Ethernet Switching Feature Guide

Modified: 2018-11-08

Copyright © 2018, Juniper Networks, Inc.

Juniper Networks, Inc.
1133 Innovation Way
Sunnyvale, California 94089
USA
408-745-2000
www.juniper.net

Juniper Networks, the Juniper Networks logo, Juniper, and Junos are registered trademarks of Juniper Networks, Inc. and/or its affiliates in the United States and other countries. All other trademarks may be property of their respective owners.

Juniper Networks assumes no responsibility for any inaccuracies in this document. Juniper Networks reserves the right to change, modify, transfer, or otherwise revise this publication without notice.

Ethernet Switching Feature Guide
Copyright © 2018 Juniper Networks, Inc. All rights reserved.

The information in this document is current as of the date on the title page.

YEAR 2000 NOTICE

Juniper Networks hardware and software products are Year 2000 compliant. Junos OS has no known time-related limitations through the year 2038. However, the NTP application is known to have some difficulty in the year 2036.

END USER LICENSE AGREEMENT

The Juniper Networks product that is the subject of this technical documentation consists of (or is intended for use with) Juniper Networks software. Use of such software is subject to the terms and conditions of the End User License Agreement ("EULA") posted at https://support.juniper.net/support/eula/. By downloading, installing or using such software, you agree to the terms and conditions of that EULA.

Table of Contents

 Copyright © 2018, Juniper Networks, Inc.

 Copyright © 2018, Juniper Networks, Inc.

List of Figures

Copyright © 2018, Juniper Networks, Inc.

List of Tables

Copyright © 2018, Juniper Networks, Inc.

Copyright © 2018, Juniper Networks, Inc.

About the Documentation

Documentation and Release Notes

To obtain the most current version of all Juniper Networks® technical documentation, see the product documentation page on the Juniper Networks website at https://www.juniper.net/documentation/.

If the information in the latest release notes differs from the information in the documentation, follow the product Release Notes.

Juniper Networks Books publishes books by Juniper Networks engineers and subject matter experts. These books go beyond the technical documentation to explore the nuances of network architecture, deployment, and administration. The current list can be viewed at https://www.juniper.net/books.

Using the Examples in This Manual

If you want to use the examples in this manual, you can use the **load merge** or the **load merge relative** command. These commands cause the software to merge the incoming configuration into the current candidate configuration. The example does not become active until you commit the candidate configuration.

If the example configuration contains the top level of the hierarchy (or multiple hierarchies), the example is a *full example*. In this case, use the **load merge** command.

If the example configuration does not start at the top level of the hierarchy, the example is a *snippet*. In this case, use the **load merge relative** command. These procedures are described in the following sections.

Merging a Full Example

To merge a full example, follow these steps:

1. From the HTML or PDF version of the manual, copy a configuration example into a text file, save the file with a name, and copy the file to a directory on your routing platform.

   For example, copy the following configuration to a file and name the file **ex-script.conf**. Copy the **ex-script.conf** file to the **/var/tmp** directory on your routing platform.

```
system {
  scripts {
    commit {
      file ex-script.xsl;
    }
  }
}
interfaces {
  fxp0 {
    disable;
    unit 0 {
      family inet {
        address 10.0.0.1/24;
      }
    }
  }
}
```

2. Merge the contents of the file into your routing platform configuration by issuing the **load merge** configuration mode command:

```
[edit]
user@host# load merge /var/tmp/ex-script.conf
load complete
```

Merging a Snippet

To merge a snippet, follow these steps:

1. From the HTML or PDF version of the manual, copy a configuration snippet into a text file, save the file with a name, and copy the file to a directory on your routing platform.

   For example, copy the following snippet to a file and name the file **ex-script-snippet.conf**. Copy the **ex-script-snippet.conf** file to the **/var/tmp** directory on your routing platform.

```
commit {
  file ex-script-snippet.xsl; }
```

 Copyright © 2018, Juniper Networks, Inc.

2. Move to the hierarchy level that is relevant for this snippet by issuing the following configuration mode command:

```
[edit]
user@host# edit system scripts
[edit system scripts]
```

3. Merge the contents of the file into your routing platform configuration by issuing the **load merge relative** configuration mode command:

```
[edit system scripts]
user@host# load merge relative /var/tmp/ex-script-snippet.conf
load complete
```

For more information about the **load** command, see CLI Explorer.

Documentation Conventions

Table 1 on page xiii defines notice icons used in this guide.

Table 1: Notice Icons

Icon	Meaning	Description
	Informational note	Indicates important features or instructions.
	Caution	Indicates a situation that might result in loss of data or hardware damage.
	Warning	Alerts you to the risk of personal injury or death.
	Laser warning	Alerts you to the risk of personal injury from a laser.
	Tip	Indicates helpful information.
	Best practice	Alerts you to a recommended use or implementation.

Table 2 on page xiv defines the text and syntax conventions used in this guide.

Table 2: Text and Syntax Conventions

Convention	Description	Examples
Bold text like this	Represents text that you type.	To enter configuration mode, type the **configure** command: user@host> **configure**
`Fixed-width text like this`	Represents output that appears on the terminal screen.	`user@host>` **show chassis alarms** `No alarms currently active`
Italic text like this	• Introduces or emphasizes important new terms. • Identifies guide names. • Identifies RFC and Internet draft titles.	• A policy *term* is a named structure that defines match conditions and actions. • *Junos OS CLI User Guide* • RFC 1997, *BGP Communities Attribute*
Italic text like this	Represents variables (options for which you substitute a value) in commands or configuration statements.	Configure the machine's domain name: [edit] root@# **set system domain-name** ***domain-name***
Text like this	Represents names of configuration statements, commands, files, and directories; configuration hierarchy levels; or labels on routing platform components.	• To configure a stub area, include the **stub** statement at the **[edit protocols ospf area area-id]** hierarchy level. • The console port is labeled **CONSOLE**.
< > (angle brackets)	Encloses optional keywords or variables.	**stub <default-metric *metric*>;**
| (pipe symbol)	Indicates a choice between the mutually exclusive keywords or variables on either side of the symbol. The set of choices is often enclosed in parentheses for clarity.	**broadcast | multicast** **(*string1* | *string2* | *string3*)**
# (pound sign)	Indicates a comment specified on the same line as the configuration statement to which it applies.	**rsvp { # Required for dynamic MPLS only**
[ ] (square brackets)	Encloses a variable for which you can substitute one or more values.	**community name members [*community-ids*]**

Indention and braces ({ })	Identifies a level in the configuration hierarchy.	[edit]
		routing-options {
		  static {
		    route default {
		      nexthop *address*;
		      retain;
		    }
		  }
		}

; (semicolon)	Identifies a leaf statement at a configuration hierarchy level.	

Copyright © 2018, Juniper Networks, Inc.

Table 2: Text and Syntax Conventions (continued)

Convention	Description	Examples
GUI Conventions		
Bold text like this	Represents graphical user interface (GUI) items you click or select.	• In the Logical Interfaces box, select **All Interfaces**. • To cancel the configuration, click **Cancel**.
> (bold right angle bracket)	Separates levels in a hierarchy of menu selections.	In the configuration editor hierarchy, select **Protocols>Ospf**.

Documentation Feedback

We encourage you to provide feedback so that we can improve our documentation. You can use either of the following methods:

- Online feedback system—Click TechLibrary Feedback, on the lower right of any page on the Juniper Networks TechLibrary site, and do one of the following:

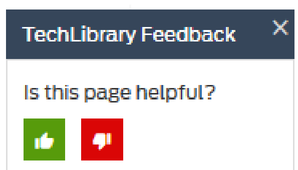

  - Click the thumbs-up icon if the information on the page was helpful to you.
  - Click the thumbs-down icon if the information on the page was not helpful to you or if you have suggestions for improvement, and use the pop-up form to provide feedback.

- E-mail—Send your comments to **techpubs-comments@juniper.net**. Include the document or topic name, URL or page number, and software version (if applicable).

Requesting Technical Support

Technical product support is available through the Juniper Networks Technical Assistance Center (JTAC). If you are a customer with an active J-Care or Partner Support Service support contract, or are covered under warranty, and need post-sales technical support, you can access our tools and resources online or open a case with JTAC.

- JTAC policies—For a complete understanding of our JTAC procedures and policies, review the *JTAC User Guide* located at https://www.juniper.net/us/en/local/pdf/resource-guides/7100059-en.pdf.
- Product warranties—For product warranty information, visit https://www.juniper.net/support/warranty/.

Copyright © 2018, Juniper Networks, Inc.

- JTAC hours of operation—The JTAC centers have resources available 24 hours a day, 7 days a week, 365 days a year.

Self-Help Online Tools and Resources

For quick and easy problem resolution, Juniper Networks has designed an online self-service portal called the Customer Support Center (CSC) that provides you with the following features:

- Find CSC offerings: https://www.juniper.net/customers/support/
- Search for known bugs: https://prsearch.juniper.net/
- Find product documentation: https://www.juniper.net/documentation/
- Find solutions and answer questions using our Knowledge Base: https://kb.juniper.net/
- Download the latest versions of software and review release notes: https://www.juniper.net/customers/csc/software/
- Search technical bulletins for relevant hardware and software notifications: https://kb.juniper.net/InfoCenter/
- Join and participate in the Juniper Networks Community Forum: https://www.juniper.net/company/communities/
- Open a case online in the CSC Case Management tool: https://www.juniper.net/cm/

To verify service entitlement by product serial number, use our Serial Number Entitlement (SNE) Tool: https://entitlementsearch.juniper.net/entitlementsearch/

Opening a Case with JTAC

You can open a case with JTAC on the Web or by telephone.

- Use the Case Management tool in the CSC at https://www.juniper.net/cm/.
- Call 1-888-314-JTAC (1-888-314-5822 toll-free in the USA, Canada, and Mexico).

For international or direct-dial options in countries without toll-free numbers, see https://www.juniper.net/support/requesting-support.html.

 Copyright © 2018, Juniper Networks, Inc.

PART 1

Bridging and VLANs

Copyright © 2018, Juniper Networks, Inc.

CHAPTER 1

Using Bridging and VLANs

Overview of Layer 2 Networking

Layer 2, also known as the Data Link Layer, is the second level in the seven-layer OSI reference model for network protocol design. Layer 2 is equivalent to the link layer (the lowest layer) in the TCP/IP network model. Layer2 is the network layer used to transfer data between adjacent network nodes in a wide area network or between nodes on the same local area network.

A *frame* is a protocol data unit, the smallest unit of bits on a Layer 2 network. Frames are transmitted to and received from devices on the same local area network (LAN). Unilke bits, frames have a defined structure and can be used for error detection, control plane activities and so forth. Not all frames carry user data. The network uses some frames to control the data link itself..

At Layer 2, *unicast* refers to sending frames from one node to a single other node, whereas *multicast* denotes sending traffic from one node to multiple nodes, and *broadcasting* refers to the transmission of frames to all nodes in a network. A *broadcast domain* is a logical division of a network in which all nodes of that network can be reached at Layer 2 by a broadcast.

Segments of a LAN can be linked at the frame level using *bridges*. Bridging creates separate broadcast domains on the LAN, creating VLANs, which are independent logical networks that group together related devices into separate network segments. The grouping of devices on a VLAN is independent of where the devices are physically located in the LAN. Without bridging and VLANs, all devices on the Ethernet LAN are in a single broadcast domain, and all the devices detect all the packets on the LAN.

Forwarding is the relaying of packets from one network segment to another by nodes in the network. On a VLAN, a frame whose origin and destination are in the same VLAN are forwarded only within the local VLAN. A network segment is a portion of a computer network wherein every device communicates using the same physical layer.

Layer 2 contains two sublayers:

- Logical link control (LLC) sublayer, which is responsible for managing communications links and handling frame traffic.
- Media access control (MAC) sublayer, which governs protocol access to the physical network medium. By using the MAC addresses that are assigned to all ports on a switch, multiple devices on the same physical link can uniquely identify one another.

  The ports, or interfaces, on a switch operate in either access mode, tagged-access, or trunk mode:

  - *Access mode* ports connect to a network device such as a desktop computer, an IP telephone, a printer, a file server, or a security camera. The port itself belongs to a single VLAN. The frames transmitted over an access interface are normal Ethernet frames. By default, all ports on a switch are in access mode.
  - *Tagged-Access mode* ports connect to a network device such as a desktop computer, an IP telephone, a printer, a file server, or a security camera. The port itself belongs to a single VLAN. The frames transmitted over an access interface are normal Ethernet frames. By default, all ports on a switch are in access mode. Tagged-access mode accommodates cloud computing, specifically scenarios including virtual machines or virtual computers. Because several virtual computers can be included on one physical server, the packets generated by one server can contain an aggregation of VLAN packets from different virtual machines on that server. To accommodate this situation, tagged-access mode reflects packets back to the physical server on the same downstream port when the destination address of the packet was learned on that downstream port. Packets are also reflected back to the physical server on the downstream port when the destination has not yet been learned. Therefore, the third interface mode, tagged access, has some characteristics of access mode and some characteristics of trunk mode:
  - *Trunk mode* ports handle traffic for multiple VLANs, multiplexing the traffic for all those VLANs over the same physical connection. Trunk interfaces are generally used to interconnect switches to other devices or switches.

    With native VLAN configured, frames that do not carry VLAN tags are sent over the trunk interface. If you have a situation where packets pass from a device to a switch in access mode, and you want to then send those packets from the switch over a trunk port, use native VLAN mode. Configure the single VLAN on the switch's port (which is in access mode) as a native VLAN. The switch's trunk port will then treat those frames differently than the other tagged packets. For example, if a trunk port has three VLANs, 10, 20, and 30, assigned to it with VLAN 10 being the native VLAN,

 Copyright © 2018, Juniper Networks, Inc.

frames on VLAN 10 that leave the trunk port on the other end have no 802.1Q header (tag). There is another native VLAN option. You can have the switch add and remove tags for untagged packets. To do this, you first configure the single VLAN as a native VLAN on a port attached to a device on the edge. Then, assign a VLAN ID tag to the single native VLAN on the port connected to a device. Last, add the VLAN ID to the trunk port. Now, when the switch receives the untagged packet, it adds the ID you specified and sends and receives the tagged packets on the trunk port configured to accept that VLAN.

Including the sublayers, Layer 2 on the QFX Series supports the following functionality:

- Unicast, multicast, and broadcast traffic.
- Bridging.
- VLAN 802.1Q—Also known as *VLAN tagging*, this protocol allows multiple bridged networks to transparently share the same physical network link by adding VLAN tags to an Ethernet frame.
- Extension of Layer 2 VLANs across multiple switches using Spanning Tree Protocol (STP) prevents looping across the network.
- *MAC learning*, including per-VLAN MAC learning and Layer 2 learning suppression–This process obtains the MAC addresses of all the nodes on a network
- Link aggregation—This process groups of Ethernet interfaces at the physical layer to form a single link layer interface, also known as a *link aggregation group (LAG)* or LAG bundle

  NOTE: **Link aggregation is not supported on NFX150 devices.**

- Storm control on the physical port for unicast, multicast, and broadcast

  NOTE: **Storm control is not supported on NFX150 devices.**

- STP support, including 802.1d, RSTP, MSTP, and Root Guard

Related Documentation

- Understanding Bridging and VLANs on Switches on page 9

Understanding Layer 2 Broadcasting on Switches

In a Layer 2 network, *broadcasting* refers to sending traffic to all nodes on a network.

Layer 2 broadcast traffic stays within a local area network (LAN) boundary; known as the *broadcast domain*. Layer 2 broadcast traffic is sent to the broadcast domain using a MAC address of FF:FF:FF:FF:FF:FF. Every device in the broadcast domain recognizes this MAC address and passes the broadcast traffic on to other devices in the broadcast

Copyright © 2018, Juniper Networks, Inc.

domain, if applicable. Broadcasting can be compared to unicasting (sending traffic to a single node) or multicasting (delivering traffic to a group of nodes simultaneously).

Layer 3 broadcast traffic, however, is sent to all devices in a network using a broadcast network address. For example, if your network address is 10.0.0.0, the broadcast network address is 10.255.255.255. In this case, only devices that belong to the 10.0.0.0 network receive the Layer 3 broadcast traffic. Devices that do not belong to this network drop the traffic.

Broadcasting is used in the following situations:

- Address Resolution Protocol (ARP) uses broadcasting to map MAC addresses to IP addresses. ARP dynamically binds the IP address (the logical address) to the correct MAC address. Before IP unicast packets can be sent, ARP discovers the MAC address used by the Ethernet interface where the IP address is configured.
- Dynamic Host Configuration Protocol (DHCP) uses broadcasting to dynamically assign IP addresses to hosts on a network segment or subnet.
- Routing protocols use broadcasting to advertise routes.

Excessive broadcast traffic can sometimes create a broadcast storm. A broadcast storm occurs when messages are broadcast on a network and each message prompts a receiving node to respond by broadcasting its own messages on the network. This, in turn, prompts further responses that create a snowball effect. The LAN is suddenly flooded with packets, creating unnecessary traffic that leads to poor network performance or even a complete loss of network service.

Related Documentation

- Overview of Layer 2 Networking on page 3
- *Understanding Storm Control*
- Understanding Bridging and VLANs on Switches on page 9

Layer 2 Learning and Forwarding for VLANs Overview

Understanding Layer 2 Forwarding Tables on Switches, Routers and NFX Series Devices

You can configure Layer 2 MAC address and VLAN learning and forwarding properties in support of Layer 2 bridging. Unicast media access control (MAC) addresses are learned to avoid flooding the packets to all the ports in a VLAN. A source MAC entry is created in its source and destination MAC tables for each MAC address learned from packets received on ports that belong to the VLAN.

When you configure a VLAN, Layer 2 address learning is enabled by default. The VLAN learns unicast media access control (MAC) addresses to avoid flooding the packets to all the ports in the VLAN. Each VLAN creates a source MAC entry in its source and destination MAC tables for each source MAC address learned from packets received on the ports that belong to the VLAN.

 Copyright © 2018, Juniper Networks, Inc.

NOTE: Traffic is not flooded back onto the interface on which it was received. However, because this “split horizon” occurs at a late stage, the packet statistics displayed by commands such as **show interfaces queue** will include flood traffic.

You can optionally disable MAC learning either for the entire device or for a specific VLAN or logical interface. You can also configure the following Layer 2 learning and forwarding properties:

- Timeout interval for MAC entries
- Static MAC entries for logical interfaces only
- Limit to the number of MAC addresses learned from a specific logical interface or from all the logical interfaces in a VLAN
- Size of the MAC address table for the VLAN
- MAC accounting for a VLAN

For more information about how to configure VLANs and virtual switches, see *Configuring a VLAN* and *Configuring a Layer 2 Virtual Switch on an EX Series Switch*.

Understanding Layer 2 Forwarding Tables on Security Devices

The SRX Series device maintains forwarding tables that contain MAC addresses and associated interfaces for each Layer 2 VLAN. When a packet arrives with a new source MAC address in its frame header, the device adds the MAC address to its forwarding table and tracks the interface at which the packet arrived. The table also contains the corresponding interface through which the device can forward traffic for a particular MAC address.

If the destination MAC address of a packet is unknown to the device (that is, the destination MAC address in the packet does not have an entry in the forwarding table), the device duplicates the packet and floods it on all interfaces in the VLAN other than the interface on which the packet arrived. This is known as *packet flooding* and is the default behavior for the device to determine the outgoing interface for an unknown destination MAC address. Packet flooding is performed at two levels: packets are flooded to different zones as permitted by configured Layer 2 security policies, and packets are also flooded to different interfaces with the same VLAN identifier within the same zone. The device learns the forwarding interface for the MAC address when a reply with that MAC address arrives at one of its interfaces.

You can specify that the SRX Series device use ARP queries and traceroute requests (which are ICMP echo requests with the time-to-live values set to 1) instead of packet flooding to locate an unknown destination MAC address. This method is considered more secure than packet flooding because the device floods ARP queries and traceroute packets—not the initial packet—on all interfaces. When ARP or traceroute flooding is used, the original packet is dropped. The device broadcasts an ARP or ICMP query to all other devices on the same subnetwork, requesting the device at the specified destination

Copyright © 2018, Juniper Networks, Inc.

IP address to send back a reply. Only the device with the specified IP address replies, which provides the requestor with the MAC address of the responder.

ARP allows the device to discover the destination MAC address for a unicast packet if the destination IP address is in the same subnetwork as the ingress IP address. (The ingress IP address refers to the IP address of the last device to send the packet to the device. The device might be the source that sent the packet or a router forwarding the packet.) Traceroute allows the device to discover the destination MAC address even if the destination IP address belongs to a device in a subnetwork beyond that of the ingress IP address.

When you enable ARP queries to locate an unknown destination MAC address, traceroute requests are also enabled. You can also optionally specify that traceroute requests not be used; however, the device can then discover destination MAC addresses for unicast packets only if the destination IP address is in the same subnetwork as the ingress IP address.

Whether you enable ARP queries and traceroute requests or ARP-only queries to locate unknown destination MAC addresses, the SRX Series device performs the following series of actions:

1. The device notes the destination MAC address in the initial packet. The device adds the source MAC address and its corresponding interface to its forwarding table, if they are not already there.

2. The device drops the initial packet.

3. The device generates an ARP query packet and optionally a traceroute packet and floods those packets out all interfaces except the interface on which the initial packet arrived.

   ARP packets are sent out with the following field values:

   - Source IP address set to the IP address of the IRB
   - Destination IP address set to the destination IP address of the original packet
   - Source MAC address set to the MAC address of the IRB
   - Destination MAC address set to the broadcast MAC address (all `0xf`)

   Traceroute (ICMP echo request or ping) packets are sent out with the following field values:

   - Source IP address set to the IP address of the original packet
   - Destination IP address set to the destination IP address of the original packet
   - Source MAC address set to the source MAC address of the original packet
   - Destination MAC address set to the destination MAC address of the original packet
   - Time-to-live (TTL) set to **1**

 Copyright © 2018, Juniper Networks, Inc.

4. Combining the destination MAC address from the initial packet with the interface leading to that MAC address, the device adds a new entry to its forwarding table.

5. The device forwards all subsequent packets it receives for the destination MAC address out the correct interface to the destination.

Related Documentation

- Layer 2 Learning and Forwarding for VLANs Overview on page 6
- *Layer 2 Transparent Mode Overview*
- *Understanding Integrated Routing and Bridging*
- *Example: Configuring an IRB Interface on a Security Device*
- *Example: Configuring the Default Learning for Unknown MAC Addresses*

Understanding Bridging and VLANs on Switches

Network switches use Layer 2 bridging protocols to discover the topology of their LAN and to forward traffic toward destinations on the LAN. This topic explains the following concepts regarding bridging and VLANs:

History of VLANs

Ethernet LANs were originally designed for small, simple networks that primarily carried text. However, over time, the type of data carried by LANs grew to include voice, graphics, and video. This more complex data, when combined with the ever-increasing speed of transmission, eventually became too much of a load for the original Ethernet LAN design. Multiple packet collisions were significantly slowing down the larger LANs.

The IEEE 802.1D-2004 standard helped evolve Ethernet LANs to cope with the higher data and transmission requirements by defining the concept of *transparent bridging* (generally called simply *bridging*). Bridging divides a single physical LAN (now called a single *broadcast domain*) into two or more virtual LANs, or VLANs. Each *VLAN* is a

Copyright © 2018, Juniper Networks, Inc.

collection of some of the LAN nodes grouped together to form individual broadcast domains.

When VLANs are grouped logically by function or organization, a significant percentage of data traffic stays within the VLAN. This relieves the load on the LAN because all traffic no longer has to be forwarded to all nodes on the LAN. A VLAN first transmits packets within the VLAN, thereby reducing the number of packets transmitted on the entire LAN. Because packets whose origin and destination are in the same VLAN are forwarded only within the local VLAN, packets that are not destined for the local VLAN are the only ones forwarded to other broadcast domains. This way, bridging and VLANs limit the amount of traffic flowing across the entire LAN by reducing the possible number of collisions and packet retransmissions within VLANs and on the LAN as a whole.

How Bridging of VLAN Traffic Works

Because the objective of the IEEE 802.1D-2004 standard was to reduce traffic and therefore reduce potential transmission collisions for Ethernet, a system was implemented to reuse information. Instead of having a switch go through a location process every time a frame is sent to a node, the transparent bridging protocol allows a switch to record the location of known nodes. When packets are sent to nodes, those destination node locations are stored in address-lookup tables called *Ethernet switching tables*. Before sending a packet, a switch using bridging first consults the switching tables to see if that node has already been located. If the location of a node is known, the frame is sent directly to that node.

Transparent bridging uses five mechanisms to create and maintain Ethernet switching tables on the switch:

- Learning
- Forwarding
- Flooding
- Filtering
- Aging

The key bridging mechanism used by LANs and VLANs is *learning*. When a switch is first connected to an Ethernet LAN or VLAN, it has no information about other nodes on the network. As packets are sent, the switch learns the embedded MAC addresses of the sending nodes and stores them in the Ethernet switching table, along with two other pieces of information—the interface (or port) on which the traffic was received on the destination node and the time the address was learned.

Learning allows switches to then do *forwarding*. By consulting the Ethernet switching table to see whether the table already contains the frame's destination MAC address, switches save time and resources when forwarding packets to the known MAC addresses. If the Ethernet switching table does not contain an entry for an address, the switch uses flooding to learn that address.

Flooding finds a particular destination MAC address without using the Ethernet switching table. When traffic originates on the switch and the Ethernet switching table does not

Copyright © 2018, Juniper Networks, Inc.

yet contain the destination MAC address, the switch first floods the traffic to all other interfaces within the VLAN. When the destination node receives the flooded traffic, it can send an acknowledgment packet back to the switch, allowing it to learn the MAC address of the node and add the address to its Ethernet switching table.

Filtering, the fourth bridging mechanism, is how broadcast traffic is limited to the local VLAN whenever possible. As the number of entries in the Ethernet switching table grows, the switch pieces together an increasingly complete picture of the VLAN and the larger LAN—it learns which nodes are in the local VLAN and which are on other network segments. The switch uses this information to filter traffic. Specifically, for traffic whose source and destination MAC addresses are in the local VLAN, filtering prevents the switch from forwarding this traffic to other network segments.

To keep entries in the Ethernet switching table current, the switch uses a fifth bridging mechanism, *aging*. Aging is the reason that the Ethernet switching table entries include timestamps. Each time the switch detects traffic from a MAC address, it updates the timestamp. A timer on the switch periodically checks the timestamp, and if it is older than a user-configured value, the switch removes the node's MAC address from the Ethernet switching table. This aging process eventually flushes unavailable network nodes out of the Ethernet switching table.

Packets Are Either Tagged or Untagged

When an Ethernet LAN is divided into VLANs, each VLAN is identified by a unique 802.1Q ID. The number of available VLANs and VLAN IDs are listed below:

- On a switch running ELS software, you can configure 4093 VLANs using VLAN IDs 1 through 4094, while VLAN IDs 0 and 4095 are reserved by Junos OS and cannot be assigned.
- On a switch running non-ELS software, you can configure 4091 VLANs using VLAN IDs 1-4094.

Ethernet packets include a tag protocol identifier (TPID) EtherType field, which identifies the protocol being transported. When a device within a VLAN generates a packet, this field includes a value of 0x8100, which indicates that the packet is a VLAN-tagged packet. The packet also has a VLAN ID field that includes the unique 802.1Q ID, which identifies the VLAN to which the packet belongs.

Junos OS switches support the TPID value 0x9100 for Q-in-Q on switches. In addition to the TPID EtherType value of 0x8100, EX Series switches that do not support the Enhanced Layer 2 Software (ELS) configuration style also support values of 0x88a8 (Provider Bridging and Shortest Path Bridging) and 0x9100 (Q-inQ).

For a simple network that has only a single VLAN, all packets include a default 802.1Q tag, which is the only VLAN membership that does not mark the packet as tagged. These packets are untagged packets.

NOTE: Q-in-Q tunnelling is not supported on NFX150 devices.

Copyright © 2018, Juniper Networks, Inc.

Switch Interface Modes—Access, Trunk, or Tagged Access

Ports, or interfaces, on a switch operate in one of three modes:

- Access mode
- Trunk mode
- Tagged-access mode

Access Mode

An interface in access mode connects a switch to a single network device, such as a desktop computer, an IP telephone, a printer, a file server, or a security camera. Access interfaces accept only untagged packets.

By default, when you boot a switch that runs Junos OS that does not support ELS and use the factory default configuration, or when you boot such a switch and do not explicitly configure a port mode, all interfaces on the switch are in access mode and accept only untagged packets from the VLAN named **default**. You can optionally configure another VLAN and use that VLAN instead of **default**.

On a switch that supports ELS, the VLAN named **default** is not supported. Therefore, on such switches, you must explicitly configure at least one VLAN, even if your network is simple and you want only one broadcast domain to exist. After you assign an interface to a VLAN, the interface functions in access mode.

For switches that run either type of software, you can also configure a trunk port or interface to accept untagged packets from a user-configured VLAN. For details about this concept (native VLAN), see "Trunk Mode and Native VLAN" on page 13.

Trunk Mode

Trunk mode interfaces are generally used to connect switches to one another. Traffic sent between switches can then consist of packets from multiple VLANs, with those packets multiplexed so that they can be sent over the same physical connection. Trunk interfaces usually accept only tagged packets and use the VLAN ID tag to determine both the packets' VLAN origin and VLAN destination.

On a switch that runs software that does not support ELS, an untagged packet is not recognized on a trunk port unless you configure additional settings on that port.

On a switch that runs Junos OS that supports ELS, a trunk port recognizes untagged control packets for protocols such as the Link Aggregation Control Protocol (LACP) and the Link Layer Discovery Protocol (LLDP). However, the trunk port does not recognize untagged data packets unless you configure additional settings on that port.

NOTE: LACP is not supported on NFX150 devices.

In the rare case where you want untagged packets to be recognized by a trunk port on switches that run either type of software, you must configure the single VLAN on a trunk

Copyright © 2018, Juniper Networks, Inc.

port as a *native VLAN*. For more information about native VLANs, see "Trunk Mode and Native VLAN" on page 13.

Trunk Mode and Native VLAN

On a switch that runs Junos OS that does not support ELS, a trunk port does not recognize packets that do not include VLAN tags, which are also known an untagged packets. On a switch that runs Junos OS that supports ELS, a trunk port recognizes untagged control packets, but it does not recognize untagged data packets. With native VLAN configured, untagged packets that a trunk port normally does not recognize are sent over the trunk interface. In a situation where packets pass from a device, such as an IP phone or printer, to a switch in access mode, and you want those packets sent from the switch over a trunk port, use native VLAN mode. Create a native VLAN by configuring a VLAN ID for it, and specify that the trunk port is a member of the native VLAN.

The switch's trunk port will then treat those packets differently than the other tagged packets. For example, if a trunk port has three VLANs, 10, 20, and 30, assigned to it with VLAN 10 being the native VLAN, packets on VLAN 10 that leave the trunk port on the other end have no 802.1Q header (tag).

There is another native VLAN option for switches that do not support ELS. You can have the switch add and remove tags for untagged packets. To do this, you first configure the single VLAN as a native VLAN on a port attached to a device on the edge. Then, assign a VLAN ID tag to the single native VLAN on the port connected to a device. Last, add the VLAN ID to the trunk port. Now, when the switch receives the untagged packet, it adds the ID you specified and sends and receives the tagged packets on the trunk port configured to accept that VLAN.

Tagged-Access Mode

Only switches that run Junos OS not using the ELS configuration style support tagged-access mode. Tagged-access mode accommodates cloud computing, specifically scenarios including virtual machines or virtual computers. Because several virtual computers can be included on one physical server, the packets generated by one server can contain an aggregation of VLAN packets from different virtual machines on that server. To accommodate this situation, tagged-access mode reflects packets back to the physical server on the same downstream port when the destination address of the packet was learned on that downstream port. Packets are also reflected back to the physical server on the downstream port when the destination has not yet been learned. Therefore, the third interface mode, tagged access, has some characteristics of access mode and some characteristics of trunk mode:

- Like access mode, tagged-access mode connects the switch to an access layer device. Unlike access mode, tagged-access mode is capable of accepting VLAN tagged packets.
- Like trunk mode, tagged-access mode accepts VLAN tagged packets from multiple VLANs. Unlike trunk port interfaces, which are connected at the core/distribution layer, tagged-access port interfaces connect devices at the access layer.

  Like trunk mode, tagged-access mode also supports native VLAN.

Copyright © 2018, Juniper Networks, Inc.

NOTE: Control packets are never reflected back on the downstream port.

Additional Advantages of Using VLANs

In addition to reducing traffic and thereby speeding up the network, VLANs have the following advantages:

- VLANs provide segmentation services traditionally provided by routers in LAN configurations, thereby reducing hardware equipment costs.
- Packets coupled to a VLAN can be reliably identified and sorted into different domains. You can contain broadcasts within parts of the network, thereby freeing up network resources. For example, when a DHCP server is plugged into a switch and starts broadcasting its presence, you can prevent some hosts from accessing it by using VLANs to split up the network.
- For security issues, VLANs provide granular control of the network because each VLAN is identified by a single IP subnetwork. All packets passing in and out of a VLAN are consistently tagged with the VLAN ID of that VLAN, thereby providing easy identification, because a VLAN ID on a packet cannot be altered. (For a switch that runs Junos OS that does not support ELS, we recommend that you avoid using 1 as a VLAN ID, because that ID is a default value.)
- VLANs react quickly to host relocation—this is also due to the persistent VLAN tag on packets.
- On an Ethernet LAN, all network nodes must be physically connected to the same network. In VLANs, the physical location of nodes is not important—you can group network devices in any way that makes sense for your organization, such as by department or business function, types of network nodes, or physical location.

Maximum VLANs and VLAN Members Per Switch

The number of VLANs supported per switch varies for each switch. Use the configuration-mode command **set vlans *vlan-name* vlan-id ?** to determine the maximum number of VLANs allowed on a switch. You cannot exceed this VLAN limit because you have to assign a specific ID number when you create a VLAN—you could overwrite one of the numbers, but you cannot exceed the limit.

You can, however, exceed the recommended VLAN member maximum for a switch.

On a switch that runs Junos OS that does not support the ELS configuration style, the maximum number of VLAN members allowed on the switch is eight times the maximum number of VLANs that the switch supports (vmember limit = vlan max * 8). If the configuration of the switch exceeds the recommended VLAN member maximum, a warning message appears when you commit the configuration. If you commit the configuration despite the warning, the commit succeeds, but there is a risk of the Ethernet switching process (eswd) failing as a result of memory allocation failure.

Copyright © 2018, Juniper Networks, Inc.

On most switches running Junos OS that supports ELS, the maximum number of VLAN members allowed on the switch is 24 times the maximum number of VLANs that the switch supports (vmember limit = vlan max * 24). If the configuration of the switch exceeds the recommended VLAN member maximum, a warning message appears in the system log (syslog).

On an EX Series switch that runs Junos OS that supports ELS, the maximum number of VLAN members allowed on the switch is as follows:

- EX4300—24 times the maximum number of VLANs that the switch supports (vmember limit = *vlan max* * 24)
- EX3400—16 times the maximum number of VLANs that the switch supports (vmember limit = *vlan max* * 16)
- EX2300—8 times the maximum number of VLANs that the switch supports (vmember limit = *vlan max* * 8)

A QFabric system supports up to 131,008 VLAN members (vmembers) on a single network node group, server node group, or redundant server node group. The number of vmembers is calculated by multiplying the maximum number of VLANs by 32.

For example, to calculate how many interfaces are required to support 4,000 VLANs, divide the maximum number of vmembers (128,000) by the number of configured VLANs (4,000). In this case, 32 interfaces are required.

On network Node groups and server Node groups, you can configure link aggregation groups (LAGs) across multiple interfaces. Each LAG and VLAN combination is considered a vmember.

NOTE: LAG is not supported on NFX150 devices.

A Virtual Chassis Fabric supports up to 512,000 vmembers. The number of vmembers is based on the number of VLANs, and the number of interfaces configured in each VLAN.

A Default VLAN Is Configured on Most Switches

Some switches running Junos OS that do not support the ELS configuration style are preconfigured with a VLAN named **default** that does not tag packets and operates only with untagged packets. On these switches, each interface already belongs to the VLAN named **default** and all traffic uses this VLAN until you configure more VLANs and assign traffic to those VLANs.

EX Series switches that run Junos OS with the ELS configuration style do not support a default VLAN. The following EX Series switches running Junos OS not supporting the ELS configuration style are not preconfigured to belong to **default** or any other VLAN:

- Modular switches, such as the EX8200 switches and EX6200 switches
- Switches that are part of a Virtual Chassis

Copyright © 2018, Juniper Networks, Inc.

The reason that these switches are not preconfigured is that the physical configuration in both situations is flexible. There is no way of knowing which line cards have been inserted in either the EX8200 switch or EX6200 switch. There is also no way of knowing which switches are included in the Virtual Chassis. Switch interfaces in these two cases must first be defined as Ethernet switching interfaces. After an interface is defined as an Ethernet switching interface, the default VLAN appears in the output from the ? help and other commands.

NOTE: When a Juniper Networks EX4500 Ethernet Switch, EX4200 Ethernet Switch, EX3300 Ethernet Switch, QFX3500 or QFX3600 switch is interconnected with other switches in a Virtual Chassis configuration, each individual switch that is included as a member of the configuration is identified with a member ID. The member ID functions as an FPC slot number. When you are configuring interfaces for a Virtual Chassis configuration, you specify the appropriate member ID (0 through 9) as the slot element of the interface name. The default factory settings for a Virtual Chassis configuration include FPC 0 as a member of the default VLAN because FPC 0 is configured as part of the ethernet-switching family. In order to include FPC 1 through FPC 9 in the default VLAN, add the ethernet-switching family to the configurations for those interfaces.

NOTE: You cannot configure a default VLAN on NFX150 devices.

Assigning Traffic to VLANs

You can assign traffic on any switch to a particular VLAN by referencing either the interface port of the traffic or the MAC addresses of devices sending traffic.

NOTE: Two logical interfaces that are configured on the same physical interface cannot be mapped to the same VLAN.

Assign VLAN Traffic According to the Interface Port Source

This method is most commonly used to assign traffic to VLANs. In this case, you specify that all traffic received on a particular switch interface is assigned to a specific VLAN. You configure this VLAN assignment when you configure the switch, by using either the VLAN number (called a VLAN ID) or by using the VLAN name, which the switch then translates into a numeric VLAN ID. This method is referred to simply as creating a VLAN because it is the most commonly used method.

Assign VLAN Traffic According to the Source MAC Address

In this case, all traffic received from a specific MAC address is forwarded to a specific egress interface (next hop) on the switch. MAC-based VLANs are either static (named MAC addresses configured one at a time) or dynamic (configured using a RADIUS server).

 Copyright © 2018, Juniper Networks, Inc.

To configure a static MAC-based VLAN on a switch that supports ELS, see *Adding a Static MAC Address Entry to the Ethernet Switching Table (CLI Procedure)*.To configure a static MAC-based VLAN on a switch that does not support ELS, see *Adding a Static MAC Address Entry to the Ethernet Switching Table (CLI Procedure)*.

For information about using 802.1X authentication to authenticate end devices and allow access to dynamic VLANs configured on a RADIUS server, see *Understanding Dynamic VLAN Assignment Using RADIUS Attributes*. You can optionally implement this feature to offload the manual assignment of VLAN traffic to automated RADIUS server databases.

Forwarding VLAN Traffic

To pass traffic within a VLAN, the switch uses Layer 2 forwarding protocols, including IEEE 802.1Q spanning-tree protocols.

To pass traffic between two VLANs, the switch uses standard Layer 3 routing protocols, such as static routing, OSPF, and RIP. The same interfaces that support Layer 2 bridging protocols also support Layer 3 routing protocols, providing multilayer switching.

To pass traffic from a single device on an access port to a switch and then pass those packets on a trunk port, use the native mode configuration previously discussed under "Trunk Mode" on page 12.

VLANs Communicate with Integrated Routing and Bridging Interfaces or Routed VLAN Interfaces

Traditionally, switches sent traffic to hosts that were part of the same broadcast domain (VLAN) but routers were needed to route traffic from one broadcast domain to another. Also, only routers performed other Layer 3 functions such as traffic engineering.

Switches that run Junos OS that supports the ELS configuration style perform inter-VLAN routing functions using an integrated routing and bridging (IRB) interface named irb, while switches that run Junos OS that does not support ELS perform these functions using a routed VLAN interface (RVI) named vlan. These interfaces detect both MAC addresses and IP addresses and route data to Layer 3 interfaces, thereby frequently eliminating the need to have both a switch and a router.

Related Documentation

- *Example: Setting Up Basic Bridging and a VLAN on Switches*
- *Understanding FCoE*
- *Interfaces Overview for Switches*
- *Understanding Multiple VLAN Registration Protocol (MVRP)*
- *Understanding Integrated Routing and Bridging*

Configuring VLANs on Switches with Enhanced Layer 2 Support

Switches use VLANs to make logical groupings of network nodes with their own broadcast domains. You can use VLANs to limit the traffic flowing across the entire LAN and reduce collisions and packet retransmissions.

NOTE: This task supports the Enhanced Layer 2 Software (ELS) configuration style. For ELS details, see *Using the Enhanced Layer 2 Software CLI*. If your switch runs software that does not support ELS, see *Configuring VLANs on Switches*.

NOTE: Starting with Junos OS Release 17.1R3, on QFX10000 switches, you cannot configure an interface with both **family ethernet-switching** and **flexible-vlan-tagging**. This configuration is not supported, and a warning will be issued if you try to commit this configuration.

NOTE: Two logical interfaces that are configured on the same physical interface cannot be mapped to the same VLAN.

For each endpoint on the VLAN, configure the following VLAN parameters on the corresponding interface:

1. Specify the description of the VLAN:

   ```
   [edit interfaces interface-name unit 0]
   user@switch# set description vlan-description
   ```

2. Specify the unique name of the VLAN:

   NOTE: Switches that run Junos OS with the ELS configuration style do not support a default VLAN. Therefore, on such switches, you must explicitly configure at least one VLAN, even if your network is simple and you want only one broadcast domain to exist.

   NOTE: On QFX5100 switches running Junos OS Release 14.1X53-D46 or earlier, when you configure an interface under a VLAN but do not specify the name of the VLAN, the system will not issue a commit error.

   ```
   [edit interfaces interface-name unit 0]
   user@switch# set family ethernet-switching vlan members vlan-name
   ```

3. Create the subnet for the VLAN:

 Copyright © 2018, Juniper Networks, Inc.

```
[edit interfaces]
user@switch# set vlan unit 0 family inet address ip-address
```

NOTE: The **family inet** option is not supported on NFX150 devices.

4. Configure the VLAN tag ID or VLAN ID list for the VLAN:

```
[edit vlans]
user@switch# set vlan-name vlan-id vlan-id-number
```

or

```
[edit vlans]
user@switch# set vlan-name vlan-id-list [vlan-ids | vlan-id--vlan-id-]
```

5. Specify a VLAN firewall filter to be applied to incoming or outgoing packets:

```
[edit vlans]
user@switch# set vlan-name filter (input | output) filter-name
```

Release History Table

Release	Description
17.1R3	Starting with Junos OS Release 17.1R3, on QFX10000 switches, you cannot configure an interface with both **family ethernet-switching** and **flexible-vlan-tagging**.

Related Documentation

- *Example: Setting Up Basic Bridging and a VLAN on Switches*
- *Configuring IRB Interfaces on Switches*
- Understanding Bridging and VLANs on Switches on page 9

Copyright © 2018, Juniper Networks, Inc.

Configuring the Native VLAN Identifier on Switches With ELS Support (CLI Procedure)

NOTE: This task uses Junos OS for EX Series switches and Junos OS for QFX3500 and QFX3600 switches with support for the Enhanced Layer 2 Software (ELS) configuration style. If your switch runs software that does not support ELS, see *Configuring the Native VLAN Identifier (CLI Procedure)*. For ELS details, see *Using the Enhanced Layer 2 Software CLI*.

Switches can receive and forward routed or bridged Ethernet frames with 802.1Q VLAN tags. Typically, trunk ports, which connect switches to each other, accept untagged control packets but do not accept untagged data packets. You can enable a trunk port to accept untagged data packets by configuring a native VLAN ID on the interface on which you want the untagged data packets to be received. The logical interface on which untagged packets are to be received must be configured with the same VLAN ID as the native VLAN ID configured on the physical interface.

To configure the native VLAN ID by using the command-line interface (CLI):

1. On the interface on which you want untagged data packets to be received, set the interface mode to **trunk**, which specifies that the interface is in multiple VLANs and can multiplex traffic between different VLANs.:

```
[edit interfaces]
user@switch# set interface-name unit logical-unit-number family
ethernet-switching interface-mode trunk
```

2. Configure the native VLAN ID:

```
[edit interfaces]
user@switch# set interface-name native-vlan-id vlan-id
```

3. Specify that the logical interface that will receive the untagged data packets is a member of the native VLAN:

```
[edit interfaces]
user@switch# set interface-name unit logical-unit-number family
ethernet-switching vlan members vlan-id
```

Related Documentation

-
- *Example: Connecting Access Switches with ELS Support to a Distribution Switch with ELS Support*
- *Example: Setting Up Basic Bridging and a VLAN for an EX Series Switch with ELS Support*
- *Example: Setting Up Basic Bridging and a VLAN on Switches*

Copyright © 2018, Juniper Networks, Inc.

Creating a Series of Tagged VLANs on NFX150 Devices

When you divide an Ethernet LAN into multiple VLANs, each VLAN is assigned a unique IEEE 802.1Q tag. This tag is associated with each frame in the VLAN, and the network nodes receiving the traffic can use the tag to identify which VLAN a frame is associated with.

Instead of configuring VLANs and 802.1Q tags one at a time for a trunk interface, you can configure a VLAN range to create a series of tagged VLANs.

When an Ethernet LAN is divided into VLANs, each VLAN is identified by a unique 802.1Q tag. The tag is applied to all frames so that the network nodes receiving the frames can detect which VLAN the frames belong to. Trunk ports, which multiplex traffic among a number of VLANs, use the tag to determine the origin of frames and where to forward them.

For example, you could configure the VLAN **employee** and specify a tag range of **10 through 12**. This creates the following VLANs and tags:

- VLAN **employee-10**, tag **10**
- VLAN **employee-11**, tag **11**
- VLAN **employee-12**, tag **12**

Creating tagged VLANs in a series has the following limitations:

- Layer 3 interfaces do not support this feature.
- Because an access interface can only support one VLAN member, access interfaces also do not support this feature.

Copyright © 2018, Juniper Networks, Inc.

To configure a series of tagged VLANs using the CLI (here, the VLAN is **employee**):

1. Configure the series (here, a VLAN series from 120 through 130):

   ```
   [edit]
   user@switch# set vlans employee vlan-id-list [ 120-130 ]
   ```

2. Associate a series of tagged VLANs when you configure an interface in one of two ways:

- Include the name of the series:

  ```
  [edit interfaces]
  user@switch# set interfaces ge-0/0/0 unit 0 family ethernet-switching vlan members employee
  ```

- Include the VLAN range:

  ```
  [edit interfaces]
  user@switch# set interfaces ge-0/0/0 unit 0 family ethernet-switching vlan members 120–130
  ```

Associating a series of tagged VLANS to an interface by name or by VLAN range the same result: VLANs **__employee_120__** through **__employee_130__** are created.

NOTE: When a series of VLANs is created using the vlan-id-list command, the VLAN names are preceded and followed by a double underscore.

Related Documentation

- Configuring VLANs on Switches with Enhanced Layer 2 Support on page 18

 Copyright © 2018, Juniper Networks, Inc.

Adding a Static MAC Address Entry to the Ethernet Switching Table on a Switch with ELS Support (CLI Procedure)

NOTE: This task uses Junos OS for EX Series switches and Junos OS for QFX3500 and QFX3600 switches with support for the Enhanced Layer 2 Software (ELS) configuration style. If your switch runs software that does not support ELS, see *Adding a Static MAC Address Entry to the Ethernet Switching Table (CLI Procedure)*. For ELS details, see *Using the Enhanced Layer 2 Software CLI*.

The Ethernet switching table, also known as the forwarding table, specifies the known locations of VLAN nodes and the addresses of devices within those nodes. There are two ways to populate the Ethernet switching table on a switch. The easiest method is to let the switch update the table with MAC addresses.

The second way to populate the Ethernet switching table is to manually insert addresses into the table. You can do this to reduce flooding and speed up the switch's automatic learning process.

Before configuring a static MAC address, be sure that you have:

- Set up the VLAN. See *Configuring VLANs for EX Series Switches with ELS Support (CLI Procedure)*.

To configure an interface to have a static MAC address:

```
[edit vlans vlan-name switch-options interface interface-name]
user@switch# set static-mac mac-address
```

Related Documentation

- Understanding Bridging and VLANs on Switches on page 9

Configuring Static ARP Entries

You can create static ARP table entries, which are explicit mappings between IP addresses and MAC addresses.

- To configure a static ARP entry:

  ```
  [edit interfaces interface-name unit logical-unit-number family inet address
  address]
  user@switch# set arp ip-address (mac | multicast-mac) mac-address
  ```

The IP address that you specify must be part of the subnet defined in the enclosing **address** statement.

To associate a multicast MAC address with a unicast IP address, use the **multicast-mac** statement.

Copyright © 2018, Juniper Networks, Inc.

Specify the MAC address as 6 hexadecimal bytes in one of the following formats: ***nnnn.nnnn.nnnn*** or ***nn:nn:nn:nn:nn:nn***; for example, **0011.2233.4455** or **00:11:22:33:44:55**.

Related Documentation

- *Understanding Static ARP Entries*
- *arp*

Troubleshooting Ethernet Switching

Problem

Description: Sometimes a MAC address entry in the switch's Ethernet switching table is not updated after the device with that MAC address has been moved from one interface to another on the switch. Typically, the switch does not wait for a MAC address expiration when a MAC move operation occurs. As soon as the switch detects the MAC address on the new interface, it immediately updates the table. Many network devices send a gratuitous ARP packet when switching an IP address from one device to another. The switch updates its ARP cache table after receipt of such gratuitous ARP messages, and then it also updates its Ethernet switching table.

Sometimes silent devices, such as syslog servers or SNMP trap receivers that receive UDP traffic but do not return acknowledgment (ACK) messages to the traffic source, fail to send gratuitous ARP packets when a device moves. If such a move occurs when the system administrator is not available to explicitly clear the affected interfaces by issuing the **clear ethernet-switching table** command, the entry for the moved device in the Ethernet switching table is not updated.

Solution

Set up the switch to handle unattended MAC address switchovers.

1. Reduce the system-wide ARP aging timer. (By default, the ARP aging timer is set at 20 minutes. The range of the ARP aging timer is from 1 through 240 minutes.)

   ```
   [edit system arp]
   user@switch# set aging-timer 3
   ```

2. Set the MAC aging timer to the same value as the ARP timer. (By default, the MAC aging timer is set to 300 seconds. The range is 60 to 1,000,000 seconds.)

   ```
   [edit protocols l2-learning]
   user@switch# set global-mac-table-aging-time 180
   ```

The ARP entry and the MAC address entry for the moved device expire within the times specified by the aging timer values. After the entries expire, the switch sends a new ARP message to the IP address of the device. The device responds to the ARP message, thereby refreshing the entries in the switch's ARP cache table and Ethernet switching table.

Related Documentation

- *arp*
- global-mac-table-aging-time on page 120

Copyright © 2018, Juniper Networks, Inc.

PART 2

MAC Addresses

Copyright © 2018, Juniper Networks, Inc.

Copyright © 2018, Juniper Networks, Inc.

CHAPTER 2

Using MAC Addresses

Introduction to the Media Access Control (MAC) Layer 2 Sublayer

This topic provides an introduction to the MAC sublayer of the data link layer (Layer 2).

In Layer 2 of a network, the Media Access Control (MAC) sublayer provides addressing and channel access control mechanisms that enable several terminals or network nodes to communicate in a network.

The MAC sublayer acts as an interface between the logical link control (LLC) Ethernet sublayer and Layer 1 (the physical layer). The MAC sublayer emulates a full-duplex logical communication channel in a multipoint network. This channel may provide unicast, multicast, or broadcast communication service. The MAC sublayer uses MAC protocols to prevent collisions.

In Layer 2, multiple devices on the same physical link can uniquely identify one another at the data link layer, by using the MAC addresses that are assigned to all ports on a switch. A MAC algorithm accepts as input a secret key and an arbitrary-length message to be authenticated, and outputs a MAC address.

A MAC address is a 12-digit hexadecimal number (48 bits in long). MAC addresses are usually written in one of these formats:

- MM:MM:MM:SS:SS:SS
- MM-MM-MM-SS-SS-SS

The first half of a MAC address contains the ID number of the adapter manufacturer. These IDs are regulated by an Internet standards body. The second half of a MAC address represents the serial number assigned to the adapter by the manufacturer.

Copyright © 2018, Juniper Networks, Inc.

Contrast MAC addressing, which works at Layer 2, with IP addressing, which runs at Layer 3 (networking and routing). One way to remember the difference is that the MAC addresses apply to a physical or virtual node, whereas IP addresses apply to the software implementation of that node. MAC addresses are typically fixed on a per-node basis, whereas IP addresses change when the node moves from one part of the network to another.

IP networks maintain a mapping between the IP and MAC addresses of a node using the Address Resolution Protocol (ARP) table. DHCP also typically uses MAC addresses when assigning IP addresses to nodes.

Related Documentation

- Overview of Layer 2 Networking on page 3
- Understanding MAC Learning on page 28

Understanding MAC Learning

MAC learning is the process of obtaining the MAC addresses of all the nodes on a network.

When a node is first connected to an Ethernet LAN or VLAN, it has no information about the other nodes on the network. As data is sent through the network, data packets include a data frame listing their source and destination MAC addresses. The data frame is forwarded to a target port, which is connected to the second device. The MAC address is learned locally at the target port, which facilitates communications for frames that later enter the target port and contain addresses previously learned from a received frame.

