JUNIPER NETWORKS | Engineering Simplicity

Junos® OS

Layer 2 Bridging, Address Learning, and Forwarding User Guide

Published
2021-04-17

junos

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

Junos® OS Layer 2 Bridging, Address Learning, and Forwarding User Guide

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

4   Configuring Layer 2 Address Learning and Forwarding

5   Configuring Layer 2 Learning and Forwarding for Bridge Domains

9 **Configuration Statements for Layer 2 Bridge Domains Functioning as Switches with Layer 2 Trunk Ports**

10 **Configuration Statements for Layer 2 Address Learning and Forwarding**

About This Guide

You can configure one or more bridge domains to perform Layer 2 bridging. A bridge domain is a set of logical interfaces that share the same flooding or broadcast characteristics. Layer 2 logical interfaces are created by defining one or more logical units on a physical interface with encapsulation as **ethernet-bridge** or **vlan-bridge**. All the member ports of the bridge domain participate in Layer 2 learning and forwarding. Like a virtual LAN (VLAN), a bridge domain spans one or more ports of multiple devices. You can configure Layer 2 MAC address and VLAN learning and forwarding properties in support of Layer 2 bridging.

Use this guide to configure, monitor, and troubleshoot Layer 2 bridging, address learning, and forwarding features on your Juniper Network devices.

1
CHAPTER

Understanding Layer 2 Bridging, Address Learning, and Forwarding

Understanding Layer 2 Bridge Domains

You can configure one or more bridge domains on MX Series routers to perform Layer 2 bridging. The Layer 2 bridging functions of the MX Series routers include integrated routing and bridging (IRB) for support for Layer 2 bridging and Layer 3 IP routing on the same interface, and virtual switches that isolate a LAN segment with its spanning-tree protocol instance and separate its VLAN ID space.

A bridge domain is a set of logical ports that share the same flooding or broadcast characteristics. Like a virtual LAN (VLAN), a bridge domain spans one or more ports of multiple devices.

On Juniper Networks MX Series 5G Universal Routing Platforms only, you can configure one or more bridge domains to perform Layer 2 bridging. Thus, MX Series routers can function as Layer 2 switches, each with multiple bridging, or broadcast, domains that participate in the same Layer 2 network. You can also configure Layer 3 routing support for a bridge domain. Integrated routing and bridging (IRB) provides support for Layer 2 bridging and Layer 3 IP routing on the same interface. IRB enables you to route packets to another routed interface or to another bridge domain that has a Layer 3 protocol configured.

You can also group one or more bridge domains within a single instance, or virtual switch. The MX Series routers also support multiple virtual switches, each of which operates independently of other virtual switches on the router. Virtual switches isolate a LAN segment with its spanning-tree protocol instance. . Thus, each virtual switch can participate in a different Layer 2 network.

In Junos OS Release 9.2 and later, bridge domains provide support for a Layer 2 trunk port. A Layer 2 trunk interface enables you to configure a single *logical interface* to represent multiple VLANs on a physical interface. You can configure a set of bridge domains and VLAN identifiers that are automatically associated with one or more Layer 2 trunk interfaces. Packets received on a trunk interface are forwarded within a bridge domain that has the same VLAN identifier. A Layer 2 trunk interface also supports IRB within a bridge domain. In addition, you can configure Layer 2 learning and forwarding properties that apply to the entire set of bridge domains.

In Junos OS Release 9.3 and later, you can configure VPLS ports in a virtual switch instead of a dedicated routing instance of type **vpls** so that the logical interfaces of the Layer 2 bridge domains in the virtual switch can handle VPLS routing instance traffic. Packets received on a Layer 2 trunk interface are forwarded within a bridge domain that has the same VLAN identifier.

RELATED DOCUMENTATION

Understanding Layer 2 Learning and Forwarding

On MX Series routers only, you can configure Layer 2 MAC address and VLAN learning and forwarding properties in support of Layer 2 bridging. The router learns unicast media access control (MAC) addresses to avoid flooding the packets to all the ports in a bridge domain. The MX Series router creates a source MAC entry in its source and destination MAC tables for each MAC address learned from packets received on ports that belong to the bridge domain. If the bridge domain receives a control protocol data unit (PDU) which does not have a corresponding protocol configured, then the control PDU is considered as an unknown multicast data packet and the packets are flooded across all the ports that are part of the same bridge domain. If the bridge domain has the protocol corresponding to the PDU configured , then the control PDU is considered as a control packet and is processed by the routing engine.

By default, Layer 2 address learning is enabled. You can disable MAC learning for the router or for a specific bridge domain or logical interfaces. You can also configure the following Layer 2 forwarding properties for an MX Series router:

- Timeout interval for MAC entries

- MAC accounting

- A limit to the number of MAC addresses learned from the logical interfaces

RELATED DOCUMENTATION

Layer 2 Bridge Domains on ACX Series Overview

A bridge domain is a set of logical interfaces that share the same flooding or broadcast characteristics. Layer 2 logical interfaces are created by defining one or more logical units on a physical interface with encapsulation as **ethernet-bridge** or **vlan-bridge**. All the member ports of the bridge domain participate in Layer 2 learning and forwarding. You can configure one or more bridge domains on ACX Series routers to perform Layer 2 bridging. The Layer 2 bridging functions of ACX Series routers include integrated routing and bridging (IRB) support for Layer 2 bridging and Layer 3 IP routing on the same interface. IRB enables you to route packets to another routed interface or to another bridge domain that has a Layer 3 protocol configured

> **NOTE**: ACX Series routers do not support the creation of bridge domains by using access and trunk ports.

You can configure E-LAN and E-LINE services by using bridge domains.

On ACX Series routers, you can configure bridge domains by using the following methods:

- Bridge domain without a **vlan-id** *number* statement

- Bridge domain with the **vlan-id** value set to *none*

- Bridge domain with a single **vlan-id**

- Bridge domain with a **vlan-id-list**

> **NOTE**: The Layer 2 CLI configurations and show commands for ACX5048 and ACX5096 routers differ compared to other ACX Series routers. For more information, see *Layer 2 Next Generation Mode for ACX Series*.

When you configure E-LAN and E-LINE services using a bridge domain without a **vlan-id** *number* statement, the bridge domain should explicitly be normalized to a service VLAN ID and TPID by configuring an input VLAN map under a logical interface. Explicit normalization is required when a logical interface's outer VLAN ID and TPID is not the same as the service VLAN ID and TPID of the service being configured using a bridge domain.

The following input VLAN map functions are supported in ACX Series routers:

- **push**—Add a new VLAN tag to the top of the VLAN stack.

- **swap**—Replace the outer VLAN tag of the VLAN stack in a frame.

- **pop**—Remove a VLAN tag from the top of the VLAN tag stack.

- **swap-swap**—Replace both the outer and inner VLAN tags of the frame.

- **push-push**—Push two VLAN tags on top of the VLAN stack.

  > **NOTE**: **push-push** does not work on ACX Series routers if the incoming packet already has a VLAN tag.

The following VLAN map functions are not supported in ACX Series routers:

- **swap-push**—Replace the outer VLAN tag of the frame and add a new VLAN tag to the top of the VLAN stack.

- **pop-swap**—Remove the outer VLAN tag of the frame and replace the inner VLAN tag of the frame.

- **pop-pop**—Remove both the outer and inner VLAN tags of the frame.

  > **NOTE**: You can configure Q-in-Q tunneling by explicitly configuring an input VLAN map with the **push** function on the ingress logical interface.

A bridge domain can also be created by using aggregated Ethernet interfaces. Aggregated Ethernet interfaces are considered as logical interfaces in a bridge domain.

The following steps outline the process for bridging a packet received over a Layer 2 logical interface:

1. When a packet is received on a physical port, it is accepted only if the VLAN identifier of the packet matches the VLAN identifier of one of the logical interfaces configured on that port.

2. If the bridge domain is configured without a **vlan-id** *number* statement, then the VLAN tags are rewritten based on the input VLAN map configured on the logical interface and normalized to a service VLAN ID.

3. If the bridge domain is configured with a normalizing VLAN identifier by using the **vlan-id** *number* statement, the VLAN tags of the received packet are compared with the normalizing VLAN identifier. If the VLAN tags of the packet are different from the normalizing VLAN identifier, the VLAN tags are rewritten as described in Table 1 on page 6.

4. If the source MAC address of the received packet is not present in the source MAC table, it is learned based on the normalizing VLAN identifier.

5. The packet is then forwarded toward one or more outbound Layer 2 logical interfaces based on the destination MAC address. A packet with a known unicast destination MAC address is forwarded only to one outbound logical interface.

6. If the bridge domain is configured without a **vlan-id** *number* statement, then for each outbound Layer 2 logical interface, the VLAN tags are rewritten based on the output VLAN map configured on that logical interface.

7. If the bridge domain is configured with a normalizing VLAN identifier by using the **vlan-id** *number* statement, for each outbound Layer 2 logical interface, the normalizing VLAN identifier configured for the bridge domain is compared with the VLAN tags configured on that logical interface. If the VLAN tags associated with an outbound logical interface do not match the normalizing VLAN identifier configured for the bridge domain, the VLAN tags are rewritten as described in Table 2 on page 7.

Table 1 on page 6 shows specific examples of how the VLAN tags of packets sent to the bridge domain are processed and translated, depending on your configuration. "–" means that the statement is not supported for the specified logical interface VLAN identifier. "No operation" means that the VLAN tags of the received packet are not translated for the specified input logical interface.

Table 1: Statement Usage and Input Rewrite Operations for VLAN Identifiers for a Bridge Domain

VLAN Identifier of Logical Interface	VLAN Configurations for Bridge Domain	
	vlan-id none	vlan-id 200
none	No operation	push 200
200	pop 200	No operation
1000	pop 1000	swap 1000 to 200
vlan-tags outer 2000 inner 300	pop 2000, pop 300	pop 2000, swap 300 to 200
vlan-tags outer 100 inner 400	pop 100, pop 400	pop 100, swap 400 to 200
vlan-id-range 10-100	–	–

Table 2 on page 7 shows specific examples of how the VLAN tags for packets sent from the bridge domain are processed and translated, depending on your configuration. "–" means that the statement is not supported for the specified logical interface VLAN identifier. "No operation" means that the VLAN tags of the outbound packet are not translated for the specified output logical interface.

Table 2: Statement Usage and Output Rewrite Operations for VLAN Identifiers for a Bridge Domain

VLAN Identifier of Logical Interface	VLAN Configurations for Bridge Domain	
	vlan-id none	vlan-id 200
none	no operation	pop 200
200	push 200	No operation
1000	push 1000	swap 200 to 1000
vlan-tags outer 2000 inner 300	push 2000, push 300	swap 200 to 300, push 2000
vlan-tags outer 100 inner 400	push 100, push 400	swap 200 to 400, push 100
vlan-id-range 10-100	–	–

Limitations on Layer 2 bridging—The following Layer 2 bridging limitations apply for ACX Series Universal Metro Routers:

- A bridge domain cannot have two or more logical interfaces that belong to the same physical interface.

- A bridge domain with dual VLAN ID tag is not supported.

- The maximum number of supported input VLAN maps with TPID **swap** is 64.

- MAC learning cannot be disabled at a logical interface level.

- MAC limit per logical interface cannot be configured.

Layer 2 Learning and Forwarding for Bridge Domains Overview

When you configure a bridge domain, Layer 2 address learning is enabled by default. The bridge domain learns unicast media access control (MAC) addresses to avoid flooding the packets to all the ports in the bridge domain. Each bridge domain 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 bridge domain.

NOTE: Traffic is not flooded back onto the interface on which it was received.

You can optionally disable MAC learning either for the entire router or for a specific bridge domain. You can also configure the following Layer 2 learning and forwarding properties:

- Static MAC entries on logical interfaces

- Size of the MAC address table for the bridge domain

2
CHAPTER

Configuring Layer 2 Bridging and Layer 3 IP Routing

Configuring a Bridge Domain

A bridge domain must include a set of logical interfaces that participate in Layer 2 learning and forwarding. You can optionally configure a VLAN identifier and a routing interface for the bridge domain to also support Layer 3 IP routing.

To enable a bridge domain, include the following statements:

```
[edit]
bridge-domains {
    bridge-domain-name {
        domain-type bridge:
        interface interface-name;
        routing-interface routing-interface-name;
        vlan-id (none | all | number);
        vlan-id-list [ vlan-id-numbers ];
        vlan-tags outer number inner number);
    }
}
```

You cannot use the slash (/) character in bridge domain names. If you do, the configuration does not commit and an error is generated.

For the **vlan-id** statement, you can specify either a valid VLAN identifier or the **none** or **all** options. For information about VLAN identifiers and VLAN tags for a bridge domain, see "Configuring VLAN Identifiers for Bridge Domains and VPLS Routing Instances" on page 28.

To include one or more logical interfaces in the bridge domain, specify an *interface-name* for an Ethernet interface you configured at the **[edit interfaces]** hierarchy level.

> **NOTE**: A maximum of 4000 active logical interfaces are supported on a bridge domain or on each mesh group in a virtual private LAN service (VPLS) instance configured for Layer 2 bridging.

By default, each bridge domain maintains a Layer 2 forwarding database that contains media access control (MAC) addresses learned from packets received on the ports that belong to the bridge domain. You can modify Layer 2 forwarding properties, including disabling MAC learning for the entire system or a bridge domain, adding static MAC addresses for specific logical interfaces, and limiting the number of MAC addresses learned by the entire system, the bridge domain, or a logical interface.

You can also configure spanning tree protocols to prevent forwarding loops. .

In Junos OS Release 8.5 and later, you can configure IGMP snooping for a bridge domain. For more information, see the Junos OS Multicast Protocols User Guide.

Integrated routing and bridging (IRB) provides simultaneous support for Layer 2 bridging and Layer 3 routing on the same interface. IRB enables you to route packets to another routed interface or to another bridge domain that has an IRB interface configured. You configure a logical routing interface by including the **irb** statement at the **[edit interfaces]** hierarchy level and include that interface in the bridge domain. For more information about how to configure a routing interface, see the Junos OS Network Interfaces Library for Routing Devices.

> **NOTE**: You can include only one routing interface in a bridge domain.

To configure a bridge domain with IRB support, include the following statements:

```
[edit]
bridge-domains {
    bridge-domain-name {
        domain-type bridge;
        interface interface-name;
        routing-interface routing-interface-name;
        service-id number;
        vlan-id (none | number);
        vlan-tags outer number inner number;
    }
}
```

For each bridge domain that you configure, specify a *bridge-domain-name*. You must also specify the value **bridge** for the **domain-type** statement.

For the **vlan-id** statement, you can specify either a valid VLAN identifier or the **none** option.

> **NOTE**: If you configure a routing interface to support IRB in a bridge domain, you cannot use the **all** option for the **vlan-id** statement.

The **vlan-tags** statement enables you to specify a pair of VLAN identifiers; an **outer** tag and an **inner** tag.

> **NOTE**: For a single bridge domain, you can include either the **vlan-id** statement or the **vlan-tags** statement, but not both.

For MC-LAG bridge domains, when the VLAN identifier is **none**, use the **service-id** statement to facilitate media access control (MAC) and Address Resolution Protocol (ARP) synchronization among MC-LAG peers.

To include one or more logical interfaces in the bridge domain, specify the interface name for each Ethernet interface to include that you configured at the **[edit interfaces]** hierarchy level.

> **NOTE**: A maximum of 4000 active logical interfaces are supported on a bridge domain or on each mesh group in a VPLS routing instance configured for Layer 2 bridging.

To associate a routing interface with a bridge domain, include the **routing-interface** *routing-interface-name* statement and specify a *routing-interface-name* you configured at the **[edit interfaces irb]** hierarchy level. You can configure only one routing interface for each bridge domain. For more information about how to configure logical and routing interfaces, see the Junos OS Network Interfaces Library for Routing Devices.

In Junos OS Release 9.0 and later, IRB interfaces are supported for multicast snooping. For more information about multicast snooping, see the *Understanding Multicast Snooping and VPLS Root Protection*.

In Junos 11.4 and later, IP multicast is supported on Layer 2 trunk ports through IRB interfaces using the Trio chipset.

In Junos OS Release 9.6 and later, in multihomed VPLS configurations, you can configure VPLS to keep a VPLS connection up if only an IRB interface is available by configuring the **irb** option for the **connectivity-type** statement at the **[edit routing-instances *routing-instance-name* protocols vpls]** hierarchy level. The **connectivity-type** statement has two options, **ce** and **irb**. The **ce** option is the default and specifies that a CE interface is required to maintain the VPLS connection. By default, if only an IRB interface is available, the VPLS connection is brought down. For more information about configuring VPNs, see the *Junos VPN Configuration Guide*.

> **NOTE**: When you configure IRB interfaces in more than one logical system on a device, all of the of the IRB logical interfaces share the same MAC address.

Integrated Bridging and Routing (IRB) interfaces are used to tie together Layer 2 switched and Layer 3 routed domains on MX routers. MX routers support classifiers and rewrite rules on the IRB interface at the **[edit class-of-service interfaces irb unit logical-unit-number]** level of the hierarchy. All types of classifiers and rewrite rules are allowed, including IEEE 802.1p.

NOTE: The IRB classifiers and rewrite rules are used only for *routed* packets; in other words, it is for traffic that originated in the Layer 2 domain and is then routed through IRB into the Layer 3 domain, or vice versa. Only IEEE classifiers and IEEE rewrite rules are allowed for pure Layer 2 interfaces within a bridge domain.

RELATED DOCUMENTATION

Understanding Layer 2 Learning and Forwarding **| 3**

Understanding Layer 2 Learning and Forwarding for Bridge Domains **| 68**

Understanding Layer 2 Learning and Forwarding for Bridge Domains Functioning as Switches with Layer 2 Trunk Ports **| 85**

Configuring a Bridge Domain on ACX Series Routers

A bridge domain must include a set of logical interfaces that participate in Layer 2 learning and forwarding.

To configure a bridge domain, include the following statements:

```
[edit]
bridge-domains {
    bridge-domain-name {
        interface interface-name;
        vlan-id (none | number);
        vlan-id-list [ vlan-id-numbers ];
    }
}
```

NOTE: The Layer 2 CLI configurations and show commands for ACX5048 and ACX5096 routers differ compared to other ACX Series routers. For more information, see *Layer 2 Next Generation Mode for ACX Series*.

You cannot use the slash (/) character in bridge domain names. If you do, the configuration does not commit and an error is generated.

For the **vlan-id** statement, you can specify either a valid VLAN identifier or **none**.

To include one or more logical interfaces in the bridge domain, specify an interface name for an Ethernet interface you configured at the **[edit bridge-domains** *bridge-domain-name*] hierarchy level.

To configure a layer 2 logical interface to be included in a bridge domain, you can either include the **encapsulation vlan-bridge** statement under the logical interface, or the **encapsulation ethernet-bridge** statement under the physical interface.

> **NOTE**: A maximum of 1000 logical interfaces can be configured on a physical interface. You can configure a maximum of 3000 bridge domains on an ACX Series router.

RELATED DOCUMENTATION

Example: Configuring Basic Layer 2 Switching on MX Series

IN THIS SECTION

This example shows how to configure Layer 2 switching with all interfaces participating in a single VLAN.

Requirements

No special configuration beyond device initialization is required before configuring this example.

This example uses an MX Series device to perform Layer 2 switching.

Overview

In this example, a single MX Series device is configured to act as a basic single-VLAN switch. Three connections are in place. The connections from the MX Series device attach to Junos OS routers, but the routers are used here for testing purposes only. In place of routers, you can use any IP networking devices.

Topology

shows the sample network.

Figure 1: Basic Layer 2 Switching

10.0.0/24 VLAN 600

shows the configuration for all of the devices in .

The section describes the steps on Device S1.

Configuration

IN THIS SECTION

CLI Quick Configuration

To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, and then copy and paste the commands into the CLI at the **[edit]** hierarchy level.

Device S1

```
set interfaces ge-2/0/0 vlan-tagging
set interfaces ge-2/0/0 encapsulation extended-vlan-bridge
set interfaces ge-2/0/0 unit 0 vlan-id 600
set interfaces ge-2/0/1 vlan-tagging
set interfaces ge-2/0/1 encapsulation extended-vlan-bridge
set interfaces ge-2/0/1 unit 0 vlan-id 600
set interfaces ge-2/0/2 vlan-tagging
set interfaces ge-2/0/2 encapsulation extended-vlan-bridge
set interfaces ge-2/0/2 unit 0 vlan-id 600
set bridge-domains customer1 domain-type bridge
set bridge-domains customer1 interface ge-2/0/0.0
set bridge-domains customer1 interface ge-2/0/2.0
set bridge-domains customer1 interface ge-2/0/1.0
```

Device R1

```
set interfaces ge-1/3/2 vlan-tagging
set interfaces ge-1/3/2 unit 0 vlan-id 600
set interfaces ge-1/3/2 unit 0 family inet address 10.0.0.1/24
```

Device R2

```
set interfaces ge-3/1/0 vlan-tagging
set interfaces ge-3/1/0 unit 0 vlan-id 600
set interfaces ge-3/1/0 unit 0 family inet address 10.0.0.2/24
```

Device R3

```
set interfaces ge-2/0/1 vlan-tagging
set interfaces ge-2/0/1 unit 0 vlan-id 600
set interfaces ge-2/0/1 unit 0 family inet address 10.0.0.3/24
```

Procedure

Step-by-Step Procedure

The following example requires that you navigate various levels in the configuration hierarchy. For information about navigating the CLI, see *Using the CLI Editor in Configuration Mode* in the Junos OS CLI User Guide.

