ports in VLAN 1 (access/trunk).

Across IEEE 802.1q trunks, Extreme switches run PVST+. The goal is to interoperate with standard IEEE STP (or RSTP or MSTP), while transparently tunneling PVST+ instance BPDUs across the MST region to potentially connect to other Extreme switches across the MST region.

On trunk ports that allow VLAN 1, PVST+ also sends PVST+ BPDUs to a Cisco-proprietary multicast MAC address (0100.0ccc.cccd) or Extreme-proprietary multicast MAC address (0304.0800.0700) depending on the configuration. By default, the PVST+ BPDUs are sent to Extreme-proprietary multicast MAC address on Extreme switches. These BPDUs are tunneled across an MST region. The PVST + BPDUs for VLAN 1 are only used for the purpose of consistency checks and that it is only the IEEE BPDUs that are used for building the VLAN 1 spanning tree. So in order to connect to the CST, it is necessary to allow VLAN 1 on all trunk ports.

For all other VLANs, PVST+ BPDUs are sent on a per-VLAN basis on the trunk ports. These BPDUs are tunneled across an MST region. Consequently, for all other VLANs, MST region appears as a logical hub. The spanning tree instances for each VLAN in one PVST+ region map directly to the corresponding instances in another PVST+ region and the spanning trees are calculated using the per-VLAN PVST+ BPDUs.

Similarly, when a PVST+ region connects to a MSTP region, from the point of view of MSTP region, the boundary bridge thinks it is connected to a standard IEEE compliant bridge sending STP BPDUs. So it joins the CIST of the MSTP region to the CST of the PVST+ region (corresponding to VLAN 1). The PVST+ BPDUs are tunneled transparently through the MSTP region. So from the Extreme bridge point of view, the MSTP region looks like a virtual hub for all VLANs except VLAN 1.

The PVST+ BPDUs are sent untagged for the native VLAN and tagged for all other VLANs on the trunk port.

On access ports, Extreme switches run classic version of IEEE STP/RSTP protocol, where the BPDUs are sent to the standard IEEE multicast address "0180.C200.0000". So if we connect a standard IEEE switch to an access port on the Extreme switch, the spanning tree instance (corresponding to the access VLAN on that port) of the Extreme switch is joined with the IEEE STP instance on the adjacent switch.

For introductory information about STP BPDUs, see the section BPDUs on page 110.

BPDU configuration notes

In order to build a spanning tree for the bridge topology, the bridges must exchange control frames. These frames are called Bridge Protocol data units (BPDU).

BPDUs are sent to a Cisco-proprietary multicast MAC address 0100.0ccc.cccd or Extreme-proprietary multicast MAC address 0304.0800.0700.By default, the PVST+ BPDUs are sent to Extreme-proprietary multicast MAC address on Extreme switches. These are called SSTP (Single Spanning Tree Protocol) BPDUs. The format of the SSTP BPDU is nearly identical to the 802.1d BPDU after the SNAP header, except that a type-length-value (TLV) field is added at the end of the BPDU. The TLV has 2 bytes for type (0x0), 2 bytes for length, and 2 bytes for the VLAN ID. See Extreme BPDU PVST+ headers/fields on page 137 and BPDU R-PVST+ header and field comparisons on page 137 for an outline of the BPDU header content.

Topology Change Notification (TCN) BPDUs are used to inform other switches of port changes. TCNs are injected into the network by a non-root switch and propagated to the root. Upon receipt of the TCN, the root switch will set a Topology Change flag in its normal BPDUs. This flag is propagated to all other switches to instruct them to rapidly age out their forwarding table entries.

In PVST+, three types of TCN BPDUs are sent out depending on the type of the link. See Extreme PVST+ TCN BPDU headers/fields on page 139 and Cisco PVST TCN BPDU headers/fields on page 140.

- Standard IEEE TCN BPDU.
- Untagged TCN BPDU sent to the Cisco/Extreme proprietary MAC address.
- Tagged TCN BPDU sent to the Cisco/Extreme proprietary MAC address.

BPDU R-PVST+ header and field comparisons

These tables outline the differences between Extreme R-PVST+ BPDU and Cisco R-PVST+ BPDU header fields.

Extreme R-PVST+ BPDU headers/fields

Header/field	Standard IEEE STP/RSTP BPDU (64B padded)	R-PVST+ untagged BPDU (64B padded)	R-PVST+ tagged BPDU (72B padded)
Source Address (MAC SA)	6B	6B	6B
Destination Address (MAC DA)	0180C2.000000 (6B)	030408.000700 (6B)	030408.000700 (6B)
Length	2B	2B	-
Type	-	-	81 00 (2B)
802.1q tag	-	-	4B
Source Service Access Point (SSAP)	42	AA 03	AA 03
Destination Service Access Point (DSAP)	42	AA	AA
Extreme Organizationally Unique Identifier (OUI)	-	02 04 08	02 04 08
PVST PID	-	01 0B	01 0B
Logical Link Control (LLC)	3B	+	+
SubNetwork Access Protocol (SNAP)	-	Yes (2B)	Yes (2B)
IEEE BPDU INFO	35B	35B	35B
Type, Length, Value (TLV) Pad Type Length	-	6B 00 (1B) 00 00 00 02	6B 00 (1B) 00 00 00 02

Header/field	Standard IEEE STP/RSTP BPDU (64B padded)	R-PVST+ untagged BPDU (64B padded)	R-PVST+ tagged BPDU (72B padded)
VLAN ID		2B	2B

Cisco R-PVST+ BPDU headers/fields

Header/field	Standard IEEE STP/RSTP BPDU (64B padded)	R-PVST+ untagged BPDU (64B padded)	R-PVST+ tagged BPDU (72B padded)
MAC SA	6B	6B	6B
MAC DA	0180C2.000000 (6B)	01000C.CCCCCD (6B)	010002.CCCCCD (6B)
Length	2B	2B	-
Type	-	-	81 00 (2B)
802.1q tag	-	-	4B
SSAP	42 03	AA 03	AA 03
DSAP	42	AA	AA
Cisco OUI	-	00 00 0C	00 00 0C
PVST PID	-	01 0B	01 0B
LLC	3B	+	+
SNAP	-	Yes	Yes
IEEE BPDU INFO	35B	35B	35B
TLV Pad Type Length VLAN ID	-	6B 00 (1B) 00 00 00 02 2B	6B 00 (1B) 00 00 00 02 2B

Sent BPDUs

On an 802.1q trunk, the PVST+ enabled switch sends the following BPDUs:

1. For all tagged VLANs on the port on which PVST+ is enabled, 802.1q tagged SSTP BPDUs are sent to the Cisco or Extreme MAC address. The 802.1q tag contains the VLAN ID. (VLAN 1 could be tagged on the port. In that case a tagged BPDU for VLAN 1 is sent). The IEEE compliant switches do not consider these BPDUs as a control packet. So they forward the frame as they would forward to any unknown multicast address on the specific VLAN.
2. If PVST+ is enabled on the untagged (native) VLAN of the port, an untagged SSTP BPDU is sent to the Extreme or Cisco MAC address on the native VLAN of the trunk. It is possible that the native VLAN on the Extreme or Cisco port is not VLAN 1. This BPDU is also forwarded on the native VLAN of the IEEE 802.1q switch just like any other frame sent to an unknown multicast address.
3. In addition to the above SSTP BPDUs, a standard IEEE BPDU (802.1d) is also sent, corresponding to the information of VLAN 1 on the Extreme or Cisco switch. This BPDU is not sent if VLAN 1 is explicitly disabled on the trunk port.

The following table lists the types of BPDUs sent in case of different port types. The numbers in the third column are the VLAN instance for which these BPDUs are sent/processed.

TABLE 13 Types of BPDUs sent for different port types

Port Configuration	Extreme or Cisco - PVST(+)	VLAN instance
Access - VLAN 1	Standard IEEE BPDU (64B)	1
Access - VLAN 100	Standard IEEE BPDU (64B)	100
Trunk - Native VLAN 1 Allowed VLANs - 1, 100, 200	Standard IEEE BPDU (64B) Extreme or Cisco untagged BPDU (68B) Extreme or Cisco tagged BPDU (72B) Extreme or Cisco tagged BPDU (72B)	1 1 100 200
Trunk - Native VLAN 100 Allowed VLANs - 1, 100, 200	Standard IEEE BPDU (64B) Extreme or Cisco untagged BPDU (68B) Extreme or Cisco tagged BPDU (72B Extreme or Cisco tagged BPDU (72B)	1 100 1 200
Trunk - Native VLAN 100 Allowed VLANs - 100	Extreme or Cisco untagged BPDU (68B)	100
Trunk - Native VLAN 100 Allowed VLANs - 100, 200	Extreme or Cisco untagged BPDU (68B) Extreme or Cisco tagged BPDU (72B)	100 200

TCN headers and fields

Since PVST+ is based on STP, and Rapid-PVST+ is based on RSTP, TCN BPDUs are sent only in PVST+ and not in Rapid-PVST+ mode.

For introductory information about STP BPDUs, see the section TCN BPDUs on page 111.

Extreme PVST+ TCN BPDU headers/fields

Header/field	Standard IEEE STP TCN BPDU (64B with padding)	PVST+ untagged TCN BPDU (64B with padding)	PVST+ tagged TCN BPDU (68B with padding)
MAC SA	6B	6B	6B
MAC DA	0180C2.000000 (6B)	030408.000700 (6B)	030408.000700 ((6B)
Length	2B	2B	-
Type	-	-	81 00 (2B)
802.1q tag	-	-	4B
SSAP	42 03	AA 03	AA 03
DSAP	42	AA	AA
Cisco OUI	-	02 04 08	02 04 08
PVST PID	-	01 0B	01 0B
LLC	3B	8B	8B
SNAP	4B	Entire BPDU with type = TCN 35B	Entire BPDU with type = TCN 35B

Cisco PVST TCN BPDU headers/fields

Header/field	Standard IEEE STP TCN BPDU (64B padded)	PVST untagged TCN BPDU (64B padded)	PVST tagged TCN BPDU (68B padded)
MAC SA	6B	6B	6B
MAC DA	0180C2.000000 (6B)	01000C.CCCCCD (6B)	01000C.CCCCCD (6B)
Length	2B	2B	-
Type	-	-	81 00 (2B)
802.1q tag	-	-	4B
SSAP	42 03	AA 03	AA 03
DSAP	42	AA	AA
Cisco OUI	-	00 00 0C	00 00 0C
PVST PID	-	01 0B	01 0B
LLC	3B	8B	8B
SNAP	-	Yes	Yes
IEEE TCN BPDU INFO	4B	Entire BPDU with type = TCN 35B	Entire BPDU with type = TCN 35B

PortFast

PortFast allows an interface to transition quickly to the forwarding state.

Consider the following when configuring PortFast:

- Do not enable PortFast on ports that connect to other devices.
- PortFast only needs to be enabled on ports that connect to workstations or PCs. Repeat this configuration for every port connected to workstations or PCs.
- Enabling PortFast on ports can cause temporary bridging loops, in both trunking and nontrunking mode.
- If BPDUs are received on a PortFast- enabled interface, the interface loses the edge port status unless it receives a **shutdown**/**no shutdown** command.
- PortFast immediately puts the interface into the forwarding state without having to wait for the standard forward time.

Edge port and automatic edge detection

Configuring the edge port feature makes a port transition directly from initialization to the forwarding state, skipping the listening and learning states.

From an interface, you can configure a device to automatically identify the edge port. The port can become an edge port if no BPDU is received. By default, automatic edge detection is disabled.

Follow these guidelines to configure a port as an edge port:

- When edge port is enabled, the port still participates in a spanning tree.
- A port can become an edge port if no BPDU is received.
- When an edge port receives a BPDU, it becomes a normal spanning tree port and is no longer an edge port.
- Because ports that are directly connected to end stations cannot create bridging loops in the network, edge ports transition directly to the forwarding state and skip the listening and learning states.

NOTE
If BPDUs are received on a port fast enabled interface, the interface loses the edge port status unless it receives a **shutdown** or **no shutdown** command.

Configuring PVST+ and R-PVST+

Enabling and configuring PVST+ globally

Use this procedure to enable and set parameters for PVST+.

You can enable PVST+ with one or more parameters configured. The parameters can be configured or changed individually by entering the commands in steps 1 and 2, running the parameter command, verifying the result, and then saving the configuration.

For more information about spanning trees and spanning tree parameters, see the introductory sections in the Spanning Tree Protocol chapter.

1. Enter global configuration mode.

```
device# configure terminal
```

2. Enable PVST+.

```
device(config)# protocol spanning-tree pvst
```

3. Configure the bridge priority for the common instance.

```
device(config-pvst)# bridge-priority 4096
```

   Valid values range from 0 through 61440 in increments of 4096. Assigning a lower priority value indicates that the bridge might become root.

   You can shut down PVST+ by entering the **shutdown** command when in PVST configuration mode.

4. Configure the forward delay parameter.

```
device(config-pvst)# forward-delay  11
```

5. Configure the hello time parameter.

```
device(config-pvst)# hello-time 2
```

6. Configure the maximum age parameter.

```
device(config-pvst)# max-age 7
```

7. Return to privileged exec mode.

```
device(config-pvst)# end
```

8. Verify the configuration.

```
device#  show spanning-tree brief
VLAN 1

 Spanning-tree Mode: PVST Protocol

      Root ID      Priority 4097
                   Address 01e0.5200.0180
                   Hello Time 2, Max Age 7, Forward Delay 11

      Bridge ID    Priority 4097
                   Address 01e0.5200.0180
                   Hello Time 2, Max Age 7, Forward Delay 11

 Interface     Role  Sts  Cost       Prio  Link-type        Edge
 ------------------------------------------------------------------

VLAN 100

 Spanning-tree Mode: PVST Protocol

      Root ID      Priority 4196
                   Address 01e0.5200.0180
                   Hello Time 2, Max Age 7, Forward Delay 11

      Bridge ID    Priority 4196
                   Address 01e0.5200.0180
                   Hello Time 2, Max Age 7, Forward Delay 11

 Interface     Role  Sts  Cost       Prio  Link-type        Edge
 ------------------------------------------------------------------
```

   Observe that the settings comply with the formula set out in the STP parameters section, as:

   (2 × (forward delay - 1)) ≥ maximum age ≥ (2 × (hello time + 1))

   or in this case: 20 ≥7 ≥ 6.

9. Save the configuration.

```
device# copy running-config startup-config
```

PVST+ configuration example

```
device# configure terminal
device(config)# protocol spanning-tree pvst
device(config-pvst)# bridge-priority 4096
device(config-pvst)# forward-delay 11
device(config-pvst)# hello-time 2
device(config-pvst)# max-age 7
device(config-pvst)# end
device# show spanning-tree brief
device# copy running-config startup-config
```

For more information about configuring PVST+ parameters, see STP parameters on page 114. PVST+, R-PVST+, and other types of spanning trees share many tasks with STP.