By default, MAC learning is enabled on the QFX and NFX Series.

Related Documentation

- Introduction to the Media Access Control (MAC) Layer 2 Sublayer on page 27
- Overview of Layer 2 Networking on page 3

Disabling MAC Learning on Devices with ELS Support

By default, MAC learning is globally enabled on all node. This topic describes how to disable MAC learning, as well as how to reenable and verify that MAC learning has been enabled or disabled.

NOTE: This task supports the Enhanced Layer 2 Software (ELS) configuration style. For ELS details, see *Using the Enhanced Layer 2 Software CLI* If your switch runs software that does not support ELS, see *Disabling MAC Learning on QFX Switches*.

Disabling dynamic MAC learning prevents a node from learning source and destination MAC addresses.

- To disable MAC learning:

 Copyright © 2018, Juniper Networks, Inc.

```
[edit vlans vlan-name switch-options interface interface-name]
user@switch# set no-mac-learning
```

- To enable MAC learning:

```
[edit vlans vlan-name switch-options interface interface-name]
user@switch# delete no-mac-learning
user@switch# deactivate no-mac-learning
```

- To verify the status of MAC learning, view the Ethernet MAC learning statistics in operational mode.

```
user@switch> show ethernet-switching table
```

```
Ethernet-switching table: 2 entries, 1 learned
  VLAN              MAC address       Type         Age Interfaces
  default           *                 Flood          - All-members
  default           00:1f:12:39:90:80 Learn         29 xe-/0/0.0
```

Related Documentation

- Understanding MAC Learning on page 28
- *Example: Disabling MAC Learning on Devices with ELS Support*
- *no-mac-learning*

Configuring MAC Notification on Switches with ELS Support (CLI Procedure)

NOTE: This task uses the Enhanced Layer 2 Software (ELS) configuration style. If your switch runs software that does not support ELS, see *Configuring Non-ELS MAC Notification* or *Configuring Non-ELS MAC Notification* . For ELS details, see *Using the Enhanced Layer 2 Software CLI*.

When a switch learns or unlearns a MAC address, SNMP notifications can be sent to the network management system at regular intervals to record the addition or removal of the MAC address. This process is known as MAC notification.

The MAC notification interval defines how often Simple Network Management Protocol (SNMP) notifications logging the addition or removal of MAC addresses on the switch are sent to the network management system.

MAC notification is disabled by default. When MAC notification is enabled, the default MAC notification interval is 30 seconds.

To enable or disable MAC notification, or to set the MAC notification interval, perform these tasks:

- Enabling MAC Notification on page 30
- Disabling MAC Notification on page 30
- Setting the MAC Notification Interval on page 30

Copyright © 2018, Juniper Networks, Inc.

Enabling MAC Notification

MAC notification is disabled by default. You need to perform this procedure to enable MAC notification.

To enable MAC notification on the switch with the default MAC notification interval of 30 seconds:

```
[edit switch-options]
user@switch# set mac-notification
```

To enable MAC notification on the switch with any other MAC notification interval (here, the MAC notification interval is set to 60 seconds):

```
[edit switch-options]
user@switch# set mac-notification notification-interval 60
```

Disabling MAC Notification

MAC notification is disabled by default. Perform this procedure only if MAC notification was previously enabled on your switch.

To disable MAC notification on the switch:

```
[edit switch-options]
user@switch# delete mac-notification
```

To disable MAC notification on a specific interface (here, the interface is ge-0/0/3):

```
[edit switch-options]
user@switch# set interface ge-0/0/3 no-mac-notification
```

Setting the MAC Notification Interval

The default MAC notification interval is 30 seconds. The procedure to change the MAC notification interval to a different interval is identical to the procedure to enable MAC notification on the switch with a nondefault value for the MAC notification interval.

To set the MAC notification interval on the switch (here, the MAC notification interval is set to 5 seconds):

```
[edit switch-options]
user@switch# set mac-notification notification-interval 5
```

Related Documentation

- Verifying That MAC Notification Is Working Properly on page 31

Copyright © 2018, Juniper Networks, Inc.

Verifying That MAC Notification Is Working Properly

Purpose Verify that MAC notification is enabled or disabled, and that the MAC notification interval is set to the specified value.

Action To verify that MAC notification is enabled or disabled on a QFX Series switch or an EX4600, and also to verify the MAC notification interval setting:

`user@switch>` **show ethernet-switching mac-notification**

```
Notification Status: Enabled
Notification Interval: 60
Notifications Sent          : 0
Notifications Table Maxsize : 256
```

The output in the **Notification Status** field shows that MAC notification is enabled. The output in the **Notification Status** field would display **Disabled** if MAC notification was disabled.

The **Notification Interval** field output shows that the MAC notification interval is set to 60 seconds.

To verify that MAC notification is enabled on an EX Series switch while also verifying the MAC notification interval setting:

`user@switch>` **show ethernet-switching mac-notification**

```
Notification Status: Enabled
Notification Interval: 30
```

The output in the **Notification Status** field shows that MAC notification is enabled. The output in the **Notification Status** field would display **Disabled** if MAC notification was disabled.

The **Notification Interval** field output shows that the MAC notification interval is set to 30 seconds.

Related Documentation

- *Configuring Non-ELS MAC Notification*
- Configuring MAC Notification on Switches with ELS Support (CLI Procedure) on page 29

Copyright © 2018, Juniper Networks, Inc.

Configuring MAC Limiting (CLI Procedure)

NOTE: This task uses Junos OS for EX Series switches and QFX3500 and QFX3600 switches with support for the Enhanced Layer 2 Software (ELS) configuration style. If your switch runs software that does not support ELS, see *Configuring MAC Limiting (CLI Procedure)*. For ELS details, see *Using the Enhanced Layer 2 Software CLI*.

This topic describes various ways of configuring a limitation on MAC addresses in packets that are received and forwarded by the switch.

For information on configuring an interface to automatically recover from a shutdown caused by MAC limiting, see *Configuring Autorecovery from the Disabled State on Secure or Storm Control Interfaces (CLI Procedure)*. If you do not configure the switch for autorecovery from the disabled condition, you can bring up the disabled interfaces by running the **clear ethernet-switching recovery-timeout** command.

The different ways of setting a MAC limit are described in the following sections:

Limiting the Number of MAC Addresses Learned by an Interface

To secure a port, you can set the maximum number of MAC addresses that can be learned by an interface:

- Set the MAC limit on an interface, and specify an action that the switch takes after the specified limit is exceeded:

  ```
  [edit switch-options]
  user@switch# set interface interface-name interface-mac-limit limit packet-action
  action
  ```

  After you set a new MAC limit for the interface, the system clears existing entries in the MAC address forwarding table associated with the interface.

Limiting the Number of MAC Addresses Learned by a VLAN

To limit the number of MAC addresses learned by a VLAN, perform both of the following steps:

1. Set the maximum number of MAC addresses that can be learned by a VLAN, and specify an action that the switch takes after the specified limit is exceeded:

   ```
   [edit vlans]
   user@switch# set vlan-name switch-options mac-table-size limit packet-action
   action
   ```


Copyright © 2018, Juniper Networks, Inc.

2. Set the maximum number of MAC addresses that can be learned by one *or* all interfaces in the VLAN, and specify an action that the switch takes after the specified limit is exceeded:

NOTE: If you specify a MAC limit and packet action for all interfaces in the VLAN *and* a specific interface in the VLAN, the MAC limit and packet action specified at the specific interface level takes precedence. Also, at the VLAN interface level, only the **drop** and **drop-and-log** options are supported.

```
[edit vlans]
user@switch# set vlan-name switch-options interface interface-name
interface-mac-limit limit packet-action action
```

```
[edit vlans]
user@switch# set vlan-name switch-options interface-mac-limit limit packet-action
action
```

After you set new MAC limits for a VLAN by using the **mac-table-size** statement or for interfaces associated with a VLAN by using the **interface-mac-limit** statement, the system clears the corresponding existing entries in the MAC address forwarding table.

NOTE: On a QFX Series Virtual Chassis, if you include the **shutdown** option at the **[edit vlans *vlan-name* switch-options interface *interface-name* interface-mac-limit packet-action]** hierarchy level and issue the **commit** operation, the system generates a commit error. The system does not generate an error if you include the **shutdown** option at the **[edit switch-options interface *interface-name* interface-mac-limit packet-action]** hierarchy level.

Related Documentation

- *Configuring Autorecovery from the Disabled State on Secure or Storm Control Interfaces (CLI Procedure)*
- *Configuring Persistent MAC Learning (CLI Procedure)*

 Copyright © 2018, Juniper Networks, Inc.

PART 3

Spanning Trees

 Copyright © 2018, Juniper Networks, Inc.

CHAPTER 3

Using Spanning Trees

Overview of Spanning-Tree Protocols

QFX Series switches provide Layer 2 loop prevention through Spanning Tree Protocol (STP), Rapid Spanning Tree Protocol (RSTP), Multiple Spanning Tree Protocol (MSTP), and VLAN Spanning Tree Protocol (VSTP). The default spanning-tree protocol on the QFX Series is RSTP. RSTP provides faster convergence times than STP. However, some legacy networks require the slower convergence times of basic STP.

NOTE: **If you are configuring an interface for any spanning tree protocol (STP, MSTP, RSTP, and VSTP), the interface all, vlan all, and vlan-group options are not available when you configure an interface with the flexible-vlan-tagging family option.**

The STP support provided for the QFX Series includes:

- IEEE 802.1d
- 802.1w RSTP
- 802.1s MSTP

If your network includes IEEE 802.1D 1998 bridges, you can remove RSTP and explicitly configure STP. When you explicitly configure STP, the QFX Series products use the IEEE 802.1D 2004 specification, force version 0. This configuration runs a version of RSTP that is compatible with the classic, basic STP. If you use virtual LANs (VLANs), you should enable VSTP and use it on your network. See "Understanding VSTP" on page 60.

You can use the same operational commands (**show spanning-tree bridge** and **show spanning-tree interface**) to check the status of your spanning-tree configuration, regardless of which spanning-tree protocol has been configured.

STP uses bridge protocol data unit (BPDU) packets to exchange information with other switches. BPDUs send hello packets out at regular intervals to exchange information across bridges and detect loops in a network topology. There are two types of BPDUs:

- Configuration BPDUs—These BPDUs contain configuration information about the transmitting switch and its ports, including switch and port MAC addresses, switch priority, port priority, and port cost.
- Topology change notification (TCN) BPDUs—When a bridge needs to signal a topology change, it starts to send TCNs on its root port. The designated bridge receives the TCN, acknowledges it, and generates another one for its own root port. The process continues until the TCN reaches the root bridge.

STP uses the information provided by the BPDUs to elect a root bridge, identify root ports for each switch, identify designated ports for each physical LAN segment, and prune specific redundant links to create a loop-free tree topology. All leaf devices calculate the best path to the root device and place their ports in blocking or forwarding states based on the best path to the root. The resulting tree topology provides a single active Layer 2 data path between any two end stations.

 Copyright © 2018, Juniper Networks, Inc.

Understanding Spanning Tree Protocols on a QFabric System

Although there is no need to run STP in a QFabric system, you can connect a QFabric system to another Layer 2 device and use STP. STP traffic can only be processed on network Node groups. Other Node groups, such redundant server Node groups and server Node groups, discard the STP bridge protocol data units (BPDUs) traffic and disable the interface automatically. Server Node groups only process host-facing protocols, whereas Network Node groups process all supported protocols.

Related Documentation

- *Understanding BPDU Protection for STP, RSTP, and MSTP*
- *Understanding MSTP*
- Understanding RSTP on page 39
- Understanding VSTP on page 60

Understanding RSTP

Juniper Networks QFX Series products use Rapid Spanning Tree Protocol (RSTP) on the network side of the QFX Series to provide quicker convergence time than the base Spanning Tree Protocol (STP) does. RSTP identifies certain links as point to point. When a point-to-point link fails, the alternate link can transition to the forwarding state, which speeds up convergence.

Although STP provides basic loop prevention functionality, it does not provide fast network convergence when there are topology changes. The STP process to determine network state transitions is slower than the RSTP process because it is timer-based. A device must reinitialize every time a topology change occurs. The device must start in the listening state and transition to the learning state and eventually to a forwarding or blocking state. When default values are used for the maximum age (20 seconds) and forward delay (15 seconds), it takes 50 seconds for the device to converge. RSTP converges faster because it uses a handshake mechanism based on point-to-point links instead of the timer-based process used by STP.

For networks with virtual LANs (VLANs), you can use VLAN Spanning Tree Protocol (VSTP), which takes the paths of each VLAN into account when calculating routes. VSTP uses RSTP by default.

An RSTP domain running from the edge outward on a QFX Series product has the following components:

- A *root port*, which is the “best path” to the root device.
- A *designated port*, which indicates that the switch is the designated bridge for the other switch connecting to this port.
- An *alternate port*, which provides an alternate root port.
- A *backup port*, which provides an alternate designated port.

Copyright © 2018, Juniper Networks, Inc.

Port assignments change through messages exchanged throughout the domain. An RSTP device generates configuration messages once per hello time interval. If an RSTP device does not receive a configuration message from its neighbor after an interval of three hello times, it determines that the connection with the neighbor is lost. When a *root port* or a *designated port* fails on a device, the device generates a configuration message with the proposal bit set. Once its neighbor device receives this message, it verifies that this configuration message is valid for that port and starts a *synchronizing* operation to ensure that all of its ports are in sync with the new information.

Similar sets of messages propagate through the network, restoring the connectivity very quickly after a topology change (in a well-designed network that uses RSTP, network convergence can take as little as 0.5 seconds). If a device does not receive an agreement to a proposal message it has sent, it returns to the original IEEE 802.D convention.

RSTP was originally defined in the IEEE 802.1w draft specification and later incorporated into the IEEE 802.1D-2004 specification.

VSTP and RSTP can be configured at the same time. If you configure VSTP and RSTP at the same time and the switch has more than 253 VLANs, VSTP is configured only for the first 253 VLANs. For the remaining VLANs, only RSTP is configured. RSTP and VSTP are the only spanning-tree protocols that can be configured at the same time on the QFX Series.

NOTE: **Using the same VLAN for RSTP and VSTP is not supported. For example, if you are configuring a VLAN under VSTP, configuring RSTP with an interface that contains the same VLAN is not supported.**

Related Documentation

- Overview of Spanning-Tree Protocols on page 38
- *Understanding MSTP*
- Understanding VSTP on page 60
- *Example: Configuring Faster Convergence and Improving Network Stability with RSTP*
- *Example: Configuring Faster Convergence and Network Stability on ELS Switches with RSTP*
- *Configuring RSTP on EX Series Switches (CLI Procedure)*

Example: Configuring Faster Convergence and Network Stability on ELS Switches with RSTP

NOTE: This example uses Junos OS for EX Series switches with support for the Enhanced Layer 2 Software (ELS) configuration style. If your switch runs software that does not support ELS, see *Example: Faster Convergence and Improved Network Stability with RSTP on EX Series Switches*. For ELS details, see *Using the Enhanced Layer 2 Software CLI*.

Copyright © 2018, Juniper Networks, Inc.

EX Series switches use Rapid Spanning Tree Protocol (RSTP) by default to provide a loop-free topology.

When switches that support redundant Routing Engines use RSTP, it is important to keep RSTP synchronized on both Routing Engines so that no loss of service occurs after a Routing Engine switchover. Nonstop bridging protocol keeps Routing Engines synchronized.

This example describes how to configure RSTP and NSB on four EX Series switches:

Requirements

This example uses the following software and hardware components:

- Junos OS Release 15.1 or later or later for EX Series switches
- Four EX Series switches

Before you configure the switches for RSTP, be sure you have:

- Installed and connected the four switches. See the hardware documentation for your switch.
- Performed the initial software configuration on all switches. See *Connecting and Configuring an EX Series Switch (CLI Procedure)*.

Overview and Topology

RSTP works by identifying certain links as point to point links and blocking other possible paths. When one of the point-to-point links fails, a designated alternate link transitions to the forwarding state and take over. Configuring nonstop bridging (NSB) on a switch with redundant Routing Engines keeps RSTP synchronized on both Routing Engines. This way, RSTP remains active immediately after a switchover because it is already synchronized to the backup Routing Engine. RSTP does not have to reconverge after a Routing Engine switchover when NSB is enabled because the neighbor devices do not detect an RSTP change on the switch. In this example, four EX Series switches are connected in the topology displayed in Figure 1 on page 42 to create a loop-free topology with NSB applied to switches with dual Routing Engines.

Figure 1: Network Topology for RSTP

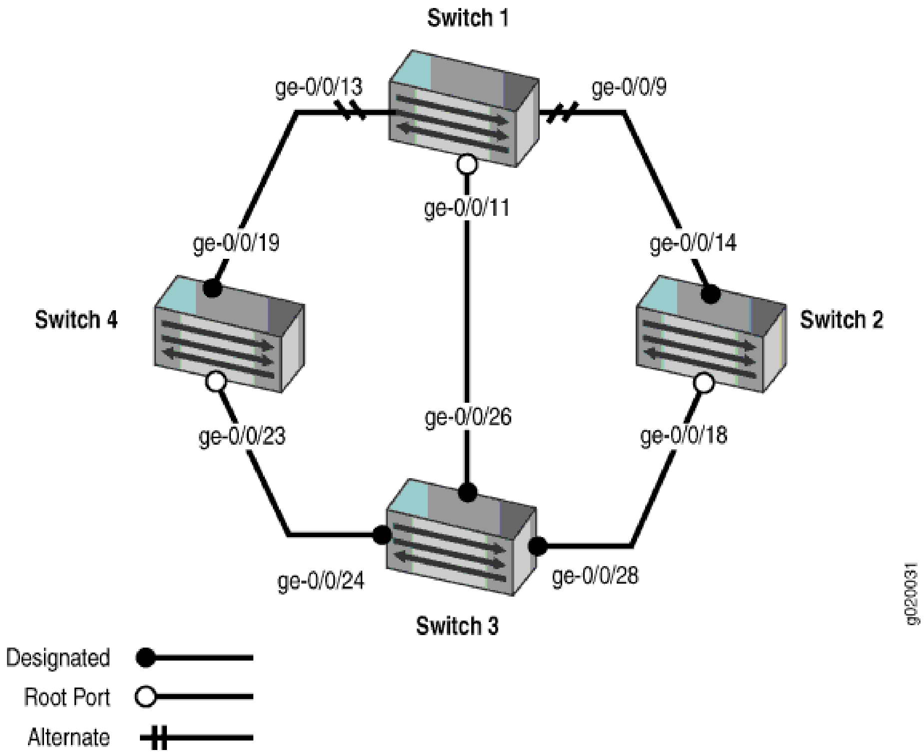

Table 3 on page 42 shows the components of the topology for this example.

NOTE: You can configure RSTP only on physical interfaces, not on logical interfaces.

Table 3: Components of the Topology for Configuring RSTP

Property	Settings
Switch 1	The following interfaces on Switch 1 are connected in this way: • **ge-0/0/9** is connected to Switch 2 • **ge-0/0/13** is connected to Switch 4 • **ge-0/0/11** is connected to Switch 3
Switch 2	The following interfaces on Switch 2 are connected in this way: • **ge-0/0/14** is connected to Switch 1 • **ge-0/0/18** is connected to Switch 3
Switch 3	The following interfaces on Switch 3 are connected in this way: • **ge-0/0/26** is connected to Switch 1 • **ge-0/0/28** is connected to Switch 2 • **ge-0/0/24** is connected to Switch 4

 Copyright © 2018, Juniper Networks, Inc.

Table 3: Components of the Topology for Configuring RSTP (continued)

Property	Settings
Switch 4	The following interfaces on Switch 4 are connected in this way: • **ge-0/0/19** is connected to Switch 1 • **ge-0/0/23** is connected to Switch 3
VLAN names and tag IDs	**voice-vlan**, tag **10** **employee-vlan**, tag **20** **guest-vlan**, tag **30** **camera-vlan**, tag **40**

This configuration example creates a loop-free topology between four EX Series switches using RSTP.

An RSTP topology contains ports that have specific roles:

- The *root port* is responsible for forwarding data to the root bridge.
- The *alternate port* is a standby port for the root port. When a root port goes down, the alternate port becomes the active root port.
- The *designated port* forwards data to the downstream network segment or device.
- The *backup port* is a backup port for the designated port. When a designated port goes down, the backup port becomes the active designated port and starts forwarding data.

NOTE: You also can create a loop-free topology between the aggregation layer and the distribution layer using redundant trunk links. For more information about configuring redundant trunk links, see *Example: Configuring Redundant Trunk Links for Faster Recovery on Devices with ELS Support*.

Configuring RSTP and Nonstop Bridging on Switch 1

CLI Quick Configuration

To quickly configure RSTP and nonstop bridging on Switch 1, copy the following commands and paste them into the switch terminal window:

```
[edit]
set vlans voice-vlan description “Voice VLAN”
set vlans voice-vlan vlan-id 10
set vlans employee-vlan description “Employee VLAN”
set vlans employee-vlan vlan-id 20
set vlans guest-vlan description “Guest VLAN”
set vlans guest-vlan vlan-id 30
set vlans camera-vlan description “Camera VLAN”
set vlans camera-vlan vlan-id 40
set interfaces ge-0/0/13 unit 0 family ethernet-switching vlan members [10 20 30 40]
set interfaces ge-0/0/9 unit 0 family ethernet-switching vlan members [10 20 30 40]
set interfaces ge-0/0/11 unit 0 family ethernet-switching vlan members [10 20 30 40]
set interfaces ge-0/0/13 unit 0 family ethernet-switching interface-mode trunk
set interfaces ge-0/0/9 unit 0 family ethernet-switching interface-mode trunk
```

Copyright © 2018, Juniper Networks, Inc.

set interfaces ge-0/0/11 unit 0 family ethernet-switching interface-mode trunk
set protocols rstp bridge-priority 16k
set protocols rstp interface all cost 1000
set protocols rstp interface all mode point-to-point

If Switch 1 includes dual Routing Engines, configure NSB. To quickly configure nonstop bridging on Switch 1, copy the following commands and paste them into the switch terminal window:

set chassis redundancy graceful switchover
set system commit synchronize
set protocols layer2-control nonstop-bridging

NOTE: NFX150 devices support only a single Routing Engine.

Step-by-Step Procedure

To configure RSTP and nonstop bridging on Switch 1:

1. Configure the VLANs **voice-vlan**, **employee-vlan**, **guest-vlan**, and **camera-vlan**:

```
[edit vlans]
user@switch1# set voice-vlan description "Voice VLAN"
user@switch1# set voice-vlan vlan-id 10
user@switch1# set employee-vlan description "Employee VLAN"
user@switch1# set employee-vlan vlan-id 20
user@switch1# set guest-vlan description "Guest VLAN"
user@switch1# set guest-vlan vlan-id 30
user@switch1# set camera-vlan description "Camera VLAN"
user@switch1# set camera-vlan vlan-id 40
```

2. Configure the VLANs on the interfaces, including support for the Ethernet switching protocol:

```
[edit interfaces]
user@switch1# set ge-0/0/13 unit 0 family ethernet-switching vlan members [10 20 30 40]
user@switch1# set ge-0/0/9 unit 0 family ethernet-switching vlan members [10 20 30 40]
user@switch1# set ge-0/0/11 unit 0 family ethernet-switching vlan members [10 20 30 40]
```

3. Configure the port mode for the interfaces:

```
[edit interfaces]
user@switch1# set ge-0/0/13 unit 0 family ethernet-switching interface-mode trunk
user@switch1# set ge-0/0/9 unit 0 family ethernet-switching interface-mode trunk
user@switch1# set ge-0/0/11 unit 0 family ethernet-switching interface-mode trunk
```

4. Configure RSTP on the switch:

```
[edit protocols]
user@switch1# rstp bridge-priority 16k
user@switch1# rstp interface all cost 1000
user@switch1# rstp interface all mode point-to-point
```


Copyright © 2018, Juniper Networks, Inc.

Step-by-Step Procedure

If Switch 1 includes dual Routing Engines, configure nonstop bridging. To configure NSB on Switch 1:

1. Enable graceful Routing Engine switchover (GRES):

   ```
   [edit chassis redundancy]
   user@switch1# set graceful-switchover
   ```

2. Configure the switch to always synchronize configuration changes between the Routing Engines:

   ```
   [edit system]
   user@switch1# set commit synchronize
   ```

   If you try to commit a configuration in which nonstop bridging is configured but synchronization of configuration changes is not configured, the configuration is not committed.

3. Enable nonstop bridging:

   ```
   [edit protocols layer2-control]
   user@switch1# set nonstop-bridging
   ```

   NOTE: This process enables NSB for all NSB-supported Layer 2 protocols on the switch, including RSTP.

Results

Check the results of the configuration:

```
user@switch1> show configuration
interfaces {
  ge-0/0/13 {
    unit 0 {
      family ethernet-switching {
        interface-mode trunk;
        vlan {
          members [10 20 30 40];
        }
      }
    }
  }
  ge-0/0/9 {
    unit 0 {
      family ethernet-switching {
        interface-mode trunk;
        vlan {
          members [10 20 30 40];
        }
      }
```

Copyright © 2018, Juniper Networks, Inc.

```
      }
    }
    ge-0/0/11 {
      unit 0 {
        family ethernet-switching {
          interface-mode trunk;
          vlan {
            members [10 20 30 40];
          }
        }
      }
    }
}
protocols {
    layer2-control {
      nonstop-bridging;
    }
    rstp {
      bridge-priority 16k;
        interface ge-0/0/13 {
          cost 1000;
          mode point-to-point;
        }
        interface ge-0/0/9 {
          cost 1000;
          mode point-to-point;
        }
        interface ge-0/0/11 {
          cost 1000;
          mode point-to-point;
        }
      }
    }
}
vlans {
    voice-vlan {
      vlan-id 10;
    }
    employee-vlan {
      vlan-id 20;
    }
    guest-vlan {
      vlan-id 30;
    }
    camera-vlan {
      vlan-id 40;
    }
}
system {
    commit synchronize;
}
chassis {
    redundancy {
      graceful-switchover;
    }
```


Copyright © 2018, Juniper Networks, Inc.

Configuring RSTP and Nonstop Bridging on Switch 2

CLI Quick Configuration

To quickly configure RSTP and nonstop bridging on Switch 2, copy the following commands and paste them into the switch terminal window:

```
[edit]
set vlans voice-vlan description “Voice VLAN”
set vlans voice-vlan vlan-id 10
set vlans employee-vlan description “Employee VLAN”
set vlans employee-vlan vlan-id 20
set vlans guest-vlan description “Guest VLAN”
set vlans guest-vlan vlan-id 30
set vlans camera-vlan description “Camera VLAN”
set vlans camera-vlan vlan-id 40
set interfaces ge-0/0/14 unit 0 family ethernet-switching vlan members [10 20 30 40]
set interfaces ge-0/0/18 unit 0 family ethernet-switching vlan members [10 20 30 40]
set interfaces ge-0/0/14 unit 0 family ethernet-switching interface-mode trunk
set interfaces ge-0/0/18 unit 0 family ethernet-switching interface-mode trunk
set protocols rstp bridge-priority 32k
set protocols rstp interface ge-0/0/14 cost 1000
set protocols rstp interface ge-0/0/14 mode point-to-point
set protocols rstp interface ge-0/0/18 cost 1000
set protocols rstp interface ge-0/0/18 mode point-to-point
```

NOTE: Starting with Junos OS Release 15.1 for EX Series and QFX Series switches with support for the Enhanced Layer 2 Software (ELS) configuration style, you can configure spanning tree parameters globally on all spanning tree interfaces. See *Configuring RSTP on EX Series Switches (CLI Procedure)* for additional information.

If Switch 2 includes dual Routing Engines, configure NSB. To quickly configure nonstop bridging on Switch 2, copy the following commands and paste them into the switch terminal window:

```
set chassis redundancy graceful switchover
set system commit synchronize
set protocols layer2-control nonstop-bridging
```

Step-by-Step Procedure

To configure RSTP and nonstop bridging on Switch 2:

1. Configure the VLANs **voice-vlan**, **employee-vlan**, **guest-vlan**, and **camera-vlan**:

```
[edit vlans]
user@switch2# set voice-vlan description “Voice VLAN”
user@switch2# set voice-vlan vlan-id 10
user@switch2# set employee-vlan description “Employee VLAN”
user@switch2# set employee-vlan vlan-id 20
user@switch2# set guest-vlan description “Guest VLAN”
user@switch2# set guest-vlan vlan-id 30
user@switch2# set camera-vlan vlan-description “Camera VLAN”
user@switch2# set camera-vlan vlan-id 40
```

Copyright © 2018, Juniper Networks, Inc.

2. Configure the VLANs on the interfaces, including support for the Ethernet switching protocol:

```
[edit interfaces]
user@switch2# set ge-0/0/14 unit 0 family ethernet-switching vlan members [10 20 30 40]
user@switch2# set ge-0/0/18 unit 0 family ethernet-switching vlan members [10 20 30 40]
```

3. Configure the port mode for the interfaces:

```
[edit interfaces]
user@switch2# set ge-0/0/14 unit 0 family ethernet-switching interface-mode trunk
user@switch2# set ge-0/0/18 unit 0 family ethernet-switching interface-mode trunk
```

4. Configure RSTP on the switch:

```
[edit protocols]
user@switch2# rstp bridge-priority 32k
user@switch2# rstp interface ge-0/0/14 cost 1000
user@switch2# rstp interface ge-0/0/14 mode point-to-point
user@switch2# rstp interface ge-0/0/18 cost 1000
user@switch2# rstp interface ge-0/0/18 mode point-to-point
```

Step-by-Step Procedure

If Switch 2 includes dual Routing Engines, configure nonstop bridging. To configure NSB on Switch 2:

1. Enable graceful Routing Engine switchover (GRES):

```
[edit chassis redundancy]
user@switch2# set graceful-switchover
```

2. Configure the switch to always synchronize configuration changes between the Routing Engines:

```
[edit system]
user@switch2# set commit synchronize
```

If you try to commit a configuration in which nonstop bridging is configured but synchronization of configuration changes is not configured, the configuration is not committed.

3. Enable nonstop bridging:

```
[edit protocols layer2-control]
user@switch2# set nonstop-bridging
```

NOTE: This process enables NSB for all NSB-supported Layer 2 protocols on the switch, including RSTP.

 Copyright © 2018, Juniper Networks, Inc.

Results Check the results of the configuration:

```
user@switch2> show configuration
interfaces {
   ge-0/0/14 {
      unit 0 {
         family ethernet-switching {
            interface-mode trunk;
            vlan {
               members [10 20 30 40];
            }
         }
      }
   }
   ge-0/0/18 {
      unit 0 {
         family ethernet-switching {
            interface-mode trunk;
            vlan {
               members [10 20 30 40];
            }
         }
      }
   }
}
protocols {
   layer2-control {
      nonstop-bridging;
   }
   rstp {
      bridge-priority 32k;
         interface ge-0/0/14 {
            cost 1000;
            mode point-to-point;
         }
         interface ge-0/0/18 {
            cost 1000;
            mode point-to-point;
         }
      }
   }
}
vlans {
   voice-vlan {
      vlan-id 10;
   }
   employee-vlan {
      vlan-id 20;
   }
   guest-vlan {
      vlan-id 30;
   }
   camera-vlan {
      vlan-id 40;
   }
```

Copyright © 2018, Juniper Networks, Inc.

```
}
system {
   commit synchronize;
}
chassis {
   redundancy {
      graceful-switchover;
   }
```

Configuring RSTP and Nonstop Bridging on Switch 3

CLI Quick Configuration

To quickly configure RSTP and nonstop bridging on Switch 3, copy the following commands and paste them into the switch terminal window:

```
[edit]
set vlans voice-vlan description “Voice VLAN”
set vlans voice-vlan vlan-id 10
set vlans employee-vlan description “Employee VLAN”
set vlans employee-vlan vlan-id 20
set vlans guest-vlan description “Guest VLAN”
set vlans guest-vlan vlan-id 30
set vlans camera-vlan description “Camera VLAN”
set vlans camera-vlan vlan-id 40
set interfaces ge-0/0/26 unit 0 family ethernet-switching vlan members [10 20 30 40]
set interfaces ge-0/0/28 unit 0 family ethernet-switching vlan members [10 20 30 40]
set interfaces ge-0/0/24 unit 0 family ethernet-switching vlan members [10 20 30 40]
set interfaces ge-0/0/26 unit 0 family ethernet-switching interface-mode trunk
set interfaces ge-0/0/28 unit 0 family ethernet-switching interface-mode trunk
set interfaces ge-0/0/24 unit 0 family ethernet-switching interface-mode trunk
set protocols rstp bridge-priority 8k
set protocols rstp interface ge-0/0/26 cost 1000
set protocols rstp interface ge-0/0/26 mode point-to-point
set protocols rstp interface ge-0/0/28 cost 1000
set protocols rstp interface ge-0/0/28 mode point-to-point
set protocols rstp interface ge-0/0/24 cost 1000
set protocols rstp interface ge-0/0/24 mode point-to-point
```

If Switch 3 includes dual Routing Engines, configure NSB. To quickly configure nonstop bridging on Switch 3, copy the following commands and paste them into the switch terminal window:

```
set chassis redundancy graceful switchover
set system commit synchronize
set protocols layer2-control nonstop-bridging
```


Copyright © 2018, Juniper Networks, Inc.

Step-by-Step Procedure

To configure RSTP and nonstop bridging on Switch 3:

1. Configure the VLANs **voice-vlan**, **employee-vlan**, **guest-vlan**, and **camera-vlan**:

```
[edit vlans]
user@switch3# set voice-vlan description "Voice VLAN"
user@switch3# set voice-vlan vlan-id 10
user@switch3# set employee-vlan description "Employee VLAN"
user@switch3# set employee-vlan vlan-id 20
user@switch3# set guest-vlan description "Guest VLAN"
user@switch3# set guest-vlan vlan-id 30
user@switch3# set camera-vlan description "Camera VLAN"
user@switch3# set camera-vlan vlan-id 40
```

2. Configure the VLANs on the interfaces, including support for the Ethernet switching protocol:

```
[edit interfaces]
user@switch3# set ge-0/0/26 unit 0 family ethernet-switching vlan members [10 20 30 40]
user@switch3# set ge-0/0/28 unit 0 family ethernet-switching vlan members [10 20 30 40]
user@switch3# set ge-0/0/24 unit 0 family ethernet-switching vlan members [10 20 30 40]
```

3. Configure the port mode for the interfaces:

```
[edit interfaces]
user@switch3# set ge-0/0/26 unit 0 family ethernet-switching interface-mode trunk
user@switch3# set ge-0/0/28 unit 0 family ethernet-switching interface-mode trunk
user@switch3# set ge-0/0/24 unit 0 family ethernet-switching interface-mode trunk
```

4. Configure RSTP on the switch:

```
[edit protocols]
user@switch3# rstp bridge-priority 8k
user@switch3# rstp interface ge-0/0/26 cost 1000
user@switch3# rstp interface ge-0/0/26 mode point-to-point
user@switch3# rstp interface ge-0/0/28 cost 1000
user@switch3# rstp interface ge-0/0/28 mode point-to-point
user@switch3# rstp interface ge-0/0/24 cost 1000
user@switch3# rstp interface ge-0/0/24 mode point-to-point
```

Step-by-Step Procedure

If Switch 3 includes dual Routing Engines, configure nonstop bridging. To configure NSB on Switch 3:

1. Enable graceful Routing Engine switchover (GRES):

```
[edit chassis redundancy]
user@switch3# set graceful-switchover
```

2. Configure the switch to always synchronize configuration changes between the Routing Engines:

```
[edit system]
```

```
user@switch3# set commit synchronize
```

If you try to commit a configuration in which nonstop bridging is configured but synchronization of configuration changes is not configured, the configuration is not committed.

3. Enable nonstop bridging:

```
[edit protocols layer2-control]
user@switch3# set nonstop-bridging
```

NOTE: This process enables NSB for all NSB-supported Layer 2 protocols on the switch, including RSTP.

Results Check the results of the configuration:

```
user@switch3> show configuration
  interfaces {
    ge-0/0/26 {
      unit 0 {
        family ethernet-switching {
          interface-mode trunk;
          vlan {
            members [10 20 30 40];
          }
        }
      }
    }
    ge-0/0/28 {
      unit 0 {
        family ethernet-switching {
          interface-mode trunk;
          vlan {
            members [10 20 30 40];
          }
        }
      }
    }
    ge-0/0/24 {
      unit 0 {
        family ethernet-switching {
          interface-mode trunk;
          vlan {
            members [10 20 30 40];
          }
        }
      }
    }
  }
```

 Copyright © 2018, Juniper Networks, Inc.

```
}
protocols {
   layer2-control {
      nonstop-bridging;
   }
   rstp {
      bridge-priority 8k;
         interface ge-0/0/26 {
            cost 1000;
            mode point-to-point;
         }
         interface ge-0/0/28 {
            cost 1000;
            mode point-to-point;
         }
         interface ge-0/0/24 {
            cost 1000;
            mode point-to-point;
         }
      }
      bridge-priority 8k;
         }
      }
   }
}
vlans {
   voice-vlan {
      vlan-id 10;
   }
   employee-vlan {
      vlan-id 20;
   }
   guest-vlan {
      vlan-id 30;
   }
   camera-vlan {
      vlan-id 40;
   }
}
system {
   commit synchronize;
}
chassis {
   redundancy {
      graceful-switchover;
   }
```

Configuring RSTP and Nonstop Bridging on Switch 4

CLI Quick Configuration

To quickly configure RSTP and nonstop bridging on Switch 4, copy the following commands and paste them into the switch terminal window:

```
[edit]
set vlans voice-vlan description “Voice VLAN”
```

Copyright © 2018, Juniper Networks, Inc.

```
set vlans voice-vlan vlan-id 10
set vlans employee-vlan description “Employee VLAN”
set vlans employee-vlan vlan-id 20
set vlans guest-vlan description “Guest VLAN”
set vlans guest-vlan vlan-id 30
set vlans camera-vlan description “Camera VLAN”
set vlans camera-vlan vlan-id 40
set interfaces ge-0/0/23 unit 0 family ethernet-switching vlan members [10 20 30 40]
set interfaces ge-0/0/19 unit 0 family ethernet-switching vlan members [10 20 30 40]
set interfaces ge-0/0/23 unit 0 family ethernet-switching interface-mode trunk
set interfaces ge-0/0/19 unit 0 family ethernet-switching interface-mode trunk
set protocols rstp bridge-priority 16k
set protocols rstp interface ge-0/0/23 cost 1000
set protocols rstp interface ge-0/0/23 mode point-to-point
set protocols rstp interface ge-0/0/19 cost 1000
set protocols rstp interface ge-0/0/19 mode point-to-point
```

If Switch 4 includes dual Routing Engines, configure NSB. To quickly configure nonstop bridging on Switch 4, copy the following commands and paste them into the switch terminal window:

```
set chassis redundancy graceful switchover
set system commit synchronize
set protocols layer2-control nonstop-bridging
```

Step-by-Step Procedure

To configure RSTP and nonstop bridging on Switch 4:

1. Configure the VLANs **voice-vlan**, **employee-vlan**, **guest-vlan**, and **camera-vlan**:

```
[edit vlans]
user@switch4# set voice-vlan description “Voice VLAN”
user@switch4# set voice-vlan vlan-id 10
user@switch4# set employee-vlan description “Employee VLAN”
user@switch4# set employee-vlan vlan-id 20
user@switch4# set guest-vlan description “Guest VLAN”
user@switch4# set guest-vlan vlan-id 30
user@switch4# set camera-vlan description “Camera VLAN”
user@switch4# set camera-vlan vlan-id 40
```

2. Configure the VLANs on the interfaces, including support for the Ethernet switching protocol:

```
[edit interfaces]
user@switch4# set ge-0/0/23 unit 0 family ethernet-switching vlan members [10 20 30
40]
user@switch4# set ge-0/0/19 unit 0 family ethernet-switching vlan members [10 20 30
40]
```

3. Configure the port mode for the interfaces:

```
[edit interfaces]
user@switch4# set ge-0/0/23 unit 0 family ethernet-switching interface-mode trunk
user@switch4# set ge-0/0/19 unit 0 family ethernet-switching interface-mode trunk
```

4. Configure RSTP on the switch:

 Copyright © 2018, Juniper Networks, Inc.

```
[edit protocols]
user@switch4# rstp bridge-priority 16k
user@switch4# rstp interface ge-0/0/23 cost 1000
user@switch4# rstp interface ge-0/0/23 mode point-to-point
user@switch4# rstp interface ge-0/0/19 cost 1000
user@switch4# rstp interface ge-0/0/19 mode point-to-point
```

Step-by-Step Procedure

If Switch 4 includes dual Routing Engines, configure nonstop bridging. To configure NSB on Switch 4:

1. Enable graceful Routing Engine switchover (GRES):

   ```
   [edit chassis redundancy]
   user@switch4# set graceful-switchover
   ```

2. Configure the switch to always synchronize configuration changes between the Routing Engines:

   ```
   [edit system]
   user@switch4# set commit synchronize
   ```

   If you try to commit a configuration in which nonstop bridging is configured but synchronization of configuration changes is not configured, the configuration is not committed.

3. Enable nonstop bridging:

   ```
   [edit protocols layer2-control]
   user@switch4# set nonstop-bridging
   ```

NOTE: This process enables NSB for all NSB-supported Layer 2 protocols on the switch, including RSTP.

Results

Check the results of the configuration:

```
user@switch4> show configuration
interfaces {
  ge-0/0/23 {
    unit 0 {
      family ethernet-switching {
        interface-mode trunk;
        vlan {
          members [10 20 30 40];
        }
      }
    }
  }
  ge-0/0/19 {
```

```
        unit 0 {
            family ethernet-switching {
                interface-mode trunk;
                vlan {
                    members [10 20 30 40];
                }
            }
        }
    }
}
protocols {
    layer2-control {
        nonstop-bridging;
    }
    rstp {
        bridge-priority 16k;
            interface ge-0/0/23 {
                cost 1000;
                mode point-to-point;
            }
            interface ge-0/0/19 {
                cost 1000;
                mode point-to-point;
            }
        }
    }
}
vlans {
    voice-vlan {
        vlan-id 10;
    }
    employee-vlan {
        vlan-id 20;
    }
    guest-vlan {
        vlan-id 30;
    }
    camera-vlan {
        vlan-id 40;
    }
}
system {
    commit synchronize;
}
chassis {
    redundancy {
        graceful-switchover;
    }
```

 Copyright © 2018, Juniper Networks, Inc.

Verification

To confirm that the configuration is working properly, perform these tasks on both Routing Engines:

- Verifying RSTP Configuration on Switch 1 on page 57
- Verifying RSTP Configuration on Switch 2 on page 57
- Verifying RSTP Configuration on Switch 3 on page 58
- Verifying RSTP Configuration on Switch 4 on page 58

Verifying RSTP Configuration on Switch 1

Purpose Verify the RSTP configuration on Switch 1.

Action Use the operational mode command:

user@switch1> **show spanning-tree interface**

```
Spanning tree interface parameters for instance 0

Interface     Port ID    Designated      Designated         Port    State  Role
                          port ID        bridge ID          Cost
ge-0/0/13     128:526      128:526  16384.0019e25040e0      1000   BLK    ALT
ge-0/0/9      128:522      128:522  32768.0019e2503d20      1000   BLK    ALT
ge-0/0/11     128:524      128:524   8192.0019e25051e0      1000   FWD    ROOT
```

Meaning The operational mode command **show spanning-tree interface** shows that **ge-0/0/13** is in a forwarding state. The other interfaces on Switch 1 are blocking.

Verifying RSTP Configuration on Switch 2

Purpose Use this procedure to verify the RSTP configuration on both Switch 2 Routing Engines.

Action Use the operational mode command:

user@switch2> **show spanning-tree interface**

```
Spanning tree interface parameters for instance 0

Interface     Port ID    Designated      Designated         Port    State  Role
                          port ID        bridge ID          Cost
ge-0/0/14      128:527      128:527   32768.0019e2503d20     1000   FWD    DESG
ge-0/0/18      128:529      128:529    8192.0019e25051e0     1000   FWD    ROOT
```

Meaning The operational mode command **show spanning-tree interface** shows that **ge-0/0/18** is in a forwarding state and is the root port.

Verifying RSTP Configuration on Switch 3

Purpose Use this procedure to verify the RSTP configuration on both Switch 3 Routing Engines.

Action Use the operational mode commands:

user@switch3> **show spanning-tree interface**

```
Spanning tree interface parameters for instance 0

Interface     Port ID    Designated      Designated         Port    State  Role
                          port ID        bridge ID          Cost
ge-0/0/26      128:539      128:539    8192.0019e25051e0     1000   FWD    DESG
ge-0/0/28      128:541      128:541    8192.0019e25051e0     1000   FWD    DESG
ge-0/0/24      128:537      128:537    8192.0019e25051e0     1000   FWD    DESG
```

Meaning The operational mode command **show spanning-tree interface** shows that no interface is the root interface.

Verifying RSTP Configuration on Switch 4

Purpose Use this procedure to verify the RSTP configuration on both Switch 4 Routing Engines.

 Copyright © 2018, Juniper Networks, Inc.

Action Use the operational mode commands:

```
user@switch4> show spanning-tree interface
Spanning tree interface parameters for instance 0

Interface     Port ID    Designated      Designated         Port    State  Role
                          port ID         bridge ID          Cost

ge-0/0/23     128:536      128:536    8192.0019e25051e0     1000   FWD     ROOT
ge-0/0/19     128:532      128:532    16384.0019e25040e0    1000   FWD     DESG
```

Meaning The operational mode command **show spanning-tree interface** shows that interface **ge-0/0/23** is the root interface and forwarding.

Release History Table

Release	Description
15.1	Starting with Junos OS Release 15.1 for EX Series and QFX Series switches with support for the Enhanced Layer 2 Software (ELS) configuration style, you can configure spanning tree parameters globally on all spanning tree interfaces.

Configuring RSTP on EX Series Switches (CLI Procedure)

The default spanning-tree protocol for EX Series switches is Rapid Spanning Tree Protocol (RSTP). RSTP provides faster convergence times than the original Spanning Tree Protocol (STP). Because RSTP is configured by default, you only need to use this procedure If another spanning-tree protocol has been configured. In that case, you can reconfigure RSTP.

To enable RSTP:

1. Disable the other configured spanning-tree protocol (MSTP):
   - To disable MSTP:

     ```
     [edit protocols]
     user@switch# set mstp disable
     ```

2. Configure RSTP
   - To enable RSTP on a specific interface:

     ```
     [edit protocols]
     user@switch# set rstp interface interface-name
     ```

   - To disable RSTP on a specific interface:

     ```
     [edit protocols]
     user@switch# set rstp interface interface-name disable
     ```

Copyright © 2018, Juniper Networks, Inc.

- To enable RSTP on a range of interfaces:

```
[edit protocols]
user@switch# set rstp interface interface-range-name
```

- To enable RSTP on all interfaces:

```
[edit protocols]
user@switch# set rstp interface all
```

Related Documentation

- show spanning-tree bridge on page 267
- show spanning-tree interface on page 273

Understanding VSTP

VLAN Spanning Tree Protocol (VSTP) enables Juniper Networks switches to run one or more Spanning Tree Protocol (STP) or Rapid Spanning Tree Protocol (RSTP) instances for each VLAN on which VSTP is enabled. For networks with multiple VLANs, VSTP improves intelligent tree spanning by defining best paths within the VLANs instead of within the entire network.

You can configure VSTP for a maximum of 509 VLANs.

VSTP and RSTP can be configured at the same time. If you configure VSTP and RSTP at the same time and the switch has more than 253 VLANs, VSTP is configured only for the first 253 VLANs. For the remaining VLANs, only RSTP is configured. RSTP and VSTP are the only spanning-tree protocols that can be configured at the same time on a switch.

NOTE: We recommend that you enable VSTP on all VLANs that could receive VSTP bridge protocol data units (BPDUs).

NOTE: Using the same VLAN for RSTP and VSTP is not supported. For example, if you are configuring a VLAN under VSTP, configuring RSTP with an interface that contains the same VLAN is not supported.

Related Documentation

- *Example: Configuring VSTP on QFX Series Switches and EX4600 Switches*

- Overview of Spanning-Tree Protocols on page 38
- Understanding RSTP on page 39
- Configuring VSTP (CLI Procedure) on page 61

- *Configuring VLAN Spanning Tree Protocol*

 Copyright © 2018, Juniper Networks, Inc.

Configuring VSTP (CLI Procedure)

NOTE: This topic applies to Junos OS for EX Series and QFX switches with support for the Enhanced Layer 2 Software (ELS) configuration style. If your switch runs software that does not support ELS, see *Configuring VSTP on EX Series Switches (CLI Procedure)*. For ELS details, see *Using the Enhanced Layer 2 Software CLI*.

The default spanning-tree protocol for EX Series switches is Rapid Spanning Tree Protocol (RSTP). VLAN Spanning Tree Protocol (VSTP) is an alternate protocol that allows EX Series switches to run one or more Spanning Tree Protocol (STP) or RSTP instances for each VLAN on which VSTP is enabled. For networks with multiple VLANs, VSTP improves intelligent tree spanning by defining best paths within the VLANs instead of within the entire network.

NOTE: On EX Series (other than EX9200) and QFX switches running Junos OS that supports ELS—VSTP can support up to 510 VLANs. However, on EX9200 switches, VSTP can support only up to 253 VLANs.

You can configure VSTP at the global level:

- For all interfaces on the switch
- For all interfaces within all VLANs
- For all interfaces within a specified VLAN
- For all interfaces within a specified VLAN group

You can configure or disable VSTP for a specific interface:

- For a specific interface on the switch
- For a specific interface within all VLANs
- For a specific interface within a specified VLAN
- For a specific interface within a specified VLAN group

NOTE:

- If you configure VSTP on an interface at both the global and the specific VLAN level, the interface configuration that is defined at the specific VLAN level overrides the interface configuration that is defined at the global level.
- If you specify VSTP to be configured on an interface that is not configured to belong to the VLAN (or VLANs), an error message is displayed.

To configure VSTP:

Copyright © 2018, Juniper Networks, Inc.

- For all interfaces within any of the following scopes:
  - For all interfaces on the switch:

    [edit protocols vstp]
    user@switch# **set interface all**

  - For all interfaces within all VLANs:

    [edit protocols vstp]
    user@switch# **set vlan all** interface all

  - For all interfaces within a specified VLAN:

    [edit protocols vstp]
    user@switch# **set vlan (*vlan-id* |*vlan-range* |*open-set-of-values*) interface all** interface all

  - For all interfaces within a specified VLAN group:

    [edit protocols vstp]
    user@switch# **set vlan-group *vlan-group-name* vlan (*vlan-id* |*vlan-range* |*open-set-of-values*) interface all**

- On a specific interface within any of the following scopes:
  - For a specific interface on the switch:

    [edit protocols vstp]
    user@switch# **set interface *interface-name***

  - For a specific interface within all VLANs:

    [edit protocols vstp]
    user@switch# **set vlan all interface *interface-name***

    CAUTION: Ensure that the interface is a member of all VLANs before you add the interface to the VSTP configuration. If the interface is not a member of all VLANs, this VSTP configuration will fail when you try to commit it.

  - For a specific interface within a specified VLAN:

    [edit protocols vstp]
    user@switch# **set vlan *vlan-id-or-vlan-range* interface *interface-name***

  - For a specific interface within a specified VLAN group:

    [edit protocols vstp]
    user@switch# **set vlan-group *vlan-group-name* vlan (*vlan-id* |*vlan-range* |*open-set-of-values*) interface *interface-name***

 Copyright © 2018, Juniper Networks, Inc.

Related Documentation

- show spanning-tree bridge on page 267
- show spanning-tree interface on page 273
- *Understanding VSTP*

Example: Configuring BPDU Protection on Switch Edge Interfaces With ELS to Prevent STP Miscalculations

EX Series and QFX Series switches provide Layer 2 loop prevention through Rapid Spanning Tree protocol (RSTP) and Multiple Spanning Tree Protocol (MSTP). All spanning-tree protocols use a special type of frame called a bridge protocol data unit (BPDU) to communicate. Other devices—PC bridging applications, for example, also use BPDUs and generate their own BPDUs. These different BPDUs are not compatible. When BPDUs generated by spanning-tree protocols are transmitted to a device that uses another type of BPDU, they can cause problems on the device. Similarly, if switches within a spanning-tree topology receive BPDUs from other devices, network outages can occur because of STP miscalculations.

This example configures BPDU protection on an EX Series switch that uses RSTP. The upstream configuration is done on the edge interfaces, where outside BPDUs are often received from other devices:

Requirements

This example uses the following software and hardware components:

- Two EX Series switches in an RSTP topology
- Junos OS Release 13.2X50-D10 or later or later for EX Series or QFX Series switches

Before you configure the interfaces on Switch 2 for BPDU protection, be sure you have:

- RSTP enabled on the switches.

NOTE: **By default, RSTP is enabled on all EX Series switches.**

Overview and Topology

The switches, being in an RSTP topology, support a loop-free network through the exchange of BPDUs. Receipt of outside BPDUs in an RSTP or MSTP topology, however, can lead to network outages by triggering an STP misconfiguration. To prevent such outages, enable BPDU protection on spanning tree interfaces that could receive outside

Copyright © 2018, Juniper Networks, Inc.

BPDUs. If an outside BPDU is received on a BPDU-protected interface, the interface shuts down to prevent the outside BPDU from accessing the spanning tree interface.

Figure 2 on page 64 shows the topology for this example. In this example, Switch 1 and Switch 2 are configured for RSTP and create a loop-free topology. The interfaces on Switch 2 are edge access ports—edge access ports frequently receive outside BPDUs generated by PC applications.