To configure Device S1:

1. Configure the device interfaces.

```
[edit interfaces]
user@S1# set interfaces ge-2/0/0 vlan-tagging
user@S1# set interfaces ge-2/0/0 encapsulation extended-vlan-bridge
user@S1# set interfaces ge-2/0/0 unit 0 vlan-id 600
user@S1# set interfaces ge-2/0/1 vlan-tagging
user@S1# set interfaces ge-2/0/1 encapsulation extended-vlan-bridge
user@S1# set interfaces ge-2/0/1 unit 0 vlan-id 600
user@S1# set interfaces ge-2/0/2 vlan-tagging
user@S1# set interfaces ge-2/0/2 encapsulation extended-vlan-bridge
user@S1# set interfaces ge-2/0/2 unit 0 vlan-id 600
```

2. Configure the bridge domain.

```
[edit interfaces]
user@S1# set bridge-domains customer1 domain-type bridge
user@S1# set bridge-domains customer1 interface ge-2/0/0.0
user@S1# set bridge-domains customer1 interface ge-2/0/2.0
user@S1# set bridge-domains customer1 interface ge-2/0/1.0
```

Results

From configuration mode, confirm your configuration by entering the **show interfaces** and **show bridge-domains** commands. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
user@S1# show interfaces
ge-2/0/0 {
```

```
        vlan-tagging;
        encapsulation extended-vlan-bridge;
        unit 0 {
            vlan-id 600;
        }
    }
    ge-2/0/1 {
        vlan-tagging;
        encapsulation extended-vlan-bridge;
        unit 0 {
            vlan-id 600;
        }
    }
    ge-2/0/2 {
        vlan-tagging;
        encapsulation extended-vlan-bridge;
        unit 0 {
            vlan-id 600;
        }
    }
```

```
user@S1# show bridge-domains
customer1 {
    domain-type bridge;
    interface ge-2/0/0.0;
    interface ge-2/0/2.0;
    interface ge-2/0/1.0;
}
```

If you are done configuring the device, enter **commit** from configuration mode.

Verification

IN THIS SECTION

Confirm that the configuration is working properly.

Confirming the MAC Address Learning

Purpose

Display Layer 2 MAC address information.

Action

- From Device S1, run the **show bridge mac-table** command.

```
user@S1> show bridge mac-table

MAC flags (S -static MAC, D -dynamic MAC, L -locally learned, C -Control MAC
          SE -Statistics enabled, NM -Non configured MAC, R -Remote PE MAC)

Routing instance : default-switch
 Bridging domain : customer1, VLAN : NA
   MAC                  MAC       Logical            NH      RTR
   address              flags     interface          Index   ID
   00:12:1e:ee:34:dd    D         ge-2/0/2.0
   00:1d:b5:5e:86:79    D         ge-2/0/0.0
   00:21:59:0f:35:2b    D         ge-2/0/1.0
```

- From Device S1, run the **show bridge mac-table extensive** command.

```
user@S1> show bridge mac-table extensive

MAC address: 00:12:1e:ee:34:dd
  Routing instance: default-switch
```

```
     Bridging domain: customer1, VLAN : NA
     Learning interface: ge-2/0/2.0
     Layer 2 flags: in_hash,in_ifd,in_ifl,in_vlan,in_rtt,kernel,in_ifbd
     Epoch: 1                          Sequence number: 0
     Learning mask: 0x00000004

  MAC address: 00:1d:b5:5e:86:79
    Routing instance: default-switch
     Bridging domain: customer1, VLAN : NA
     Learning interface: ge-2/0/0.0
     Layer 2 flags: in_hash,in_ifd,in_ifl,in_vlan,in_rtt,kernel,in_ifbd
     Epoch: 1                          Sequence number: 0
     Learning mask: 0x00000004

  MAC address: 00:21:59:0f:35:2b
    Routing instance: default-switch
     Bridging domain: customer1, VLAN : NA
     Learning interface: ge-2/0/1.0
     Layer 2 flags: in_hash,in_ifd,in_ifl,in_vlan,in_rtt,kernel,in_ifbd
     Epoch: 3                          Sequence number: 0
     Learning mask: 0x00000004
```

Meaning

The output shows that the MAC addresses have been learned.

Making Sure That the Attached Devices Can Reach Each Other

Purpose

Verify connectivity.

Action

```
user@R1> ping 10.0.0.2
PING 10.0.0.2 (10.0.0.2): 56 data bytes
64 bytes from 10.0.0.2: icmp_seq=0 ttl=64 time=1.178 ms
64 bytes from 10.0.0.2: icmp_seq=1 ttl=64 time=1.192 ms
64 bytes from 10.0.0.2: icmp_seq=2 ttl=64 time=1.149 ms
^C
```

```
--- 10.0.0.2 ping statistics ---
3 packets transmitted, 3 packets received, 0% packet loss
round-trip min/avg/max/stddev = 1.149/1.173/1.192/0.018 ms
```

```
user@R1> ping 10.0.0.3
PING 10.0.0.3 (10.0.0.3): 56 data bytes
64 bytes from 10.0.0.3: icmp_seq=0 ttl=64 time=1.189 ms
64 bytes from 10.0.0.3: icmp_seq=1 ttl=64 time=1.175 ms
64 bytes from 10.0.0.3: icmp_seq=2 ttl=64 time=1.178 ms
64 bytes from 10.0.0.3: icmp_seq=3 ttl=64 time=1.133 ms
^C
--- 10.0.0.3 ping statistics ---
4 packets transmitted, 4 packets received, 0% packet loss
round-trip min/avg/max/stddev = 1.133/1.169/1.189/0.021 ms
```

```
user@R2> ping 10.0.0.3
PING 10.0.0.3 (10.0.0.3): 56 data bytes
64 bytes from 10.0.0.3: icmp_seq=0 ttl=64 time=0.762 ms
64 bytes from 10.0.0.3: icmp_seq=1 ttl=64 time=0.651 ms
64 bytes from 10.0.0.3: icmp_seq=2 ttl=64 time=0.722 ms
64 bytes from 10.0.0.3: icmp_seq=3 ttl=64 time=0.705 ms
^C
--- 10.0.0.3 ping statistics ---
4 packets transmitted, 4 packets received, 0% packet loss
round-trip min/avg/max/stddev = 0.651/0.710/0.762/0.040 ms
```

Meaning

The output shows that the attached devices have established Layer 3 connectivity, with Device S1 doing transparent Layer 2 bridging.

Checking the Bridge Domain

Purpose

Display bridge domain information.

Action

```
user@S1> show bridge domain extensive


Routing instance: default-switch
Bridge domain: customer1                       State: Active
Bridge VLAN ID: NA
Interfaces:
    ge-2/0/0.0
    ge-2/0/1.0
    ge-2/0/2.0
Total MAC count: 3
```

Meaning

The output shows that bridge domain is active.

Checking the Bridge Statistics

Purpose

Display bridge statistics.

Action

```
user@S1> show bridge statistics


   Local interface: ge-2/0/0.0, Index: 65543
     Broadcast packets:                    0
     Broadcast bytes  :                    0
     Multicast packets:                   80
     Multicast bytes  :                 8160
     Flooded packets  :                    0
     Flooded bytes    :                    0
     Unicast packets  :                    1
     Unicast bytes    :                   64
     Current MAC count:                    1 (Limit 1024)
   Local interface: ge-2/0/2.0, Index: 324
     Broadcast packets:                    0
```

```
      Broadcast bytes  :                        0
      Multicast packets:                       80
      Multicast bytes  :                     8160
      Flooded packets  :                        1
      Flooded bytes    :                       74
      Unicast packets  :                       52
      Unicast bytes    :                     4332
      Current MAC count:                        1 (Limit 1024)
   Local interface: ge-2/0/1.0, Index: 196613
      Broadcast packets:                        2
      Broadcast bytes  :                      128
      Multicast packets:                        0
      Multicast bytes  :                        0
      Flooded packets  :                        1
      Flooded bytes    :                       93
      Unicast packets  :                       51
      Unicast bytes    :                     4249
      Current MAC count:                        1 (Limit 1024)
```

Meaning

The output shows that bridge domain interfaces are sending and receiving packets.

Checking the Bridge Flooding

Purpose

Display bridge flooding information.

Action

```
user@S1> show bridge flood extensive

Name: __juniper_private1__
CEs: 0
VEs: 0
Name: default-switch
CEs: 3
VEs: 0
Bridging domain: customer1
```

```
    Flood route prefix: 0x30003/51
    Flood route type: FLOOD_GRP_COMP_NH
    Flood route owner: __all_ces__
    Flood group name: __all_ces__
    Flood group index: 1
    Nexthop type: comp
    Nexthop index: 568
      Flooding to:
      Name            Type           NhType         Index
      __all_ces__     Group           comp            562
         Composition: split-horizon
         Flooding to:
         Name            Type           NhType         Index
         ge-2/0/0.0      CE              ucst            524
         ge-2/0/1.0      CE              ucst            513
         ge-2/0/2.0      CE              ucst            523

    Flood route prefix: 0x30005/51
    Flood route type: FLOOD_GRP_COMP_NH
    Flood route owner: __re_flood__
    Flood group name: __re_flood__
    Flood group index: 65534
    Nexthop type: comp
    Nexthop index: 565
      Flooding to:
      Name            Type           NhType         Index
      __all_ces__     Group           comp            562
         Composition: split-horizon
         Flooding to:
         Name            Type           NhType         Index
         ge-2/0/0.0      CE              ucst            524
         ge-2/0/1.0      CE              ucst            513
         ge-2/0/2.0      CE              ucst            523
```

Meaning

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 bridge domain 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.

Checking Layer 2 Learning

Purpose

Display Layer 2 learning information for all the interfaces.

Action

```
user@S1>  show l2-learning interface

Routing Instance Name : default-switch
Logical Interface flags (DL -disable learning, AD -packet action drop,
                  LH - MAC limit hit, DN - Interface Down )
Logical            BD         MAC        STP          Logical
Interface          Name       Limit      State        Interface flags
ge-2/0/2.0                    0
                   custom..   1024       Forwarding
Routing Instance Name : default-switch
Logical Interface flags (DL -disable learning, AD -packet action drop,
                  LH - MAC limit hit, DN - Interface Down )
Logical            BD         MAC        STP          Logical
Interface          Name       Limit      State        Interface flags
ge-2/0/0.0                    0
                   custom..   1024       Forwarding
Routing Instance Name : default-switch
Logical Interface flags (DL -disable learning, AD -packet action drop,
                  LH - MAC limit hit, DN - Interface Down )
Logical            BD         MAC        STP          Logical
Interface          Name       Limit      State        Interface flags
ge-2/0/1.0                    0
                   custom..   1024       Forwarding
```

RELATED DOCUMENTATION

Understanding OSPF Areas

Examples: Configuring OSPF Stub and Not-So-Stubby Areas

Configuring VLAN Identifiers for Bridge Domains and VPLS Routing Instances

For a bridge domain that is performing Layer 2 switching only, you do not have to specify a VLAN identifier.

For a bridge domain that is performing Layer 3 IP routing, you must specify either a VLAN identifier or dual VLAN identifier tags.

For a VPLS routing instance, you must specify either a VLAN identifier or dual VLAN identifier tags.

You can configure VLAN identifiers for a bridge domain or a VPLS routing instance in the following ways:

- By using the **input-vlan-map** and the **output-vlan-map** statements at the **[edit interfaces *interface-name*]** or **[edit logical-systems *logical-system-name* interfaces *interface-name*]** hierarchy level to configure VLAN mapping. For information about configuring input and output VLAN maps to stack and rewrite VLAN tags in incoming or outgoing frames, see the Junos OS Network Interfaces Library for Routing Devices.

- By using either the **vlan-id** statement or the **vlan-tags** statement to configure a normalizing VLAN identifier. This topic describes how normalizing VLAN identifiers are processed and translated in a bridge domain or a VPLS routing instance.

The **vlan-id** and **vlan-tags** statements are used to specify the normalizing VLAN identifier under the bridge domain or VPLS routing instance. The normalizing VLAN identifier is used to perform the following functions:

- Translate, or normalize, the VLAN tags of received packets received into a learn VLAN identifier.

- Create multiple learning domains that each contain a learn VLAN identifier. A learning domain is a MAC address database to which MAC addresses are added based on the learn VLAN identifier.

> **NOTE**: You cannot configure VLAN mapping using the **input-vlan-map** and **output-vlan-map** statements if you configure a normalizing VLAN identifier for a bridge domain or VPLS routing instance using the **vlan-id** or **vlan-tags** statements.

To configure a VLAN identifier for a bridge domain, include either the **vlan-id** or the **vlan-tags** statement at the **[edit interfaces *interface-name* unit *logic-unit-number* family bridge]** or **[edit logical-systems *logical-system-name* interfaces *interface-name* unit *logic-unit-number* family bridge]** hierarchy level, and then include that logical interface in the bridge domain configuration. For more information about configuring a bridge domain, see "Configuring a Bridge Domain" on page 11.

For a VPLS routing instance, include either the **vlan-id** or **vlan-tags** statement at the **[edit interfaces** *interface-name* **unit** *logic-unit-number***]** or **[edit logical-systems** *logical-system-name* **interfaces** *interface-name* **unit** *logic-unit-number***]** hierarchy level, and then include that logical interface in the VPLS routing instance configuration. For more information about configuring a VPLS routing instance, see the Junos OS VPNs Library for Routing Devices.

> **NOTE**: The maximum number of Layer 2 interfaces that you can associate with a bridge domain or a VPLS instance on MX Series routers is 4000.

> **NOTE**: For a single bridge domain or VPLS routing instance, you can include either the **vlan-id** or the **vlan-tags** statement, but not both. If you do not configure a **vlan-id**, **vlan-tags**, or **vlan-id-list** [*vlan-id-numbers*] for the bridge domain or the VPLS routing instance, the Layer 2 packets received are forwarded to the outbound Layer 2 interface without having the VLAN tag modified unless an **output-vlan-map** is configured on the Layer 2 interface. This results in a frame being forwarded to a Layer 2 interface with a VLAN tag that is different from what is configured for the Layer 2 interface. Note that a frame received from the Layer 2 interface is still required to match the VLAN tag(s) specified in the interface configuration. The invalid configuration may cause a Layer 2 loop to occur.

The VLAN tags associated with the inbound logical interface are compared with the normalizing VLAN identifier. If the tags are different, they are rewritten as described in Table 3 on page 32. The source MAC address of a received packet is learned based on the normalizing VLAN identifier.

> **NOTE**: You do not have to specify a VLAN identifier for a bridge domain that is performing Layer 2 switching only. To support Layer 3 IP routing, you must specify either a VLAN identifier or a pair of VLAN tags. However, you cannot specify the same VLAN identifier for more than one bridge domain within a routing instance. Each bridge domain must have a unique VLAN identifier.

If the VLAN tags associated with the outbound logical interface and the normalizing VLAN identifier are different, the normalizing VLAN identifier is rewritten to match the VLAN tags of the outbound logical interface, as described in Table 4 on page 34.

For the packets sent over the VPLS routing instance to be tagged by the normalizing VLAN identifier, include one of the following configuration statements:

- **vlan-id** *number* to tag all packets that are sent over the VPLS virtual tunnel (VT) interfaces with the VLAN identifier.

- **vlan-tags outer** *number* **inner** *number* to tag all packets sent over the VPLS VT interfaces with dual outer and inner VLAN tags.

Use the **vlan-id none** statement to have the VLAN tags removed from packets associated with an inbound logical interface when those packets are sent over VPLS VT interfaces. Note that those packets might still be sent with other customer VLAN tags.

The **vlan-id all** statement enables you to configure bridging for several VLANs with a minimum amount of configuration. Configuring this statement creates a learning domain for:

- Each inner VLAN, or learn VLAN, identifier of a logical interface configured with two VLAN tags

- Each VLAN, or learn VLAN, identifier of a logical interface configured with one VLAN tag

We recommend that you do not use customer VLAN IDs in a VPLS routing instance because customer VLAN IDs are used for learning only.

You should use the service VLAN ID in a VPLS routing instance, as in the following configuration:

```
[edit]
interface ge-1/1/1 {
    vlan-tagging;
    unit 1 {
        vlan-id s1; /* Service vlan */
        encapsulation vlan-vpls;
        input-vlan-map pop; /* Pop the service vlan on input */
        output-vlan-map push; /* Push the service vlan on output */
    }
}
interface ge-1/1/2 {
    encapsulation ethernet-vpls;
    unit 0;
}
routing-instance {
    V1 {
        instance-type vpls;
        vlan-id all;
        interface ge-1/1/1.1;
        interface ge-1/1/2.0;
    }
}
```

> **NOTE**: If you configure the **vlan-id all** statement in a VPLS routing instance, we recommend using the **input-vlan-map pop** and **output-vlan-map push** statements on the logical interface to pop the service VLAN ID on input and push the service VLAN ID on output and in this way limit the impact of doubly-tagged frames on scaling. You cannot use the native **vlan- id** statement when the **vlan-id all** statement is included in the configuration.

The **vlan-id-list [** *vlan-id-numbers* **]** statement enables you to configure bridging for multiple VLANs on a trunk interface. Configuring this statement creates a learning domain for:

- Each VLAN listed: **vlan-id-list [100 200 300]**

- Each VLAN in a range: **vlan-id-list [100-200]**

- Each VLAN in a list and range combination: **vlan-id-list [50, 100-200, 300]**

The following steps outline the process for bridging a packet received over a Layer 2 logical interface when you specify a normalizing VLAN identifier using either the **vlan-id** *number* or **vlan-tags** statement for a bridge domain or a VPLS routing instance:

1. When a packet is received on a physical port, it is accepted only if the VLAN identifier of the packet matches the VLAN identifier of one of the logical interfaces configured on that port.
2. The VLAN tags of the received packet are then compared with the normalizing VLAN identifier. If the VLAN tags of the packet are different from the normalizing VLAN identifier, the VLAN tags are rewritten as described in Table 3 on page 32.
3. If the source MAC address of the received packet is not present in the source MAC table, it is learned based on the normalizing VLAN identifier.
4. The packet is then forwarded toward one or more outbound Layer 2 logical interfaces based on the destination MAC address. A packet with a known unicast destination MAC address is forwarded only to one outbound logical interface. For each outbound Layer 2 logical interface, the normalizing VLAN identifier configured for the bridge domain or VPLS routing instance is compared with the VLAN tags configured on that logical interface. If the VLAN tags associated with an outbound logical interface do not match the normalizing VLAN identifier configured for the bridge domain or VPLS routing instance, the VLAN tags are rewritten as described in Table 4 on page 34.

The tables below show how VLAN tags are applied for traffic sent to and from the bridge domain, depending on how the **vlan-id** and **vlan-tags** statements are configured for the bridge domain and on how VLAN identifiers are configured for the logical interfaces in a bridge domain or VPLS routing instance. Depending on your configuration, the following rewrite operations are performed on VLAN tags:

- **pop**—Remove a VLAN tag from the top of the VLAN tag stack.

- **pop-pop**—Remove both the outer and inner VLAN tags of the frame.

- **pop-swap**—Remove the outer VLAN tag of the frame and replace the inner VLAN tag of the frame.

- **swap**—Replace the VLAN tag of the frame.

- **push**—Add a new VLAN tag to the top of the VLAN stack.

- **push-push**—Push two VLAN tags in front of the frame.

- **swap-push**—Replace the VLAN tag of the frame and add a new VLAN tag to the top of the VLAN stack.

- **swap-swap**—Replace both the outer and inner VLAN tags of the frame.

Table 3 on page 32 shows specific examples of how the VLAN tags for packets sent to the bridge domain are processed and translated, depending on your configuration. "–" means that the statement is not supported for the specified logical interface VLAN identifier. "No operation" means that the VLAN tags of the received packet are not translated for the specified input logical interface.

Table 3: Statement Usage and Input Rewrite Operations for VLAN Identifiers for a Bridge Domain

VLAN Identifier of Logical Interface	VLAN Configurations for Bridge Domain			
	vlan-id none	vlan-id 200	vlan-id all	vlan tags outer 100 inner 300
none	No operation	push 200	–	push 100, push 300
200	pop 200	No operation	No operation	swap 200 to 300, push 100
1000	pop 1000	swap 1000 to 200	No operation	swap 1000 to 300, push 100

Table 3: Statement Usage and Input Rewrite Operations for VLAN Identifiers for a Bridge Domain *(Continued)*

VLAN Identifier of Logical Interface	VLAN Configurations for Bridge Domain			
	vlan-id none	vlan-id 200	vlan-id all	vlan tags outer 100 inner 300
vlan-tags outer 2000 inner 300	pop 2000, pop 300	pop 2000, swap 300 to 200	pop 2000	swap 2000 to 100
vlan-tags outer 100 inner 400	pop 100, pop 400	pop 100, swap 400 to 200	pop 100	swap 400 to 300
vlan-id-range 10-100	–	–	No operation	–
vlan-tags outer 200 inner-range 10-100	–	–	pop 200	–

Table 4 on page 34 shows specific examples of how the VLAN tags for packets sent from the bridge domain are processed and translated, depending on your configuration. "–" means that the statement is not supported for the specified logical interface VLAN identifier. "No operation" means that the VLAN tags of the outbound packet are not translated for the specified output logical interface.

Table 4: Statement Usage and Output Rewrite Operations for VLAN Identifiers for a Bridge Domain

VLAN Identifier of Logical Interface	VLAN Configurations for Bridge Domain			
	vlan-id none	vlan-id 200	vlan-id all	vlan tags outer 100 inner 300
none	no operation	pop 200	–	pop 100, pop 300
200	push 200	No operation	No operation	pop 100, swap 300 to 200
1000	push 1000	swap 200 to 1000	No operation	pop 100, swap 300 to 1000
vlan-tags outer 2000 inner 300	push 2000, push 300	swap 200 to 300, push 2000	push 2000	swap 100 to 2000
vlan-tags outer 100 inner 400	push 100, push 400	swap 200 to 400, push 100	push 100	swap 300 to 400
vlan-id-range 10-100	–	–	No operation	–
vlan-tags outer 200 inner-range 10-100	–	–	push 200	–

Configuring VLAN Identifiers for Bridge Domains in ACX Series

You can configure VLAN identifiers for a bridge domain for normalization in the following ways:

- Configure VLAN mapping by using the **input-vlan-map** and the **output-vlan-map** statements at the **[edit interfaces *interface-name*]** hierarchy level.

- Configure an implicit normalizing VLAN identifier under the bridge domain by using the **vlan-id** statement at the **[edit bridge-domains *bridge-domain-name*]** hierarchy level.

> **NOTE**: You cannot configure VLAN mapping by using the **input-vlan-map** and **output-vlan-map** statements if you configure a normalizing VLAN identifier for a bridge domain by using the **vlan-id** statement.

You can use the **vlan-id-list [*vlan-id-numbers*]** statement to configure bridging for multiple VLANs. Configuring this statement creates a bridge domain for:

- Each VLAN listed—for example, **vlan-id-list [100 200 300]**

- Each VLAN in a range—for example, **vlan-id-list [100-200]**

- Each VLAN in a list and range combination—for example, **vlan-id-list [50, 100-200, 300]**

Configuring Bridge Domains as Switches for Layer 2 Trunk Ports

You can configure a set of bridge domains that are associated with a Layer 2 trunk port. The set of bridge domains function as a switch. Packets received on a trunk interface are forwarded within a bridge domain that has the same VLAN identifier. A trunk interface also provides support for IRB, which provides support for Layer 2 bridging and Layer 3 IP routing on the same interface.

To configure a Layer 2 trunk port and set of bridge domains, include the following statements:

```
[edit interfaces]
interface-name {
    unit number {
        family bridge {
            interface-mode access;
            vlan-id number;
        }
    }
}
interface-name {
    native-vlan-id number;
    unit number {
        family bridge {
            interface-mode trunk;
            vlan-id-list [ vlan-id-numbers ];
        }
    }
}
[edit bridge-domains]
bridge-domain-name {
    vlan-id number;
    vlan-id-list [ vlan-id-numbers ];
```

```
        .   .   .   .
    }
```

For **interface-mode trunk**, you can include the **vlan-id-list** statement.

You must configure a bridge domain and VLAN identifier for each VLAN associated with the trunk interface. You can configure one or more trunk or access interfaces at the **[edit interfaces]** hierarchy level. An access interface enables you to accept packets with no VLAN identifier. For more information about configuring trunk and access interfaces, see the Interfaces User Guide for Security Devices.

RELATED DOCUMENTATION

Configuring a Bridge Domain | 11

3
CHAPTER

Configuring Layer 2 Virtual Switches

Understanding Layer 2 Virtual Switches

On MX Series routers only, you can group one or more bridge domains to form a virtual switch to isolate a LAN segment with its spanning-tree protocol instance and separate its VLAN ID space. A bridge domain consists of a set of logical ports that share the same flooding or broadcast characteristics. Like a virtual LAN, a bridge domain spans one or more ports of multiple devices. You can configure multiple virtual switches, each of which operates independently of the other virtual switches on the routing platform. Thus, each virtual switch can participate in a different Layer 2 network.

You can configure a virtual switch to participate only in Layer 2 bridging and optionally to perform Layer 3 routing. In addition, you can configure one of three Layer 2 control protocols—Spanning-Tree Protocol, Rapid Spanning-Tree Protocol (RSTP), or Multiple Spanning-Tree Protocol (MSTP)—to prevent forwarding loops. For more information about how to configure Layer 2 logical ports on an interface, see the Junos OS Network Interfaces Library for Routing Devices.

In Junos OS Release 9.2 and later, you can associate one or more logical interfaces configured as trunk interfaces with a virtual switch. A trunk interface, or Layer 2 trunk port, enables you to configure a *logical interface* to represent multiple VLANs on the physical interface. Packets received on a trunk interface are forwarded within a bridge domain that has same VLAN identifier. For more information about how to configure trunk interfaces, see the Junos OS Network Interfaces Library for Routing Devices.

You can also configure Layer 2 forwarding and learning properties for the virtual switch as well as any bridge domains that belong to a virtual switch. .

For more information about configuring a routing instance for Layer 2 VPN, see the Junos OS VPNs Library for Routing Devices. .

RELATED DOCUMENTATION

Configuring a Layer 2 Control Protocol Routing Instance

Understanding Layer 2 Learning and Forwarding for Bridge Domains | 68

Understanding Layer 2 Protocol Tunneling

Configuring a Layer 2 Virtual Switch

A Layer 2 virtual switch, which isolates a LAN segment with its spanning-tree protocol instance and separates its VLAN ID space, filters and forwards traffic only at the data link layer. Layer 3 routing is not

performed. Each bridge domain consists of a set of logical ports that participate in Layer 2 learning and forwarding. A virtual switch represents a Layer 2 network.

Two main types of interfaces are used in virtual switch hierarchies:

- Layer 2 logical interface—This type of interface uses the VLAN-ID as a virtual circuit identifier and the scope of the VLAN-ID is local to the interface port. This type of interface is often used in service-provider-centric applications.

- Access or trunk interface—This type of interface uses a VLAN-ID with global significance. The access or trunk interface is implicitly associated with bridge domains based on VLAN membership. Access or trunk interfaces are typically used in enterprise-centric applications.

> **NOTE**: The difference between access interfaces and trunk interfaces is that access interfaces can be part of one VLAN only and the interface is normally attached to an end-user device (packets are implicitly associated with the configured VLAN). In contrast, trunk interfaces multiplex traffic from multiple VLANs and usually interconnect switches.

To configure a Layer 2 virtual switch, include the following statements:

```
[edit]
routing-instances {
    routing-instance-name (
        instance-type virtual-switch;
        bridge-domains {
            bridge-domain-name {
                domain-type bridge;
                interface interface-name;
                vlan-id (all | none | number); # Cannot be used with 'vlan-tags'
statement
                vlan-id-list [ vlan-id-numbers ];
                vlan-tags outer number inner number; # Cannot be used with 'vlan-
id' statement
            }
        }
        protocols {
            mstp {
                ...mstp-configuration ...
            }
        }
```

```
        }
    }
```

To enable a virtual switch, you must specify **virtual-switch** as the **instance-type**.

For each bridge domain that you configure for the virtual switch, specify a *bridge-domain-name*. You must also specify the value **bridge** for the **domain-type** statement.

For the **vlan-id** statement, you can specify either a valid VLAN identifier or the **none** or **all** options.

The **all** option is not supported with IRB.

> **NOTE**: You do not have to specify a VLAN identifier for a bridge domain. However, you cannot specify the same VLAN identifier for more than one bridge domain within a virtual switch. Each bridge domain within a virtual switch must have a unique VLAN identifier.

> **NOTE**: For a single bridge domain, you can include either the **vlan-id** statement or the **vlan-tags** statement, but not both. The **vlan-id** statement, **vlan-id-list** statement, and **vlan-tags** statement are mutually exclusive.

The **vlan-id-list** statement allows you to automatically create multiple bridge-domains for each vlan-id in the list.

To specify one or more logical interfaces to include in the bridge domain, specify an *interface-name* for an Ethernet interface you configured at the **[edit interfaces]** hierarchy level. For more information, see the Junos OS Network Interfaces Library for Routing Devices.

RELATED DOCUMENTATION

Configuring a Virtual Switch Routing Instance on MX Series Routers

On MX Series routers only, use the **virtual-switch** routing instance type to isolate a LAN segment with its spanning-tree instance and to separate its VLAN ID space. A bridge domain consists of a set of ports that share the same flooding or broadcast characteristics. Each virtual switch represents a Layer 2 network. You can optionally configure a virtual switch to support Integrated Routing and Bridging (IRB), which facilitates simultaneous Layer 2 bridging and Layer 3 IP routing on the same interface. You can also configure Layer 2 control protocols to provide loop resolution. Protocols supported include the Spanning-Tree Protocol (STP), Rapid Spanning-Tree Protocols (RSTP), Multiple Spanning-Tree Protocol (MSTP), and VLAN Spanning-Tree Protocol (VSTP).

To create a routing instance for a virtual switch, include at least the following statements in the configuration:

```
[edit]
routing-instances {
    routing-instance-name
        instance-type virtual-switch;
        bridge-domains {
            bridge-domain-name {
                domain-type bridge;
                interface interface-name;
                vlan-id (all | none | number);
                vlan-tags outer number inner number;
            }
        }
        protocols {
            (rstp | mstp | vstp) {
                ...stp-configuration ...
            }
        }
    }
}
```

For more information about configuring virtual switches, see "Configuring a Layer 2 Virtual Switch " on page 39.

Configuring Integrated Routing and Bridging for a Bridge Domain in a Layer 2 Virtual Switch

Integrated routing and bridging (IRB) provides simultaneous support for Layer 2 bridging and Layer 3 IP routing on the same interface. IRB enables you to route local packets to another routed interface or to another bridge domain that has a Layer 3 protocol configured. You configure a logical routing interface by including the **irb** statement at **[edit interfaces]** hierarchy level and include that interface in the bridge domain. For more information about how to configure a routing interface, see the Junos OS Network Interfaces Library for Routing Devices.

> **NOTE**: You can include only one routing interface in a bridge domain.

To configure a virtual switch with IRB support, include the following statements:

```
[edit]
routing-instances {
    routing-instance-name {
        instance-type virtual-switch;
        bridge-domains {
            bridge-domain-name {
                domain-type bridge;
                interface interface-name;
                routing-interface routing-interface-name;
                vlan-id (none | number);
                vlan-tags outer number inner number;
            }
        }
    }
}
```

To enable a virtual switch, you must specify **virtual-switch** as the **instance-type**. The **instance-type virtual-switch** statement is not supported at the **[edit logical-systems *logical-system-name*]** hierarchy level.

For each bridge domain that you configure for the virtual switch, specify a *bridge-domain-name*. You must also specify the value **bridge** for the **domain-type** statement.

For the **vlan-id** statement, you can specify either a valid VLAN identifier or the **none** option.

> **NOTE**: For a single bridge domain, you can include either the **vlan-id** statement or the **vlan-tags** statement, but not both.

To include one or more logical interfaces in the bridge domain, specify the *interface-name* for each Ethernet interface to include that you configured at the **[edit interfaces irb]** hierarchy level.

To associate a routing interface with a bridge domain, include the **routing-interface** *routing-interface-name* statement and specify a *routing-interface-name* you configured at the **[edit interfaces irb]** hierarchy level. You can configure only one routing interface for each bridge domain. For more information about how to configure logical and routing interfaces, see the Junos OS Network Interfaces Library for Routing Devices.

> **NOTE**: If you configure a routing interface to support IRB in a bridge domain, you cannot use the **all** option for the **vlan-id** statement.

RELATED DOCUMENTATION

Configuring Integrated Routing and Bridging in ACX Series

Integrated routing and bridging (IRB) provides simultaneous support for Layer 2 bridging and Layer 3 routing on the same interface. IRB enables you to route packets to another routed interface or to another bridge domain that has an IRB interface configured. You configure a logical routing interface by including the **irb** statement at the **[edit interfaces]** hierarchy level and include that interface in the bridge domain. For more information about how to configure a routing interface, see the Junos OS Network Interfaces Library for Routing Devices.

> **NOTE**: You can include only one routing interface in a bridge domain.

The following are the list of features supported for IRB:

- Family **inet**, **inet6**, and **iso** are supported on an IRB interface.

- Routing protocols supported on an IRB interface are BGP, ISIS, OSPF, RIP, IGMP, and PIM.

- DHCP Relay with option 82 is supported on an IRB interface.

- IRB can be added in a VRF routing instance.

- VRRP is supported on an IRB inteface.

- Bidirectional Forwarding Detection (BFD) protocol is supported on an IRB interface.

- The following Class-of-Service configurations are supported on an IRB interface:

  - The IRB classifiers and rewrite on routed packets.

  - Fixed classifier can be applied on an IRB logical interface.

  - Firewall filters (multifield filter) can be used to assign forwarding class and loss priority. You should define a family inet or inet6 filter and apply it as the input filter on an IRB logical interface under family inet.

    > **NOTE**: **physical-interface-filter** is not supported for family inet6 filter on IRB logical interface.

  - Re-write can be applied only at the IRB interface level.

  - dscp, inet-precedence, ieee-802.1, and ieee-802.1ad values can be rewritten.

ACX routers do not support MPLS families on IRB.

IRB can be configured under the following hierarchies:

- [**edit intefaces irb interface_type**] hierarchy level

  - disable—Disables the interface

  - gratuitous-arp-reply—Enables gratuitous ARP reply

  - hold-time—Hold time for link up and link down

  - mtu—Maximum transmit packet size (256..9192)

  - no-gratuitous-arp-reply—Does not enable gratuitous ARP reply

  - no-gratuitous-arp-request—Ignores gratuitous ARP request

- [**edit interfaces irb.***unit* **family (inet | inet6 | iso)**] hierarchy level

- [**edit bridge-domains routing-interface interface irb.***unit*] hierarchy level

- [**edit routing-instances instance-type vrf**] hierarchy level

- [**edit protocols (***bgp* | *isis* | *ospf* | *rip* | *igmp* | *pim***) interface irb.***unit*] hierarchy level

- [**edit class-of-service interfaces irb**]] hierarchy level

In ACX5048 and ACX5096 routers, you can configure IRB at the [**edit vlans** *vlan-name*] **l3-interface irb.***unit*; level.

> **NOTE**: The Layer 2 CLI configurations and show commands for ACX5048 and ACX5096 routers differ compared to other ACX Series routers. For more information, see *Layer 2 Next Generation Mode for ACX Series*.

To configure a bridge domain with IRB support, include the following statements:

```
[edit]
bridge-domains {
    bridge-domain-name {
        domain-type bridge;
        interface interface-name;
        routing-interface routing-interface-name;
        vlan-id (none | number);
        vlan-tags outer number inner number;
    }
}
```

For each bridge domain that you configure, specify a *bridge-domain-name*. You must also specify the value **bridge** for the **domain-type** statement.

For the **vlan-id** statement, you can specify either a valid VLAN identifier or the **none** option.

The **vlan-tags** statement enables you to specify a pair of VLAN identifiers; an **outer** tag and an **inner** tag.

> **NOTE**: For a single bridge domain, you can include either the **vlan-id** statement or the **vlan-tags** statement, but not both.

To include one or more logical interfaces in the bridge domain, specify the *interface-name* for each Ethernet interface to include that you configured at the [**edit interfaces**] hierarchy level.

> **NOTE**: A maximum of 4000 active logical interfaces are supported on a bridge domain configured for Layer 2 bridging.

To associate a routing interface with a bridge domain, include the **routing-interface** *routing-interface-name* statement and specify a *routing-interface-name* you configured at the **[edit interfaces irb]** hierarchy level. You can configure only one routing interface for each bridge domain. For more information about how to configure logical and routing interfaces, see the Junos OS Network Interfaces Library for Routing Devices.

In Junos OS Release 9.0 and later, IRB interfaces are supported for multicast snooping. For more information about multicast snooping, see the Junos OS Multicast Protocols User Guide.

> **NOTE**: When you configure multiple IRB logical interfaces, all the IRB logical interfaces share the same MAC address.

The following is a sample configuration for IRB over bridge domain:

```
[edit]
interfaces {
    ge-1/0/0 {
        encapsulation flexible-ethernet-services;
            flexible-vlan-tagging;
            unit 0 {
                encapsulation vlan-bridge;
                vlan-id 100;
            }
        }
    }
    ge-1/0/1 {
        encapsulation flexible-ethernet-services;
            flexible-vlan-tagging;
            unit 0 {
                encapsulation vlan-bridge;
                vlan-id 100;
            }
        }
    }
    irb {
        unit 0 {
```

```
            family inet {
                address 10.0.1.2/24 {
                }
            }
        }
    }
}
bridge-domains {
    bd {
        domain-type bridge;
        vlan-id none;
        interface ge-1/0/0.0;
        interface ge-1/0/1.0;
        routing-interface irb.0;
    }
}
```

Configuring VPLS Ports in a Virtual Switch

In Junos OS Release 9.3 and later, you can configure VPLS ports in a virtual switch so that the logical interfaces of the Layer 2 bridge domains in the virtual switch can handle VPLS routing instance traffic. VPLS configuration no longer requires a dedicated routing instance of type **vpls**. Packets received on a Layer 2 trunk interface are forwarded within a bridge domain that has the same VLAN identifier.

A trunk interface is implicitly associated with bridge domains based on VLAN membership. Whereas access interfaces can be part of one VLAN only, trunk interfaces multiplex traffic from multiple VLANs and usually interconnect switches. A Layer 2 trunk port also supports IRB.

To configure VPLS ports in a virtual switch, perform the following tasks:

1. To configure the Layer 2 trunk ports that you will associate with the bridge domains in the virtual switch, include the following statements in the configuration:

```
[edit]
interfaces {
    interface-name {
        unit logical-unit-number { # Call this 'L2-trunk-port-A'
            family bridge {
                interface-mode trunk;
                vlan-id-list [ vlan-id-numbers ] ; # Trunk mode VLAN
membership for this interface
```

```
                }
            }
        }
        .
        .
        .
    interface-name {
        unit logical-unit-number { # Call this 'L2-trunk-port-B'
            family bridge {
                interface-mode trunk;
                vlan-id-list [ vlan-id-numbers ] ; # Trunk mode VLAN
 membership for this interface
            }
        }
    }
}
```

To configure a logical interface as a trunk port, include the **interface-mode** statement and the **trunk** option at the **[edit interfaces *interface-name* unit *logical-unit-number* family bridge]** hierarchy level.

To configure all the VLAN identifiers to associate with a Layer 2 trunk port, include the **vlan-id-list [*vlan-id-numbers*]** statement at the **[edit interfaces *interface-name* unit *logical-unit-number* family bridge]** hierarchy level.

Each of the logical interfaces "*L2-trunk-port-A*" and "*L2-trunk-port-B*" accepts packets tagged with any VLAN ID specified in the respective **vlan-id-list** statements.

2. To configure a virtual switch consisting of a set of bridge domains that are associated with one or more logical interfaces configured as a trunk ports, include the following statements in the configuration:

```
 [edit]
 routing-instance {
     routing-instance-name
         instance-type virtual-switch;
         interface L2-trunk-port-A; # Include one trunk port
         interface L2-trunk-port-B; # Include the other trunk port
         bridge-domains {
             bridge-domain-name-0 {
                 domain-type bridge;
                 vlan-id number;
             }
             bridge-domain-name-1 {
```

```
            domain-type bridge;
            vlan-id number;
        }
    }
    protocols {
        vpls {
            vpls-id number;
            ... vpls-configuration ...
        }
    }
}
}
```

To begin configuring a virtual switch, include the **instance-type** statement and the **virtual-switch** option at the **[edit routing-instances *routing-instance-name*]** hierarchy level.

To configure a virtual switch consisting of a set of bridge domains that are associated with one or more logical interfaces configured as a trunk ports, you must identify each logical interface by including the **interface *interface-name*** statement at the **[edit routing-instances *routing-instance-name*]** hierarchy level.

For each VLAN configured for a trunk port, you must configure a bridge-domain that includes the trunk port logical interface and uses a VLAN identifier within the range carried by that trunk interface. To configure, include the **domain-type bridge**, **vlan-id *number***, and statements at the **[edit routing-instances *routing-instance-name* bridge-domain *bridge-domain-name*]** hierarchy level.

RELATED DOCUMENTATION

Configuring a Bridge Domain | **11**

Configuring a Layer 2 Virtual Switch with a Layer 2 Trunk Port

You can associate one or more Layer 2 trunk interfaces with a virtual switch. A Layer 2 trunk interface enables you to configure a logical interface to represent multiple VLANs on the physical interface. Within the virtual switch, you configure a bridge domain and VLAN identifier for each VLAN identifier configured on the trunk interfaces. Packets received on a trunk interface are forwarded within a bridge

domain that has the same VLAN identifier. Each virtual switch you configure operates independently and can participate in a different Layer 2 network.

A virtual switch configured with a Layer 2 trunk port also supports IRB within a bridge domain. IRB provides simultaneous support for Layer 2 bridging and Layer 3 IP routing on the same interface. Only an interface configured with the **interface-mode (access | trunk)** statement can be associated with a virtual switch. An access interface enables you to accept packets with no VLAN identifier. For more information about configuring trunk and access interfaces, see the Junos OS Network Interfaces Library for Routing Devices.

In addition, you can configure Layer 2 learning and forwarding properties for the virtual switch.

To configure a virtual switch with a Layer 2 trunk interface, include the following statements:

```
[edit]
routing-instances {
    routing-instance-name {
        instance-type virtual-switch;
        interface interface-name;
        bridge-domains {
            bridge-domain-name {
                vlan-id number;
            }
        }
    }
}
```

NOTE: You must configure a bridge domain and VLAN identifier for each VLAN identifier configured for the trunk interface.

Layer 2 trunk ports are used in two distinct types of virtual switch configuration. One method is called service provider style and the other is called enterprise style. The two methods can be confusing because both methods involve configuring interfaces known as trunk interfaces. However, both types of configuration are distinct.

Service provider style and enterprise style each have benefits and drawbacks.

- Service provider style—Offers more control, but requires more care in configuration. Service providers can use all bridging features in any shape or size, but for large bridged designs, customization requirements quickly grow.

- Enterprise style—Offers a single Layer 2 network connected by simple bridges. Easier to use, but more limited in function. Configuration is simple and straightforward and condensed.

> **NOTE**: The terms "service provider style" and "enterprise style" do not imply any limitations based on organization type or size. Any large enterprise may use service-provider-style configurations and a small regional service provider is free to use enterprise style. The differences apply only to the configuration styles.

The easiest way to understand the differences in configuration of the two styles is to compare them using the same interfaces and VLAN IDs.

You can configure multiple bridge domains between the same pair of Ethernet interfaces, for example, **xe-0/0/1** and **xe-0/0/2**. If there are two bridge domains needed, you can configure one bridge domain as **VLAN-100** and the other as **VLAN-200**. However, the configuration requirements are different when implementing service provider style or enterprise style. Here is a look at both styles using the same interfaces and VLANs.

Service provider style involves configuring the values for three main parameters, plus the bridge domains to connect them:

- VLAN tagging—Configure the bridged physical interfaces with **vlan-tagging** to allow them to operate in IEEE 802.1Q mode, also known as a trunk interface.

- Extended VLAN Bridge—Configure the physical interface with the encapsulation statement type **extended-vlan-bridge** to allow bridging on each logical interface.

- Logical unit—Configure a logical unit for each bridged VLAN ID. In most cases, you configure the unit number to be the same as the VLAN ID (that is, unit 100 = VLAN ID 100).

- Bridge domains—Configure the VLAN bridge domains to associate the logical interfaces with the correct VLAN IDs.

Here is the service provider style configuration showing two interfaces used for bridging across two bridge domains, VLAN ID 100 and 200.