Enabling and configuring PVST+ on an interface

Follow these steps to enable and configure PVST+ on an interface.

The ports and parameters can be configured individually on a system by:

1. Entering the commands in steps 1, and 2

2. Running the relevant addition steps and parameter commands
3. Verifying the result
4. Saving the configuration

For detailed descriptions of the parameters and features, see the sections STP parameters and STP features.

1. Enter global configuration mode.

```
device# configure terminal
```

2. Enable PVST+.

```
device(config)# protocol spanning-tree pvst
```

3. Enter interface configuration mode.

```
device(config-pvst)# interface ethernet 0/3
```

4. Enable spanning tree on the interface.

```
device(conf-if-eth-0/3)# no spanning-tree shutdown
```

5. Configure the interface link type.

```
device(conf-if-eth-0/3)# spanning-tree link-type point-to-point
```

6. Specify the port priority to influence the selection of root or designated ports.

```
device(conf-if-eth-0/3)# spanning-tree priority 64
```

   The range is from 0 through 240 in increments of 16. The default value is 128.

7. Configure the path cost for spanning tree calculations on the interface.

```
device(conf-if-eth-0/3)# spanning-tree cost 10000
```

   The lower the path cost means a greater chance that the interface becomes the root port. The range is 1 through 200000000. The default path cost is assigned as per the port speed.

8. Configure the path cost for spanning tree calculations a specific VLAN.

```
device(conf-if-eth-0/3)# spanning-tree vlan 10 cost 10000
```

   The lower the path cost means a greater chance that the interface becomes the root port. The range is 1 through 200000000. The default path cost is assigned as per the port speed.

9. Enable root guard on the interface.

```
device(conf-if-eth-0/3)# spanning-tree guard root
```

   Root guard protects the root bridge from malicious attacks and unintentional misconfigurations where a bridge device that is not intended to be the root bridge becomes the root bridge.

10. Enable BPDU guard on the interface.

```
device(conf-if-eth-0/3)# spanning-tree port-fast bpdu-guard
```

    BPDU guard removes a node that reflects BPDUs back in the network. It enforces the STP domain borders and keeps the active topology predictable by not allowing any network devices behind a BPDU guard-enabled port to participate in STP.

11. Enable BPDU filtering on the interface.

```
device(conf-if-eth-0/3)# spanning-tree port-fast bpdu-filter
```

    BPDU filtering allows you to avoid transmitting BPDUs on ports that are connected to an end system.

12. Return to privileged EXEC mode.

```
device(conf-if-eth-0/3)# exit
```

13. Verify the configuration.

```
device# show spanning-tree brief

 Spanning-tree Mode: PVST Protocol

      Root ID      Priority 4096
                   Address 768e.f805.5800
                   Hello Time 8, Max Age 25, Forward Delay 20

      Bridge ID    Priority 4096
                   Address 768e.f805.5800
                   Hello Time 8, Max Age 25, Forward Delay 20

 Interface    Role  Sts  Cost       Prio  Link-type     Edge
 ------------------------------------------------------------
 Eth 0/3      DES   FWD  200000     64    P2P            No
```

    Observe that the settings comply with the formula set out in the STP parameters section, as:

    (2 × (forward delay - 1)) ≥ maximum age ≥ (2 × (hello time + 1))

    or in this case :38 ≥ 25 ≥ 18.

14. Save the configuration.

```
device# copy running-config startup-config
```

PVST+ on an interface configuration example

```
device# configure terminal
device(config)# protocol spanning-tree pvst
device(conf-pvst)# interface ethernet 0/3
device(conf-if-eth-0/3)# no spanning-tree shutdown
device(conf-if-eth-0/3)# spanning-tree link-type point-to-point
device(conf-if-eth-0/3)# spanning-tree priority 64
device(conf-if-eth-0/3)# spanning-tree cost 10000
device(conf-if-eth-0/3)# spanning-tree vlan 10 cost 10000
device(conf-if-eth-0/3)# spanning-tree guard root
device(conf-if-eth-0/3)# spanning-tree port-fast bpdu-guard
device(conf-if-eth-0/3)# exit
device# show spanning-tree
device# copy running-config startup-config
```

Enabling and configuring PVST+ on a system

Follow the steps to configure PVST+ on a system.

The ports and parameters can be configured individually on a system by:

1. Entering the commands in steps 1, and 2
2. Running the relevant addition steps and parameter commands

3. Verifying the result
4. Saving the configuration

For detailed descriptions of the parameters and features, see the sections STP parameters and STP features.

1. Enter global configuration mode.

```
device# configure terminal
```

2. Enable PVST+.

```
device(config)# protocol spanning-tree pvst
```

3. Configure the bridge priority for the common instance.

```
device(config-pvst)# bridge-priority 4096
```

   Valid values range from 0 through 61440 in multiples of 4096. Assigning a lower priority value indicates that the bridge might become root.

4. Configure the forward delay parameter.

```
device(config-pvst)# forward-delay  15
```

5. Configure the hello time parameter.

```
device(config-pvst)# hello-time 2
```

6. Configure the maximum age parameter.

```
device(config-pvst)# max-age 20
```

7. Add VLANs.
   a) Configure VLAN 100 with a priority of 0.

```
device(config-pvst)# vlan 100 priority 0
```

      The bridge priority in configured in multiples of 4096.

   b) Configure VLAN 201 with a priority of 12288.

```
device(config-pvst)# vlan 201 priority 12288
```

   c) Configure VLAN 301 with a priority of 20480.

```
device(config-pvst)# vlan 301 priority 20480
```

8. Set the switching characteristics for interface 0/3.
   a) Enter interface configuration mode.

```
device(config)# interface ethernet 0/3
```

   b) Set the switching characteristics of the interface.

```
device(conf-if-eth-0/3)# switchport
```

   c) Set the interface mode to trunk.

```
device(conf-if-eth-0/3)# switchport mode trunk
```

   d) Add VLAN 100 as a member VLAN.

```
device(conf-if-eth-0/3)# switchport trunk allowed vlan add 100
```

   e) Add VLAN 201 as a member VLAN.

```
device(conf-if-eth-0/3)# switchport trunk allowed vlan add 201
```

   f) Add VLAN 301 as a member VLAN.

```
device(conf-if-eth-0/3)# switchport trunk allowed vlan add 301
```

   g) Enable spanning tree on the interface.

```
device(conf-if-eth-0/3)# no spanning-tree shutdown
```

   h) Return to privileged EXEC mode.

```
device(conf-if-eth-0/3)# exit
```

9. Set the switching characteristics for interface 0/4.
   a) Enter interface configuration mode.

```
device(config)# interface ethernet 0/4
```

   b) Set the switching characteristics of the interface.

```
device(conf-if-eth-0/4)# switchport
```

   c) Set the interface mode to trunk.

```
device(conf-if-eth-0/4)# switchport mode trunk
```

   d) Add VLAN 100 as a member VLAN.

```
device(conf-if-eth-0/4)# switchport trunk allowed vlan add 100
```

   e) Add VLAN 201 as a member VLAN.

```
device(conf-if-eth-0/4)# switchport trunk allowed vlan add 201
```

   f) Add VLAN 301 as a member VLAN.

```
device(conf-if-eth-0/4)# switchport trunk allowed vlan add 301
```

   g) Enable spanning tree on the interface.

```
device(conf-if-eth-0/4)# no spanning-tree shutdown
```

   h) Return to privileged EXEC mode.

```
device(conf-if-eth-0/4)# exit
```

10. To interoperate with switches other than VDX switches in PVST+ mode, you must configure the interface that is connected to that switch.
   a) Enter interface configuration mode for the port that interoperates with a VDX device.

```
device(config)# interface ethernet 0/12
```

   b) Specify the MAC address for the device.

```
device(conf-if-eth-0/12)# spanning-tree bpdu-mac 0100.0ccc.cccd
```

   c) Enable spanning tree on the interface.

```
device(conf-if-eth-0/12)# no spanning-tree shutdown
```

   d) Return to privileged EXEC mode.

```
device(conf-if-eth-0/12)# end
```

11. Verify the configuration.

```
device# show spanning-tree

VLAN 1

 Spanning-tree Mode: PVST Protocol

 Root Id: 0001.01e0.5200.0180 (self)
 Bridge Id: 0001.01e0.5200.0180

 Root Bridge Forward Delay: 15; Hello Time: 2; Max Age: 20; Max-hops: 20
 Configured Forward Delay: 15; Hello Time: 2; Max Age: 20; Max-hops: 20
 Number of topology change(s): 0

 Bpdu-guard errdisable timeout: disabled
 Bpdu-guard errdisable timeout interval: 300 sec

 Port Et 0/3 enabled
    Ifindex: 201351168; Id: 8001; Role: Disabled; State: Disabled
    Designated Path Cost: 0
    Configured Path Cost: 20000000
    Designated Port Id: 0; Port Priority: 128
    Designated Bridge: 0000.0000.0000.0000
    Number of forward-transitions: 0
    Version: Per-VLAN Spanning Tree Protocol - Received None - Sent STP
    Portfast: off
    Configured Root guard: off; Operational Root guard: off
    Bpdu-guard: off
    Link-type: point-to-point
    Received BPDUs: 0; Sent BPDUs: 0

 Port Et 0/4 enabled
    Ifindex: 201359360; Id: 8002; Role: Disabled; State: Disabled
    Designated Path Cost: 0
    Configured Path Cost: 20000000
    Designated Port Id: 0; Port Priority: 128
    Designated Bridge: 0000.0000.0000.0000
    Number of forward-transitions: 0
    Version: Per-VLAN Spanning Tree Protocol - Received None - Sent STP
    Portfast: off
    Configured Root guard: off; Operational Root guard: off
    Bpdu-guard: off
    Link-type: point-to-point
    Received BPDUs: 0; Sent BPDUs: 0

VLAN 100

 Spanning-tree Mode: PVST Protocol

 Root Id: 0064.01e0.5200.0180 (self)
 Bridge Id: 0064.01e0.5200.0180

 Root Bridge Forward Delay: 15; Hello Time: 2; Max Age: 20; Max-hops: 20
 Configured Forward Delay: 15; Hello Time: 2; Max Age: 20; Max-hops: 20
 Number of topology change(s): 0

 Bpdu-guard errdisable timeout: disabled
 Bpdu-guard errdisable timeout interval: 300 sec

 Port Et 0/3 enabled
    Ifindex: 201351168; Id: 8001; Role: Disabled; State: Disabled
    Designated Path Cost: 0
    Configured Path Cost: 20000000
    Designated Port Id: 0; Port Priority: 128
    Designated Bridge: 0000.0000.0000.0000
    Number of forward-transitions: 0
    Version: Per-VLAN Spanning Tree Protocol - Received None - Sent STP
    Portfast: off
    Configured Root guard: off; Operational Root guard: off
    Bpdu-guard: off
```

```
    Link-type: point-to-point
    Received BPDUs: 0; Sent BPDUs: 0

 Port Et 0/4 enabled
    Ifindex: 201359360; Id: 8002; Role: Disabled; State: Disabled
    Designated Path Cost: 0
    Configured Path Cost: 20000000
    Designated Port Id: 0; Port Priority: 128
    Designated Bridge: 0000.0000.0000.0000
    Number of forward-transitions: 0
    Version: Per-VLAN Spanning Tree Protocol - Received None - Sent STP
    Portfast: off
    Configured Root guard: off; Operational Root guard: off
    Bpdu-guard: off
    Link-type: point-to-point
    Received BPDUs: 0; Sent BPDUs: 0

VLAN 201

 Spanning-tree Mode: PVST Protocol

 Root Id: 30c9.01e0.5200.0180 (self)
 Bridge Id: 30c9.01e0.5200.0180

 Root Bridge Forward Delay: 15; Hello Time: 2; Max Age: 20; Max-hops: 20
 Configured Forward Delay: 15; Hello Time: 2; Max Age: 20; Max-hops: 20
 Number of topology change(s): 0

 Bpdu-guard errdisable timeout: disabled
 Bpdu-guard errdisable timeout interval: 300 sec

 Port Et 0/3 enabled
    Ifindex: 201351168; Id: 8001; Role: Disabled; State: Disabled
    Designated Path Cost: 0
    Configured Path Cost: 20000000
    Designated Port Id: 0; Port Priority: 128
    Designated Bridge: 0000.0000.0000.0000
    Number of forward-transitions: 0
    Version: Per-VLAN Spanning Tree Protocol - Received None - Sent STP
    Portfast: off
    Configured Root guard: off; Operational Root guard: off
    Bpdu-guard: off
    Link-type: point-to-point
    Received BPDUs: 0; Sent BPDUs: 0

 Port Et 0/4 enabled
    Ifindex: 201359360; Id: 8002; Role: Disabled; State: Disabled
    Designated Path Cost: 0
    Configured Path Cost: 20000000
    Designated Port Id: 0; Port Priority: 128
    Designated Bridge: 0000.0000.0000.0000
    Number of forward-transitions: 0
    Version: Per-VLAN Spanning Tree Protocol - Received None - Sent STP
    Portfast: off
    Configured Root guard: off; Operational Root guard: off
    Bpdu-guard: off
    Link-type: point-to-point
    Received BPDUs: 0; Sent BPDUs: 0

VLAN 301

 Spanning-tree Mode: PVST Protocol

 Root Id: 512d.01e0.5200.0180 (self)
 Bridge Id: 512d.01e0.5200.0180

 Root Bridge Forward Delay: 15; Hello Time: 2; Max Age: 20; Max-hops: 20
 Configured Forward Delay: 15; Hello Time: 2; Max Age: 20; Max-hops: 20
 Number of topology change(s): 0

 Bpdu-guard errdisable timeout: disabled
 Bpdu-guard errdisable timeout interval: 300 sec
```

```
Port Et 0/3 enabled
   Ifindex: 201351168; Id: 8001; Role: Disabled; State: Disabled
   Designated Path Cost: 0
   Configured Path Cost: 20000000
   Designated Port Id: 0; Port Priority: 128
   Designated Bridge: 0000.0000.0000.0000
   Number of forward-transitions: 0
   Version: Per-VLAN Spanning Tree Protocol - Received None - Sent STP
   Portfast: off
   Configured Root guard: off; Operational Root guard: off
   Bpdu-guard: off
   Link-type: point-to-point
   Received BPDUs: 0; Sent BPDUs: 0

Port Et 0/4 enabled
   Ifindex: 201359360; Id: 8002; Role: Disabled; State: Disabled
   Designated Path Cost: 0
   Configured Path Cost: 20000000
   Designated Port Id: 0; Port Priority: 128
   Designated Bridge: 0000.0000.0000.0000
   Number of forward-transitions: 0
   Version: Per-VLAN Spanning Tree Protocol - Received None - Sent STP
   Portfast: off
   Configured Root guard: off; Operational Root guard: off
   Bpdu-guard: off
   Link-type: point-to-point
   Received BPDUs: 0; Sent BPDUs: 0
```

Observe that the settings comply with the formula set out in the STP parameters section, as:

(2 × (forward delay - 1)) ≥ maximum age ≥ (2 × (hello time + 1))

or in this case: 28 ≥ 20 ≥ 6.