This example configures interface **ge-0/0/5** and interface **ge-0/0/6** as edge ports on Switch 2, and then configures BPDU protection on those ports. With BPDU protection enabled, these interfaces shut down when they encounter an outside BPDU sent by the PCs connected to Switch 2.

Figure 2: BPDU Protection Topology

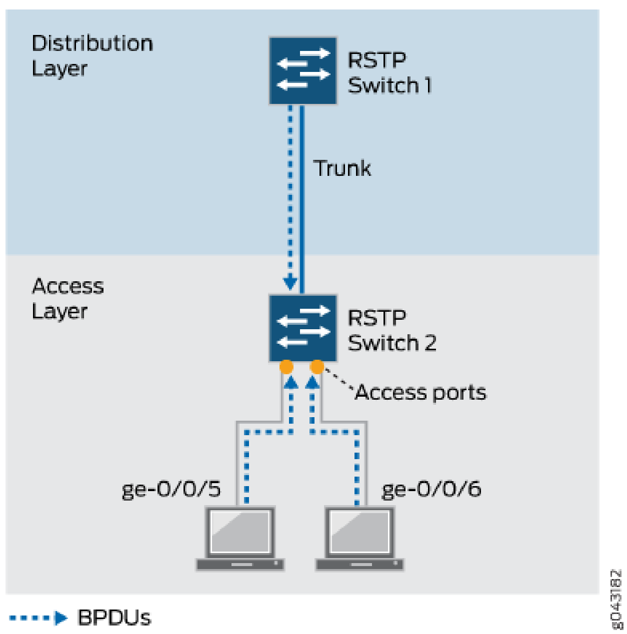

Table 4 on page 64 shows the components that will be configured for BPDU protection.

Table 4: Components of the Topology for Configuring BPDU Protection on EX Series Switches

Property	Settings
Switch 1 (Distribution Layer)	Switch 1 is connected to Switch 2 on a trunk interface.
Switch 2 (Access Layer)	Switch 2 has these access ports that require BPDU protection: • **ge-0/0/5** • **ge-0/0/6**

This configuration example uses RSTP topology. You also can configure BPDU protection for MSTP topologies at the [**edit protocols mstp**] hierarchy level.

 Copyright © 2018, Juniper Networks, Inc.

Configuration

To configure BPDU protection on two access interfaces:

CLI Quick Configuration

Quickly configure RSTP on the two Switch 2 interfaces, and then configure BPDU protection on all edge ports on Switch 2 by copying the following commands and pasting them into the switch terminal window:

NOTE: This example configures BPDU protection on specific interfaces. Starting with Junos OS Release 15.1 for EX Series and QFX Series switches with support for the Enhanced Layer 2 Software (ELS) configuration style, you can also configure BPDU protection globally on all spanning tree interfaces. See *Configuring BPDU Protection on Switch Spanning Tree Interfaces* for additional information.

```
[edit]
set protocols rstp interface ge-0/0/5 edge
set protocols rstp interface ge-0/0/6 edge
set protocols rstp bpdu-block-on-edge
```

Step-by-Step Procedure

To configure RSTP on the two Switch 2 interfaces, and then configure BPDU protection:

1. Configure RSTP on interface **ge-0/0/5** and interface **ge-0/0/6**, and configure them as edge ports:

```
[edit protocols rstp]
user@switch# set interface ge-0/0/5 edge
user@switch# set interface ge-0/0/6 edge
```

2. Configure BPDU protection on all edge ports on this switch:

```
[edit protocols rstp]
user@switch# set bpdu-block-on-edge
```

Results

Check the results of the configuration:

```
user@switch> show configuration protocols rstp
interface ge-0/0/5 {
  edge;
}
interface ge-0/0/6 {
  edge;
}
bpdu-block-on-edge;
```

Copyright © 2018, Juniper Networks, Inc.

Verification

To confirm that the configuration is working properly:

- Displaying the Interface State Before BPDU Protection Is Triggered on page 66
- Verifying That BPDU Protection Is Working Correctly on page 66

Displaying the Interface State Before BPDU Protection Is Triggered

Purpose Before BPDUs can be received from PCs connected to interface **ge-0/0/5** and interface **ge-0/0/6**, confirm the interface state.

Action Use the operational mode command:

```
user@switch>  show spanning-tree interface

Spanning tree interface parameters for instance 0

Interface     Port ID    Designated      Designated         Port    State  Role
                          port ID        bridge ID          Cost
ge-0/0/0      128:513      128:513  32768.0019e2503f00     20000  BLK    DIS
ge-0/0/1      128:514      128:514  32768.0019e2503f00     20000  BLK    DIS
ge-0/0/2      128:515      128:515  32768.0019e2503f00     20000  BLK    DIS
ge-0/0/3      128:516      128:516  32768.0019e2503f00     20000  FWD    DESG
ge-0/0/4      128:517      128:517  32768.0019e2503f00     20000  FWD    DESG
ge-0/0/5      128:518      128:518  32768.0019e2503f00     20000  FWD    DESG
ge-0/0/6      128:519      128:519  32768.0019e2503f00     20000  FWD    DESG
[output truncated]
```

Meaning The output from the operational mode command **show spanning-tree interface** shows that **ge-0/0/5** and interface **ge-0/0/6** are ports in a forwarding state.

Verifying That BPDU Protection Is Working Correctly

Purpose In this example, the PCs connected to Switch 2 start sending BPDUs to interface **ge-0/0/5** and interface **ge-0/0/6**. Verify that BPDU protection is working on the interfaces.

 Copyright © 2018, Juniper Networks, Inc.

Action Use the operational mode command:

```
user@switch> show spanning-tree interface
```

```
Spanning tree interface parameters for instance 0

Interface     Port ID    Designated      Designated         Port    State  Role
                          port ID         bridge ID          Cost
ge-0/0/0      128:513     128:513  32768.0019e2503f00     20000  BLK    DIS
ge-0/0/1      128:514     128:514  32768.0019e2503f00     20000  BLK    DIS
ge-0/0/2      128:515     128:515  32768.0019e2503f00     20000  BLK    DIS
ge-0/0/3      128:516     128:516  32768.0019e2503f00     20000  FWD    DESG
ge-0/0/4      128:517     128:517  32768.0019e2503f00     20000  FWD    DESG
ge-0/0/5      128:518     128:518  32768.0019e2503f00     20000  BLK    DIS
(Bpdu-Incon)
ge-0/0/6      128:519     128:519  32768.0019e2503f00     20000  BLK    DIS
(Bpdu-Incon)
ge-0/0/7      128:520       128:1  16384.00aabbcc0348     20000  FWD    ROOT
ge-0/0/8      128:521     128:521  32768.0019e2503f00     20000  FWD    DESG
[output truncated]
```

Meaning When BPDUs are sent from the PCs to interface **ge-0/0/5** and interface **ge-0/0/6** on Switch 2, the output from the operational mode command **show spanning-tree interface** shows that the interfaces have transitioned to a BPDU inconsistent state. The BPDU inconsistent state causes the interfaces to shut down.

Disabling the BPDU protection configuration on an interface does not automatically reenable the interface. However, if the disable-timeout statement has been included in the BPDU configuration, the interface does return to service after the timer expires. Otherwise, you must use the operational mode command **clear error bpdu** to unblock and reenable the interface.

If the PCs connected to Switch 2 send BPDUs to the interfaces again, BPDU protection is triggered once more and the interfaces transition back to the BPDU inconsistent state, causing them to shut down. In such cases, you need to find and repair the misconfiguration on the PCs that is sending BPDUs to Switch 2.

Release History Table

Release	Description
15.1	Starting with Junos OS Release 15.1 for EX Series and QFX Series switches with support for the Enhanced Layer 2 Software (ELS) configuration style, you can also configure BPDU protection globally on all spanning tree interfaces.

Related Documentation

- *Example: Faster Convergence and Improved Network Stability with RSTP on EX Series Switches*
- *Example: Configuring Loop Protection to Prevent Interfaces from Transitioning from Blocking to Forwarding in a Spanning Tree on non-ELS EX Series Switches*

Copyright © 2018, Juniper Networks, Inc.

- *Example: Configuring Root Protection to Enforce Root Bridge Placement in Spanning Trees on non-ELS EX Series Switches*

Copyright © 2018, Juniper Networks, Inc.

Configuring BPDU Protection on Switch Spanning Tree Interfaces

All spanning-tree protocols use a special type of frame called bridge protocol data units (BPDUs) to communicate with each other. Other devices in the network, such as PCs, generate their own BPDUs that are not compatible with the spanning-tree BPDUs. When BPDUs generated by other devices are transmitted to switches on which spanning-tree protocols are configured, a misconfiguration can occur in the spanning tree and a network outage can occur. Therefore, it is necessary to protect an interface in a spanning-tree topology from BPDUs generated from other devices.

On the ACX Series routers, MX Series routers QFX Series switches, and EX Series switches, you can configure BPDU protection to ignore BPDU received on interfaces where none are expected. If a BPDU is received on a blocked interface, the interface is disabled and stops forwarding frames. By default, all BPDUs are accepted and processed on all interfaces.

NOTE: This topic applies to Junos OS for EX Series and QFX switches with support for the Enhanced Layer 2 Software (ELS) configuration style. For ELS details, see *Using the Enhanced Layer 2 Software CLI*.

You can configure BPDU protection to ignore BPDU received on interfaces where none are expected. If a BPDU is received on a blocked interface, the interface is disabled and stops forwarding frames. By default, all BPDUs are accepted and processed on all interfaces.

Copyright © 2018, Juniper Networks, Inc.

To configure BPDU protection for spanning-tree instance interfaces:

- On a specific spanning-tree interface:
  1. Enable BPDU protection on a specified spanning-tree interface:

     ```
     [edit protocols layer2-control bpdu-block ]
     user@switcht# set interface (aex | (ge-fpc/pic/port | xe-fpc/pic/port)
     ```

     If a BPDU is received on the interface, the system will disable the interface and stop forwarding frames out the interface until the bridging process is restarted.

  2. (Optional) Configure the amount of time the system waits before *automatically* unblocking this interface after it has received a BPDU.

     ```
     [edit protocols layer2-contorl bpdu-block interface interface-name]
     user@switch# set disable-timeout seconds
     ```

     The range of the *seconds* option value is from 10 through 3600 seconds (one hour). A *seconds* option value of **0** is allowed, but this results in the default behavior (the interface is blocked until the interface is cleared).

  3. Verify the configuration of BPDU blocking for individual interfaces:

     ```
     [edit]
     interfaces {
       ge-fpc/pic/port { # VLAN encapsulation on a Gigabit Ethernet.
         encapsulation (ethernet-bridge | ethernett-vpls | extended-vlan-bridge |
           extended-vlan-vpls | vlan-bridge| vlan-vpls);
       }
       xe-fpc/pic/port { # VLAN encapsulation on 10-Gigabit Ethernet.
         encapsulation (ethernet-bridge | ethernett-vpls | extended-vlan-bridge |
           extended-vlan-vpls | vlan-bridge| vlan-vpls);
       }
       ae-X { # VLAN encapsulation
         encapsulation (ethernet-vpls vlan-vpls); # on aggregated Ethernet.
         ...
       }
       ae-X { # Extended VLAN encapsulation
         vlan-tagging; # on aggregated Ethernet.
         encapsulation extended-vlan-vpls;
         unit logical-unit-number {
           vlan-id number;
           ......
         }
         ......
       }
     }
     protocols
       layer2-control {
         bpdu-block
           interface interface-name;
           disable-timeout seconds;
     ```

 Copyright © 2018, Juniper Networks, Inc.

```
    }
  }
```

- To disable BPDU protection for a specific spanning-tree interface

  ```
  [edit protocols layer2-control bpdu-block interface interface-name]
  user@switch# set disable-timeout seconds
  ```

Related Documentation

- *Understanding Root Protection for Spanning-Tree Instance Interfaces in a Layer 2 Switched Network*

Unblocking a Switch Interface That Receives BPDUs in Error (CLI Procedure)

EX Series and QFX Series switches use bridge protocol data unit (BPDU) protection on interfaces to prevent them from receiving BPDUs that could trigger a spanning-tree misconfiguration. If BPDUs are received on a BPDU-protected interface, the interface either shuts down or transitions to a blocking state and stops forwarding frames. In the latter scenario, after the misconfiguration that triggered the BPDUs being sent to an interface is fixed in the topology, the interface can be unblocked and returned to service.

NOTE: This topic applies to Junos OS for EX Series and QFX switches with support for the Enhanced Layer 2 Software (ELS) configuration style. For switches that do not support ELS, see *Unblocking an Interface on non-ELS EX Series Switches That Receives BPDUs in Error (CLI Procedure)*. For ELS details, see *Using the Enhanced Layer 2 Software CLI*.

To unblock an interface and return it to service using the CLI:

- Automatically unblock an interface by configuring a timer that expires:

  ```
  [edit protocol layer 2]
  user@switch# set protocols layer2-control bpdu-block disable-timeout 30
  ```

  All interfaces on the switch will be reenabled (unblocked) after the timer expires. However, once an interface on the switch receives a new spanning-tree protocol BPDU, the interface returns to the blocked state.

- Manually unblock an interface using the operational mode command:

  ```
  user@switch> clear error bpdu interface ge-0/0/6
  ```

  This command will only reenable an interface but the BPDU configuration for the interface will continue to exist unless you remove the BPDU configuration explicitly.

Related Documentation

- *Example: Configuring BPDU Protection on Switch Edge Interfaces With ELS to Prevent STP Miscalculations*

Copyright © 2018, Juniper Networks, Inc.

- *Example: Configuring BPDU Protection on Interfaces to Prevent STP Miscalculations on EX Series Switches*
- *Understanding BPDU Protection for STP, RSTP, and MSTP*

 Copyright © 2018, Juniper Networks, Inc.

PART 4

Configuration Statements and Operational Commands

 Copyright © 2018, Juniper Networks, Inc.

CHAPTER 4

VLAN Configuration Statements

description (VLAN)

Syntax

```
description text-description;
```

Hierarchy Level

```
[edit vlans vlan-name]
```

Release Information

Statement introduced in Junos OS Release 9.0 for EX Series switches.
Statement introduced in Junos OS Release 11.1 for the QFX Series.
Option **text-description** enhanced from supporting up to 128 characters to supporting up to 256 characters in Junos OS Release 10.2 for EX Series switches.

Description

Provide a textual description for the VLAN. The text has no effect on the operation of the VLAN or switch.

Options

text-description—Text to describe the interface. It can contain letters, numbers, and hyphens (-) and can be up to 255 characters long. If the text includes spaces, enclose the entire text in quotation marks.

Required Privilege Level

routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.
system—To view this statement in the configuration.
system-control—To add this statement to the configuration.

Related Documentation

- Understanding Bridging and VLANs on Switches on page 9
- Understanding Bridging and VLANs on Switches on page 9

- *Example: Setting Up Basic Bridging and a VLAN on Switches*
- *Example: Setting Up Basic Bridging and a VLAN for an EX Series Switch*
- *Example: Setting Up Basic Bridging and a VLAN for an EX Series Switch with ELS Support*

- **show vlans on page 231**

Copyright © 2018, Juniper Networks, Inc.

dhcp-relay

Syntax

```
dhcp-relay {
  active-server-group server-group-name;
  authentication {
    password password-string;
    username-include {
      circuit-type;
      delimiter delimiter-character;
      domain-name domain-name-string;
      interface-name;
      logical-system-name;
      mac-address;
      option-60;
      option-82 <circuit-id> <remote-id>;
      routing-instance-name;
      user-prefix user-prefix-string;
    }
  }
  dhcpv6 {
    active-server-group server-group-name;
    authentication {
    password password-string;
      username-include {
        circuit-type;
        client-id;
        delimiter delimiter-character;
        domain-name domain-name-string;
        interface-name;
        logical-system-name;
        relay-agent-interface-id;
        relay-agent-remote-id;
        relay-agent-subscriber-id;
        routing-instance-name;
        user-prefix user-prefix-string;
      }
    dynamic-profile profile-name {
      aggregate-clients (merge | replace);
      use-primary primary-profile-name;
    }
    }
  }
    group group-name {
      active-server-group server-group-name;
      authentication {
        ...
      }
      dynamic-profile profile-name {
        ...
      }
      interface interface-name {
        exclude;
        liveness-detection {
```

```
            failure-action (clear-binding | clear-binding-if-interface-up | log-only);
            method {
              bfd {
                version (0 | 1 | automatic);
                minimum-interval milliseconds;
                minimum-receive-interval milliseconds;
                multiplier number;
                no-adaptation;
                transmit-interval {
                  minimum-interval milliseconds;
                  threshold milliseconds;
                }
                detection-time {
                  threshold milliseconds;
                }
                session-mode(automatic | multihop | singlehop);
                holddown-interval milliseconds;
              }
            }
          }
          overrides {
            ...
          }
          service-profile dynamic-profile-name;
          trace;
          upto upto-interface-name;
        }
        service-profile dynamic-profile-name;
      }
      overrides {
        ...
      }
      relay-agent-interface-id {
        ...
      }
    service-profile dynamic-profile-name;
    liveness-detection {
      failure-action (clear-binding | clear-binding-if-interface-up | log-only);
      method {
        bfd {
          version (0 | 1 | automatic);
          minimum-interval milliseconds;
          minimum-receive-interval milliseconds;
          multiplier number;
          no-adaptation;
          transmit-interval {
            minimum-interval milliseconds;
            threshold milliseconds;
          }
          detection-time {
            threshold milliseconds;
          }
          session-mode(automatic | multihop | singlehop);
          holddown-interval milliseconds;
        }
```


Copyright © 2018, Juniper Networks, Inc.

```
        }
    }
    overrides {
        allow-snooped-clients;
        interface-client-limit number;
        no-allow-snooped-clients;
        no-bind-on-request;
        send-release-on-delete;
    }
    relay-agent-interface-id {
        prefix prefix;
        use-interface-description (logical | device):
    }
    server-group {
        server-group-name {
            server-ip-address;
        }
    }
dynamic-profile profile-name {
    aggregate-clients (merge | replace);
    use-primary primary-profile-name;
}
forward-snooped-clients (all-interfaces | configured-interfaces |
    non-configured-interfaces);
group group-name {
    active-server-group server-group-name;
    authentication {
        ...
    }
    dynamic-profile profile-name {
        ...
    }
    interface interface-name {
        exclude;
        liveness-detection {
            failure-action (clear-binding | clear-binding-if-interface-up | log-only);
            method {
                bfd {
                    version (0 | 1 | automatic);
                    minimum-interval milliseconds;
                    minimum-receive-interval milliseconds;
                    multiplier number;
                    no-adaptation;
                    transmit-interval {
                        minimum-interval milliseconds;
                        threshold milliseconds;
                    }
                    detection-time {
                        threshold milliseconds;
                    }
                    session-mode(automatic | multihop | singlehop);
                    holddown-interval milliseconds;
                }
            }
        }
```

```
        overrides {
            ...
        }
        service-profile dynamic-profile-name;
        trace;
        upto upto-interface-name;
    }
    overrides {
        ...
    }
    relay-option-82 {
        ...
    }
    service-profile dynamic-profile-name;
}
liveness-detection {
    failure-action (clear-binding | clear-binding-if-interface-up | log-only);
    method {
        bfd {
            version (0 | 1 | automatic);
            minimum-interval milliseconds;
            minimum-receive-interval milliseconds;
            multiplier number;
            no-adaptation;
            transmit-interval {
                minimum-interval milliseconds;
                threshold milliseconds;
            }
            detection-time {
                threshold milliseconds;
            }
            session-mode(automatic | multihop | singlehop);
            holddown-interval milliseconds;
        }
    }
}
overrides {
    allow-snooped-clients;
    always-write-giaddr;
    always-write-option-82;
    client-discover-match <option60-and-option82>;
    disable-relay;
    interface-client-limit number;
    layer2-unicast-replies;
    no-allow-snooped-clients;
    no-bind-on-request;
    proxy-mode;
    replace-ip-source-with;
    send-release-on-delete;
    trust-option-82;
}
relay-option-82 {
    circuit-id {
        prefix prefix;
        use-interface-description (logical | device);
```

 Copyright © 2018, Juniper Networks, Inc.

```
    }
  }
  server-group {
    server-group-name {
      server-ip-address;
    }
  }
  service-profile dynamic-profile-name;
}
```

Hierarchy Level

```
[edit forwarding-options],
[edit vlans forwarding-options]
```

Release Information Statement introduced in Junos OS Release 11.3 for the QFX Series.

Description Configure extended Dynamic Host Configuration Protocol (DHCP) relay and DHCPv6 relay options on the switch and enable the switch to function as a DHCP relay agent. A DHCP relay agent forwards DHCP request and reply packets between a DHCP client and a DHCP server.

DHCP relay supports the attachment of dynamic profiles and also interacts with the local AAA Service Framework to use back-end authentication servers, such as RADIUS, to provide subscriber authentication. You can attach dynamic profiles and configure authentication support on a global basis or for a specific group of interfaces.

The extended DHCP and DHCPv6 relay agent options configured with the **dhcp-relay** and **dhcpv6** statements are incompatible with the DHCP/BOOTP relay agent options configured with the **bootp** statement. As a result, the extended DHCP or DHCPv6 relay agent and the DHCP/BOOTP relay agent cannot both be enabled on the router at the same time.

The remaining statements are explained separately.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- *Configuring DHCP and BOOTP*

filter (VLANs)

Syntax

```
filter (input | output) filter-name;
```

Hierarchy Level

```
[edit vlans vlan-name]
```

```
[edit vlans vlan-name forwarding-options]
```

Release Information

Statement introduced in Junos OS Release 9.0 for EX Series switches.
Statement introduced in Junos OS Release 11.1 for the QFX Series.
Statement introduced in Junos OS Release 13.2X50-D10 for EX Series switches.

Description

Apply a firewall filter to traffic entering or exiting a VLAN.

Default

All incoming traffic is accepted unmodified to the VLAN, and all outgoing traffic is sent unmodified from the VLAN.

Options

filter-name —Name of a firewall filter defined in a **filter** statement.

- **input**—Apply a firewall filter to VLAN ingress traffic.
- **output**—Apply a firewall filter to VLAN egress traffic.

Required Privilege Level

system—To view this statement in the configuration.
system-control—To add this statement to the configuration.
interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- *Example: Configuring Firewall Filters for Port, VLAN, and Router Traffic on EX Series Switches*
- *Configuring Firewall Filters*
- *Configuring Firewall Filters (CLI Procedure)*
- *Overview of Firewall Filters*
- *Firewall Filters for EX Series Switches Overview*
- *Configuring VLANs for EX Series Switches with ELS Support (CLI Procedure)*

Copyright © 2018, Juniper Networks, Inc.

forwarding-options

Syntax

```
forwarding-options {
  dhcp-relay {
    active-server-group server-group-name;
    authentication {
      password password-string;
      username-include {
        circuit-type;
        delimiter delimiter-character;
        domain-name domain-name-string;
        interface-name;
        logical-system-name;
        mac-address;
        option-60;
        option-82 <circuit-id> <remote-id>;
        routing-instance-name;
        user-prefix user-prefix-string;
      }
    }
    dhcpv6 {
      active-server-group server-group-name;
      authentication {
        password password-string;
        username-include {
          circuit-type;
          client-id;
          delimiter delimiter-character;
          domain-name domain-name-string;
          interface-name;
          logical-system-name;
          relay-agent-interface-id;
          relay-agent-remote-id;
          relay-agent-subscriber-id;
          routing-instance-name;
          user-prefix user-prefix-string;
        }
      }
      dynamic-profile profile-name {
        aggregate-clients (merge | replace);
        use-primary primary-profile-name;
      }
      group group-name {
        active-server-group server-group-name;
        authentication {
          ...
        }
        dynamic-profile profile-name {
          ...
        }
        interface interface-name {
          exclude;
          liveness-detection {
```

```
            failure-action (clear-binding | clear-binding-if-interface-up | log-only);
            method {
              bfd {
                version (0 | 1 | automatic);
                minimum-interval milliseconds;
                minimum-receive-interval milliseconds;
                multiplier number;
                no-adaptation;
                transmit-interval {
                  minimum-interval milliseconds;
                  threshold milliseconds;
                }
                detection-time {
                  threshold milliseconds;
                }
                session-mode(automatic | multihop | singlehop);
                holddown-interval milliseconds;
              }
            }
          }
          overrides {
            ...
          }
          service-profile dynamic-profile-name;
          trace;
          upto upto-interface-name;
        }
        service-profile dynamic-profile-name;
      }
      overrides {
        ...
      }
      relay-agent-interface-id {
        ...
      }
    service-profile dynamic-profile-name;
    liveness-detection {
      failure-action (clear-binding | clear-binding-if-interface-up | log-only);
      method {
        bfd {
          version (0 | 1 | automatic);
          minimum-interval milliseconds;
          minimum-receive-interval milliseconds;
          multiplier number;
          no-adaptation;
          transmit-interval {
            minimum-interval milliseconds;
            threshold milliseconds;
          }
          detection-time {
            threshold milliseconds;
          }
          session-mode(automatic | multihop | singlehop);
          holddown-interval milliseconds;
        }
```

 Copyright © 2018, Juniper Networks, Inc.

```
        }
      }
      overrides {
        allow-snooped-clients;
        interface-client-limit number;
        no-allow-snooped-clients;
        no-bind-on-request;
        send-release-on-delete;
      }
      relay-agent-interface-id {
        prefix prefix;
        use-interface-description (logical | device):
      }
      server-group {
        server-group-name {
          server-ip-address;
        }
      }
    dynamic-profile profile-name {
      aggregate-clients (merge | replace);
      use-primary primary-profile-name;
    }
    forward-snooped-clients (all-interfaces | configured-interfaces |
      non-configured-interfaces);
    group group-name {
      active-server-group server-group-name;
      authentication {
        ...
      }
      dynamic-profile profile-name {
        ...
      }
      interface interface-name {
        exclude;
        liveness-detection {
          failure-action (clear-binding | clear-binding-if-interface-up | log-only);
          method {
            bfd {
              version (0 | 1 | automatic);
              minimum-interval milliseconds;
              minimum-receive-interval milliseconds;
              multiplier number;
              no-adaptation;
              transmit-interval {
                minimum-interval milliseconds;
                threshold milliseconds;
              }
              detection-time {
                threshold milliseconds;
              }
              session-mode(automatic | multihop | singlehop);
              holddown-interval milliseconds;
            }
          }
        }
```

```
      overrides {
        ...
      }
      service-profile dynamic-profile-name;
      trace;
      upto upto-interface-name;
    }
    overrides {
      ...
    }
    relay-option-60 {
      ...
    }
    relay-option-82 {
      ...
    }
    service-profile dynamic-profile-name;
  }
  liveness-detection {
    failure-action (clear-binding | clear-binding-if-interface-up | log-only);
    method {
      bfd {
        version (0 | 1 | automatic);
        minimum-interval milliseconds;
        minimum-receive-interval milliseconds;
        multiplier number;
        no-adaptation;
        transmit-interval {
          minimum-interval milliseconds;
          threshold milliseconds;
        }
        detection-time {
          threshold milliseconds;
        }
        session-mode(automatic | multihop | singlehop);
        holddown-interval milliseconds;
      }
    }
  }
  overrides {
    allow-snooped-clients;
    always-write-giaddr;
    always-write-option-82;
    client-discover-match <option60-and-option82>;
    disable-relay;
    interface-client-limit number;
    layer2-unicast-replies;
    no-allow-snooped-clients;
    no-bind-on-request;
    proxy-mode;
    replace-ip-source-with;
    send-release-on-delete;
    trust-option-82;
  }
  relay-option-82 {
```

 Copyright © 2018, Juniper Networks, Inc.

```
        circuit-id {
            prefix prefix;
            use-interface-description (logical | device);
        }
    }
    server-group {
        server-group-name {
            server-ip-address;
        }
    }
    service-profile dynamic-profile-name;
}
dhcp-security {
    arp-inspection;
    group group-name {
        interface interface-name {
            static-ip ip-address {
                mac mac-address;
            }
        }
        overrides {
            no-option82;
            trusted;
            untrusted;
        }
    }
    ip-source-guard;
    no-dhcp-snooping;
    option-82 {
        circuit-id {
            prefix {
                host-name;
                logical-system-name;
                routing-instance-name;
            }
            use-interface-description (device | logical);
            use-vlan-id;
        }
        remote-id {
            host-name hostname;
            use-interface-description (device | logical);
            use-string string;
        }
        vendor-id {
            use-string string;
        }
    }
}
fip-security {
    examine-vn2vf {
    }
    examine-vn2vn {
        beacon-period milliseconds;
    }
    fc-map fc-map-value;
```

```
    interface interface-name {
      (fcoe-trusted | no-fcoe-trusted;)
    }
  }
}
```

Hierarchy Level

```
[edit]
[edit vlans]
```

Release Information

Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 11.3 for QFX Series switches.

Description

Configure traffic forwarding.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Copyright © 2018, Juniper Networks, Inc.

interface (VLANs)

List of Syntax

Syntax (QFX Series, QFabric, NFX Series and EX4600) on page 89
Syntax (EX Series and SRX210) on page 89

Syntax (QFX Series, QFabric, NFX Series and EX4600)

```
interface interface-name {
  mapping (native (push | swap) | tag (push | swap));
}
```

Syntax (EX Series and SRX210)

```
interface interface-name {
  egress;
  ingress;
  mapping (native (push | swap) | policy | tag (push | swap));
  pvlan-trunk;
}
```

QFX Series, QFabric, NFX Series and EX4600

```
[edit vlans vlan-name]
```

EX Series and SRX210

```
[edit vlans vlan-name],
[edit vlans vlan-name],
[edit vlans vlan-name vlan-id number],
[edit vlans vlan-name vlan-id number],
[edit vlans vlan-name vlan-id-list number]
```

Release Information

Statement introduced in Junos OS Release 9.3 for EX Series switches.
Statement introduced in Junos OS Release 11.1 for the QFX Series.

Description

For a specific VLAN, configure an interface.

Options

interface-name—Name of the Ethernet interface

The remaining statement is explained separately. See CLI Explorer.

Required Privilege Level

routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.
system—To view this statement in the configuration.system-control—To add this statement to the configuration.

Related Documentation

- *Example: Setting Up Basic Bridging and a VLAN on Switches*
- *Configuring VLANs on Switches*
- *Configuring VLANs for EX Series Switches (CLI Procedure)*
- *Configuring Q-in-Q Tunneling on EX Series Switches (CLI Procedure)*

- *Configuring Q-in-Q Tunneling on EX Series Switches with ELS Support (CLI Procedure)*

Copyright © 2018, Juniper Networks, Inc.

interface-mac-limit

Syntax

```
interface-mac-limit {
  limit
  disable;
  packet-action ;
}
```

Hierarchy Level

[edit bridge-domains *bridge-domain-name* bridge-options],
[edit bridge-domains *bridge-domain-name* bridge-options interface *interface-name*],
[edit logical-systems *logical-system-name* bridge-domains *bridge-domain-name* bridge-options],
[edit logical-systems *logical-system-name* bridge-domains *bridge-domain-name* bridge-options interface *interface-name*],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* bridge-domains *bridge-domain-name* bridge-options],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* bridge-domains *bridge-domain-name* bridge-options interface *interface-name*],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* switch-options],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* switch-options interface *interface-name*],
[edit logical-systems *logical-system-name* switch-options],
[edit logical-systems *logical-system-name* switch-options interface *interface-name*],
[edit routing-instances *routing-instance-name* bridge-domains *bridge-domain-name* bridge-options],
[edit routing-instances *routing-instance-name* bridge-domains *bridge-domain-name* bridge-options interface *interface-name*],
[edit routing-instances *routing-instance-name* switch-options],
[edit routing-instances *routing-instance-name* switch-options interface *interface-name*],
[edit switch-options],
[edit switch-options],
[edit switch-options interface *interface-name*],
[edit switch-options interface *interface-name*],
[edit vlans *vlan-name* switch-options],
[edit vlans *vlan-name* switch-options interface *interface-name*]

Release Information

Statement introduced in Junos OS Release 8.4.
Support for the **switch-options** statement added in Junos OS Release 9.2.
Support for top-level configuration for the **virtual-switch** type of routing instance added in Junos OS Release 9.2. In Junos OS Release 9.1 and earlier, the routing instances hierarchy supported this statement only for a VPLS instance or a bridge domain configured within a virtual switch.
Support for logical systems added in Junos OS Release 9.6.
[edit switch-options], [edit switch-options interface *interface-name*], [edit vlans *vlan-name* switch-options], and [edit vlans *vlan-name* switch-options interface *interface-name*] hierarchy levels introduced in Junos OS Release 12.3R2 for EX Series switches.

Description Configure a limit to the number of MAC addresses that can be learned from a bridge domain, VLAN, virtual switch, or set of bridge domains or VLANs.

NOTE: For multichassis link aggregation (MC-LAG) peers in active-active mode, configuring the **interface-mac-limit** statement or changing the **interface-mac-limit** configuration when traffic is flowing can cause the MAC entries to be out of synchronization between the two MC-LAG peers, which might result in flooding. To avoid flooding, you must either halt traffic forwarding and then configure the **interface-mac-limit** statement or use the **commit at** configuration statement to commit the changes at the same time in both the peer nodes.

Alternatively, if flooding does occur, you can clear the bridge MAC table on both the routers or switches by using the **clear bridge mac-table** command. Running this command ensures that the MAC entries are re-learned and in synchronization between both the peers.

Default The default MAC limit varies with the platform.

Options **disable**—Disables the global interface-mac-limit configuration on an interface and sets the maximum interface-mac-limit that is permitted on the device.

limit—Sets the maximum number of MAC addresses learned from an interface.

Range: 1 through <default MAC limit> MAC addresses per interface. Range is platform specific.

If you configure both **disable** and **limit**, disable takes precedence and packet-action is set to **none**. The remaining statement is explained separately.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- *Understanding Layer 2 Learning and Forwarding for Bridge Domains*
- Layer 2 Learning and Forwarding for VLANs Overview on page 6
- *Understanding Layer 2 Learning and Forwarding for Bridge Domains Functioning as Switches with Layer 2 Trunk Ports*
- *Layer 2 Learning and Forwarding for VLANs Acting as a Switch for a Layer 2 Trunk Port*

 Copyright © 2018, Juniper Networks, Inc.

interface-mode

Syntax

```
interface-mode (access | trunk <inter-switch-link>);
```

Hierarchy Level

[edit interfaces *interface-name* unit *logical-unit-number* family bridge],
[edit interfaces *interface-name* unit *logical-unit-number* family ethernet-switching],
[edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number* family bridge]

Release Information

Statement introduced in Junos OS Release 9.2.
Statement introduced in Junos OS Release 13.2X50-D10 for EX Series switches.
Statement introduced in Junos OS Release 13.2 for the QFX Series.
Statement introduced in Junos OS Release 15.1.
inter-switch-link option introduced in Junos OS Release 14.2 for MX240, MX480, and MX960 routers in enhanced LAN mode.

Description

NOTE: This statement supports the Enhanced Layer 2 Software (ELS) configuration style. If your switch runs software that does not support ELS, see port-mode. For ELS details, see *Using the Enhanced Layer 2 Software CLI*.

QFX3500 and QFX3600 standalone switches—Determine whether the logical interface accepts or discards packets based on VLAN tags. Specify the **trunk** option to accept packets with a VLAN ID that matches the list of VLAN IDs specified in the **vlan-id** or **vlan-id-list** statement, then forward the packet within the bridge domain or VLAN configured with the matching VLAN ID. Specify the **access** option to accept packets with no VLAN ID, then forward the packet within the bridge domain or VLAN configured with the VLAN ID that matches the VLAN ID specified in the **vlan-id** statement.

NOTE: On MX Series routers, if you want IGMP snooping to be functional for a bridge domain, then you should not configure interface-mode and irb for that bridge. Such a configuration commit succeeds, but IGMP snooping is not functional, and a message informing the same is displayed. For more information, see *Configuring a Trunk Interface on a Bridge Network*.

Options

access—Configure a logical interface to accept untagged packets. Specify the VLAN to which this interface belongs using the **vlan-id** statement.

trunk—Configure a single logical interface to accept packets tagged with any VLAN ID specified with the **vlan-id** or **vlan-id-list** statement.

Copyright © 2018, Juniper Networks, Inc.

trunk inter-switch-link—For a private VLAN, configure the InterSwitch Link protocol (ISL) on a trunk port of the primary VLAN in order to connect the switches composing the PVLAN to each other. You do not need to configure an ISL when a PVLAN is configured on a single switch. This configuration specifies whether the particular interface assumes the role of interswitch link for the PVLAN domains of which it is a member. This option is supported only on MX240, MX480, and MX960 routers in enhanced LAN mode.

Required Privilege Level

interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- *Configuring Access Mode on a Logical Interface*
- *Configuring a Logical Interface for Trunk Mode*
- *Example: Connecting Access Switches with ELS Support to a Distribution Switch with ELS Support*
- *Tunnel Services Overview*
- *Tunnel Interface Configuration on MX Series Routers Overview*

mac (Static MAC-Based VLANs)

Syntax

```
mac mac-address {
    next-hop interface-name;
}
```

Hierarchy Level

```
[edit ethernet-switching-options static vlan vlan-name]
```

Description

Specify the MAC address to add to the Ethernet switching table.

The remaining statement is explained separately. See CLI Explorer.

Options

mac-address—MAC address

Required Privilege Level

system—To view this statement in the configuration.
system-control—To add this statement to the configuration.

Related Documentation

- *Adding a Static MAC Address Entry to the Ethernet Switching Table (CLI Procedure)*

Copyright © 2018, Juniper Networks, Inc.

members

Syntax

```
members [(all | names | vlan-ids)];
```

Hierarchy Level

```
[edit interfaces (QFX Series) interface-name unit 0 family ethernet-switching vlan]
```

Release Information

Statement introduced in Junos OS Release 9.0 for EX Series switches.
Statement updated with enhanced **?** (CLI completion feature) functionality in Junos OS Release 9.5 for EX Series switches.
Statement introduced in Junos OS Release 11.1 for the QFX Series.

Description

For trunk interfaces, configure the VLANs for which the interface can carry traffic.

TIP: To display a list of all configured VLANs on the system, including VLANs that are configured but not committed, type ? after **vlan** or **vlans** in your configuration mode command line. Note that only one VLAN is displayed for a VLAN range.

NOTE: The number of VLANs supported per switch varies for each model. Use the configuration-mode command **set vlans id vlan-id ?** to determine the maximum number of VLANs allowed on a switch. You cannot exceed this VLAN limit because each VLAN is assigned an ID number when it is created. You can, however, exceed the recommended VLAN member maximum.

On an EX Series switch that runs Junos OS that does not support the Enhanced Layer 2 Software (ELS) configuration style, the maximum number of VLAN members allowed on the switch is 8 times the maximum number of VLANs the switch supports (vmember limit = vlan max * 8). If the switch configuration exceeds the recommended VLAN member maximum, you see a warning message when you commit the configuration. If you ignore the warning and commit such a configuration, the configuration succeeds but you run the risk of crashing the Ethernet switching process (**eswd**) due to memory allocation failure.

On an EX Series switch that runs Junos OS that supports ELS, the maximum number of VLAN members allowed on the switch is 24 times the maximum number of VLANs the switch supports (vmember limit = vlan max * 24). If the configuration of one of these switches exceeds the recommended VLAN member maximum, a warning message appears in the system log (syslog).

Copyright © 2018, Juniper Networks, Inc.

Options

all—Specifies that this trunk interface is a member of all VLANs that are configured on this switch. When a new VLAN is configured on the switch, this trunk interface automatically becomes a member of the VLAN.

NOTE: Since VLAN members are limited, specifying **all** could cause the number of VLAN members to exceed the limit at some point.

NOTE: Each VLAN that is configured must have a specified VLAN ID when you attempt to commit the configuration; otherwise, the configuration commit fails. Also, **all** cannot be the name of a VLAN on the switch.

names—Name of one or more VLANs. VLAN IDs are applied automatically in this case.

vlan-ids—Numeric identifier of one or more VLANs. For a series of tagged VLANs, specify a range; for example, **10–20** or **10–20 23 27–30**.

NOTE: Each configured VLAN must have a specified VLAN ID to successfully commit the configuration; otherwise, the configuration commit fails.

Required Privilege Level

routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.
interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- *Configuring Gigabit Ethernet Interfaces (CLI Procedure)*
- *Configuring Gigabit Ethernet Interfaces (CLI Procedure) for EX Series Switches with ELS support*
- *Configuring Gigabit Ethernet Interfaces (J-Web Procedure)*
- *Configuring VLANs for EX Series Switches (CLI Procedure)*
- *Configuring VLANs for EX Series Switches with ELS Support (CLI Procedure)*
- *Example: Setting Up Basic Bridging and a VLAN on Switches*
- Understanding Bridging and VLANs on Switches on page 9
- **show ethernet-switching interfaces on page 198**
- **show vlans on page 231**

Copyright © 2018, Juniper Networks, Inc.

native-vlan-id

Syntax

```
native-vlan-id vlan-id;
```

Hierarchy Level (QFX Series and EX4600)

For platforms without ELS:

[edit interfaces (QFX Series) *interface-name* unit 0 family ethernet-switching]

For platforms with ELS:

[edit interfaces (QFX Series) *interface-name*]

Hierarchy Level (ACX Series, EX Series, SRX Series, M Series, MX Series, and T Series)

[edit interfaces ge-*fpc/pic/port*],
[edit interfaces *interface-name*]

Hierarchy Level (SRX Series)

[edit interfaces*interface-name*]

Release Information

Statement introduced in Junos OS Release 8.3.
Statement introduced in Junos OS Release 9.0 for EX Series switches.
Statement introduced in Junos OS Release 9.5 for SRX Series.
Statement introduced in Junos OS Release 11.1 for the QFX Series.
Statement introduced in Junos OS Release 12.2 for ACX Series Universal Metro Routers.
Statement introduced in Junos OS Release 12.3R2 for EX Series switches.
Statement introduced in Junos OS Release 13.2X51-D20 for the QFX Series.

Description

Configure the VLAN identifier to associate with untagged packets received on the physical interface of a trunk mode interface for the following:

- QFX Series and EX4600
- M Series routers with Gigabit Ethernet IQ PICs with SFP and Gigabit Ethernet IQ2 PICs with SFP configured for 802.1Q flexible VLAN tagging
- MX Series routers with Gigabit Ethernet DPCs and MICs, Tri-Rate Ethernet DPCs and MICs, and 10-Gigabit Ethernet DPCs and MICs and MPCs configured for 802.1Q flexible VLAN tagging
- T4000 routers with 100-Gigabit Ethernet Type 5 PIC with CFP
- EX Series switches with Gigabit Ethernet, 10-Gigabit Ethernet, 40-Gigabit Ethernet, and aggregated Ethernet interfaces

The logical interface on which untagged packets are received must be configured with the same VLAN ID as the native VLAN ID configured on the physical interface, otherwise the untagged packets are dropped. To configure the logical interface, include the **vlan-id** statement (matching the **native-vlan-id** statement on the physical interface) at the **[edit interfaces *interface-name* unit *logical-unit-number*] hierarchy level**.

When the **native-vlan-id** statement is included with the *flexible-vlan-tagging* statement, untagged packets are accepted on the same mixed VLAN-tagged port and on the interfaces that are configured for Q-in-Q tunneling.

When the **native-vlan-id** statement is combined with the interface-mode statement, untagged packets are accepted and forwarded within the bridge domain or VLAN that is configured with the matching VLAN ID.

To configure the logical interface, include the **vlan-id** statement (matching the **native-vlan-id** statement on the physical interface) at the **[edit interfaces *interface-name* unit *logical-unit-number*] hierarchy level**.

NOTE: Starting in Junos OS Release 17.1R1, you can send untagged traffic without a native VLAN ID to the remote end of the network. To do this, remove the native VLAN ID from the untagged traffic configuration by setting the no-native-vlan-insert statement. If you do not configure this statement, the native VLAN ID is added to the untagged traffic.

Default By default, the untagged packets are dropped. That is, if you do not configure the **native-vlan-id** option, the untagged packets are dropped.

Options *vlan-id*—Numeric identifier of the VLAN.

Range: 1 through 4094

number—VLAN ID number.

Range: (ACX Series routers, SRX Series devices and EX Series switches) 0 through 4094.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.
interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

 Copyright © 2018, Juniper Networks, Inc.

Related Documentation

- *Configuring Gigabit Ethernet Interfaces (CLI Procedure)*
- *Configuring Gigabit Ethernet Interfaces (J-Web Procedure)*
- Understanding Bridging and VLANs on Switches on page 9
- *Enabling VLAN Tagging*
- *Configuring Access Mode on a Logical Interface*
- Configuring the Native VLAN Identifier on Switches With ELS Support (CLI Procedure) on page 20
- *Understanding Interfaces*
- *Understanding Q-in-Q Tunneling and VLAN Translation*
- *no-native-vlan-insert*
- *Sending Untagged Traffic Without VLAN ID to Remote End*
- **show ethernet-switching interfaces on page 198**
- **show vlans on page 231**
- *flexible-vlan-tagging*
- *Junos OS Network Interfaces Configuration Guide*

packet-action

Syntax

packet-action *action*;

Hierarchy Level

[edit bridge-domains *bridge-domain-name* bridge-options interface *interface-name* interface-mac-limit *limit*],
[edit bridge-domains *bridge-domain-name* bridge-options interface-mac-limit *limit*],
[edit logical-systems *logical-system-name* bridge-domains *bridge-domain-name* bridge-options interface *interface-name* interface-mac-limit *limit*],
[edit logical-systems *logical-system-name* bridge-domains *bridge-domain-name* bridge-options interface-mac-limit *limit*],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* bridge-domains *bridge-domain-name* bridge-options interface *interface-name* interface-mac-limit *limit*],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* bridge-domains *bridge-domain-name* bridge-options interface-mac-limit *limit*],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* switch-options interface *interface-name* interface-mac-limit *limit*],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* switch-options interface-mac-limit *limit*],
[edit logical-systems *logical-system-name* switch-options interface-mac-limit *limit*],
[edit protocols l2-learning global-mac-limit *limit*],
[edit routing-instances *routing-instance-name* bridge-domains *bridge-domain-name* bridge-options interface *interface-name* interface-mac-limit *limit*],
[edit routing-instances *routing-instance-name* bridge-domains *bridge-domain-name* bridge-options interface-mac-limit *limit*],
[edit routing-instances *routing-instance-name* protocols evpn interface-mac-limit (VPLS)],
[edit routing-instances *routing-instance-name* protocols evpn interface *interface-name* interface-mac-limit (VPLS)],
[edit routing-instances *routing-instance-name* protocols evpn mac-table-size *limit*],
[edit routing-instances *routing-instance-name* switch-options interface *interface-name* interface-mac-limit *limit*],
[edit routing-instances *routing-instance-name* switch-options interface-mac-limit *limit*],
[edit switch-options interface-mac-limit *limit*],
[edit switch-options interface *interface-name* interface-mac-limit *limit*],
[edit switch-options interface-mac-limit *limit*],
[edit switch-options interface *interface-name* interface-mac-limit *limit*],
[edit switch-options interface-mac-limit *limit*],
[edit switch-options mac-table-size *limit*],
[edit switch-options interface *interface-name* interface-mac-limit *limit*],
[edit vlans *vlan-name* switch-options interface *interface-name* interface-mac-limit *limit*],
[edit vlans *vlan-name* switch-options interface-mac-limit *limit*],
[edit vlans *vlan-name* switch-options mac-table-size *limit*]
[edit vlans *vlan-name* switch-options interface-mac-limit *limit*],
[edit vlans *vlan-name* switch-options interface *interface-name* interface-mac-limit *limit*],
[edit vlans *vlan-name* switch-options mac-table-size *limit*]

Release Information

Statement introduced in Junos OS Release 8.4.
Support for the **switch-options** statement added in Junos OS Release 9.2.
Support for top-level configuration for the **virtual-switch** type of routing instance added in Junos OS Release 9.2. In Junos OS Release 9.1 and earlier, the routing instances hierarchy

Copyright © 2018, Juniper Networks, Inc.

supported this statement only for a VPLS instance or a bridge domain configured within a virtual switch.
Support for logical systems added in Junos OS Release 9.6.
[edit switch-options interface *interface-name* interface-mac-limit *limit*], [edit switch-options interface-mac-limit *limit*], [edit switch-options mac-table-size *limit*], [edit vlans *vlan-name* switch-options interface *interface-name* interface-mac-limit *limit*], [edit vlans *vlan-name* switch-options interface-mac-limit *limit*], and [edit vlans *vlan-name* switch-options mac-table-size *limit*] hierarchy levels introduced in Junos OS Release 12.3R2 for EX Series switches.
Support for EVPNs introduced in Junos OS Release 13.2 on MX Series 5G Universal Routing Platforms.
Support at the **[edit switch-options interface *interface-name* interface-mac-limit *limit*]** hierarchy level and hierarchy levels under **[edit vlans *vlan-name*]** introduced in Junos OS Release 13.2X50-D10 for EX Series switches and Junos OS Release 13.2 for the QFX Series.

Description Specify the action taken when packets with new source MAC addresses are received after the MAC address limit is reached. If this statement is not configured, packets with new source MAC addresses are forwarded by default.

NOTE: The **packet-action** statement is not supported on the QFX10002-60C switch.

Default

NOTE: On a QFX Series Virtual Chassis, if you include the **shutdown** option at the **[edit vlans *vlan-name* switch-options interface *interface-name* interface-mac-limit packet-action]** hierarchy level and issue the **commit** operation, the system generates a commit error. The system does not generate an error if you include the **shutdown** option at the **[edit switch-options interface *interface-name* interface-mac-limit packet-action]** hierarchy level.

Disabled. The default is for packets for new source MAC addresses to be forwarded after the MAC address limit is reached.

Options

drop—Drop packets with new source MAC addresses, and do not learn the new source MAC addresses.

> NOTE: On QFX10000 switches, if you include the drop option, you cannot configure unicast reverse-path forwarding (URFP) on integrated routing and bridging (IRB) and MAC limiting on the same interface. If you have an MC-LAG configuration, you cannot configure MAC limiting on the interchassis link (ICL) interface.

drop-and-log—(EX Series switches and QFX Series only) Drop packets with new source MAC addresses, and generate an alarm, an SNMP trap, or a system log entry.

log—(EX Series switches and QFX Series only) Hold packets with new source MAC addresses, and generate an alarm, an SNMP trap, or a system log entry.

none—(EX Series switches and QFX Series only) Forward packets with new source MAC addresses, and learn the new source MAC address.

shutdown—(EX Series switches and QFX Series only) Disable the specified interface, and generate an alarm, an SNMP trap, or a system log entry.

Required Privilege Level

routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- *Configuring EVPN Routing Instances*
- *Configuring EVPN Routing Instances on EX9200 Switches*
- Configuring MAC Limiting (CLI Procedure) on page 32
- *Configuring Persistent MAC Learning (CLI Procedure)*
- *Understanding Layer 2 Learning and Forwarding for Bridge Domains*
- Layer 2 Learning and Forwarding for VLANs Overview on page 6
- *Understanding Layer 2 Learning and Forwarding for Bridge Domains Functioning as Switches with Layer 2 Trunk Ports*
- Layer 2 Learning and Forwarding for VLANs Overview on page 6
- *Layer 2 Learning and Forwarding for VLANs Acting as a Switch for a Layer 2 Trunk Port*

 Copyright © 2018, Juniper Networks, Inc.

port-mode

Syntax

```
port-mode (access | tagged-access | trunk);
```

Hierarchy Level

```
[edit interfaces (QFX Series) interface-name unit logical-unit-number family
  ethernet-switching]
```

Release Information

Statement introduced in Junos OS Release 9.0 for EX Series switches.
Statement introduced in Junos OS Release 11.1 for the QFX Series.

Description

NOTE: This statement does not support the Enhanced Layer 2 Software (ELS) configuration style. If your switch runs software that supports ELS, see interface-mode. For ELS details, see *Using the Enhanced Layer 2 Software CLI*.

Configure whether an interface on the switch operates in access, tagged access, or trunk mode.

Default

All switch interfaces are in access mode.

Options

access—Have the interface operate in access mode. In this mode, the interface can be in a single VLAN only. Access interfaces typically connect to network devices, such as PCs, printers, IP telephones, and IP cameras.

tagged-access—Have the interface operate in tagged-access mode. In this mode, the interface can be in multiple VLANs. Tagged access interfaces typically connect to network devices, such as PCs, printers, IP telephones, and IP cameras.

trunk—Have the interface operate in trunk mode. In this mode, the interface can be in multiple VLANs and can multiplex traffic between different VLANs. Trunk interfaces typically connect to other switches and to routers on the LAN.

NOTE: The number of VLANs supported per switch varies for each model. Use the configuration-mode command **set vlans id vlan-id ?** to determine the maximum number of VLANs allowed on a switch. You cannot exceed this VLAN limit because each VLAN is assigned an ID number when it is created. You can, however, exceed the recommended VLAN member maximum. To determine the maximum number of VLAN members allowed on a switch, multiply the VLAN maximum for the switch times 8 (vmember limit = vlan max * 8).

Copyright © 2018, Juniper Networks, Inc.

If a switch configuration exceeds the recommended VLAN member maximum, you see a warning message when you commit the configuration. If you ignore the warning and commit such a configuration, the configuration succeeds but you run the risk of crashing the Ethernet switching process (eswd) due to memory allocation failure.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- *Example: Connecting an EX Series Access Switch to a Distribution Switch*
- *Configuring Gigabit Ethernet Interfaces (CLI Procedure)*
- *Configuring VLANs for EX Series Switches (CLI Procedure)*
- *Junos OS Ethernet Interfaces Configuration Guide*

service-id

Syntax service-id *number*;

Hierarchy Level [edit switch-options]
[edit vlans *vlan-name*]

Release Information Statement introduced in Junos OS Release 12.3R2 for EX Series switches and MX Series routers.
Statement introduced in Junos OS Release 13.2 for the QFX Series.