```
[edit]
interfaces {
    xe-0/0/1 {
        vlan-tagging;
        encapsulation extended-vlan-bridge;
        unit 100 {
            vlan-id 100;
        }
```

```
        unit 200 {
            vlan-id 200;
        }
    }
    xe-0/0/2 {
        vlan-tagging;
        encapsulation extended-vlan-bridge;
        unit 100 {
            vlan-id 100;
        }
        unit 200 {
            vlan-id 200;
        }
    }
}

bridge-domains {
    VLAN-100 {
        vlan-id 100;
        interface xe-0/0/1.100;
        interface xe-0/0/2.100;
    }
    VLAN-200 {
        vlan-id 200;
        interface xe-0/0/1.200;
        interface xe-0/0/2.200;
    }
}
```

Note that each physical interface has VLAN tagging enabled as well as extended VLAN bridge encapsulation. There are many more parameters that can be configured in service provider style.

In contrast, enterprise style involves configuring the values for three *different* parameters, plus the bridge domains to connect them:

- Family— Configure each bridged physical interface with the family type **bridge**.

- Interface mode—Configure logical interface so that the physical interface operates as either an untagged access port (not shown in this topic) or as an IEEE 801Q **trunk**.

- VLAN ID—Configure each logical interface with a VLAN ID to determine with which bridge the interface belongs.

- Bridge domain—Configure the VLAN bridge domains to associate with the correct VLAN IDs.

Here is the enterprise style configuration showing the same two interfaces used for bridging across the same two bridge domains, VLAN ID 100 and 200.

```
[edit]
interfaces {
    xe-0/0/1 {
        unit 0 {
            family bridge {
                interface-mode trunk;
                vlan-id-list [ 100 200 ];
            }
        }
    }
    xe-0/0/2 {
        unit 0 {
            family bridge {
                interface-mode trunk;
                vlan-id-list [ 100 200 ];
            }
        }
    }
}

bridge-domains {
    VLAN-100 {
        vlan-id 100;
    }
    VLAN-200 {
        vlan-id 200;
    }
}
```

In exchange for simplicity, enterprise style does not allow you to configure VLAN tagging options or encapsulation type. You do not create a separate logical interface for each VLAN ID.

NOTE: You can configure more parameters in each style. These further parameters are beyond the scope of this basic configuration topic.

RELATED DOCUMENTATION

Understanding Layer 2 Learning and Forwarding for Bridge Domains Functioning as Switches with Layer 2 Trunk Ports **| 85**

4
CHAPTER

Configuring Layer 2 Address Learning and Forwarding

Configuring the MAC Table Timeout Interval

The MAC table aging process ensures that a router tracks only active MAC addresses on the network and is able to flush out address that are no longer used.

You can configure the MAC table aging time, the maximum time that an entry can remain in the MAC table before it "ages out," on all bridge domains, one or all VPLS instances, or one or all Ethernet virtual private network (EVPNs) instances on the router. This configuration can influence efficiency of network resource use by affecting the amount of traffic that is flooded to all interfaces because when traffic is received for MAC addresses no longer in the Ethernet routing table, the router floods the traffic to all interfaces.

Depending on how long you want to keep a MAC address in a MAC table before it expires, you can either increase or decrease the aging timer. By default, the timeout interval for all entries in the MAC table is 300 seconds. You can modify the timeout interval for MAC table entries on an MX Series router. You cannot modify the timeout interval for a virtual switch.

NOTE: The timeout interval applies only to dynamically learned MAC addresses. This value does not apply to configured static MAC addresses, which never time out.

The range for *seconds* is from 10 through 1,000,000.

You can modify the timeout interval for a router(at the global level) or on a per-domain basis (bridge domain).

- To modify the timeout interval for the MAC table for a router:

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

- To modify the timeout interval for a bridge domain:

```
[edit bridge-domain bridge-domain-name bridge-options];
user@host# set mac-table-aging-time time;
```

- To modify the timeout for a VPLS or an Ethernet virtual private network (EVPN) instance within a bridge domain:

```
[edit routing-instance routing-instance-name protocols vpls];
[edit routing-instance routing-instance-name protocols evpn];
user@host#  set mac-table-aging-time time;
```

RELATED DOCUMENTATION

Understanding Layer 2 Learning and Forwarding | 3

Enabling MAC Accounting | 58

Limiting the Number of MAC Addresses Learned from Each Logical Interface | 59

Disabling Layer 2 Learning and Forwarding

Enabling MAC Accounting

By default, MAC accounting is disabled. On MX Series routers, you can enable packet accounting either for the router as a whole or for a specific bridge domain. After you enable packet accounting, the Junos OS maintains packet counters for each MAC address learned.

To enable MAC accounting for an MX Series router, include the **global-mac-statistics** statement at the **[edit protocols l2-learning]** hierarchy level:

```
[edit protocols l2-learning]
global-mac-statistics;
```

RELATED DOCUMENTATION

Understanding Layer 2 Learning and Forwarding | 3

Configuring the MAC Table Timeout Interval | 57

Limiting the Number of MAC Addresses Learned from Each Logical Interface | 59

Disabling Layer 2 Learning and Forwarding

Limiting the Number of MAC Addresses Learned from Each Logical Interface

You can configure a limit to the number of MAC addresses learned from the logical interfaces on an MX Series router.

To configure a limit to the total number of MAC addresses that can be learned from the logical interfaces, include the **global-mac-limit** *limit* statement at the **[edit protocols l2-learning]** hierarchy level:

```
[edit]
protocols {
    l2-learning {
        global-mac-limit limit;
    }
}
```

The default limit to the number of MAC addresses that can be learned the router as a whole is 393,215. The range that you can configure for the router as a whole is 20 through 1,048,575.

After the configured MAC address limit is reached, the default is for packets to be forwarded. You can specify that the packets be dropped by including the **packet-action drop** statement at the **[edit protocols l2-learning global-mac-limit]** hierarchy level:

```
[edit]
protocols {
    l2-learning {
        global-mac-limit limit {
            packet-action drop;
        }
    }
}
```

You can also configure a limit to the number of MAC address learned from all the interfaces in a bridge domain or from a specific logical interface only. .

> **NOTE**: The behavior is different for some configurations. For aggregated Ethernet interfaces and label-switched interfaces, the behavior is to learn all the new MAC addresses even when the limit has been reached. The excess addresses are later deleted. The learning limit does not apply

to bridge domain trunk ports, because they have no counters for the individual domains, and those domains might have different MAC learning limits.

NOTE: When static MAC addresses are configured, the learning limit is the configured limit minus the number of static addresses.

NOTE: On MX Series routers running Junos OS Release 8.4 and later, statistics for an aged destination MAC entry are not retained. In addition, source and destination statistics are reset during a MAC move. In previous releases, only source statistics were reset during a MAC move.

Disabling Layer 2 Learning and Forwarding

Disabling dynamic MAC learning on an MX Series router or an EX Series switch prevents all the logical interfaces on the router or switch from learning source and destination MAC addresses.

To disable MAC learning for an MX Series router or an EX Series switch, include the **global-no-mac-learning** statement at the **[edit protocols l2-learning]** hierarchy level:

```
[edit protocols l2-learning]
global-no-mac-learning;
```

For information about how to configure a virtual switch, see "Configuring a Layer 2 Virtual Switch " on page 39.

RELATED DOCUMENTATION

Example: Loop Detection Using the MAC Move Approach

This example shows how to detect loops using the MAC move approach.

Requirements

This example requires the following hardware and software components:

- MX Series 3D Universal Edge Routers

- Junos OS Release 13.2 running on all the devices

Overview

When a MAC address appears on a different physical interface or within a different unit of the same physical interface and if this behavior occurs frequently, it is considered a MAC move.

Configuration errors at the network can force traffic into never ending circular paths. Once there are loops in the Layer 2 network, one of the symptoms is frequent MAC moves, which can be used for rectification of the problem. When it is observed that a source MAC address is moving among the ports, interface is blocked based on the configured `action-priority` for the interface. If the **action-priority** value configured for interfaces is the same, the last interface for the bridge domain on which the MAC address move occurred is blocked.

Configuration

CLI Quick Configuration

To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, and then copy and paste the commands into the CLI at the **[edit]** hierarchy level.

```
set interfaces ge-1/0/4 vlan-tagging
set interfaces ge-1/0/4 encapsulation flexible-ethernet-services
set interfaces ge-1/0/4 unit 10 encapsulation vlan-bridge
set interfaces ge-1/0/4 unit 10 vlan-id 10
set interfaces ge-1/0/4 unit 11 encapsulation vlan-bridge
set interfaces ge-1/0/4 unit 11 vlan-id 11
set interfaces ge-1/0/5 unit 0 family bridge interface-mode trunk
set interfaces ge-1/0/5 unit 0 family bridge vlan-id-list 10-12
set interfaces ge-1/0/6 unit 0 family bridge interface-mode trunk
set interfaces ge-1/0/6 unit 0 family bridge vlan-id-list 10-12
set bridge-domains bd10 vlan-id 10
set bridge-domains bd10 enable-mac-move-action
set bridge-domains bd10 bridge-options interface ge-1/0/5.0 action-priority 1
set bridge-domains bd10 bridge-options interface ge-1/0/6.0 action-priority 5
set bridge-domains bd11 vlan-id 11
set bridge-domains bd11 enable-mac-move-action
set bridge-domains bd12 vlan-id 12
```

In the previous example, all the interfaces, including the trunk interfaces in bd10 and bd11 will be monitored. If there are frequent MAC moves detected within interfaces ge-1/0/5 and ge-1/0/6, interface ge-1/0/5 is blocked. The blocking for trunk interfaces is such that data traffic only for a VLAN (on which the MAC move is detected) will be blocked and not for all the VLANs in the trunk. No action will be taken if a frequent MAC move is observed in **bd12**.

Configuring Loop Detection Using the MAC Move Approach

Step-by-Step Procedure

The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see *Using the CLI Editor in Configuration Mode*.

To configure loop detection using the MAC address move approach:

1. Configure the interfaces.

```
[edit interfaces]
user@host# set ge-1/0/4 vlan-tagging
user@host# set ge-1/0/4 encapsulation flexible-ethernet-services
user@host# set ge-1/0/4 unit 10 encapsulation vlan-bridge
user@host# set ge-1/0/4 unit 10 vlan-id 10
user@host# set ge-1/0/4 unit 11 encapsulation vlan-bridge
user@host# set ge-1/0/4 unit 11 vlan-id 11
user@host# set ge-1/0/5 unit 0 family bridge interface-mode trunk
user@host# set ge-1/0/5 unit 0 family bridge vlan-id-list 10-12
user@host# set ge-1/0/6 unit 0 family bridge interface-mode trunk
user@host# set ge-1/0/6 unit 0 family bridge vlan-id-list 10-12
```

2. Configure the bridge domain parameters.

```
[edit bridge-domains]
user@host# set bd10 vlan-id 10
user@host# set bd10 enable-mac-move-action
user@host# set bd10 bridge-options interface ge-1/0/5.0 action-priority 1
user@host# set bd10 bridge-options interface ge-1/0/6.0 action-priority 5
user@host# set bd11 vlan-id 11
user@host# set bd11 enable-mac-move-action
user@host# set bd12 vlan-id 12
```

Results

From configuration mode, confirm your configuration by entering **show interfaces** and **show bridge-domains** commands. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
user@host# show interfaces
ge-1/0/4 {
    vlan-tagging;
    encapsulation flexible-ethernet-services;
    unit 10 {
        encapsulation vlan-bridge;
        vlan-id 10;
    }
    unit 11 {
        encapsulation vlan-bridge;
        vlan-id 11;
    }
}
ge-1/0/5 {
    unit 0 {
        family bridge {
            interface-mode trunk;
            vlan-id-list 10-12;
        }
    }
}
ge-1/0/6 {
    unit 0 {
        family bridge {
            interface-mode trunk;
            vlan-id-list 10-12;
        }
    }
}
```

```
user@host# show bridge-domains
bridge-domains {
    bd10 {
        vlan-id 10;
```

```
        bridge-options {
            interface ge-1/0/5.0 {
                action-priority 1;
            }
            interface ge-1/0/6.0 {
                action-priority 5
            }
        }
        enable-mac-move-action;
    }
    bd11 {
        vlan-id 11;
        enable-mac-move-action;
    }
    bd12 {
        vlan-id 12;
    }
}
```

If you are done configuring the device, enter **commit** from configuration mode.

Verification

IN THIS SECTION

- Verifying That the Logical Interfaces Blocked Due to MAC Move Are Displayed | **65**

Verifying That the Logical Interfaces Blocked Due to MAC Move Are Displayed

Purpose

Ensure that the current set of logical interfaces blocked due to a MAC move, if any, are displayed.

Action

From operational mode, enter the **show l2-learning mac-move-buffer active** command.

```
user@host# show l2-learning mac-move-buffer active
MAC Address: 00:00:00:00:01:01, VLAN Id: 0
    Time Rec : 2012-06-25 06:23:41   Bridge Domain: bd10
    Prev IFL : ge-1/0/5.0             New IFL: ge-1/0/6.0
    IFBD     : ge-1/0/6.0:10          Blocked  : YES
```

Meaning

As a result of MAC move detection, one of the involved interface bridge domains will be blocked. The output shows that the ge-1/0/6 logical interface is blocked.

RELATED DOCUMENTATION

bridge-domains

Understanding Layer 2 Learning and Forwarding | 3

5
CHAPTER

Configuring Layer 2 Learning and Forwarding for Bridge Domains

Understanding Layer 2 Learning and Forwarding for Bridge Domains

When you configure a bridge domain, Layer 2 address learning is enabled by default. The bridge domain learns unicast media access control (MAC) addresses to avoid flooding the packets to all the ports in the bridge domain. Each bridge domain 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 bridge domain.

> **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 router or for a specific bridge domain or *logical interface*. You can also configure the following Layer 2 learning and forwarding properties:

- 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 bridge domain

- Size of the MAC address table for the bridge domain

- MAC accounting for a bridge domain

RELATED DOCUMENTATION

Understanding Layer 2 Learning and Forwarding | 3

Configuring Static MAC Addresses for Logical Interfaces in a Bridge Domain

You can manually add static MAC entries for the logical interfaces in a bridge domain. You can specify one or more static MAC addresses for each logical interface.

To add a static MAC address for a logical interface in a bridge domain, include the **static-mac** *mac-address* statement at the **[edit bridge-domains** *bridge-domain-name* **bridge-options interface** *interface-name***]** hierarchy level.

```
[edit]
bridge-domains {
    bridge-domain-name {
        domain-type bridge;
        bridge-options {
            interface interface-name {
                static-mac mac-address {
                    <vlan-id number>;
                }
            }
        }
    }
}
```

You can optionally specify a VLAN identifier for the static MAC address by using the **vlan-id** statement. To specify a VLAN identifier for a static MAC address, you must use the **all** option when configuring a VLAN identifier for the bridge domain.

> **NOTE**: If a static MAC address you configure for a logical interface appears on a different logical interface, packets sent to that interface are dropped.

RELATED DOCUMENTATION

Disabling MAC Learning for a Bridge Domain or Logical Interface **| 80**

Configuring the Size of the MAC Address Table for a Bridge Domain **| 71**

Limiting MAC Addresses Learned from an Interface in a Bridge Domain **| 73**

Enabling MAC Accounting for a Bridge Domain **| 79**

Configuring Static MAC Addresses for Logical Interfaces in a Bridge Domain in ACX Series

You can manually add static MAC entries for the logical interfaces in a bridge domain. You can specify one or more static MAC addresses for each logical interface.

To add a static MAC address for a logical interface in a bridge domain, include the **static-mac** *mac-address* statement at the **[edit bridge-domains** *bridge-domain-name* **bridge-options interface** *interface-name***]** hierarchy level.

```
[edit]
bridge-domains {
    bridge-domain-name {
        bridge-options {
            interface interface-name {
                static-mac mac-address {
                }
            }
        }
    }
}
```

RELATED DOCUMENTATION

Layer 2 Bridge Domains on ACX Series Overview **| 3**

Q-in-Q Tunneling on ACX Series Overview **| 91**

Layer 2 Learning and Forwarding for Bridge Domains Overview **| 8**

Configuring a Bridge Domain on ACX Series Routers **| 14**

Configuring Q-in-Q Tunneling on ACX Series **| 92**

Configuring VLAN Identifiers for Bridge Domains in ACX Series **| 35**

Disabling MAC Learning for Bridge Domains on ACX Series **| 81**

Configuring the Size of the MAC Address Table for Bridge Domains in ACX Series **| 72**

Configuring the Size of the MAC Address Table for a Bridge Domain

You can modify the size of the MAC address table for each bridge domain. The default table size is 5120 addresses. The minimum you can configure is 16 addresses, and the maximum is 1,048,575 addresses.

If the MAC table limit is reached, new addresses can no longer be added to the table. Unused MAC addresses are removed from the MAC address table automatically. This frees space in the table, allowing new entries to be added.

To modify the size of the MAC table, include the **mac-table-size** *limit* statement at the **[edit bridge-domains** *bridge-domain-name* **bridge-options]** hierarchy level:

```
[edit]
bridge-domains {
    bridge-domain-name {
        domain-type bridge;
        bridge-options {
            mac-table-size limit {
                packet-action drop;
            }
        }
    }
}
```

RELATED DOCUMENTATION

Disabling MAC Learning for a Bridge Domain or Logical Interface | 80

Configuring Static MAC Addresses for Logical Interfaces in a Bridge Domain | 68

Limiting MAC Addresses Learned from an Interface in a Bridge Domain | 73

Enabling MAC Accounting for a Bridge Domain | 79

Configuring the Size of the MAC Address Table for Bridge Domains in ACX Series

You can modify the size of the MAC address table for each bridge domain. The default table size is 5120 addresses per bridge domain. The minimum you can configure is 1 address, and the maximum is 32,000 addresses.

If the MAC table limit is reached, new addresses can no longer be added to the table.

> **NOTE**: Unused MAC addresses are removed from the MAC address table automatically. This frees space in the table thereby allowing new entries to be added.

To modify the size of the MAC table, include the **mac-table-size** *limit* statement at the **[edit bridge-domains** *bridge-domain-name* **bridge-options]** hierarchy level:

```
[edit]
bridge-domains {
    bridge-domain-name {
        bridge-options {
            mac-table-size limit {
                packet-action drop;
            }
        }
    }
}
```

> **NOTE**: The **mac-table-size** CLI statement is not supported on ACX5048 and ACX5096 routers.

RELATED DOCUMENTATION

Layer 2 Bridge Domains on ACX Series Overview | **3**

Q-in-Q Tunneling on ACX Series Overview | **91**

Layer 2 Learning and Forwarding for Bridge Domains Overview | **8**

Configuring a Bridge Domain on ACX Series Routers | **14**

Limiting MAC Addresses Learned from an Interface in a Bridge Domain

You can configure a limit on the number of MAC addresses learned from a specific bridge domain or from a specific logical interface that belongs to a bridge domain.

To configure a limit for the number of MAC addresses learned from each logical interface in a bridge domain, include the **interface-mac-limit** *limit* statement at the **[edit bridge-domains** *bridge-domain-name* **bridge-options]** hierarchy level:

```
[edit]
bridge-domains {
    bridge-domain-name {
        domain-type bridge;
        interface interface-name;
        bridge-options {
            interface-mac-limit limit;
        }
    }
}
```

To limit the number of MAC addresses learned from a specific logical interface in a bridge domain or an entire bridge domain, include the **interface-mac-limit** *limit* statement at the **[edit bridge-domains** *bridge-domain-name* **bridge-options interface** *interface-name*] or **[edit bridge-domains** *bridge-domain-name* **bridge-options]** hierarchy level:

```
[edit]
bridge-domains {
    bridge-domain-name {
        domain-type bridge;
        interface interface-name;
        bridge-options {
```

```
            interface-mac-limit limit{
                packet-action drop;
            }
            interface interface-name {
                interface-mac-limit limit{
                    packet-action drop;
                }
            }
        }
    }
}
```

For an access port, the default limit on the maximum number of MAC addresses that can be learned on an access port is 1024. Because an access port can be configured in only one bridge domain in a network topology, the default limit is 1024 addresses, which is same as the limit for MAC addresses learned on a logical interface in a bridge domain (configured by including the **interface-mac-limit** *limit* statement at the **[edit bridge-domains** *bridge-domain-name* **bridge-options interface** *interface-name*] or **[edit bridge-domains** *bridge-domain-name* **bridge-options]** hierarchy level.

For a trunk port, the default limit on the maximum number of MAC addresses that can be learned on a trunk port is 8192. Because a trunk port can be associated with multiple bride domains, the default limit is the same as the limit for MAC addresses learned on a logical interface in a virtual switch instance (configured by including the **interface-mac-limit** *limit* statement at the **[edit routing-instances** *routing-instance-name* **switch- options interface** *interface-name*] for a virtual switch instance).

The value you configure for a specific logical interface overrides any value you specify for the entire bridge domain at the **[edit bridge-domains** *bridge-domain-name* **bridge-options]** hierarchy level.

The default limit to the number of MAC addresses that can be learned on a logical interface is 1024. The range that you can configure for a specific logical interface is 1 through 131,071.

