Junos® OS

Multichassis Link Aggregation Feature Guide for EX Series, MX Series, and QFX Series Devices

Modified: 2019-06-24

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

Junos® OS Multichassis Link Aggregation Feature Guide for EX Series, MX Series, and QFX Series Devices

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

List of Figures

List of Tables

About the Documentation

Documentation and Release Notes

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

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

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

Using the Examples in This Manual

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

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

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

Merging a Full Example

To merge a full example, follow these steps:

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

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

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

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

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

Merging a Snippet

To merge a snippet, follow these steps:

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

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

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

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

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

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

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

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

Documentation Conventions

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

Table 1: Notice Icons

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

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

Table 2: Text and Syntax Conventions

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

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

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

GUI Conventions

Table 2: Text and Syntax Conventions (continued)

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

Documentation Feedback

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

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

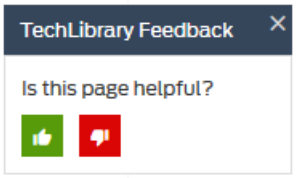

  - Click the thumbs-up icon if the information on the page was helpful to you.
  - Click the thumbs-down icon if the information on the page was not helpful to you or if you have suggestions for improvement, and use the pop-up form to provide feedback.

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

Requesting Technical Support

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

- JTAC policies—For a complete understanding of our JTAC procedures and policies, review the *JTAC User Guide* located at https://www.juniper.net/us/en/local/pdf/resource-guides/7100059-en.pdf.
- Product warranties—For product warranty information, visit https://www.juniper.net/support/warranty/.
- JTAC hours of operation—The JTAC centers have resources available 24 hours a day, 7 days a week, 365 days a year.

Self-Help Online Tools and Resources

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

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

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

Creating a Service Request with JTAC

You can create a service request with JTAC on the Web or by telephone.

- Visit https://myjuniper.juniper.net.
- Call 1-888-314-JTAC (1-888-314-5822 toll-free in the USA, Canada, and Mexico).

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

CHAPTER 1

Overview

Understanding Multichassis Link Aggregation Groups

Layer 2 networks are increasing in scale mainly because of technologies such as virtualization. Protocol and control mechanisms that limit the disastrous effects of a topology loop in the network are necessary. The Spanning Tree Protocol (STP) is the primary solution to this problem because it provides a loop-free Layer 2 environment. STP has gone through a number of enhancements and extensions, and even though it scales to very large network environments, it still only provides one active path from one device to another, regardless of how many actual connections might exist in the network. Although STP is a robust and scalable solution to redundancy in a Layer 2 network, the single logical link creates two problems: At least half of the available system bandwidth is off-limits to data traffic, and network topology changes occur. The Rapid Spanning Tree Protocol (RSTP) reduces the overhead of the rediscovery process and allows a Layer 2 network to reconverge faster, but the delay is still high.

Link aggregation (IEEE 802.3ad) solves some of these problems by enabling users to use more than one link connection between switches. All physical connections are considered one logical connection. The problem with standard link aggregation is that the connections are point to point.

Multichassis link aggregation groups (MC-LAGs) enable a client device to form a logical LAG interface between two MC-LAG peers. An MC-LAG provides redundancy and load balancing between the two MC-LAG peers, multihoming support, and a loop-free Layer 2 network without running STP.

On one end of an MC-LAG, there is an MC-LAG client device, such as a server, that has one or more physical links in a link aggregation group (LAG). This client device uses the link as a LAG. On the other side of the MC-LAG, there can be a maximum of two MC-LAG peers. Each of the MC-LAG peers has one or more physical links connected to a single client device.

The MC-LAG peers use the Inter-Chassis Control Protocol (ICCP) to exchange control information and coordinate with each other to ensure that data traffic is forwarded properly.

The Link Aggregation Control Protocol (LACP) is a subcomponent of the IEEE 802.3ad standard. LACP is used to discover multiple links from a client device connected to an MC-LAG peer. LACP must be configured on both MC-LAG peers for an MC-LAG to work correctly.

NOTE: You must specify a service identifier (service-id) at the global level; otherwise, multichassis link aggregation will not work.