12. Save the configuration.

```
device# copy running-config startup-config
```

Enable PVST+ on a system configuration example

```
device# configure terminal
device(config)# protocol spanning-tree pvst
device(config-pvst)# bridge-priority 4096
device(config-pvst)# forward-delay  15
device(config-pvst)# hello-time 2
device(config-pvst)# max-age 20
device(config-pvst)# vlan 100 priority 0
device(config-pvst)# vlan 201 priority 12288
device(config-pvst)# vlan 301 priority 20480
device(config-pvst)# interface ethernet 0/3
device(conf-if-eth-0/3)# switchport
device(conf-if-eth-0/3)# switchport mode trunk
device(conf-if-eth-0/3)# switchport trunk allowed vlan add 100
device(conf-if-eth-0/3)# switchport trunk allowed vlan add 201
device(conf-if-eth-0/3)# switchport trunk allowed vlan add 301
device(conf-if-eth-0/3)# no spanning-tree shutdown
device(conf-if-eth-0/3)# exit
device(config)# interface ethernet 0/4
device(conf-if-eth-0/4)# switchport
device(conf-if-eth-0/4)# switchport mode trunk
device(conf-if-eth-0/4)# switchport trunk allowed vlan add 100
device(conf-if-eth-0/4)# switchport trunk allowed vlan add 201
device(conf-if-eth-0/4)# switchport trunk allowed vlan add 301
device(conf-if-eth-0/4)# no spanning-tree shutdown
device(conf-if-eth-0/4)# end
device# show spanning-tree
device# copy running-config startup-config
```

Enabling and configuring R-PVST+ globally

Use this procedure to enable the Rapid Per-VLAN Spanning Tree Protocol Plus (R-PVST+).

1. Enter global configuration mode.

```
device# configure terminal
```

2. Enable R-PVST+.

```
device(config)# protocol spanning-tree rpvst
```

3. Configure the bridge priority for the common instance.

```
device(config-rpvst)# bridge-priority 4096
```

   Valid priority values range from 0 through 61440 in multiples of 4096. Assigning a lower priority value indicates that the bridge might become root.

4. Configure the forward delay parameter.

```
device(config-rpvst)# forward-delay  20
```

5. Configure the hello time parameter.

```
device(config-rpvst)# hello-time 22
```

6. Configure the maximum age parameter.

```
device(config-rpvst)# max-age 8
```

7. Set the transmit hold count for the bridge.

```
device(config-rpvst)# transmit-holdcount 9
```

   This command configures the maximum number of BPDUs transmitted per second before pausing for 1 second. The range is 1 through 10. The default is 6.

8. Return to privileged exec mode.

```
device(config-rpvst)# end
```

9. Verify the configuration.

```
device# show spanning-tree brief
VLAN 1

 Spanning-tree Mode: Rapid PVST Protocol

      Root ID      Priority 4096
                   Address 01e0.5200.0180
                   Hello Time 2, Max Age 7, Forward Delay 11

      Bridge ID    Priority 32769
                   Address 01e0.5200.0180
                   Hello Time 8, Max Age 22, Forward Delay 20, Tx-HoldCount 9
                   Migrate Time 3 sec

 Interface     Role  Sts  Cost       Prio  Link-type        Edge
 -------------------------------------------------------------------
```

   Observe that the settings comply with the formula set out in the STP parameters section, as:

   (2 × (forward delay - 1)) ≥ maximum age ≥ (2 × (hello time + 1))

   or in this case: 20 ≥ 7 ≥ 6.

10. Save the configuration.

```
device# copy running-config startup-config
```

R-PVST+ configuration example

```
device# configure terminal
device(config)# protocol spanning-tree rpvst
device(config-rpvst)# bridge-priority 4096
device(config-rpvst)# forward-delay 20
device(config-rpvst)# hello-time 22
device(config-rpvst)# max-age 8
device(config-rpvst)# transmit-holdcount 9
device(config-rpvst)# end
device# show spanning-tree brief
device# copy running-config startup-config
```

For more information about configuring parameters, see the section STP parameter configuration.

Enabling and configuring R-PVST+ on an interface

Follow these steps to enable and configure R-PVST+ on an interface.

The ports and parameters can be configured individually on a system by:

1. Entering the commands in steps 1-3
2. Running the relevant addition steps and parameter commands
3. Verifying the result
4. Saving the configuration

For detailed descriptions of the parameters and features, see the sections STP parameters and STP features.

1. Enter global configuration mode.

```
device# configure terminal
```

2. Enable R-PVST+.

```
device(config)# protocol spanning-tree rpvst
```

3. Enter interface configuration mode.

```
device(config-rpvst)# interface ethernet 0/3
```

4. Enable the spanning tree on the interface.

```
device(conf-if-eth-0/3)# no spanning-tree shutdown
```

5. Configure the interface link type.

```
device(conf-if-eth-0/3)# spanning-tree link-type point-to-point
```

6. Specify the port priority to influence the selection of root or designated ports.

```
device(conf-if-eth-0/3)# spanning-tree priority 64
```

   The range of priority values is from 0 through 240 in multiples of 16. The default value is 128.

7. Configure the path cost for spanning tree calculations on the interface.

```
device(conf-if-eth-0/3)# spanning-tree cost 200000
```

   The lower the path cost means a greater chance that the interface becomes the root port. The range is 1 through 200000000. The default path cost is assigned as per the port speed.

8. Configure the path cost for spanning tree calculations a specific VLAN.

```
device(conf-if-eth-0/3)# spanning-tree vlan 10 cost 10000
```

   The lower the path cost means a greater chance that the interface becomes the root port. The range is 1 through 200000000. The default path cost is assigned as per the port speed.

9. Enable automatic edge detection on the interface.

```
device(conf-if-eth-0/3)# spanning-tree autoedge
```

   You use this command to automatically identify the edge port. A port becomes an edge port if it receives no BPDUs. By default, automatic edge detection is disabled.

10. Enable root guard on the interface.

```
device(conf-if-eth-0/3)# spanning-tree guard root
```

    Root guard protects the root bridge from malicious attacks and unintentional misconfigurations where a bridge device that is not intended to be the root bridge becomes the root bridge.

11. Enable the spanning tree on the edge port.

```
device(conf-if-eth-0/3)# spanning-tree edgeport
```

    If BPDUs are received on a port fast enabled interface, the interface loses the edge port status unless it receives a **shutdown** or **no shutdown** command.

12. Enable BPDU guard on the interface.

```
device(conf-if-eth-0/3)# spanning-tree edgeport bpdu-guard
```

    BPDU guard removes a node that reflects BPDUs back in the network. It enforces the STP domain borders and keeps the active topology predictable by not allowing any network devices behind a BPDU guard-enabled port to participate in STP.

13. Return to privileged EXEC mode.

```
device(conf-if-eth-0/3)# exit
```

14. Verify the configuration.

```
device# show spanning-tree brief

 Spanning-tree Mode: Rapid PVST Protocol

      Root ID      Priority 4096
                   Address 768e.f805.5800
                   Hello Time 8, Max Age 25, Forward Delay 20

      Bridge ID    Priority 4096
                   Address 768e.f805.5800
                   Hello Time 8, Max Age 25, Forward Delay 20

 Interface    Role  Sts  Cost       Prio  Link-type     Edge
 -------------------------------------------------------------
 Eth 0/3      DES   FWD  200000     128   P2P            No
```

Observe that the settings comply with the formula set out in the STP parameters section, as:

(2 × (forward delay - 1)) ≥ maximum age ≥ (2 × (hello time + 1))

or in this case: 38 ≥ 25 ≥ 18.

15. Save the configuration.

```
device# copy running-config startup-config
```

R-PVST+ on an interface configuration example

```
device# configure terminal
device(config)# protocol spanning-tree rpvst
device(config-rpvst)# interface ethernet 0/3
device(conf-if-eth-0/3)# no spanning-tree shutdown
device(conf-if-eth-0/3)# spanning-tree link-type point-to-point
device(conf-if-eth-0/3)# spanning-tree priority 64
device(conf-if-eth-0/3)# spanning-tree cost 200000
device(conf-if-eth-0/3)# spanning-tree vlan 10 cost 10000
device(conf-if-eth-0/3)# spanning-tree autoedge
device(conf-if-eth-0/3)# spanning-tree guard root
device(conf-if-eth-0/3)# spanning-tree edgeport
device(conf-if-eth-0/3)# spanning-tree edgeport bpdu-guard
device(conf-if-eth-0/3)# exit
device# show spanning-tree
device# copy running-config startup-config
```

Enabling and configuring R-PVST+ on a system

Follow the steps to configure R-PVST+ on a system.

The ports and parameters can be configured individually by:

1. Entering the commands in steps 1 and 2
2. Running the relevant addition steps and parameter commands
3. Verifying the result
4. Saving the configuration

For detailed descriptions of the parameters and features, see the sections STP parameters and STP features.

1. Enter global configuration mode.

```
device# configure terminal
```

2. Enable R-PVST+.

```
device(config)# protocol spanning-tree rpvst
```

   You can shut down R-PVST+ by entering the **shutdown** command when in `rpvst` configuration mode.

3. Configure the bridge priority for the common instance.

```
device(config-rpvst)# bridge-priority 4096
```

   Valid values range from 0 through 61440 in increments of 4096. Assigning a lower priority value indicates that the bridge might become root.

4. Configure the forward delay parameter.

```
device(config-rpvst)# forward-delay  20
```

5. Configure the hello time parameter.

```
device(config-rpvst)# hello-time 8
```

6. Configure the maximum age parameter.

```
device(config-rpvst)# max-age 22
```

7. Specify the transmit hold count.

```
device(config-rpvst)# transmit-holdcount 5
```

   This command configures the maximum number of BPDUs transmitted per second. The range of values is 1 through 10.

8. Configure VLANs.
   a) Configure VLAN 100 with a priority of 0.

```
device(config-rpvst)# vlan 100 priority 0
```

      Valid priority values range from 0 through 61440 in multiples of 4096.

   b) Configure VLAN 201 with a priority of 12288.

```
device(config-rpvst)# vlan 201 priority 12288
```

   c) Configure VLAN 301 with a priority of 20480.

```
device(config-rpvst)# vlan 301 priority 20480
```

9. Set the switching characteristics for interface 0/3.
   a) Enter interface configuration mode.

```
device(config-rpvst)# interface ethernet 0/3
```

   b) Set the switching characteristics of the interface.

```
device(conf-if-eth-0/3)# switchport
```

   c) Set the interface mode to trunk.

```
device(conf-if-eth-0/3)# switchport mode trunk
```

   d) Add VLAN 100 as a member VLAN.

```
device(conf-if-eth-0/3)# switchport trunk allowed vlan add 100
```

   e) Add VLAN 201 as a member VLAN.

```
device(conf-if-eth-0/3)# switchport trunk allowed vlan add 201
```

   f) Add VLAN 301 as a member VLAN.

```
device(conf-if-eth-0/3)# switchport trunk allowed vlan add 301
```

   g) Enable spanning tree on the interface.

```
device(conf-if-eth-0/3)# no spanning-tree shutdown
```

   h) Return to privileged EXEC mode.

```
device(conf-if-eth-0/3)# exit
```

10. Set the switching characteristics for interface 0/4.
    a) Enter interface configuration mode.

```
device(config-rpvst)# interface ethernet 0/4
```

    b) Set the switching characteristics of the interface.

```
device(conf-if-eth-0/4)# switchport
```

    c) Set the interface mode to trunk.

```
device(conf-if-eth-0/4)# switchport mode trunk
```

    d) Add VLAN 100 as a member VLAN.

```
device(conf-if-eth-0/4)# switchport trunk allowed vlan add 100
```

    e) Add VLAN 201 as a member VLAN.

```
device(conf-if-eth-0/4)# switchport trunk allowed vlan add 201
```

    f) Add VLAN 301 as a member VLAN.

```
device(conf-if-eth-0/4)# switchport trunk allowed vlan add 301
```

    g) Enable spanning tree on the interface.

```
device(conf-if-eth-0/4)# no spanning-tree shutdown
```

    h) Return to privileged EXEC mode.

```
device(conf-if-eth-0/4)# exit
```

11. To interoperate with switches other than VDX switches in R-PVST+ mode, you must configure the interface that is connected to that switch.
    a) Enter interface configuration mode for the port that interoperates with a VDX switch.

```
device(config)# interface ethernet 0/12
```

    b) Specify the MAC address for the device.

```
device(conf-if-eth-0/12)# spanning-tree bpdu-mac 0100.0ccc.cccd
```

    c) Enable spanning tree on the interface.

```
device(conf-if-eth-0/12)# no spanning-tree shutdown
```