Description Specify a service identifier for each multichassis aggregated Ethernet interface that belongs to a link aggregation group (LAG).

Options **number**—A number that identifies a particular service.

Range: 1 through 65535

Required Privilege Level system—To view this statement in the configuration.
system control—To add this statement to the configuration.

 Copyright © 2018, Juniper Networks, Inc.

switch-options (VLANs)

List of Syntax

Syntax (EX Series, MX Series, QFX Series and NFX Series) on page 105
Syntax (SRX Series) on page 105

Syntax (EX Series, MX Series, QFX Series and NFX Series)

```
switch-options {
  interface interface-name {
    interface-mac-limit limit {
      packet-action drop;
    }
    mac-pinning
    no-mac-learning;
    static-mac static-mac-address {
      vlan-id number;
    }
  }
  interface-mac-limit limit {
    packet-action drop;
  }
  mac-statistics;
  mac-ip-table-size number;
  mac-table-size limit {
    packet-action drop;
  }
  no-mac-learning;
  service-id number;
  vtep-source-interface
}
```

Syntax (SRX Series)

```
switch-options {
  interface interface-name {
    encapsulation-type;
    ignore-encapsulation-mismatch;
    pseudowire-status-tlv;
    static-mac mac-address {
      vlan-id vlan-id;
    }
  }
  mac-table-aging-time seconds;
  mac-table-size {
    number;
    packet-action drop;
  }
}
```

EX Series, MX Series, QFX Series and NFX Series

```
[edit ],
[edit logical-systems logical-system-name routing-instances routing-instance-name vlans
  vlan-name],
[edit routing-instances routing-instance-name vlans vlan-name],
[edit vlans vlan-name]
```

SRX Series

```
[edit vlans vlans-name]
```

Release Information

Statement modified in Junos OS Release 9.5.
Statement introduced in Junos OS Release 12.3R2 for EX Series switches and MX Series routers.
Statement introduced in Junos OS Release 13.2 for the QFX Series.
Statement (mac-pinning) introduced in Junos OS 16.2 for MX Series routers.
mac-ip-table-size statement introduced in Junos OS 17.4 Release for MX Series routers and EX9200 switches.

Description

Configure Layer 2 learning and forwarding properties for a VLAN or a virtual switch.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- *Ethernet Switching and Layer 2 Transparent Mode Overview*

Copyright © 2018, Juniper Networks, Inc.

static-mac

Syntax

```
static-mac mac-address;
```

```
static-mac mac-address {
  vlan-id number;
}
```

Hierarchy Level

```
[edit vlans vlan-name switch-options interface interface-name]
```

```
[edit bridge-domains bridge-domain-name bridge-options interface interface-name],
```

```
[edit logical-systems logical-system-name bridge-domains bridge-domain-name
  bridge-options interface interface-name],
[edit logical-systems logical-system-name routing-instances routing-instance-name
  bridge-domains bridge-domain-name bridge-options interface interface-name],
```

```
[edit routing-instances routing-instance-name bridge-domains bridge-domain-name
  bridge-options interface interface-name],
[edit routing-instances routing-instance-name protocols evpn interface interface-name]
```

Release Information

Statement introduced in Junos OS Release 8.4.
Statement modified in Junos OS Release 9.5.
Support for logical systems added in Junos OS Release 9.6.
[edit vlans *vlan-name* switch-options interface *interface name*] hierarchy level introduced in Junos OS Release 12.3R2 for EX Series switches.
Statement introduced in Junos OS Release 13.2 for EX Series switches.
Statement introduced in Junos OS Release 13.2 for the QFX Series.
Support for EVPNs added in Junos OS Release 13.2 for MX 3D Series routers. The **vlan-id** option is not available for EVPNs.
[edit vlans *vlan-name* switch-options interface *interface name*] hierarchy level introduced in Junos OS Release 13.2 for the QFX Series.

Description

Configure a static MAC address for a logical interface in a bridge domain or VLAN.

The **vlan-id** option can be specified for **static-macs** only if **vlan-id all** is configured for the bridging domain or VLAN.

Options

mac-address—MAC address

vlan-id ***number***—(Optional) VLAN identifier to associate with static MAC address.

Required Privilege Level

routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Copyright © 2018, Juniper Networks, Inc.

Related Documentation

- *Configuring EVPN Routing Instances*
- *Understanding Layer 2 Learning and Forwarding for Bridge Domains*
- Layer 2 Learning and Forwarding for VLANs Overview on page 6
- Adding a Static MAC Address Entry to the Ethernet Switching Table on a Switch with ELS Support (CLI Procedure) on page 23
- *Understanding VLANs on Security Devices*

 Copyright © 2018, Juniper Networks, Inc.

vlan-id

Syntax

```
vlan-id (all | none | number);
```

VLANs and Bridge Domain VLANs

For platforms without ELS:

```
[edit vlans vlan-name vlan-range]
```

For platforms without ELS and with ELS:

```
[edit vlans vlan-name]
```

For ELS platforms only:

```
[edit interfaces interface-name unit number]
[edit vlans vlan-name vlan-id-list]
```

```
[edit vlans vlan-name],
[edit logical-systems logical-system-name vlansvlan-name],
[edit logical-systems logical-system-name routing-instances routing-instance-name
   vlansvlan-name],
[edit routing-instances routing-instance-name vlans vlan--name]
```

802.1Q Tagging

```
[edit vlans vlan-name]
```

VLAN ID to Rewrite

```
[edit interfaces interface-name unit logical-unit-number input-vlan-map],
[edit interfaces interface-name unit logical-unit-number output-vlan-map],
[edit logical-systems logical-system-name interfaces interface-name unit logical-unit-number
   input-vlan-map],
[edit logical-systems logical-system-name interfaces interface-name unit logical-unit-number
   output-vlan-map]
```

VLAN Tagging and Layer 3 Subinterfaces

```
[edit interfaces interface-name unit logical-unit-number]
```

Release Information

Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 8.4.
Statement introduced in Junos OS Release 9.0 for EX Series switches.
Statement introduced in Junos OS Release 9.2 for EX Series switches VLAN tagging and Layer 3 subinterfaces.
Support for Layer 2 trunk ports added in Junos OS Release 9.2.
Support for SRX 5600, and SRX 5800 devices added in Junos OS Release 9.6.
Support for logical systems added in Junos OS Release 9.6.
Statement introduced in Junos OS Release 11.1 for the QFX Series.
Statement introduced in Junos OS Release 12.3R2 for EX Series switches.

Description For VLANs, specify a VLAN identifier (VID) to include in the packets sent to and from the VLAN, or a VPLS routing instance.

NOTE: When configuring a VLAN identifier for provider backbone bridge (PBB) routing instances, dual-tagged VIDs and the none option are not permitted.

For 802.1Q tagging, configure an 802.1Q tag to apply to all traffic that originates on the VLAN.

The number zero is reserved for priority tagging and the number 4095 is also reserved.

For VLAN ID to Rewrite Gigabit Ethernet IQ and 10-Gigabit Ethernet IQ2, 10-Gigabit Ethernet LAN/WAN PIC, and IQ2-E interfaces and aggregated Ethernet using Gigabit Ethernet IQ interfaces, specify the line VLAN identifiers to be rewritten at the input or output interface.

You cannot include the **vlan-id** statement with the **swap** statement, **swap-push** statement, **push-push** statement, or **push-swap** statement at the **[edit interfaces *interface-name* unit *logical-unit-number* output-vlan-map]** hierarchy level. If you include any of those statements in the output VLAN map, the VLAN ID in the outgoing frame is rewritten to the **vlan-id** statement that you include at the **[edit interfaces *interface-name* unit *logical-unit-number*]** hierarchy level.

 Copyright © 2018, Juniper Networks, Inc.

Default For 802.1Q Tagging on EX Series and SRX Series, If you use the default factory configuration, all traffic originating on the VLAN is untagged and has a VLAN identifier of 1.

For VLANs on a QFX3500 and QFX3500 switch, if you use the default factory configuration, all traffic originating on the VLAN is untagged and has a VLAN identifier of 1. The number zero is reserved for priority tagging and the number 4093 is also reserved.

On a QFX5100 switch, if you use the default factory configuration, all traffic originating on the VLAN is untagged and has a VLAN identifier of 1. The number zero is reserved for priority tagging and the number 4093 is also reserved.

NOTE: You can only create up 4090 VLANs on a QFX5100 switch. If you create more than 4090 VLANs, the interfaces associated with the extra VLANs are not displayed in the **show vlans** command output. For example, if you create 4094 VLANs, the extra VLANs will not have interfaces associated with the VLANs. The order in which you configure the extra VLANs determines which interfaces are missing from the **show vlans** command output.

For VLAN tagging and Layer 3 subinterfaces, bind an 802.1Q VLAN tag ID to a logical interface.

NOTE: The VLAN tag ID cannot be configured on logical interface unit 0. The logical unit number must be 1 or higher.

Copyright © 2018, Juniper Networks, Inc.

Options For VLANs:

number—A valid VLAN identifier. If you configure multiple VLANs with a valid VLAN identifier, you must specify a unique VLAN identifier for each. However, you can use the same VLAN identifier for VLANs that belong to different virtual switches. Use this option to send single tagged frames with the specified VLAN identifier over VPLS VT interfaces.

NOTE: If you specify a VLAN identifier, you cannot also use the **all** option. They are mutually exclusive.

all—Specify that the VLAN spans all the VLAN identifiers configured on the member logical interfaces.

NOTE: You cannot specify the **all** option if you include a routing interface in the VLAN.

none—Specify to enable shared VLAN learning or to send untagged frames over VPLS VT interfaces.

NOTE: Multichassis link aggregation (MC-LAG) does not support the **none** option with the **vlan-id** statement with VLANs.

For 802.1Q Tagging:

number —VLAN tag identifier

Range:

- 1 through 4094 (all switches except EX8200 Virtual Chassis)
- 1 through 4092 (EX8200 Virtual Chassis only)

Default: 1

Required Privilege Level

routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.
system—To view this statement in the configuration.
system-control—To add this statement to the configuration.
interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

 Copyright © 2018, Juniper Networks, Inc.

Related Documentation

- *Example: Connecting Access Switches with ELS Support to a Distribution Switch with ELS Support*
- *Example: Configuring a Private VLAN on a Single Switch with ELS Support*
- *Creating a Private VLAN on a Single Switch with ELS Support (CLI Procedure)*
- *Creating a Private VLAN Spanning Multiple EX Series Switches (CLI Procedure)*
- *Example: Configuring VLANs on Security Devices*
- *Example: Configuring Interfaces and Routing Instances for a User Logical Systems*
- *Rewriting the VLAN Tag on Tagged Frames*
- *Binding VLAN IDs to Logical Interfaces*
- vlan-tagging on page 116
- *Example: Configuring Layer 3 Subinterfaces for a Distribution Switch and an Access Switch*
- *Configuring Gigabit Ethernet Interfaces (CLI Procedure)*
- *Configuring Gigabit Ethernet Interfaces (CLI Procedure) for EX Series Switches with ELS support*
- *Configuring Gigabit Ethernet Interfaces (J-Web Procedure)*
- *Configuring a Layer 3 Subinterface (CLI Procedure)*
- *Configuring Q-in-Q Tunneling on EX Series Switches with ELS Support (CLI Procedure)*
- *Junos OS Ethernet Interfaces Configuration Guide*

Copyright © 2018, Juniper Networks, Inc.

vlan-id-list

Syntax

```
vlan-id-list [ vlan-id-numbers ];
```

Hierarchy Level

[edit bridge-domains *bridge-domain-name*],
[edit logical-systems *logical-system-name* bridge-domains *bridge-domain-name*],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name*
bridge-domains *bridge-domain-name*],
[edit routing-instances *routing-instance-name* bridge-domains *bridge-domain-name*],
[edit interfaces *interface-name* unit 0],
[edit interfaces *interface-name* unit *logical-unit-number*],
[edit vlans *vlan-name*]

Release Information

Statement introduced in Junos OS Release 9.4.
Support for logical systems added in Junos OS Release 9.6.
Statement introduced in Junos OS Release 12.3R2 for EX Series switches.
Statement introduced in Junos OS Release 13.2 for the QFX Series.

Description

Specify a VLAN identifier list to use for a bridge domain or VLAN in trunk mode.

Specify the **trunk** option in the **interface-mode** statement to accept packets with a VLAN ID that matches the list of VLAN IDs specified in the **vlan-id-list** statement to forward the packet within the bridge domain or VLAN configured with the matching VLAN ID. Specify the **access** option to accept packets with no VLAN ID to forward the packet within the bridge domain or VLAN configured with the VLAN ID that matches the VLAN ID specified in the **vlan-id** statement.

This statement also enables you to bind a logical interface to a list of VLAN IDs, thereby configuring the logical interface to receive and forward a frame with a tag that matches the specified VLAN ID list.

WARNING: On some EX and QFX Series switches, you can apply no more than eight VLAN identifier lists to a physical interface.

Options

vlan-id-numbers—Valid VLAN identifiers. You can combine individual numbers with range lists by including a hyphen.

Range: 0 through 4095

NOTE: On EX Series switches and the QFX Series, the range is 0 through 4094.

Copyright © 2018, Juniper Networks, Inc.

Required Privilege Level
routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation
- *Configuring a Bridge Domain*
- *Configuring a VLAN*
- *Configuring VLAN Identifiers for Bridge Domains and VPLS Routing Instances*
- *Configuring VLAN Identifiers for VLANs and VPLS Routing Instances*

vlan-rewrite

Syntax
```
vlan-rewrite translate (200 500 | 201 501)
```

Hierarchy Level
[edit interfaces *interface-name* unit *number* family bridge interface-mode trunk]
[edit interfaces *interface-name* unit *number* family ethernet-switching interface-mode trunk]

Release Information
Statement introduced in Junos OS Release 9.4.
Statement introduced in Junos OS Release 12.3R2 for EX Series switches.

Description
Translates an incoming VLAN to a bridge-domain VLAN, corresponding counter translation at egress. Supports translation of VLAN 200 to VLAN 500 and VLAN 201 to VLAN 501. Other valid VLANs pass through without translation.

Options
translate 200 500—Translates incoming packets with VLAN 200 to 500.

translate 201 501—Translates incoming packets with VLAN 201 to 501.

translate 202 502—Translates incoming packets with VLAN 202 to 502.

Required Privilege Level
interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation
- *Rewriting a VLAN Tag and Adding a New Tag*

Copyright © 2018, Juniper Networks, Inc.

vlan-tagging

Syntax

```
vlan-tagging;
```

Syntax (QFX Series, NFX Series, and EX4600)

```
vlan-tagging;
```

Syntax (SRX Series Interfaces)

vlan-tagging native-vlan-id *vlan-id*;

Hierarchy Level

[edit interfaces *interface-name*],
[edit logical-systems *logical-system-name* interfaces *interface-name*]

QFX Series, NFX Series, and EX4600 Interfaces

[edit interfaces (QFX Series) *interface-name*]
[edit interfaces (QFX Series) interface-range *interface-range-name*]

SRX Series Interfaces

[edit interfaces *interface*]

Release Information

Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 9.0 for EX Series switches.
Statement introduced in Junos OS Release 9.5.
Statement introduced in Junos OS Release 11.3 for the QFX Series.
Statement introduced in Junos OS Release 12.2 for ACX Series Universal Metro Routers.
Statement introduced in Junos OS Release 13.2 for PTX Series Routers.
Statement introduced in Junos OS Release 14.1X53-D10 for the QFX Series.

Description

For Fast Ethernet and Gigabit Ethernet interfaces, aggregated Ethernet interfaces configured for VPLS, and pseudowire subscriber interfaces, enable the reception and transmission of 802.1Q VLAN-tagged frames on the interface.

NOTE: For QFX Series configure VLAN identifier for untagged packets received on the physical interface of a trunk mode interface. Enable VLAN tagging. The platform receives and forwards single-tag frames with 802.1Q VLAN tags.

On EX Series switches except for EX4300 and EX9200 switches, the **vlan-tagging** and **family ethernet-switching** statements cannot be configured on the same interface. Interfaces on EX2200, EX3200, EX3300, EX4200, and EX4500 switches are set to **family ethernet-switching** by the default factory configuration. EX6200 and EX8200 switch interfaces do not have a default **family** setting.

Copyright © 2018, Juniper Networks, Inc.

Default VLAN tagging is disabled by default.

Options **native-vlan-id**— (SRX Series)Configures a VLAN identifier for untagged packets. Enter a number from 0 through 4094.

> NOTE: The **native-vlan-id** can be configured only when either **flexible-vlan-tagging** mode or **interface-mode** trunk is configured.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- *802.1Q VLANs Overview*
- *Configuring a Layer 3 Subinterface (CLI Procedure)*
- *Configuring Tagged Aggregated Ethernet Interfaces*
- *Example: Configuring Layer 3 Subinterfaces for a Distribution Switch and an Access Switch*
- *vlan-id*
- *Configuring a Layer 3 Logical Interface*
- *Configuring VLAN Tagging*

Copyright © 2018, Juniper Networks, Inc.

vlan-tags

Syntax

```
vlan-tags outer number inner number;
```

Hierarchy Level

```
[edit bridge-domains bridge-domain-name],
[edit logical-systems logical-system-name bridge-domains bridge-domain-name],
[edit logical-systems logical-system-name routing-instances routing-instance-name
   bridge-domains bridge-domain-name],
[edit routing-instances routing-instance-name bridge-domains bridge-domain-name]
[edit vlans vlan-name]
```

Release Information

Statement introduced in Junos OS Release 8.4.
Support for logical systems added in Junos OS Release 9.6.
Statement introduced in Junos OS Release 12.3R2 for EX Series switches.
Statement introduced in Junos OS Release 13.2X51-D10 for QFX Series switches.

Description

Specify dual VLAN identifier tags for a bridge domain, VLAN, or VPLS routing instance.

Options

outer ***number***—A valid VLAN identifier.

inner ***number***—A valid VLAN identifier.

Required Privilege Level

routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- *Configuring a Bridge Domain*
- *Configuring a VLAN*
- *Configuring VLAN Identifiers for Bridge Domains and VPLS Routing Instances*
- *Configuring VLAN Identifiers for VLANs and VPLS Routing Instances*
- *Configuring a Layer 2 Virtual Switch* .
- *Configuring a Layer 2 Virtual Switch on an EX Series Switch*

Copyright © 2018, Juniper Networks, Inc.

CHAPTER 5

MAC Address Configuration Statements

global-mac-table-aging-time

Syntax

```
global-mac-table-aging-time seconds;
```

Hierarchy Level

```
[edit protocols l2-learning]
```

Release Information

Statement introduced in Junos OS Release 9.2.
Statement modified in Junos OS Release 9.5.
Support for logical systems added in Junos OS Release 9.6.

Description

Configure the timeout interval for entries in the MAC table.

Default

300 seconds

Options

seconds—Time elapsed before MAC table entries are timed out and entries are deleted from the table.

Range: For MX Series routers: 10 through 1 million; for EX Series and QFX Series switches: 60 through 1 million; for SRX devices: 10 through 64,000 seconds

Required Privilege Level

routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- *Configuring the MAC Table Timeout Interval*
- *Configuring MAC Table Aging on Switches*
- *Example: Configuring VLANs on Security Devices*

Copyright © 2018, Juniper Networks, Inc.

mac-limit

List of Syntax

Syntax (QFX Series and EX4600) on page 121
Syntax (SRX Series and EX Series) on page 121

Syntax (QFX Series and EX4600)

```
mac-limit number;
```

Syntax (SRX Series and EX Series)

```
mac-limit limit action action;
```

Hierarchy Level

```
[edit vlans vlan-name]
```

Release Information

Statement introduced in Junos OS Release 9.0 for EX Series switches.
Statement introduced in Junos OS Release 11.1 for the QFX Series.
The short description of **interface-mac-limit** at the CLI command hierarchy is changed from **Maximum number of MAC addresses per interface (1..16383)** to **Maximum number of MAC addresses per interface (1..5120)** at the **[edit vlans vlan-name switch-options]** hierarchy level from Junos OS Release 18.2R1.

Description

Specify the maximum number of MAC addresses to be associated with a VLAN—the default is **unlimited**, which can leave the network vulnerable to flooding. Change **unlimited** to any number from 2 to the switch's maximum VLAN MAC limit. The maximum number of MAC addresses allowed in a switching table per VLAN varies depending on the EX Series switch. To see the maximum number of MAC addresses per VLAN allowed on your switch, issue the **set vlans *vlan-name* mac-limit ?** configuration-mode command.

NOTE: Do not set the **mac-limit** value to 1. The first learned MAC address is often inserted into the forwarding database automatically—for instance, for a routed VLAN interface (RVI), the first MAC address inserted into the forwarding database is the MAC address of the RVI. For aggregated Ethernet bundles (LAGs) using LACP, the first MAC address inserted into the forwarding database in the Ethernet switching table is the source address of the protocol packet. In these cases, the switch does not learn MAC addresses other than the automatic address when **mac-limit** is set to 1, and this causes problems with MAC learning and forwarding.

When the MAC limit set by this statement is reached, no more MAC addresses are added to the Ethernet switching table. You can also, optionally, have a system log entry generated when the limit is exceeded by adding the option **action log**.

NOTE: When you reconfigure the number of MAC addresses, the Ethernet switching table is not automatically cleared. Therefore, if you reduce the

Copyright © 2018, Juniper Networks, Inc.

number of addresses from the default (unlimited) or a previously set limit, you could already have more entries in the table than the new limit allows. Previous entries remain in the table after you reduce the number of addresses, so you should clear the Ethernet switching table for a specified interface, MAC address, or VLAN when you reduce the MAC limit. Use the command clear ethernet-switching table to clear existing MAC addresses from the table before using the **mac-limit** configuration statement.

Default The MAC limit is disabled, so entries are unlimited.

Options QFX Series and EX4600:

number—Maximum number of MAC addresses.

Range: 1 through 32768

NOTE: This statement is not supported on QFabric systems.

EX Series:

limit—Maximum number of MAC addresses.

Range: 1 through *switch maximum*

SRX Series:

number—Maximum number of MAC addresses.

Range: 1 through 5120

action—**Log** is the only action available. Configure **action log** to add a message to the system log when the mac-limit value is exceeded. A typical logged message looks like this:

May 5 06:18:31 bmp-199p1-dev edwd[5665]: ESWD_VLAN_MAC_LIMIT_EXCEEDED: vlan default mac 00:1f:12:37:af:5b (tag 40). vlan limit exceeded

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- show vlans on page 231
- *Example: Setting Up Basic Bridging and a VLAN on Switches*
- *Configuring MAC Table Aging on Switches*
- Understanding Bridging and VLANs on Switches on page 9

 Copyright © 2018, Juniper Networks, Inc.

mac-notification

Syntax

```
mac-notification {
  notification-interval seconds;
}
```

Hierarchy Level

```
[edit ethernet-switching-options]
[edit switch-options]
```

Release Information

Statement introduced in Junos OS Release 9.6 for EX Series switches.
Statement introduced in Junos OS Release 11.1 for the QFX Series.
Hierarchy level **[edit switch-options]** added in Junos OS Release 14.1X53-D10 for EX Series and QFX Series.

Description

Enable MAC notification for a switch. If you configure this statement without setting a notification interval, MAC notification is enabled with the default MAC notification interval of 30 seconds.

The remaining statement is explained separately. See CLI Explorer.

Default

MAC notification is disabled by default.

Required Privilege Level

system—To view this statement in the configuration.
system-control—To add this statement to the configuration.
routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- *Configuring Non-ELS MAC Notification*
- Configuring MAC Notification on Switches with ELS Support (CLI Procedure) on page 29

mac-statistics

Syntax

mac-statistics;

Hierarchy Level

[edit bridge-domains *bridge-domain-name* bridge-options],
[edit logical-systems *logical-system-name* bridge-domains *bridge-domain-name* bridge-options],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* bridge-domains *bridge-domain-name* bridge-options],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* switch-options],
[edit logical-systems *logical-system-name* switch-options],
[edit routing-instances *routing-instance-name* bridge-domains *bridge-domain-name* bridge-options],
[edit routing-instances *routing-instance-name* switch-options],
[edit routing-instances *routing-instance-name* protocols evpn],
[edit switch-options],
[edit switch-options],
[edit vlans *vlan-name* switch-options]

Release Information

Statement introduced in Junos OS Release 8.4.
Support for the **switch-options** statement added in Junos OS Release 9.2.
Support for top-level configuration for the **virtual-switch** type of routing instance added in Junos OS Release 9.2. In Junos OS Release 9.1 and earlier, the routing instances hierarchy supported this statement only for a VPLS instance or a bridge domain configured within a virtual switch.
Support for logical systems added in Junos OS Release 9.6.
[edit switch-options] and [edit vlans *vlan-name* switch-options] hierarchy levels introduced in Junos OS Release 12.3R2 for EX Series switches.
Support for EVPNs added in Junos OS Release 13.2 for MX 3D Series routers.
[edit switch-options] and [edit vlans *vlan-name* switch-options] hierarchy levels introduced in Junos OS Release 13.2 for the QFX Series.

Description

(MX Series routers, EX Series switches, and QFX Series only) For bridge domains or VLANs, enable MAC accounting either for a specific bridge domain or VLAN, or for a set of bridge domains or VLANs associated with a Layer 2 trunk port.

Default

disabled

Required Privilege Level

routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- *Understanding Layer 2 Learning and Forwarding for Bridge Domains*
- Layer 2 Learning and Forwarding for VLANs Overview on page 6

Copyright © 2018, Juniper Networks, Inc.

- *Understanding Layer 2 Learning and Forwarding for Bridge Domains Functioning as Switches with Layer 2 Trunk Ports*
- *Layer 2 Learning and Forwarding for VLANs Acting as a Switch for a Layer 2 Trunk Port*
- *Configuring EVPN Routing Instances*
- *Configuring EVPN Routing Instances on EX9200 Switches*

mac-table-size

Syntax

```
mac-table-size limit {
  packet-action drop;
}
```

Hierarchy Level

```
[edit bridge-domains bridge-domain-name bridge-options],
[edit logical-systems logical-system-name bridge-domains bridge-domain-name
  bridge-options],
[edit logical-systems logical-system-name routing-instances routing-instance-name
  bridge-domains bridge-domain-name bridge-options],
[edit logical-systems logical-system-name routing-instances routing-instance-name
  switch-options],
[edit logical-systems logical-system-name switch-options],
[edit routing-instances routing-instance-name bridge-domains bridge-domain-name
  bridge-options],
[edit routing-instances routing-instance-name switch-options],
[edit switch-options],
[edit switch-options],
[edit vlans vlan-name switch-options]
```

Release Information

Statement introduced in Junos OS Release 8.4.
Support for the **switch-options** statement added in Junos OS Release 9.2.
Support for top-level configuration for the **virtual-switch** type of routing instance added in Junos OS Release 9.2. In Junos OS Release 9.1 and earlier, the routing instances hierarchy supported this statement only for a VPLS instance or a bridge domain configured within a virtual switch.
Support for logical systems added in Junos OS Release 9.6.
[edit switch-options] and **[edit vlans *vlan-name* switch-options]** hierarchy levels introduced in Junos OS Release 12.3R2 for EX Series switches.
Support at the **[edit vlans *vlan-name* switch-options** hierarchy level introduced in Junos OS Release 13.2 for the QFX Series.

Description

Modify the size of the MAC address table for the bridge domain or VLAN, a set of bridge domains or VLANs associated with a trunk port, or a virtual switch. The default is 5120 MAC addresses.

NOTE: For multichassis link aggregation (MC-LAG) peers in active-active mode, configuring the **mac-table-size** statement or changing the **mac-table-size** configuration when traffic is flowing can cause the MAC entries to be out of synchronization between the two MC-LAG peers, which might result in flooding. To avoid flooding, you must either halt traffic forwarding and then configure the **mac-table-size** statement or use the **commit at** configuration statement to commit the changes at the same time in both the peer nodes.

 Copyright © 2018, Juniper Networks, Inc.

Alternatively, if flooding does occur, you can clear the bridge MAC table on both the routers by using the **clear bridge mac-table** command. Running this command ensures that the MAC entries are re-learned and in synchronization between both the peers.

Options

limit—Specify the maximum number of addresses in the MAC address table.

Range: 16 through 1,048,575 MAC addresses

Default: 5120 MAC addressesThere is no default MAC address limit for the **mac-table-size** statement at the **[edit switch-options]** hierarchy level. The number of MAC addresses that can be learned is only limited by the platform, 65,535 MAC addresses for EX Series switches and 1,048,575 MAC addresses for other devices.

The remaining statement is explained separately. See CLI Explorer.

Required Privilege Level

routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- *Understanding Layer 2 Learning and Forwarding for Bridge Domains*
- Layer 2 Learning and Forwarding for VLANs Overview on page 6
- *Understanding Layer 2 Learning and Forwarding for Bridge Domains Functioning as Switches with Layer 2 Trunk Ports*
- *Layer 2 Learning and Forwarding for VLANs Acting as a Switch for a Layer 2 Trunk Port*

notification-interval

Syntax

notification-interval *seconds*;

Hierarchy Level

[edit ethernet-switching-options mac-notification]
[edit switch-options mac-notification]

Release Information

Statement introduced in Junos OS Release 9.6 for EX Series switches.
Statement introduced in Junos OS Release 11.1 for the QFX Series.
Hierarchy level **[edit switch-options]** added in Junos OS Release 14.1X53-D10 for EX Series and QFX Series.

Description

Configure the MAC notification interval for a switch.

The MAC notification interval is the amount of time the switch waits before sending learned or unlearned MAC address SNMP notifications to the network management server. For instance, if the MAC notification interval is set to 10, all of the MAC address addition and removal SNMP notifications will be sent to the network management system every 10 seconds.

Options

seconds—The MAC notification interval, in seconds.

Range: 1 through 60

Default: 30

Required Privilege Level

system—To view this statement in the configuration.
system-control—To add this statement to the configuration.
routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- *Configuring Non-ELS MAC Notification*
- *Configuring Non-ELS MAC Notification*
- Configuring MAC Notification on Switches with ELS Support (CLI Procedure) on page 29

Copyright © 2018, Juniper Networks, Inc.

CHAPTER 6

STP Configuration Statements

- revision-level on page 177
- rstp on page 178
- traceoptions (Spanning Tree) on page 181
- vlan (MSTP) on page 185
- vlan (VSTP) on page 187
- vlan-group on page 189
- vstp on page 190

 Copyright © 2018, Juniper Networks, Inc.

bpdu-block

Syntax

```
bpdu-block {
  interface (interface-name disable | all);
  disable-timeout seconds;
}
```

Hierarchy Level

```
[edit protocols layer2-control ]
```

```
[edit ethernet-switching-options]
```

Release Information

Statement introduced in Junos OS Release 9.1 for EX Series switches.
Statement introduced in Junos OS Release 13.2X50-D10 for EX Series switches.

Description

Configures bridge protocol data unit (BPDU) protection on a specified interface or on all interfaces. If the interface receives incompatible BPDUs, it is disabled.

If the disable-timeout statement is included in the BPDU configuration, the interface is automatically reenabled after the timer expires. Otherwise, you must use the operational mode command **clear ethernet-switching bpdu-error interface** to unblock and reenable the interface.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- *Understanding BPDU Protection for Spanning-Tree Instance Interfaces*
- *Configuring BPDU Protection for Individual Spanning-Tree Instance Interfaces*
- *Understanding BPDU Protection for STP, RSTP, and MSTP*
- *Example: Configuring BPDU Protection on Interfaces to Prevent STP Miscalculations on EX Series Switches*
- *Unblocking an Interface on non-ELS EX Series Switches That Receives BPDUs in Error (CLI Procedure)*
- *Understanding BPDU Protection for STP, RSTP, and MSTP*
- show spanning-tree bridge on page 267
- show spanning-tree interface on page 273
- *clear ethernet-switching bpdu-error interface*

Copyright © 2018, Juniper Networks, Inc.

bpdu-block-on-edge

Syntax

bpdu-block-on-edge;

Hierarchy Level

[edit logical-systems *logical-system-name* protocols (mstp | rstp | vstp)],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols (mstp | rstp | vstp)],
[edit protocols (mstp | rstp| vstp)],
[edit routing-instances *routing-instance-name* protocols (mstp | rstp | vstp)]

Release Information

Statement introduced in Junos OS Release 9.1 for EX Series switches.
Statement introduced in Junos OS Release 9.4 for additional devices.
Statement updated in Junos OS Release 9.4 for EX Series switches to add VSTP support.
Statement updated in Junos OS Release 11.1 for EX Series switches to change blocking behavior to port shutdown.
Support for logical systems added in Junos OS Release 9.6.
Statement introduced in Junos OS Release 17.1 for ACX Series routers.

Description

Configure bridge protocol data unit (BPDU) protection on all edge ports of a switch. When the **bpdu-block-on-edge** statement is configured and the interface encounters an incompatible BPDU, the interface shuts down.

If the disable-timeout statement is included in the BPDU configuration, the interface is automatically reenabled after the timer expires. Otherwise, you must use the operational mode command **clear ethernet-switching bpdu-error interface** to unblock and reenable the interface.

Default

Not enabled.

Required Privilege Level

routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- *Understanding BPDU Protection for Spanning-Tree Instance Interfaces*
- *Understanding BPDU Protection for STP, RSTP, and MSTP*
- *BPDU Protection on All Edge Ports of the Bridge*
- *Configuring BPDU Protection on ACX Router, EX Switch and MX Router Edge Ports*
- *Configuring BPDU Protection on Switch Spanning Tree Interfaces*
- *Example: Configuring Network Regions for VLANs with MSTP on non-ELS EX Series Switches*
- *Example: Faster Convergence and Improved Network Stability with RSTP on EX Series Switches*

Copyright © 2018, Juniper Networks, Inc.

- *Example: Configuring BPDU Protection on Edge Interfaces to Prevent STP Miscalculations on non-ELS EX Series Switches*
- rstp on page 178
- *mstp*
- vstp on page 190
- show spanning-tree bridge on page 267
- show spanning-tree interface on page 273
- *clear ethernet-switching bpdu-error interface*

bpdu-timeout-action

Syntax

```
bpdu-timeout-action {
  block;
  log;
}
```

Hierarchy Level

[edit logical-systems *logical-system-name* protocols (mstp | rstp | vstp)],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols
(mstp | rstp | vstp)],

Switches supporting ELS

[edit protocols mstp interface (all | *interface-name*)],
[edit protocols mstp interface (all | *interface-name*) arp-on-stp],
[edit protocols rstp interface (all | *interface-name*)],
[edit protocols rstp interface (all | *interface-name*) arp-on-stp],
[edit protocols stp interface (all | *interface-name*)],
[edit protocols stp interface (all | *interface-name*) arp-on-stp],
[edit protocols vstp vlan *vlan-id* interface (all | *interface-name*)]
[edit protocols vstp vlan *vlan-id* interface (all | *interface-name*) arp-on-stp]

Release Information

Statement introduced in Junos OS Release 9.1 for EX Series switches.
Statement introduced in Junos OS Release 9.4 for other devices.
Statement updated in Junos OS Release 9.4 for EX Series switches to add VSTP support.
Support for logical systems added in Junos OS Release 9.6.
Statement introduced in Junos OS Release 17.1 for ACX Series routers.

Description

Provide STP loop protection for a given STP family protocol interface. Configure the BPDU timeout action on a specific interface. You must configure at least one action (**log**, **block**, or both).

The remaining statements are explained separately. See CLI Explorer.

Default

If the **bpdu-timeout-action** statement is not configured, an interface that stops receiving BPDUs will transition to the designated port (forwarding) state, creating a potential loop.

Options

log—The interface logs the fact that it has not received BPDUs during the timeout interval.

block—The interface is blocked and the fact that the interface has not received BPDUs during the timeout interval is logged.

Required Privilege Level

routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- *Understanding Loop Protection for Spanning-Tree Instance Interfaces*
- *Understanding VSTP*

 Copyright © 2018, Juniper Networks, Inc.

- *Configuring Loop Protection for a Spanning-Tree Instance Interface*
- *Example: Enabling Loop Protection for Spanning-Tree Protocols*
- *Example: Configuring Network Regions for VLANs with MSTP on non-ELS EX Series Switches*
- *Example: Faster Convergence and Improved Network Stability with RSTP on EX Series Switches*
- *Example: Configuring Loop Protection to Prevent Interfaces from Transitioning from Blocking to Forwarding in a Spanning Tree on non-ELS EX Series Switches*
- show spanning-tree bridge on page 267
- show spanning-tree interface on page 273
- rstp on page 178
- *mstp*
- vstp on page 190

Copyright © 2018, Juniper Networks, Inc.

bridge-priority

Syntax

bridge-priority *priority*;

Hierarchy Level

[edit logical-systems *logical-system-name* protocols (mstp | rstp)],
[edit logical-systems *logical-system-name* protocols mstp msti *msti-id*],
[edit logical-systems *logical-system-name* protocols vstp vlan *vlan-id*],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols (mstp | rstp)],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols mstp msti *msti-id*],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols vstp vlan *vlan-id*],

[edit routing-instances *routing-instance-name* protocols (mstp | rstp)],
[edit routing-instances *routing-instance-name* protocols mstp msti *msti-id*],
[edit routing-instances *routing-instance-name* protocols vstp vlan *vlan-id*]

[edit protocols mstp],
[edit protocols mstp msti *msti-id*],
[edit protocols rstp],
[edit protocols stp],
[edit protocols vstp vlan *vlan-id*]

Release Information

Statement introduced in Junos OS Release 8.4.
Statement introduced in Junos OS Release 9.0 for EX Series switches.
Statement updated in Junos OS Release 9.4 for EX Series switches to add VSTP support.
Support for logical systems added in Junos OS Release 9.6.
Statement introduced in Junos OS Release 13.2X50-D10 for EX Series switches.

Description

Configures the bridge priority, which determines which bridge is elected as the root bridge. If two bridges have the same path cost to the root bridge, the bridge priority determines which bridge becomes the designated bridge for a LAN segment.

Default

32,768

Options

priority—The bridge priority can be set only in increments of 4096.

Range: 0 through 61,440

Default: 32,768

Required Privilege Level

routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- *Understanding MSTP*
- *Understanding VSTP*

 Copyright © 2018, Juniper Networks, Inc.

- *Understanding Bridge Priority for Election of Root Bridge and Designated Bridge*
- *Example: Configuring Network Regions for VLANs with MSTP on Switches*
- *Example: Configuring Network Regions for VLANs with MSTP on non-ELS EX Series Switches*
- show spanning-tree bridge on page 267
- show spanning-tree interface on page 273

configuration-name

Syntax

```
configuration-name configuration-name;
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols mstp],
[edit logical-systems logical-system-name routing-instances routing-instance-name
  protocols mstp],
```

```
[edit protocols mstp],
```

```
[edit routing-instances routing-instance-name protocols mstp]
```

Release Information

Statement introduced in Junos OS Release 8.4.
Statement introduced in Junos OS Release 9.0 for EX Series switches.
Support for logical systems added in Junos OS Release 9.6.
Statement introduced in Junos OS Release 13.2X50-D10 for EX Series switches.

Description

Specify the configuration name, which is the MSTP region name carried in the MSTP BPDUs.

Required Privilege Level

routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- *Understanding MSTP*
- *Understanding BPDUs Used for Exchanging Information Among Bridges*
- *Configuring Multiple Spanning Tree Protocol*
- *Configuring MSTP on Switches*
- *Example: Configuring Network Regions for VLANs with MSTP on non-ELS EX Series Switches*
- *Example: Faster Convergence and Improved Network Stability with RSTP on EX Series Switches*
- show spanning-tree bridge on page 267
- show spanning-tree interface on page 273
- *Example: Configuring Network Regions for VLANs with MSTP on Switches*
- *Example: Configuring Faster Convergence and Network Stability on ELS Switches with RSTP*
- *Understanding MSTP*

 Copyright © 2018, Juniper Networks, Inc.

cost

Syntax

cost *cost*;

Hierarchy Level

[edit logical-systems *logical-system-name* protocols (mstp | rstp | vstp) interface *interface-name*],
[edit logical-systems *logical-system-name* protocols mstp msti *msti-id* interface *interface-name*],
[edit logical-systems *logical-system-name* protocols vstp vlan *vlan-id* interface *interface-name*],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols (mstp | rstp | vstp) interface *interface-name*],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols mstp msti *msti-id* interface *interface-name*],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols vstp vlan *vlan-id* interface *interface-name*],

[edit protocols (mstp | rstp | vstp) interface *interface-name*],
[edit protocols mstp msti *msti-id* interface *interface-name*],
[edit protocols vstp vlan *vlan-id* interface *interface-name*],

[edit routing-instances *routing-instance-name* protocols (mstp | rstp | vstp) interface *interface-name*],
[edit routing-instances *routing-instance-name* protocols mstp msti *msti-id* interface *interface-name*],
[edit routing-instances *routing-instance-name* protocols vstp vlan *vlan-id* interface *interface-name*]

Release Information

Statement introduced in Junos OS Release 8.4.
Statement introduced in Junos OS Release 9.0 for EX Series switches.
Statement updated in Junos OS Release 9.4 for EX Series switches to add VSTP support.
Support for logical systems added in Junos OS Release 9.6.

Description

For Spanning Tree Protocol (STP), Rapid Spanning Tree Protocol (RSTP), VLAN Spanning Tree Protocol (VSTP), or Multiple Spanning Tree Protocol (MSTP), configure the link cost to control which bridge is the designated bridge and which interface is the designated interface.

Default

The link cost is determined by the link speed.

Options

cost—(Optional) Link cost associated with the port.

Range: 1 through 200,000,000

Required Privilege Level

routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Copyright © 2018, Juniper Networks, Inc.

Related Documentation

- *Understanding Spanning-Tree Instance Interfaces*
- *Understanding STP*
- *Understanding RSTP*
- *Understanding MSTP*
- *Understanding VSTP*
- show spanning-tree bridge on page 267
- show spanning-tree interface on page 273

 Copyright © 2018, Juniper Networks, Inc.

disable

Syntax

```
disable;
```

Hierarchy Level

[edit logical-systems *logical-system-name* protocols mstp],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name*
protocols mstp],

[edit protocols mpls],
[edit protocols mpls interface (all | *interface-name*)],

[edit protocols mstp],
[edit protocols mstp (all | *interface-name*)],
[edit protocols mstp (all | *interface-name*) arp-on-stp],
[edit protocols mstp msti *msti-id* vlan (*vlan-id* | *vlan-name*) interface (all | *interface-name*)],
[edit protocols mstp msti *msti-id* vlan (*vlan-id* | *vlan-name*) interface *interface-name*
arp-on-stp],
[edit protocols rstp],
[edit protocols rstp interface (all | *interface-name*)],
[edit protocols rstp interface (all | *interface-name*) arp-on-stp],
[edit protocols stp],
[edit protocols stp interface (all | *interface-name*)],
[edit protocols stp interface (all | *interface-name*) arp-on-stp],
[edit protocols vstp],
[edit protocols vstp vlan *vlan-id* interface (all | *interface-name*)]
[edit protocols vstp vlan *vlan-id* interface (all | *interface-name*) arp-on-stp]
[edit protocols mstp interface *interface-name*,
[edit protocols mstp msti *msti-id* vlan (all *vlan-id* | *vlan-name*) interface *interface-name*],
[edit protocols rstp interface *interface-name*],
[edit protocols vstp vlan *vlan-id* interface *interface-name*]

[edit routing-instances *routing-instance-name* protocols mstp]

[edit ethernet-switching-options bpdu-block interface]

Release Information

Statement introduced in Junos OS Release 9.0 for EX Series switches.
Statement introduced in Junos OS Release 9.1.
Statement updated in Junos OS Release 9.4 for EX Series switches to add VSTP support.
Statement introduced in Junos OS Release 12.2 for EX Series switches.
Statement introduced in Junos OS Release 13.2X50-D10 for EX Series switches.
Support for logical systems added in Junos OS Release 9.6.
Statement updated in Junos OS Release 15.1 for EX Series switches.

Description

Disable the entire MSTP, RSTP, or VSTP instance on a specific interface.

NOTE: **You cannot disable spanning tree parameters for all interfaces.**

Copyright © 2018, Juniper Networks, Inc.

Default Not enabled

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- *Understanding RSTP*
- *Understanding VSTP*
- *Configuring Multiple Spanning Tree Protocol*
- *Configuring RSTP on EX Series Switches (CLI Procedure)*
- *Configuring MSTP on Switches*
- *Disabling MSTP*
- *Example: Configuring BPDU Protection on Interfaces to Prevent STP Miscalculations on EX Series Switches*
- *Example: Configuring Network Regions for VLANs with MSTP on Switches*
- *Example: Configuring Faster Convergence and Network Stability on ELS Switches with RSTP*
- *interface*
- show spanning-tree bridge on page 267
- show spanning-tree interface on page 273

 Copyright © 2018, Juniper Networks, Inc.

bpdu-block

Syntax

```
bpdu-block {
  interface (interface-name disable | all);
  disable-timeout seconds;
}
```

Hierarchy Level

```
[edit protocols layer2-control ]
```

```
[edit ethernet-switching-options]
```

Release Information

Statement introduced in Junos OS Release 9.1 for EX Series switches.
Statement introduced in Junos OS Release 13.2X50-D10 for EX Series switches.

Description

Configures bridge protocol data unit (BPDU) protection on a specified interface or on all interfaces. If the interface receives incompatible BPDUs, it is disabled.

If the disable-timeout statement is included in the BPDU configuration, the interface is automatically reenabled after the timer expires. Otherwise, you must use the operational mode command **clear ethernet-switching bpdu-error interface** to unblock and reenable the interface.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- *Understanding BPDU Protection for Spanning-Tree Instance Interfaces*
- *Configuring BPDU Protection for Individual Spanning-Tree Instance Interfaces*
- *Understanding BPDU Protection for STP, RSTP, and MSTP*
- *Example: Configuring BPDU Protection on Interfaces to Prevent STP Miscalculations on EX Series Switches*
- *Unblocking an Interface on non-ELS EX Series Switches That Receives BPDUs in Error (CLI Procedure)*
- *Understanding BPDU Protection for STP, RSTP, and MSTP*
- show spanning-tree bridge on page 267
- show spanning-tree interface on page 273
- *clear ethernet-switching bpdu-error interface*

bpdu-block-on-edge

Syntax

```
bpdu-block-on-edge;
```

Hierarchy Level

[edit logical-systems *logical-system-name* protocols (mstp | rstp | vstp)],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols (mstp | rstp | vstp)],
[edit protocols (mstp | rstp| vstp)],
[edit routing-instances *routing-instance-name* protocols (mstp | rstp | vstp)]

Release Information

Statement introduced in Junos OS Release 9.1 for EX Series switches.
Statement introduced in Junos OS Release 9.4 for additional devices.
Statement updated in Junos OS Release 9.4 for EX Series switches to add VSTP support.
Statement updated in Junos OS Release 11.1 for EX Series switches to change blocking behavior to port shutdown.
Support for logical systems added in Junos OS Release 9.6.
Statement introduced in Junos OS Release 17.1 for ACX Series routers.

Description

Configure bridge protocol data unit (BPDU) protection on all edge ports of a switch. When the **bpdu-block-on-edge** statement is configured and the interface encounters an incompatible BPDU, the interface shuts down.

If the disable-timeout statement is included in the BPDU configuration, the interface is automatically reenabled after the timer expires. Otherwise, you must use the operational mode command **clear ethernet-switching bpdu-error interface** to unblock and reenable the interface.

Default

Not enabled.

Required Privilege Level

routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- *Understanding BPDU Protection for Spanning-Tree Instance Interfaces*
- *Understanding BPDU Protection for STP, RSTP, and MSTP*
- *BPDU Protection on All Edge Ports of the Bridge*
- *Configuring BPDU Protection on ACX Router, EX Switch and MX Router Edge Ports*
- *Configuring BPDU Protection on Switch Spanning Tree Interfaces*
- *Example: Configuring Network Regions for VLANs with MSTP on non-ELS EX Series Switches*
- *Example: Faster Convergence and Improved Network Stability with RSTP on EX Series Switches*

 Copyright © 2018, Juniper Networks, Inc.

- *Example: Configuring BPDU Protection on Edge Interfaces to Prevent STP Miscalculations on non-ELS EX Series Switches*
- rstp on page 178
- *mstp*
- vstp on page 190
- show spanning-tree bridge on page 267
- show spanning-tree interface on page 273
- *clear ethernet-switching bpdu-error interface*

Copyright © 2018, Juniper Networks, Inc.

bpdu-timeout-action

Syntax

```
bpdu-timeout-action {
  block;
  log;
}
```

Hierarchy Level

[edit logical-systems *logical-system-name* protocols (mstp | rstp | vstp)],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols (mstp | rstp | vstp)],

Switches supporting ELS

[edit protocols mstp interface (all | *interface-name*)],
[edit protocols mstp interface (all | *interface-name*) arp-on-stp],
[edit protocols rstp interface (all | *interface-name*)],
[edit protocols rstp interface (all | *interface-name*) arp-on-stp],
[edit protocols stp interface (all | *interface-name*)],
[edit protocols stp interface (all | *interface-name*) arp-on-stp],
[edit protocols vstp vlan *vlan-id* interface (all | *interface-name*)]
[edit protocols vstp vlan *vlan-id* interface (all | *interface-name*) arp-on-stp]

Release Information

Statement introduced in Junos OS Release 9.1 for EX Series switches.
Statement introduced in Junos OS Release 9.4 for other devices.
Statement updated in Junos OS Release 9.4 for EX Series switches to add VSTP support.
Support for logical systems added in Junos OS Release 9.6.
Statement introduced in Junos OS Release 17.1 for ACX Series routers.

Description

Provide STP loop protection for a given STP family protocol interface. Configure the BPDU timeout action on a specific interface. You must configure at least one action (**log**, **block**, or both).

The remaining statements are explained separately. See CLI Explorer.

Default

If the **bpdu-timeout-action** statement is not configured, an interface that stops receiving BPDUs will transition to the designated port (forwarding) state, creating a potential loop.

Options

log—The interface logs the fact that it has not received BPDUs during the timeout interval.

block—The interface is blocked and the fact that the interface has not received BPDUs during the timeout interval is logged.

Required Privilege Level

routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- *Understanding Loop Protection for Spanning-Tree Instance Interfaces*
- *Understanding VSTP*

Copyright © 2018, Juniper Networks, Inc.

- *Configuring Loop Protection for a Spanning-Tree Instance Interface*
- *Example: Enabling Loop Protection for Spanning-Tree Protocols*
- *Example: Configuring Network Regions for VLANs with MSTP on non-ELS EX Series Switches*
- *Example: Faster Convergence and Improved Network Stability with RSTP on EX Series Switches*
- *Example: Configuring Loop Protection to Prevent Interfaces from Transitioning from Blocking to Forwarding in a Spanning Tree on non-ELS EX Series Switches*
- show spanning-tree bridge on page 267
- show spanning-tree interface on page 273
- rstp on page 178
- *mstp*
- vstp on page 190

bridge-priority

Syntax

bridge-priority *priority*;

Hierarchy Level

[edit logical-systems *logical-system-name* protocols (mstp | rstp)],
[edit logical-systems *logical-system-name* protocols mstp msti *msti-id*],
[edit logical-systems *logical-system-name* protocols vstp vlan *vlan-id*],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols (mstp | rstp)],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols mstp msti *msti-id*],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols vstp vlan *vlan-id*],

[edit routing-instances *routing-instance-name* protocols (mstp | rstp)],
[edit routing-instances *routing-instance-name* protocols mstp msti *msti-id*],
[edit routing-instances *routing-instance-name* protocols vstp vlan *vlan-id*]

[edit protocols mstp],
[edit protocols mstp msti *msti-id*],
[edit protocols rstp],
[edit protocols stp],
[edit protocols vstp vlan *vlan-id*]

Release Information

Statement introduced in Junos OS Release 8.4.
Statement introduced in Junos OS Release 9.0 for EX Series switches.
Statement updated in Junos OS Release 9.4 for EX Series switches to add VSTP support.
Support for logical systems added in Junos OS Release 9.6.
Statement introduced in Junos OS Release 13.2X50-D10 for EX Series switches.

Description

Configures the bridge priority, which determines which bridge is elected as the root bridge. If two bridges have the same path cost to the root bridge, the bridge priority determines which bridge becomes the designated bridge for a LAN segment.

Default

32,768

Options

priority—The bridge priority can be set only in increments of 4096.

Range: 0 through 61,440

Default: 32,768

Required Privilege Level

routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- *Understanding MSTP*
- *Understanding VSTP*

Copyright © 2018, Juniper Networks, Inc.

- *Understanding Bridge Priority for Election of Root Bridge and Designated Bridge*
- *Example: Configuring Network Regions for VLANs with MSTP on Switches*
- *Example: Configuring Network Regions for VLANs with MSTP on non-ELS EX Series Switches*
- show spanning-tree bridge on page 267
- show spanning-tree interface on page 273

Copyright © 2018, Juniper Networks, Inc.

configuration-name

Syntax

```
configuration-name configuration-name;
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols mstp],
[edit logical-systems logical-system-name routing-instances routing-instance-name
  protocols mstp],
```

```
[edit protocols mstp],
```

```
[edit routing-instances routing-instance-name protocols mstp]
```

Release Information

Statement introduced in Junos OS Release 8.4.
Statement introduced in Junos OS Release 9.0 for EX Series switches.
Support for logical systems added in Junos OS Release 9.6.
Statement introduced in Junos OS Release 13.2X50-D10 for EX Series switches.

Description

Specify the configuration name, which is the MSTP region name carried in the MSTP BPDUs.