After the MAC address limit is reached, the default is for any incoming packets with a new source MAC address to be forwarded. You can specify that the packets be dropped by including the **packet-action drop** statement. To specify that packets be dropped for the entire bridge domain, include the **packet-action drop** statement at the **[edit bridge-domains** *bridge-domain-name* **bridge-options interface-mac-limit** *limit*] hierarchy level:

```
[edit bridge-domains bridge-domain-name bridge-options interface-mac-limit limit]
packet-action drop;
```

To specify that the packets be dropped for a specific logical interface in a bridge domain, include the **packet-action drop** statement at the **[edit bridge-domains *bridge-domain-name* bridge-options interface *interface-name* interface-mac-limit *limit*]** hierarchy level:

```
[edit bridge-domains bridge-domain-name bridge-options interface interface-name
interface-mac-limit limit]
packet-action drop;
```

NOTE: The behavior is different for some configurations. For aggregated Ethernet interfaces and label-switched interfaces, the behavior is to learn all the new MAC addresses even when the limit has been reached. The excess addresses are later deleted. The learning limit does not apply to bridge domain trunk ports, because they have no counters for the individual domains, and those domains might have different MAC learning limits.

NOTE: When static MAC addresses are configured, the learning limit is the configured limit minus the number of static addresses.

NOTE:

You can also configure a limit to the number of MAC addresses learned for an MX Series router.

RELATED DOCUMENTATION

Disabling MAC Learning for a Bridge Domain or Logical Interface | 80

Configuring Static MAC Addresses for Logical Interfaces in a Bridge Domain | 68

Configuring the Size of the MAC Address Table for a Bridge Domain | 71

Enabling MAC Accounting for a Bridge Domain | 79

Configuring MAC Address Limits on a Logical Interface

IN THIS SECTION

You can configure a limit on the number of MAC addresses learned from a specific logical interface. This feature allows the MAC address table space to be distributed among different logical interfaces, thereby avoiding congestion. The MAC address limit can be applied for both VLAN and VPLS routing instances and by default the MAC limit depends on the profile configured. You can limit the number of MAC addresses learned for a bridge domain and a logical interface at the same time.

> **NOTE**: MAC address limiting is supported only on ACX5000 line of routers.

Configuring MAC Address Limit

You can configure the MAC Address limit by using the **set protocols l2-learning global-no-hw-mac-learning** CLI command.

The following configuration example enables limiting MAC address learning on logical interfaces:

```
[edit protocols]
l2-learning {
    global-no-hw-mac-learning;
}
```

Configuring MAC Address Limit for VLANs

To configure a limit for the number of MAC addresses learned on each logical interface in a VLAN, include the **interface-mac-limit** *limit* statement at the [**edit vlans** *vlan-name*] hierarchy level. To limit the MAC addresses learned on a specific logical interface of the VLAN, include the **interface-mac-limit** *limit* statement at the [**edit vlans** *vlan-name* **interface** *interface-name*] hierarchy level. To limit the MAC addresses learned on each of the logical interfaces of the VLAN, include the **interface-mac-limit** *limit* statement at the [**edit vlans** *vlan-name* **switch-options**] hierarchy level.

The following example configures a limit for the number of MAC addresses learned on a logical interface in a VLAN:

```
[edit vlans]
vlan10 {
    interface ge-0/0/3.1;
    interface ge-0/0/1.5;
    switch-options {
        interface-mac-limit {
            10;
        }
    }
    interface ge-0/0/1.5 {
        interface-mac-limit {
            20;
        }
    }
}
```

Configuring MAC Address Limit for VPLS

To configure a limit for the number of MAC addresses learned on each logical interface in a VPLS routing instance, include the **interface-mac-limit** *limit* statement at the [**edit routing-instances** *routing-instance-name* **protocols vpls**] hierarchy level. To limit the MAC addresses learned on a specific logical interface of the VPLS instance, include the **interface-mac-limit** *limit* statement at the [**edit routing-instances** *routing-instance-name* **protocols vpls interface** *interface-name*] hierarchy level.

The following is an example to configure a limit for the number of MAC addresses learned on a logical interface in VPLS routing instance:

```
[edit routing-instance]
v1 {
    protocols {
        vpls {
            interface-mac-limit {
                10;
            }
            interface ge-0/0/1.3 {
                interface-mac-limit {
                    20;
                }
            }
        }
    }
}
```

If you have configured an interface MAC address limit for the logical interface in a bridge domain and a global MAC address limit for a bridge domain, then the interface MAC address limit is considered. The following example shows two MAC address limits configured on the interface ge-0/0/3.5 with the global value as 50 and local value as 30. In this case, the MAC address limit of 30 is considered for the interface ge-0/0/3.5 in the bridge domain.

```
vlan20 {
    interface ge-0/0/1.5;
    interface ge-0/0/3.5;
    switch-options {
        interface-mac-limit {
            50;
        }
        interface ge-0/0/1.5;
        interface ge-0/0/3.5 {
            interface-mac-limit {
                30;
            }
        }
    }
}
```

CLI Commands to Configure MAC Address Limiting

The following CLI commands are used for configuring MAC address limiting:

- **set protocols l2-learning global-no-hw-mac-learning**—Command to change the hardware-based MAC learning to software-based MAC learning mode.

- **set vlans** *vlan-name* **switch-options interface-mac-limit** *limit*—Command to configure the MAC address limit for each logical interface in a VLAN. The limit is applied to all logical interfaces belonging to the VLAN for which a separate interface MAC address limit is not configured.

- **set vlans** *vlan-name* **switch-options interface** *interface-name* **interface-mac-limit** *limit*—Command to configure the interface MAC address limit for a logical interface in a VLAN. The limit is applied to a specific logical interface in the VLAN for which it is configured.

- **set routing-instances** *routing-instance-name* **protocols vpls interface-mac-limit** *limit*—Command to configure the MAC address limit for each logical interface in the VPLS routing instance. This limit is applied to all logical interfaces belonging to the VPLS for which a separate interface MAC address limit is not configured.

- **set routing-instances** *routing-instance-name* **protocols vpls interface** *interface-name* **interface-mac-limit** *limit*—Command to configure the interface MAC address limit for a logical interface in the VPLS. This limit is applied to a specific logical interface in the VPLS for which it is configured.

Enabling MAC Accounting for a Bridge Domain

By default, MAC accounting is disabled. You can enable packet counting for a bridge domain. When you enable packet accounting, the Junos OS maintains packet counters for each MAC address learned on the interfaces in the bridge domain.

To enable MAC accounting for a bridge domain, include the **mac-statistics** statement at the **[edit bridge-domains** *bridge-domain-name* **bridge-options]** hierarchy level:

```
[edit bridge-domains bridge-domain-name bridge-options]
mac-statistics;
```

Disabling MAC Learning for a Bridge Domain or Logical Interface

You can disable MAC learning for all logical interfaces in a specified bridge domain, or for a specific logical interface in a bridge domain. Disabling dynamic MAC learning prevents the specified interfaces from learning source MAC addresses.

To disable MAC learning for all logical interfaces in a bridge domain in a virtual switch, include the **no-mac-learning** statement at the **[edit bridge-domains** *bridge-domain-name* **bridge-options]** hierarchy level:

```
[edit]
bridge-domains {
    bridge-domain-name {
        domain-type bridge;
        interface interface-name;
        bridge-options {
            no-mac-learning;
        }
    }
}
```

To disable MAC learning for a specific logical interface in a bridge domain, include the **no-mac-learning** statement at the **[edit bridge-domains** *bridge-domain-name* **bridge-options interface** *interface-name*] hierarchy level.

```
[edit]
bridge-domains {
    bridge-domain-name {
        domain-type bridge;
        interface interface-name;
        bridge-options {
            interface interface-name {
```

```
            no-mac-learning;
        }
      }
    }
}
```

> **NOTE**: When you disable MAC learning, source MAC addresses are not dynamically learned, and any packets sent to these source addresses are flooded into the bridge domain.

> **NOTE**: When you gather interfaces into a bridge domain, the **no-mac-learn-enable** statement at the **[edit interfaces** *interface-name* **gigether-options ethernet-switch-profile]** hierarchy level is not supported. You must use the **no-mac-learning** statement at the **[edit bridge-domains** *bridge-domain-name* **bridge-options interface** *interface-name*] hierarchy level to disable MAC learning on an interface in a bridge domain.

> **NOTE**: When MAC learning is disabled for a VPLS routing instance, traffic is not load balanced and only one of the equal-cost next hops is used.

RELATED DOCUMENTATION

Disabling MAC Learning for Bridge Domains on ACX Series

You can disable MAC learning in a bridge domain. Disabling dynamic MAC learning prevents the bridge domain from learning source MAC addresses.

To disable MAC learning in a bridge domain, include the **no-mac-learning** statement at the **[edit bridge-domains** *bridge-domain-name* **bridge-options]** hierarchy level:

```
[edit]
bridge-domains {
    bridge-domain-name {
        bridge-options {
            no-mac-learning;
        }
    }
}
```

NOTE: When you disable MAC learning, source MAC addresses are not dynamically learned, and any packets sent to these source addresses are flooded into the bridge domain.

RELATED DOCUMENTATION

Preventing Communication Among Customer Edge Devices as ACX Routers

In a bridge domain, when a frame is received from a CE interface, it is flooded to the other CE interfaces and all of the provider edge (PE) interfaces if the destination MAC address is not learned or if the frame is either broadcast or multicast. If the destination MAC address is learned on another CE device, such a

frame is unicasted to the CE interface on which the MAC address is learned. This might not be desirable if the service provider does not want CE devices to communicate with each other directly.

To prevent CE devices from communicating directly, include the **no-local-switching** statement at the **[edit bridge-domains** *bridge-domain-name*] hierarchy level. Configure the logical interfaces in the bridge domain as core-facing (PE interfaces) by including the **core-facing** statement at the **[edit interfaces** *interface-name***unit** *logical-unit-number* **family** *family*] hierarchy level to specify that the VLAN is physically connected to a core-facing ISP router and ensures that the network does not improperly treat the interface as a client interface. When specified, traffic from one CE interface is not forwarded to another CE interface.

For the **no-local-switching** option , integrated routing and bridging (IRB) configured on a bridge domain with this option enabled is not treated as a designated CE or PE interface. Traffic arriving from a CE or PE interface can navigate towards IRB and traffic that reaches in the input direction to the IRB can pass out of a CE or PE interface. The disabling of local switching achieves the functionality of split-horizon in a bridge domain. If no-local-switching is configured in a bridge domain, , then traffic cannot flow between CE and CE interfaces. This stoppage of trafic flow includes known unicast and multicast, unknown unicast and multicast, and broadcast traffic. However, traffic continues to be transmitted between CE and PE interfaces, and PE and PE interfaces..

RELATED DOCUMENTATION

no-local-switching

6
CHAPTER

Configuring Layer 2 Learning and Forwarding for Bridge Domains Functioning as Switches with Layer 2 Trunk Ports

Understanding Layer 2 Learning and Forwarding for Bridge Domains Functioning as Switches with Layer 2 Trunk Ports

Layer 2 learning is enabled by default. A set of bridge domains, configured to function as a switch with a Layer 2 trunk port, learns unicast media access control (MAC) addresses to avoid flooding packets to the trunk port.

> **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 Layer 2 learning for the entire set of bridge domains as well as modify the following Layer 2 learning and forwarding properties:

- Limit the number of MAC addresses learned from the Layer 2 trunk port associated with the set of bridge domains

- Modify the size of the MAC address table for the set of bridge domains

- Enable MAC accounting for the set of bridge domains

RELATED DOCUMENTATION

Understanding Layer 2 Learning and Forwarding | 3

Limiting MAC Addresses Learned from a Layer 2 Trunk Port

You can configure a limit on the number of MAC addresses learned from a trunk port or from a specific trunk or access interface.

To limit the number of MAC addresses learned through a trunk port associated with a set of bridge domains, include the **interface-mac-limit** *limit* statement at the **[edit switch-options]** hierarchy level:

```
[edit]
switch-options {
    interface-mac-limit limit;
}
```

To limit the number of MAC addresses learned from a specific logical interface configured as an access interface or a trunk interface, include the **interface-mac-limit** *limit* statement at the **[edit switch-options interface *interface-name*]** hierarchy level:

```
[edit]
switch-options {
    interface interface-name {
        interface-mac-limit limit;
    }
}
```

The default value for the number MAC addresses that can be learned from a logical interface is 1024. You can specify a limit either for a set of bridge domains or for a specific logical interface in the range from 1 through 131,071. The value you configure for a specific logical interface overrides any value you specify for the set of bridge domains.

After the specified MAC address limit is reached, the default is for any incoming packets with a new source MAC address to be forwarded. You can specify that the packets be dropped for the entire virtual switch after the MAC address limit is reached by including the **packet-action drop** statement at the **[edit switch-options interface-mac-limit *limit*]** hierarchy level:

```
[edit switch-options interface interface-name interface-mac-limit limit]
packet-action drop;
```

To specify that the packets be dropped from a specific logical interface in a set of bridge domains with a trunk port after the MAC address limit is reached, include the **packet-action drop** statement at the **[edit routing-instances *routing-instance-name* interface *interface-name* interface-mac-limit *limit*]** hierarchy level:

```
[edit routing-instances routing-instance-name interface interface-name interface-
mac-limit limit]
packet-action drop;
```

Configuring the Size of the MAC Address Table for a Set of Bridge Domains

You can modify the size of the MAC address table for a set of bridge domains. The minimum you can configure is 16 addresses, and the maximum is 1,048,575 addresses. The default table size is 5120 addresses.

If the MAC table limit is reached, new addresses can no longer be added to the table. Unused MAC addresses are removed from the MAC address table automatically. This frees space in the table, allowing new entries to be added to the table.

To modify the size of the MAC table for a set of bridge domains, include the **mac-table-size** statement at the **[edit switch-options]** hierarchy level:

```
[edit switch-options]
mac-table-size limit;
```

Enabling MAC Accounting for a Set of Bridge Domains

By default, MAC accounting is disabled. You can enable packet counting for a set of bridge domains. After you enable packet accounting, the Junos OS maintains packet counters for each MAC address learned on the trunk port associated with the set of bridge domains.

To enable MAC accounting for a set of bridge domains, include the **mac-statistics** statement at the **[edit switch-options]** hierarchy level:

```
[edit switch-options]
mac-statistics;
```

RELATED DOCUMENTATION

Disabling MAC Learning for a Set of Bridge Domains

By default, MAC learning is enabled for a set of bridge domains. You can disable MAC learning for a set of bridge domains. Disabling dynamic MAC learning prevents the Layer 2 trunk port associated with the set of bridge domains from learning source and destination MAC addresses. When you disable MAC learning, source MAC addresses are not dynamically learned, and any packets sent to these source addresses are flooded into the switch.

To disable MAC learning for a set of bridge domains, include the **no-mac-learning** statement at the **[edit switch-options]** hierarchy level:

```
[edit switch-options]
no-mac-learning;
```

RELATED DOCUMENTATION

7
CHAPTER

Configuring Q-in-Q Tunneling

Q-in-Q Tunneling on ACX Series Overview

Q-in-Q tunneling allows service providers to create a Layer 2 Ethernet connection between two customer sites. Providers can segregate different customers' VLAN traffic on a link (for example, if the customers use overlapping VLAN IDs) or bundle different customer VLANs into a single service VLAN. Service providers can use Q-in-Q tunneling to isolate customer traffic within a single site or to enable customer traffic flows across geographic locations.

Q-in-Q tunneling adds a service VLAN tag before the customer's 802.1Q VLAN tags. The Juniper Networks Junos operating system implementation of Q-in-Q tunneling supports the IEEE 802.1ad standard.

In Q-in-Q tunneling, as a packet travels from a customer VLAN (C-VLAN) to a service provider's VLAN (S-VLAN), another 802.1Q tag for the appropriate S-VLAN is added before the C-VLAN tag. The C-VLAN tag remains and is transmitted through the network. As the packet exits from the S-VLAN space, in the downstream direction, the S-VLAN 802.1Q tag is removed.

In ACX Series routers, you can configure Q-in-Q tunneling by explicitly configuring an input VLAN map with **push** function on customer facing interfaces in a bridge domain.

You can configure Q-in-Q tunneling on aggregated Ethernet interface by configuring input and output VLAN map.

RELATED DOCUMENTATION

Configuring Q-in-Q Tunneling on ACX Series

To configure Q-in-Q tunneling, you need to configure the logical interface connected to the customer network (user-to-network interfaces (UNI)) and the logical interface connected to the service provider network (network-to-network interface (NNI)).

The following is an example to configure a logical interface connected to a customer network:

```
[edit]
    interface ge-1/0/1 {
        flxible-vlan-tagging;
            encapsulation flexible-ethernet-services;
        unit 0 {
            encapsulation vlan-bridge;
                vlan-id-list 10-20;
                input-vlan-map {
                push;
                vlan-id 500;
                }
                output-vlan-map pop;
            }
    }
```

The following is an example to configure a logical interface connected to a service provider network:

```
[edit]
interface ge-1/0/2; {
    flexible-vlan-tagging;
        encapsulation flexible-ethernet-services;
        unit 0 {
        encapsulation vlan-bridge;
        vlan-id 500;
    }
}
```

The following is an example to configure the bridge domain:

```
[edit]
bridge-domains {
```

```
    qnq-stag-500{
        interface ge-1/0/1;
        interface ge-1/0/2;
    }
}
```

You can configure Q-in-Q tunneling on aggregated Ethernet interface connected to the customer network (UNI) and the logical interface connected to the service provider network (NNI).

RELATED DOCUMENTATION

8

CHAPTER

Configuration Statements for Layer 2 Bridge Domains

action-priority

Syntax

```
action-priority number;
```

Hierarchy Level

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

Description

Configure the action priority value for an interface in a bridge domain on MAC move detection. This priority value is used to determine which interface should be blocked when a throttled MAC move is detected between two interfaces. The priority value can be between 0 and 7 inclusive. A higher value means lower priority. For example, if a MAC address move occurs between two interfaces with the action priority value set to 5 and 6, the interface with value 5 as the action priority value is blocked.

Default

4

Required Privilege Level

routing—To view this statement in the configuration.

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

Release Information

Statement introduced in Junos OS Release 13.2.

bridge-domains

IN THIS SECTION

Syntax

```
bridge-domains {
    bridge-domain-name {
        bridge-options {
            ...bridge-options-configuration...
        }
        domain-type bridge;
        interface interface-name;
        no-irb-layer-2-copy;
        no-local-switching;
        routing-interface routing-interface-name;
        vlan-id (all | none | number);
        vlan-id-list [ vlan-id-numbers ];
        vlan-tags outer number inner number;
        bridge-options {
            interface interface-name {
                mac-pinning
                static-mac mac-address;
            }
            interface-mac-limit limit;
            mac-statistics;
            mac-table-size limit;
            no-mac-learning;
        }
    }
}
```

Hierarchy Level

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

Description

(MX Series routers only) Configure a domain that includes a set of logical ports that share the same flooding or broadcast characteristics in order to perform Layer 2 bridging.

Options

bridge-domain-name—Name of the bridge domain.

> **NOTE**: You cannot use the slash (/) character as part of the bridge domain name. If you do, the configuration will not commit.

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.

Release Information

Statement introduced in Junos OS Release 8.4.

Support for logical systems added in Junos OS Release 9.6.

Support for the **no-irb-layer-2-copy** statement added in Junos OS Release 10.2.

RELATED DOCUMENTATION

Configuring a Bridge Domain

Configuring a Layer 2 Virtual Switch

bridge-options

Syntax

```
bridge-options {
    interface interface-name;
        static-mac static-mac-address;
    }
    global-mac-ip-limit limit;
    interface-mac-ip-limit limit;
    interface-mac-limit limit;
        packet-action drop;
    }
    mac-pinning
    mac-statistics;
    mac-ip-table-size limit;
    mac-table-size limit;
    mac-table-aging-time time;
    no-mac-learning;
}
```

Hierarchy Level

```
[edit 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]
```

Description

(MX Series routers only) Configure Layer 2 learning and forwarding properties for a bridge domain 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.

Release Information

Statement introduced in Junos OS Release 8.4.

Support for logical systems added in Junos OS Release 9.6.

Statment (mac-pinning) introduced in Junos OS Release 16.2.

global-mac-ip limit, **interface-mac-ip-limit**, and **mac-ip-table-size** statements introduced in Junos OS Release 17.4R1.

RELATED DOCUMENTATION

Understanding Layer 2 Learning and Forwarding for Bridge Domains | **68**

disable-action

Syntax

```
disable-action;
```

Hierarchy Level

```
[edit protocols l2-learning global-mac-move]
```

Description

(MX Series routers only) Disable the MAC move action feature globally. MAC move detection configuration does exist, but the action is disabled.

Required Privilege Level

routing—To view this statement in the configuration.

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

Release Information

Statement introduced in Junos OS Release 13.2.

Configuring MAC Move Parameters

domain-type (Bridge Domains)

IN THIS SECTION

Syntax

```
domain-type bridge;
```

ACX Series and MX Series

```
[edit 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]
```

SRX Series

```
[edit bridge-domains bridge-domain-name]
```

Description

Define the domain type **bridge** for a Layer 2 bridge domain.

NOTE: There is only one domain type **bridge**, that can be configured on SRX Series devices. Domain type **bridge** is not enabled by default. An SRX Series device operates in the Layer 2 transparent mode when all physical bridge domains on the device are partitioned into logical bridge domains.

NOTE: Starting with Junos OS Release 15.1X49-D10 and Junos OS Release 17.3R1, the CLI **domain-type** is not available.

NOTE: Starting in Junos OS Release 15.1X49-D10 and Junos OS Release 17.3R1, the hierarchy **[edit bridge-domains bridge-domain-name]** is renamed to **[edit vlans vlan-name]**. For detailed information about the modified hierarchies, see *Enhanced Layer 2 CLI Configuration Statement and Command Changes for Security Devices*.