    d) Return to privileged EXEC mode.

```
device(conf-if-eth-0/12)# end
```

12. Verify the configuration.

```
device# show spanning-tree

VLAN 1

 Spanning-tree Mode: Rapid PVST Protocol

 Root Id: 0001.01e0.5200.0180 (self)
 Bridge Id: 0001.01e0.5200.0180

 Root Bridge Forward Delay: 15; Hello Time: 2; Max Age: 20; Max-hops: 20
 Configured Forward Delay: 20; Hello Time: 8; Max Age: 22; Max-hops: 20
 Tx-HoldCount 5
 Number of topology change(s): 0

 Bpdu-guard errdisable timeout: disabled
 Bpdu-guard errdisable timeout interval: 300 sec

 Port Et 0/3 enabled
    Ifindex: 201351168; Id: 8001; Role: Disabled; State: Disabled
    Designated Path Cost: 0
    Configured Path Cost: 20000000
    Designated Port Id: 0; Port Priority: 128
    Designated Bridge: 0000.0000.0000.0000
    Number of forward-transitions: 0
    Version: Per-VLAN Spanning Tree Protocol - Received None - Sent STP
    Portfast: off
    Configured Root guard: off; Operational Root guard: off
    Bpdu-guard: off
    Link-type: point-to-point
    Received BPDUs: 0; Sent BPDUs: 0

 Port Et 0/4 enabled
    Ifindex: 201359360; Id: 8002; Role: Disabled; State: Disabled
    Designated Path Cost: 0
    Configured Path Cost: 20000000
    Designated Port Id: 0; Port Priority: 128
    Designated Bridge: 0000.0000.0000.0000
    Number of forward-transitions: 0
    Version: Per-VLAN Spanning Tree Protocol - Received None - Sent STP
    Portfast: off
    Configured Root guard: off; Operational Root guard: off
    Bpdu-guard: off
    Link-type: point-to-point
    Received BPDUs: 0; Sent BPDUs: 0

VLAN 100

 Spanning-tree Mode: Rapid PVST Protocol

 Root Id: 0064.01e0.5200.0180 (self)
 Bridge Id: 0064.01e0.5200.0180

 Root Bridge Forward Delay: 15; Hello Time: 2; Max Age: 20; Max-hops: 20
 Configured Forward Delay: 20; Hello Time: 8; Max Age: 22; Max-hops: 20
 Tx-HoldCount 5
 Number of topology change(s): 0

 Bpdu-guard errdisable timeout: disabled
 Bpdu-guard errdisable timeout interval: 300 sec

 Port Et 0/3 enabled
    Ifindex: 201351168; Id: 8001; Role: Disabled; State: Disabled
    Designated Path Cost: 0
    Configured Path Cost: 20000000
    Designated Port Id: 0; Port Priority: 128
    Designated Bridge: 0000.0000.0000.0000
    Number of forward-transitions: 0
    Version: Per-VLAN Spanning Tree Protocol - Received None - Sent STP
    Portfast: off
```

```
    Configured Root guard: off; Operational Root guard: off
    Bpdu-guard: off
    Link-type: point-to-point
    Received BPDUs: 0; Sent BPDUs: 0

 Port Et 0/4 enabled
    Ifindex: 201359360; Id: 8002; Role: Disabled; State: Disabled
    Designated Path Cost: 0
    Configured Path Cost: 20000000
    Designated Port Id: 0; Port Priority: 128
    Designated Bridge: 0000.0000.0000.0000
    Number of forward-transitions: 0
    Version: Per-VLAN Spanning Tree Protocol - Received None - Sent STP
    Portfast: off
    Configured Root guard: off; Operational Root guard: off
    Bpdu-guard: off
    Link-type: point-to-point
    Received BPDUs: 0; Sent BPDUs: 0

VLAN 201

 Spanning-tree Mode: Rapid PVST Protocol

 Root Id: 30c9.01e0.5200.0180 (self)
 Bridge Id: 30c9.01e0.5200.0180

 Root Bridge Forward Delay: 15; Hello Time: 2; Max Age: 20; Max-hops: 20
 Configured Forward Delay: 20; Hello Time: 8; Max Age: 22; Max-hops: 20
 Tx-HoldCount 5
 Number of topology change(s): 0

 Bpdu-guard errdisable timeout: disabled
 Bpdu-guard errdisable timeout interval: 300 sec

 Port Et 0/3 enabled
    Ifindex: 201351168; Id: 8001; Role: Disabled; State: Disabled
    Designated Path Cost: 0
    Configured Path Cost: 20000000
    Designated Port Id: 0; Port Priority: 128
    Designated Bridge: 0000.0000.0000.0000
    Number of forward-transitions: 0
    Version: Per-VLAN Spanning Tree Protocol - Received None - Sent STP
    Portfast: off
    Configured Root guard: off; Operational Root guard: off
    Bpdu-guard: off
    Link-type: point-to-point
    Received BPDUs: 0; Sent BPDUs: 0

 Port Et 0/4 enabled
    Ifindex: 201359360; Id: 8002; Role: Disabled; State: Disabled
    Designated Path Cost: 0
    Configured Path Cost: 20000000
    Designated Port Id: 0; Port Priority: 128
    Designated Bridge: 0000.0000.0000.0000
    Number of forward-transitions: 0
    Version: Per-VLAN Spanning Tree Protocol - Received None - Sent STP
    Portfast: off
    Configured Root guard: off; Operational Root guard: off
    Bpdu-guard: off
    Link-type: point-to-point
    Received BPDUs: 0; Sent BPDUs: 0

VLAN 301

 Spanning-tree Mode: Rapid PVST Protocol

 Root Id: 512d.01e0.5200.0180 (self)
 Bridge Id: 512d.01e0.5200.0180

 Root Bridge Forward Delay: 15; Hello Time: 2; Max Age: 20; Max-hops: 20
 Configured Forward Delay: 20; Hello Time: 8; Max Age: 22; Max-hops: 20
 Tx-HoldCount 5
```

```
Number of topology change(s): 0

Bpdu-guard errdisable timeout: disabled
Bpdu-guard errdisable timeout interval: 300 sec

Port Et 0/3 enabled
   Ifindex: 201351168; Id: 8001; Role: Disabled; State: Disabled
   Designated Path Cost: 0
   Configured Path Cost: 20000000
   Designated Port Id: 0; Port Priority: 128
   Designated Bridge: 0000.0000.0000.0000
   Number of forward-transitions: 0
   Version: Per-VLAN Spanning Tree Protocol - Received None - Sent STP
   Portfast: off
   Configured Root guard: off; Operational Root guard: off
   Bpdu-guard: off
   Link-type: point-to-point
   Received BPDUs: 0; Sent BPDUs: 0

Port Et 0/4 enabled
   Ifindex: 201359360; Id: 8002; Role: Disabled; State: Disabled
   Designated Path Cost: 0
   Configured Path Cost: 20000000
   Designated Port Id: 0; Port Priority: 128
   Designated Bridge: 0000.0000.0000.0000
   Number of forward-transitions: 0
   Version: Per-VLAN Spanning Tree Protocol - Received None - Sent STP
   Portfast: off
   Configured Root guard: off; Operational Root guard: off
   Bpdu-guard: off
   Link-type: point-to-point
   Received BPDUs: 0; Sent BPDUs: 0
```

Observe that the settings comply with the formula set out in the STP parameters section, as:

(2 × (forward delay - 1)) ≥ maximum age ≥ (2 × (hello time + 1))

or in this case: 28 ≥ 20 ≥ 6.

13. Save the configuration.

```
device# copy running-config startup-config
```

Enable R-PVST+ on a system configuration example

```
device# configure terminal
device(config)# protocol spanning-tree rpvst
device(config-rpvst)# bridge-priority 4096
device(config-rpvst)# forward-delay  20
device(config-rpvst)# hello-time 8
device(config-rpvst)# max-age 22
device(config-rpvst)# transmit-holdcount 5
device(config-rpvst)# vlan 100 priority 0
device(config-rpvst)# vlan 201 priority 12288
device(config-rpvst)# vlan 301 priority 20480
device(config-rpvst)# interface ethernet 0/3
device(conf-if-eth-0/3)# switchport
device(conf-if-eth-0/3)# switchport mode trunk
device(conf-if-eth-0/3)# switchport trunk allowed vlan add 100
device(conf-if-eth-0/3)# switchport trunk allowed vlan add 201
device(conf-if-eth-0/3)# switchport trunk allowed vlan add 301
device(conf-if-eth-0/3)# no spanning-tree shutdown
device(conf-if-eth-0/3)# exit
device(config)# interface ethernet 0/4
device(conf-if-eth-0/4)# switchport
device(conf-if-eth-0/4)# switchport mode trunk
device(conf-if-eth-0/4)# switchport trunk allowed vlan add 100
device(conf-if-eth-0/4)# switchport trunk allowed vlan add 201
device(conf-if-eth-0/4)# switchport trunk allowed vlan add 301
device(conf-if-eth-0/4)# no spanning-tree shutdown
device(conf-if-eth-0/4)# end
device# show spanning-tree
device# copy running-config startup-config
```

Clearing spanning tree counters

Follow these steps to clear spanning tree counters on all interfaces or on the specified interface.

1. Clear spanning tree counters on all interfaces.

   ```
   device# clear spanning-tree counter
   ```

2. Clear spanning tree counters on a specified Ethernet interface.

   ```
   device# clear spanning-tree counter interface ethernet 0/3
   ```

3. Clear spanning tree counters on a specified port channel interface.

   ```
   device# clear spanning-tree counter interface port-channel 12
   ```

   Port channel interface numbers range from 1 through 64.

Clearing spanning tree-detected protocols

Follow these steps to restart the protocol migration process.

These commands force a spanning tree renegotiation with neighboring devices on either all interfaces or on a specified interface.

1. Restart the spanning tree migration process on all interfaces.

   ```
   device# clear spanning-tree detected-protocols
   ```

2. Restart the spanning tree migration process on a specific Ethernet interface.

   ```
   device# clear spanning-tree detected-protocols interface ethernet 0/3
   ```

3. Restart the spanning tree migration process on a specific port channel interface.

```
device# clear spanning-tree detected-protocols port-channel 12
```

   Port channel interface numbers range from 1 through 64.

Shutting down PVST+ or R-PVST+

Follow these steps to o shut down PVST+, or R-PVST+ either globally, on a specific interface, or a specific VLAN.

1. Enter global configuration mode.

```
device# configure terminal
```

2. Shut down PVST+ or R-PVST+.
   - Shut down PVST+ or R-PVST+ globally and return to privileged EXEC mode.

```
device(config)# protocol spanning-tree pvst
device(config-pvst)# shutdown
device(config-pvst)# end
```

   - Shut down PVST+ or R-PVST+ on a specific interface and return to privileged EXEC mode.

```
device(config)# interface ethernet 1/2
device(conf-if-eth-1/2)# spanning-tree shutdown
device(conf-if-eth-1/2)# end
```

   - Shut down PVST+ or R-PVST+ on a specific VLAN. and return to privileged EXEC mode.

```
device(config)# vlan 10
device(config-vlan-10)# spanning-tree shutdown
device(config-vlan-10)# end
```

3. Verify the configuration.

```
device# show spanning-tree
device#
```

4. Save the running configuration to the startup configuration.

```
device# copy running-config startup-config
```

Shut down PVST+ or R-PVST+ configuration example

```
device# configure terminal
device(config)# vlan 10
device(config-vlan-10)# spanning-tree shutdown
device(config-vlan-10)# end
device# show spanning-tree
device# copy running-config startup-config
```

NOTE

Shutting down PVST+ on a VLAN is used in this example.

802.1s Multiple Spanning Tree Protocol

MSTP overview

IEEE 802.1s Multiple STP (MSTP) helps create multiple loop-free active topologies on a single physical topology.

MSTP uses RSTP to group VLANs into separate spanning-tree instance. Each instance has its own spanning-tree topology independent of other spanning tree instances, which allows multiple forwarding paths, permits load balancing, and facilitates the movement of data traffic. A failure in one instance does not affect other instances. By enabling the MSTP, you are able to more effectively utilize the physical resources present in the network and achieve better load balancing of VLAN traffic.

The MSTP evolved as a compromise between the two extremes of the RSTP and R-PVST+, it was standardized as IEEE 802.1s and later incorporated into the IEEE 802.1Q-2003 standard. The MSTP configures a meshed topology into a loop-free, tree-like topology. When the link on a bridge port goes up, an MSTP calculation occurs on that port. The result of the calculation is the transition of the port into either a forwarding or blocking state. The result depends on the position of the port in the network and the MSTP parameters. All the data frames are forwarded over the spanning tree topology calculated by the protocol.

NOTE

Multiple switches must be configured consistently with the same MSTP configuration to participate in multiple spanning tree instances. A group of interconnected switches that have the same MSTP configuration is called an MSTP region.
MSTP is backward compatible with the STP and the RSTP.

Common Spanning Tree (CST)

The single Spanning Tree instance used by the Extreme device, and other vendor devices to interoperate with MSTP bridges. This spanning tree instance stretches across the entire network domain (including PVST, PVST+ and MSTP regions). It is associated with VLAN 1 on the Extreme device.

Internal Spanning Tree (IST)

An MSTP bridge must handle at least these two instances: one IST and one or more MSTIs (Multiple Spanning Tree Instances). Within each MST region, the MSTP maintains multiple spanning-tree instances. Instance 0 is a special instance known as IST, which extends CST inside the MST region. IST always exists if the device runs MSTP. Besides IST, this implementation supports up to 31 MSTIs.

Common Internal Spanning Tree (CIST)

The single spanning tree calculated by STP (including PVST+) and RSTP (including R-PVST+) and the logical continuation of that connectivity through MSTP bridges and regions, calculated by MSTP to ensure that all LANs in the bridged LAN are simply and fully connected

Multiple Spanning Tree Instance (MSTI)

One of a number of spanning trees calculated by the MSTP within an MST Region, to provide a simply and fully connected active topology for frames classified as belonging to a VLAN that is mapped to the MSTI by the MST configuration table used by the MST bridges of that MST region.

The Extreme implementation supports up to 32 spanning tree instances in an MSTP enabled bridge that can support up to 32 different Layer 2 topologies. The spanning tree algorithm used by the MSTP is the RSTP, which provides quick convergence.

By default all configured VLANs including the default VLAN are assigned to and derive port states from CIST until explicitly assigned to MSTIs.

MST regions

MST regions are clusters of bridges that run multiple instances of the MSTP protocol. Multiple bridges detect that they are in the same region by exchanging their configuration (instance to VLAN mapping), name, and revision-level. Therefore, if you need to have two bridges in the same region, the two bridges must have identical configurations, names, and revision-levels. Also, one or more VLANs can be mapped to one MST instance (IST or MSTI) but a VLAN cannot be mapped to multiple MSTP instances

MSTP regions

MSTP introduces a hierarchical way of managing device domains using regions. Devices that share common MSTP configuration attributes belong to a region. The MSTP configuration determines the MSTP region where each device resides. The common MSTP configuration attributes are as follows:

- Alphanumeric configuration name (32 bytes)
- Configuration revision number (2 bytes)
- 4096-element table that maps each of the VLANs to an MSTP instance

Region boundaries are determined by the above attributes. An MSTI is an RSTP instance that operates inside an MSTP region and determines the active topology for the set of VLANs mapping to that instance. Every region has a CIST that forms a single spanning tree instance which includes all the devices in the region. The difference between the CIST instance and the MSTP instance is that the CIST instance operates across the MSTP region and forms a loop-free topology across regions, while the MSTP instance operates only within a region. The CIST instance can operate using the RSTP only if all the devices across the regions support the RSTP. However, if any of the devices operate using the STP, the CIST instance reverts to the STP.

Each region is viewed logically as a single STP or a single RSTP bridge to other regions.

NOTE
Extreme supports 32 MSTP instances and one MSTP region.

For more information about spanning trees, see the introductory sections in the Spanning Tree Protocol chapter.

MSTP guidelines and restrictions

Follow these restrictions and guidelines when configuring the MSTP:

- Create VLANs before mapping them to the MSTP instances.
- The Extreme implementation of the MSTP supports up to 32 MSTP instances and one MSTP region.
- The MSTP **force-version** option is not supported.
- You must create VLANs before mapping them to the MSTP instances.

- For two or more switches to be in the same the MSTP region, they must have the same VLAN-to-instance map, the same configuration revision number, and the same region name.
- MSTP is backward compatible with the STP and the RSTP.
- Only one MSTP region can be configured on a bridge.
- A maximum of 4090 VLANs can be configured across the 32 MSTP instances.
- MSTP and topology groups cannot be configured together.
- MSTP configured over MCT VLANs is not supported.

Default MSTP configuration

As well as the defaults listed in the section Understanding the default STP configuration on page 111 there are defaults taht apply only to MSTP configurations.