Figure 1: Basic MC-LAG Topology

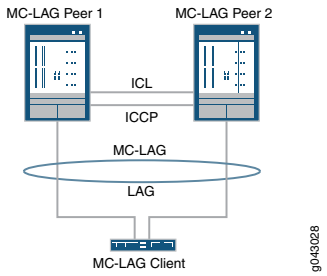

The following sections provide information regarding the functional behavior of multichassis link aggregation, configuration guidelines, and best practices.

Benefits of MC-LAGs

Multichassis link aggregation groups (MC-LAGs) provide redundancy and load balancing between a maximum of two switches, multihoming support for client devices such as servers, and a loop-free Layer 2 network without running Spanning Tree Protocol (STP).

ICCP and ICL

The MC-LAG peers use the Inter-Chassis Control Protocol (ICCP) to exchange control information and coordinate with each other to ensure that data traffic is forwarded properly. ICCP replicates control traffic and forwarding states across the MC-LAG peers and communicates the operational state of the MC-LAG members. Because ICCP uses TCP/IP to communicate between the peers, the two peers must be connected to each other. ICCP messages exchange MC-LAG configuration parameters and ensure that both peers use the correct LACP parameters.

The interchassis link (ICL), also known as the interchassis link-protection link (ICL-PL), is used to forward data traffic across the MC-LAG peers. This link provides redundancy when a link failure (for example, an MC-LAG trunk failure) occurs on one of the active links. The ICL can be a single physical Ethernet interface or an aggregated Ethernet interface.

You can configure multiple ICLs between MC-LAG peers. Each ICL can learn up to 512K MAC addresses. You can configure additional ICLs for virtual switch instances.

NOTE: DHCP snooping, dynamic ARP inspection (DAI), and IP source guard are not supported on the ICL or MC-LAG interfaces. Consequently, incoming address resolution protocol replies on the ICL are discarded. However, ARP entries can be populated on the ICL interface through ICCP exchanges from a remote MC-LAG peer.

BEST PRACTICE: We recommend that you use separate ports and choose different Flexible PIC Concentrators (FPCs) for the interchassis link (ICL) and Inter-Chassis Control Protocol (ICCP) interfaces. Although you can use a single link for the ICCP interface, an aggregated Ethernet interface is preferred.

When configuring ICCP and ICL, we recommend that you:

- Configure an aggregated Ethernet interface to be used for the ICL interface.
- Configure an aggregated Ethernet interface to be used for the ICCP interface.
- Configure the IP address for the management port (fxp0).

  When you configure backup liveness detection, this out-of-band channel is established between the peers through the management network
- Use the peer loopback address to establish ICCP peering. Doing so avoids any direct link failure between MC-LAG peers. As long as the logical connection between the peers remains up, ICCP stays up.
- Configure the ICCP liveness-detection interval (the Bidirectional Forwarding Detection (BFD) timer) to be at least 8 seconds if you have configured ICCP connectivity through an IRB interface. A liveness-detection interval of 8 seconds or more allows for graceful Routing Engine switchover (GRES) to work seamlessly. By default, ICCP liveness detection uses multihop BFD, which runs in centralized mode.

  This recommendation does not apply if you configured ICCP connectivity through a dedicated physical interface. In this case, you can configure single-hop BFD.
- Configure a session establishment hold time for ICCP. Doing this results in faster ICCP connection between the MC-LAG peers and also prevents any delay during convergence.

  

  BEST PRACTICE: On QFX and EX Series switches, the default session establishment hold time is 300 seconds. However, the session establishment time must be at least 100 seconds higher than the init delay time. You can optionally update the session establishment time to be 340 seconds and the init delay time to be 240 seconds.
- Configure a hold-down timer on the ICL member links that is greater than the configured BFD timer for the ICCP interface. This prevents the ICL from being advertised as being down before the ICCP link is down. If the ICL goes down before the ICCP link, this causes a flap of the MC-LAG interface on the status-control standby node, which leads to a delay in convergence.
- Starting with Junos OS Release 15.1 on MX Series routers, configure the backup liveness detection feature to implement faster failover of data traffic during an MC-LAG peer reboot. Configure the backup-liveness-detectionstatement on the management interface (fxp0) only.