Required Privilege Level

routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- *Understanding MSTP*
- *Understanding BPDUs Used for Exchanging Information Among Bridges*
- *Configuring Multiple Spanning Tree Protocol*
- *Configuring MSTP on Switches*
- *Example: Configuring Network Regions for VLANs with MSTP on non-ELS EX Series Switches*
- *Example: Faster Convergence and Improved Network Stability with RSTP on EX Series Switches*
- show spanning-tree bridge on page 267
- show spanning-tree interface on page 273
- *Example: Configuring Network Regions for VLANs with MSTP on Switches*
- *Example: Configuring Faster Convergence and Network Stability on ELS Switches with RSTP*
- *Understanding MSTP*

 Copyright © 2018, Juniper Networks, Inc.

cost

Syntax

```
cost cost;
```

Hierarchy Level

[edit logical-systems *logical-system-name* protocols (mstp | rstp | vstp) interface *interface-name*],
[edit logical-systems *logical-system-name* protocols mstp msti *msti-id* interface *interface-name*],
[edit logical-systems *logical-system-name* protocols vstp vlan *vlan-id* interface *interface-name*],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols (mstp | rstp | vstp) interface *interface-name*],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols mstp msti *msti-id* interface *interface-name*],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols vstp vlan *vlan-id* interface *interface-name*],

[edit protocols (mstp | rstp | vstp) interface *interface-name*],
[edit protocols mstp msti *msti-id* interface *interface-name*],
[edit protocols vstp vlan *vlan-id* interface *interface-name*],

[edit routing-instances *routing-instance-name* protocols (mstp | rstp | vstp) interface *interface-name*],
[edit routing-instances *routing-instance-name* protocols mstp msti *msti-id* interface *interface-name*],
[edit routing-instances *routing-instance-name* protocols vstp vlan *vlan-id* interface *interface-name*]

Release Information

Statement introduced in Junos OS Release 8.4.
Statement introduced in Junos OS Release 9.0 for EX Series switches.
Statement updated in Junos OS Release 9.4 for EX Series switches to add VSTP support.
Support for logical systems added in Junos OS Release 9.6.

Description

For Spanning Tree Protocol (STP), Rapid Spanning Tree Protocol (RSTP), VLAN Spanning Tree Protocol (VSTP), or Multiple Spanning Tree Protocol (MSTP), configure the link cost to control which bridge is the designated bridge and which interface is the designated interface.

Default

The link cost is determined by the link speed.

Options

cost—(Optional) Link cost associated with the port.

Range: 1 through 200,000,000

Required Privilege Level

routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Copyright © 2018, Juniper Networks, Inc.

Related Documentation

- *Understanding Spanning-Tree Instance Interfaces*
- *Understanding STP*
- *Understanding RSTP*
- *Understanding MSTP*
- *Understanding VSTP*
- show spanning-tree bridge on page 267
- show spanning-tree interface on page 273

 Copyright © 2018, Juniper Networks, Inc.

disable

Syntax

```
disable;
```

Hierarchy Level

[edit logical-systems *logical-system-name* protocols mstp],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name*
  protocols mstp],

[edit protocols mpls],
[edit protocols mpls interface (all | *interface-name*)],

[edit protocols mstp],
[edit protocols mstp (all | *interface-name*)],
[edit protocols mstp (all | *interface-name*) arp-on-stp],
[edit protocols mstp msti *msti-id* vlan (*vlan-id* | *vlan-name*) interface (all | *interface-name*)],
[edit protocols mstp msti *msti-id* vlan (*vlan-id* | *vlan-name*) interface *interface-name*
  arp-on-stp],
[edit protocols rstp],
[edit protocols rstp interface (all | *interface-name*)],
[edit protocols rstp interface (all | *interface-name*) arp-on-stp],
[edit protocols stp],
[edit protocols stp interface (all | *interface-name*)],
[edit protocols stp interface (all | *interface-name*) arp-on-stp],
[edit protocols vstp],
[edit protocols vstp vlan *vlan-id* interface (all | *interface-name*)]
[edit protocols vstp vlan *vlan-id* interface (all | *interface-name*) arp-on-stp]
[edit protocols mstp interface *interface-name*,
[edit protocols mstp msti *msti-id* vlan (all *vlan-id* | *vlan-name*) interface *interface-name*],
[edit protocols rstp interface *interface-name*],
[edit protocols vstp vlan *vlan-id* interface *interface-name*]

[edit routing-instances *routing-instance-name* protocols mstp]

[edit ethernet-switching-options bpdu-block interface]

Release Information

Statement introduced in Junos OS Release 9.0 for EX Series switches.
Statement introduced in Junos OS Release 9.1.
Statement updated in Junos OS Release 9.4 for EX Series switches to add VSTP support.
Statement introduced in Junos OS Release 12.2 for EX Series switches.
Statement introduced in Junos OS Release 13.2X50-D10 for EX Series switches.
Support for logical systems added in Junos OS Release 9.6.
Statement updated in Junos OS Release 15.1 for EX Series switches.

Description

Disable the entire MSTP, RSTP, or VSTP instance on a specific interface.

NOTE: **You cannot disable spanning tree parameters for all interfaces.**

Default Not enabled

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- *Understanding RSTP*
- *Understanding VSTP*
- *Configuring Multiple Spanning Tree Protocol*
- *Configuring RSTP on EX Series Switches (CLI Procedure)*
- *Configuring MSTP on Switches*
- *Disabling MSTP*
- *Example: Configuring BPDU Protection on Interfaces to Prevent STP Miscalculations on EX Series Switches*
- *Example: Configuring Network Regions for VLANs with MSTP on Switches*
- *Example: Configuring Faster Convergence and Network Stability on ELS Switches with RSTP*
- *interface*
- show spanning-tree bridge on page 267
- show spanning-tree interface on page 273

Copyright © 2018, Juniper Networks, Inc.

disable-timeout

Syntax

```
disable-timeout seconds;
```

Hierarchy Level

[edit protocols layer2-control bpdu-block]

[edit ethernet-switching-options bpdu-block]

Release Information

Statement introduced in Junos OS Release 9.1 for EX Series switches.
Statement introduced in Junos OS Release 9.4.

Description

For interfaces configured for BPDU protection, specify the amount of time an interface receiving BPDUs is disabled.

Configure the timeout value to periodically check to see if an interface is still disabled with BPDU blocking. If this option is not configured, the interface is not periodically checked and remains disabled.

Default

The disable timeout is not enabled.

Options

seconds——Amount of time, in seconds, the interface receiving BPDUs is disabled. Once the timeout expires, the interface is brought back into service.

Range: 10 through 3600

Default: If this option is not configured, the interface is not periodically checked and remains disabled.

Required Privilege Level

routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.
system—To view this statement in the configuration.
system-control—To add this statement to the configuration.

Related Documentation

- *Understanding BPDU Protection for Spanning-Tree Instance Interfaces*
- *Understanding BPDU Protection for STP, RSTP, and MSTP*
- *Configuring BPDU Protection for Individual Spanning-Tree Instance Interfaces*
- *Example: Configuring Network Regions for VLANs with MSTP on non-ELS EX Series Switches*
- *Example: Faster Convergence and Improved Network Stability with RSTP on EX Series Switches*
- *Example: Configuring BPDU Protection on Interfaces to Prevent STP Miscalculations on EX Series Switches*
- *Example: Configuring Network Regions for VLANs with MSTP on Switches*

Copyright © 2018, Juniper Networks, Inc.

- *Example: Configuring BPDU Protection on Switch Edge Interfaces With ELS to Prevent STP Miscalculations*
- show spanning-tree bridge on page 267
- show spanning-tree interface on page 273

 Copyright © 2018, Juniper Networks, Inc.

edge

Syntax

```
edge;
```

Hierarchy Level

[edit logical-systems *logical-system-name* protocols (mstp | rstp | vstp) interface *interface-name*],
[edit logical-systems *logical-system-name* protocols mstp msti *msti-id* interface *interface-name*],
[edit logical-systems *logical-system-name* protocols vstp vlan *vlan-id* interface *interface-name*],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols (mstp | rstp | vstp) interface *interface-name*],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols mstp msti *msti-id* interface *interface-name*],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols vstp vlan *vlan-id* interface *interface-name*],

[edit protocols (mstp | rstp | vstp) interface (all | *interface-name*)],)arp-on-stp],
[edit protocols mstp msti *msti-id* interface (all | *interface-name*)],)arp-on-stp],
[edit protocols vstp vlan *vlan-id* interface (all | *interface-name*)],)arp-on-stp],

[edit routing-instances *routing-instance-name* protocols (mstp | rstp | vstp) interface *interface-name*]
[edit routing-instances *routing-instance-name* protocols mstp msti *msti-id* interface *interface-name*],
[edit routing-instances *routing-instance-name* protocols vstp vlan *vlan-id* interface *interface-name*]

Release Information

Statement introduced in Junos OS Release 8.4.
Statement introduced in Junos OS Release 9.0 for EX Series switches.
Statement updated in Junos OS Release 9.4 for EX Series switches to add VSTP support.
Support for logical systems added in Junos OS Release 9.6.
Statement introduced in Junos OS Release 13.2X50-D10 for EX Series switches.

Description

For Rapid Spanning Tree Protocol (RSTP), VLAN Spanning Tree Protocol (VSTP), or Multiple Spanning Tree Protocol (MSTP), configure interfaces as edge ports or edge interfaces. Edge ports do not expect to receive BPDUs. If a BPDU is received, the port becomes a nonedge port and the Edge interfaces immediately transition to a forwarding state.

NOTE: Although the **edge** configuration statement appears in the **[edit protocols stp interface (all |** ***interface-name*****)]** or **[edit protocols rstp force-version stp interface (all |** ***interface-name*****)]** hierarchy on the switch, this statement has no effect on the switch operation if you configure it.

Copyright © 2018, Juniper Networks, Inc.

Default Edge interfaces are not enabled.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- *Understanding STP*
- *Understanding VSTP*
- *Understanding Spanning-Tree Instance Interfaces*
- *Configuring a Spanning-Tree Instance Interface as an Edge Port for Faster Convergence*
- *Example: Configuring Network Regions for VLANs with MSTP on Switches*
- *Example: Configuring Faster Convergence and Network Stability on ELS Switches with RSTP*
- *Example: Configuring Network Regions for VLANs with MSTP on non-ELS EX Series Switches*
- *Example: Faster Convergence and Improved Network Stability with RSTP on EX Series Switches*
- show spanning-tree bridge on page 267
- show spanning-tree interface on page 273

force-version (IEEE 802.1D STP)

Syntax

```
force-version stp;
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols (rstp | vstp)],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols
  (rstp | vstp)],
```

```
[edit protocols (rstp | vstp)],
[edit routing-instances routing-instance-name protocols (rstp | vstp)]
```

Release Information Statement introduced in Junos OS Release 8.4.
Statement introduced in Junos OS Release 9.4 for EX Series switches.
Support for logical systems added in Junos OS Release 9.6.
Statement introduced in Junos OS Release 13.2X50-D10 for EX Series switches.
Support for logical systems added in Junos OS Release 9.6.

Description Forces the spanning-tree version to be the original IEEE 803.1D STP protocol.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

 Copyright © 2018, Juniper Networks, Inc.

forward-delay

Syntax

```
forward-delay seconds;
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols (mstp | rstp)],
[edit logical-systems logical-system-name protocols vstp vlan vlan-id],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols
  (mstp | rstp)],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols
  vstp vlan vlan-id],
```

```
[edit protocols (mstp | rstp)],
[edit protocols vstp vlan vlan-id],
[edit routing-instances routing-instance-name protocols (mstp | rstp)],
```

```
[edit routing-instances routing-instance-name protocols vstp vlan vlan-id]
```

Release Information

Statement introduced in Junos OS Release 8.4.
Statement introduced in Junos OS Release 9.0 for EX Series switches.
Statement updated in Junos OS Release 9.4 for EX Series switches to add VSTP support.
Support for logical systems added in Junos OS Release 9.6.
Statement introduced in Junos OS Release 13.2X50-D10 for EX Series switches.

Description

For Spanning Tree Protocol (STP), Rapid Spanning Tree Protocol (RSTP), VLAN Spanning Tree Protocol (VSTP), or Multiple Spanning Tree Protocol (MSTP), specifies the length of time an STP bridge port remains in the listening and learning states before transitioning to the forwarding state.

Options

seconds—(Optional) Number of seconds the bridge port remains in the listening and learning states.

Range: 4 through 30

Default: 15 seconds

Required Privilege Level

routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- *Understanding Forward Delay Before Ports Transition to Forwarding State*
- *Understanding STP*
- *Understanding MSTP*
- *Understanding VSTP*
- *Example: Configuring Network Regions for VLANs with MSTP on Switches*

- *Example: Configuring Faster Convergence and Network Stability on ELS Switches with RSTP*
- *Example: Faster Convergence and Improved Network Stability with RSTP on EX Series Switches*
- *Example: Configuring Network Regions for VLANs with MSTP on non-ELS EX Series Switches*
- show spanning-tree bridge on page 267
- show spanning-tree interface on page 273

Copyright © 2018, Juniper Networks, Inc.

hello-time

Syntax

```
hello-time seconds;
```

Hierarchy Level

[edit logical-systems *logical-system-name* protocols (mstp | rstp)],
[edit logical-systems *logical-system-name* protocols vstp vlan *vlan-id*],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols (mstp | rstp)],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols vstp vlan *vlan-id*],

[edit protocols (mstp | rstp)],
[edit protocols vstp vlan *vlan-id*],

[edit routing-instances *routing-instance-name* protocols (mstp | rstp)],
[edit routing-instances *routing-instance-name* protocols vstp vlan *vlan-id*]

Release Information

Statement introduced in Junos OS Release 8.4.
Statement introduced in Junos OS Release 9.0 for EX Series switches.
Statement updated in Junos OS Release 9.4 for EX Series switches to add VSTP support.
Support for logical systems added in Junos OS Release 9.6.
Statement introduced in Junos OS Release 13.2X50-D10 for EX Series switches.
Support for logical systems added in Junos OS Release 9.6.

Description

For Spanning Tree Protocol (STP), Rapid Spanning Tree Protocol (RSTP), VLAN Spanning Tree Protocol (VSTP), or Multiple Spanning Tree Protocol (MSTP), specifies the number of seconds between transmissions of configuration BPDUs by the root bridge.

Default

2 seconds

Options

seconds—(Optional) Number of seconds between transmissions of configuration BPDUs.

Range: 1 through 10

Default: 2 seconds

Required Privilege Level

routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- *Understanding MSTP*
- *Understanding VSTP*
- *Example: Configuring Network Regions for VLANs with MSTP on Switches*
- *Example: Configuring Network Regions for VLANs with MSTP on non-ELS EX Series Switches*

Copyright © 2018, Juniper Networks, Inc.

- *Example: Configuring Faster Convergence and Network Stability on ELS Switches with RSTP*
- *Example: Faster Convergence and Improved Network Stability with RSTP on EX Series Switches*
- show spanning-tree bridge on page 267
- show spanning-tree interface on page 273

interface (BPDU)

Syntax

```
interface (all | [interface-name]) {
  drop;
}
```

Hierarchy Level

- For platforms with ELS CLI:

  `[edit protocols layer2-control]`

- For platforms with Original CLI:

  `[edit ethernet-switching-options]`

Release Information Statement introduced in Junos OS Release 11.1 for the QFX Series.

Description Apply BPDU protection to all interfaces or one or more interfaces.

Options **all**—All interfaces.

interface-name—Name of the interface.

drop—Drops xSTP BPDUs.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- *Example: Configuring Network Regions for VLANs with MSTP*
- *Example: Configuring Faster Convergence and Improving Network Stability with RSTP*
- *Understanding BPDU Protection for STP, RSTP, and MSTP*
- show spanning-tree bridge on page 267
- show spanning-tree interface on page 273

 Copyright © 2018, Juniper Networks, Inc.

interface (Spanning Tree)

Syntax

```
interface interface-name {
  bpdu-timeout-action {
    alarm;
    block;
  }
  cost cost;
  edge;
  mode (p2p | shared);
  no-root-port;
  priority interface-priority;
}
```

Syntax

```
interface interface-name {
  arp-on-stp;
  bpdu-timeout-action
    block;
    log;
  cost cost;
  disable;
  edge;
  mode mode;
  no-root-port;
  priority priority;
}
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols (mstp | rstp | vstp)],
[edit logical-systems logical-system-name protocols vstp vlan vlan-id],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols
  (mstp | rstp | vstp)],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols
  vstp vlan vlan-id],
```

```
[edit protocols (mstp | rstp | vstp)],
[edit protocols vstp vlan vlan-id],
```

```
[edit protocols (mstp | rstp | vstp)],
[edit protocols vstp vlan vlan-id],
[edit protocols vstp vlan-group group group-name vlan (vlan-id | vlan-range
  |open-set-of-values)
```

```
[edit routing-instances routing-instance-name protocols (mstp | rstp | vstp)],
[edit routing-instances routing-instance-name protocols vstp vlan vlan-id]
```

Release Information

Statement introduced in Junos OS Release 8.4.
Statement introduced in Junos OS Release 9.0 for EX Series switches.
Statement updated in Junos OS Release 9.4 for EX Series switches to add VSTP support.

Copyright © 2018, Juniper Networks, Inc.

Support for logical systems added in Junos OS Release 9.6.
Statement introduced in Junos OS Release 11.1 for the QFX Series.
Statement introduced in Junos OS Release 13.2X50-D10 for EX Series switches.
Statement updated in Junos OS Release 15.1 for EX Series switches to support configuration of spanning tree parameters globally on all interfaces.

NOTE: You cannot disable spanning tree parameters globally on all interfaces.

Description Configures the interface to participate in the RSTP, MSTP, or VSTP instance.

The **edge**, **mode**, and **no-root-port** options are not available at the [**edit protocols mstp msti** ***msti-id***] hierarchy level.

Options ***interface-name***—Name of a Gigabit Ethernet or 10-Gigabit Ethernet interface.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- *Understanding Spanning-Tree Instance Interfaces*
- *Understanding RSTP*
- *Understanding MSTP*
- *Understanding VSTP*
- *Configuring RSTP on EX Series Switches (CLI Procedure)*
- *Configuring MSTP on Switches*
- *Example: Configuring Network Regions for VLANs with MSTP on non-ELS EX Series Switches*
- *Example: Faster Convergence and Improved Network Stability with RSTP on EX Series Switches*
- show spanning-tree bridge on page 267
- show spanning-tree interface on page 273

 Copyright © 2018, Juniper Networks, Inc.

max-age

Syntax

```
max-age seconds;
```

Hierarchy Level

[edit logical-systems *logical-system-name* protocols (mstp | rstp)],
[edit logical-systems *logical-system-name* protocols vstp vlan *vlan-id*],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols (mstp | rstp)],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols vstp vlan *vlan-id*],

[edit protocols mstp],
[edit protocols rstp],
[edit protocols stp],
[edit protocols vstp vlan *vlan-id*]

[edit routing-instances *routing-instance-name* protocols (mstp | rstp)],
[edit routing-instances *routing-instance-name* protocols vstp vlan *vlan-id*]

Release Information

Statement introduced in Junos OS Release 8.4.
Statement introduced in Junos OS Release 9.0 for EX Series switches.
Statement updated in Junos OS Release 9.4 for EX Series switches to add VSTP support.
Support for logical systems added in Junos OS Release 9.6.
Statement introduced in Junos OS Release 13.2X50-D10 for EX Series switches.

Description

For Spanning Tree Protocol (STP), Rapid Spanning Tree Protocol (RSTP), VLAN Spanning Tree Protocol (VSTP), or Multiple Spanning Tree Protocol (MSTP), specifies the maximum expected arrival time of hello BPDUs.

Default

20 seconds

Options

seconds—(Optional) Number of seconds expected between hello BPDUs.

seconds—The maximum age of received protocol BPDUs.
Range: 6 through 40
Default: 20 seconds

Required Privilege Level

routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- *Understanding STP*
- *Understanding MSTP*
- *Understanding VSTP*

Copyright © 2018, Juniper Networks, Inc.

- *Example: Configuring Network Regions for VLANs with MSTP on Switches*
- *Example: Configuring Network Regions for VLANs with MSTP on non-ELS EX Series Switches*
- *Example: Configuring Faster Convergence and Network Stability on ELS Switches with RSTP*
- *Example: Faster Convergence and Improved Network Stability with RSTP on EX Series Switches*
- show spanning-tree bridge on page 267
- show spanning-tree interface on page 273

Copyright © 2018, Juniper Networks, Inc.

max-hops

Syntax

max-hops *hops*;

Hierarchy Level

[edit logical-systems *logical-system-name* protocols mstp],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols mstp],

[edit protocols mstp],

[edit routing-instances *routing-instance-name* protocols mstp]

Release Information

Statement introduced in Junos OS Release 8.4.
Statement introduced in Junos OS Release 9.0 for EX Series switches.
Support for logical systems added in Junos OS Release 9.6.
Statement introduced in Junos OS Release 13.2X50-D10 for EX Series switches.

Description

Configure the maximum number of hops a BPDU can be forwarded in the MSTP region.

Options

hops—(Optional) Number of hops the BPDU can be forwarded.

Range: 1 through 255

Default: 19 hops20 hops

Required Privilege Level

routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- *Understanding MSTP*
- *Configuring Multiple Spanning Tree Protocol*
- *Configuring MSTP on Switches*
- *Example: Configuring Network Regions for VLANs with MSTP on Switches*
- *Example: Configuring Network Regions for VLANs with MSTP on non-ELS EX Series Switches*
- *Understanding MSTP*
- show spanning-tree bridge on page 267
- show spanning-tree interface on page 273

mode

Syntax

```
mode (p2p | shared);
```

Hierarchy Level

[edit logical-systems *logical-system-name* protocols (mstp | rstp | vstp) interface *interface-name*],
[edit logical-systems *logical-system-name* protocols vstp vlan *vlan-id* interface *interface-name*],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols (mstp | rstp | vstp) interface *interface-name*],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols vstp vlan *vlan-id* interface *interface-name*],

[edit protocols (mstp | rstp | vstp) interface *interface-name*],
[edit protocols vstp vlan *vlan-id* interface *interface-name*],
[edit protocols mstp interface (all | *interface-name*) arp-on-stp],
[edit protocols mstpmsti *msti-id* interface *interface-name*) arp-on-stp],
[edit protocols rstp interface (all | *interface-name*) arp-on-stp],

[edit routing-instances *routing-instance-name* protocols (mstp | rstp | vstp) interface *interface-name*],
[edit routing-instances *routing-instance-name* protocols vstp vlan *vlan-id* interface *interface-name*]
[edit protocols stp interface (all | (all | *interface-name*)],
[edit protocols stp interface (all | *interface-name*) arp-on-stp],
[edit protocols vstp vlan *vlan-id* interface (all | *interface-name*) arp-on-stp]

Release Information

Statement introduced in Junos OS Release 8.4.
Statement introduced in Junos OS Release 9.0 for EX Series switches.
Statement updated in Junos OS Release 9.4 for EX Series switches to add VSTP support.
Support for logical systems added in Junos OS Release 9.6.
Statement introduced in Junos OS Release 13.2X50-D10 for EX Series switches.

Description

For Spanning Tree Protocol (STP), Rapid Spanning Tree Protocol (RSTP), VLAN Spanning Tree Protocol (VSTP), or Multiple Spanning Tree Protocol (MSTP), configures the link mode to identify point-to-point links.

Default

When the link is configured as full-duplex, the default link mode is **p2p**. When the link is configured half-duplex, the default link mode is **shared**.

NOTE: For EX4300 switches, the interfaces operate in full-duplex mode only.

Options

p2p—The link is point to point.

shared—The link is shared media.

 Copyright © 2018, Juniper Networks, Inc.

Required Privilege Level

routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- *Understanding Spanning-Tree Instance Interfaces*
- *Understanding STP*
- *Understanding VSTP*
- *Example: Configuring Network Regions for VLANs with MSTP on Switches*
- *Example: Configuring Faster Convergence and Network Stability on ELS Switches with RSTP*
- *Example: Configuring Network Regions for VLANs with MSTP on non-ELS EX Series Switches*
- *Example: Faster Convergence and Improved Network Stability with RSTP on EX Series Switches*
- show spanning-tree bridge on page 267
- show spanning-tree interface on page 273

msti

Syntax

```
msti msti-id {
  bridge-priority priority;
  vlan (vlan-id | vlan-range|open-set-of-values);
  interface (interface-name | all) {
    cost cost;
    edge;
    priority interface-priority;
  }
}
```

Syntax

```
msti msti-id {
  vlan (vlan-id | vlan-name);
  interface interface-name {
    disable-timeout;
    cost cost;
    priority priority;
  }
}
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols mstp],
[edit logical-systems logical-system-name routing-instances routing-instance-name
  protocols mstp],
```

```
[edit protocols mstp],
```

```
[edit routing-instances routing-instance-name protocols mstp]
```

Release Information

Statement introduced in Junos OS Release 8.4.
Statement introduced in Junos OS Release 9.0 for EX Series switches.
Support for logical systems added in Junos OS Release 9.6.
Statement introduced in Junos OS Release 13.2X50-D10 for EX Series switches.
Statement updated in Junos OS Release 15.1 for EX Series switches to support configuration of spanning tree parameters globally on all interfaces.

NOTE: You cannot disable spanning tree parameters globally on all interfaces.

Description

Configures the Multiple Spanning Tree Instance (MSTI) identifier for Multiple Spanning Tree Protocol (MSTP). MSTI IDs are local to each region, so you can reuse the same MSTI ID in different regions.

Default

MSTI is disabled.

 Copyright © 2018, Juniper Networks, Inc.

Options ***msti-id***—MSTI instance identifier.

Range: 1 through 64

Range: 1 through 4094. The Common Instance Spanning Tree (CIST) is always MSTI 0.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- *Understanding MSTP*
- *Configuring Multiple Spanning Tree Protocol*
- *Configuring MSTP on Switches*
- *Configuring MSTP Instances on a Physical Interface*
- *Example: Configuring Network Regions for VLANs with MSTP on Switches*
- *Example: Configuring Network Regions for VLANs with MSTP on non-ELS EX Series Switches*
- *Understanding MSTP*
- show spanning-tree bridge on page 267
- show spanning-tree interface on page 273

Copyright © 2018, Juniper Networks, Inc.

no-root-port

Syntax

```
no-root-port;
```

Hierarchy Level

[edit logical-systems *logical-system-name* protocols (mstp | rstp | vstp) interface *interface-name*],
[edit logical-systems *logical-system-name* protocols vstp vlan *vlan-id* interface *interface-name*],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols (mstp | rstp | vstp) interface *interface-name*],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols vstp vlan *vlan-id* interface *interface-name*],

[edit protocols (mstp | rstp | vstp) interface *interface-name*],
[edit protocols vstp vlan *vlan-id* interface *interface-name*],
[edit protocols mstp interface (all | *interface-name*) arp-on-stp],
[edit protocols rstp interface (all | *interface-name*) arp-on-stp],
[edit protocols stp interface (all | *interface-name*) arp-on-stp],

[edit routing-instances *routing-instance-name* protocols (mstp | rstp | vstp) interface *interface-name*],
[edit routing-instances *routing-instance-name* protocols vstp vlan *vlan-id* interface *interface-name*]
[edit protocols vstp vlan *vlan-id* interface (all | *interface-name*) arp-on-stp]

Release Information

Statement introduced in Junos OS Release 9.1.
Statement updated in Junos OS Release 9.4 for EX Series switches to add VSTP support.
Support for logical systems added in Junos OS Release 9.6.
Statement introduced in Junos OS Release 13.2X50-D10 for EX Series switches.
Statement introduced in Junos OS Release 17.1 for ACX Series routers.

Description

Configures an interface to be a spanning-tree designated port. If the bridge receives superior STP bridge protocol data units (BPDUs) on a root-protected interface, that interface transitions to a root-prevented STP state (inconsistency state) and the interface is blocked. This blocking prevents a bridge from being elected the root bridge. When the bridge stops receiving superior STP BPDUs on the root-protected interface, interface traffic is no longer blocked.

Required Privilege Level

routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- *Understanding VSTP*
- *Understanding Root Protection for Spanning-Tree Instance Interfaces in a Layer 2 Switched Network*
- *Enabling Root Protection for a Spanning-Tree Instance Interface*

Copyright © 2018, Juniper Networks, Inc.

- *Example: Configuring Root Protection to Enforce Root Bridge Placement in Spanning Trees on non-ELS EX Series Switches*
- *Example: Configuring Network Regions for VLANs with MSTP on non-ELS EX Series Switches*
- *Example: Faster Convergence and Improved Network Stability with RSTP on EX Series Switches*
- show spanning-tree bridge on page 267
- show spanning-tree interface on page 273

priority (Protocols STP)

Syntax

priority *interface-priority*;

Hierarchy Level

[edit logical-systems *logical-system-name* protocols (mstp | rstp | vstp) interface *interface-name*],
[edit logical-systems *logical-system-name* protocols mstp msti *msti-id* interface *interface-name*],
[edit logical-systems *logical-system-name* protocols vstp vlan *vlan-id* interface *interface-name*],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols (mstp | rstp | vstp) interface *interface-name*],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols mstp msti *msti-id* interface *interface-name*],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols vstp vlan *vlan-id* interface *interface-name*],

[edit protocols (mstp | rstp | vstp) interface *interface-name*],
[edit protocols mstp msti *msti-id* interface *interface-name*],
[edit protocols vstp *vlan-id* interface *interface-name*],
[edit protocols mstp ;interface (all | *interface-name*) arp-on-stp],
[edit protocols mstpmsti *msti-id* interface *interface-name* arp-on-stp],
[edit protocols rstp interface (all | *interface-name*) arp-on-stp],
[edit protocols stp interface (all | *interface-name*) arp-on-stp],
[edit protocols vstp vlan *vlan-id* interface (all | *interface-name*) arp-on-stp]

[edit routing-instances *routing-instance-name* protocols (mstp | rstp | vstp) interface *interface-name*],
[edit routing-instances *routing-instance-name* protocols mstp msti *msti-id* interface *interface-name*],
[edit routing-instances *routing-instance-name* protocols vstp vlan *vlan-id* interface *interface-name*]

Release Information

Statement introduced in Junos OS Release 8.4.
Statement introduced in Junos OS Release 9.0 for EX Series switches.
Statement updated in Junos OS Release 9.4 for EX Series switches to add VSTP support.
Support for logical systems added in Junos OS Release 9.6.
Statement introduced in Junos OS Release 13.2X50-D10 for EX Series switches.

Description

For Spanning Tree Protocol (STP), Rapid Spanning Tree Protocol (RSTP), VLAN Spanning Tree Protocol (VSTP), or Multiple Spanning Tree Protocol (MSTP), specifies the interface priority to control which interface is elected as the root port. The interface priority must be set in increments of 16.

Default

The default value is 128.

Copyright © 2018, Juniper Networks, Inc.

Options *priority*—(Optional) Interface priority. The interface priority must be set in increments of 16.

Range: 0 through 240

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- *Understanding Spanning-Tree Instance Interfaces*
- *Understanding STP*
- *Understanding VSTP*
- *Configuring a Spanning-Tree Instance Interface as an Edge Port for Faster Convergence*
- *Example: Configuring Network Regions for VLANs with MSTP on non-ELS EX Series Switches*
- *Example: Faster Convergence and Improved Network Stability with RSTP on EX Series Switches*
- show spanning-tree bridge on page 267
- show spanning-tree interface on page 273

protocols (STP Type)

Syntax

```
protocols {
  mstp { ... }
  rstp { ... }
  vstp { ... }
}
```

Hierarchy Level

[edit],

[edit logical-systems *logical-system-name*],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name*],

[edit routing-instances *routing-instance-name*]

Release Information

Statement introduced in Junos OS Release 8.4.
Statement introduced in Junos OS Release 13.2X50-D10 for EX Series switches.
Support for logical systems added in Junos OS Release 9.6.

Description

Configure the Spanning Tree Protocol type as MSTP, RSTP, or VSTP.

Options

mstp—Configure the protocol as Multiple Spanning Tree.

rstp—Configure the protocol as Rapid Spanning Tree.

vstp—Configure the protocol as VLAN Spanning Tree.

Required Privilege Level

routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- *Configuring RSTP on EX Series Switches (CLI Procedure)*
- *Configuring MSTP on Switches*
- *Configuring MSTP Instances on a Physical Interface*
- *Configuring VLAN Spanning Tree Protocol on Switches*
- *Configuring Rapid Spanning Tree Protocol*
- *Configuring Multiple Spanning Tree Protocol*
- *Configuring VLAN Spanning Tree Protocol*
- *Understanding MSTP*

Copyright © 2018, Juniper Networks, Inc.

revision-level

Syntax

```
revision-level revision-level;
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols mstp],
[edit logical-systems logical-system-name routing-instances routing-instance-name
   protocols mstp],
```

```
[edit protocols mstp],
[edit routing-instances routing-instance-name protocols mstp]
```

Release Information

Statement introduced in Junos OS Release 8.4.
Statement introduced in Junos OS Release 9.0 for EX Series switches.
Support for logical systems added in Junos OS Release 9.6.
Statement introduced in Junos OS Release 13.2X50-D10 for EX Series switches.

Description

For Multiple Spanning Tree Protocol (MSTP), set the revision number of the MSTP configuration.

Options

revision-level—Configure the revision number of the MSTP region configuration.

Range: 0 through 65,535

Required Privilege Level

routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- *Understanding MSTP*
- *Configuring Multiple Spanning Tree Protocol*
- *Example: Configuring Network Regions for VLANs with MSTP on Switches*
- *Example: Configuring Network Regions for VLANs with MSTP on non-ELS EX Series Switches*
- show spanning-tree bridge on page 267
- show spanning-tree interface on page 273

rstp

List of Syntax

MX Series on page 178
EX Series on page 178
ELS Versions: EX Series, QFX Series, NFX Series on page 179

MX Series

```
rstp {
  bpdu-block-on-edge;
  bpdu-destination-mac-address provider-bridge-group;
  bridge-priority priority;
  extended-system-id;
  force-version stp;
  forward-delay seconds;
  hello-time seconds;
  max-age seconds;
  interface interface-name {
    bpdu-timeout-action {
      alarm;
      block;
    }
    cost cost;
    edge;
    mode (p2p | shared);
    no-root-port;
    priority interface-priority;
  }
  priority-hold-time seconds;
  traceoptions {
    file filename <files number> <size size> <world-readable | no-world-readable>;
    flag flag <flag-modifier> <disable>;
  }
}
```

EX Series

```
rstp {
  bpdu-block-on-edge;
  bridge-priority priority; priority;
  disable;
  forward-delay seconds;
  hello-time seconds;
  interface (all | interface-name) {
    arp-on-stp;
    bpdu-timeout-action {
      block;
      log;
    }
    cost cost;
    edge;
    mode mode;
    no-root-port;
    priority priority;
  }
  max-age seconds;
```

 Copyright © 2018, Juniper Networks, Inc.

```
  traceoptions {
    file filename <files number > <size size> <no-stamp | no-world-readable |
      world-readable>;
    flag flag;
  }
}
```

ELS Versions: EX Series, QFX Series, NFX Series

```
rstp {
  bpdu-block-on-edge;
  bpdu-destination-mac-address provider-bridge-group;
  bridge-priority priority;
  disable;
  extended-system-id;
  force-version stp;
  forward-delay seconds;
  hello-time seconds;
  max-age seconds;
  priority-hold-time seconds;
  traceoptions {
    file filename <files number> <size size> <world-readable | no-world-readable>;
    flag flag <flag-modifier> <disable>;
  }
}
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols],
[edit logical-systems logical-system-name routing-instances routing-instance-name
  protocols],
```

```
[edit protocols],
```

```
[edit routing-instances routing-instance-name protocols]
```

Release Information

Statement introduced in Junos OS Release 8.4.
Statement introduced in Junos OS Release 9.0 for EX Series switches.
bpdu-block-on-edge statement added in Junos OS Release 9.4.
bpdu-timeout-action statement added in Junos OS Release 9.4.
Support for logic systems added in Junos OS Release 9.6.
Statement introduced in Junos OS Release 13.2X50-D10 for EX Series switches.
Statement updated in Junos OS Release 15.1 for EX Series and QFX Series switches to support configuration of spanning tree parameters globally on all interfaces.

NOTE: You cannot disable spanning tree parameters globally on all interfaces.

Description Configure Rapid Spanning Tree Protocol (RSTP). RSTP is defined in the IEEE 802.1D-2004 specification and is used to prevent loops in Layer 2 networks, which results in shorter convergence times than those provided by basic Spanning Tree Protocol (STP).

VSTP and RSTP can be configured concurrently. You can selectively configure up to 253 VLANs using VSTP; the remaining VLANs will be configured using RSTP. VSTP and RSTP are the only spanning-tree protocols that can be configured concurrently on the switch.

BEST PRACTICE: **Configure RSTP when you configure VSTP. RSTP overhead is minimal and this configuration ensures that a spanning-tree protocol is running on all VLANs on your switch, even when your switch is supporting more than 253 VLANs.**

The remaining statements are explained separately. See CLI Explorer.

Default RSTP is enabled on all Ethernet switching interfaces.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- *Configuring Rapid Spanning Tree Protocol*
- *Understanding RSTP*
- *Configuring RSTP on EX Series Switches (CLI Procedure)*
- *Example: Faster Convergence and Improved Network Stability with RSTP on EX Series Switches*
- show spanning-tree bridge on page 267
- show spanning-tree interface on page 273

 Copyright © 2018, Juniper Networks, Inc.

traceoptions (Spanning Tree)

Syntax

```
traceoptions {
  file filename &lt;files number&gt; &lt;size size&gt; &lt;world-readable | no-world-readable&gt;;
  flag flag &lt;flag-modifier&gt; &lt;disable&gt;;
}
```

Hierarchy Level

[edit logical-systems *logical-system-name* protocols (mstp | rstp | vstp)],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols (mstp | rstp | vstp)],

[edit protocols (mstp | rstp | vstp | vstp vlan *vlan-id*)],

[edit routing-instances *routing-instance-name* protocols (mstp | rstp | vstp)]

Release Information

Statement introduced in Junos OS Release 8.4.
Statement introduced in Junos OS Release 9.0 for EX Series switches.
Statement updated in Junos OS Release 9.4 for EX Series switches to add VSTP support.
Support for logical systems added in Junos OS Release 9.6.
Statement introduced in Junos OS Release 13.2X50-D10 for EX Series switches.

Description

Sets protocol-level tracing options for for spanning-tree protocols MPLS, MVRP, STP, RSTP, MSTP, and VSTP.

Default

The default STP protocol-level trace options are inherited from the global **traceoptions** statement. Traceoptions is disabled.

Options

disable—(Optional) Disable the tracing operation. One use of this option is to disable a single operation when you have defined a broad group of tracing operations, such as **all**.

file ***filename***—Name of the file to receive the output of the tracing operation. Enclose the name in quotation marks. We recommend that you place STP tracing output in the file **/var/log/stp-log**.

files ***number***—(Optional) Maximum number of trace files. When a trace file named ***trace-file*** reaches its maximum size, it is renamed ***trace-file*.0**, then ***trace-file*.1**, and so on, until the maximum number of trace files is reached. Then, the oldest trace file is overwritten.

If you specify a maximum number of files, you must also specify a maximum file size with the **size** option.

Range: 2 through 1000 files

Default: 1 trace file only

flag—Tracing operation to perform. To specify more than one tracing operation, include multiple **flag** statements. The following are the STP-specific tracing options:

- **all**—Trace all operations.
- **all-failures**—Trace all failure conditions.
- **bpdu**—Trace BPDU reception and transmission.
- **bridge-detection-state-machine**—Trace the bridge detection state machine.
- **error**—Trace all failure conditions.
- **events**—Trace events of the protocol state machine.
- **pdu**—Trace PDUs that were received and sent.
- **port-information-state-machine**—Trace the port information state machine.
- **port-migration-state-machine**—Trace the port migration state machine.
- **port-receive-state-machine**—Trace the port receive state machine.
- **port-role-transit-state-machine**—Trace the port role transit state machine.
- **port-role-select-state-machine**—Trace the port role selection state machine.
- **port-state-transit-state-machine**—Trace the port state transit state machine.
- **port-transmit-state-machine**—Trace the port transmit state machine.
- **ppmd**—Trace the state and events for the ppmd process.
- **socket**—Trace socket activity.
- **state-machine**—Trace state machine information.
- **state-machine-variables**—Trace when the state machine variables change.
- **timers**—Trace protocol timers.
- **topology-change-state-machine**—Trace the topology change state machine.

The following are the global tracing options:

- **all**—All tracing operations.
- **config-internal**—Trace configuration internals.
- **general**—Trace general events.
- **normal**—All normal events.

Default: If you do not specify this option, only unusual or abnormal operations are traced.

Copyright © 2018, Juniper Networks, Inc.

- **parse**—Trace configuration parsing.
- **policy**—Trace policy operations and actions.
- **regex-parse**—Trace regular-expression parsing.
- **route**—Trace routing table changes.
- **state**—Trace state transitions.
- **task**—Trace protocol task processing.
- **timer**—Trace protocol task timer processing.

no-world-readable—(Optional) Prevent any user from reading the log file.

replace—(Optional) Replace an existing trace file if there is one.

Default: If you do not include this option, tracing output is appended to an existing trace file.

size *size*—(Optional) Maximum size of each trace file, in kilobytes (KB) or megabytes (MB). When a trace file named ***trace-file*** reaches this size, it is renamed ***trace-file*.0**. When the ***trace-file*** again reaches its maximum size, ***trace-file*.0** is renamed ***trace-file*.1** and ***trace-file*** is renamed ***trace-file*.0**. This renaming scheme continues until the maximum number of trace files is reached. Then the oldest trace file is overwritten.

If you specify a maximum file size, you must also specify a maximum number of trace files with the files option.

Syntax: ***x*k** to specify KB, ***x*m** to specify MB, or ***x*g** to specify GB

Range: 10 KB through the maximum file size supported on your system

Default: 1 MB

world-readable—(Optional) Allow any user to read the log file.

Required Privilege Level

routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- *Understanding RSTP*
- *Understanding STP*
- *Understanding VSTP*
- *Understanding Multiple VLAN Registration Protocol (MVRP)*
- *Understanding Spanning-Tree Protocol Trace Options*
- *Configuring Tracing Spanning-Tree Operations*
- *Understanding MSTP*
- *Example: Configuring Network Regions for VLANs with MSTP on non-ELS EX Series Switches*
- *Example: Faster Convergence and Improved Network Stability with RSTP on EX Series Switches*
- *Example: Tracing Spanning-Tree Protocol Operations*
- show spanning-tree bridge on page 267
- show spanning-tree interface on page 273

 Copyright © 2018, Juniper Networks, Inc.

vlan (MSTP)

Syntax

```
vlan vlan-id;
```

EX Series

```
vlan (all | vlan-id | vlan-name) {
  bridge-priority priority;
  forward-delay seconds;
  hello-time seconds;
  interface interface-name {
    bpdu-timeout-action {
      block;
      log;
    }
    cost  cost;
    disable;
    edge;
    mode mode;
    no-root-port;
    priority priority;
  }
  max-age seconds;
  traceoptions {
    file filename <files number > <size size> <no-stamp | world-readable |
      no-world-readable>;
    flag flag;
  }
}
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols mstp msti msti-id],
[edit logical-systems logical-system-name routing-instances routing-instance-name
  protocols mstp msti msti-id],
```

```
[edit protocols mstp msti msti-id],
[edit protocols mstpmsti msti-id]
[edit protocols vstp]
```

```
[edit routing-instances routing-instance-name protocols mstp msti msti-id]
```

Release Information

Statement introduced in Junos OS Release 8.4.
Statement introduced in Junos OS Release 9.0 for EX Series switches.
Statement updated with enhanced **?** (CLI completion feature) functionality in Junos OS Release 9.5 for EX Series switches.
Support for logical systems added in Junos OS Release 9.6.
Statement introduced in Junos OS Release 13.2X50-D10 for EX Series switches.

Description

Configure the VLANs for a Multiple Spanning Tree Instance (MSTI) or VSTP instance.

Copyright © 2018, Juniper Networks, Inc.

NOTE: When you configure VSTP with the **set prototocol vstp vlan all** command, vlan-id 1 is excluded to be compatible with Cisco PVST+. If you want vlan-id 1 to be included in VSTP, you must set it separately with the **set protocol vstp vlan 1** command.

TIP: To display a list of all configured VLANs on the system, including VLANs that are configured but not committed, type **?** after **vlan** or **vlans** in your configuration mode command line. Note that only one VLAN is displayed for a VLAN range.

Options

vlan-name—Name of the VLAN.

vlan-id—The VLAN identifier associated with the MSTI.

vlan-id-range—Range of VLAN identifiers associated with the MSTI in the form ***minimum-vlan-id-maximum-vlan-id***. VLAN identifier ranges are not supported for VSTP.

Range: 1 through 4096

Required Privilege Level

routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- *Understanding MSTP*
- *Understanding VSTP*
- *Configuring Multiple Spanning Tree Protocol*
- *Example: Configuring Network Regions for VLANs with MSTP on non-ELS EX Series Switches*

 Copyright © 2018, Juniper Networks, Inc.

vlan (VSTP)

Syntax

```
vlan vlan-id {
  bridge-priority priority;
  forward-delay seconds;
  hello-time seconds;
  max-age seconds;
  interface interface-name {
    cost cost;
    edge;
    mode (p2p | shared);
    no-root-port;
    priority interface-priority;
  }
}
```

EX Series

```
vlan (all | vlan-id | vlan-name) {
  bridge-priority priority;
  forward-delay seconds;
  hello-time seconds;
  interface (all | interface-name) {
    bpdu-timeout-action {
      block;
      log;
    }
    cost  cost;
    disable;
    edge;
    mode mode;
    no-root-port;
    priority priority;
  }
  max-age seconds;
  traceoptions {
    file filename <files number > <size size> <no-stamp | world-readable |
      no-world-readable>;
    flag flag;
  }
}
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols vstp],
[edit protocols vstp]
```

Release Information

Statement introduced in Junos OS Release 9.0.
Statement introduced in Junos OS Release 9.4 for EX Series switches.
Support for logical systems added in Junos OS Release 9.6.
Statement introduced in Junos OS Release 13.2X50-D10 for EX Series switches.

Description

Configure VSTP VLAN parameters.

Copyright © 2018, Juniper Networks, Inc.

TIP: To display a list of all configured VLANs on the system, including VLANs that are configured but not committed, type ? after **vlan** or **vlans** in your configuration mode command line. Note that only one VLAN is displayed for a VLAN range.

Options **all**—All VLANs.

vlan-id—Numeric VLAN identifier.

vlan-range—Name of the VLAN range.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- *Understanding VSTP*
- *Configuring VLAN Spanning Tree Protocol*

 Copyright © 2018, Juniper Networks, Inc.

vlan-group

Syntax

```
vlan-group group group-name {
vlan (vlan-id |vlan-group |all) {
}
```

Hierarchy Level

```
[edit protocols vstp]
```

Release Information Statement introduced in Junos OS Release 15.1 for EX Series switches.

Description Configure VLAN group for Spanning Tree Protocol (VSTP). VSTP is used to prevent loops in Layer 2 networks on a per-VLAN basis.

> BEST PRACTICE: **Configure RSTP when you configure VSTP. RSTP overhead is minimal and this configuration ensures that a spanning-tree protocol is running on all VLANs on your switch, even when your switch is supporting more than the maximum number of allowed VSTP VLANs.**

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- vstp on page 190
- show spanning-tree bridge on page 267
- show spanning-tree interface on page 273
- *Understanding VSTP*

vstp

List of Syntax

MX Series on page 190
EX Series, QFX Series, QFabric on page 190
ELS versions: EX Series, QFX Series, NFX Series on page 192

MX Series

```
vstp {
  bpdu-block-on-edge;
  force-version stp;
  interface interface-name {
    bpdu-timeout-action {
      alarm;
      block;
    }
    cost cost;
    edge;
    mode (p2p | shared);
    no-root-port;
    priority interface-priority;
  }
  priority-hold-time seconds;
  vlan vlan-id {
    bridge-priority priority;
    forward-delay seconds;
    hello-time seconds;
    max-age seconds;
    interface interface-name {
      access-trunk
      bpdu-timeout-action {
        alarm;
        block;
      }
      cost cost;
      edge;
      mode (p2p | shared);
      no-root-port;
      priority interface-priority;
    }
  }
  traceoptions {
    file filename <files number> <size size> <world-readable | no-world-readable>;
    flag flag <flag-modifier> <disable>;
  }
}
```

EX Series, QFX Series, QFabric

```
vstp {
  bpdu-block-on-edge;
  disable;
  force-version stp;
  vlan (all | vlan-id | vlan-name) {
    bridge-priority priority;
    forward-delay seconds;
```

 Copyright © 2018, Juniper Networks, Inc.

```
    hello-time seconds;
    interface (all | interface-name) {
      arp-on-stp;
      bpdu-timeout-action {
        block;
        log;
      }
      cost  cost;
      disable;
      edge;
      mode mode;
      no-root-port;
      priority priority;
    }
    max-age seconds;
    traceoptions {
      file filename <files number > <size size> <no-stamp | no-world-readable |
        world-readable>;
      flag flag;
    }
  }
}
```

ELS versions: EX Series, QFX Series, NFX Series

```
vstp {
  bpdu-block-on-edge;
  disable;
  force-version stp;
  interface (interface-name disable | interface-range-name | all ){
    bpdu-timeout-action {
      alarm;
      block;
    }
    cost cost;
    edge;
    mode (p2p | shared);
    no-root-port;
    priority interface-priority;
  }
priority-hold-time seconds;
  vlan (vlan-id | all){
    bridge-priority priority;
    forward-delay seconds;
    hello-time seconds;
    max-age seconds;
    interface (interface-name disable | interface-range-name | all ){
      bpdu-timeout-action {
        alarm;
        block;
      }
      cost cost;
      edge;
      mode (p2p | shared);
      no-root-port;
      priority interface-priority;
    }
  }
  traceoptions {
    file filename <files number> <size size> <world-readable | no-world-readable>;
    flag flag <flag-modifier> <disable>;
  }
  vlan-group group group-name {
    vlansvlan-name (vlan-id |vlan-range | open-set-of-values) {
      interface all;
      interface interface-name {
        disable;
      }
    }
  }
```


Copyright © 2018, Juniper Networks, Inc.

Hierarchy Level

[edit protocols]

[edit logical-systems *logical-system-name* protocols],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols],

[edit protocols],

[edit routing-instances *routing-instance-name* protocols]

Release Information

Statement introduced in Junos OS Release 9.0.
Statement introduced in Junos OS Release 9.4 for EX Series switches.
bpdu-block-on-edge statement added in Junos OS Release 9.4.
bpdu-timeout-action statement added in Junos OS Release 9.4.
Support for logical systems added in Junos OS Release 9.6.
Statement introduced in Junos OS Release 13.2X50-D10 for EX Series switches.
Statement updated in Junos OS Release 15.1 for EX Series switches to support configuration of spanning tree parameters globally on all interfaces.

NOTE: You cannot disable spanning tree parameters globally on all interfaces.

Description Configures VLAN Spanning Tree Protocol (VSTP). VSTP is used to prevent loops in Layer 2 networks on a per-VLAN basis.

You can have a maximum of 253 VSTP VLANs per switch.

If the number of VLANs on your switch exceeds the VSTP VLAN limit, you must use the vlan statement to specify which VLANs or VLAN groups use VSTP. You also cannot use the **vlan all** option to configure VSTP when your switch has more than the maximum allowed VSTP VLANs. To ensure all VLANs are running a spanning-tree protocol, run RSTP for networks with large numbers of VLANs .

NOTE: When you configure VSTP with the **set protocol vstp vlan all** command, VLAN ID 1 is not set; it is excluded so that the configuration is compatible with Cisco PVST+. If you want VLAN ID 1 to be included in the VSTP configuration on your switch, you must set it separately with the **set protocol vstp vlan 1** command.

NOTE: Option **vlan all** is not supported in Junos OS Release 13.2X50.

BEST PRACTICE: Configure RSTP when you configure VSTP. RSTP overhead is minimal and this configuration ensures that some spanning tree protocol is running on all VLANs on your switch, even when your switch has more than the maximum number of allowed VSTP VLANs.

The remaining statements are explained separately. See CLI Explorer.

Default VSTP is not enabled by default.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- *Understanding VSTP*
- *Configuring VLAN Spanning Tree Protocol*
- Configuring VSTP (CLI Procedure) on page 61
- *Example: Configuring VSTP on QFX Series Switches and EX4600 Switches*
- **show spanning-tree bridge on page 267**
- **show spanning-tree interface on page 273**

 Copyright © 2018, Juniper Networks, Inc.

CHAPTER 7

Bridging and VLANs Monitoring Commands

clear ethernet-switching table

Syntax

clear ethernet-switching table
<interface *interface-name*>
<mac *mac-address*>
<management-vlan>
<persistent-mac <*interface* | *mac-address*>>
<vlan *vlan-name*>

Syntax (QFX Series)

clear ethernet-switching table
<interface *interface-name*>
<mac *mac-address*>
<persistent-mac <*interface* | *mac-address*>>
<vlan *vlan-name*>

Release Information

Command introduced in Junos OS Release 9.3 for EX Series switches.
Command introduced in Junos OS Release 11.1 for the QFX Series.

Description

NOTE: On a QFabric system, using this command on an FCoE-enabled VLAN when FCoE sessions are active can cause traffic flooding and FCoE traffic drop. The FCoE sessions are not terminated and the traffic reconverges after a short period of time.

Clear learned entries, which are media access control (MAC) addresses, in the Ethernet switching table (also called the forwarding database table).

Options

none—Clear learned entries in the Ethernet switching table, except for persistent MAC addresses.

interface *interface-name*—(Optional) Clear all learned MAC addresses for the specified interface from the Ethernet switching table.

mac *mac-address*—(Optional) Clear the specified learned MAC address from the Ethernet switching table.

management-vlan—(Optional) Clear all MAC addresses learned for the management VLAN from the Ethernet switching table. Note that you do not specify a VLAN name because only one management VLAN exists.

persistent-mac <*interface* | *mac-address*>—(Optional) Clear all MAC addresses, including persistent MAC addresses. Use the **interface** option to clear all MAC addresses on an interface, or use the **mac-address** option to clear all entries for a specific MAC address.