Parameter	Default setting
Cisco interoperability	Disabled
Device priority (when mapping a VLAN to an MSTP instance)	32768
Maximum hops	20 hops
Revision number	0

Interoperability with PVST+ and R-PVST+

Since Extreme or other vendor devices enabled with PVST+ and R-PVST+ send IEEE STP BPDUs in addition to the PVST and R-PVST BPDUs, the VLAN 1 spanning tree joins the Common Spanning Tree (CST) of the network and thus interoperates with MSTP. The IEEE compliant devices treat the BPDUs addressed to the Extreme proprietary multicast MAC address as an unknown multicast address and flood them over the active topology for the particular VLAN.

MSTP global level parameters

To configure a switch for MSTP, first you set the region name and the revision on each switch that is being configured for MSTP. You must then create an MSTP Instance and assign an ID. VLANs are then assigned to MSTP instances. These instances must be configured on all switches that interoperate with the same VLAN assignments.

Each of the steps used to configure and operate MSTP are described in the following:

NOTE
The MSTP Region and Revision global parameters are enabled for interface level parameters as described below.

- Set the MSTP region name — Each switch that is running MSTP is configured with a name. It applies to the switch which can have many different VLANs that can belong to many different MSTP regions. The default MSTP name is "NULL".
- Set the MSTP revision number — Each switch that is running MSTP is configured with a revision number. It applies to the switch, which can have many different VLANs that can belong to many different MSTP regions.
- Enabling and disabling Cisco interoperability — While in MSTP mode, use the **cisco-interoperability** command to enable or disable the ability to interoperate with certain legacy Cisco switches. If Cisco interoperability is required on any switch in the network, then all switches in the network must be compatible, and therefore enabled by means of this command. By default the Cisco interoperability is disabled.

- The parameters you would normally set when you configure STP are applicable to MSTP. Before you configure MSTP parameters see the sections explaining bridge parameters, the error disable timeout parameter and the port-channel path cost parameter in the STP section of this guide.

MSTP interface level parameters

Edge port and automatic edge detection

Configuring the edge port feature makes a port transition directly from initialization to the forwarding state, skipping the listening and learning states.

From an interface, you can configure a device to automatically identify the edge port. The port can become an edge port if no BPDU is received. By default, automatic edge detection is disabled.

Follow these guidelines to configure a port as an edge port:

- When edge port is enabled, the port still participates in a spanning tree.
- A port can become an edge port if no BPDU is received.
- When an edge port receives a BPDU, it becomes a normal spanning tree port and is no longer an edge port.
- Because ports that are directly connected to end stations cannot create bridging loops in the network, edge ports transition directly to the forwarding state and skip the listening and learning states.

  NOTE
  If BPDUs are received on a port fast enabled interface, the interface loses the edge port status unless it receives a **shutdown** or **no shutdown** command.

BPDU guard

In an STP environment, switches, end stations, and other Layer 2 devices use BPDUs to exchange information that STP will use to determine the best path for data flow.

In a valid configuration, edge port-configured interfaces do not receive BPDUs. If an edge port-configured interface receives a BPDU, an invalid configuration exists, such as the connection of an unauthorized device. The BPDU Guard provides a secure response to invalid configurations because the administrator must manually put the interface back in service.

BPDU guard removes a node that reflects BPDUs back in the network. It enforces the STP domain borders and keeps the active topology predictable by not allowing any network devices behind a BPDU guard-enabled port to participate in STP.

In some instances, it is unnecessary for a connected device, such as an end station, to initiate or participate in an STP topology change. In this case, you can enable the STP BPDU guard feature on the Extreme device port to which the end station is connected. The STP BPDU guard shuts down the port and puts it into an "error disabled" state. This disables the connected device's ability to initiate or participate in an STP topology. A log message is then generated for a BPDU guard violation, and a message is displayed to warn the network administrator of an invalid configuration.

The BPDU Guard provides a secure response to invalid configurations because the administrator must manually put the interface back in service with the **no shutdown** command if error disable recovery is not enabled by enabling the **errdisable-timeout** command. The interface can also be automatically configured to be enabled after a timeout. However, if the offending BPDUs are still being received, the port is disabled again.

Expected behavior in an interface context

When BPDU Guard is enabled on an interface, the device is expected to put the interface in Error Disabled state when BPDU is received on the port when edge-port and BPDU guard is enabled on the switch interface. When the port ceases to receive the BPDUs, it does not automatically switch to edge port mode, you must configure **error disable timeout** or **no shutdown** on the port to move the port back into edge port mode.

Restricted role

Configuring restricted role on a port causes the port not to be selected as root port for the CIST or any MSTI, even if it has the best spanning tree priority vector.

Restricted role ports are selected as an alternate port after the root port has been selected. It is configured by a network administrator to prevent bridges external to a core region of the network influencing the spanning tree active topology, possibly because those bridges are not under the full control of the administrator. It will protect the root bridge from malicious attack or even unintentional misconfigurations where a bridge device which is not intended to be root bridge, becomes root bridge causing severe bottlenecks in data path. These types of mistakes or attacks can be avoided by configuring 'restricted-role' feature on ports of the root bridge . This feature is similar to the "root-guard" feature which is proprietary implementation of Cisco for STP and RSTP but had been adapted in the 802.1Q standard as "restricted-role". The "restricted-role" feature if configured on an incorrect port can cause lack of spanning tree connectivity.

Expected behavior in an interface context

When this feature is enabled on an interface the device is expected to prevent a port configured with restricted-role feature from assuming the role of a Root port. Such a port is expected to assume the role of an Alternate port instead, once Root port is selected.

Restricted TCN

TCN BPDUs are used to inform other switches of port changes.

Configuring "restricted TCN" on a port causes the port not to propagate received topology change notifications and topology changes originated from a bridge external to the core network to other ports. It is configured by a network administrator to prevent bridges external to a core region of the network from causing MAC address flushing in that region, possibly because those bridges are not under the full control of the administrator for the attached LANs. If configured on an incorrect port it can cause temporary loss of connectivity after changes in a spanning trees active topology as a result of persistent incorrectly learned station location information.

Expected behavior in an interface context

When this feature is enabled on an interface, the device is expected to prevent propagation of topology change notifications from a port configured with the Restricted TCN feature to other ports. In this manner, the device prevents TCN propagation from causing MAC flushes in the entire core network.

Configuring MSTP

Enabling and configuring MSTP globally

Follow this procedure to configure the Multiple Spanning Tree Protocol.

1. Enter global configuration mode.

```
device# configure terminal
```

2. Enable MSTP.

```
device(config)# protocol spanning-tree mstp
```

This command creates a context for MSTP. MSTP is automatically enabled. All MSTP specific CLI commands can be issued only from this context. Entering **no protocol spanning-tree mstp** deletes the context and all the configurations defined within the context.

3. Specify the region name.

```
device(config-mstp)# region kerry
```

4. Specify the revision number.

```
device(config-mstp)# revision 1
```

5. Configure an optional description of the MSTP instance.

```
device(config-mstp)# description kerry switches
```

6. Specify the maximum hops for a BPDU to prevent the messages from looping indefinitely on the interface.

```
device(config-mstp)# max-hops 25
```

Setting this parameter prevents messages from looping indefinitely on the interface. The range is 1 through 40 hops while the default is 20 .

7. Map VLANs to MSTP instances and set the instance priority.
   a) Map VLANs 7 and 8 to instance 1.

```
device(config-mstp)# instance 1 vlan 7,8
```

   b) Map VLANs 21, 22, and 23 to instance 2.

```
device(config-mstp)# instance 2 vlan 21-23
```

   c) Set the priority of instance 1.

```
device(config-mstp)# instance 1 priority 4096
```

This command can be used only after the VLAN is created. VLAN instance mapping is removed from the configuration if the underlying VLANs are deleted.

8. Configure a bridge priority for the CIST bridge.

```
device(config-mstp)# bridge-priority 4096
```

The range is 0 through 61440 in increments of 4096. The default is 32768.

9. Set the error disable parameters.
   a) Enable the timer to bring the port out of error disable state.

```
device(config-mstp)# error-disable-timeout enable
```

   b) Specify the time in seconds it takes for an interface to time out..

```
device(config-mstp)# error-disable-timeout interval 60
```

   The range is fro 10 to 1000000 seconds with a default of 300 seconds.
10. Configure forward delay.
   a) Specify the bridge forward delay.

```
device(config-mstp)# forward-delay 15
```

   This command allows you to specify how long an interface remains in the listening and learning states before it begins forwarding. This command affects all MSTP instances. The range of values is from 4 to 30 seconds with a default of 15 seconds.
11. Configure hello time.

```
device(config-mstp)# hello-time 2
```

   The hello time determines how often the switch interface broadcasts hello BPDUs to other devices. The range is from 1 through 10 seconds with a default of 2 seconds.
12. Configure the maximum age.

```
device(config-mstp)# max-age 20
```

   You must set the **max-age** so that it is greater than the **hello-time**. The range is 6 through 40 seconds with a default of 20 seconds.
13. Specify the port-channel path cost.

```
device(config-mstp)# port-channel path-cost custom
```

   This command allows you to control the path cost of a port channel according to bandwidth.
14. Specify the transmit hold count.

```
device(config-mstp)# transmit-holdcount 5
```

   The transmit hold count is used to limit the maximum number of MSTP BPDUs that the bridge can transmit on a port before pausing for 1 second. The range is from 1 to 10 seconds with a default of 6 seconds.
15. Configure Cisco interoperability.

```
device(config-mstp)# cisco-interoperability enable
```

   This command enables the ability to interoperate with certain legacy Cisco switches. The default is Cisco interoperability is disabled.
16. Return to privileged exec mode.

```
device(config-mstp)# end
```

17. Verify the configuration. The following is an example configuration.

```
device# show spanning-tree mst-config

Spanning-tree Mode: Multiple Spanning Tree Protocol

 CIST Root Id: 8000.001b.ed9f.1700
 CIST Bridge Id: 8000.768e.f80a.6800
 CIST Reg Root Id: 8000.001b.ed9f.1700

 CIST Root Path Cost: 0; CIST Root Port: Eth 1/2
 CIST Root Forward Delay: 15; Hello Time: 2; Max Age: 20; Max-hops: 19
 Configured Forward Delay: 15; Hello Time: 2; Max Age: 20; Max-hops: 20;
 Tx-HoldCount: 6
 Number of topology change(s): 139; Last change occurred 00:03:36 ago on Eth 1/2

 Bpdu-guard errdisable timeout: disabled
 Bpdu-guard errdisable timeout interval: 300 sec
 Migrate Time: 3 sec

  Name           : kerry
  Revision Level : 1
  Digest         : 0x9357EBB7A8D74DD5FEF4F2BAB50531AA

  Instance        VLAN
  --------        ----
   0:             1
   1:             7,8
   2:             21-23
```

Observe that the settings comply with the formula set out in the STP parameters section, as:

(2 × (forward delay - 1)) ≥ maximum age ≥ (2 × (hello time + 1))

or in this case: 28 ≥ 20 ≥ 6.

18. Save the configuration.

```
device# copy running-config startup-config
```

MSTP configuration example

```
device# configure terminal
device(config)# protocol spanning-tree mstp
device(config-mstp)# region kerry
device(config-mstp)# revision 1
device(config-mstp)# description kerry switches
device(config-mstp)# max-hops 20
device(config-mstp)# instance 1 vlan 7,8
device(config-mstp)# instance 2 vlan 21-23
device(config-mstp)# instance 1 priority 4096
device(config-mstp)# bridge-priority 4096
device(config-mstp)# error-disable-timeout enable
device(config-mstp)# error-disable-timeout interval 60
device(config-mstp)# forward-delay 16
device(config-mstp)# hello-time 5
device(config-mstp)# max-age 16
device(config-mstp)# port-channel path-cost custon
device(config-mstp)# transmit-holdcount 5
device(config-mstp)# cisco-interoperability enable
device(config-mstp)# end
device# show spanning-tree mst
device# copy running-config startup-config
```

Enabling and configuring MSTP on an interface

Follow these steps to configure and enable MSTP on an Ethernet interface.

The parameters can be configured individually on an interface by:

1. Entering the commands in Steps 1 through Step 3 for the target interface
2. Running the relevant parameter command
3. Verifying the result
4. Saving the configuration

For detailed descriptions of the parameters and features, see the sections STP parameters and STP features.

1. Enter configuration mode.

   ```
   device# configure terminal
   ```

2. Enable MSTP.

   ```
   device(config)# protocol spanning-tree mstp
   ```

3. Enter interface configuration mode.

   ```
   device(config-mstp)# interface ethernet 0/5
   ```

4. Enable the interface.

   ```
   device(conf-if-eth-0/5)# no shutdown
   ```

5. Configure the restricted role feature for the port.

   ```
   device(conf-if-eth-0/5)# spanning-tree restricted-role
   ```

   This command keeps a port from becoming a root.

6. Restrict topology change notifications (TCN) BPDUs for an MSTP instance.

   ```
   device(conf-if-eth-0/5)# spanning-tree instance 5 restricted-tcn
   ```

   This prevents the port from propagating received TCNs and topology changes originating from a bridge, external to the core network, to other ports.

7. Enable auto detection of an MSTP edge port.

   ```
   device(conf-if-eth-0/5)# spanning-tree autoedge
   ```

   Enabling this feature allows the system to automatically identify the edge port. The port can become an edge port if no BPDU is received. By default, automatic edge detection is disabled.

8.
   ```
   device(conf-if-eth-0/5)# spanning-tree edgeport
   ```

   Enabling edge port allows the port to quickly transition to the forwarding state. By default, automatic edge detection is disabled.

9. Enable BPDU guard on the port

   ```
   device(conf-if-eth-0/5)# spanning-tree edgeport bpdu-guard
   ```

   BPDU guard removes a node that reflects BPDUs back in the network. It enforces the STP domain borders and keeps the active topology predictable by not allowing any network devices behind a BPDU guard-enabled port to participate in STP.

10. Set the path cost of a port.

```
device(conf-if-eth-0/5)# spanning-tree cost 200000
```

    The path cost range is from 1 to 200000000. Leaving the default adjusts path cost relative to changes in the bandwidth. A lower path cost indicates greater likelihood of becoming root port.

11. Configure the link type.

```
device(conf-if-eth-0/5)# spanning-tree link-type point-to-point
```

    The options are point-to-point or shared.

12. Enable port priority.

```
device(conf-if-eth-0/5)# spanning-tree priority 128
```

    The range is from 0 to 240 in increments of 16 with a default of 32. A lower priority indicates greater likelihood of becoming root port.