Failure Handling

Configuring ICCP adjacency over an aggregated interface with child links on multiple FPCs mitigates the possibility of a split-brain state. A split-brain occurs when ICCP adjacency is lost between the MC-LAG peers. To work around this problem, enable backup liveness detection. With backup liveness detection enabled, the MC-LAG peers establish an out-of-band channel through the management network in addition to the ICCP channel.

During a split-brain state, both active and standby peers change LACP system IDs. Because both MC-LAG peers change the LACP system ID, the customer edge (CE) device accepts the LACP system ID of the first link that comes up and brings down other links carrying different LACP system IDs. When the ICCP connection is active, both of the MC-LAG peers use the configured LACP system ID. If the LACP system ID is changed during failures, the server that is connected over the MC-LAG removes these links from the aggregated Ethernet bundle.

When the ICL is operationally down and the ICCP connection is active, the LACP state of the links with status control configured as standby is set to the standby state. When the LACP state of the links is changed to standby, the server that is connected over the MC-LAG makes these links inactive and does not use them for sending data.

Recovery from the split-brain state occurs automatically when the ICCP adjacency comes up between MC-LAG peers.

If only one physical link is available for ICCP, then ICCP might go down due to link failure or FPC failure, while the peer is still up. This results in a split-brain state. If you do not set a special configuration to avoid this situation, the MC-LAG interfaces change the LACP sytem ID to their local defaults, thus ensuring that only one link (the first) comes up from the downstream device. A convergence delay results from thte LACP state changes on both active and standby peers.

Table 3 on page 25 describes the different ICCP failure scenarios for EX9200 switches. The dash means that the item is not applicable.

Table 3: ICCP Failure Scenarios for EX9200 Switches

ICCP Connection Status	ICL Status	Backup Liveness Peer Status	Action on Multichassis Aggregated Ethernet Interface with Status Set to Standby	Action on Multichassis Aggregated Ethernet Interface with Status Set to Standby and Prefer Status Control Set to Active
Down	Down or Up	Not configured	LACP system ID is changed to default value.	Not applicable. Liveness detection must be configured.
Down	Down or Up	Active	LACP system ID is changed to default value.	No change in LACP system ID.
Down	Down or Up	Inactive	No change in LACP system ID.	No change in LACP system ID.

Table 3: ICCP Failure Scenarios for EX9200 Switches (continued)

ICCP Connection Status	ICL Status	Backup Liveness Peer Status	Action on Multichassis Aggregated Ethernet Interface with Status Set to Standby	Action on Multichassis Aggregated Ethernet Interface with Status Set to Standby and Prefer Status Control Set to Active
Up	Down	–	LACP state is set to standby. MUX state moves to waiting state.	LACP status is set to standby. MUX state moves to waiting status.

Table 4 on page 26 describes the different ICCP failure scenarios for QFX Series switches. The dash means that the item is not applicable.

Table 4: ICCP Failure Scenarios for QFX Series Switches

ICCP Connection Status	ICL Status	Backup Liveness Peer Status	Action on Multichassis Aggregated Ethernet Interface with Status Set to Standby
Down	Down or Up	Not configured	LACP system ID is changed to default value.
Down	Down or Up	Active	LACP system ID is changed to default value.
Down	Down or Up	Inactive	No change in LACP system ID.
Up	Down	–	LACP state is set to standby. MUX state moves to waiting state.

Configure the **master-only** statement on the IP address of the fxp0 interface for backup liveness detection on both the master and backup Routing Engines. This ensures that the connection is not reset during GRES in the remote peer.

For example, on the master Routing Engine:

```
user@switch-re1 > show configuration interfaces fxp0 | display inheritance no-comments
unit 0 {
    family inet {
        address 10.8.2.31/8;
        address 10.8.2.33/8 {
            master-only;
        }
    }
}
```

For example, on the backup Routing Engine:

```
user@switch1-re1 > show configuration interfaces fxp0 | display inheritance no-comments
unit 0 {
    family inet {
        address 10.8.2.32/8;
        address 10.8.2.33/8 {
```

```
        master-only;
      }
    }
}
```

The master Routing Engine services both 10.8.2.31 and 10.8.2.33. Configure 10.8.2.33 in a backup-liveness-detection configuration on the peer node.