 Copyright © 2018, Juniper Networks, Inc.

Use this command whenever you move a device in your network that has a persistent MAC address on the switch. If you move the device to another port on the switch and do not clear the persistent MAC address from the original port it was learned on, then the new port will not learn the MAC address and the device will not be able to connect. If the original port is down when you move the device, then the new port will learn the MAC address and the device can connect—however, unless you cleared the MAC address on the original port, when the port comes back up, the system reinstalls the persistent MAC address in the forwarding table for that port. If this occurs, the address is removed from the new port and the device loses connectivity.

vlan ***vlan-name***—(Optional) Clear all MAC addresses learned for the specified VLAN from the Ethernet switching table.

Required Privilege Level view

Related Documentation

- show ethernet-switching table on page 206

List of Sample Output clear ethernet-switching table on page 197

Output Fields This command produces no output.

Sample Output

clear ethernet-switching table

```
user@switch> clear ethernet-switching table
```

show ethernet-switching interfaces

Syntax

```
show ethernet-switching interfaces
<brief | detail | summary>
<interface interface-name>
```

Release Information

Command introduced in Junos OS Release 9.0 for EX Series switches.
In Junos OS Release 9.6 for EX Series switches, the following updates were made:

- **Blocking** field output was updated.
- The default view was updated to include information about 802.1Q tags.
- The **detail** view was updated to include information on VLAN mapping.

Command introduced in Junos OS Release 11.1 for the QFX Series.
In Junos OS Release 11.1 for EX Series switches, the **detail** view was updated to include reflective relay information.
Command introduced in Junos OS Release 14.1X53-D20 for OCX Series switches.

Description

Display information about switched Ethernet interfaces.

Options

none—(Optional) Display brief information for Ethernet-switching interfaces.

brief | detail | summary—(Optional) Display the specified level of output.

interface *interface-name*—(Optional) Display Ethernet-switching information for a specific interface.

Required Privilege Level

view

Related Documentation

- Troubleshooting Ethernet Switching on page 24Understanding Bridging and VLANs on Switches on page 9
- *Example: Setting Up Basic Bridging and a VLAN on Switches*
- *Example: Setting Up Bridging with Multiple VLANs*
- *Understanding FCoE*
- *Interfaces Overview for Switches*
- **show ethernet-switching mac-learning-log on page 254**
- **show ethernet-switching table on page 206**
- *Configuring Autorecovery From the Disabled State on Secure or Storm Control Interfaces (CLI Procedure)*

List of Sample Output

show ethernet-switching interfaces on page 201

 Copyright © 2018, Juniper Networks, Inc.

Output Fields For QFX Series, QFabric, NFX Series, EX4600 and OCX1100:

Table 5 on page 199 lists the output fields for the **show ethernet-switching interfaces** command on QFX Series, QFabric, NFX Series, EX4600 and OCX1100. Output fields are listed in the approximate order in which they appear.

Table 5: show ethernet-switching interfaces Output Fields

Field Name	Field Description	Level of Output
Interface	Name of a switching interface.	All levels
State	Interface state. Values are **up** or **down**.	none, **brief**, **detail**, **summary**
VLAN members	Name of a VLAN.	none, **brief**, **detail**, **summary**
Blocking	Forwarding state of the interface: • **blocked**—Traffic is not being forwarded on the interface. • **unblocked**—Traffic is forwarded on the interface. • **MAC limit exceeded**—The interface is temporarily disabled because of a MAC limiting error. The disabled interface is automatically restored to service when the disable timeout expires. • **MAC move limit exceeded**—The interface is temporarily disabled because of a MAC move limiting error. The disabled interface is automatically restored to service when the disable timeout expires. • **Storm control in effect** —The interface is temporarily disabled because of a storm control error. The disabled interface is automatically restored to service when the disable timeout expires. • **Storm control shutdown in effect** —The interface is temporarily disabled because of a storm control shutdown error. The disabled interface is automatically restored to service when the disable timeout expires.	none, **brief**, **detail**, **summary**
Index	VLAN index internal to Junos OS software.	**detail**
untagged | tagged	Specifies whether the interface forwards IEEE802.1Q-tagged or untagged traffic.	**detail**

Copyright © 2018, Juniper Networks, Inc.

Output fields for EX Series:

Table 6 on page 200 lists the output fields for the **show ethernet-switching interfaces** command on EX Series switches. Output fields are listed in the approximate order in which they appear.

Table 6: show ethernet-switching interfaces Output Fields

Field Name	Field Description	Level of Output
Interface	Name of a switching interface.	none, **brief**, **detail**, **summary**
Index	VLAN index internal to Junos OS.	**detail**
State	Interface state. Values are **up** and **down**.	none, **brief**, **detail**
Port mode	The **access** mode is the port mode default and works with a single VLAN. Port mode can also be **trunk**, which accepts tagged packets from multiple VLANs on other switches. The third port mode value is **tagged-access**, which accepts tagged packets from access devices.	**detail**
Reflective Relay Status	Reflective relay allows packets to use the same interface for both upstream and downstream traffic. When reflective relay has been configured, the status displayed is always **enabled** . When reflective relay is not configured, this entry does not appear in the command output.	**detail**
Ether type for the interface	Ether type is a two-octet field in an Ethernet frame used to indicate which protocol is encapsulated in the payload of an incoming Ethernet packet. Both 802.1Q packets and Q-in-Q packets use this field. The output displayed for this particular field indicates the interface's Ether type, which is used to match the Ether type of incoming 802.1Q packets and Q-in-Q packets. The indicated Ether type field is also added to the interface's outgoing 802.1Q and Q-in-Q packets.	**detail**
VLAN membership	Names of VLANs that belong to this interface.	none, **brief**, **detail**,
Tag	Number of the 802.1Q tag.	none, **brief**, **detail**,
Tagging	Specifies whether the interface forwards 802.1Q **tagged** or **untagged** traffic.	none, **brief**, **detail**,

Copyright © 2018, Juniper Networks, Inc.

Table 6: show ethernet-switching interfaces Output Fields (continued)

Field Name	Field Description	Level of Output
Blocking	The forwarding state of the interface: • **unblocked**—Traffic is forwarded on the interface. • **blocked**—Traffic is not being forwarded on the interface. • **Disabled by bpdu control**—The interface is disabled due to receiving BPDUs on a protected interface. If the **disable-timeout** statement has been included in the BPDU configuration, the interface automatically returns to service after the timer expires. • **blocked by RTG**—The specified redundant trunk group is disabled. • **blocked by STP**—The interface is disabled due to a spanning-tree protocol error. • **MAC limit exceeded**—The interface is temporarily disabled due to a MAC limit error. The disabled interface is automatically restored to service when the disable timeout expires. • **MAC move limit exceeded**—The interface is temporarily disabled due to a MAC move limit error. The disabled interface is automatically restored to service when the disable timeout expires. • **Storm control in effect**—The interface is temporarily disabled due to a storm control error. The disabled interface is automatically restored to service when the disable timeout expires.	none, **brief**, **detail**,
Number of MACs learned on IFL	Number of MAC addresses learned by this interface.	**detail**
mapping	When mapping is configured, the status is one of the following C-VLAN to S-VLAN mapping types: • **dot1q-tunneled**—The interface maps all traffic to the S-VLAN (all-in-one bundling). • **native**—The interface maps untagged and priority tagged packets to the S-VLAN. • **push**—The interface maps packets to a firewall filter to an S-VLAN. • **policy-mapped**—The interface maps packets to a specifically defined S-VLAN. • *integer*—The interface maps packets to the specified S-VLAN. When mapping is not configured, this entry does not appear in the command output.	**detail**

Sample Output for QFX Series Switches, QFabric, NFX Series, EX4600 and OCX1100

show ethernet-switching interfaces

user@switch> **show ethernet-switching interfaces**

```
Interface    State    VLAN members           Blocking
xe-0/0/0.0   up       T1122                  unblocked
xe-0/0/1.0   down     default                - MAC limit exceeded
xe-0/0/2.0   down     default                - MAC move limit exceeded
xe-0/0/3.0   down     default                - Storm control in effect
xe-0/0/4.0   down     default                unblocked
```

```
xe-0/0/5.0  down     default               unblocked
xe-0/0/6.0  down     default               unblocked
xe-0/0/7.0  down     default               unblocked
xe-0/0/8.0  down     default               unblocked
xe-0/0/9.0  up       T111                  unblocked
xe-0/0/10.0 down     default               unblocked
xe-0/0/11.0 down     default               unblocked
xe-0/0/12.0 down     default               unblocked
xe-0/0/13.0 down     default               unblocked
xe-0/0/14.0 down     default               unblocked
xe-0/0/15.0 down     default               unblocked
xe-0/0/16.0 down     default               unblocked
xe-0/0/17.0 down     default               unblocked
xe-0/0/18.0 down     default               unblocked
xe-0/0/19.0 up       T111                  unblocked
xe-0/1/0.0  down     default               unblocked
xe-0/1/1.0  down     default               unblocked
xe-0/1/2.0  down     default               unblocked
xe-0/1/3.0  down     default               unblocked
```

show ethernet-switching interfaces summary

`user@switch>` **show ethernet-switching interfaces summary**

```
xe-0/0/0.0
xe-0/0/1.0
xe-0/0/2.0
xe-0/0/3.0
xe-0/0/8.0
xe-0/0/10.0
xe-0/0/11.0
```

show ethernet-switching interfaces brief

`user@switch>` **show ethernet-switching interfaces brief**

```
Interface   State    VLAN members          Blocking
xe-0/0/0.0  down     default               unblocked
xe-0/0/1.0  down     employee-vlan         unblocked
xe-0/0/2.0  down     employee-vlan         unblocked
xe-0/0/3.0  down     employee-vlan         unblocked
xe-0/0/8.0  down     employee-vlan         unblocked
xe-0/0/10.0 down     default               unblocked
xe-0/0/11.0 down     employee-vlan         unblocked
```

show ethernet-switching interfaces detail

`user@switch>` **show ethernet-switching interfaces detail**

```
Interface: xe-0/0/0.0 Index: 65
  State: down
  VLANs:
    default                untagged     unblocked

Interface: xe-0/0/1.0 Index: 66
  State: down
  VLANs:
    employee-vlan          untagged     unblocked
```

 Copyright © 2018, Juniper Networks, Inc.

```
Interface: xe-0/0/2.0 Index: 67
  State: down
  VLANs:
    employee-vlan          untagged     unblocked

Interface: xe-0/0/3.0 Index: 68
  State: down
  VLANs:
    employee-vlan          untagged     unblocked

Interface: xe-0/0/8.0 Index: 69
  State: down
  VLANs:
    employee-vlan          untagged     unblocked

Interface: xe-0/0/10.0 Index: 70
  State: down
  VLANs:
    default                untagged     unblocked

Interface: xe-0/0/11.0 Index: 71
  State: down
  VLANs:
    employee-vlan          tagged       unblocked
```

show ethernet-switching interfaces interface-name

user@switch> **show ethernet-switching interfaces xe-0/0/0.0**

```
 Interface    State    VLAN members          Blocking
xe-0/0/0.0   down     default               unblocked
```

Sample Output for EX Series Switches

show ethernet-switching interfaces

user@switch> **show ethernet-switching interfaces**

```

Interface     State  VLAN members        Tag   Tagging  Blocking

ae0.0         up      default                   untagged unblocked
ge-0/0/2.0    up     vlan300           300    untagged blocked by RTG (rtggroup)
ge-0/0/3.0    up      default                            blocked by STP
ge-0/0/4.0    down    default                            MAC limit exceeded
ge-0/0/5.0    down    default                            MAC move limit exceeded
ge-0/0/6.0    down    default                            Storm control in effect
ge-0/0/7.0    down    default                            unblocked
ge-0/0/13.0   up      default                   untagged unblocked
ge-0/0/14.0   up      vlan100             100   tagged   unblocked
                      vlan200             200   tagged   unblocked
ge-0/0/15.0   up      vlan100             100   tagged   blocked by STP
                      vlan200             200   tagged   blocked by STP
ge-0/0/16.0   down    default                   untagged unblocked
ge-0/0/17.0   down    vlan100             100   tagged   Disabled by bpdu-control

                      vlan200             200   tagged   Disabled by bpdu-control
```

Copyright © 2018, Juniper Networks, Inc.

show ethernet-switching interfaces ge-0/0/15 brief

user@switch> **show ethernet-switching interfaces ge-0/0/15 brief**

```
Interface    State  VLAN members    Tag   Tagging        Blocking

ge-0/0/15.0   up     vlan100        100   tagged          blocked by STP
                     vlan200        200   tagged          blocked by STP
```

show ethernet-switching interfaces ge-0/0/2 detail (Blocked by RTG rtggroup)

user@switch> **show ethernet-switching interfaces ge-0/0/2 detail**

```
Interface: ge-0/0/2.0, Index: 65, State: up, Port mode: Access
Ether type for the interface: 0X8100
VLAN membership:
    vlan300, 802.1Q Tag: 300, untagged, msti-id: 0, blocked by RTG(rtggroup)
Number of MACs learned on IFL: 0
```

show ethernet-switching interfaces ge-0/0/15 detail (Blocked by STP)

user@switch> **show ethernet-switching interfaces ge-0/0/15 detail**

```
Interface: ge-0/0/15.0, Index: 70, State: up, Port mode: Trunk
Ether type for the interface: 0X8100
VLAN membership:
    vlan100, 802.1Q Tag: 100, tagged, msti-id: 0, blocked by STP
    vlan200, 802.1Q Tag: 200, tagged, msti-id: 0, blocked by STP

Number of MACs learned on IFL: 0
```

show ethernet-switching interfaces ge-0/0/17 detail (Disabled by bpdu-control)

user@switch> **show ethernet-switching interfaces ge-0/0/17 detail**

```
Interface: ge-0/0/17.0, Index: 71, State: down, Port mode: Trunk
Ether type for the interface: 0X8100
VLAN membership:
    vlan100, 802.1Q Tag: 100, tagged, msti-id: 1, Disabled by bpdu-control
    vlan200, 802.1Q Tag: 200, tagged, msti-id: 2, Disabled by bpdu-control
Number of MACs learned on IFL: 0
```

show ethernet-switching interfaces detail (C-VLAN to S-VLAN Mapping)

user@switch>**show ethernet-switching interfaces ge-0/0/6.0 detail**

```
Interface: ge-0/0/6.0, Index: 73, State: up, Port mode: Access
Ether type for the interface: 0X8100
VLAN membership:
    map, 802.1Q Tag: 134, Mapped Tag: native, push, dot1q-tunneled, unblocked
    map, 802.1Q Tag: 134, Mapped Tag: 20, push, dot1q-tunneled, unblocked
```

show ethernet-switching interfaces detail (Reflective Relay Is Configured)

user@switch1> **show ethernet-switching interfaces ge-7/0/2 detail**

 Copyright © 2018, Juniper Networks, Inc.

```
Interface: ge-7/0/2, Index: 66, State: down, Port mode: Tagged-access
Ether type for the interface: 0X8100
Reflective Relay Status: Enabled
Ether type for the interface: 0x8100
VLAN membership:
    VLAN_Purple VLAN_Orange VLAN_Blue, 802.1Q Tag: 450, tagged, unblocked
Number of MACs learned on IFL: 0
```

show ethernet-switching table

List of Syntax

Syntax (QFX Series, QFabric, NFX Series and EX4600) on page 206
Syntax (EX Series) on page 206
Syntax (EX Series, MX Series and QFX Series) on page 206
Syntax (SRX Series) on page 206

Syntax (QFX Series, QFabric, NFX Series and EX4600)

```
show ethernet-switching table
<brief | detail | extensive | summary>
<interface interface-name>
<management-vlan>
<sort-by (name | tag)>
<vlan vlan-name>
```

Syntax (EX Series)

```
show ethernet-switching table
<brief | detail | extensive | summary>
<interface interface-name>
<management-vlan>
<persistent-mac <interface interface-name>>
<sort-by (name | tag)>
<vlan vlan-name>
```

Syntax (EX Series, MX Series and QFX Series)

```
show ethernet-switching table
<brief | count | detail | extensive | summary>
<address>
<instance instance-name>
<interface interface-name>
isid isid
<logical-system logical-system-name>
<persistent-learning (interface interface-name | mac mac-address)>
<address>
<vlan-id (all-vlan | vlan-id)>
<vlan-name (all | vlan-name)>
```

Syntax (SRX Series)

```
show ethernet-switching table (brief |detail |extensive) interface  interface-name
```

Release Information

Command introduced in Junos OS Release 9.0 for EX Series switches.
Command introduced in Junos OS Release 9.5 for SRX Series.
Options **summary**, **management-vlan**, and **vlan** ***vlan-name*** introduced in Junos OS Release 9.6 for EX Series switches.
Option **sort-by** and field name **tag** introduced in Junos OS Release 10.1 for EX Series switches.
Command introduced in Junos OS Release 11.1 for the QFX Series.
Output for private VLANs introduced in Junos OS Release 12.1 for the QFX Series.
Option **persistent-mac** introduced in Junos OS Release 11.4 for EX Series switches.
Command introduced in Junos OS Release 12.3R2.

 Copyright © 2018, Juniper Networks, Inc.

Command introduced in Junos OS Release 12.3R2 for EX Series switches.
Options **logical-system**, **persistent-learning**, and **summary** introduced in Junos OS Release 13.2X50-D10 (ELS).

Description Displays the Ethernet switching table.

(MX Series routers,EX Series switches only) Displays Layer 2 MAC address information.

Options For QFX Series, QFabric, NFX Series and EX4600:

none—(Optional) Display brief information about the Ethernet switching table.

brief | detail | extensive | summary—(Optional) Display the specified level of output.

interface ***interface-name***—(Optional) Display the Ethernet switching table for a specific interface.

management-vlan—(Optional) Display the Ethernet switching table for a management VLAN.

persistent-mac <interface ***interface-name*****>**—(Optional) Display the persistent MAC addresses learned for all interfaces or a specified interface. You can use this command to view entries that you want to clear for an interface that you intentionally disabled.

sort-by (*name* | *tag*)—(Optional) Display VLANs in ascending order of VLAN IDs or VLAN names.

vlan ***vlan-name***—(Optional) Display the Ethernet switching table for a specific VLAN.

For EX Series, MX Series and QFX Series:

none—Display all learned Layer 2 MAC address information.

brief | count | detail | extensive | summary—(Optional) Display the specified level of output.

address—(Optional) Display the specified learned Layer 2 MAC address information.

instance ***instance-name***—(Optional) Display learned Layer 2 MAC addresses for the specified routing instance.

interface ***interface-name***—(Optional) Display learned Layer 2 MAC addresses for the specified interface.

isid ***isid***—(Optional) Display learned Layer 2 MAC addresses for the specified ISID.

logical-system ***logical-system-name***—(Optional) Display Ethernet-switching statistics information for the specified logical system.

persistent-learning (interface ***interface-name*** **| mac** ***mac-address*****)**—(Optional) Display dynamically learned MAC addresses that are retained despite device restarts and

Copyright © 2018, Juniper Networks, Inc.

interface failures for a specified interface, or information about a specified MAC address.

vlan-id (all-vlan | ***vlan-id*****)**—(Optional) Display learned Layer 2 MAC addresses for all VLANs or for the specified VLAN.

vlan-name (all | ***vlan-name*****)**—(Optional) Display learned Layer 2 MAC addresses for all VLANs or for the specified VLAN.

For SRX Series:

- **none**—(Optional) Display brief information about the Ethernet switching table.
- **brief** | **detail** | **extensive**—(Optional) Display the specified level of output.
- **interface-name**—(Optional) Display the Ethernet switching table for a specific interface.

Additional Information

When Layer 2 protocol tunneling is enabled, the tunneling MAC address 01:00:0c:cd:cd:d0 is installed in the MAC table. When the Cisco Discovery Protocol (CDP), Spanning Tree Protocol (STP), or VLAN Trunk Protocol (VTP) is configured for Layer 2 protocol tunneling on an interface, the corresponding protocol MAC address is installed in the MAC table.

Required Privilege Level

view

Related Documentation

- *Example: Setting Up Basic Bridging and a VLAN on Switches*
- *Example: Setting Up Bridging with Multiple VLANs*
- *Example: Setting Up Basic Bridging and a VLAN for an EX Series Switch*
- *Example: Setting Up Bridging with Multiple VLANs for EX Series Switches*
- *Example: Setting Up Q-in-Q Tunneling on EX Series Switches*
- clear ethernet-switching table on page 196
- show ethernet-switching mac-learning-log on page 254

List of Sample Output

 Copyright © 2018, Juniper Networks, Inc.

show ethernet-switching table interface (QFX Series, QFabric, NFX Series and EX460) on page 219
show ethernet-switching table (EX Series switches) on page 219
show ethernet-switching table brief (EX Series switches) on page 220
show ethernet-switching table detail (EX Series switches) on page 220
show ethernet-switching table extensive (EX Series switches) on page 221
show ethernet-switching table persistent-mac (EX Series switches) on page 221
show ethernet-switching table persistent-mac interface ge-0/0/16.0 (EX Series switches) on page 222
show ethernet-switching table (EX Series, MX Series and QFX Series) on page 222
show ethernet-switching table brief on page 223
show ethernet-switching table count on page 224
show ethernet-switching table extensive on page 225
show ethernet-switching table detail (SRX Series) on page 226
show ethernet-switching table extensive (SRX Series) on page 227
show ethernet-switching table interface ge-0/0/1 (SRX Series) on page 228

Output Fields

For QFX Series, QFabric, NFX Series and EX4600:

The following table lists the output fields for the **show ethernet-switching table** command on QFX Series, QFabric, NFX Series and EX4600. Output fields are listed in the approximate order in which they appear.

Table 7: show ethernet-switching table Output Fields

Field Name	Field Description	Level of Output
VLAN	Name of a VLAN.	All levels
MAC address	MAC address associated with the VLAN.	All levels
Type	Type of MAC address: • **static**—The MAC address is manually created. • **learn**—The MAC address is learned dynamically from a packet's source MAC address. • **flood**—The MAC address is unknown and flooded to all members.	All levels
Age	Time remaining before the entry ages out and is removed from the Ethernet switching table.	All levels
Interfaces	Interface associated with learned MAC addresses or with the **All-members** option (flood entry).	All levels
Learned	For learned entries, the time at which the entry was added to the Ethernet switching table.	**detail**, **extensive**

For EX Series switches:

The following table lists the output fields for the **show ethernet-switching table** command on EX Series switches. Output fields are listed in the approximate order in which they appear.

Copyright © 2018, Juniper Networks, Inc.

Table 8: show ethernet-switching table Output Fields

Field Name	Field Description	Level of Output
VLAN	The name of a VLAN.	All levels
Tag	The VLAN ID tag name or number.	**extensive**
MAC or **MAC address**	The MAC address associated with the VLAN.	All levels
Type	The type of MAC address. Values are: • **static**—The MAC address is manually created. • **learn**—The MAC address is learned dynamically from a packet's source MAC address. • **flood**—The MAC address is unknown and flooded to all members. • **persistent**—The learned MAC addresses that will persist across restarts of the switch or interface-down events.	All levels except **persistent-mac**
Type	The type of MAC address. Values are: • **installed**—addresses that are in the Ethernet switching table. • **uninstalled**—addresses that could not be installed in the table or were uninstalled in an interface-down event and will be reinstalled in the table when the interface comes back up.	**persistent-mac**
Age	The time remaining before the entry ages out and is removed from the Ethernet switching table.	All levels
Interfaces	Interface associated with learned MAC addresses or **All-members** (flood entry).	All levels
Learned	For learned entries, the time which the entry was added to the Ethernet switching table.	**detail**, **extensive**
Nexthop index	The next-hop index number.	**detail**, **extensive**
persistent-mac	**installed** indicates MAC addresses that are in the Ethernet switching table and **uninstalled** indicates MAC addresses that could not be installed in the table or were uninstalled in an interface-down event (and will be reinstalled in the table when the interface comes back up).	

For EX Series, MX Series and QFX Series:

The table describes the output fields for the **show ethernet-switching table** command on EX Series, MX Series and QFX Series. Output fields are listed in the approximate order in which they appear.

Table 9: show ethernet-switching table Output fields

Field Name	Field Description
Routing instance	Name of the routing instance.
VLAN name	Name of the VLAN.

Copyright © 2018, Juniper Networks, Inc.

Table 9: show ethernet-switching table Output fields (continued)

Field Name	Field Description
MAC address	MAC address or addresses learned on a logical interface.
MAC flags	Status of MAC address learning properties for each interface: • **S**—Static MAC address is configured. • **D**—Dynamic MAC address is configured. • **L**—Locally learned MAC address is configured. • **SE**—MAC accounting is enabled. • **NM**—Non-configured MAC. • **R**—Locally learned MAC address is configured.
Age	This field is not supported.
Logical interface	Name of the logical interface.
Active source	IP address of remote entity on which MAC address is learned.
MAC count	Number of MAC addresses learned on the specific routing instance or interface.
Learning interface	Name of the logical interface on which the MAC address was learned.
Learning VLAN	VLAN ID of the routing instance or VLAN in which the MAC address was learned.
Layer 2 flags	Debugging flags signifying that the MAC address is present in various lists.
Epoch	Spanning-tree-protocolepoch number identifying when the MAC address was learned. Used for debugging.
Sequence number	Sequence number assigned to this MAC address. Used for debugging.
Learning mask	Mask of the Packet Forwarding Engines where this MAC address was learned. Used for debugging.
IPC generation	Creation time of the logical interface when this MAC address was learned. Used for debugging.

For SRX Series:

Table 10 on page 211 lists the output fields for the **show ethernet-switching table** command. Output fields are listed in the approximate order in which they appear.

Table 10: show ethernet-switching table Output Fields

Field Name	Field Description
VLAN	The name of a VLAN.

Copyright © 2018, Juniper Networks, Inc.

Table 10: show ethernet-switching table Output Fields (continued)

Field Name	Field Description
MAC address	The MAC address associated with the VLAN.
Type	The type of MAC address. Values are: • static—The MAC address is manually created. • learn—The MAC address is learned dynamically from a packet's source MAC address. • flood—The MAC address is unknown and flooded to all members.
Age	The time remaining before the entry ages out and is removed from the Ethernet switching table.
Interfaces	Interface associated with learned MAC addresses or All-members (flood entry).
Learned	For learned entries, the time which the entry was added to the Ethernet switching table.

Sample Output

show ethernet-switching table (Enhanced Layer 2 Software on QFX Series, QFabric, NFX Series and EX460)

user@switch> **show ethernet-switching table**

```
MAC flags (S - static MAC, D - dynamic MAC, L - locally learned, P - Persistent
static
           SE - statistics enabled, NM - non configured MAC, R - remote PE MAC,
O - ovsdb MAC)


Ethernet switching table : 2 entries, 2 learned
Routing instance : default-switch
    Vlan                MAC                 MAC         Age    Logical
    name                address             flags              interface
    vlan1               b0:c6:9a:ca:3c:01   D             -   ae1.0

    vlan1               b0:c6:9a:ca:3c:03   D             -   ae1.0


MAC flags (S - static MAC, D - dynamic MAC, L - locally learned, P - Persistent
static
           SE - statistics enabled, NM - non configured MAC, R - remote PE MAC,
O - ovsdb MAC)

Ethernet switching table : 2 entries, 2 learned

Routing instance : default-switch
    Vlan                MAC                 MAC         Age    Logical
    name                address             flags              interface
    vlan10              b0:c6:9a:ca:3c:01   D             -   ae1.0

    vlan10              b0:c6:9a:ca:3c:03   D             -   ae1.0


MAC flags (S - static MAC, D - dynamic MAC, L - locally learned, P - Persistent
static
```

 Copyright © 2018, Juniper Networks, Inc.

```
           SE - statistics enabled, NM - non configured MAC, R - remote PE MAC,
O - ovsdb MAC)


Ethernet switching table : 2 entries, 2 learned
Routing instance : default-switch
    Vlan                MAC                 MAC         Age    Logical
    name                address             flags              interface
    vlan2               b0:c6:9a:ca:3c:01   D             -   ae1.0

    vlan2               b0:c6:9a:ca:3c:03   D             -   ae1.0


MAC flags (S - static MAC, D - dynamic MAC, L - locally learned, P - Persistent
static
           SE - statistics enabled, NM - non configured MAC, R - remote PE MAC,
O - ovsdb MAC)


Ethernet switching table : 2 entries, 2 learned
Routing instance : default-switch
    Vlan                MAC                 MAC         Age    Logical
    name                address             flags              interface
    vlan3               b0:c6:9a:ca:3c:01   D             -   ae1.0

    vlan3               b0:c6:9a:ca:3c:03   D             -   ae1.0


MAC flags (S - static MAC, D - dynamic MAC, L - locally learned, P - Persistent
static
           SE - statistics enabled, NM - non configured MAC, R - remote PE MAC,
O - ovsdb MAC)


Ethernet switching table : 2 entries, 2 learned
Routing instance : default-switch
    Vlan                MAC                 MAC         Age    Logical
    name                address             flags              interface
    vlan4               b0:c6:9a:ca:3c:01   D             -   ae1.0

    vlan4               b0:c6:9a:ca:3c:03   D             -   ae1.0
```

show ethernet-switching table (QFX Series, QFabric, NFX Series and EX460)

user@switch> **show ethernet-switching table**

```
Ethernet-switching table: 57 entries, 17 learned
  VLAN              MAC address       Type         Age Interfaces
  F2                *                 Flood          - All-members
  F2                00:00:05:00:00:03 Learn          0 xe-0/0/44.0
  F2                00:19:e2:50:7d:e0 Static         - Router
  Linux             *                 Flood          - All-members
  Linux             00:19:e2:50:7d:e0 Static         - Router
  Linux             00:30:48:90:54:89 Learn          0 xe-0/0/47.0
  T1                *                 Flood          - All-members
  T1                00:00:05:00:00:01 Learn          0 xe-0/0/46.0
  T1                00:00:5e:00:01:00 Static         - Router
  T1                00:19:e2:50:63:e0 Learn          0 xe-0/0/46.0
  T1                00:19:e2:50:7d:e0 Static         - Router
  T10               *                 Flood          - All-members
```

Copyright © 2018, Juniper Networks, Inc.

```
  T10               00:00:5e:00:01:09 Static         - Router
  T10               00:19:e2:50:63:e0 Learn          0 xe-0/0/46.0
  T10               00:19:e2:50:7d:e0 Static         - Router
  T111              *                 Flood          - All-members
  T111              00:19:e2:50:63:e0 Learn          0 xe-0/0/15.0
  T111              00:19:e2:50:7d:e0 Static         - Router
  T111              00:19:e2:50:ac:00 Learn          0 xe-0/0/15.0
  T2                *                 Flood          - All-members
  T2                00:00:5e:00:01:01 Static         - Router
  T2                00:19:e2:50:63:e0 Learn          0 xe-0/0/46.0
  T2                00:19:e2:50:7d:e0 Static         - Router
  T3                *                 Flood          - All-members
  T3                00:00:5e:00:01:02 Static         - Router
  T3                00:19:e2:50:63:e0 Learn          0 xe-0/0/46.0
  T3                00:19:e2:50:7d:e0 Static         - Router
  T4                *                 Flood          - All-members
  T4                00:00:5e:00:01:03 Static         - Router
  T4                00:19:e2:50:63:e0 Learn          0 xe-0/0/46.0
[output truncated]
```

show ethernet-switching table (Private VLANs on QFX Series, QFabric, NFX Series and EX460)

user@switch> **show ethernet-switching table**

```
Ethernet-switching table: 10 entries, 3 learned
  VLAN              MAC address       Type         Age Interfaces
  pvlan             *                 Flood          - All-members
  pvlan             00:10:94:00:00:02 Replicated        - xe-0/0/28.0
  pvlan             00:10:94:00:00:35 Replicated        - xe-0/0/46.0
  pvlan             00:10:94:00:00:46 Replicated        - xe-0/0/4.0
  c2                *                 Flood          - All-members
  c2                00:10:94:00:00:02 Learn          0 xe-0/0/28.0
  c1                *                 Flood          - All-members
  c1                00:10:94:00:00:46 Learn          0 xe-0/0/4.0
  __pvlan_pvlan_xe-0/0/46.0__ *       Flood          - All-members
  __pvlan_pvlan_xe-0/0/46.0__ 00:10:94:00:00:35 Learn        0 xe-0/0/46.0
```

show ethernet-switching table (Junos Fusion Data Center with EVPN on QFX Series switches)

user@switch> **show ethernet-switching table**

```
MAC flags (S - static MAC, D - dynamic MAC, L - locally learned, P - Persistent
static
           SE - statistics enabled, NM - non configured MAC, R - remote PE MAC,
O - ovsdb MAC)

Ethernet switching table : 30 entries, 30 learned
Routing instance : default-switch
   Vlan                MAC                 MAC      Logical                Active

   name                address             flags    interface              source

   v100                 00:31:46:e8:f9:d6   D        vtep.32768
192.168.2.22
   v100                 7c:e2:ca:e2:75:7c   D        vtep.32771
192.168.4.44
   v100                 7c:e2:ca:e4:05:9a   D        vtep.32770
192.168.3.33
   v101                 00:31:46:e8:f9:d6   D        vtep.32768
```

 Copyright © 2018, Juniper Networks, Inc.

```
192.168.2.22
   v101              7c:e2:ca:e2:75:7c   D        vtep.32771
192.168.4.44
   v101              7c:e2:ca:e4:05:9a   D        vtep.32770
192.168.3.33
   v102              00:31:46:e8:f9:d6   D        vtep.32768
192.168.2.22
   v102              7c:e2:ca:e2:75:7c   D        vtep.32771
192.168.4.44
   v102              7c:e2:ca:e4:05:9a   D        vtep.32770
192.168.3.33
   v103              00:31:46:e8:f9:d6   D        vtep.32768
192.168.2.22
   v103              7c:e2:ca:e2:75:7c   D        vtep.32771
192.168.4.44
   v103              7c:e2:ca:e4:05:9a   D        vtep.32770
192.168.3.33
   v3001             00:31:46:e8:f9:d6   D        vtep.32768
192.168.2.22
   v3001             28:c0:da:6a:9f:c2   DL       ae11.0
   v3001             7c:e2:ca:e2:75:7c   D        vtep.32771
192.168.4.44
   v3001             7c:e2:ca:e4:05:9a   D        vtep.32770
192.168.3.33
   v3002             00:31:46:e8:f9:d6   D        vtep.32768
192.168.2.22
   v3002             7c:e2:ca:e2:75:7c   D        vtep.32771
192.168.4.44
   v3002             7c:e2:ca:e4:05:9a   D        vtep.32770
192.168.3.33
   v3003             00:31:46:e8:f9:d6   D        vtep.32768
192.168.2.22
   v3003             28:c0:da:6a:9f:c2   DL       ae11.0
   v3003             7c:e2:ca:e2:75:7c   D        vtep.32771
192.168.4.44
   v3003             7c:e2:ca:e4:05:9a   D        vtep.32770
192.168.3.33
   v3004             00:31:46:e8:f9:d6   D        vtep.32768
192.168.2.22
   v3004             7c:e2:ca:e2:75:7c   D        vtep.32771
192.168.4.44
   v3004             7c:e2:ca:e4:05:9a   D        vtep.32770
192.168.3.33
   v3005             00:31:46:e8:f9:d6   D        vtep.32768
192.168.2.22
   v3005             28:c0:da:6a:9f:c2   DL       ae11.0
   v3005             7c:e2:ca:e2:75:7c   D        vtep.32771
192.168.4.44
   v3005             7c:e2:ca:e4:05:9a   D        vtep.32770
192.168.3.33
```

show ethernet-switching table brief (QFX Series, QFabric, NFX Series and EX460)

user@switch> **show ethernet-switching table brief**

```
Ethernet-switching table: 57 entries, 17 learned
  VLAN              MAC address       Type         Age Interfaces
  F2                *                 Flood          - All-members
  F2                00:00:05:00:00:03 Learn          0 xe-0/0/44.0
  F2                00:19:e2:50:7d:e0 Static         - Router
```

Copyright © 2018, Juniper Networks, Inc.

```
  Linux             *                 Flood          - All-members
  Linux             00:19:e2:50:7d:e0 Static         - Router
  Linux             00:30:48:90:54:89 Learn          0 xe-0/0/47.0
  T1                *                 Flood          - All-members
  T1                00:00:05:00:00:01 Learn          0 xe-0/0/46.0
  T1                00:00:5e:00:01:00 Static         - Router
  T1                00:19:e2:50:63:e0 Learn          0 xe-0/0/46.0
  T1                00:19:e2:50:7d:e0 Static         - Router
  T10               *                 Flood          - All-members
  T10               00:00:5e:00:01:09 Static         - Router
  T10               00:19:e2:50:63:e0 Learn          0 xe-0/0/46.0
  T10               00:19:e2:50:7d:e0 Static         - Router
  T111              *                 Flood          - All-members
  T111              00:19:e2:50:63:e0 Learn          0 xe-0/0/15.0
  T111              00:19:e2:50:7d:e0 Static         - Router
  T111              00:19:e2:50:ac:00 Learn          0 xe-0/0/15.0
  T2                *                 Flood          - All-members
  T2                00:00:5e:00:01:01 Static         - Router
  T2                00:19:e2:50:63:e0 Learn          0 xe-0/0/46.0
  T2                00:19:e2:50:7d:e0 Static         - Router
  T3                *                 Flood          - All-members
  T3                00:00:5e:00:01:02 Static         - Router
  T3                00:19:e2:50:63:e0 Learn          0 xe-0/0/46.0
  T3                00:19:e2:50:7d:e0 Static         - Router
  T4                *                 Flood          - All-members
  T4                00:00:5e:00:01:03 Static         - Router
  T4                00:19:e2:50:63:e0 Learn          0 xe-0/0/46.0
[output truncated]
```

show ethernet-switching table detail (QFX Series, QFabric, NFX Series and EX4600)

`user@switch>` **show ethernet-switching table detail**

```
Ethernet-switching table: 57 entries, 17 learned
  F2, *
    Interface(s): xe-0/0/44.0
    Type: Flood
    Nexthop index: 0

  F2, 00:00:05:00:00:03
    Interface(s): xe-0/0/44.0
    Type: Learn, Age: 0, Learned: 2:03:09
    Nexthop index: 0

  F2, 00:19:e2:50:7d:e0
    Interface(s): Router
    Type: Static
    Nexthop index: 0

  Linux, *
    Interface(s): xe-0/0/47.0
    Type: Flood
    Nexthop index: 0

  Linux, 00:19:e2:50:7d:e0
    Interface(s): Router
    Type: Static
    Nexthop index: 0

  Linux, 00:30:48:90:54:89
```

 Copyright © 2018, Juniper Networks, Inc.

```
    Interface(s): xe-0/0/47.0
    Type: Learn, Age: 0, Learned: 2:03:08
    Nexthop index: 0

  T1, *
    Interface(s): xe-0/0/46.0
    Type: Flood
    Nexthop index: 0

  T1, 00:00:05:00:00:01
    Interface(s): xe-0/0/46.0
    Type: Learn, Age: 0, Learned: 2:03:07
    Nexthop index: 0

  T1, 00:00:5e:00:01:00
    Interface(s): Router
    Type: Static
    Nexthop index: 0

  T1, 00:19:e2:50:63:e0
    Interface(s): xe-0/0/46.0
    Type: Learn, Age: 0, Learned: 2:03:07
    Nexthop index: 0

  T1, 00:19:e2:50:7d:e0
    Interface(s): Router
    Type: Static
    Nexthop index: 0

  T10, *
    Interface(s): xe-0/0/46.0
    Type: Flood
    Nexthop index: 0

  T10, 00:00:5e:00:01:09
    Interface(s): Router
    Type: Static
    Nexthop index: 0

  T10, 00:19:e2:50:63:e0
    Interface(s): xe-0/0/46.0
    Type: Learn, Age: 0, Learned: 2:03:08
    Nexthop index: 0

  T10, 00:19:e2:50:7d:e0
    Interface(s): Router
    Type: Static
    Nexthop index: 0

  T111, *
    Interface(s): xe-0/0/15.0
    Type: Flood
    Nexthop index: 0
[output truncated]
```

show ethernet-switching table extensive (QFX Series, QFabric, NFX Series and EX460)

`user@switch>` **show ethernet-switching table extensive**

Copyright © 2018, Juniper Networks, Inc.

```
Ethernet-switching table: 57 entries, 17 learned
  F2, *
    Interface(s): xe-0/0/44.0
    Type: Flood
    Nexthop index: 0

  F2, 00:00:05:00:00:03
    Interface(s): xe-0/0/44.0
    Type: Learn, Age: 0, Learned: 2:03:09
    Nexthop index: 0

  F2, 00:19:e2:50:7d:e0
    Interface(s): Router
    Type: Static
    Nexthop index: 0

  Linux, *
    Interface(s): xe-0/0/47.0
    Type: Flood
    Nexthop index: 0

  Linux, 00:19:e2:50:7d:e0
    Interface(s): Router
    Type: Static
    Nexthop index: 0

  Linux, 00:30:48:90:54:89
    Interface(s): xe-0/0/47.0
    Type: Learn, Age: 0, Learned: 2:03:08
    Nexthop index: 0

  T1, *
    Interface(s): xe-0/0/46.0
    Type: Flood
    Nexthop index: 0

  T1, 00:00:05:00:00:01
    Interface(s): xe-0/0/46.0
    Type: Learn, Age: 0, Learned: 2:03:07
    Nexthop index: 0

  T1, 00:00:5e:00:01:00
    Interface(s): Router
    Type: Static
    Nexthop index: 0

  T1, 00:19:e2:50:63:e0
    Interface(s): xe-0/0/46.0
    Type: Learn, Age: 0, Learned: 2:03:07
    Nexthop index: 0

  T1, 00:19:e2:50:7d:e0
    Interface(s): Router
    Type: Static
    Nexthop index: 0

  T10, *
    Interface(s): xe-0/0/46.0
    Type: Flood
    Nexthop index: 0
```

 Copyright © 2018, Juniper Networks, Inc.

```
  T10, 00:00:5e:00:01:09
    Interface(s): Router
    Type: Static
    Nexthop index: 0

  T10, 00:19:e2:50:63:e0
    Interface(s): xe-0/0/46.0
    Type: Learn, Age: 0, Learned: 2:03:08
    Nexthop index: 0

  T10, 00:19:e2:50:7d:e0
    Interface(s): Router
    Type: Static
    Nexthop index: 0

  T111, *
    Interface(s): xe-0/0/15.0
    Type: Flood
    Nexthop index: 0
[output truncated]
```

show ethernet-switching table interface (QFX Series, QFabric, NFX Series and EX460)

user@switch> **show ethernet-switching table interface xe-0/0/1**

```
Ethernet-switching table: 1 unicast entries
  VLAN              MAC address       Type         Age Interfaces
  V1                *                 Flood          - All-members
  V1                00:00:05:00:00:05 Learn          0 xe-0/0/1.0
```

show ethernet-switching table (EX Series switches)

user@switch> **show ethernet-switching table**

```
Ethernet-switching table: 57 entries, 15 learned, 2 persistent
  VLAN              MAC address       Type         Age Interfaces
  F2                *                 Flood          - All-members
  F2                00:00:05:00:00:03 Learn          0 ge-0/0/44.0
  F2                00:19:e2:50:7d:e0 Static         - Router
  Linux             *                 Flood          - All-members
  Linux             00:19:e2:50:7d:e0 Static         - Router
  Linux             00:30:48:90:54:89 Learn          0 ge-0/0/47.0
  T1                *                 Flood          - All-members
  T1                00:00:05:00:00:01 Persistent     0 ge-0/0/46.0
  T1                00:00:5e:00:01:00 Static         - Router
  T1                00:19:e2:50:63:e0 Persistent     0 ge-0/0/46.0
  T1                00:19:e2:50:7d:e0 Static         - Router
  T10               *                 Flood          - All-members
  T10               00:00:5e:00:01:09 Static         - Router
  T10               00:19:e2:50:63:e0 Learn          0 ge-0/0/46.0
  T10               00:19:e2:50:7d:e0 Static         - Router
  T111              *                 Flood          - All-members
  T111              00:19:e2:50:63:e0 Learn          0 ge-0/0/15.0
  T111              00:19:e2:50:7d:e0 Static         - Router
  T111              00:19:e2:50:ac:00 Learn          0 ge-0/0/15.0
  T2                *                 Flood          - All-members
  T2                00:00:5e:00:01:01 Static         - Router
  T2                00:19:e2:50:63:e0 Learn          0 ge-0/0/46.0
```

Copyright © 2018, Juniper Networks, Inc.

```
  T2               00:19:e2:50:7d:e0 Static         - Router
  T3               *                 Flood          - All-members
  T3               00:00:5e:00:01:02 Static         - Router
  T3               00:19:e2:50:63:e0 Learn          0 ge-0/0/46.0
  T3               00:19:e2:50:7d:e0 Static         - Router
  T4               *                 Flood          - All-members
  T4               00:00:5e:00:01:03 Static         - Router
  T4               00:19:e2:50:63:e0 Learn          0 ge-0/0/46.0
[output truncated]
```

show ethernet-switching table brief (EX Series switches)

user@switch> **show ethernet-switching table brief**

```
Ethernet-switching table: 57 entries, 15 learned, 2 persistent entries
  VLAN             MAC address       Type         Age Interfaces
  F2               *                 Flood          - All-members
  F2               00:00:05:00:00:03 Learn          0 ge-0/0/44.0
  F2               00:19:e2:50:7d:e0 Static         - Router
  Linux            *                 Flood          - All-members
  Linux            00:19:e2:50:7d:e0 Static         - Router
  Linux            00:30:48:90:54:89 Learn          0 ge-0/0/47.0
  T1               *                 Flood          - All-members
  T1               00:00:05:00:00:01 Persistent     0 ge-0/0/46.0
  T1               00:00:5e:00:01:00 Static         - Router
  T1               00:19:e2:50:63:e0 Persistent     0 ge-0/0/46.0
  T1               00:19:e2:50:7d:e0 Static         - Router
  T10              *                 Flood          - All-members
  T10              00:00:5e:00:01:09 Static         - Router
  T10              00:19:e2:50:63:e0 Learn          0 ge-0/0/46.0
  T10              00:19:e2:50:7d:e0 Static         - Router
  T111             *                 Flood          - All-members
  T111             00:19:e2:50:63:e0 Learn          0 ge-0/0/15.0
  T111             00:19:e2:50:7d:e0 Static         - Router
  T111             00:19:e2:50:ac:00 Learn          0 ge-0/0/15.0
  T2               *                 Flood          - All-members
  T2               00:00:5e:00:01:01 Static         - Router
  T2               00:19:e2:50:63:e0 Learn          0 ge-0/0/46.0
  T2               00:19:e2:50:7d:e0 Static         - Router
  T3               *                 Flood          - All-members
  T3               00:00:5e:00:01:02 Static         - Router
  T3               00:19:e2:50:63:e0 Learn          0 ge-0/0/46.0
  T3               00:19:e2:50:7d:e0 Static         - Router
  T4               *                 Flood          - All-members
  T4               00:00:5e:00:01:03 Static         - Router
  T4               00:19:e2:50:63:e0 Learn          0 ge-0/0/46.0
[output truncated]
```

show ethernet-switching table detail (EX Series switches)

user@switch> **show ethernet-switching table detail**

```
Ethernet-switching table: 5 entries, 2 learned entries
 VLAN: default, Tag: 0, MAC: *, Interface: All-members
  Interfaces:
   ge-0/0/11.0, ge-0/0/20.0, ge-0/0/30.0, ge-0/0/36.0, ge-0/0/3.0
  Type: Flood
  Nexthop index: 1307
```

 Copyright © 2018, Juniper Networks, Inc.

```
VLAN: default, Tag: 0, MAC: 00:1f:12:30:b8:83, Interface: ge-0/0/3.0
  Type: Learn, Age: 0, Learned: 20:09:26
  Nexthop index: 1315

VLAN: v1, Tag: 101, MAC: *, Interface: All-members
 Interfaces:
  ge-0/0/31.0
 Type: Flood
 Nexthop index: 1313

VLAN: v1, Tag: 101, MAC: 00:1f:12:30:b8:89, Interface: ge-0/0/31.0
 Type: Learn, Age: 0, Learned: 20:09:25
 Nexthop index: 1312

VLAN: v2, Tag: 102, MAC: *, Interface: All-members
 Interfaces:
  ae0.0
 Type: Flood
 Nexthop index: 1317
```

show ethernet-switching table extensive (EX Series switches)

user@switch> **show ethernet-switching table extensive**

```
Ethernet-switching table: 3 entries, 1 learned, 5 persistent entries

    VLAN: v1, Tag: 10, MAC: *, Interface: All-members
    Interfaces:
                ge-0/0/14.0, ge-0/0/1.0, ge-0/0/2.0, ge-0/0/3.0, ge-0/0/4.0,
                ge-0/0/5.0, ge-0/0/6.0, ge-0/0/7.0, ge-0/0/8.0, ge-0/0/10.0,
                ge-0/0/0.0
    Type: Flood
    Nexthop index: 567

    VLAN: v1, Tag: 10, MAC: 00:21:59:c6:93:22, Interface: Router
    Type: Static
    Nexthop index: 0

    VLAN: v1, Tag: 10, MAC: 00:21:59:c9:9a:4e, Interface: ge-0/0/14.0
    Type: Learn, Age: 0, Learned: 18:40:50
    Nexthop index: 564
```

show ethernet-switching table persistent-mac (EX Series switches)

user@switch> **show ethernet-switching table persistent-mac**

```
VLAN               MAC address       Type       Interface
default            00:10:94:00:00:02 installed    ge-0/0/42.0
default            00:10:94:00:00:03 installed    ge-0/0/42.0
default            00:10:94:00:00:04 installed    ge-0/0/42.0
default            00:10:94:00:00:05 installed    ge-0/0/42.0
default            00:10:94:00:00:06 installed    ge-0/0/42.0
default            00:10:94:00:05:02 uninstalled ge-0/0/16.0
default            00:10:94:00:06:03 uninstalled ge-0/0/16.0
default            00:10:94:00:07:04 uninstalled ge-0/0/16.0
```

show ethernet-switching table persistent-mac interface ge-0/0/16.0 (EX Series switches)

```
VLAN              MAC address       Type        Interface
default           00:10:94:00:05:02 uninstalled ge-0/0/16.0
default           00:10:94:00:06:03 uninstalled ge-0/0/16.0
default           00:10:94:00:07:04 uninstalled ge-0/0/16.0
```

show ethernet-switching table (EX Series, MX Series and QFX Series)

user@host> **show ethernet-switching table**

```
MAC flags (S - static MAC, D - dynamic MAC, L - locally learned
           SE - statistics enabled, NM - non configured MAC, R - remote PE MAC)

Routing instance : default-switch
    Vlan                MAC                 MAC         Age    Logical
    name                address             flags              interface
    VLAN101             88:e0:f3:bb:07:f0   D             -   ae20.0


MAC flags (S - static MAC, D - dynamic MAC, L - locally learned
           SE - statistics enabled, NM - non configured MAC, R - remote PE MAC)

Routing instance : default-switch
    Vlan                MAC                 MAC         Age    Logical
    name                address             flags              interface
    VLAN102             88:e0:f3:bb:07:f0   D             -   ae20.0


MAC flags (S - static MAC, D - dynamic MAC, L - locally learned
           SE - statistics enabled, NM - non configured MAC, R - remote PE MAC)

Routing instance : default-switch
    Vlan                MAC                 MAC         Age    Logical
    name                address             flags              interface
    VLAN103             88:e0:f3:bb:07:f0   D             -   ae20.0


MAC flags (S - static MAC, D - dynamic MAC, L - locally learned
           SE - statistics enabled, NM - non configured MAC, R - remote PE MAC)

Routing instance : default-switch
    Vlan                MAC                 MAC         Age    Logical
    name                address             flags              interface
    VLAN104             88:e0:f3:bb:07:f0   D             -   ae20.0


MAC flags (S - static MAC, D - dynamic MAC, L - locally learned
           SE - statistics enabled, NM - non configured MAC, R - remote PE MAC)

Routing instance : default-switch
    Vlan                MAC                 MAC         Age    Logical
    name                address             flags              interface
    VLAN1101            00:1f:12:32:f5:c1   D             -   ae0.0


MAC flags (S - static MAC, D - dynamic MAC, L - locally learned
           SE - statistics enabled, NM - non configured MAC, R - remote PE MAC)
```

 Copyright © 2018, Juniper Networks, Inc.

```
Routing instance : default-switch
    Vlan                MAC                 MAC         Age    Logical
    name                address             flags              interface
    VLAN1102            00:1f:12:32:f5:c1   D             -   ae0.0


MAC flags (S - static MAC, D - dynamic MAC, L - locally learned
           SE - statistics enabled, NM - non configured MAC, R - remote PE MAC)

Routing instance : default-switch
    Vlan                MAC                 MAC         Age    Logical
    name                address             flags              interface
    VLAN1103            00:1f:12:32:f5:c1   D             -   ae0.0


MAC flags (S - static MAC, D - dynamic MAC, L - locally learned
           SE - statistics enabled, NM - non configured MAC, R - remote PE MAC)

Routing instance : default-switch
    Vlan                MAC                 MAC         Age    Logical
    name                address             flags              interface
    VLAN1104            00:1f:12:32:f5:c1   D             -   ae0.0


MAC flags (S - static MAC, D - dynamic MAC, L - locally learned
           SE - statistics enabled, NM - non configured MAC, R - remote PE MAC)

Routing instance : default-switch
    Vlan                MAC                 MAC         Age    Logical
    name                address             flags              interface
    VLAN1105            00:1f:12:32:f5:c1   D             -   ae0.0


MAC flags (S - static MAC, D - dynamic MAC, L - locally learned
           SE - statistics enabled, NM - non configured MAC, R - remote PE MAC)

Routing instance : default-switch
    Vlan                MAC                 MAC         Age    Logical
    name                address             flags              interface
    VLAN1106            00:1f:12:32:f5:c1   D             -   ae0.0
[...output truncated...]
```

show ethernet-switching table brief

`user@host>` **show ethernet-switching table brief**

```
MAC flags (S - static MAC, D - dynamic MAC, L - locally learned
           SE - statistics enabled, NM - non configured MAC, R - remote PE MAC)

Routing instance : default-switch
    Vlan                MAC                 MAC         Age    Logical
    name                address             flags              interface
    VLAN101             88:e0:f3:bb:07:f0   D             -   ae20.0


MAC flags (S - static MAC, D - dynamic MAC, L - locally learned
           SE - statistics enabled, NM - non configured MAC, R - remote PE MAC)

Routing instance : default-switch
    Vlan                MAC                 MAC         Age    Logical
```