13. Return to privileged exec mode.

```
device(conf-if-eth-0/5)# end
```

14. Verify the configuration.

```
device# show spanning-tree interface ethernet 0/5

Spanning-tree Mode: Multiple Spanning Tree Protocol

 Root Id: 8000.001b.ed9f.1700
 Bridge Id: 8000.01e0.5200.011d

 Port Eth 0/5 enabled
    Ifindex: 411271175; Id: 8002; Role: Designated; State: Forwarding
    Designated External Path Cost: 0; Internal Path Cost: 20000000
    Configured Path Cost: 200000
    Designated Port Id: 8002; Port Priority: 128
    Designated Bridge: 8000.01e0.5200.011d
    Number of forward-transitions: 1
    Version: Multiple Spanning Tree Protocol - Received MSTP - Sent MSTP
    Edgeport: yes; AutoEdge: yes; AdminEdge: no; EdgeDelay: 3 sec
    Restricted-role is enabled
    Restricted-tcn is enabled
    Boundary: no
    Bpdu-guard: on
    Link-type: point-to-point
    Received BPDUs: 86; Sent BPDUs: 1654
```

15. Save the configuration.

```
device# copy running-config startup-config
```

Enable MSTP on an interface configuration example

```
device# configure terminal
device(config)# protocol spanning-tree mstp
device(config-mstp)# interface ethernet 0/5
device(conf-if-eth-0/5)# no shutdown
device(conf-if-eth-0/5)# spanning-tree restricted-role
device(conf-if-eth-0/5)# spanning-tree instance 5 restricted-tcn
device(conf-if-eth-0/5)# spanning-tree autoedge
device(conf-if-eth-0/5)# spanning-tree edgeport
device(conf-if-eth-0/5)# spanning-tree edgeport bpdu-guard
device(conf-if-eth-0/5)# spanning-tree cost 200000
device(conf-if-eth-0/5)# spanning-tree link-type point-to-point
device(conf-if-eth-0/5)# spanning-tree priority 128
device(conf-if-eth-0/5)# end
device# show spanning-tree interface ethernet 0/5
device# copy running-config startup-config
```

Enabling MSTP on a VLAN

1. Enter configuration mode.

   ```
   device# configure terminal
   ```

2. Enter the protocol command to enable MSTP configuration.

   ```
   device(config)# protocol spanning-tree mstp
   ```

3. Map a VLAN to an MSTP instance.

   ```
   device(config-mstp)# instance 5 vlan 300
   ```

4. Return to privileged EXEC mode.

   ```
   device(config-mstp)# end
   ```

5. Verify the configuration.

```
device# show spanning-tree mst

 Spanning-tree Mode: Multiple Spanning Tree Protocol

 CIST Root Id: 8000.609c.9f5d.4800 (self)
 CIST Bridge Id: 8000.609c.9f5d.4800
 CIST Reg Root Id: 8000.609c.9f5d.4800 (self)

 CIST Root Forward Delay: 15; Hello Time: 2; Max Age: 20; Max-hops: 20
 Configured Forward Delay: 15; Hello Time: 2; Max Age: 20; Max-hops: 20;
 Tx-HoldCount: 6
 Number of topology change(s): 0

 Bpdu-guard errdisable timeout: disabled
 Bpdu-guard errdisable timeout interval: 300 sec
 Migrate Time: 3 sec

  Name           : NULL
  Revision Level : 0
  Digest         : 0xD5FF4C3F6C18E2F27AF3A8300297ABAA

  Instance        VLAN
  --------        ----
   0:             1
   5:             100
```

   Observe that the settings comply with the formula set out in the STP parameters section, as:

   (2 × (forward delay - 1)) ≥ maximum age ≥ (2 × (hello time + 1))

   or in this case: 28 ≥ 20 ≥ 6.

6. Save the configuration.

```
device# copy running-config startup-config
```

Enable spanning tree on a VLAN configuration example

```
device# configure terminal
device(config)# protocol spanning-tree mstp
device(config-mstp)# instance 5 vlan 300
device(config-mstp)# end
device# show spanning-tree mst
device# copy running-config startup-config
```

Configuring a basic MSTP

Follow these steps to configure a basic MSTP.

1. Enter global configuration mode.

```
device# configure terminal
```

2. Enable MSTP.

```
device(config)# protocol spanning-tree mstp
```

3. Specify the region name.

```
device(config-mstp)# region connemara
```

4. Specify the revision number.

```
device(config-mstp)# revision 1
```

5. Map MSTP instances to VLANs.
   a) Map instance 1 to VLANs 2 and 3.

```
device(config-mstp)# instance 1 vlan 2,3
```

   b) Map instance 2 to VLANs 4, 5, and 6.

```
device(config-mstp)# instance 2 vlan 4-6
```

6. Set a priority for an instance.

```
device(conf-Mstp)# instance 1 priority 28672
```

   The priority ranges from 0 through 61440 and the value must be in multiples of 4096.

7. Specify the maximum hops for a BPDU.

```
device(conf-Mstp)# max-hops 25
```

   This prevents the messages from looping indefinitely on an interface

8. Return to privileged EXEC mode.

```
device(conf-Mstp)# end
```

9. Verify the configuration.

```
device# show spanning-tree mst

 Spanning-tree Mode: Multiple Spanning Tree Protocol

 CIST Root Id: 8000.609c.9f5d.4800 (self)
 CIST Bridge Id: 8000.609c.9f5d.4800
 CIST Reg Root Id: 8000.609c.9f5d.4800 (self)

 CIST Root Forward Delay: 15; Hello Time: 2; Max Age: 20; Max-hops: 20
 Configured Forward Delay: 15; Hello Time: 2; Max Age: 20; Max-hops: 25;
 Tx-HoldCount: 6
 Number of topology change(s): 0

 Bpdu-guard errdisable timeout: disabled
 Bpdu-guard errdisable timeout interval: 300 sec
 Migrate Time: 3 sec

  Name           : connemara
  Revision Level : 1
  Digest         : 0xD5FF4C3F6C18E2F27AF3A8300297ABAA

  Instance        VLAN
  --------        ----
   0:             1,7,8,9
   1:             2,3
   2:             4-6
```

NOTE

Observe that the settings comply with the formula set out in the STP parameters section, as:
(2 × (forward delay - 1)) ≥ maximum age ≥ (2 × (hello time + 1))
or in this case: 28 ≥ 20 ≥ 6.

```
device# show running-config | begin spanning-tree
protocol spanning-tree mstp
 instance 1 vlan 2,3
 instance 1 priority 28672
 instance 2 vlan 4-6
 region connemars
 revision 1
 max-hops 25
 !
...
```

10. Save the configuration

```
device# copy running-config startup-config
```

Basic MSTP configuration example

```
device# configure terminal
device(config)# protocol spanning-tree mstp
device(config-mstp)# region connemara
device(config-mstp)# revision 1
device(config-mstp)# instance 1 vlan 2,3
device(config-mstp)# instance 2 vlan 4-6
device(conf-Mstp)# instance 1 priority 28582
device(conf-Mstp)# max-hops 25
device(conf-Mstp)# end
device# show spanning-tree mst
device# copy running-config startup-config
```

Clearing spanning tree counters

Follow these steps to clear spanning tree counters on all interfaces or on the specified interface.

1. Clear spanning tree counters on all interfaces.

```
device# clear spanning-tree counter
```

2. Clear spanning tree counters on a specified Ethernet interface.

```
device# clear spanning-tree counter interface ethernet 0/3
```

3. Clear spanning tree counters on a specified port channel interface.

```
device# clear spanning-tree counter interface port-channel 12
```

   Port channel interface numbers range from 1 through 64.

Clearing spanning tree-detected protocols

Follow these steps to restart the protocol migration process.

These commands force a spanning tree renegotiation with neighboring devices on either all interfaces or on a specified interface.

1. Restart the spanning tree migration process on all interfaces.

```
device# clear spanning-tree detected-protocols
```

2. Restart the spanning tree migration process on a specific Ethernet interface.

```
device# clear spanning-tree detected-protocols interface ethernet 0/3
```

3. Restart the spanning tree migration process on a specific port channel interface.

```
device# clear spanning-tree detected-protocols port-channel 12
```

   Port channel interface numbers range from 1 through 64.

Shutting down MSTP

Follow these steps to o shut down MSTP either globally, on a specific interface, or a specific VLAN.

1. Enter global configuration mode.

```
device# configure terminal
```

2. Shut down MSTP.
   - Shut down MSTP globally and return to privileged EXEC mode.

     ```
     device(config)# protocol spanning-tree mstp
     device(config-mstp)# shutdown
     device(config-mstp)# end
     ```

   - Shut down MSTP on a specific interface and return to privileged EXEC mode.

     ```
     device(config)# interface ethernet 1/2
     device(conf-if-eth-1/2)# spanning-tree shutdown
     device(conf-if-eth-1/2)# end
     ```

   - Shut down MSTP on a specific VLAN and return to privileged EXEC mode.

     ```
     device(config)# vlan 10
     device(config-vlan-10)# spanning-tree shutdown
     device(config-vlan-10)# end
     ```

3. Verify the configuration.

   ```
   device# show spanning-tree
   device#
   ```

4. Save the running configuration to the startup configuration.

   ```
   device# copy running-config startup-config
   ```

Shut down MSTP configuration example

```
device# configure terminal
device(config)# vlan 10
device(config-vlan-10)# spanning-tree shutdown
device(config-stp)# end
device# show spanning-tree
device# copy running-config startup-config
```

NOTE

Shutting down MSTP on a VLAN is used in this example.

Topology Groups

Topology Groups

A topology group is a named set of VLANs and bridge-domains that share a Layer 2 control protocol. Topology groups simplify configuration and enhance scalability of Layer 2 protocols by allowing you to run a single instance of a Layer 2 protocol on multiple VLANs and bridge-domains. One instance of the Layer 2 protocol controls all the VLANs and bridge-domains.

You can use topology groups with the following Layer 2 protocols:

- Per VLAN Spanning Tree (PVST+)
- Rapid per VLAN Spanning tree (R-PVST+)

Master VLAN, member VLANs, and bridge-domains

Each topology group contains a master VLAN and can contain one or more member VLANs and bridge-domains. A definition for each of these VLAN types follows:

- Master VLAN—The master VLAN contains the configuration information for the Layer 2 protocol. For example, if you plan to use the topology group for Rapid per VLAN Spanning tree (R-PVST), the topology group's master VLAN contains the R-PVST configuration information.
- Member VLANs—The member VLANs are additional VLANs that share ports with the master VLAN. The Layer 2 protocol settings for the ports in the master VLAN apply to the same ports in the member VLANs. A change to the master VLAN's Layer 2 protocol configuration or Layer 2 topology affects all the member VLANs. Member VLANs do not independently run a Layer 2 protocol.
- Member bridge domains—The member bridge domains are similar to VLANs that share ports with the master VLAN. The Layer 2 protocol settings for the ports in the master VLAN apply to the same ports in the bridge domains. A change to the master VLAN's Layer 2 protocol configuration or Layer 2 topology affects all the bridge domains. Bridge domains do not independently run a Layer 2 protocol. In a bridge domain, a single port can have multiple logical interfaces. In this scenario, all the logical interfaces on that port (and bridge domain) will follow the state of master VLAN port.

When a Layer 2 topology change occurs, resulting in a change of port state in the master VLAN, the same port state is applied to all the member VLANs and bridge-domains belonging to the topology group on that port. For example, if you configure a topology group whose master VLAN contains ports 1/1 and 1/2, a Layer 2 state change on port 1/1 applies to port 1/1 in all the member VLANs and bridge-domains that contain that port. However, the state change does not affect port 1/1 in VLANs that are not members of the topology group.

Control ports and free ports

A port in a topology group can be a control port or a free port:

- A **control port** is a port in the master VLAN and, therefore, is controlled by the Layer 2 protocol configured in the master VLAN. The same port in all the member VLANs and bridge-domains is controlled by the master VLAN's Layer 2 protocol. Each member VLAN and bridge-domain must contain all of the control ports. All other ports in the member VLAN and bridge-domain are "free ports."
- **Free ports** are not controlled by the master VLAN's Layer 2 protocol. The master VLAN can contain free ports. (In this case, the Layer 2 protocol is disabled on those ports.) In addition, any ports in the member VLANs and bridge-domains that are not also in the master VLAN are free ports.

  NOTE
  Because free ports are not controlled by the master port's Layer 2 protocol, they are always in the forwarding state.

Configuration considerations

The configuration considerations are as follows:

- You can configure up to 128 topology groups. A VLAN or bridge-domain cannot controlled by more than one topology group. You can configure up to 4K VLANs or bridge-domain as members of topology group.
- The topology group must contain a master VLAN. The group can also contain individual member VLANs and or member bridge-domains. You must configure the member VLANs or member bridge-domains before adding them to the topology group. Bridge-domains cannot be configured as a master VLAN.
- You cannot delete a master VLAN from the topology group when the member VLANs or bridge-domains are in the topology group.
- The control port membership must match the master VLAN when adding a member VLAN or member bridge-domain.
- If a VLAN enabled with the PVST+ or R-PVST+ protocol is added as a member VLAN of a topology group, the protocol is disabled. The member VLAN is added to the topology group. If the VLAN is removed from the topology group, the protocol is disabled, and you must enable the protocol if required.
- Enabling STP on an interface is only allowed if both master VLAN and member VLAN or bridge-domains are configured on the interface across all topology groups.
- You cannot remove the master VLAN or member VLAN or bridge-domains from an STP enabled interface.
- Topology group configuration is allowed only with PVST+ and R-PVST+ spanning tree configurations.

Configuring a topology group

Follow this procedure to configure a topology group. Extreme SLX devices support creating 128 topology groups in a system.

1. Enter the **configure terminal** command to access global configuration mode.

```
device# configure terminal
device(config)#
```

2. Enter the **topology-group** command to create a topology group at the global configuration level.

```
device(config)# topology-group 1
device(conf-topo-group-1)#
```

NOTE

The **no topology-group** command deletes an existing topology group.

Configuring a master VLAN

Follow this procedure to configure a master VLAN in a topology group.

Before configuring a master VLAN, you should have configured a topology group.

1. Enter the **configure terminal** command to access global configuration mode.

```
device# configure terminal
device(config)#
```

2. Enter the **topology-group** command to create a topology group at the global configuration level.

```
device(config)# topology-group 1
device(conf-topo-group-1)#
```

3. Enter the **master-vlan** command to configure a master VLAN in the topology group.

```
device(conf-topo-group-1)# master-vlan 100
```

NOTE

The **no master-vlan** command removes an existing master VLAN from the topology group.

Adding member VLANs

Follow this procedure to add member VLANs to a topology group. Member VLANs follow the master VLAN protocol states and also no L2 protocol will be running on the member VLANs.

Before adding a member VLAN, you should have created a topology group and configured the master VLAN for that group. The VLAN should not be part of any other topology group. All control ports of master VLAN must also be configured for the member VLAN.