For example, on the backup Routing Engine:

```
user@switch2 > show configuration protocols iccp
local-ip-addr 10.2.2.2;
peer 10.1.1.1 {
   session-establishment-hold-time 340;
   redundancy-group-id-list 1;
   backup-liveness-detection {
      backup-peer-ip 10.8.2.33;
   }
   liveness-detection {
      minimum-interval 500;
      multiplier 3;
      single-hop;
   }
}
```

Multichassis Link Protection

Multichassis link protection provides link protection between the two MC-LAG peers that host an MC-LAG. If the ICCP connection is up and the ICL comes up, the peer configured as standby brings up the multichassis aggregated Ethernet interfaces shared with the peer. Multichassis protection must be configured on each MC-LAG peer that is hosting an MC-LAG.

MC-AE Statement Options

The following options are available:

- MC-AE-ID

  Specifies which MC-LAG group the aggregated Ethernet interface belongs to.

- Redundancy groups

  Uses ICCP to associate multiple chassis that perform similar redundancy functions and to establish a communication channel so that applications on peering chassis can send messages to each other.

BEST PRACTICE: We recommend that you configure only one redundancy group between MC-LAG nodes. The redundancy group represents the domain of high availability between the MC-LAG nodes. One redundancy group is sufficient between a pair of MC-LAG nodes. If you are using logical

systems, then configure one redundancy group between MC-LAG nodes in each logical system.

- Init Delay Time

  Specifies the number of seconds by which to delay bringing the MC-LAG interface back to the up state when the MC-LAG peer is rebooted. By delaying the startup of the interface until after protocol convergence, you can prevent packet loss during the recovery of failed links and devices.

  

  NOTE: On QFX and EX Series switches, the default session establishment hold time is 300 seconds. However, the session establishment time must be at least 100 seconds higher than the init delay time. You can optionally update the session establishment time to be 340 seconds and the init delay time to be 240 seconds.

- Chassis ID

  Specifies that LACP uses the chassis ID to calculate the port number of the MC-LAG physical member links. Each MC-LAG peer should have a unique chassis ID.

- Mode

  Indicates whether an MC-LAG is in active-standby mode or active-active mode. Chassis that are in the same group must be in the same mode.

  In active-active mode, all member links are active on the MC-LAG. In this mode, media access control (MAC) addresses learned on one MC-LAG peer are propagated to the other MC-LAG peer. Active-active mode is a simple and deterministic design and is easier to troubleshoot than active-standby mode.

  

  NOTE: Active-active mode is not supported on Dense Port Concentrator (DPC) line cards. Instead, use active-standby mode.

  In active-active MC-LAG topologies, network interfaces are categorized into three interface types, as follows:

  - S-Link—Single-homed link (S-Link) terminating on an MC-LAG peer device
  - MC-Link—MC-LAG link
  - ICL—Inter-chassis link

  Depending on the incoming and outgoing interface types, some constraints are added to the Layer 2 forwarding rules for MC-LAG configurations. The following data traffic forwarding rules apply.

  

  NOTE: If only one MC-LAG member link is in the up state, it is considered an S-Link.

- When an MC-LAG network receives a packet from a local MC-Link or S-Link, the packet is forwarded to other local interfaces, including S-Links and MC-Links based on the normal Layer 2 forwarding rules and on the configuration of the **mesh-group** and **no-local-switching** statements. If MC-Links and S-Links are in the same mesh group and their **no-local-switching** statements are enabled, the received packets are only forwarded upstream and not sent to MC-Links and S-Links.
- The following circumstances determine whether or not an ICL receives a packet from a local MC-Link or S-Link:
  - If the peer MC-LAG network device has S-Links or MC-LAGs that do not reside on the local MC-LAG network device
  - Whether or not interfaces on two peering MC-LAG network devices are allowed to talk to each other
- When an MC-LAG network receives a packet from the ICL, the packet is forwarded to all local S-Links and active MC-LAGs that do not exist in the MC-LAG network from which the packet was sent.