Required Privilege Level

routing—To view this statement in the configuration.

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

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.

RELATED DOCUMENTATION

Ethernet Switching and Layer 2 Transparent Mode Overview

Configuring a Bridge Domain | **11**

Configuring a Layer 2 Virtual Switch | **39**

enable-mac-move-action

Syntax

```
enable-mac-move-action;
```

Hierarchy Level

```
[edit bridge-domains bridge-domain-name]
```

Description

Enable the MAC move action feature at the bridge domain level. This statement blocks the logical interface for the bridge domain when a MAC move is detected on the interface.

Required Privilege Level

routing—To view this statement in the configuration.

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

Release Information

Statement introduced in Junos OS Release 13.2.

Statement supported in Junos OS Release 14.2 for MX104 Router.

RELATED DOCUMENTATION

Configuring a Bridge Domain | **11**

interface

Syntax

```
interface interface-name;
```

Hierarchy Level

```
[edit 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]
```

Description

(MX Series routers and EX Series switches only) Specify the logical interfaces to include in the bridge domain, VLAN, VPLS instance, or virtual switch.

Options

interface-name—Name of a logical interface. For more information about how to configure logical interfaces, see the Junos OS Network Interfaces Library for Routing Devices.

Required Privilege Level

routing—To view this statement in the configuration.

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

Release Information

Statement introduced in Junos OS Release 8.4.

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.

RELATED DOCUMENTATION

Configuring a Layer 2 Virtual Switch on an EX Series Switch

Tunnel Services Overview

Tunnel Interface Configuration on MX Series Routers Overview

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]
```

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.

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.

21

RELATED DOCUMENTATION

Understanding Layer 2 Learning and Forwarding for Bridge Domains

Layer 2 Learning and Forwarding for VLANs Overview

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

mac-statistics

IN THIS SECTION

- Syntax | **112**
- Hierarchy Level | **112**
- Description | **113**
- Default | **113**
- Required Privilege Level | **113**
- Release Information | **113**

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]
```

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.

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.

RELATED DOCUMENTATION

Understanding Layer 2 Learning and Forwarding for Bridge Domains

Layer 2 Learning and Forwarding for VLANs Overview

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

IN THIS SECTION

- Syntax | **115**
- Hierarchy Level | **115**
- Description | **115**
- Options | **116**
- Required Privilege Level | **116**
- Release Information | **116**

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]
```

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.

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 addresses

  There 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.

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.

RELATED DOCUMENTATION

Understanding Layer 2 Learning and Forwarding for Bridge Domains **| 68**

Layer 2 Learning and Forwarding for VLANs Overview

Understanding Layer 2 Learning and Forwarding for Bridge Domains Functioning as Switches with Layer 2 Trunk Ports **| 85**

Layer 2 Learning and Forwarding for VLANs Acting as a Switch for a Layer 2 Trunk Port

mac-table-aging-time

IN THIS SECTION

- Syntax | **117**
- Hierarchy Level | **118**
- Description | **118**
- Options | **118**
- Required Privilege Level | **119**
- Release Information | **119**

Syntax

```
mac-table-aging-time time;
```

Hierarchy Level

```
[edit logical-systems logical-system-name routing-instances routing-instance-
name protocols vpls],
[edit routing-instances routing-instance-name protocols vpls]
[edit bridge-options],
[edit routing-instances routing-instance-name protocols evpn]
```

> **NOTE**: For MX Series routers, the configuration statement is supported at the **[bridge-options]**, **[protocols vpls]**, and **[protocols evpn]** hierarchy levels only.

Description

Modify the timeout interval for the MAC table.

For MX Series routers, you can use the **global-mac-table-aging-time** statement at the **[edit protocols l2-learning]** hierarchy level to configure the timeout interval at the global level or use the **mac-table-aging-time** to configure the timeout interval for a bridge domain or for a specific VPLS or EVPN instance. If multiple timeout interval values are configured on a router, the router determines the timeout interval value in the following order of priority:

- Timeout interval configured at the VPLS or EVPN instance

- Timeout interval configured for the bridge domain

- Global timeout interval configured on the router

> **NOTE**: For MX Series routers, the timeout interval configuration feature is supported on routers with MPCs only.

Options

time—Specify the number of seconds to wait between MAC table clearings.

- **Range:** 10 through 1,000,000 seconds

- **Default:** 300 seconds

Required Privilege Level

routing—To view this statement in the configuration.

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

Release Information

Statement introduced in Junos OS Release 7.4.

no-irb-layer-2-copy

Syntax

```
no-irb-layer-2-copy;
```

Hierarchy Level

```
[edit bridge-domains],
[edit logical-routers logical-router-name bridge-domains],
[edit routing-instances routing-instance-name bridge-domains]
```

Description

If you include this statement when using port mirroring with Integrated Routing and Bridging (IRB), then the packet is mirrored as a Layer 3 packet. By default, the packet is mirrored as a Layer 2 packet. This statement is also supported if a routing instance is set to type VPLS.

Usage Guidelines

See "Configuring a Bridge Domain" on page 11

Required Privilege Level

view-level—To view this statement in the configuration.

control-level—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.2.

no-mac-learning

IN THIS SECTION

Syntax

```
no-mac-learning;
```

QFX Series and EX4600

For QFX Series and EX4600 platforms without ELS:

```
[edit ethernet-switching-options interfaces interface-name]
```

For QFX Series and EX4600 platforms with ELS:

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

QFX Series per VLAN

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

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

EX Series Q-in-Q Interfaces

```
[edit ethernet-switching-options interfaces interface-name]
```

EX Series and SRX Series Q-inQ Vlans

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

ACX Series, MX Series, EX Series with ELS support, M Series, T Series

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

Description

For QFX Series, EX Series switches and SRX Series devices, disables MAC address learning for the specified VLAN.

For QFX Series and EX4600, disable MAC address learning for the specified interface. Disabling MAC address learning on an interface disables learning for all the VLANs of which that interface is a member.

For EX Series switches' Q-in-Q interfaces, disables MAC address learning for the specified interface. Disabling MAC address learning on an interface disables learning for all the VLANs of which that interface is a member.

For MX Series routers and EX Series switches with ELS support, disables MAC learning for a virtual switch, for a bridge domain or VLAN, for a specific logical interface in a bridge domain or VLAN, or for a set of bridge domains or VLANs associated with a Layer 2 trunk port. On platforms that support EVPNs, you can disable MAC learning on an EVPN.

NOTE: When MAC learning is disabled for a VPLS routing instance, traffic is not load-balanced and only one of the equal-cost next hops is used.

Default

MAC learning is enabled.

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.

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 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 12.3 R2 for EX Series switches.

Support for EVPNs added in Junos OS Release 13.2 for MX 3D Series routers.

Hierarchy levels **[edit switch-options interface *interface-name*]** and **[edit vlans *vlan-name* switch-options]** introduced in Junos OS Release 13.2X50-D10 for EX Series switches.

RELATED DOCUMENTATION

Configuring EVPN Routing Instances

Configuring EVPN Routing Instances on EX9200 Switches

Understanding Layer 2 Learning and Forwarding for Bridge Domains

packet-action

IN THIS SECTION

- Syntax | **125**
- Hierarchy Level | **125**
- Description | **127**
- Default | **127**
- Options | **127**
- Required Privilege Level | **128**
- Release Information | **128**

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]
```

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.

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 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.

RELATED DOCUMENTATION

Configuring EVPN Routing Instances

Configuring EVPN Routing Instances on EX9200 Switches

Configuring MAC Limiting (ELS)

Configuring Persistent MAC Learning (ELS)

Understanding Layer 2 Learning and Forwarding for Bridge Domains

Layer 2 Learning and Forwarding for VLANs Overview

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

Layer 2 Learning and Forwarding for VLANs Acting as a Switch for a Layer 2 Trunk Port

reopen-time

IN THIS SECTION

Syntax

```
reopen-time seconds;
```

Hierarchy Level

```
[edit protocols l2-learning global-mac-move]
```

Description

(MX Series routers only) Configure the value for the reopen timer.

Default

180 seconds

Options

seconds—Time duration after which the port is unblocked.

Required Privilege Level

routing—To view this statement in the configuration.

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

Release Information

Statement introduced in Junos OS Release 13.2.

RELATED DOCUMENTATION

| *Configuring MAC Move Parameters*

routing-interface

Syntax

```
routing-interface routing-interface-name;
```

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]
```

Description

(MX Series routers only) Specify a routing interface to include in a bridge domain or a VPLS routing instance.

When you configure **routing-interface irb.*x***, the VPLS connection comes up, even if no customer edge (CE) interfaces are configured. This works with one site configured, but not when multiple sites (multisite) are configured.

Options

routing-interface-name—Name of the routing interface to include in the bridge domain or the VPLS routing instance. The format of the routing interface name is **irb.*x***, where *x* is the unit number of the routing interface you configured at the **[edit interfaces irb]** hierarchy level. For more information about how to configure a routing interface, see the Junos OS Network Interfaces Library for Routing Devices.

> **NOTE**: You can specify only one routing interface for each bridge domain or VPLS instance.

Required Privilege Level

routing—To view this statement in the configuration.

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

Release Information

Statement introduced in Junos OS Release 8.4.

Support for logical systems added in Junos OS Release 9.6.

service-id

IN THIS SECTION

Syntax

```
service-id number;
```

Hierarchy Level

```
[edit bridge-domains bridge-domain-name]
```

Description

Specify a service identifier to include in the packets sent to and from the multichassis link aggregation (MC-LAG) bridge domain when the VLAN identifier is set to **none**. This configuration facilitates media access control (MAC) and Address Resolution Protocol (ARP) synchronization among MC-LAG peers.

> **NOTE**: The VLAN identifier none is supported only for IPv4 traffic.

Options

number—A valid service identifier. You must configure the same service identifier within the bridge domains of MC-LAG peers.

Required Privilege Level

routing—To view this statement in the configuration.

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

Release Information

Statement introduced in Junos OS Release 13.2

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]
```

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.

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.

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.

RELATED DOCUMENTATION

Configuring EVPN Routing Instances

Understanding Layer 2 Learning and Forwarding for Bridge Domains

Layer 2 Learning and Forwarding for VLANs Overview

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

vlan-id-list

IN THIS SECTION

- Syntax | **137**
- Hierarchy Level | **138**
- Description | **138**
- Options | **138**
- Required Privilege Level | **139**
- Release Information | **139**

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]
```

Description

Specify a VLAN identifier list to use for a bridge domain or VLAN in trunk mode. VLAN identifier list can be used on C-VLAN interfaces in Q–in–Q tunneling for EX and QFX Series switches.

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, if VLAN identifier list (**vlan-id-list**) is used for Q-in-Q tunnelling, 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.

Required Privilege Level

routing—To view this statement in the configuration.

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

Release Information

Statement introduced in Junos OS Release 9.4.

Support for logical systems added in Junos OS Release 9.6.

RELATED DOCUMENTATION

vlan-tags

IN THIS SECTION

-

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]
```

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.

Release Information

Statement introduced in Junos OS Release 8.4.

Support for logical systems added in Junos OS Release 9.6.

RELATED DOCUMENTATION

Configuring a Bridge Domain | **11**

Configuring a VLAN

Configuring VLAN Identifiers for Bridge Domains and VPLS Routing Instances | **28**

Configuring VLAN Identifiers for VLANs and VPLS Routing Instances

Configuring a Layer 2 Virtual Switch | **39**

Configuring a Layer 2 Virtual Switch on an EX Series Switch

9
CHAPTER

Configuration Statements for Layer 2 Bridge Domains Functioning as Switches with Layer 2 Trunk Ports

switch-options

Syntax

```
switch-options {
    interface interface-name {
        mac-pinning;
        interface-mac-limit limit;
    }
    interface-mac-limit limit {
        packet-action drop;
    }
    mac-statistics;
    mac-table-size limit {
        packet-action drop;
    }
    no-mac-learning;
    route-distinguisher (as-number:id | ip-address:id);
     service-id number; number;
    vrf-target {
        community;
        auto
        import community-name;
        export community-name;
    }
    vrf-import[ policy-names ];
```

```
    vrf-export[ policy-names ];
}
```

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]
```

Description

Configure Layer 2 learning and forwarding properties for a set of bridge domains.

Options

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.

Release Information

Statement introduced in Junos OS Release 9.2.

Support for logical systems added in Junos OS Release 9.6.

14.1x53-D10

interface-mac-limit

IN THIS SECTION

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]
```

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.

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.

RELATED DOCUMENTATION

Understanding Layer 2 Learning and Forwarding for Bridge Domains

Layer 2 Learning and Forwarding for VLANs Overview

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

mac-statistics

IN THIS SECTION

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]
```

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.

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.

RELATED DOCUMENTATION

Understanding Layer 2 Learning and Forwarding for Bridge Domains

Layer 2 Learning and Forwarding for VLANs Overview

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]
```

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.
>
> 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 addresses

  There 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.

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.

RELATED DOCUMENTATION

Understanding Layer 2 Learning and Forwarding for Bridge Domains | **68**

Layer 2 Learning and Forwarding for VLANs Overview

Understanding Layer 2 Learning and Forwarding for Bridge Domains Functioning as Switches with Layer 2 Trunk Ports | **85**

Layer 2 Learning and Forwarding for VLANs Acting as a Switch for a Layer 2 Trunk Port

no-mac-learning

IN THIS SECTION

- Syntax | **154**
- QFX Series and EX4600 | **154**
- QFX Series per VLAN | **154**

Syntax

```
no-mac-learning;
```

QFX Series and EX4600

For QFX Series and EX4600 platforms without ELS:

```
[edit ethernet-switching-options interfaces interface-name]
```

For QFX Series and EX4600 platforms with ELS:

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

QFX Series per VLAN

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

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

EX Series Q-in-Q Interfaces

```
[edit ethernet-switching-options interfaces interface-name]
```

EX Series and SRX Series Q-inQ Vlans

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

ACX Series, MX Series, EX Series with ELS support, M Series, T Series

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

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

Description

For QFX Series, EX Series switches and SRX Series devices, disables MAC address learning for the specified VLAN.

For QFX Series and EX4600, disable MAC address learning for the specified interface. Disabling MAC address learning on an interface disables learning for all the VLANs of which that interface is a member.

For EX Series switches' Q-in-Q interfaces, disables MAC address learning for the specified interface. Disabling MAC address learning on an interface disables learning for all the VLANs of which that interface is a member.

For MX Series routers and EX Series switches with ELS support, disables MAC learning for a virtual switch, for a bridge domain or VLAN, for a specific logical interface in a bridge domain or VLAN, or for a set of bridge domains or VLANs associated with a Layer 2 trunk port. On platforms that support EVPNs, you can disable MAC learning on an EVPN.

> **NOTE**: When MAC learning is disabled for a VPLS routing instance, traffic is not load-balanced and only one of the equal-cost next hops is used.

Default

MAC learning is enabled.

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.

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 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 12.3 R2 for EX Series switches.

Support for EVPNs added in Junos OS Release 13.2 for MX 3D Series routers.

Hierarchy levels **[edit switch-options interface** *interface-name*] and **[edit vlans** *vlan-name* **switch-options]** introduced in Junos OS Release 13.2X50-D10 for EX Series switches.

RELATED DOCUMENTATION

Configuring EVPN Routing Instances

Configuring EVPN Routing Instances on EX9200 Switches

Understanding Layer 2 Learning and Forwarding for Bridge Domains

Layer 2 Learning and Forwarding for VLANs Overview

Understanding Layer 2 Learning and Forwarding for Bridge Domains Functioning as Switches with Layer 2 Trunk Ports

Understanding Bridging and VLANs on Switches

Understanding Q-in-Q Tunneling and VLAN Translation

Understanding Q-in-Q Tunneling and VLAN Translation

Configuring Q-in-Q Tunneling on EX Series Switches

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]
```

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.

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 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.

RELATED DOCUMENTATION

Configuring EVPN Routing Instances

Configuring EVPN Routing Instances on EX9200 Switches

Configuring MAC Limiting (ELS)

Configuring Persistent MAC Learning (ELS)

Understanding Layer 2 Learning and Forwarding for Bridge Domains

Layer 2 Learning and Forwarding for VLANs Overview

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

Layer 2 Learning and Forwarding for VLANs Acting as a Switch for a Layer 2 Trunk Port

10
CHAPTER

Configuration Statements for Layer 2 Address Learning and Forwarding

l2-learning

Syntax (MX Series, QFX Series, EX Series)

```
l2-learning {
    global-le-bridge-domain-aging-time;
    global-mac-ip-limit number;
    global-mac-ip-table-aging-time seconds;
    global-mac-limit limit;
    global-mac-statistics;
    global-mac-table-aging-time seconds;
    global-no-mac-learning;
    global-mac-move;
}
```

Syntax (SRX Series)

```
l2-learning {
    global-mac-limit limit {
    packet-action-drop
    }
```

```
    global-mac-table-aging-time seconds;
    global-mode (switching | transparent-bridge) ;
    global-no-mac-learning;
}
```

Hierarchy Level

```
[edit protocols]
```

Description

Configure Layer 2 address learning and forwarding properties globally.

The remaining statements are explained separately. See CLI Explorer.

Options

global-le-bridge-domain-aging-time	Specify the aging time of LE bridge-domain. The MAC address is learnt after next hop(NH) and bridge-domain(BD), also called NHBD. This aging time delays the deletion of NHBD. Configuring lesser time, in seconds, results in faster deletion of NHBD.
	• **Range:** 120 to 1000000 seconds

Required Privilege Level

routing—To view this statement in the configuration.

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

Release Information

Statement introduced in Junos OS Release 8.4.

Statement modified in Junos OS Release 9.5. Support for global mode added in Junos OS Release 15.1X49-D40.

global-le-bridge-domain-aging-time option introduced in Junos OS Release 14.2R5 for the MX Series.

global-mac-ip-limit and **global-mac-ip-table-aging-time** options introduced in Junos OS Release 17.4R1 for MX Series routers and EX9200 switches.

RELATED DOCUMENTATION

Understanding Layer 2 Learning and Forwarding

global-mac-table-aging-time

global-mac-limit (Protocols)

global-no-mac-learning

global-mode (Protocols)

global-mac-limit

IN THIS SECTION

- Syntax | **167**
- Hierarchy Level | **167**
- Description | **167**
- Default | **167**
- Options | **167**
- Required Privilege Level | **168**
- Release Information | **168**

Syntax

```
global-mac-limit limit {
    packet-action drop;
}
```

Hierarchy Level

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

Description

(MX Series routers and EX Series switches only) Limit the number of media access control (MAC) addresses learned from the logical interfaces on the router or switch.

Default

(MX Series) 393,215 MAC addresses

(EX9200) 524,287 MAC addresses

Options

limit—Number of MAC addresses that can be learned systemwide.

- **Range:** 20 through 1,048,575

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.

Release Information

Statement introduced in Junos OS Release 8.4.

Support for logical systems added in Junos OS Release 9.6.

RELATED DOCUMENTATION

Limiting the Number of MAC Addresses Learned from Each Logical Interface | 59

global-mac-move

IN THIS SECTION

Syntax

```
global-mac-move {
    cooloff-time seconds;
    disable-action;
    exclusive-mac virtual-mac-mac-address/mask;
    interface-recovery-time seconds;
    notification-time seconds;
    reopen-time seconds;
    statistical-approach-wait-time seconds;
    threshold-count count;
    threshold-time seconds;
    virtual-mac mac-address /mask;
}
```

Hierarchy Level

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

Description

Set parameters for media access control (MAC) address move reporting.

Default

By default, MAC moves notify every second, with a threshold time of 1 second and a threshold count of 50.

Required Privilege Level

view-level—To view this statement in the configuration.

control-level—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.4.

Support for logical systems added in Junos OS Release 9.6.

Support for disable-action and reopen-time added in Junos OS Release 13.2.

Support for exclusive-mac added in Junos OS Release 14.1X53-D45.

Statements **cooloff-time**, **interface-recovery-time**, **statistical-approach-wait-time**, and **virtual-mac** moved from vpls-mac-move to global-mac-move hierarchy level in Junos OS Release 17.4.

RELATED DOCUMENTATION

Configuring MAC Move Parameters

MAC Moves Loop Prevention in VPLS Network Overview

Example: Configuring Loop Prevention in VPLS Network Due to MAC Moves

virtual-mac

global-mac-statistics

IN THIS SECTION

- Syntax | **171**
- Hierarchy Level | **171**
- Description | **171**
- Default | **171**
- Options | **171**
- Required Privilege Level | **171**
- Release Information | **172**

Syntax

```
global-mac-statistics;
```

Hierarchy Level

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

Description

(MX Series routers and EX Series switches only) Enable MAC accounting for the entire router or switch.

Default

disabled

Options

Required Privilege Level

routing—To view this statement in the configuration.

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

Release Information

Statement introduced in Junos OS Release 9.2.

Support for logical systems added in Junos OS Release 9.6.

global-mac-table-aging-time

IN THIS SECTION

Syntax

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

Hierarchy Level

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

Description

Configure the timeout interval for entries in the MAC table.

> **NOTE**: The **global-mac-table-aging-time** statement appears in the Junos OS CLI for devices that support the Enhanced Layer 2 Software (ELS) configuration style. If your device runs software that does *not* support ELS, use the *mac-table-aging-time* statement, which appears in the **[edit ethernet-switching-options]** and the **[edit vlans]** hierarchies. For ELS details, see *Using the Enhanced Layer 2 Software CLI*.

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.

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.

RELATED DOCUMENTATION

Configuring the MAC Table Timeout Interval | **57**

Configuring MAC Table Aging on Switches

Example: Configuring VLANs on Security Devices

global-no-mac-learning

IN THIS SECTION

- Syntax | **174**
- Hierarchy Level | **175**
- Description | **175**
- Default | **175**
- Required Privilege Level | **175**
- Release Information | **175**

Syntax

```
global-no-mac-learning;
```

Hierarchy Level

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

Description

Disable MAC learning on the entire device.

Default

MAC learning is enabled.

Required Privilege Level

routing—To view this statement in the configuration.

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

Release Information

Statement introduced in Junos OS Release 9.2.

Statement modified for SRX Series in Junos OS Release 9.5.

Support for logical systems added in Junos OS Release 9.6.

RELATED DOCUMENTATION

Disabling Layer 2 Learning and Forwarding

Understanding Q-in-Q Tunneling and VLAN Translation

Example: Configuring VLANs on Security Devices

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]
```

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.