1. Enter the **configure terminal** command to access global configuration mode.

```
device# configure terminal
device(config)#
```

2. Enter the **topology-group** command to create a topology group at the global configuration level.

```
device(config)# topology-group 1
device(conf-topo-group-1)#
```

3. Enter the **master-vlan** command to configure a master VLAN in the topology group.

```
device(conf-topo-group-1)# master-vlan 100
```

4. Enter the **member-vlan** command to add member VLANs to the topology group.

```
device(conf-topo-group-1)# member-vlan add 200-201
```

NOTE
The **member-vlan remove** command removes an existing member VLAN from the topology group.

```
device(conf-topo-group-1)# member-vlan remove 200
```

Adding member bridge-domains

Follow this procedure to add member bridge domains to a topology group. Member bridge domains follow the master VLAN protocol states and also no L2 protocol will be running on the bridge domains.

Before adding a bridge domain, you should have created a topology group and configured the master VLAN for that group. The bridge-domain should not be part of any other topology group. All control ports of master VLAN must also be configured for the member bridge-domain.

1. Enter the **configure terminal** command to access global configuration mode.

```
device# configure terminal
device(config)#
```

2. Enter the **topology-group** command to create a topology group at the global configuration level.

```
device(config)# topology-group 1
device(conf-topo-group-1)#
```

3. Enter the **master-vlan** command to configure a master VLAN in the topology group.

```
device(conf-topo-group-1)# master-vlan 100
```

4. Enter the **member-bridge-domain** command to add member bridge-domains to the topology group.

```
device(conf-topo-group-1)# member-bridge-domain add 300
```

NOTE
The **member-bridge-domain remove** command removes an existing member bridge-domain from the topology group.

```
device(conf-topo-group-1)# member-bridge-domain remove 1
```

The example adds 300 as member bridge-domain to the topology group.

```
device# configure terminal
device(config)# topology-group 1
device(conf-topo-group-1)# master-vlan 100
device(conf-topo-group-1)# member-bridge-domain add 300
```

Replacing a master VLAN

For replacing the existing master VLAN of a topology group, use the **master-vlan** command with the new master VLAN.

To avoid temporary loops when the master VLAN is replaced by another VLAN, the following recommendation is made:

- Control ports for both the old and the new master VLAN must match.
- The new master VLAN and the old master VLAN must have same ports in the blocking state to avoid the possibility of temporary loops.

If the recommendation is not followed, and a new master VLAN is configured with a different convergence, the configuration is still accepted.

NOTE
The master VLAN replacement is accepted if both the old and the new master VLANs are spanning-tree disabled.

Displaying topology group information

Follow the procedure to display topology group information for a specified group.

Before displaying the topology group information, you should have configured a topology group and defined the master VLAN.

Enter the **show topology-group** command to display the group information.

```
device# show topology-group 1
Topology Group 1
==============================
Master VLAN : 100
L2 Protocol: R-PVST
Member VLANs : 200 300
Member Bridge-domains: 10
Control Ports : eth 2/1, eth 2/2, po10
Free Ports : VLAN: 200 –eth 2/3, po11
Bridge-domain: 10 –eth 2/3.20, po11.10
```

The example displays information about topology group 1.

The **show running-config** command displays topology group configurations.

```
device# show running-config
topology-group 1
   master-vlan 100
   member-vlan add 200 300
   member-bridge-domain add 10
```

Loop Detection

LD protocol overview

The loop detection (LD) protocol is an Extreme proprietary protocol used to detect and break Layer 2 loops caused by misconfigurations, thereby preventing packet storms.

Layer 2 networks rely on learning and flooding to build their forwarding databases. Because of the flooding nature of these networks, any loops can be disastrous as they cause broadcast storms.

> **ATTENTION**
> The LD feature should be used only as a tool to detect loops in the network. It should not be used to replace other Layer 2 protocols such as STP.

This feature provides support for the following:

- Strict and loose modes
- Multi-Chassis Trunk (MCT)
- Breakout ports

LD protocol data units (PDUs) are initiated and received on the native device. Loop detection and action on the port state is also done on the same native device. Intermediate devices in the network must be capable of flooding unknown Layer 2 unicast PDUs on the VLAN through which they are received.

Strict mode

In what is referred to as *strict mode*, LD is configured on an interface. If the LD PDU is sent on an interface and received on the same interface, that port is shut down by LD. Strict mode overcomes specific hardware issues that cause packets to be echoed back to the input port. The following figure illustrates strict mode.

FIGURE 19 Strict mode

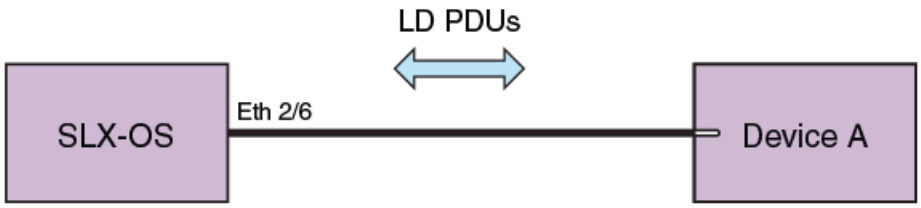

If the user provides a VLAN, then the PDUs are tagged accordingly. Otherwise PDUs are sent untagged. With a LAG, PDUs are sent out on the port-channel interface. If Device A has a loop (for example, a LAG is not configured), then the PDU is flooded back to SLX-OS, which detects the loop. In case of a loop, the port-channel interface is shut down. The following figure illustrates LD on a LAG.

FIGURE 20 LD on a LAG

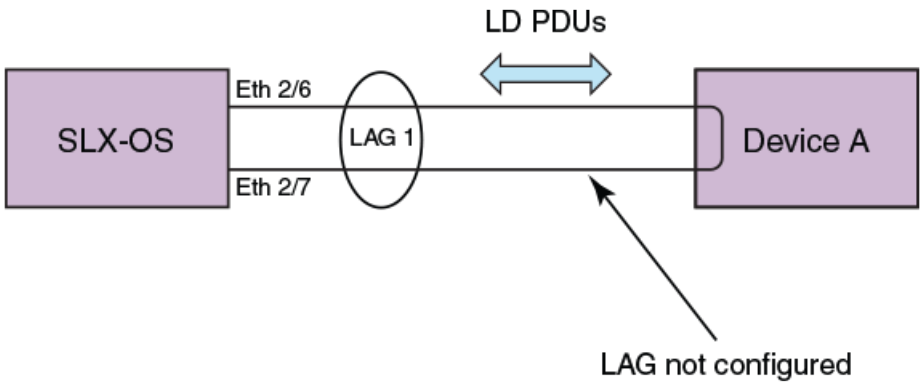

LD supports 256 instances of strict mode.

Loose mode

In what is referred to as *loose mode*, LD is configured on a VLAN. If a VLAN in the device receives an LD PDU that originated from the same device on that VLAN, this is considered to be a loop and the receiving port is shut down. In loose mode, LD works at the VLAN level and takes action at the link level. The following figure illustrates loose mode, with LD on a VLAN.

FIGURE 21 Loose mode: LD on a VLAN

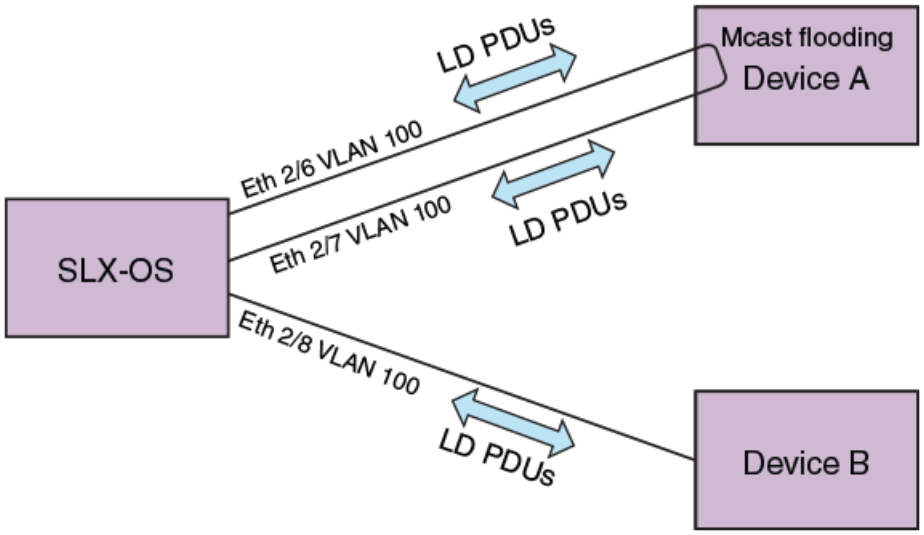

SLX-OS generates the LD PDUs on the VLAN. if Device A has a loop, PDUs are flooded back to SLX-OS, which detects the loop. SLX-OS then shuts down the receiving port on the VLAN.

LD supports 256 instances of loose mode, which means that it can be enabled on 256 VLANs.

LD PDU format

The following figure illustrates the format of the LD PDU in bytes.

FIGURE 22 LD PDU format in bytes

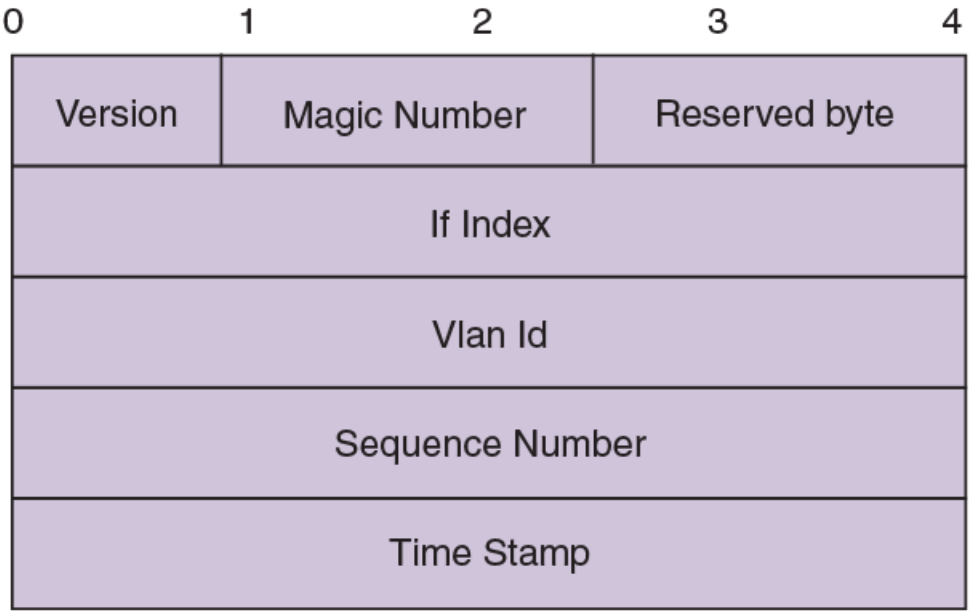

Parameter	Definition
Version	LD protocol version (1 by default)
Magic Number	0x13EF; used to differentiate between LD multicast PDUs and other multicast PDUs
Reserved byte	For future use
If Index	Index of the source port; populated only in strict mode
Vlan Id	VLAN ID
Sequence Number	Reserved for future enhancements
Time Stamp	Reserved for future enhancements

LD PDU transmission

Each LD-enabled interface or VLAN on a device continually transmits Layer 2 LD PDUs at a 1-second default hello-timer interval, with the destination MAC address as the multicast address. The multicast MAC address is derived from the system MAC address of the device with the multicast bit (8) and the local bit (7) set.

For example, if the MAC address is 00E0.5200.1800, then the multicast MAC address is 03E0.5200.1800. In the case of a LAG port-channel, LD PDUs are sent out one of the ports of the LAG as chosen by hardware.

LD PDU reception

When the LD PDU is received and is generated by the same device, the PDU is processed. If the PDU is generated by another device, then the PDU is flooded.

If a port is already blocked by any other Layer 2 protocol such as STP, then the LD PDUs are neither sent for LD processing nor flooded on that port.

LD parameters

This section discusses the various global protocol-level, interface level, and VLAN-level parameters that are used to control and process LD PDUs.

Protocol level

hello-interval

hello-interval is the rate at which the LD PDUs are transmitted by an LD-enabled interface or VLAN, which is 1000 milliseconds by default. Lowering the hello-interval below the default increases the PDU transmission rate, providing faster loop detection and also removing transient loops that last less than one second. On the other hand, increasing the interval above the default (for example, to 100 milliseconds) can increase the steady-state CPU load.

shutdown-time

shutdown-time is the duration after which an interface that is shut down by LD is automatically reenabled. The range is from 0 through 1440 minutes. The default is 0 minutes, which means that the interface is not automatically reenabled.

> ATTENTION
> Changing this value can cause repeated interface flapping when a loop is persistent in the network.

raslog-duration

raslog-duration is the interval between RASLog messages when a port is shut down by LD to prevent flooding of these messages. The range is from 10 through 1440 minutes. The default is 10.

Interface level

In strict mode, the parameters in this section are configurable at the interface level, and the configuration is specific to an interface. The following figure illustrates strict mode configuration.

FIGURE 23 Strict mode configuration

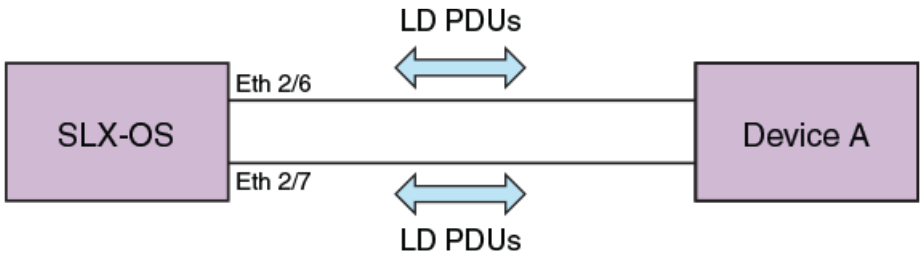

shutdown-disable

By default, the device shuts down the interface if a loop is detected. Configuring **shutdown-disable** means that the interface shutdown is disabled and LD never brings down such interface. If a loop is already detected by LD and the port is in shutdown state, then configuring **shutdown-disable** is not effective until the port is back up.

vlan-association

Although user can enable LD on an interface without specifying a VLAN, the **vlan-association** keyword is used to specify a VLAN associated with the interface.

VLAN level

In loose mode, the user can configure LD under a VLAN. In this case, LD PDUs are flooded on the VLAN. The following figure illustrates loose mode configuration.

FIGURE 24 Loose mode configuration

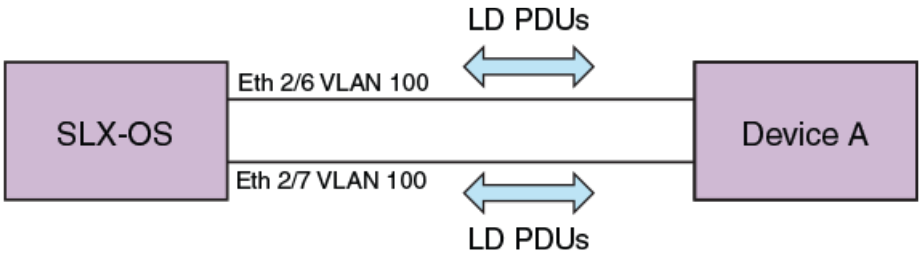

LD PDU processing

As long as LD PDUs are not received, there is no loop. If an LD PDU is received, then there is a loop that is present in the network.