In active-standby mode, only one of the MC-LAG peers is active at any given time. The other MC-LAG peer is in backup (standby) mode. The active MC-LAG peer uses Link Aggregation Control Protocol (LACP) to advertise to client devices that its child link is available for forwarding data traffic. Active-standby mode should be used if you are interested in redundancy only. If you require both redundancy and load sharing across member links, use active-active mode.

NOTE: Active-standby mode is not supported on EX4300 and QFX Series switches.

- Status Control

  Specifies whether a node becomes active or goes into standby mode when an ICL failure occurs. If one node is active, the other node must be standby.

BEST PRACTICE: We recommend that you configure **prefer-status-control-active** statement with the mc-ae **status-control active** configuration. Do not configure the **prefer-status-control-active** statement with the mc-ae **status-control standby** configuration.

NOTE: On EX9200 and QFX Series switches, if you configure both nodes as **prefer-status-control-active**, you must also configure ICCP peering using the peer's loopback address to make sure that the ICCP session does not go down because of physical link failures. Additionally, you must configure backup liveness detection on both of the MC-LAG nodes.

NOTE: On EX9200 switches, the **prefer-status-control-active** statement was added in Junos OS Release 13.2R1.

- Events ICCP-Peer-Down Force-ICL-Down

  Forces the ICL to be down if the peer of this node goes down.

- Events ICCP-Peer-Down Prefer-Status-Control-Active

  Allows the LACP system ID to be retained during a reboot, which provides better convergence after a failover.

Multichassis Link Aggregation Group (MC-LAG) Configuration Synchronization

MC-LAG configuration synchronization enables you to easily propagate, synchronize, and commit configurations from one MC-LAG peer to another. You can log into any one of the MC-LAG peers to manage both MC-LAG peers, thus having a single point of management. You can also use configuration groups to simplify the configuration process. You can create one configuration group for the local MC-LAG peer, one for the remote MC-LAG peer, and one for the global configuration, which is essentially a configuration that is common to both MC-LAG peers.

In addition, you can create conditional groups to specify when a configuration is synchronized with another MC-LAG peer. You can enable the **peers-synchronize** statement at the **[edit system commit]** hierarchy to synchronize the configurations and commits across the MC-LAG peers by default. NETCONF over SSH provides a secure connection between the MC-LAG peers, and Secure Copy Protocol (SCP) copies the configurations securely between them.

Multichassis Link Aggregation Group (MC-LAG) Configuration Consistency Check

Configuration consistency check uses the Inter-Chassis Control Protocol (ICCP) to exchange MC-LAG configuration parameters (chassis ID, service ID, and so on) and checks for any configuration inconsistencies across MC-LAG peers. An example of an inconsistency is configuring identical chassis IDs on both peers instead of configuring unique chassis IDs on both peers. When there is an inconsistency, you are notified and can take action to resolve it. Only committed MC-LAG parameters are checked for consistency.

Enhanced Convergence

Starting with Junos OS Release 14.2R3 on MX Series routers, enhanced convergence improves Layer 2 and Layer 3 convergence time when a multichassis aggregated Ethernet (MC-AE) link goes down or comes up in a bridge domain or VLAN. Starting with Junos OS Release 18.1R1, the number of vmembers has increased to 128k, and the number of ARP and ND entries has increased to 96k when enabling the **enhanced-convergence** statement. Starting with Junos OS Release 19.1R1, the number of number of ARP and ND entries has increased to 256,000 when enabling the **enhanced-convergence** and **arp-enhanced-scale** statements. Enhanced convergence improves Layer 2 and Layer 3

convergence time during multichassis aggregated Ethernet (MC-AE) link failures and restoration scenarios

When enhanced convergence is enabled, the MAC address, ARP or ND entries learned over the MC-AE interfaces are programmed in the forwarding table with the MC-AE link as the primary next-hop and with ICL as the backup next-hop. With this enhancement, during an MC-AE link failure or restoration, only the next-hop information in the forwarding table is updated and there is no flushing and relearning of the MAC address, ARP or ND entry. This process improves traffic convergence during MC-AE link failure or restoration because the convergence involves only next-hop repair in the forwarding plane, with the traffic being fast rerouted from the MC-AE link to the ICL.