Copyright © 2018, Juniper Networks, Inc.

```
    name                address             flags            interface
    VLAN102             88:e0:f3:bb:07:f0   D             -   ae20.0


MAC flags (S - static MAC, D - dynamic MAC, L - locally learned
           SE - statistics enabled, NM - non configured MAC, R - remote PE MAC)

Routing instance : default-switch
    Vlan                MAC                 MAC         Age    Logical
    name                address             flags              interface
    VLAN103             88:e0:f3:bb:07:f0   D             -   ae20.0


MAC flags (S - static MAC, D - dynamic MAC, L - locally learned
           SE - statistics enabled, NM - non configured MAC, R - remote PE MAC)

Routing instance : default-switch
    Vlan                MAC                 MAC         Age    Logical
    name                address             flags              interface
    VLAN104             88:e0:f3:bb:07:f0   D             -   ae20.0


MAC flags (S - static MAC, D - dynamic MAC, L - locally learned
           SE - statistics enabled, NM - non configured MAC, R - remote PE MAC)

Routing instance : default-switch
    Vlan                MAC                 MAC         Age    Logical
    name                address             flags              interface
    VLAN1101            00:1f:12:32:f5:c1   D             -   ae0.0
[...output truncated...]
```

show ethernet-switching table count

user@host> **show ethernet-switching table count**

```
0 MAC address learned in routing instance default-switch VLAN VLAN1000
ae26.0:1000

1 MAC address learned in routing instance default-switch VLAN VLAN101
ae20.0:101

  MAC address count per learn VLAN within routing instance:
    Learn VLAN ID      MAC count      Static MAC count
           101              1                 0

1 MAC address learned in routing instance default-switch VLAN VLAN102
ae20.0:102

  MAC address count per learn VLAN within routing instance:
    Learn VLAN ID      MAC count      Static MAC count
           102              1                 0

1 MAC address learned in routing instance default-switch VLAN VLAN103
ae20.0:103

  MAC address count per learn VLAN within routing instance:
    Learn VLAN ID      MAC count      Static MAC count
           103              1                 0

1 MAC address learned in routing instance default-switch VLAN VLAN104
```

 Copyright © 2018, Juniper Networks, Inc.

```
ae20.0:104

  MAC address count per learn VLAN within routing instance:
    Learn VLAN ID      MAC count      Static MAC count
           104              1                 0

0 MAC address learned in routing instance default-switch VLAN VLAN105
ae20.0:105

0 MAC address learned in routing instance default-switch VLAN VLAN106
ae20.0:106

0 MAC address learned in routing instance default-switch VLAN VLAN107
ae20.0:107

0 MAC address learned in routing instance default-switch VLAN VLAN108
ae20.0:108

0 MAC address learned in routing instance default-switch VLAN VLAN109
ae20.0:109

0 MAC address learned in routing instance default-switch VLAN VLAN110
ae20.0:110

1 MAC address learned in routing instance default-switch VLAN VLAN1101
ae0.0:1101

  MAC address count per learn VLAN within routing instance:
    Learn VLAN ID      MAC count      Static MAC count
          1101              1                 0

1 MAC address learned in routing instance default-switch VLAN VLAN1102
ae0.0:1102

  MAC address count per learn VLAN within routing instance:
    Learn VLAN ID      MAC count      Static MAC count
          1102              1                 0
[...output truncated...]
```

show ethernet-switching table extensive

user@host> **show ethernet-switching table extensive**

```

MAC address: 88:e0:f3:bb:07:f0
  Routing instance: default-switch
VLAN ID: 101
   Learning interface: ae20.0
   Layer 2 flags: in_hash,in_ifd,in_ifl,in_vlan,in_rtt,kernel,in_ifbd
   Epoch: 0                            Sequence number: 2
   Learning mask: 0x00000008

MAC address: 88:e0:f3:bb:07:f0
  Routing instance: default-switch
VLAN ID: 102
   Learning interface: ae20.0
   Layer 2 flags: in_hash,in_ifd,in_ifl,in_vlan,in_rtt,kernel,in_ifbd
   Epoch: 0                            Sequence number: 2
   Learning mask: 0x00000008

```

Copyright © 2018, Juniper Networks, Inc.

```
MAC address: 88:e0:f3:bb:07:f0
  Routing instance: default-switch
VLAN ID: 103
   Learning interface: ae20.0
   Layer 2 flags: in_hash,in_ifd,in_ifl,in_vlan,in_rtt,kernel,in_ifbd
   Epoch: 0                            Sequence number: 2
   Learning mask: 0x00000008

MAC address: 88:e0:f3:bb:07:f0
  Routing instance: default-switch
VLAN ID: 104
   Learning interface: ae20.0
   Layer 2 flags: in_hash,in_ifd,in_ifl,in_vlan,in_rtt,kernel,in_ifbd
   Epoch: 0                            Sequence number: 2
   Learning mask: 0x00000008

MAC address: 00:1f:12:32:f5:c1
  Routing instance: default-switch
VLAN ID: 1101
   Learning interface: ae0.0
   Layer 2 flags: in_hash,in_ifd,in_ifl,in_vlan,in_rtt,kernel,in_ifbd
   Epoch: 0                            Sequence number: 2
   Learning mask: 0x00000008

MAC address: 00:1f:12:32:f5:c1
  Routing instance: default-switch
VLAN ID: 1102
   Learning interface: ae0.0
   Layer 2 flags: in_hash,in_ifd,in_ifl,in_vlan,in_rtt,kernel,in_ifbd
   Epoch: 0                            Sequence number: 2
   Learning mask: 0x00000008

MAC address: 00:1f:12:32:f5:c1
  Routing instance: default-switch
VLAN ID: 1103
   Learning interface: ae0.0
   Layer 2 flags: in_hash,in_ifd,in_ifl,in_vlan,in_rtt,kernel,in_ifbd
   Epoch: 0                            Sequence number: 2
   Learning mask: 0x00000008

MAC address: 00:1f:12:32:f5:c1
  Routing instance: default-switch
VLAN ID: 1104
   Learning interface: ae0.0
   Layer 2 flags: in_hash,in_ifd,in_ifl,in_vlan,in_rtt,kernel,in_ifbd
   Epoch: 0                            Sequence number: 2
   Learning mask: 0x00000008
```

Sample Output

show ethernet-switching table detail (SRX Series)

`user@host>` **show ethernet-switching table detail**

```
Ethernet-switching table: 57 entries, 17 learned
F2, *
Interface(s): ge-0/0/44.0
Type: Flood
F2, 00:00:5E:00:53:AC
Interface(s): ge-0/0/44.0
```

 Copyright © 2018, Juniper Networks, Inc.

```
Type: Learn, Age: 0, Learned: 2:03:09
F2, 00:00:5E:00:53:AA
Interface(s): Router
Type: Static
Linux, *
Interface(s): ge-0/0/47.0
Type: Flood
Linux, 00:00:5E:00:53:AB
Interface(s): Router
Type: Static
Linux, 00:00:5E:00:53:AC
Interface(s): ge-0/0/47.0
Type: Learn, Age: 0, Learned: 2:03:08
T1, *
Interface(s): ge-0/0/46.0
Type: Flood
T1, 00:00:5E:00:53:AD
Interface(s): ge-0/0/46.0
Type: Learn, Age: 0, Learned: 2:03:07
T1, 00:00:5E:00:53:AE
Interface(s): Router
Type: Static
T1, 00:00:5E:00:53:AF
Interface(s): ge-0/0/46.0
Type: Learn, Age: 0, Learned: 2:03:07
T1, 00:00:5E:00:53:AG
Interface(s): Router
Type: Static
T10, *
Interface(s): ge-0/0/46.0
Type: Flood
T10, 00:00:5E:00:53:AH
Interface(s): Router
Type: Static
T10, 00:00:5E:00:53:AI
Interface(s): ge-0/0/46.0
Type: Learn, Age: 0, Learned: 2:03:08
T10, 00:00:5E:00:53:AJ
Interface(s): Router
Type: Static
T111, *
Interface(s): ge-0/0/15.0
Type: Flood
[output truncated]
```

Sample Output

show ethernet-switching table extensive (SRX Series)

user@host> **show ethernet-switching table extensive**

```
Ethernet-switching table: 57 entries, 17 learned
F2, *
Interface(s): ge-0/0/44.0
Type: Flood
F2, 00:00:5E:00:53:AC
Interface(s): ge-0/0/44.0
Type: Learn, Age: 0, Learned: 2:03:09
F2, 00:00:5E:00:53:AA
Interface(s): Router
```

Copyright © 2018, Juniper Networks, Inc.

```
Type: Static
Linux, *
Interface(s): ge-0/0/47.0
Type: Flood
Linux, 00:00:5E:00:53:AB
Interface(s): Router
Type: Static
Linux, 00:00:5E:00:53:AC
Interface(s): ge-0/0/47.0
Type: Learn, Age: 0, Learned: 2:03:08
T1, *
Interface(s): ge-0/0/46.0
Type: Flood
T1, 00:00:5E:00:53:AD
Interface(s): ge-0/0/46.0
Type: Learn, Age: 0, Learned: 2:03:07
T1, 00:00:5E:00:53:AE
Interface(s): Router
Type: Static
T1, 00:00:5E:00:53:AF
Interface(s): ge-0/0/46.0
Type: Learn, Age: 0, Learned: 2:03:07
T1, 00:00:5E:00:53:AG
Interface(s): Router
Type: Static
T10, *
Interface(s): ge-0/0/46.0
Type: Flood
T10, 00:00:5E:00:53:AH
Interface(s): Router
Type: Static
T10, 00:00:5E:00:53:AI
Interface(s): ge-0/0/46.0
Type: Learn, Age: 0, Learned: 2:03:08
T10, 00:00:5E:00:53:AJ
Interface(s): Router
Type: Static
T111, *
Interface(s): ge-0/0/15.0
Type: Flood
[output truncated]
```

Sample Output

show ethernet-switching table interface ge-0/0/1 (SRX Series)

user@host> **show ethernet-switching table interface ge-0/0/1**

```
Ethernet-switching table: 1 unicast entries
VLAN     MAC address         Type     Age Interfaces
V1       *                   Flood     - All-members
V1       00:00:5E:00:53:AF   Learn     0 ge-0/0/1.0
```

 Copyright © 2018, Juniper Networks, Inc.

show system statistics arp

Syntax

```
show system statistics arp
```

Release Information
Command introduced in Junos OS Release 9.6 for EX Series switches.

Description
Display system-wide Address Resolution Protocol (ARP) statistics.

Required Privilege Level
view

Related Documentation

- *Example: Configuring Proxy ARP on an EX Series Switch*
- *Verifying That Proxy ARP Is Working Correctly*

show system statistics arp

`user@switch>` **show system statistics arp**

```
arp:
         90060 datagrams received
         34 ARP requests received
         610 ARP replies received
         0 resolution request  received
         0 unrestricted proxy requests
         0 restricted proxy requests
         0 received proxy requests
         0 unrestricted proxy requests not proxied
         0 restricted proxy requests not proxied
         0 datagrams with bogus interface
         0 datagrams with incorrect length
         0 datagrams for non-IP protocol
         0 datagrams with unsupported op code
         0 datagrams with bad protocol address length
         0 datagrams with bad hardware address length
         0 datagrams with multicast source address
         0 datagrams with multicast target address
         0 datagrams with my own hardware address
         0 datagrams for an address not on the interface
         0 datagrams with a broadcast source address
         294 datagrams with source address duplicate to mine
         89113 datagrams which were not for me
         0 packets discarded waiting for resolution
         0 packets sent after waiting for resolution
         309 ARP requests sent
         35 ARP replies sent
         0 requests for memory denied
         0 requests dropped on entry
         0 requests dropped during retry
         0 requests dropped due to interface deletion
         0 requests on unnumbered interfaces
         0 new requests on unnumbered interfaces
         0 replies for from unnumbered interfaces
```

Copyright © 2018, Juniper Networks, Inc.

```
        0 requests on unnumbered interface with non-subnetted donor
        0 replies from unnumbered interface with non-subnetted donor
```

 Copyright © 2018, Juniper Networks, Inc.

show vlans

List of Syntax

Syntax (EX Series and QFX Series Switches) on page 231
Syntax (EX Series with ELS Switches and MX Routers) on page 231
Syntax (SRX Devices) on page 231

Syntax (EX Series and QFX Series Switches)

```
show vlans
<brief | detail | extensive>
<dot1q-tunneling>
<management-vlan>
<sort-by (tag | name)>
<vlan-range-name>
<summary>
<vlan-name>
<vlan-range-name>
```

Syntax (EX Series with ELS Switches and MX Routers)

```
show vlans
<brief | detail | extensive>
<instance instance-name>
<logical-system logical-system-name>
<operational>
<vlan-name)>
<interface interface-name>
```

Syntax (SRX Devices)

```
show vlans
<brief | detail | extensive>
<interface interface-name>
<logical-system (logical-system | all)>
<operational>
```

Release Information

Command introduced in Junos OS Release 8.4.
Command introduced in Junos OS Release 9.0 for EX Series switches.
Command introduced in Junos OS Release 11.1 for the QFX Series.
Option **dot1q-tunneling** added in Junos OS Release 12.1 for the QFX Series.
Command introduced in Junos OS Release 12.3R2 for EX Series switches.
Option **interface** introduced in Junos OS Release 13.2X50-D10 (ELS).

Description

Display information about VLANs configured on bridged Ethernet interfaces. For interfaces configured to support a VoIP VLAN and a data VLAN, the **show vlans** command displays both tagged and untagged membership for those VLANs.

NOTE: When a series of VLANs is created using the **vlan-range** statement, such VLAN names are preceded and followed by a double underscore. For example, a series of VLANs using the VLAN range 1 through 3 and the base VLAN name **marketing** would be displayed as **__marketing_1__**, **__marketing_2__**, and **__marketing_3__**.

Copyright © 2018, Juniper Networks, Inc.

> NOTE: To display an 802.1X supplicant successfully authenticated in multiple-supplicant mode with dynamic VLAN movement, use the **show vlans** ***vlan-name*** **extensive** operational mode command, where ***vlan-name*** is the dynamic VLAN.

Options For EX Series and QFX Series switches:

none—Display information for all VLANs. VLAN information is displayed by VLAN name in ascending order.

brief | detail | extensive—(Optional) Display the specified level of output.

dot1q-tunneling—(Optional) Display VLANs with the Q-in-Q tunneling feature enabled.

management-vlan—(Optional) Display management VLANs.

sort-by (tag | name)—(Optional) Display VLANs in ascending order of VLAN IDs or VLAN names.

vlan-range-name—(Optional) Display VLANs in ascending order of VLAN range names.

summary—(Optional) Display the total number of VLANs and counts of VLANs by type—for example, the number of dynamic, 802.1Q-tagged, and Q-in-Q tunneled VLANs.

vlan-range-name—(Optional) Display information for the specified VLAN range. To display information for all members of the VLAN range, specify the base VLAN name—for example, **employee** for a VLAN range that includes **__employee_1__** through **__employee_10__**.

For EX Series with ELS Switches and MX Routers:

none—Display information for all VLANs.

brief | detail | extensive—(Optional) Display the specified level of output.

instance ***instance-name***—(Optional) Display information for the specified routing instance.

logical-system ***logical-system-name***—(Optional) Display Ethernet-switching statistics information for the specified logical system.

operational—(Optional) Display information for the operational routing instances.

vlan-name**)**— (Optional) Display information about the specified VLAN.

interface ***interface-name***—(Optional) Display information about the specified interface.

For SRX devices:

none—Display information for all VLANs.

Copyright © 2018, Juniper Networks, Inc.

brief | detail | extensive—(Optional) Display the specified level of output.

interface *interface-name*— (Optional) Display information about a specific interface.

logical system—(Optional) Display name of the logical system or all.

operational—(Optional) Display information for the operational switching instances.

Required Privilege Level

view

Related Documentation

- *Example: Setting Up Basic Bridging and a VLAN on Switches*
- *Example: Setting Up Bridging with Multiple VLANs*
- *Example: Setting Up Basic Bridging and a VLAN for an EX Series Switch*
- *Example: Setting Up Bridging with Multiple VLANs for EX Series Switches*
- *Example: Configuring a Private VLAN on a Single EX Series Switch*
- *Example: Configuring a Private VLAN Spanning Multiple EX Series Switches*
- *Example: Setting Up Q-in-Q Tunneling on EX Series Switches*
- Understanding Bridging and VLANs on Switches on page 9
- **show ethernet-switching interfaces on page 198**

List of Sample Output

Copyright © 2018, Juniper Networks, Inc.

show vlans (SRX devices) on page 250
show vlans brief (SRX devices) on page 250
show vlans detail (SRX devices) on page 250

Output Fields

Table 11 on page 234 lists the output fields for the **show vlans** command. Output fields are listed in the approximate order in which they appear.

Table 11: show vlans Output Fields

Field Name	Field Description	Level of Output
Name	Name of a VLAN.	none, **brief**
Tag	802.1Q tag applied to this VLAN. If **none** is displayed, no tag is applied.	All levels
Interfaces	Interface associated with learned MAC addresses or **All-members** option (flood entry). An asterisk (*****) beside the interface indicates that the interface is **UP**.	All levels
Address	IP address.	none, **brief**
Ports Active /Total	Number of interfaces associated with a VLAN: **Active** indicates interfaces that are **UP**, and **Total** indicates interfaces that are active and inactive.	**brief**
VLAN	Name of a VLAN.	**detail, extensive**
Admin state	State of the interface. Values are: **enabled**—The interface is turned on, and the physical link is operational and can pass packets.	**detail,extensive**
MAC learning Status	Indicates if MAC learning is disabled.	**detail, extensive**
Description	Description for the VLAN.	**detail,extensive**
Primary IP	Primary IP address associated with a VLAN.	**detail**
Number of interfaces	Number of interfaces associated with a VLAN. Both the total number of interfaces and the number of active interfaces associated with a VLAN are displayed.	**detail, extensive**
STP	Spanning tree associated with a VLAN.	**detail,extensive**
Tagged interfaces	Tagged interfaces with which a VLAN is associated.	**detail,extensive**
Untagged interfaces	Untagged interfaces with which a VLAN is associated.	**detail. extensive**
Dot1q Tunneling Status	Indicates if Q-in-Q tunneling is enabled.	**extensive**
Customer VLAN ranges	List of customer VLAN (C-VLAN) ranges associated with this service VLAN (S-VLAN).	**extensive**

Copyright © 2018, Juniper Networks, Inc.

Table 11: show vlans Output Fields (continued)

Field Name	Field Description	Level of Output
Private VLAN Mode	The private VLAN mode for this VLAN. Values include **Primary**, **Isolated**, and **Community**.	**extensive**
Primary VLAN	Primary VLAN tag for this secondary VLAN.	**extensive**
Internal Index	VLAN index internal to Junos OS software.	**extensive**
Origin	Manner in which the VLAN was created: **static** or **learn**.	**extensive**
Protocol	Port-based VLAN or MAC-based VLAN. MAC-based protocol is displayed when VLAN assignment is done either statically or dynamically through 802.1X,	**extensive**
IP addresses	IP address associated with a VLAN.	**extensive**
Number of MAC entries	For MAC-based VLANs created either statically or dynamically, the MAC addresses associated with an interface.	**extensive**
Number of mapping rules	Number of mapping rules for Q-in-Q tunneling (**Push**) and VLAN translation (**Swap**).	
Secondary VLANs	Secondary VLANs associated with a primary VLAN.	**extensive**
Isolated VLANs	Isolated VLANs associated with a primary VLAN.	**extensive**
Community VLANs	Community VLANs associated with a primary VLAN.	**extensive**
VLANs summary	VLAN counts: • **Total**—Total number of VLANs on the switch. • **Configured VLANs**—Number of VLANs that are based on user-configured settings. • **Internal VLANs**—Number of VLANs created by the system with no explicit configuration or protocol—for example, the **default** VLAN and the VLAN created when a trunk interface is not configured with native VLAN membership. • **Temporary VLANs**—Number of VLANs from the previous configuration that the system retains for a limited time after restart. Temporary VLANs are converted into one of the other types of VLAN, or are removed from the system if the current configuration does not require them.	All levels

Copyright © 2018, Juniper Networks, Inc.

Table 11: show vlans Output Fields (continued)

Field Name	Field Description	Level of Output

Dot1q VLANs summary	802.1Q VLAN counts:
	• **Total**—Total number of 802.1Q-tagged and untagged VLANs on the switch.
	• **Tagged VLANs**—Number of 802.1Q-tagged VLANs.
	• **Untagged VLANs**—Number of untagged 802.1Q VLANs.
	• **Private VLAN**—Counts of the following kinds of 802.1Q private VLANs (PVLANs):
	  ◦ **Primary VLANs**—Number of primary forwarding private VLANs.
	  ◦ **Community VLANs**—Number of community transporting and forwarding private VLANs.
	  ◦ **Isolated VLANs**—Number of isolated receiving and forwarding private VLANs.
	  ◦ **Inter–switch–isolated VLANs**—Number of inter-switch isolated receiving and forwarding private VLANs.
		All levels

Dot1q Tunneled VLANs summary	Q-in-Q-tunneled VLAN counts: • **Total**—Total number of Q-in-Q-tunneled VLANs on the switch. • **Private VLAN**—Counts of primary, community, and isolated Q-in-Q-tunneled private VLANs (PVLANs).	All levels

Sample Output (EX Series and QFX Series Switches)

show vlans (EX Series and QFX Series)

user@switch> **show vlans**

```

Name           Tag     Interfaces
default        None
                       xe-0/0/34.0, xe-0/0/33.0, xe-0/0/32.0, xe-0/0/31.0,
                       xe-0/0/30.0, xe-0/0/29.0, xe-0/0/28.0, xe-0/0/27.0,
                       xe-0/0/26.0, xe-0/0/25.0, xe-0/0/19.0, xe-0/0/18.0,
                       xe-0/0/17.0, xe-0/0/16.0, xe-0/0/15.0, xe-0/0/14.0,
                       xe-0/0/13.0, xe-0/0/11.0, xe-0/0/9.0, xe-0/0/8.0,
                       xe-0/0/3.0, xe-0/0/2.0, xe-0/0/1.0
v0001          1
                       xe-0/0/24.0, xe-0/0/23.0, xe-0/0/22.0, xe-0/0/21.0
v0002          2
                       None
v0003          3
                       None
v0004          4
                       None
v0005          5
                       None
```

show vlans (Private VLANs on EX and QFX Series)

user@switch> **show vlans**

 Copyright © 2018, Juniper Networks, Inc.

```
Name           Tag     Interfaces
__pvlan_pvlan_xe-0/0/46.0__
                       xe-0/0/44.0*, xe-0/0/46.0*
c1
                       xe-0/0/4.0*, xe-0/0/44.0*
c2
                       xe-0/0/28.0*, xe-0/0/44.0*
default
                       None
pvlan          500
                       xe-0/0/4.0*, xe-0/0/28.0*, xe-0/0/44.0*, xe-0/0/46.0*
```

show vlans brief (EX and QFX Series)

user@switch> **show vlans brief**

```
                                           Ports
Name           Tag     Address             Active/Total
default        None                        0/23
v0001          1                           0/4
v0002          2                           0/0
v0003          3                           0/0
v0004          4                           0/0
v0005          5                           0/0
v0006          6                           0/0
v0007          7                           0/0
v0008          8                           0/0
v0009          9                           0/0
v0010          10                          0/2
v0011          11                          0/0
v0012          12                          0/0
v0013          13                          0/0
v0014          14                          0/0
v0015          15                          0/0
v0016          16                          0/0
```

show vlans detail (EX Series and QFX Series)

user@switch> **show vlans detail**

```
VLAN: default, Tag: Untagged, Admin state: Enabled
  Description: None
  Primary IP: None, Number of interfaces: 23 (Active = 0)
  STP: None, RTG: None
  Untagged interfaces: xe-0/0/34.0, xe-0/0/33.0, xe-0/0/32.0, xe-0/0/31.0,
  xe-0/0/30.0, xe-0/0/29.0, xe-0/0/28.0, xe-0/0/27.0, xe-0/0/26.0,
  xe-0/0/25.0, xe-0/0/19.0, xe-0/0/18.0, xe-0/0/17.0, xe-0/0/16.0,
  xe-0/0/15.0, xe-0/0/14.0, xe-0/0/13.0, xe-0/0/11.0, xe-0/0/9.0, xe-0/0/8.0,
  xe-0/0/3.0, xe-0/0/2.0, xe-0/0/1.0,
  Tagged interfaces: None

VLAN: v0001, Tag: 802.1Q Tag 1, Admin state: Enabled
  Description: None
  Primary IP: None, Number of interfaces: 4 (Active = 0)
  Dot1q Tunneling Status: Enabled
  STP: None, RTG: None
  Untagged interfaces: None
  Tagged interfaces: xe-0/0/24.0, xe-0/0/23.0, xe-0/0/22.0, xe-0/0/21.0,
```

Copyright © 2018, Juniper Networks, Inc.

```
VLAN: v0002, Tag: 802.1Q Tag 2, Admin state: Enabled
  Description: None
  Primary IP: None, Number of interfaces: 0 (Active = 0)
  STP: None, RTG: None
  Untagged interfaces: None
  Tagged interfaces: None

VLAN: v0003, Tag: 802.1Q Tag 3, Admin state: Enabled
  Description: None
  Primary IP: None, Number of interfaces: 0 (Active = 0)
  STP: None, RTG: None
  Untagged interfaces: None
  Tagged interfaces: None

VLAN: vlan4000, 802.1Q Tag: Untagged, Admin State: Enabled
MAC learning Status: Disabled
Number of interfaces: 0 (Active = 0)
```

show vlans extensive (for a PVLAN spanning multiple switches)

user@switch> **show vlans extensive**

```
VLAN: COM1, Created at: Tue May 11 18:16:05 2010
802.1Q Tag: 100, Internal index: 3, Admin State: Enabled, Origin: Static
Private VLAN Mode: Community, Primary VLAN: primary
Protocol: Port Mode, Mac aging time: 300 seconds
Number of interfaces: Tagged 3 (Active = 3), Untagged  1 (Active = 1)
      ge-0/0/20.0*, tagged, trunk
      ge-0/0/22.0*, tagged, trunk, pvlan-trunk
      ge-0/0/23.0*, tagged, trunk, pvlan-trunk
      ge-0/0/7.0*, untagged, access

VLAN: __pvlan_primary_ge-0/0/0.0__, Created at: Tue May 11 18:16:05 2010
Internal index: 5, Admin State: Enabled, Origin: Static
Private VLAN Mode: Isolated, Primary VLAN: primary
Protocol: Port Mode, Mac aging time: 300 seconds
Number of interfaces: Tagged 3 (Active = 3), Untagged  1 (Active = 1)
      ge-0/0/20.0*, tagged, trunk
      ge-0/0/22.0*, tagged, trunk, pvlan-trunk
      ge-0/0/23.0*, tagged, trunk, pvlan-trunk
      ge-0/0/0.0*, untagged, access


VLAN: __pvlan_primary_ge-0/0/2.0__, Created at: Tue May 11 18:16:05 2010
Internal index: 6, Admin State: Enabled, Origin: Static
Private VLAN Mode: Isolated, Primary VLAN: primary
Protocol: Port Mode, Mac aging time: 300 seconds
Number of interfaces: Tagged 3 (Active = 3), Untagged  1 (Active = 0)
      ge-0/0/20.0*, tagged, trunk
      ge-0/0/22.0*, tagged, trunk, pvlan-trunk
      ge-0/0/23.0*, tagged, trunk, pvlan-trunk
      ge-0/0/2.0, untagged, access


VLAN: __pvlan_primary_isiv__, Created at: Tue May 11 18:16:05 2010
802.1Q Tag: 50, Internal index: 7, Admin State: Enabled, Origin: Static
Private VLAN Mode: Inter-switch-isolated, Primary VLAN: primary
Protocol: Port Mode, Mac aging time: 300 seconds
Number of interfaces: Tagged 3 (Active = 3), Untagged  0 (Active = 0)
```


Copyright © 2018, Juniper Networks, Inc.

```
      ge-0/0/20.0*, tagged, trunk
      ge-0/0/22.0*, tagged, trunk, pvlan-trunk
      ge-0/0/23.0*, tagged, trunk, pvlan-trunk


VLAN: community2, Created at: Tue May 11 18:16:05 2010
802.1Q Tag: 20, Internal index: 8, Admin State: Enabled, Origin: Static
Private VLAN Mode: Community, Primary VLAN: primary
Protocol: Port Mode, Mac aging time: 300 seconds
Number of interfaces: Tagged 3 (Active = 3), Untagged  2 (Active = 2)
      ge-0/0/20.0*, tagged, trunk
      ge-0/0/22.0*, tagged, trunk, pvlan-trunk
      ge-0/0/23.0*, tagged, trunk, pvlan-trunk
      ge-0/0/1.0*, untagged, access
      ge-1/0/6.0*, untagged, access


VLAN: primary, Created at: Tue May 11 18:16:05 2010
802.1Q Tag: 10, Internal index: 2, Admin State: Enabled, Origin: Static
Private VLAN Mode: Primary
Protocol: Port Mode, Mac aging time: 300 seconds
Number of interfaces: Tagged 3 (Active = 3), Untagged  5 (Active = 4)
      ge-0/0/20.0*, tagged, trunk
      ge-0/0/22.0*, tagged, trunk, pvlan-trunk
      ge-0/0/23.0*, tagged, trunk, pvlan-trunk
      ge-0/0/0.0*, untagged, access
      ge-0/0/1.0*, untagged, access
      ge-0/0/2.0, untagged, access
      ge-0/0/7.0*, untagged, access
      ge-1/0/6.0*, untagged, access

Secondary VLANs: Isolated 2, Community  2, Inter-switch-isolated  1
  Isolated VLANs :
      __pvlan_primary_ge-0/0/0.0__
      __pvlan_primary_ge-0/0/2.0__
  Community VLANs :
      COM1
      community2
  Inter-switch-isolated VLAN :
      __pvlan_primary_isiv__
```

show vlans extensive (Port-Based on EX Series and QFX Series)

`user@switch>` **show vlans extensive**

```
VLAN: default, created at Mon Feb  4 12:13:47 2008
  Tag: None, Internal index: 0, Admin state: Enabled, Origin: static
  Description: None
  Customer VLAN ranges:
                   1-4100
    Protocol: Port based
  IP addresses: None
  STP: None, RTG: None.
  Number of interfaces: Tagged 0 (Active = 0), Untagged 23 (Active = 0)
    xe-0/0/34.0 (untagged, access)
    xe-0/0/33.0 (untagged, access)
    xe-0/0/32.0 (untagged, access)
    xe-0/0/31.0 (untagged, access)
    xe-0/0/30.0 (untagged, access)
```

Copyright © 2018, Juniper Networks, Inc.

```
    xe-0/0/29.0 (untagged, access)
    xe-0/0/28.0 (untagged, access)
    xe-0/0/27.0 (untagged, access)
    xe-0/0/26.0 (untagged, access)
    xe-0/0/25.0 (untagged, access)
    xe-0/0/19.0 (untagged, access)
    xe-0/0/18.0 (untagged, access)
    xe-0/0/17.0 (untagged, access)
    xe-0/0/16.0 (untagged, access)
    xe-0/0/15.0 (untagged, access)
    xe-0/0/14.0 (untagged, access)
    xe-0/0/13.0 (untagged, access)
    xe-0/0/11.0 (untagged, access)
    xe-0/0/9.0 (untagged, access)
    xe-0/0/8.0 (untagged, access)
    xe-0/0/3.0 (untagged, access)
    xe-0/0/2.0 (untagged, access)
    xe-0/0/1.0 (untagged, access)

Secondary VLANs: Isolated 1, Community  1
  Isolated VLANs :
      __pvlan_pvlan_xe-0/0/3.0__
  Community VLANs :
      comm1

VLAN: v0001, created at Mon Feb  4 12:13:47 2008
  Tag: 1, Internal index: 1, Admin state: Enabled, Origin: static
  Description: None
  Protocol: Port based, Layer 3 interface: None
  IP addresses: None
  STP: None, RTG: None.
  Number of interfaces: Tagged 4 (Active = 0), Untagged 0 (Active = 0)
    xe-0/0/24.0 (tagged, trunk)
    xe-0/0/23.0 (tagged, trunk)
    xe-0/0/22.0 (tagged, trunk)
    xe-0/0/21.0 (tagged, trunk)

VLAN: v0002, created at Mon Feb  4 12:13:47 2008
  Tag: 2, Internal index: 2, Admin state: Enabled, Origin: static
  Description: None
  Protocol: Port based, Layer 3 interface: None
  IP addresses: None
  STP: None, RTG: None.
  Number of interfaces: Tagged 0 (Active = 0), Untagged 0 (Active = 0)
    None

VLAN: v0003, created at Mon Feb  4 12:13:47 2008
  Tag: 3, Internal index: 3, Admin state: Enabled, Origin: static
  Description: None
  Protocol: Port based, Layer 3 interface: None
  IP addresses: None
  STP: None, RTG: None.
  Number of interfaces: Tagged 0 (Active = 0), Untagged 0 (Active = 0)
    None
```

show vlans extensive (MAC-based)

`user@switch>` **show vlans extensive**

Copyright © 2018, Juniper Networks, Inc.

```
VLAN: default, Created at: Thu May 15 13:43:09 2008
Internal index: 3, Admin State: Enabled, Origin: Static
Protocol: Port Mode, Mac aging time: 300 seconds
Number of interfaces: Tagged 0 (Active = 0), Untagged  2 (Active = 2)
      ge-0/0/0.0*, untagged, access
      ge-0/0/14.0*, untagged, access

VLAN: vlan_dyn, Created at: Thu May 15 13:43:09 2008
Internal index: 4, Admin State: Enabled, Origin: Static
Protocol: Port Mode
Number of interfaces: Tagged 0 (Active = 0), Untagged  0 (Active = 0)
Protocol: MAC Based
Number of MAC entries: 6
      ge-0/0/0.0*
                  00:00:00:00:00:02 (untagged)
                  00:00:00:00:00:03 (untagged)
                  00:00:00:00:00:04 (untagged)
                  00:00:00:00:00:05 (untagged)
                  00:00:00:00:00:06 (untagged)
                  00:00:00:00:00:07 (untagged)
```

show vlans (Q-in-Q Tunneling on EX Series and QFX Series)

user@switch> **show vlans dot1q-tunneling**

```
Name           Tag     Interfaces
sv100          100     xe-0/0/4.0*, xe-0/0/15.0*
```

show vlans extensive (Q-in-Q Tunneling on EX Series and QFX Series)

user@switch> **show vlans sv100 extensive**

```
VLAN: sv100, Created at: Sat Sep 10 12:53:52 2011
802.1Q Tag: 100, Internal index: 2, Admin State: Enabled, Origin: Static
Dot1q Tunneling Status: Enabled
Customer VLAN ranges:
                     10-20
                     40-50
Protocol: Port Mode
Number of interfaces: Tagged 1 (Active = 1), Untagged  0 (Active = 0)
      ge-0/0/0.0, tagged, trunk

Number of mapping rules:
      Push 1 (Active = 0), Policy 0 (Active = 0), Swap 0 (Active = 0)

      xe-0/0/3.0*, 300, push
```

show vlans extensive (Q-in-Q Tunneling and L2TP on EX Series and QFX Series)

user@switch> **show vlans v1 extensive**

```
VLAN: v1, Created at: Fri Mar 2 05:07:38 2012
802.1Q Tag: 100, Internal index: 4, Admin State: Enabled, Origin: Static
Dot1q Tunneling status: Enabled
Layer2 Protocol Tunneling status: Enabled
```

Copyright © 2018, Juniper Networks, Inc.

show vlans sort-by tag (EX Series and QFX Series)

user@switch> **show vlans sort-by tag**

```
Name           Tag     Interfaces
default                       None
__vlan-x_1__   1
                       None
__vlan-x_2__   2
                       None
__vlan-x_3__   3
                       None
__vlan-x_4__   4
                       None
__vlan-x_5__   5
                       None
__vlan-x_6__   6
                       None
__vlan-x_7__   7
                       None
__vlan-x_8__   8
                       None
__vlan-x_9__   9
                       None
__vlan-x_10__  10
                       None
__vlan-x_11__  11
                       None
__vlan-x_12__  12
                       None
__vlan-x_13__  13
                       None
__vlan-x_14__  14
                       None
__vlan-x_15__  15
                       None
__vlan-x_16__  16
                       None
__vlan-x_17__  17
                       None
__vlan-x_18__  18
                       None
__vlan-x_19__  19
                       None
__vlan-x_20__  20
                       None
```

show vlans sort-by name (EX Series and QFX Series)

user@switch> **show vlans sort-by employee**

```

Name             Tag   Interfaces

__employee_120__ 120
                       xe-0/0/22.0*
__employee_121__ 121
                       xe-0/0/22.0*
__employee_122__ 122
                       xe-0/0/22.0*
```


Copyright © 2018, Juniper Networks, Inc.

```
__employee_123__ 123
                        xe-0/0/22.0*
__employee_124__ 124
                        xe-0/0/22.0*
__employee_125__ 125
                        xe-0/0/22.0*
__employee_126__ 126
                        xe-0/0/22.0*
__employee_127__ 127
                        xe-0/0/22.0*
__employee_128__ 128
                        xe-0/0/22.0*
__employee_129__ 129
                        xe-0/0/22.0*
__employee_130__ 130
                        xe-0/0/22.0*
```

show vlans tag (EX Series and QFX Series)

`user@switch>` **show vlans employee**

```

Name             Tag    Interfaces

__employee_120__ 120
                        xe-0/0/22.0*
__employee_121__ 121
                        xe-0/0/22.0*
__employee_122__ 122
                        xe-0/0/22.0*
__employee_123__ 123
                        xe-0/0/22.0*
__employee_124__ 124
                        xe-0/0/22.0*
__employee_125__ 125
                        xe-0/0/22.0*
__employee_126__ 126
                        xe-0/0/22.0*
__employee_127__ 127
                        xe-0/0/22.0*
__employee_128__ 128
                        xe-0/0/22.0*
__employee_129__ 129
                        xe-0/0/22.0*
__employee_130__ 130
                        xe-0/0/22.0*
```

show vlans sort-by tag (EX Series)

`user@switch>` **show vlans sort-by tag**

```
Name           Tag      Interfaces
default                 None
__vlan-x_1__   1
                        None
__vlan-x_2__   2
                        None
__vlan-x_3__   3
                        None
```

Copyright © 2018, Juniper Networks, Inc.

```
__vlan-x_4__    4
                       None
__vlan-x_5__    5
                       None
__vlan-x_6__    6
                       None
__vlan-x_7__    7
                       None
__vlan-x_8__    8
                       None
__vlan-x_9__    9
                       None
__vlan-x_10__   10
                       None
__vlan-x_11__   11
                       None
__vlan-x_12__   12
                       None
__vlan-x_13__   13
                       None
__vlan-x_14__   14
                       None
__vlan-x_15__   15
                       None
__vlan-x_16__   16
                       None
__vlan-x_17__   17
                       None
__vlan-x_18__   18
                       None
__vlan-x_19__   19
                       None
__vlan-x_20__   20
                       None
```

show vlans employee (vlan-range-name)

user@switch> **show vlans employee**

```
Name                 Tag    Interfaces

__employee_120__  120
                       ge-0/0/22.0*
__employee_121__  121
                       ge-0/0/22.0*
__employee_122__  122
                       ge-0/0/22.0*
__employee_123__  123
                       ge-0/0/22.0*
__employee_124__  124
                       ge-0/0/22.0*
__employee_125__  125
                       ge-0/0/22.0*
__employee_126__  126
                       ge-0/0/22.0*
__employee_127__  127
                       ge-0/0/22.0*
__employee_128__  128
                       ge-0/0/22.0*
```

 Copyright © 2018, Juniper Networks, Inc.

```
__employee_129__ 129
                       ge-0/0/22.0*
__employee_130__ 130
                       ge-0/0/22.0*
```

show vlans summary (EX Series)

user@switch> **show vlans summary**

```
VLANs summary:
         Total: 8,   Configured VLANs: 5
         Internal VLANs: 1,  Temporary VLANs: 0

     Dot1q VLANs summary:
         Total: 8,  Tagged VLANs: 2, Untagged VLANs: 6
         Private VLAN:
             Primary VLANs: 2,  Community VLANs: 2, Isolated VLANs: 3

     Dot1q Tunneled VLANs summary:
         Total: 0
         Private VLAN:
             Primary VLANs: 0,  Community VLANs: 0, Isolated VLANs: 0

     Dynamic VLANs:
        Total: 2,  Dot1x: 2, MVRP: 0
```

Sample Output: EX Series with ELS Switches and MX Routers

show vlans brief (EX Series Switch)

user@switch> **show vlans brief**

```
Routing instance        VLAN name             Tag          Interfaces
default-switch          c1                    20
                                                           ge-0/0/0.0*
                                                           ge-1/0/0.0*
                                                           ge-2/0/0.0*
default-switch          c2                    30
                                                           ge-0/0/0.0*
                                                           ge-2/0/0.0*
default-switch          default               1

default-switch          iso                   10
                                                           ge-0/0/1.0*
default-switch          iso1                  50
                                                           ge-0/0/0.0*
                                                           ge-2/0/0.0*
default-switch          pri                   100
                                                           ge-0/0/0.0*
                                                           ge-1/0/0.0*
                                                           ge-2/0/0.0*
```

show vlans brief (MX Routers)

user@host> **show vlans brief**

```
Routing instance        VLAN name             Tag          Interfaces
VPLS-1                  __VPLS-1__               all
```

```
                                                                ae1.0
VPLS-2                   __VPLS-2__                all
                                                                ae3.0
                                                                ge-3/1/2.0
                                                                vt-3/3/10.1048576
default-switch           VLAN1000                  1000
                                                                ae26.0
default-switch           VLAN101                   101
                                                                ae20.0
default-switch           VLAN102                   102
                                                                ae20.0
default-switch           VLAN103                   103
                                                                ae20.0
default-switch           VLAN104                   104
                                                                ae20.0
default-switch           VLAN105                   105
                                                                ae20.0
default-switch           VLAN106                   106
                                                                ae20.0
default-switch           VLAN107                   107
                                                                ae20.0
default-switch           VLAN108                   108
                                                                ae20.0
[...output truncated...]
```

show vlans detail (EX Series Switch)

user@switch> **show vlans detail**

```
Routing instance: default-switch
  VLAN Name: c1                            State: Active
Tag: 20
PVLAN type : Community
Internal index: 16, Generation Index: 21, Origin: Static
MAC aging time: 300 seconds
Interfaces:
    ge-0/0/0.0*,tagged,trunk
    ge-1/0/0.0*,tagged,trunk
    ge-2/0/0.0*,tagged,trunk
Number of interfaces: Tagged 3    , Untagged 0
Total MAC count: 0

Routing instance: default-switch
  VLAN Name: c2                            State: Active
Tag: 30
PVLAN type : Community
Internal index: 17, Generation Index: 22, Origin: Static
MAC aging time: 300 seconds
Interfaces:
    ge-0/0/0.0*,tagged,trunk
    ge-2/0/0.0*,tagged,trunk
Number of interfaces: Tagged 2    , Untagged 0
Total MAC count: 0

Routing instance: default-switch
  VLAN Name: default                       State: Active
Tag: 1
Internal index: 5, Generation Index: 5, Origin: Static
MAC aging time: 300 seconds
Number of interfaces: Tagged 0    , Untagged 0
```

 Copyright © 2018, Juniper Networks, Inc.

```
Total MAC count: 0

Routing instance: default-switch
  VLAN Name: iso                            State: Active
Tag: 10
Internal index: 14, Generation Index: 19, Origin: Static
MAC aging time: 300 seconds
Interfaces:
    ge-0/0/1.0*,untagged,access
Number of interfaces: Tagged 0    , Untagged 1
Total MAC count: 0

Routing instance: default-switch
  VLAN Name: iso1                           State: Active
Tag: 50
PVLAN type : Isolated
Internal index: 15, Generation Index: 20, Origin: Static
MAC aging time: 300 seconds
Interfaces:
    ge-0/0/0.0*,tagged,trunk
    ge-2/0/0.0*,tagged,trunk
Number of interfaces: Tagged 2    , Untagged 0
Total MAC count: 0

Routing instance: default-switch
  VLAN Name: pri                            State: Active
Tag: 100
PVLAN type : Primary
Isolated VLAN :
vlan-id : 50 vlan name : iso1
Community VLAN :
vlan-id : 20 vlan name : c1
vlan-id : 30 vlan name : c2
Internal index: 9, Generation Index: 14, Origin: Static
MAC aging time: 300 seconds
Interfaces:
    ge-0/0/0.0*,tagged,trunk
    ge-1/0/0.0*,tagged,trunk
    ge-2/0/0.0*,tagged,trunk
Number of interfaces: Tagged 3    , Untagged 0
Total MAC count: 0
```

show vlans detail (MX Routers)

user@host> **show vlans detail**

```
Routing instance: VPLS-1
  VLAN Name: __VPLS-1__                     State: Active
Tag: all
Internal index: 2, Generation Index:  , Origin: Dynamic
Interfaces:
    ae1.0,tagged
Number of interfaces: Tagged 1    , Untagged 0
Total MAC count: 0

Routing instance: VPLS-2
  VLAN Name: __VPLS-2__                     State: Active
Tag: all
Internal index: 3, Generation Index:  , Origin: Dynamic
```

Copyright © 2018, Juniper Networks, Inc.

```
Interfaces:
    ae3.0,tagged
    ge-3/1/2.0,tagged
    vt-3/3/10.1048576,tagged
Number of interfaces: Tagged 3    , Untagged 0
Total MAC count: 4

Routing instance: default-switch
  VLAN Name: VLAN1000                        State: Active
Tag: 1000
Internal index: 4, Generation Index: 1, Origin: Static
Layer 3 interface: irb.1000
Interfaces:
    ae26.0,tagged,trunk
Number of interfaces: Tagged 1    , Untagged 0
Total MAC count: 0

Routing instance: default-switch
  VLAN Name: VLAN101                         State: Active
Tag: 101
Internal index: 5, Generation Index: 2, Origin: Static
Layer 3 interface: irb.101
Interfaces:
    ae20.0,tagged,trunk
Number of interfaces: Tagged 1    , Untagged 0
Total MAC count: 1

Routing instance: default-switch
  VLAN Name: VLAN102                         State: Active
Tag: 102
Internal index: 6, Generation Index: 3, Origin: Static
Layer 3 interface: irb.102
Interfaces:
    ae20.0,tagged,trunk
Number of interfaces: Tagged 1    , Untagged 0
Total MAC count: 1
[...output truncated...]
```

show vlans extensive (EX Series Switch)

`user@switch>` **show vlans extensive**

```
Routing instance: default-switch
  VLAN Name: c1                              State: Active
Tag: 20
PVLAN type : Community
Internal index: 16, Generation Index: 21, Origin: Static
MAC aging time: 300 seconds
Interfaces:
    ge-0/0/0.0*,tagged,trunk
    ge-1/0/0.0*,tagged,trunk
    ge-2/0/0.0*,tagged,trunk
Number of interfaces: Tagged 3    , Untagged 0
Total MAC count: 0

Routing instance: default-switch
  VLAN Name: c2                              State: Active
Tag: 30
PVLAN type : Community
Internal index: 17, Generation Index: 22, Origin: Static
```

 Copyright © 2018, Juniper Networks, Inc.

```
MAC aging time: 300 seconds
Interfaces:
    ge-0/0/0.0*,tagged,trunk
    ge-2/0/0.0*,tagged,trunk
Number of interfaces: Tagged 2    , Untagged 0
Total MAC count: 0

Routing instance: default-switch
  VLAN Name: default                        State: Active
Tag: 1
Internal index: 5, Generation Index: 5, Origin: Static
MAC aging time: 300 seconds
Number of interfaces: Tagged 0    , Untagged 0
Total MAC count: 0

Routing instance: default-switch
  VLAN Name: iso                            State: Active
Tag: 10
Internal index: 14, Generation Index: 19, Origin: Static
MAC aging time: 300 seconds
Interfaces:
    ge-0/0/1.0*,untagged,access
Number of interfaces: Tagged 0    , Untagged 1
Total MAC count: 0

Routing instance: default-switch
  VLAN Name: iso1                           State: Active
Tag: 50
PVLAN type : Isolated
Internal index: 15, Generation Index: 20, Origin: Static
MAC aging time: 300 seconds
Interfaces:
    ge-0/0/0.0*,tagged,trunk
    ge-2/0/0.0*,tagged,trunk
Number of interfaces: Tagged 2    , Untagged 0
Total MAC count: 0

Routing instance: default-switch
  VLAN Name: pri                            State: Active
Tag: 100
PVLAN type : Primary
Isolated VLAN :
vlan-id : 50 vlan name : iso1
Community VLAN :
vlan-id : 20 vlan name : c1
vlan-id : 30 vlan name : c2
Internal index: 9, Generation Index: 14, Origin: Static
MAC aging time: 300 seconds
Interfaces:
    ge-0/0/0.0*,tagged,trunk
    ge-1/0/0.0*,tagged,trunk
    ge-2/0/0.0*,tagged,trunk
Number of interfaces: Tagged 3    , Untagged 0
Total MAC count: 0
```

show vlans extensive (MX Routers)

`user@host>` **show vlans extensive**

Copyright © 2018, Juniper Networks, Inc.

```
Routing instance: default-switch
  VLAN Name: VLAN_10                        State: Active
Tag: 10
Internal index: 2, Generation Index: 1, Origin: Static
MAC aging time: 300 seconds
Interfaces:
    ge-1/0/3.0*,tagged,trunk
Number of interfaces: Tagged 1    , Untagged 0
Total MAC count: 0

Routing instance: default-switch
  VLAN Name: VLAN_20                        State: Active
Tag: 20
Internal index: 3, Generation Index: 2, Origin: Static
MAC aging time: 300 seconds
Interfaces:
    ge-1/0/3.0*,tagged,trunk
Number of interfaces: Tagged 1    , Untagged 0
Total MAC count: 0
```

Sample Output (SRX Devices)

show vlans (SRX devices)

user@host> **show vlans**

```
Routing instance        VLAN name             Tag          Interfaces
default-switch          vlan-22                  22

default-switch          vlan-333                 333
                                                           ge-0/0/3.0*
                                                           ge-0/0/4.0*
default-switch          default               1

default-switch          vlan100               100
                                                           ge-0/0/1.0*
```

show vlans brief (SRX devices)

user@host> **show vlans brief**

```
Routing instance        VLAN name             Tag          Interfaces
default-switch          vlan-22                  22

default-switch          vlan-333                 333
                                                           ge-0/0/3.0*
                                                           ge-0/0/4.0*
default-switch          default               1

default-switch          vlan100               100
                                                           ge-0/0/1.0*
```

show vlans detail (SRX devices)

user@host> **show vlans detail**

```
Routing instance: default-switch
  VLAN Name: vlan-22                           State: Active
```


Copyright © 2018, Juniper Networks, Inc.

```
Tag: 22
Internal index: 2, Generation Index: 1, Origin: Static
MAC aging time: 300 seconds
VXLAN Enabled : No
Number of interfaces: Tagged 0    , Untagged 0
Total MAC count: 0

Routing instance: default-switch
  VLAN Name: vlan-333                          State: Active
Tag: 333
Internal index: 3, Generation Index: 2, Origin: Static
MAC aging time: 300 seconds
VXLAN Enabled : No
Interfaces:
    ge-0/0/3.0*,tagged,trunk
    ge-0/0/4.0*,tagged,trunk
Number of interfaces: Tagged 2    , Untagged 0
Total MAC count: 0
```

Copyright © 2018, Juniper Networks, Inc.

Copyright © 2018, Juniper Networks, Inc.

CHAPTER 8

MAC Address Operational Commands

show ethernet-switching mac-learning-log

Syntax

```
show ethernet-switching mac-learning-log
```

Release Information

Command introduced in Junos OS Release 9.0 for EX Series switches.
Command introduced in Junos OS Release 9.5 for SRX Series devices.
Command introduced in Junos OS Release 11.1 for the QFX Series.

Description

Displays the event log of learned MAC addresses.

Required Privilege Level

view

Related Documentation

- show ethernet-switching table on page 206
- show ethernet-switching interfaces on page 198
- show ethernet-switching table on page 206
- show ethernet-switching interfaces on page 198

- *Example: Setting Up Basic Bridging and a VLAN for an EX Series Switch*
- *Example: Setting Up Bridging with Multiple VLANs for EX Series Switches*
- *Example: Connecting an EX Series Access Switch to a Distribution Switch*

List of Sample Output

show ethernet-switching mac-learning-log (EX Series switch) on page 256
show ethernet-switching mac-learning-log (QFX Series Switches, QFabric, NFX Series Devices and EX4600) on page 256
show ethernet-switching mac-learning-log (SRX Series devices) on page 257

Output Fields

Output fields for EX Series switches:

The following table lists the output fields for the **show ethernet-switching mac-learning-log** command. Output fields are listed in the approximate order in which they appear.