If the if-index field in the received LD PDU is valid, then it is considered to be operating in strict mode. If the port on which the LD PDU was received is same as one encoded in the PDU (with a match for VLAN ID if a VLAN is associated), the port is shut down. For an MCT, if a strict mode LD PDU is received on an ICL interface, and the PDU is originated by another interface, then the ICL interface is not shut down. Instead, the sender interface is shut down. In addition, for strict mode the required interfaces should be configured with LD, or else the PDUs will not get processed

If the if-index field in the received LD PDU is invalid, then it is considered to be operating in loose mode. Based on VLAN ID information present in the received LD PDU, the receiving interface is shut down. If the receiving interface is an MCT ICL interface, the LD PDU is dropped.

In the case of a LAG (port-channel) interface, if the sent LD PDU is received on the port-channel, then the port-channel interface is shut down.

If the **shutdown-disable** option is configured for the particular interface, then the port drops the received PDU without processing it.

The re-enablement of the LD shut down port depends on the **shutdown-time** configuration. For manual recovery, either flap the port or clear the loop through the command line.

Configuration considerations

On an external switch that is unaware of LD or where LD is not configured, there may be some ACL rules applied to interfaces to permit traffic from known MAC addresses, and at the last of these rules there is an ACL deny-any rule to block all unknown MAC addresses. If this interface is part of a loop, LD enabled on SLX-OS will not be able to detect and break the loop.

LD use cases

In an MCT configuration, LD runs independently on both nodes. With loose mode the user must enable loop detection for the same VLANs on both nodes in the MCT cluster. MCT strict mode and loose mode use cases are detailed below.

MCT strict mode

The following figure illustrates a use case for MCT strict mode, followed by a sequence of events.

Strict mode LD is enabled on the MCT 1 cluster client edge port (CCEP) interface that connects to the Client.

FIGURE 25 MCT strict mode

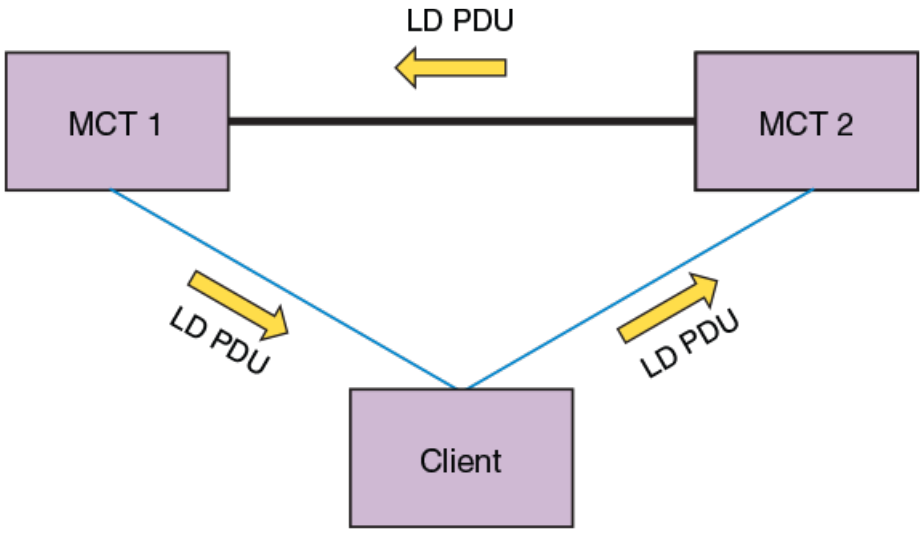

1. MCT 1 generates LD PDUs.
2. If the Client has the LAG interface configured to support LD, the Client drops the PDUs and there is no loop.
3. If there is a misconfiguration, the Client floods the PDUs, reaching MCT 2.
4. MCT 1 then identifies the interface information encoded in the PDUs, shutting down the interface on which the packets were generated.

MCT loose mode

The following figures illustrates two use cases for MCT loose mode, followed by a sequence of events.

FIGURE 26 MCT loose mode: Use case 1

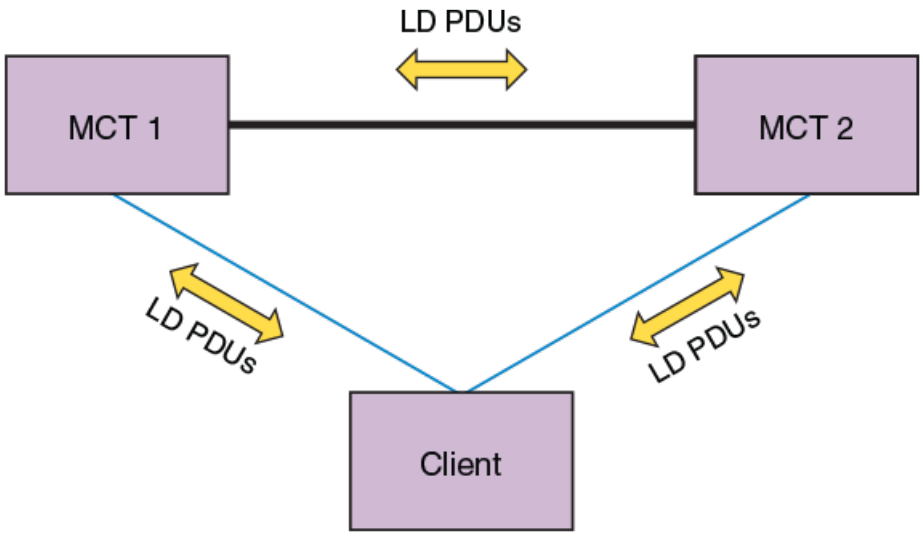

Use case 1: LD enabled on VLAN *x* on MCT 1

1. MCT 1 sends LD PDUs on VLAN x on all the interfaces that are part of the CCEP, client edge port (CEP), and ICL interface.
2. If the Client has LD configured on the LAG interface, then it drops the PDUs and no loop exists. If there is a misconfiguration, the Client floods the PDUs and they reach MCT 2.
3. MCT 2 floods the PDUs back to MCT 1, where the loop is detected. With loose mode no information about the interface that transmitted the PDU is encoded in the PDU, so normally the receiving interface is shut down. Because in this case the PDU is received on the ICL interface, that interface is not shut down.
4. MCT 1 receives the loop detection PDUs on the CCEP interface as well, as the packets were flooded in the VLAN in the following sequence: MCT 1 > MCT 2 > Client > MCT 1. In this case the receiving CCEP is shut down to break the loop. For MCT 2 to forward the PDUs in this case it must be the designated forwarder (DF) for that VLAN.

FIGURE 27 MCT loose mode: Use case 2

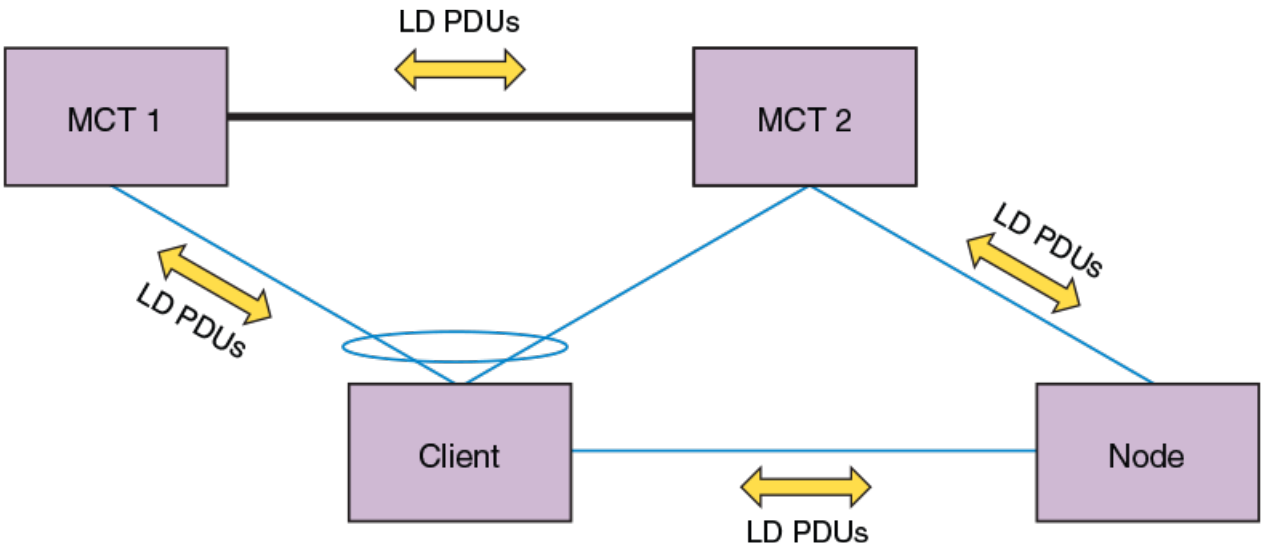

Use case 2: LD enabled on VLAN *x* on MCT 1 and MCT 2

1. Both MCT 1 and MCT 2 will flood the PDUs in VLAN *x* on all the interfaces that are part of the CCEP, CEP, and ICL interface.
2. Assuming PDUs from MCT 1 take the path MCT 1 > MCT 2 > Node > Client > MCT 1, then the receiving CCEP interface is shut down. For MCT 2 to forward the PDUs in this case, it must be the DF for that VLAN.
3. Assuming PDUs from MCT 2 take the path MCT 2 > MCT 1 > Client > Node > MCT 2, then the receiving CEP interface is shut down
4. If PDUs from MCT 2 take the path MCT 2 > Node > Client > MCT 2, then the receiving CCEP interface is shut down.
5. Multiple interfaces can be shut down in this case, depending on the sequence in which loops are detected.
6. In addition, to avoid CCEP interfaces from being shut down over a CEP interface, thje user can configure a CCEP port not to be shut down.

Configuring LD protocol

Follow these steps to configure loop detection (LD) protocol globally and at the interface and VLAN level.

1. Enter global configuration mode.

```
device# configure terminal
```

2. Enter the **protocol loop-detection** command to enable loop detection, enter Protocol Loop Detect configuration mode, and configure a variety of global options.

```
device(config)# protocol loop-detection
```

3. (Optional) Enter the **hello-interval** command to change the hello interval from the default.

```
device(config-loop-detect)# hello-interval 2000
```

4. (Optional) Enter the **shutdown-time** command to change from the default the interval after which an interface that is shut down by loop detection (LD) protocol is automatically reenabled.

```
device(config-loop-detect)# shutdown-tiome 20
```

5. (Optional) Enter the **raslog-duration** command to change from default the interval between RASLog messages that are sent when a port is disabled by the loop detection (LD) protocol.

```
device(config-loop-detect)# raslog-duration 20
```

6. Enable LD at the interface level.
   a) In global configuration mode, specify an interface (either an Ethernet interface or a port-channel interface).

```
device(config)# interface ethernet 2/6
```

   b) In interface subtype configuration mode, enter the **loop-detection** command.

```
device(conf-if-eth-2/6)# loop-detection
```

7. Enable LD at the VLAN level.
   a) In global configuration mode, create a VLAN.

```
device(config)# vlan 5
```

   b) In VLAN configuration mode, enter the **loop-detection** command.

```
device(config-vlan-5)# loop-detection
```

8. Associate the VLAN with an interface.
   a) In global configuration mode, specify an interface (either an Ethernet interface or a port-channel interface).

```
device(config)# interface ethernet 2/6
```

   b) In interface subtype configuration mode, enter the **loop-detection vlan** command and specify a VLAN. (The VLAN must already be created.)

```
device(conf-if-eth-2/6)# loop-detection vlan 5
```

9. (Optional) Disable the shutting down of an interface (Ethernet or port-channel) as a result of the loop detection (LD) protocol.
   a) In global configuration mode, specify an interface (either an Ethernet interface or a port-channel interface).

```
device(config)# interface ethernet 2/6
```

   b) In interface subtype configuration mode, enter the **loop-detection shutdown-disable** command.

```
device(conf-if-eth-2/6)# loop-detection shutdown-disable
```

10. (Optional) Disable the shutting down of an interface (Ethernet or port-channel) as a result of the loop detection (LD) protocol.
    a) In global configuration mode, specify an interface (either an Ethernet interface or a port-channel interface).

```
device(config)# interface ethernet 2/6
```

    b) In interface subtype configuration mode, enter the **loop-detection shutdown-disable** command.

```
device(conf-if-eth-2/6)# loop-detection shutdown-disable
```

11. Confirm the LD configuration, using the **show loop-detection** command with a variety of options.
    a) To display LD information at the system level, enter the **show loop-detection** command as in the following example.

```
device# show loop-detection
Strict Mode:
------------------------

Number of loop-detection instances enabled: 1

Interface: eth 2/6
        Enabled on VLANs: 100
        Shutdown Disable:   No
        Interface status: UP
        Auto enable in: Never

Packet Statistics:
vlan            sent            rcvd        disable-count
100             100             0           0

Loose Mode:
------------------------

Number of LD instances:   2
Disabled Ports:                  2/7

Packet Statistics:
vlan            sent            rcvd        disable-count
100             100             0           0
```

    b) To display ports disabled by LD, enter the **show loop-detection disabled-ports** command as in the following example.

```
device# show loop-detection disabled-ports
Ports disabled by loop detection
-----------------------------------------
port       age(min)       disable cause
2/6        5              Disabled by Self
```

    c) To display global LD configuration values, enter the **show loop-detection globals** command.

```
device# show loop-detection globals
Loop Detection:                 Disabled
Shutdown-time (minutes):        0
Hello-time (msec):              1000
Raslog-duration (minutes):      10
```

12. Use the **clear loop-detection** command in privileged EXEC mode with a variety of options to reenable ports that were disabled by LD and clear the LD statistics.
    a) To enable LD-disabled ports and clear LD statistics on all interfaces, enter the **clear loop-detection** command.

```
device# clear loop-detection
```

    b) To enable LD-disabled ports and clear LD statistics on an Ethernet interface, enter the **clear loop-detection interface ethernet** command.

```
device# clear loop-detection interface ethernet 2/6
```

    c) To enable LD-disabled ports and clear LD statistics on a port-channel interface, enter the **clear loop-detection interface port-channel** command.

```
device# clear loop-detection interface port-channel 20
```

    d) To enable LD-disabled ports and clear LD statistics on a VLAN, enter the **clear loop-detection interface vlan** command.

```
device# clear loop-detection interface vlan 10
```