If you have configured an IRB interface over an MC-AE interface that has enhanced convergences enabled, then you must configure enhanced convergence on the IRB interface as well. Enhanced convergence must be enabled for both Layer 2 and Layer 3 interfaces.

IPv6 Neighbor Discovery Protocol

Neighbor Discovery Protocol (NDP) is an IPv6 protocol that enables nodes on the same link to advertise their existence to their neighbors and to learn about the existence of their neighbors. NDP is built on top of Internet Control Message Protocol version 6 (ICMPv6). It replaces the following IPv4 protocols: Router Discovery (RDISC), Address Resolution Protocol (ARP), and ICMPv4 redirect.

You can use NDP in a multichassis link aggregation group (MC-LAG) active-active configuration on switches.

NDP on MC-LAGs uses the following message types:

- Neighbor solicitation (NS)—Messages used for address resolution and to test reachability of neighbors.

  A host can verify that its address is unique by sending a neighbor solicitation message destined to the new address. If the host receives a neighbor advertisement in reply, the address is a duplicate.

- Neighbor advertisement (NA)—Messages used for address resolution and to test reachability of neighbors. Neighbor advertisements are sent in response to neighbor solicitation messages.

Load Balancing

Load balancing of network traffic between MC-LAG peers is 100 percent local bias. Load balancing of network traffic between multiple LAG members in a local MC-LAG node is achieved through a standard LAG hashing algorithm.

Layer 2 Unicast Features Supported

The following Layer 2 unicast features, learning and aging, are supported:

- Learned MAC addresses are propagated across MC-LAG peers for all of the VLANs that are spawned across the peers.

- Aging of MAC addresses occurs when the MAC address is not seen on both of the peers.
- MAC addresses learned on single-homed links are propagated across all of the VLANs that have MC-LAG links as members.

NOTE: MAC learning is disabled on the ICL. Consequently, source MAC addresses cannot be learned locally on the ICL. However, MAC addresses from a remote MC-LAG node can be installed on the ICL interface. For example, the MAC address for a single-homed client on a remote MC-LAG node can be installed on the ICL interface of the local MC-LAG node.

VLANs

Use the following best practice for configuring VLANs:

BEST PRACTICE: We recommend that you limit the scope of VLANs and configure them only where they are necessary. Configure the MC-AE trunk interfaces with only the VLANs that are necessary for the access layer. This limits the broadcast domain and reduces the STP load on aggregation and access switches.

Layer 2 Multicast Features Supported

The following Layer 2 multicast features, unknown unicast and IGMP snooping, are supported:

- Flooding happens on all links across peers if both peers have virtual LAN membership. Only one of the peers forwards traffic on a given MC-LAG link.
- Known and unknown multicast packets are forwarded across the peers by adding the ICL port as a multicast router port.
- IGMP membership learned on MC-LAG links is propagated across peers.

  You must configure the **multichassis-lag-replicate-state** statement for Internet Group Management Protocol (IGMP) snooping to work properly in an MC-LAG environment.
- During an MC-LAG peer reboot, known multicast traffic is flooded until the IGMP snooping state is synchronized with the peer.

IGMP Snooping on an Active-Active MC-LAG

Internet Group Management Protocol (IGMP) snooping controls multicast traffic in a switched network. When IGMP snooping is not enabled, the Layer 2 device broadcasts multicast traffic out of all of its ports, even if the hosts on the network do not want the multicast traffic. With IGMP snooping enabled, a Layer 2 device monitors the IGMP join and leave messages sent from each connected host to a multicast router. This enables the Layer 2 device to keep track of the multicast groups and associated member ports. The Layer 2 device uses this information to make intelligent decisions and to forward

multicast traffic to only the intended destination hosts. IGMP uses Protocol Independent Multicast (PIM) to route the multicast traffic. PIM uses distribution trees to determine which traffic is forwarded.