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.

Understanding Layer 2 Learning and Forwarding for Bridge Domains

Layer 2 Learning and Forwarding for VLANs Overview

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-time

IN THIS SECTION

- Syntax | **179**
- Hierarchy Level | **179**
- Description | **180**
- Default | **180**
- Options | **180**
- Required Privilege Level | **180**
- Release Information | **180**

Syntax

```
notification-time seconds;
```

Hierarchy Level

```
[edit protocols l2-learning global-mac-move]
```

Description

(MX Series routers only) Configure the notification time value for MAC move reports that a MAC address moves before counting against the threshold values.

Default

1 second

Options

seconds—Time elapsed before MAC move reports are generated.

Required Privilege Level

routing—To view this statement in the configuration.

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

Release Information

Statement introduced in Junos OS Release 9.4.

Support for logical systems added in Junos OS Release 9.6.

RELATED DOCUMENTATION

Configuring MAC Move Parameters

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]
```

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.

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 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.

RELATED DOCUMENTATION

threshold-count

IN THIS SECTION

Syntax

```
threshold-count count;
```

Hierarchy Level

```
[edit protocols l2-learning global-mac-move]
```

Description

(MX Series routers only) Configure the threshold count value for MAC move reports.

Default

50

Options

count—Number of MAC moves needed in the notification time to generate a MAC move report.

Required Privilege Level

routing—To view this statement in the configuration.

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

Release Information

Statement introduced in Junos OS Release 9.4.

Support for logical systems added in Junos OS Release 9.6.

Configuring MAC Move Parameters

threshold-time

IN THIS SECTION

Syntax

```
threshold-time seconds;
```

Hierarchy Level

```
[edit protocols l2-learning global-mac-move]
```

Description

(MX Series routers only) Configure the threshold time value for MAC move reports when the MAC address moves at least a specified number of times (threshold count) in the configured interval.

Default

1 second

Options

seconds—Timer threshold before MAC move reports are generated.

Required Privilege Level

routing—To view this statement in the configuration.

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

Release Information

Statement introduced in Junos OS Release 9.4.

Support for logical systems added in Junos OS Release 9.6.

RELATED DOCUMENTATION

Configuring MAC Move Parameters

11

CHAPTER

Operational Mode Commands for Layer 2 Bridge Domains

clear bridge mac-table

Syntax

```
clear bridge mac-table
<bridge-domain (all | bridge-domain-name)>
<instance instance-name>
<interface interface-name>
<learning-vlan id (all-vlan | learning-vlan-id)>
<mac-address>
```

Description

(MX Series routers only) Clear learned Layer 2 address information from the media access control (MAC) address table.

Options

none   Clear all learned Layer 2 address information from the MAC address table.

bridge-domain (all | *bridge-domain-name*)	(Optional) Clear learned Layer 2 MAC addresses for all bridging domains or for the specified bridging domain.
instance *instance-name*	(Optional) Clear learned Layer 2 MAC addresses for the specified routing instance.
interface *interface-name*	(Optional) Clear learned Layer 2 MAC addresses for the specified interface.
learning-vlan-id (all-vlan | *learning-vlan-id*)	(Optional) Clears learned Layer 2 MAC addresses for all VLANs or for the specified VLAN.
mac-address	(Optional) Clear the specified learned Layer 2 address from the MAC address table.

Required Privilege Level

clear

Output Fields

When you enter this command, you are provided feedback on the status of your request.

Sample Output

clear bridge mac-table

```
user@host>  clear bridge mac-table
```

Release Information

Command introduced in Junos OS Release 8.4.

clear interfaces mac-database

Syntax

```
clear interfaces mac-database (interface-name | aex) <mac-address mac-address>
```

Description

Clear learned media access control (MAC) addresses from the hardware and MAC database for Gigabit Ethernet IQ2 interfaces or aggregated Ethernet interfaces. Static MAC addresses configured by the operator are not cleared.

Options

interface-name   Name of a physical or logical interface. When you clear a physical interface, all learned MAC addresses on all the logical interfaces under the physical interface are cleared.

ae*x*   Name of aggregated Ethernet interface.

mac-address *mac-address*	(Optional) Clear only the specified MAC address.

Required Privilege Level

view

Output Fields

This command produces no output.

Sample Output

clear interfaces mac-database

```
user@host> clear interfaces mac-database ge-0/0/0.0
```

Release Information

Command introduced in Junos OS Release 8.3.

Support for statement with the **ae**x option introduced in Junos OS Release 15.1.

clear interfaces mac-database statistics

IN THIS SECTION

- Syntax | **194**

Syntax

```
clear interfaces mac-database statistics (interface-name |all)
```

Description

Clear statistics that are collected for every MAC address, including policer statistics, on a physical or logical interface or all interfaces.

Options

(*interface-name* | all)   Clear MAC database statistics for the specified physical or logical gigabit or 10-Gigabit Ethernet interface. Specify **all** to clear the MAC database statistics for all interfaces.

Required Privilege Level

view

Output Fields

This command produces no output.

Sample Output

clear interfaces mac-database statistics (Gigabit Ethernet)

```
user@host> clear interfaces mac-database statistics ge-0/1/0
```

Release Information

Command introduced in Junos OS Release 8.3.

show bridge domain

Syntax

```
show bridge domain
<brief | detail | extensive>
<bridge-domain (all | domain-name)>
<instance instance-name>
<operational>
```

Description

(MX Series routers only) Display bridge domain information.

Options

none	Display information for all bridge domains.
brief | detail | extensive	(Optional) Display the specified level of output.
bridge-domain (all | *domain-name*)	(Optional) Display information about all bridge domains or the specified bridge domain.
instance *instance-name*	(Optional) Display information for the specified routing instance.
operational	(Optional) Display information for the operational routing instances.

Required Privilege Level

view

Sample Output

show bridge domain

```
user@host> show bridge domain
Instance                      Bridging Domain       Type
          Primary Table                                       Active
vs1                           vlan100               bridge
          bridge.0                                            2
vs1                           vlan200               bridge
          bridge.0                                            0
```

show bridge domain brief

```
user@host> show bridge domain brief
Instance                      Bridging Domain       Type
          Primary Table                                       Active
vs1                           vlan100               bridge
          bridge.0                                            2
vs1                           vlan200               bridge
          bridge.0                                            0
```

show bridge domain detail

```
user@host> show bridge domain detail
Routing Instance:vs1
  Bridging Domain:vlan100
  Router ID: 0.0.0.0
  Type: bridge              State: Active
  Interfaces:
    ge-11/0/3.0
    ge-11/1/4.100
    ge-11/1/1.100
    ge-11/1/0.100
    xe-10/2/0.100
    xe-10/0/0.100
  Tables:
     bridge.0                : 2 macs (2 active)
```

```
Routing Instance:vs1
  Bridging Domain:vlan200
  Router ID: 0.0.0.0
  Type: bridge              State: Active
  Interfaces:
    ge-11/1/0.200
    ge-11/1/1.200
    ge-11/1/4.200
    xe-10/0/0.200
    xe-10/2/0.200
  Tables:
     bridge.0                : 0 macs (0 active)
```

Release Information

Command introduced in Junos OS Release 8.4.

show bridge flood

IN THIS SECTION

- Syntax | **199**
- Description | **199**
- Options | **199**
- Required Privilege Level | **200**
- Output Fields | **200**
- Sample Output | **200**
- Release Information | **208**

Syntax

```
show bridge flood
<brief | detail | extensive>
<bridge-domain domain-name>
<event-queue>
<instance instance-name>
<route (all-ce-flood | all ve-flood | alt-root-flood | bd-flood | mlp-flood | re-
flood)>
```

Description

(MX Series routers only) Display bridging flooding information.

Options

none	Display all bridging flooding information for all bridging domains.
brief | detail | extensive	(Optional) Display the specified level of output.
bridge-domain *domain-name*	(Optional) Display bridging flooding information for the specified bridge domain.
event-queue	(Optional) Display the queue of pending bridge flood events.
instance *instance-name*	(Optional) Display bridging flooding information for the specified routing instance.
route (all-ce-flood | all ve-flood | alt-root-flood | bd-flood | mlp-flood | re-flood)	(Optional) Display the following: • **all-ce-flood**—Display the route for flooding traffic to all customer edge routers if **no-local-switching** is enabled. • **all-ve-flood**—Display the route for flooding traffic to all VPLS edge routers if **no-local-switching** is enabled.

- **alt-root-flood**—Display the Spanning Tree Protocol (STP) alt-root flooding route used for the interface.

- **bd-flood**—Display the route for flooding traffic of a bridge domain if **no-local-switching** is not enabled.

- **mlp-flood**—Display the route for flooding traffic to MAC learning chips.

- **re-flood**—Display the route for Routing Engine flooding to all interfaces.

Required Privilege Level

view

Output Fields

to be provided

Sample Output

show bridge flood

```
user@host> show bridge flood
Name: __juniper_private1__
CEs: 0
VEs: 0
Flood Routes:
  Prefix    Type           Owner               NhType         NhIndex
  0x36/16   MLP_FLOOD      __vs1+vlan100__     flood          426
  0x3a/16   MLP_FLOOD      __vs1+vlan200__     flood          428
Name: vs1::vlan100
CEs: 6
VEs: 0
Flood Routes:
  Prefix    Type           Owner               NhType         NhIndex
```

```
 0x35/16   ALL_FLOOD    __vs1+vlan100__        flood      425
 0x35/16   RE_FLOOD     __vs1+vlan100__        flood      425
 0x3780/17 ALT_ROOT_RT  ge-11/0/3.0            flood      425
 0x3b80/17 ALT_ROOT_RT  ge-11/1/4.100          flood      425
 0x3c80/17 ALT_ROOT_RT  ge-11/1/1.100          flood      425
 0x3d80/17 ALT_ROOT_RT  ge-11/1/0.100          flood      425
 0x3e80/17 ALT_ROOT_RT  xe-10/2/0.100          flood      425
 0x3f80/17 ALT_ROOT_RT  xe-10/0/0.100          flood      425
Name: vs1::vlan200
CEs: 5
VEs: 0
Flood Routes:
  Prefix    Type         Owner                 NhType     NhIndex
  0x39/16   ALL_FLOOD    __vs1+vlan200__        flood      427
  0x39/16   RE_FLOOD     __vs1+vlan200__        flood      427
  0x4180/17 ALT_ROOT_RT  ge-11/1/0.200          flood      427
  0x4080/17 ALT_ROOT_RT  ge-11/1/1.200          flood      427
  0x4280/17 ALT_ROOT_RT  ge-11/1/4.200          flood      427
  0x4480/17 ALT_ROOT_RT  xe-10/0/0.200          flood      427
  0x4380/17 ALT_ROOT_RT  xe-10/2/0.200          flood      427
```

show bridge flood brief

```
user@host> show bridge flood brief
Name                 Active CEs       Active VEs
__juniper_private1__ 0                0
vs1::vlan100         6                0
vs1::vlan200         5                0
```

show bridge flood detail

```
user@host> show bridge flood detail
Name: __juniper_private1__
CEs: 0
VEs: 0
Flood Routes:
  Prefix    Type         Owner                 NhType     NhIndex
  0x36/16   MLP_FLOOD    __vs1+vlan100__        flood      426
  0x3a/16   MLP_FLOOD    __vs1+vlan200__        flood      428
```

```
Name: vs1::vlan100
CEs: 6
VEs: 0
Flood Routes:
  Prefix     Type          Owner               NhType        NhIndex
  0x35/16    ALL_FLOOD     __vs1+vlan100__     flood         425
  0x35/16    RE_FLOOD      __vs1+vlan100__     flood         425
  0x3780/17  ALT_ROOT_RT   ge-11/0/3.0         flood         425
  0x3b80/17  ALT_ROOT_RT   ge-11/1/4.100       flood         425
  0x3c80/17  ALT_ROOT_RT   ge-11/1/1.100       flood         425
  0x3d80/17  ALT_ROOT_RT   ge-11/1/0.100       flood         425
  0x3e80/17  ALT_ROOT_RT   xe-10/2/0.100       flood         425
  0x3f80/17  ALT_ROOT_RT   xe-10/0/0.100       flood         425
Name: vs1::vlan200
CEs: 5
VEs: 0
Flood Routes:
  Prefix     Type          Owner               NhType        NhIndex
  0x39/16    ALL_FLOOD     __vs1+vlan200__     flood         427
  0x39/16    RE_FLOOD      __vs1+vlan200__     flood         427
  0x4180/17  ALT_ROOT_RT   ge-11/1/0.200       flood         427
  0x4080/17  ALT_ROOT_RT   ge-11/1/1.200       flood         427
  0x4280/17  ALT_ROOT_RT   ge-11/1/4.200       flood         427
  0x4480/17  ALT_ROOT_RT   xe-10/0/0.200       flood         427
  0x4380/17  ALT_ROOT_RT   xe-10/2/0.200       flood         427
```

show bridge flood extensive

```
user@host> show bridge flood extensive
Name: __juniper_private1__
CEs: 0
VEs: 0
  Flood route prefix: 0x36/16
  Flood route type: MLP_FLOOD
  Flood route owner: __vs1+vlan100__
  Nexthop type: flood
  Nexthop index: 426
    Interfaces Flooding to:
    Name              Type          NhType          Index
    lc-11/0/0.32769   LC
    lc-10/2/0.32769   LC
```

```
      lc-10/0/0.32769  LC
      lc-11/1/0.32769  LC


   Flood route prefix: 0x3a/16
   Flood route type: MLP_FLOOD
   Flood route owner: __vs1+vlan200__
   Nexthop type: flood
   Nexthop index: 428
     Interfaces Flooding to:
     Name             Type         NhType         Index
     lc-10/0/0.32769  LC
     lc-10/2/0.32769  LC
     lc-11/1/0.32769  LC
Name: vs1::vlan100
CEs: 6
VEs: 0


   Flood route prefix: 0x35/16
   Flood route type: ALL_FLOOD
   Flood route owner: __vs1+vlan100__
   Nexthop type: flood
   Nexthop index: 425
     Interfaces Flooding to:
     Name             Type         NhType         Index
     ge-11/0/3.0      CE
     ge-11/1/4.100    CE
     ge-11/1/1.100    CE
     ge-11/1/0.100    CE
     xe-10/2/0.100    CE
     xe-10/0/0.100    CE


   Flood route prefix: 0x35/16
   Flood route type: RE_FLOOD
   Flood route owner: __vs1+vlan100__
   Nexthop type: flood
   Nexthop index: 425
     Interfaces Flooding to:
     Name             Type         NhType         Index
     ge-11/0/3.0      CE
     ge-11/1/4.100    CE
     ge-11/1/1.100    CE
     ge-11/1/0.100    CE
     xe-10/2/0.100    CE
```

```
   xe-10/0/0.100    CE


Flood route prefix: 0x3780/17
Flood route type: ALT_ROOT_RT
Flood route owner: ge-11/0/3.0
Nexthop type: flood
Nexthop index: 425
  Interfaces Flooding to:
  Name              Type          NhType         Index
  ge-11/0/3.0       CE
  ge-11/1/4.100     CE
  ge-11/1/1.100     CE
  ge-11/1/0.100     CE
  xe-10/2/0.100     CE
  xe-10/0/0.100     CE


Flood route prefix: 0x3b80/17
Flood route type: ALT_ROOT_RT
Flood route owner: ge-11/1/4.100
Nexthop type: flood
Nexthop index: 425
  Interfaces Flooding to:
  Name              Type          NhType         Index
  ge-11/0/3.0       CE
  ge-11/1/4.100     CE
  ge-11/1/1.100     CE
  ge-11/1/0.100     CE
  xe-10/2/0.100     CE
  xe-10/0/0.100     CE


Flood route prefix: 0x3c80/17
Flood route type: ALT_ROOT_RT
Flood route owner: ge-11/1/1.100
Nexthop type: flood
Nexthop index: 425
  Interfaces Flooding to:
  Name              Type          NhType         Index
  ge-11/0/3.0       CE
  ge-11/1/4.100     CE
  ge-11/1/1.100     CE
  ge-11/1/0.100     CE
  xe-10/2/0.100     CE
  xe-10/0/0.100     CE
```

```
Flood route prefix: 0x3d80/17
Flood route type: ALT_ROOT_RT
Flood route owner: ge-11/1/0.100
Nexthop type: flood
Nexthop index: 425
  Interfaces Flooding to:
  Name              Type          NhType          Index
  ge-11/0/3.0       CE
  ge-11/1/4.100     CE
  ge-11/1/1.100     CE
  ge-11/1/0.100     CE
  xe-10/2/0.100     CE
  xe-10/0/0.100     CE


Flood route prefix: 0x3e80/17
Flood route type: ALT_ROOT_RT
Flood route owner: xe-10/2/0.100
Nexthop type: flood
Nexthop index: 425
  Interfaces Flooding to:
  Name              Type          NhType          Index
  ge-11/0/3.0       CE
  ge-11/1/4.100     CE
  ge-11/1/1.100     CE
  ge-11/1/0.100     CE
  xe-10/2/0.100     CE
  xe-10/0/0.100     CE


Flood route prefix: 0x3f80/17
Flood route type: ALT_ROOT_RT
Flood route owner: xe-10/0/0.100
Nexthop type: flood
Nexthop index: 425
  Interfaces Flooding to:
  Name              Type          NhType          Index
  ge-11/0/3.0       CE
  ge-11/1/4.100     CE
  ge-11/1/1.100     CE
  ge-11/1/0.100     CE
  xe-10/2/0.100     CE
  xe-10/0/0.100     CE
Name: vs1::vlan200
```

```
CEs: 5
VEs: 0

  Flood route prefix: 0x39/16
  Flood route type: ALL_FLOOD
  Flood route owner: __vs1+vlan200__
  Nexthop type: flood
  Nexthop index: 427
    Interfaces Flooding to:
    Name            Type         NhType         Index
    ge-11/1/0.200   CE
    ge-11/1/1.200   CE
    ge-11/1/4.200   CE
    xe-10/0/0.200   CE
    xe-10/2/0.200   CE


  Flood route prefix: 0x39/16
  Flood route type: RE_FLOOD
  Flood route owner: __vs1+vlan200__
  Nexthop type: flood
  Nexthop index: 427
    Interfaces Flooding to:
    Name            Type         NhType         Index
    ge-11/1/0.200   CE
    ge-11/1/1.200   CE
    ge-11/1/4.200   CE
    xe-10/0/0.200   CE
    xe-10/2/0.200   CE


  Flood route prefix: 0x4180/17
  Flood route type: ALT_ROOT_RT
  Flood route owner: ge-11/1/0.200
  Nexthop type: flood
  Nexthop index: 427
    Interfaces Flooding to:
    Name            Type         NhType         Index
    ge-11/1/0.200   CE
    ge-11/1/1.200   CE
    ge-11/1/4.200   CE
    xe-10/0/0.200   CE
    xe-10/2/0.200   CE


  Flood route prefix: 0x4080/17
```

```
Flood route type: ALT_ROOT_RT
Flood route owner: ge-11/1/1.200
Nexthop type: flood
Nexthop index: 427
  Interfaces Flooding to:
  Name              Type           NhType          Index
  ge-11/1/0.200     CE
  ge-11/1/1.200     CE
  ge-11/1/4.200     CE
  xe-10/0/0.200     CE
  xe-10/2/0.200     CE


Flood route prefix: 0x4280/17
Flood route type: ALT_ROOT_RT
Flood route owner: ge-11/1/4.200
Nexthop type: flood
Nexthop index: 427
  Interfaces Flooding to:
  Name              Type           NhType          Index
  ge-11/1/0.200     CE
  ge-11/1/1.200     CE
  ge-11/1/4.200     CE
  xe-10/0/0.200     CE
  xe-10/2/0.200     CE


Flood route prefix: 0x4480/17
Flood route type: ALT_ROOT_RT
Flood route owner: xe-10/0/0.200
Nexthop type: flood
Nexthop index: 427
  Interfaces Flooding to:
  Name              Type           NhType          Index
  ge-11/1/0.200     CE
  ge-11/1/1.200     CE
  ge-11/1/4.200     CE
  xe-10/0/0.200     CE
  xe-10/2/0.200     CE


Flood route prefix: 0x4380/17
Flood route type: ALT_ROOT_RT
Flood route owner: xe-10/2/0.200
Nexthop type: flood
Nexthop index: 427
```

```
   Interfaces Flooding to:
   Name                 Type            NhType          Index
   ge-11/1/0.200        CE
   ge-11/1/1.200        CE
   ge-11/1/4.200        CE
   xe-10/0/0.200        CE
   xe-10/2/0.200        CE
```

Release Information

Command introduced in Junos OS Release 8.4.

show bridge mac-table

IN THIS SECTION

Syntax

show bridge mac-table
<age>
<brief | count | detail | extensive>