Table 12: show ethernet-switching mac-learning-log Output Fields

Field Name	Field Description
Date and Time	Timestamp when the MAC address was added or deleted from the log.
vlan_name	VLAN name. A value defined by the user for all user-configured VLANs.
MAC	Learned MAC address.
Deleted | Added	MAC address deleted or added to the MAC learning log.

Copyright © 2018, Juniper Networks, Inc.

Table 12: show ethernet-switching mac-learning-log Output Fields (continued)

Field Name	Field Description
Blocking	The forwarding state of the interface: • **blocked**—Traffic is not being forwarded on the interface. • **unblocked**—Traffic is forwarded on the interface.
Flags	Displays the MAC address flags in which the MAC event occurred. This option is for debugging purposes.

Output fields for QFX Series switches, QFabric, NFX Series devices and EX4600:

Table 13 on page 255 lists the output fields for the **show ethernet-switching mac-learning-log** command. Output fields are listed in the approximate order in which they appear.

Table 13: show ethernet-switching mac-learning-log Output Fields

Field Name	Field Description
Date and Time	Timestamp in UTC when the MAC operation occurred.
vlan_name	VLAN name. A value defined by the user for all user-configured VLANs. The name of the VLAN on which the MAC is learned.
MAC	Learned MAC address.
Event op	MAC address that are added, learned, deleted, changed or moved from one interface to another interface.
Interface Name	The name of the interface on which the MAC address is learned. When a MAC address is moved, there is another field with the name of the interface. The log displays the name of the interface from where the MAC address moved, and the name of the interface to where the MAC address moved.
Flags	Displays the MAC address flags in which the MAC event occurred. This option is for debugging purposes.

Output fields for SRX Series devices:

Table 14 on page 255 lists the output fields for the show ethernet-switching mac-learning-log command on SRX Series devices. Output fields are listed in the approximate order in which they appear.

Table 14: show ethernet-switching-mac-learning-log Output Fields

Field Name	Field Description
Date and Time	Timestamp when the MAC address was added or deleted from the log.
VLAN-IDX	VLAN index. An internal value assigned by Junos OS for each VLAN.
MAC	Learned MAC address.
Deleted | Added	MAC address deleted or added to the MAC learning log.

Copyright © 2018, Juniper Networks, Inc.

Table 14: show ethernet-switching-mac-learning-log Output Fields (continued)

Field Name	Field Description
Blocking	The forwarding state of the interface: • blocked—Traffic is not being forwarded on the interface. • unblocked—Traffic is forwarded on the interface.

Sample Output

show ethernet-switching mac-learning-log (EX Series switch)

user@switch> **show ethernet-switching mac-learning-log**

```
Mon Feb 25 08:07:05 2008
 vlan_name v1 mac 00:00:00:00:00:00 was deleted
Mon Feb 25 08:07:05 2008
 vlan_name v9 mac 00:00:00:00:00:00 was deleted
Mon Feb 25 08:07:05 2008
 vlan_name HR_vlan mac 00:00:00:00:00:00 was deleted
Mon Feb 25 08:07:05 2008
 vlan_name v3 mac 00:00:00:00:00:00 was deleted
Mon Feb 25 08:07:05 2008
 vlan_name v12 mac 00:00:00:00:00:00 was deleted
Mon Feb 25 08:07:05 2008
 vlan_name v13 mac 00:00:00:00:00:00 was deleted
Mon Feb 25 08:07:05 2008
 vlan_name sales_vlan mac 00:00:00:00:00:00 was deleted
Mon Feb 25 08:07:05 2008
 vlan_name employee1 mac 00:00:00:00:00:00 was deleted
Mon Feb 25 08:07:05 2008
 vlan_name employee2 mac 00:00:00:00:00:00 was deleted
Mon Feb 25 08:07:05 2008
 vlan_name v3 mac 00:00:00:00:00:00 was added
Mon Feb 25 08:07:05 2008
 vlan_name HR_vlan mac 00:00:00:00:00:00 was added
Mon Feb 25 08:07:05 2008
 vlan_name employee2 mac 00:00:00:00:00:00 was added
Mon Feb 25 08:07:05 2008
 vlan_name employee1 mac 00:00:00:00:00:00 was added
Mon Feb 25 08:07:05 2008
 vlan_name employee2 mac 00:00:05:00:00:05 was learned
Mon Feb 25 08:07:05 2008
 vlan_name employee1 mac 00:30:48:90:54:89 was learned
Mon Feb 25 08:07:05 2008
 vlan_name HR_vlan mac 00:00:5e:00:01:00 was learned
Mon Feb 25 08:07:05 2008
 vlan_name sales_vlan mac 00:00:5e:00:01:08 was learned
[output truncated]
```

show ethernet-switching mac-learning-log (QFX Series Switches, QFabric, NFX Series Devices and EX4600)

user@switch> **show ethernet-switching mac-learning-log**

```
Mon Jun 30 13:49:49 2014 vlan_name v11+11 mac 00:10:94:00:00:02 was learned on
ge-1/0/22.0 with flags: 0x2001f     << MAC address that as dynamically learned
Mon Jun 30 13:50:29 2014 vlan_name v11+11 mac 00:10:94:00:00:02 was deleted from
 ge-1/0/22.0 with flags: 0x1080  << MAC address that was deleted
```

 Copyright © 2018, Juniper Networks, Inc.

```
Mon Jun 30 13:51:28 2014 vlan_name v11+11 mac 00:00:00:01:01:01 was added to
ge-1/0/22.0 with flags: 0x2013f        << Static MAC address that was added
Mon Jun 30 13:51:46 2014 vlan_name v11+11 mac 00:00:00:01:01:01 was deleted from
 ge-1/0/22.0 with flags: 0x1120   << delete of Static MAC address that was deleted
Mon Jun 30 13:52:03 2014 vlan_name v11+11 mac 00:10:94:00:00:02 was learned on
ge-1/0/22.0 with flags: 0x2001f    << MAC address that was dynamically learned
Mon Jun 30 13:52:11 2014 vlan_name v11+11 mac 00:10:94:00:00:02 was moved from
ge-1/0/22.0 to ge-1/0/21.0 with flags: 0x2101f    << MAC address that was moved
Mon Jun 30 13:54:24 2014 vlan_name v11+11 mac 00:10:94:00:00:02 was changed on
ge-1/0/21.0 with flags: 0x2113f   << MAC address that changed from a dynamic
address to a static address
```

show ethernet-switching mac-learning-log (SRX Series devices)

`user@host>` **show ethernet-switching mac-learning-log**

```
Wed Mar 18 08:07:05 2009
vlan_idx 7 mac 00:00:5E:00:53:00 was deleted
Wed Mar 18 08:07:05 2009
vlan_idx 9 mac 00:00:5E:00:53:00 was deleted
Wed Mar 18 08:07:05 2009
vlan_idx 10 mac 00:00:5E:00:53:00 was deleted
Wed Mar 18 08:07:05 2009
vlan_idx 11 mac 00:00:5E:00:53:00 was deleted
Wed Mar 18 08:07:05 2009
vlan_idx 12 mac 00:00:5E:00:53:00 was deleted
Wed Mar 18 08:07:05 2009
vlan_idx 13 mac 00:00:5E:00:53:00 was deleted
Wed Mar 18 08:07:05 2009
vlan_idx 14 mac 00:00:5E:00:53:00 was deleted
Wed Mar 18 08:07:05 2009
vlan_idx 15 mac 00:00:5E:00:53:00 was deleted
Wed Mar 18 08:07:05 2009
vlan_idx 16 mac 00:00:5E:00:53:00 was deleted
Wed Mar 18 08:07:05 2009
vlan_idx 4 mac 00:00:5E:00:53:00 was added
Wed Mar 18 08:07:05 2009
vlan_idx 6 mac 00:00:5E:00:53:00 was added
Wed Mar 18 08:07:05 2009
vlan_idx 7 mac 00:00:5E:00:53:00 was added
Wed Mar 18 08:07:05 2009
vlan_idx 9 mac 00:00:5E:00:53:00 was added
Wed Mar 18 08:07:05 2009
vlan_idx 10 mac 00:00:5E:00:53:00 was added
Wed Mar 18 08:07:05 2009
vlan_idx 11 mac 00:00:5E:00:53:00 was added
Wed Mar 18 08:07:05 2009
vlan_idx 12 mac 00:00:5E:00:53:00 was added
Wed Mar 18 08:07:05 2009
vlan_idx 13 mac 00:00:5E:00:53:00 was added
Wed Mar 18 08:07:05 2009
vlan_idx 14 mac 00:00:5E:00:53:00 was added
Wed Mar 18 08:07:05 2009
vlan_idx 15 mac 00:00:5E:00:53:00 was added
Wed Mar 18 08:07:05 2009
vlan_idx 16 mac 00:00:5E:00:53:00 was added
Wed Mar 18 08:07:05 2009
vlan_idx 5 mac 00:00:5E:00:53:00 was added
Wed Mar 18 08:07:05 2009
vlan_idx 18 mac 00:00:5E:00:53:AA was learned
```

Copyright © 2018, Juniper Networks, Inc.

```
Wed Mar 18 08:07:05 2009
vlan_idx 5 mac 00:00:5E:00:53:AB was learned
Wed Mar 18 08:07:05 2009
vlan_idx 6 mac 00:00:5E:00:53:AC was learned
Wed Mar 18 08:07:05 2009
vlan_idx 16 mac 00:00:5E:00:53:AD was learned
Wed Mar 18 08:07:05 2009
vlan_idx 7 mac 00:00:5E:00:53:AE was learned
Wed Mar 18 08:07:05 2009
vlan_idx 8 mac 00:00:5E:00:53:AF was learned
Wed Mar 18 08:07:05 2009
vlan_idx 12 mac 00:00:5E:00:53:AG was learned
[output truncated]
```

 Copyright © 2018, Juniper Networks, Inc.

show ethernet-switching statistics mac-learning

Syntax

```
show ethernet-switching statistics mac-learning
<brief | detail>
<interface interface-name>
```

Release Information Command introduced in Junos OS Release 9.4 for EX Series switches.
Command introduced in Junos OS Release 11.1 for the QFX Series.

Description Display media access control (MAC) learning statistics.

> NOTE: For the QFX Series, this command is not supported in Enhanced Layer 2 Software (ELS).

Options **none**—(Optional) Display MAC learning statistics for all interfaces.

brief | detail—(Optional) Display the specified level of output. The default is **brief**.

interface *interface-name*—(Optional) Display MAC learning statistics for the specified interface.

Required Privilege Level view

Related Documentation

- show ethernet-switching table on page 206
- show ethernet-switching interfaces on page 198
- show ethernet-switching mac-learning-log on page 254
- show ethernet-switching table on page 206
- show ethernet-switching interfaces on page 198
- *Example: Setting Up Basic Bridging and a VLAN on Switches*

List of Sample Output

show ethernet-switching statistics mac-learning on page 260
show ethernet-switching statistics mac-learning detail on page 261
show ethernet-switching statistics mac-learning interface ge-0/0/28 detail on page 261
show ethernet-switching statistics mac-learning interface on page 261
show ethernet-switching statistics mac-learning detail (QFX Series) on page 261

Output Fields Table 15 on page 260 lists the output fields for the **show ethernet-switching statistics mac-learning** command. Output fields are listed in the approximate order in which they appear.

Copyright © 2018, Juniper Networks, Inc.

Table 15: show ethernet-switching statistics mac-learning Output Fields

Field Name	Field Description	Level of Output
Interface	Name of the interface for which statistics are being reported. (Displayed in the output under the heading **Interface**.)	All levels
Learning message from local packets	MAC learning message generated due to packets coming in on the management interface. (Displayed in the output under the heading **Local pkts**.)	All levels
Learning message from transit packets	MAC learning message generated due to packets coming in on network interfaces. (Displayed in the output under the heading **Transit pkts**.)	All levels
Learning message with error	MAC learning messages received with errors (Displayed under the heading **Error**): • **Invalid VLAN**—The VLAN of the packet does not exist. • **Invalid MAC**—The MAC address is either NULL or a multicast MAC address. • **Security violation**—The MAC address is not an allowed MAC address. • **Interface down**—The MAC address is learned on an interface that is down. • **Incorrect membership**—The MAC address is learned on an interface that is not a member of the VLAN. • **Interface limit**—The number of MAC addresses learned on the interface has exceeded the limit. • **MAC move limit**—This MAC address has moved among multiple interfaces too many times in a given interval. • **VLAN limit**—The number of MAC addresses learned on the VLAN has exceeded the limit. • **VLAN membership limit**—The number of MAC addresses learned on the interface as a member of the specified VLAN (VLAN membership MAC limit) has exceeded the limit. • **Invalid VLAN index**—The VLAN of the packet, although configured, does not yet exist in the kernel. • **Interface not learning**—The MAC address is learned on an interface that does not yet allow learning—for example, the interface is blocked. • **No nexthop**—The MAC address is learned on an interface that does not have a unicast next hop. • **MAC learning disabled**—The MAC address is learned on an interface on which MAC learning has been disabled. • **Others**—The message contains some other error.	All levels

Sample Output

show ethernet-switching statistics mac-learning

```
user@switch> show ethernet-switching statistics mac-learning

Learning stats: 0 learn msg rcvd, 0 error
  Interface         Local pkts         Transit pkts       Error
  ge-0/0/0.0        0                  0                  0
  ge-0/0/1.0        0                  0                  0
  ge-0/0/2.0        0                  0                  0
  ge-0/0/3.0        0                  0                  0
```


Copyright © 2018, Juniper Networks, Inc.

show ethernet-switching statistics mac-learning detail

user@switch> **show ethernet-switching statistics mac-learning detail**

```
Learning stats: 0 learn msg rcvd, 0 error

    Interface: ge-0/0/0.0
    Learning message from local packets:   0
    Learning message from transit packets: 1
    Learning message with error:           0
      Invalid VLAN:             0       Invalid MAC:                0
      Security violation:       0       Interface down:             0
      Incorrect membership:     0       Interface limit:            0
      MAC move limit:           0       VLAN limit:                 0
      Invalid VLAN index:       0       Interface not learning:     0
      No nexthop:               0       MAC learning disabled:      0
      Others:                   0

    Interface: ge-0/0/1.0
    Learning message from local packets:   0
    Learning message from transit packets: 2
    Learning message with error:           0
      Invalid VLAN:             0       Invalid MAC:                0
      Security violation:       0       Interface down:             0
      Incorrect membership:     0       Interface limit:            0
      MAC move limit:           0       VLAN limit:                 0
      Invalid VLAN index:       0       Interface not learning:     0
      No nexthop:               0       MAC learning disabled:      0
      Others:                   0
```

show ethernet-switching statistics mac-learning interface ge-0/0/28 detail

user@switch> **show ethernet-switching statistics mac-learning interface ge-0/0/28 detail**

```

    Interface: ge-0/0/28.0
    Learning message from local packets:   0
    Learning message from transit packets: 5
    Learning message with error:           0
      Invalid VLAN:             0       Invalid MAC:                0
      Security violation:       0       Interface down:             0
      Incorrect membership:     0       Interface limit:            0
      MAC move limit:           0       VLAN limit:                 0
                                        VLAN membership limit:      20
      Invalid VLAN index:       0       Interface not learning:     0
      No nexthop:               0       MAC learning disabled:      0
      Others:                   0
```

show ethernet-switching statistics mac-learning interface

user@switch> **show ethernet-switching statistics mac-learning interface ge-0/0/1**

```
Interface           Local pkts          Transit pkts        Error
ge-0/0/1.0          0                   1                   1
```

show ethernet-switching statistics mac-learning detail (QFX Series)

user@switch> **show ethernet-switching statistics mac-learning detail**

Copyright © 2018, Juniper Networks, Inc.

```
Learning stats: 0 learn msg rcvd, 0 error

    Interface: xe-0/0/0.0
    Learning message from local packets:   0
    Learning message from transit packets: 1
    Learning message with error:           0
      Invalid VLAN:             0       Invalid MAC:                 0
      Security violation:       0       Interface down:              0
      Incorrect membership:     0       Interface limit:             0
      MAC move limit:           0       VLAN limit:                  0
      Invalid VLAN index:       0       Interface not learning:      0
      No nexthop:               0       MAC learning disabled:       0
      Others:                   0

    Interface: xe-0/0/1.0
    Learning message from local packets:   0
    Learning message from transit packets: 2
    Learning message with error:           0
      Invalid VLAN:             0       Invalid MAC:                 0
      Security violation:       0       Interface down:              0
      Incorrect membership:     0       Interface limit:             0
      MAC move limit:           0       VLAN limit:                  0
      Invalid VLAN index:       0       Interface not learning:      0
      No nexthop:               0       MAC learning disabled:       0
      Others:                   0
```


Copyright © 2018, Juniper Networks, Inc.

CHAPTER 9

Spanning Tree Monitoring Commands

clear error bpdu interface

List of Syntax

MX Series on page 264
QFX Series, EX Series, NFX Series on page 264

MX Series

```
clear error bpdu interface interface-name
```

QFX Series, EX Series, NFX Series

```
clear error bpdu interface (all | interface-name)
```

Release Information

Command introduced in Junos OS Release 9.4.
Command introduced in Junos OS Release 13.2X50-D10 for EX Series switches.
Command supports **all** option in Junos OS Release 15.1 for EX Series switches.

Description

Clear a bridge protocol data unit (BPDU) error condition caused by the detection of a possible bridging loop from Spanning Tree Protocol (STP) operation.

Required Privilege Level

clear

Related Documentation

- *Configuring BPDU Protection on ACX Router, EX Switch and MX Router Edge Ports*
- *Unblocking a Switch Interface That Receives BPDUs in Error (CLI Procedure)*

List of Sample Output

clear error bpdu interface ge-1/1/1 on page 264

Output Fields

When you enter this command, you are provided feedback on the status of your request.

Sample Output

clear error bpdu interface ge-1/1/1

```
user@host> clear error bpdu interface ge-1/1/1
```

 Copyright © 2018, Juniper Networks, Inc.

clear spanning-tree statistics

List of Syntax

Syntax on page 265
Syntax (EX Series Switches and the QFX Series) on page 265

Syntax

```
clear spanning-tree statistics
<interface interface-name>
<logical-system logical-system-name>
```

Syntax (EX Series Switches and the QFX Series)

```
clear spanning-tree statistics
<interface interface-name>
```

Release Information

Command introduced in Junos OS Release 8.4.
Command introduced in Junos OS Release 9.0 for EX Series switches.
Command introduced in Junos OS Release 11.1 for the QFX Series.

Description

Clear Spanning Tree Protocol statistics for all interfaces or a specified interface.

Options

none—Reset STP counters for all interfaces for all routing instances.

interface *interface-name*—(Optional) Clear STP statistics for the specified interface only.

logical-system *logical-system-name*—(Optional) Clear STP statistics on a particular logical system.

NOTE: The **logical-system** option is not available on QFabric systems.

Required Privilege Level

clear

Related Documentation

- *Understanding STP*
- **show spanning-tree statistics on page 285**
- **show spanning-tree bridge on page 267**
- **show spanning-tree interface on page 273**

List of Sample Output

clear spanning-tree statistics on page 266

Output Fields

This command produces no output.

Sample Output

clear spanning-tree statistics

```
user@switch> clear spanning—tree statistics
```

 Copyright © 2018, Juniper Networks, Inc.

show spanning-tree bridge

List of Syntax

Syntax on page 267
Syntax (QFX Series and EX Series) on page 267

Syntax

```
show spanning-tree bridge
<brief | detail>
<msti msti-id>
<routing-instance routing-instance-name>
<vlan-id vlan-id>
```

Syntax (QFX Series and EX Series)

```
show spanning-tree bridge
<brief | detail>
<msti msti-id>
<vlan-id vlan-id>
```

Release Information

Command introduced in Junos OS Release 8.4.
Command introduced in Junos OS Release 9.0 for EX Series switches.
Command introduced in Junos OS Release 11.1 for the QFX Series.

Description

Displays the configured or calculated spanning-tree protocol (can be either STP, RSTP, MSTP, or VSTP) parameters.

Options

none—(Optional) Display brief STP bridge information for all multiple spanning-tree instances (MSTIs).

brief | detail—(Optional) Display the specified level of output.

msti *msti-id*—(Optional) Display spanning-tree protocol bridge information for the specified MSTI or Common and Internal Spanning Tree (CIST). Specify **0** for CIST. Specify a value from **1** through **4094** for an MSTI.

routing-instance *routing-instance-name*—(Optional) Display STP bridge information for the specified routing instance.

vlan-id *vlan-id*—(Optional) Display spanning-tree protocol bridge information for the specified VLAN. Specify a VLAN tag identifier from **1** through **4094**.

Required Privilege Level

view

Related Documentation

- *Understanding STP*
- *Understanding RSTP*
- *Understanding VSTP*

- *Example: Configuring Network Regions for VLANs with MSTP on non-ELS EX Series Switches*
- show spanning-tree interface on page 273

List of Sample Output

Output Fields

Table 16 on page 268 lists the output fields for the **show spanning-tree bridge** command. Output fields are listed in the approximate order in which they appear.

Table 16: show spanning-tree bridge Output Fields

Field Name	Field Description
Routing instance name	Name of the routing instance under which the bridge is configured.
Enabled protocol	Spanning Tree Protocol type enabled.
Root ID	Bridge ID of the elected spanning-tree root bridge. The bridge ID consists of a configurable bridge priority and the MAC address of the bridge.
Root cost	Calculated cost to reach the root bridge from the bridge where the command is entered.
Root port	Interface that is the current elected root port for this bridge.
CIST regional root	Bridge ID of the elected MSTP regional root bridge.
CIST internal root cost	Calculated cost to reach the regional root bridge from the bridge where the command is entered.
Hello time	Configured number of seconds between transmissions of configuration BPDUs.
Maximum age	Configured maximum expected arrival time of hello bridge protocol data units (BPDUs).
Forward delay	How long an STP bridge port remains in the listening and learning states before transitioning to the forwarding state.
Hop count	Configured maximum number of hops a BPDU can be forwarded in the MSTP region.
Message age	Number of elapsed seconds since the most recent BPDU was received.

 Copyright © 2018, Juniper Networks, Inc.

Table 16: show spanning-tree bridge Output Fields (continued)

Field Name	Field Description
Number of topology changes	Total number of STP topology changes detected since the routing device last booted.
Time since last topology change	Number of elapsed seconds since the most recent topology change.
Bridge ID (Local)	Locally configured bridge ID. The bridge ID consists of a configurable bridge priority and the MAC address of the bridge.
Extended system ID	System identifier.
MSTI regional root	Bridge ID of the elected MSTP regional root bridge.

Sample Output

show spanning-tree bridge routing-instance

user@host> **show spanning-tree bridge routing-instance vs1 detail**

```
STP bridge parameters
Routing instance name                  : vs1
Enabled protocol                       : MSTP

STP bridge parameters for CIST
  Root ID                              : 32768.00:13:c3:9e:c8:80
  Root cost                            : 0
  Root port                            : ge-10/2/0
  CIST regional root                   : 32768.00:13:c3:9e:c8:80
  CIST internal root cost              : 22000
  Hello time                           : 2 seconds
  Maximum age                          : 20 seconds
  Forward delay                        : 15 seconds
  Hop count                            : 18
  Message age                          : 0
  Number of topology changes           : 1
  Time since last topology change      : 1191 seconds
  Local parameters
    Bridge ID                          : 32768.00:90:69:0b:7f:d1
    Extended system ID                 : 1

STP bridge parameters for MSTI 1
  MSTI regional root                   : 32769.00:13:c3:9e:c8:80
  Root cost                            : 22000
  Root port                            : ge-10/2/0
  Hello time                           : 2 seconds
  Maximum age                          : 20 seconds
  Forward delay                        : 15 seconds
  Hop count                            : 18
  Number of topology changes           : 1
  Time since last topology change      : 1191 seconds
  Local parameters
    Bridge ID                          : 32769.00:90:69:0b:7f:d1
    Extended system ID                 : 1
```

Copyright © 2018, Juniper Networks, Inc.

```
STP bridge parameters for MSTI 2
  MSTI regional root                 : 32770.00:13:c3:9e:c8:80
  Root cost                          : 22000
  Root port                          : ge-10/2/0
  Hello time                         : 2 seconds
  Maximum age                        : 20 seconds
  Forward delay                      : 15 seconds
  Hop count                          : 18
  Number of topology changes         : 1
  Time since last topology change    : 1191 seconds
  Local parameters
    Bridge ID                        : 32770.00:90:69:0b:7f:d1
    Extended system ID               : 1
```

show spanning-tree bridge msti

user@host> **show spanning-tree bridge msti 1 routing-instance vs1 detail**

```
STP bridge parameters
Routing instance name                : vs1
Enabled protocol                     : MSTP

STP bridge parameters for MSTI 1
  MSTI regional root                 : 32769.00:13:c3:9e:c8:80
  Root cost                          : 22000
  Root port                          : xe-10/2/0
  Hello time                         : 2 seconds
  Maximum age                        : 20 seconds
  Forward delay                      : 15 seconds
  Hop count                          : 18
  Number of topology changes         : 1
  Time since last topology change    : 1191 seconds
  Local parameters
    Bridge ID                        : 32769.00:90:69:0b:7f:d1
    Extended system ID               : 1
```

show spanning-tree bridge vlan-id (MSTP)

user@host> **show spanning-tree bridge vlan-id 1 101 routing-instance vs1 detail**

```
STP bridge parameters
Routing instance name                : vs1
Enabled protocol                     : MSTP

STP bridge parameters for CIST
  Root ID                            : 32768.00:13:c3:9e:c8:80
  Root cost                          : 0
  Root port                          : xe-10/2/0
  CIST regional root                 : 32768.00:13:c3:9e:c8:80
  CIST internal root cost            : 22000
  Hello time                         : 2 seconds
  Maximum age                        : 20 seconds
  Forward delay                      : 15 seconds
  Hop count                          : 18
  Message age                        : 0
  Number of topology changes         : 0
  Local parameters
    Bridge ID                        : 32768.00:90:69:0b:7f:d1
```


Copyright © 2018, Juniper Networks, Inc.

```
    Extended system ID              : 1
    Hello time                      : 2 seconds
    Maximum age                     : 20 seconds
    Forward delay                   : 15 seconds
    Path cost method                : 32 bit
    Maximum hop count               : 20
```

show spanning-tree bridge (RSTP)

user@host> **show spanning-tree bridge**

```
STP bridge parameters
Routing instance name               : GLOBAL
Enabled protocol                    : RSTP
  Root ID                           : 28672.00:90:69:0b:3f:d0
  Hello time                        : 2 seconds
  Maximum age                       : 20 seconds
  Forward delay                     : 15 seconds
  Message age                       : 0
  Number of topology changes        : 58
  Time since last topology change   : 14127 seconds
  Local parameters
    Bridge ID                       : 28672.00:90:69:0b:3f:d0
    Extended system ID              : 0

 STP bridge parameters for bridge VLAN 10
   Root ID                         : 28672.00:90:69:0b:3f:d0
   Hello time                      : 2 seconds
   Maximum age                     : 20 seconds
   Forward delay                   : 15 seconds
   Message age                     : 0
   Number of topology changes      : 58
   Time since last topology change : 14127 seconds
   Local parameters
     Bridge ID                     : 28672.00:90:69:0b:3f:d0
     Extended system ID            : 0

 STP bridge parameters for bridge VLAN 20
   Root ID                         : 28672.00:90:69:0b:3f:d0
   Hello time                      : 2 seconds
   Maximum age                     : 20 seconds
   Forward delay                   : 15 seconds
   Message age                     : 0
   Number of topology changes      : 58
   Time since last topology change : 14127 seconds
   Local parameters
     Bridge ID                     : 28672.00:90:69:0b:3f:d0
     Extended system ID            : 0
```

show spanning-tree bridge vlan-id (RSTP)

user@host> **show spanning-tree bridge vlan-id 10**

```
STP bridge parameters
Routing instance name               : GLOBAL
Enabled protocol                    : RSTP

 STP bridge parameters for VLAN 10
   Root ID                         : 28672.00:90:69:0b:3f:d0
```

Copyright © 2018, Juniper Networks, Inc.

```
Hello time                      : 2 seconds
Maximum age                     : 20 seconds
Forward delay                   : 15 seconds
Message age                     : 0
Number of topology changes      : 58
Time since last topology change : 14127 seconds
Local parameters
  Bridge ID                     : 28672.00:90:69:0b:3f:d0
  Extended system ID            : 0
```

 Copyright © 2018, Juniper Networks, Inc.

show spanning-tree interface

List of Syntax

Syntax on page 273
Syntax (EX Series Switches and QFX Series Switches) on page 273
Syntax (EX Series Switches) on page 273

Syntax

```
show spanning-tree interface
<brief | detail>
<msti msti-id>
<routing-instance routing-instance-name>
<vlan-id vlan-id>
```

Syntax (EX Series Switches and QFX Series Switches)

```
show spanning-tree interface
<brief | detail>
<msti msti-id>
<vlan-id vlan-id>
```

Syntax (EX Series Switches)

```
show spanning-tree interface
<brief | detail>
<interface-name interface-name>
<msti msti-id>
<vlan-id vlan-id>
```

Release Information

Command introduced in Junos OS Release 8.4.
Command introduced in Junos OS Release 9.0 for EX Series switches.
Command introduced in Junos OS Release 11.1 for the QFX Series.

Description

Display the configured or calculated interface-level spanning-tree protocol (can be either STP, RSTP, or MSTP) parameters. In **brief** mode, will not display interfaces that are administratively disabled or do not have a physical link.

Options

none—Display brief STP interface information.

brief | detail—(Optional) Display the specified level of output.

interface-name *interface-name*—(Optional) Name of an interface.

msti *msti-id*—(Optional) Display STP interface information for the specified MST instance.

routing-instance *routing-instance-name*—(Optional) Display STP interface information for the specified routing instance.

vlan-id *vlan-id*—(Optional) Display STP interface information for the specified VLAN.

Required Privilege Level

view

Copyright © 2018, Juniper Networks, Inc.

Related Documentation

- show spanning-tree bridge on page 267
- *Example: Configuring Network Regions for VLANs with MSTP on non-ELS EX Series Switches*
- *Understanding STP*
- *Understanding RSTP*
- *Understanding MSTP*
- *Understanding VSTP*

List of Sample Output

Output Fields

Table 17 on page 274 lists the output fields for the **show spanning-tree interface** command. Output fields are listed in the approximate order in which they appear.

Table 17: show spanning-tree Interface Output Fields

Field Name	Field Description
Interface name	Interface configured to participate in the STP, RSTP, VSTP, or MSTP instance.
Port ID	Logical interface identifier configured to participate in the MSTP or VSTP instance.
Designated port ID	Port ID of the designated port for the LAN segment to which this interface is attached.
Designated bridge ID	Bridge ID of the designated bridge for the LAN segment to which this interface is attached.
Port Cost	Configured cost for the interface.
Port State	STP port state: forwarding (**FWD**), blocking (**BLK**), listening, learning, or disabled.
Port Role	MSTP, VSTP, or RSTP port role: designated (**DESG**), backup (**BKUP**), alternate (**ALT**), (**ROOT**), or Root Prevented (**Root-Prev**).

Copyright © 2018, Juniper Networks, Inc.

Table 17: show spanning-tree Interface Output Fields (continued)

Field Name	Field Description
Link type	MSTP, VSTP, or RSTP link type. Shared or point-to-point (pt-pt) and edge or nonedge.
Alternate	Identifies the interface as an MSTP, VSTP, or RSTP alternate root port (**Yes**) or nonalternate root port (**No**).
Boundary Port	Identifies the interface as an MSTP regional boundary port (**Yes**) or nonboundary port (**No**).
Edge delay while expiry count	Number of times the edge delay timer expired on that interface.
Rcvd info while expiry count	Number of times the rcvd info timer expired on that interface.

Sample Output

show spanning-tree interface

user@host> **show spanning-tree interface routing-instance vs1 detail**

```
Spanning tree interface parameters for instance 0

Interface     Port ID    Designated      Designated         Port    State   Role
                          port ID        bridge ID          Cost
ae1              128:1        128:1  32768.0090690b47d1      1000   FWD     DESG
ge-2/1/2         128:2        128:2  32768.0090690b47d1     20000   FWD     DESG
ge-2/1/5         128:3        128:3  32768.0090690b47d1     29999   FWD     DESG
ge-2/2/1         128:4       128:26  32768.0013c39ec880     20000   FWD     ROOT
xe-9/2/0         128:5        128:5  32768.0090690b47d1      2000   FWD     DESG
xe-9/3/0         128:6        128:6  32768.0090690b47d1      2000   FWD     DESG

Spanning tree interface parameters for instance 1

Interface     Port ID    Designated      Designated         Port    State   Role
                          port ID        bridge ID          Cost
ae1              128:1        128:1  32769.0090690b47d1      1000   FWD     DESG
ge-2/1/2         128:2        128:2  32769.0090690b47d1     20000   FWD     DESG
ge-2/1/5         128:3        128:3  32769.0090690b47d1     29999   FWD     DESG
ge-2/2/1         128:4       128:26  32769.0013c39ec880     20000   FWD     ROOT
xe-9/2/0         128:5        128:5  32769.0090690b47d1      2000   FWD     DESG
xe-9/3/0         128:6        128:6  32769.0090690b47d1      2000   FWD     DESG

Spanning tree interface parameters for instance 2

Interface     Port ID    Designated      Designated         Port    State   Role
                          port ID        bridge ID          Cost
ae1              128:1        128:1  32770.0090690b47d1      1000   FWD     DESG
ge-2/1/2         128:2        128:2  32770.0090690b47d1     20000   FWD     DESG
ge-2/1/5         128:3        128:3  32770.0090690b47d1     29999   FWD     DESG
ge-2/2/1         128:4       128:26  32770.0013c39ec880     20000   FWD     ROOT
xe-9/2/0         128:5        128:5  32770.0090690b47d1      2000   FWD     DESG
xe-9/3/0         128:6        128:6  32770.0090690b47d1      2000   FWD     DESG
```

Copyright © 2018, Juniper Networks, Inc.

show spanning-tree interface (QFX Series)

user@host> **show spanning-tree interface routing-instance vs1 detail**

```
Spanning tree interface parameters for instance 0

Interface    Port ID    Designated      Designated         Port    State  Role
                         port ID        bridge ID          Cost
ae1              128:1        128:1  32768.0090690b47d1      1000   FWD    DESG
ge-2/1/2         128:2        128:2  32768.0090690b47d1     20000   FWD    DESG
ge-2/1/5         128:3        128:3  32768.0090690b47d1     29999   FWD    DESG
ge-2/2/1         128:4       128:26  32768.0013c39ec880     20000   FWD    ROOT
xe-9/2/0         128:5        128:5  32768.0090690b47d1      2000   FWD    DESG
xe-9/3/0         128:6        128:6  32768.0090690b47d1      2000   FWD    DESG

Spanning tree interface parameters for instance 1

Interface    Port ID    Designated      Designated         Port    State  Role
                         port ID        bridge ID          Cost
ae1              128:1        128:1  32769.0090690b47d1      1000   FWD    DESG
ge-2/1/2         128:2        128:2  32769.0090690b47d1     20000   FWD    DESG
ge-2/1/5         128:3        128:3  32769.0090690b47d1     29999   FWD    DESG
ge-2/2/1         128:4       128:26  32769.0013c39ec880     20000   FWD    ROOT
xe-9/2/0     128:5        128:5  32769.0090690b47d1      2000   FWD    DESG
xe-9/3/0     128:6        128:6  32769.0090690b47d1      2000   FWD    DESG

Spanning tree interface parameters for instance 2

Interface    Port ID    Designated      Designated         Port    State  Role
                         port ID        bridge ID          Cost
ae1              128:1        128:1  32770.0090690b47d1      1000   FWD    DESG
ge-2/1/2         128:2        128:2  32770.0090690b47d1     20000   FWD    DESG
ge-2/1/5         128:3        128:3  32770.0090690b47d1     29999   FWD    DESG
ge-2/2/1         128:4       128:26  32770.0013c39ec880     20000   FWD    ROOT
xe-9/2/0     128:5        128:5  32770.0090690b47d1      2000   FWD    DESG
xe-9/3/0     128:6        128:6  32770.0090690b47d1      2000   FWD    DESG
```

show spanning-tree interface (EX Series)

user@switch> **show spanning-tree interface**

```

Spanning tree interface parameters for instance 0

Interface    Port ID    Designated  Designated      Port    State  Role
                         port ID    bridge ID       Cost
ge-0/0/0.0   128:513  128:513    8192.0019e2500340    1000    FWD    DESG
ge-0/0/2.0   128:515  128:515    8192.0019e2500340    1000    BLK    DIS
ge-0/0/4.0   128:517  128:517    8192.0019e2500340    1000    FWD    DESG
ge-0/0/23.0 128:536  128:536    8192.0019e2500340    1000    FWD    DESG


Spanning tree interface parameters for instance 1

Interface    Port ID    Designated  Designated      Port    State  Role
                         port ID    bridge ID       Cost
ge-0/0/0.0   128:513  128:513    8193.0019e2500340    1000    FWD    DESG
ge-0/0/2.0   128:515  128:515    8193.0019e2500340    1000    BLK    DIS
ge-0/0/4.0   128:517  128:517    8193.0019e2500340    1000    FWD    DESG
ge-0/0/23.0 128:536  128:536    8193.0019e2500340    1000    FWD    DESG
```

 Copyright © 2018, Juniper Networks, Inc.

```
Spanning tree interface parameters for instance 2

Interface    Port ID    Designated  Designated     Port    State  Role
                         port ID    bridge ID      Cost
ge-0/0/0.0  128:513  128:1   8194.001b549fd000  1000  FWD    ROOT
ge-0/0/2.0  128:515  128:515  32770.0019e2500340  4000  BLK    DIS
ge-0/0/4.0  128:517  128:1    16386.001b54013080  1000  BLK    ALT
ge-0/0/23.0 128:536  128:536  32770.0019e2500340  1000  FWD    DESG
```

show spanning-tree interface detail

user@host> **show spanning-tree interface routing-instance vs1 detail**

```
Spanning tree interface parameters for instance 0

Interface name              : ae1
Port identifier             : 128.1
Designated port ID          : 128.1
Port cost                   : 1000
Port state                  : Forwarding
Designated bridge ID        : 32768.00:90:69:0b:47:d1
Port role                   : Designated
Link type                   : Pt-Pt/NONEDGE
Boundary port                : No

Interface name              : ge-2/1/2
Port identifier             : 128.2
Designated port ID          : 128.2
Port cost                   : 20000
Port state                  : Forwarding
Designated bridge ID        : 32768.00:90:69:0b:47:d1
Port role                   : Designated
Link type                   : Pt-Pt/NONEDGE
Boundary port                : No

Interface name              : ge-2/1/5
Port identifier             : 128.3
Designated port ID          : 128.3
Port cost                   : 29999
Port state                  : Forwarding
Designated bridge ID        : 32768.00:90:69:0b:47:d1
Port role                   : Designated
Link type                   : Pt-Pt/NONEDGE
Boundary port                : No

Interface name              : ge-2/2/1
Port identifier             : 128.4
Designated port ID          : 128.26
Port cost                   : 20000
Port state                  : Forwarding
Designated bridge ID        : 32768.00:13:c3:9e:c8:80
Port role                   : Root
Link type                   : Pt-Pt/NONEDGE
Boundary port                : No

Interface name              : xe-9/2/0
Port identifier             : 128.5
Designated port ID          : 128.5
```

Copyright © 2018, Juniper Networks, Inc.

```
Port cost                       : 2000
Port state                      : Forwarding
Designated bridge ID            : 32768.00:90:69:0b:47:d1
Port role                       : Designated
Link type                       : Pt-Pt/NONEDGE
Boundary port                    : No

Interface name                  : xe-9/3/0
Port identifier                 : 128.6
Designated port ID              : 128.6
Port cost                       : 2000
Port state                      : Forwarding
Designated bridge ID            : 32768.00:90:69:0b:47:d1
Port role                       : Designated
Link type                       : Pt-Pt/NONEDGE
Boundary port                    : No


Spanning tree interface parameters for instance 1

Interface name                  : ae1
Port identifier                 : 128.1
Designated port ID              : 128.1
Port cost                       : 1000
Port state                      : Forwarding
Designated bridge ID            : 32768.00:90:69:0b:47:d1
Port role                       : Designated
Link type                       : Pt-Pt/NONEDGE
Boundary port                    : No

Interface name                  : ge-2/1/2
Port identifier                 : 128.2
Designated port ID              : 128.2
Port cost                       : 20000
Port state                      : Forwarding
Designated bridge ID            : 32768.00:90:69:0b:47:d1
Port role                       : Designated
Link type                       : Pt-Pt/NONEDGE
Boundary port                    : No

Interface name                  : ge-2/1/5
Port identifier                 : 128.3
Designated port ID              : 128.3
Port cost                       : 29999
Port state                      : Forwarding
Designated bridge ID            : 32768.00:90:69:0b:47:d1
Port role                       : Designated
Link type                       : Pt-Pt/NONEDGE
Boundary port                    : No

Interface name                  : ge-2/2/1
Port identifier                 : 128.4
Designated port ID              : 128.26
Port cost                       : 20000
Port state                      : Forwarding
Designated bridge ID            : 32768.00:13:c3:9e:c8:80
Port role                       : Root
Link type                       : Pt-Pt/NONEDGE
Boundary port                    : No
```

 Copyright © 2018, Juniper Networks, Inc.

```
...
```

show spanning-tree interface detail (EX Series)

user@switch> **show spanning-tree interface detail**

```
Spanning tree interface parameters for instance 0

Interface name     : ge-1/0/0.0
Port identifier    : 128.625
Designated port ID     : 128.625
Port cost      : 20000
Port state     : Blocking
Designated bridge ID   : 32768.00:19:e2:50:95:a0
Port role      : Disabled
Link type      : Pt-Pt/EDGE
Boundary port       : NA
Edge delay while expiry count  : 0
Rcvd info while expiry count   : 0

Interface name     : ge-1/0/1.0
Port identifier    : 128.626
Designated port ID     : 128.626
Port cost      : 20000
Port state     : Blocking
Designated bridge ID   : 32768.00:19:e2:50:95:a0
Port role      : Disabled
Link type      : Pt-Pt/NONEDGE
Boundary port       : NA
Edge delay while expiry count  : 0
Rvcd info while expiry count   : 0

Interface name     : ge-1/0/2.0
Port identifier    : 128.627
Designated port ID     : 128.627
Port cost      : 20000
Port state     : Blocking
Designated bridge ID   : 32768.00:19:e2:50:95:a0
Port role      : Disabled
Link type      : Pt-Pt/NONEDGE
Boundary port       : NA
Edge delay while expiry count  : 0
Rvcd info while expiry count   : 0

Interface name     : ge-1/0/10.0
Port identifier    : 128.635
Designated port ID     : 128.635
Port cost      : 20000
Port state     : Blocking
Designated bridge ID   : 32768.00:19:e2:50:95:a0
Port role      : Disabled
Link type      : Pt-Pt/NONEDGE
Boundary port       : NA
Edge delay while expiry count  : 0
Rvcd info while expiry count   : 0

Interface name     : ge-1/0/20.0
Port identifier    : 128.645
Designated port ID     : 128.645
```

Copyright © 2018, Juniper Networks, Inc.

```
Port cost      : 20000
Port state     : Blocking
Designated bridge ID   : 32768.00:19:e2:50:95:a0
Port role      : Disabled
Link type      : Pt-Pt/NONEDGE
Boundary port       : NA
Edge delay while expiry count  : 0
Rvcd info while expiry count   : 0
[output truncated]
```

show spanning-tree interface msti

user@host> **show spanning-tree interface msti 1 routing-instance vs1 detail**

```
 Spanning tree interface parameters for instance 1

Interface     Port ID    Designated      Designated         Port    State  Role
                          port ID        bridge ID          Cost
xe-7/0/0        128:1         128:1  32769.0090690b4fd1      2000   FWD    DESG
ge-5/1/0        128:2         128:2  32769.0090690b4fd1     20000   FWD    DESG
ge-5/1/1        128:3         128:3  32769.0090690b4fd1     20000   FWD    DESG
ae1             128:4         128:1  32769.0090690b47d1     10000   BLK    ALT
ge-5/1/4        128:5         128:3  32769.0090690b47d1     20000   BLK    ALT
xe-7/2/0        128:6         128:6  32769.0090690b47d1      2000   FWD    ROOT
```

show spanning-tree interface vlan-id

user@host> **show spanning-tree interface vlan-id 101 routing-instance vs1 detail**

```
Spanning tree interface parameters for instance 0

Interface     Port ID    Designated      Designated         Port    State  Role
                          port ID        bridge ID          Cost
ge-11/0/5       128:1         128:1  32768.0090690b7fd1     20000   FWD    DESG
ge-11/0/6       128:2         128:1  32768.0090690b7fd1     20000   BLK    BKUP
ge-11/1/0       128:3         128:2  32768.0090690b4fd1     20000   BLK    ALT
ge-11/1/1       128:4         128:3  32768.0090690b4fd1     20000   BLK    ALT
ge-11/1/4       128:5         128:1  32768.0090690b47d1     20000   BLK    ALT
xe-10/0/0       128:6         128:5  32768.0090690b4fd1      2000   BLK    ALT
xe-10/2/0       128:7         128:4  32768.0090690b47d1      2000   FWD    ROOT
```

show spanning-tree interface (VSTP)

user@host> **show spanning-tree interface**

```
Spanning tree interface parameters for instance 0

Interface     Port ID    Designated      Designated         Cost    State  Role
                          port ID        bridge ID
ge-1/0/1        128:1         128:1  28672.0090690b3fe0     20000   FWD    DESG
ge-1/0/2        128:2         128:2  28672.0090690b3fe0     20000   FWD    DESG

Spanning tree interface parameters for VLAN 10

Interface     Port ID    Designated      Designated         Cost    State  Role
                          port ID        bridge ID
ge-1/0/1        128:1         128:1  28672.0090690b3fe0     20000   FWD    DESG
ge-1/0/2        128:2         128:2  28672.0090690b3fe0     20000   FWD    DESG
```


Copyright © 2018, Juniper Networks, Inc.

```
Spanning tree interface parameters for VLAN 20

Interface     Port ID    Designated      Designated         Cost     State  Role
                          port ID        bridge ID
ge-1/0/1         128:1        128:1  28672.0090690b3fe0     20000  FWD    DESG
ge-1/0/2         128:2        128:2  28672.0090690b3fe0     20000  FWD    DESG
```

show spanning-tree interface vlan-id (VSTP)

`user@host>` **show spanning-tree interface vlan-id 10**

```
Spanning tree interface parameters for VLAN 10

Interface     Port ID    Designated      Designated         Cost     State  Role
                          port ID        bridge ID
ge-1/0/1         128:1        128:1  28672.0090690b3fe0     20000  FWD    DESG
ge-1/0/2         128:2        128:2  28672.0090690b3fe0     20000  FWD    DESG
```

show spanning-tree interface brief (EX Series)

`user@switch>` **show spanning-tree interface brief**

```
Spanning tree interface parameters for instance 0

Interface     Port ID    Designated  Designated     Port    State  Role
     port ID    bridge ID  Cost
ge-1/0/0.0  128:625  128:625  32768.0019e25095a0 20000  BLK    DIS
ge-1/0/1.0  128:626  128:626  32768.0019e25095a0 20000  BLK    DIS
ge-1/0/2.0  128:627  128:627  32768.0019e25095a0 20000  BLK    DIS
ge-1/0/10.0 128:635  128:635  32768.0019e25095a0 20000  BLK    DIS
ge-1/0/20.0 128:645  128:645  32768.0019e25095a0 20000  BLK    DIS
ge-1/0/30.0 128:655  128:655  32768.0019e25095a0 20000  BLK    DIS
```

show spanning-tree interface ge-1/0/0 (EX Series)

`user@switch>` **show spanning-tree interface ge-1/0/0**

```

Interface     Port ID    Designated  Designated     Port    State  Role
     port ID    bridge ID  Cost
ge-1/0/0.0 128:625  128:625  32768.0019e25095a0 20000  BLK    DIS
```

Copyright © 2018, Juniper Networks, Inc.

show spanning-tree mstp configuration

List of Syntax

Syntax on page 282
Syntax (EX Series Switch and the QFX Series) on page 282

Syntax

```
show spanning-tree mstp configuration
<brief | detail>
<routing-instance routing-instance-name>
```

Syntax (EX Series Switch and the QFX Series)

```
show spanning-tree mstp configuration
<brief | detail>
```

Release Information

Command introduced in Junos OS Release 8.4.
Command introduced in Junos OS Release 9.0 for EX Series switches.
Command introduced in Junos OS Release 11.1 for the QFX Series.

Description

Display the MSTP configuration.

Options

none—Display MSTP configuration information.

brief | detail—(Optional) Display the specified level of output.

routing-instance ***routing-instance-name***—(Optional) Display MSTP configuration information for the specified routing instance.

Required Privilege Level

view

Related Documentation

- *Understanding MSTP*
- *Example: Configuring Network Regions for VLANs with MSTP on non-ELS EX Series Switches*
- show spanning-tree bridge on page 267
- show spanning-tree statistics on page 285
- *mstp*

List of Sample Output

show spanning-tree mstp configuration detail on page 283
show spanning-tree mstp configuration detail (QFX Series) on page 283
show spanning-tree mstp configuration (EX Series) on page 283

Output Fields

Table 18 on page 283 lists the output fields for the **show spanning-tree mstp configuration** command. Output fields are listed in the approximate order in which they appear.

 Copyright © 2018, Juniper Networks, Inc.

Table 18: show spanning-tree mstp configuration Output Fields

Field Name	Field Description
Context id	Internally generated identifier.
Region name	MSTP region name carried in the MSTP BPDUs.
Revision	Revision number of the MSTP configuration.
Configuration digest	Numerical value derived from the VLAN-to-instance mapping table.
MSTI	MST instance identifier.
Member VLANs	VLAN identifiers associated with the MSTI.

Sample Output

show spanning-tree mstp configuration detail

`user@host>` **show spanning-tree mstp configuration routing-instance vs1 detail**

```
MSTP configuration information
Context identifier     : 1
Region name            : henry
Revision               : 3
Configuration digest   : 0x6da4b5c4fd587757eef35675365e1


MSTI     Member VLANs
   0 0-99,101-199,201-4094
   1 100
   2 200
```

show spanning-tree mstp configuration detail (QFX Series)

`user@lf0>` **show spanning-tree mstp configuration routing-instance vs1 detail**

```
MSTP configuration information
Context identifier     : 1
Region name            : henry
Revision               : 3
Configuration digest   : 0x6da4b5c4fd587757eef35675365e1


MSTI     Member VLANs
   0 0-99,101-199,201-4094
   1 100
   2 200
```

show spanning-tree mstp configuration (EX Series)

`user@host>` **show spanning-tree mstp configuration**

```
MSTP configuration information
Context identifier     : 0
```

Copyright © 2018, Juniper Networks, Inc.

```
Region name            : region1
Revision               : 0
Configuration digest   : 0xc92e7af9febb44d8df928b87f16b


MSTI     Member VLANs
   0 0-100,105-4094
   1 101-102
   2 103-104
```

 Copyright © 2018, Juniper Networks, Inc.

show spanning-tree statistics

List of Syntax

Syntax on page 285
Syntax (EX Series and QFX Series) on page 285
Syntax (EX Series) on page 285

Syntax

```
show spanning-tree statistics
<brief | detail>
<interface interface-name>
<routing-instance routing-instance-name>
```

Syntax (EX Series and QFX Series)

```
show spanning-tree statistics
<brief | detail>
<interface interface-name | vlan vlan-id>
```

Syntax (EX Series)

```
show spanning-tree statistics
interface interface-name
vlan vlan-id
<brief | detail>
```

Release Information

Command introduced in Junos OS Release 8.4.
Command introduced in Junos OS Release 9.0 for EX Series switches.
Command introduced in Junos OS Release 11.1 for QFX Series switches.

Description

Display STP statistics.

Options

none—Display brief STP statistics.

brief | detail—(Optional) Display the specified level of output.

interface *interface-name*—(Optional) Display STP statistics for the specified interface.

routing-instance *routing-instance-name*—(Optional) Display STP statistics for the specified routing instance.

Required Privilege Level

view

List of Sample Output

show spanning-tree statistics routing-instance on page 286
show spanning-tree statistics interface routing-instance detail on page 286

Output Fields

Table 19 on page 286 lists the output fields for the **show spanning-tree statistics** command. Output fields are listed in the approximate order in which they appear.

Copyright © 2018, Juniper Networks, Inc.

Table 19: show spanning-tree statistics Output Fields

Field Name	Field Description
Message type	Type of message being counted.
BPDUs sent	Total number of BPDUs sent.
BPDUs received	Total number of BPDUs received.
BPDUs sent in last interval	Number of BPDUs sent within a specified interval.
BPDUs received in last interval	Number of BPDUs received within a specified interval.
Interface	Interface for which the statistics are being displayed.
Next BPDU transmission	Number of seconds until the next BPDU is scheduled to be sent.

Sample Output

show spanning-tree statistics routing-instance

`user@host>` **show spanning-tree statistics routing-instance vs1 detail**

```
Routing instance level STP statistics
Message type                    : bpdus
BPDUs sent                      : 1396
BPDUs received                  : 1027
BPDUs sent in last interval     : 5        (duration: 4 sec)
BPDUs received in last interval: 4        (duration: 4 sec)
```

show spanning-tree statistics interface routing-instance detail

`user@host>` **show spanning-tree statistics interface ge-11/1/4 routing-instance vs1 detail**

```
Interface    BPDUs sent   BPDUs received      Next BPDU
                                             transmission
ge-11/1/4            7           190                0
```

 Copyright © 2018, Juniper Networks, Inc.