```
<bridge-domain (all | bridge-domain-name)>
<global-count>
<instance instance-name>
<interface interface-name>
<mac-address>
<instance instance-name>
<vlan-id (all-vlan | vlan-id)>
```

Description

(MX Series routers only) Display Layer 2 MAC address information.

Options

none	Display all learned Layer 2 MAC address information.
age	(Optional) Display age of a single mac-address.
brief | count | detail | extensive	(Optional) Display the specified level of output.
bridge-domain (all | bridge-domain-name)	(Optional) Display learned Layer 2 MAC addresses for all bridging domains or for the specified bridging domain.
global-count	(Optional) Display the total number of learned Layer 2 MAC addresses on the system.
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.
mac-address	(Optional) Display the specified learned Layer 2 MAC address information.
vlan-id (all-vlan | vlan-id)	(Optional) Display learned Layer 2 MAC addresses for all VLANs or for the specified VLAN.

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

Output Fields

Table 5 on page 210 describes the output fields for the **show bridge mac-table** command. Output fields are listed in the approximate order in which they appear.

Table 5: show bridge mac-table Output Fields

Field Name	Field Description
Age	Age of a single mac-address.
Routing instance	Name of the routing instance.
Bridging domain	Name of the bridging domain.
MAC address	MAC address or addresses learned on a logical interface.

Table 5: show bridge mac-table Output Fields *(Continued)*

Field Name	Field Description
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. • **C**—Control MAC address is configured. • **SE**—MAC accounting is enabled. • **NM**—Non-configured MAC. • **R**—Remote PE MAC address is configured. • **P**—MAC Pinned interface is configured
Logical interface	Name of the logical interface.
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 bridge domain in which the MAC address was learned.
VXLAN ID/VXLAN	VXLAN Network Identifier (VNI).
Layer 2 flags	Debugging flags signifying that the MAC address is present in various lists.
Epoch	Spanning Tree Protocol epoch number identifying when the MAC address was learned. Used for debugging.

Table 5: show bridge mac-table Output Fields *(Continued)*

Field Name	Field Description
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.

Sample Output

show bridge mac-table

```
user@host> show bridge mac-table
MAC flags (S -static MAC, D -dynamic MAC, L -locally learned, C -Control MAC
          SE -Statistics enabled, NM -Non configured MAC, R -Remote PE MAC)


Routing instance : default-switch
 Bridging domain : test1, VLAN : 1
   MAC                  MAC       Logical          NH       RTR
   address              flags     interface        Index    ID
   01:00:0c:cc:cc:cc    S,NM      NULL
   01:00:0c:cc:cc:cd    S,NM      NULL
   01:00:0c:cd:cd:d0    S,NM      NULL
   64:87:88:6a:17:d0    D         ae0.1
   64:87:88:6a:17:f0    D         ae0.1
```

show bridge mac-table (with Layer 2 Services over GRE Interfaces)

```
user@host> show bridge mac-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
 Bridging domain : vlan-1, VLAN : 1
   MAC                 MAC       Logical
   address             flags     interface
   00:01:01:00:01:f7   D,SE      gr-1/2/10.0
   00:03:00:32:01:f7   D,SE      gr-1/2/10.0
   00:00:21:11:11:10   DL        ge-1/0/0.0
   00:00:21:11:11:11   DL        ge-1/1/0.0



Routing instance : default-switch
 Bridging domain : vlan-2, VLAN : 2
   MAC                 MAC       Logical
   address             flags     interface
   00:02:01:33:01:f7   D,SE      gr-1/2/10.1
   00:00:21:11:21:10   DL        ge-1/0/0.1
   00:00:21:11:21:11   DL        ge-1/1/0.1
```

show bridge mac-table (with VXLAN enabled)

```
user@host> show bridge mac-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
 Bridging domain : vlan-1, VLAN : 1
  VXLAN: Id : 100, Multicast group: 233.252.0.1
   MAC                 MAC       Logical
   address             flags     interface
   00:01:01:00:01:f7   D,SE      vtep.1052010
   00:03:00:32:01:f7   D,SE      vtep.1052011
   00:00:21:11:11:10   DL        ge-1/0/0.0
   00:00:21:11:11:11   DL        ge-1/1/0.0



Routing instance : default-switch
 Bridging domain : vlan-2, VLAN : 2, VXLAN : 200
  VXLAN: Id : 200, Multicast group: 233.252.0.2
```

```
   MAC                  MAC      Logical
   address              flags    interface
   00:02:01:33:01:f7    D,SE     vtep.1052010
   00:04:00:14:01:f7    D,SE     vtep.1052011
   00:00:21:11:21:10    DL       ge-1/0/0.1
   00:00:21:11:21:11    DL       ge-1/1/0.1
```

show bridge mac-table age (for GE interface)

```
user@host> show vpls mac-table age 00:02:03:aa:bb:1a instance vpls_instance_1
MAC Entry Age information
Current Age: 4 seconds
```

show bridge mac-table age (for AE interface)

```
user@host> show vpls mac-table age 00:02:03:aa:bb:1a instance vpls_instance_1
MAC Entry Age information
Current Age on FPC1: 102 seconds
Current Age on FPC2: 94 seconds
```

show bridge mac-table count

```
user@host> show bridge mac-table count
2 MAC address learned in routing instance vs1 bridge domain vlan100

  MAC address count per interface within routing instance:
    Logical interface        MAC count
    ge-11/0/3.0                      1
    ge-11/1/4.100                    0
    ge-11/1/1.100                    0
    ge-11/1/0.100                    0
    xe-10/2/0.100                    1
    xe-10/0/0.100                    0


  MAC address count per learn VLAN within routing instance:
    Learn VLAN ID            MAC count
              0                      2
```

```
0 MAC address learned in routing instance vs1 bridge domain vlan200

  MAC address count per interface within routing instance:
    Logical interface        MAC count
    ge-11/1/0.200                    0
    ge-11/1/1.200                    0
    ge-11/1/4.200                    0
    xe-10/0/0.200                    0
    xe-10/2/0.200                    0


  MAC address count per learn VLAN within routing instance:
    Learn VLAN ID            MAC count
                0                    0
```

show bridge mac-table detail

```
user@host> show bridge mac-table detail
MAC address: 00:00:00:19:1c:db
  Routing instance: vs1
   Bridging domain: vlan100
   Learning interface: ge-11/0/3.0     Learning VLAN: 0
   Layer 2 flags: in_ifd, in_ifl, in_vlan, kernel
   Epoch: 4                            Sequence number: 0
   Learning mask: 0x800               IPC generation: 0


MAC address: 00:00:00:59:3a:2f
  Routing instance: vs1
   Bridging domain: vlan100
   Learning interface: xe-10/2/0.100   Learning VLAN: 0
   Layer 2 flags: in_ifd, in_ifl, in_vlan, kernel
   Epoch: 7                            Sequence number: 0
   Learning mask: 0x400               IPC generation: 0
```

show bridge mac-table instance pbb-evpn

```
user@host> show bridge mac-table instance pbb-evpn
Routing instance : pbb-evpn
 Bridging domain : isid-bd10000, ISID : 10000
    MAC                      MAC      Logical          NH     RTR
```

```
    address                flags    interface        Index   ID
    00:19:e2:b0:76:eb   D        cbp.1000
    aa:bb:cc:dd:ee:f2   DC                        1048576 1048576
    aa:bb:cc:dd:ee:f3   DC                        1048575 1048575
```

show bridge mac-table

```
user@host>run show bridge mac-table
MAC flags (S -static MAC, D -dynamic MAC, L -locally learned, C -Control MAC
O -OVSDB MAC, SE -Statistics enabled, NM -Non configured MAC, R -Remote PE MAC,
P -Pinned MAC)

Routing instance : VS-541
Bridging domain : 541, VLAN : 541
MAC MAC Logical NH RTR
address flags interface Index ID
00:00:01:00:00:01 DPRC xe-0/0/3.0
00:00:02:00:00:01 DP xe-0/0/3.0
```

Release Information

Command introduced in Junos OS Release 8.4.

Support for PBB-EVPN instance added in Junos OS Release 16.1

MAC Flag P to indicate a MAC Pinned interface introduced in Junos OS 16.2

show bridge statistics

IN THIS SECTION

- Syntax | 217
- Description | 217

Syntax

```
show bridge statistics
<bridge-domain domain-name>
<instance instance-name>
```

Description

(MX Series routers only) Display bridge statistics.

Options

none	Display bridge statistics for all bridge domains in all routing instances.
bridge-domain *domain-name*	(Optional) Display statistics for the specified bridge domain.
instance *instance-name*	(Optional) Display statistics for the specified routing instance.

Required Privilege Level

view

Sample Output

show bridge statistics

```
user@host> show bridge statistics
Information for routing instance:

Routing instance : __juniper_private1__
    Index: 1                          Sequence number: 0
    MAC limit: 5000                   MACs learned: 0
    Static MACs learned: 0            Non config Static MACs learned: 0
    Handle: 0x829e800


Information for routing instance:

Routing instance : vs1
 Bridging domain : vlan100
    Index: 3                          Sequence number: 0
    MAC limit: 5120                   MACs learned: 2
    Static MACs learned: 0            Non config Static MACs learned: 0
    Handle: 0x829e400
    Flags: Bridge instance, Config defined, VLAN : 100
    Local interface: ge-11/0/3.0, Index: 79
      Broadcast packets:                      1
      Broadcast bytes  :                     65
      Multicast packets:                      0
      Multicast bytes  :                      0
      Flooded packets  :                      0
      Flooded bytes    :                      0
      Unicast packets  :              358624489
      Unicast bytes    :            23310592305
      Current MAC count:                      1 (Limit 1024)
    Local interface: ge-11/1/4.100, Index: 84
      Broadcast packets:                      0
      Broadcast bytes  :                      0
      Multicast packets:                      0
      Multicast bytes  :                      0
      Flooded packets  :                      0
      Flooded bytes    :                      0
      Unicast packets  :                      0
      Unicast bytes    :                      0
```

```
      Current MAC count:                    0 (Limit 1024)
   Local interface: ge-11/1/1.100, Index: 86
     Broadcast packets:                     0
     Broadcast bytes  :                     0
     Multicast packets:                     0
     Multicast bytes  :                     0
     Flooded packets  :                     0
     Flooded bytes    :                     0
     Unicast packets  :                     0
     Unicast bytes    :                     0
     Current MAC count:                     0 (Limit 1024)
   Local interface: ge-11/1/0.100, Index: 87
     Broadcast packets:                     0
     Broadcast bytes  :                     0
     Multicast packets:                     0
     Multicast bytes  :                     0
     Flooded packets  :                     0
     Flooded bytes    :                     0
     Unicast packets  :                     0
     Unicast bytes    :                     0
     Current MAC count:                     0 (Limit 1024)
   Local interface: xe-10/2/0.100, Index: 88
     Broadcast packets:                     0
     Broadcast bytes  :                     0
     Multicast packets:                     0
     Multicast bytes  :                     0
     Flooded packets  :                     0
     Flooded bytes    :                     0
     Unicast packets  :             358627393
     Unicast bytes    :           23310781065
     Current MAC count:                     1 (Limit 1024)
   Local interface: xe-10/0/0.100, Index: 89
     Broadcast packets:                     0
     Broadcast bytes  :                     0
     Multicast packets:                     0
     Multicast bytes  :                     0
     Flooded packets  :                     0
     Flooded bytes    :                     0
     Unicast packets  :                     0
     Unicast bytes    :                     0
     Current MAC count:                     0 (Limit 1024)


 Information for routing instance:
```

```
Routing instance : vs1
 Bridging domain : vlan200
    Index: 4                        Sequence number: 0
    MAC limit: 5120                 MACs learned: 0
    Static MACs learned: 0          Non config Static MACs learned: 0
    Handle: 0x829e600
    Flags: Bridge instance, Config defined, VLAN : 200
    Local interface: ge-11/1/0.200, Index: 90
      Broadcast packets:                    0
      Broadcast bytes  :                    0
      Multicast packets:                    0
      Multicast bytes  :                    0
      Flooded packets  :                    0
      Flooded bytes    :                    0
      Unicast packets  :                    0
      Unicast bytes    :                    0
      Current MAC count:                    0 (Limit 1024)
    Local interface: ge-11/1/1.200, Index: 91
      Broadcast packets:                    0
      Broadcast bytes  :                    0
      Multicast packets:                    0
      Multicast bytes  :                    0
      Flooded packets  :                    0
      Flooded bytes    :                    0
      Unicast packets  :                    0
      Unicast bytes    :                    0
      Current MAC count:                    0 (Limit 1024)
    Local interface: ge-11/1/4.200, Index: 92
      Broadcast packets:                    0
      Broadcast bytes  :                    0
      Multicast packets:                    0
      Multicast bytes  :                    0
      Flooded packets  :                    0
      Flooded bytes    :                    0
      Unicast packets  :                    0
      Unicast bytes    :                    0
      Current MAC count:                    0 (Limit 1024)
    Local interface: xe-10/0/0.200, Index: 93
      Broadcast packets:                    0
      Broadcast bytes  :                    0
      Multicast packets:                    0
      Multicast bytes  :                    0
```

```
    Flooded packets  :                        0
    Flooded bytes    :                        0
    Unicast packets  :                        0
    Unicast bytes    :                        0
    Current MAC count:                        0 (Limit 1024)
 Local interface: xe-10/2/0.200, Index: 94
    Broadcast packets:                        4
    Broadcast bytes  :                      260
    Multicast packets:                        0
    Multicast bytes  :                        0
    Flooded packets  :                        0
    Flooded bytes    :                        0
    Unicast packets  :                        0
    Unicast bytes    :                        0
    Current MAC count:                        0 (Limit 1024)
```

Release Information

Command introduced in Junos OS Release 8.4.

12
CHAPTER

Operational Mode Commands for Layer 2 Learning

clear l2-learning mac-move-buffer

Syntax

```
clear l2-learning mac-move-buffer
<active>
```

Description

Clear the MAC move buffer entries.

Options

none Clear the MAC move buffer entries.

active (Optional) Unblock the interfaces that were blocked by the MAC move action feature. This allows the user to keep the **reopen-time** configured to a large value, but when the looping error is fixed, the user can manually release the blocking.

Required Privilege Level

clear

Output Fields

When you enter this command, the MAC move buffer entries are deleted.

Sample Output

clear l2-learning mac-move-buffer

```
user@host>  clear l2-learning mac-move-buffer
```

clear l2-learning mac-move-buffer active

```
user@host>  clear l2-learning mac-move-buffer active
```

Release Information

Command introduced in Junos OS Release 13.2.

show l2-learning global-information

IN THIS SECTION

Syntax

```
show l2-learning global-information
```

Description

Display Layer 2 learning process-related information for the entire device.

Options

This command has no options.

Required Privilege Level

view

Output Fields

Table 6 on page 226 describes the output fields for the **show l2-learning global-information** command. Output fields are listed in the approximate order in which they appear.

Table 6: show l2-learning global-information Output Fields

Field Name	Field Description
MAC aging interval	Configured timeout interval, in seconds, for all MAC table entries.
MAC learning	Status of MAC learning: **Enabled** or **Disabled**.
MAC statistics	Status of MAC accounting: **Enabled** or **Disabled**.
MAC limit Count	Configured maximum limit on the number of MAC addresses that can be learned.
MAC limit hit flag	Status of the learned MAC limit hit flag: **Enabled** (the learned MAC exceeds the global MAC limit) or **Disabled** (the learned MAC does not exceed the global MAC limit).
MAC packet action drop	Status of action to drop packets after the configured MAC address limit is reached: **Enabled** (packets are dropped) or **Disabled** (packets are forwarded).

Sample Output

show l2-learning global-information

```
user@host> show l2-learning global-information
Global Configuration:

MAC aging interval    : 300
MAC learning          : Enabled
MAC statistics        : Disabled
MAC limit Count       : 393215
MAC limit hit flag    : Disabled
MAC packet action drop: Disabled
```

Release Information

Command introduced in Junos OS Release 8.4.

show l2-learning global-mac-count

Syntax

show l2-learning global-mac-count

Description

(MX Series routers only) Display the total number of dynamic and static MAC addresses learned for the entire router.

Options

This command has no options.

Required Privilege Level

view

Output Fields

Displays the total number of dynamic and static MAC addresses learned for the entire router.

Sample Output

show l2-learning global-mac-count

```
user@host> show l2-learning global-mac-count
100 dynamic and static MAC addresses learned globally
```

Release Information

Command introduced in Junos OS Release 9.3.

show l2-learning instance

Syntax

show l2-learning instance

Description

Display Layer 2 learning properties for all the configured routing instances.

Options

This command has no options.

Required Privilege Level

view

Output Fields

describes the output fields for the **show l2-learning instance** command. Output fields are listed in the approximate order in which they appear.

Table 7: show l2-learning instance Output Fields

Field Name	Field Description
Routing Instance	Name of routing instance.
Bridging Domain	Name of bridging domain. On MX Series routers you can use the **show l2-learning instance <extensive>** command option to display the Bridge Service-id information which includes the Config Service ID and the Active Service ID.
Index	Number associated with the routing instance or bridging domain.
Logical System	Name of logical system or **Default** if no logical system is configured.
Routing instance flags	Status of Layer 2 learning properties for each routing instance: • **DL**—MAC learning is disabled. • **SE**—MAC accounting is enabled. • **AD**—Packets are dropped after MAC address limit is reached. • **LH**—The maximum number of MAC addresses has been learned on the routing instance. The routing instance is not able to learn any additional MAC addresses.
MAC limit	Maximum number of MAC addresses that can be learned from each interface in the routing instance or bridging domain.

Sample Output

show l2-learning instance

```
user@host> show l2-learning instance
Information for routing instance:

Routing Instance flags (DL -disable learning, SE -stats enabled,
                        AD -packet action drop, LH -mac limit hit)

Routing                 Bridging      Index   Logical         Routing      MAC
Instance                Domain                System          flags        limit
__juniper_private1__                     1    Default                      5000
vs1                     vlan100          3    Default                      5120
vs1                     vlan200          4    Default                      5120
```

Release Information

(MX Series routers only) Command introduced in Junos OS Release 8.4.

show l2-learning interface

IN THIS SECTION

Syntax

show l2-learning interface

Description

(MX Series routers only) Display Layer 2 learning information for all the interfaces.

Options

This command has no options.

Required Privilege Level

view

Output Fields

Table 8 on page 232 describes the output fields for the **show l2-learning interface** command. Output fields are listed in the approximate order in which they appear.

Table 8: show l2-learning interfaceOutput Fields

Field Name	Field Description
Logical interface	Name of the logical interface.
Index	Index of the interface.

Table 8: show l2-learning interfaceOutput Fields *(Continued)*

Field Name	Field Description
Routing Instance	Number of the routing instance to which the interface belongs.
Interface device	Value of the order in which the Junos OS finds and initializes the interface.
Logical interface flags	Status of Layer 2 learning properties for each interface: • **DL**—MAC learning is disabled. • **SE**—MAC accounting is enabled. • **AD**—Packets are dropped after the MAC interface limit is reached. • **MAC limit**—Maximum number of MAC addresses that can be learned from the interface. • **MP**—MAC Pinning enabled.

Sample Output

show l2-learning interface

```
user@host> showl2-learning interface
Information for interface family:

Logical Interface flags (DL -disable learning, SE -stats enabled,
                    AD -packet action drop, LH -mac limit hit)

Logical             Index  Routing   Interface  Logical            MAC
interface                  instance  device     Interface flags    limit
ge-11/0/3.0          79         3       136                        1024
ge-11/1/4.100        84         3       150                        1024
ge-11/1/1.100        86         3       147                        1024
ge-11/1/0.100        87         3       146                        1024
xe-10/2/0.100        88         3       144                        1024
```

```
xe-10/0/0.100          89          3          129                          1024
ge-11/1/0.200          90          4          146                          1024
ge-11/1/1.200          91          4          147                          1024
ge-11/1/4.200          92          4          150                          1024
xe-10/0/0.200          93          4          129                          1024
xe-10/2/0.200          94          4          144                          1024
```

show l2 learning-interface

```
user@host>  run show l2-learning interface
Routing Instance Name : default-switch
Logical Interface flags (DL -disable learning, AD -packet action drop,
                    LH - MAC limit hit, DN - Interface Down,  MP - MAC Pinning
enabled)
Logical            BD          MAC        STP          Logical
Interface          Name        Limit      State        Interface flags
ae0.0                          8192                     MP
```

Release Information

Command introduced in Junos OS Release 8.4.

Added sample output to indicate an EVPN MAC Pinned interface, introduced in Junos OS 16.2R1.

show l2-learning mac-move-buffer

IN THIS SECTION

Syntax

show l2-learning mac-move-buffer
<brief | detail | extensive>
<active>

Description

(MX Series routers only) Display action as a result of configuring the MAC address move feature.

Options

none	Display action as a result of the MAC address move feature.
brief | detail | extensive	(Optional) Display the specified level of output.
active	(Optional) Display the set of interfaces blocked as a result of the MAC address move action.

Required Privilege Level

view

Output Fields

Display action as a result of the MAC address move feature.

Sample Output

show l2-learning mac-move-buffer active

```
user@host> show l2-learning mac-move-buffer active
MAC Address: 00:00:00:00:01:01, VLAN Id: 0
    Time Rec : 2012-06-25 06:23:41    Bridge Domain: bd10
    Prev IFL : ge-1/0/5.0             New IFL: ge-1/0/6.0
    IFBD     : ge-1/0/6.0:10          Blocked  : YES
```

show l2-learning mac-move-buffer extensive

```
user@host> show l2-learning mac-move-buffer extensive | display xml
<l2ald-mac-move-buffer>
<l2ald-mac-move-entry junos:style="extensive">
<l2ald-mac-address>aa:00:00:00:02:00</l2ald-mac-address>
<l2ald-learn-vlan-id>0<l2ald-learn-vlan-id>
<l2ald-mac-move-time-rec>11:27:57</l2ald-mac-move-time-rec>
<l2ald-mac-move-bridge-domain>bd</l2ald-mac-move-bridge-domain>
<l2ald-mac-move-from-ifl>ge-1/0/5.200</l2ald-mac-move-from-ifl>
<l2ald-mac-move-to-ifl>ge-1/0/6.200</l2ald-mac-move-to-ifl>
<l2ald-mac-move-to-ifbd>ge-1/0/6.200</l2ald-mac-move-to-ifbd>
<l2ald-mac-move-is-blocked>Yes</l2ald-mac-move-is-blocked>
</l2ald-mac-move-entry>
```

Release Information

Command introduced in Junos OS Release 9.4.

13

Knowledge Base