NOTE: You must enable Protocol Independent Multicast (PIM) on the IRB interface to avoid multicast duplication.

In an active-active MC-LAG configuration, IGMP snooping replicates the Layer 2 multicast routes so that each MC-LAG peer has the same routes. If a device is connected to an MC-LAG peer by way of a single-homed interface, IGMP snooping replicates the join message to its IGMP snooping peer. If a multicast source is connected to an MC-LAG by way of a Layer 3 device, the Layer 3 device passes this information to the IRB or the routed VLAN interface (RVI) that is configured on the MC-LAG. The first hop designated router is responsible for sending the register and register-stop messages for the multicast group. The last hop designated router is responsible for sending PIM join and leave messages toward the rendezvous point and source for the multicast group. The routing device with the smallest preference metric forwards traffic on transit LANs.

NOTE: You must configure the ICL interface as a router-facing interface (by configuring the multicast-router-interface statement) for multicast forwarding to work in an MC-LAG environment. For the scenario in which traffic arrives by way of a Layer 3 interface, PIM and IGMP must be enabled on the IRB or RVI interface configured on the MC-LAG peers. You must enable PIM on the IRB or RVI interface to avoid multicast duplication.

VRRP Active-Standby Support

The Juniper Networks Junos operating system (Junos OS) supports active-active MC-LAGs by using VRRP in active-standby mode. VRRP in active-standby mode enables Layer 3 routing over the multichassis aggregated Ethernet interfaces on the MC-LAG peers. In this mode, the MC-LAG peers act as virtual routers. The peers share the virtual IP address that corresponds to the default route configured on the host or server connected to the MC-LAG. This virtual IP address (of the IRB or RVI interface) maps to either of the VRRP MAC addresses or to the logical interfaces of the MC-LAG peers. The host or server uses the VRRP MAC address to send any Layer 3 upstream packets. At any time, one of the VRRP devices is the master (active), and the other is a backup (standby). Usually, a VRRP backup node does not forward incoming packets. However, when VRRP over IRB or RVI is configured in an MC-LAG active-active environment, both the VRRP master and the VRRP backup forward Layer 3 traffic arriving on the multichassis aggregated Ethernet interface. If the master fails, all the traffic shifts to the multichassis aggregated Ethernet interface on the backup.

NOTE: You must configure VRRP on both MC-LAG peers for both the active and standby members to accept and route packets. Additionally, you must configure the VRRP backup device to send and receive ARP requests.

Routing protocols run on the primary IP address of the IRB or RVI interface, and both of the MC-LAG peers run routing protocols independently. The routing protocols use the primary IP address of the IRB or RVI interface and the IRB or RVI MAC address to communicate with the MC-LAG peers. The IRB or RVI MAC address of each MC-LAG peer is replicated on the other MC-LAG peer and is installed as a MAC address that has been learned on the ICL.

NOTE: If you are using the VRRP over IRB or RVI method to enable Layer 3 functionality, you must configure static ARP entries for the IRB or RVI interface of the remote MC-LAG peer to allow routing protocols to run over the IRB or RVI interfaces.

MAC Address Management

If an MC-LAG is configured to be active-active, upstream and downstream traffic could go through different MC-LAG peer devices. Because the MAC address is learned only on one of the MC-LAG peers, traffic in the reverse direction could be going through the other MC-LAG peer and flooding the network unnecessarily. Also, a single-homed client's MAC address is learned only on the MC-LAG peer that it is attached to. If a client attached to the peer MC-LAG network device needs to communicate with that single-homed client, then traffic would be flooded on the peer MC-LAG network device. To avoid unnecessary flooding, whenever a MAC address is learned on one of the MC-LAG peers, the address is replicated to the other MC-LAG peer. The following conditions are applied when MAC address replication is performed:

NOTE: Gratuitous ARP requests are not sent when the MAC address on the IRB or RVI interface changes.

- MAC addresses learned on an MC-LAG of one MC-LAG peer must be replicated as learned on the same MC-LAG of the other MC-LAG peer.
- MAC addresses learned on single-homed customer edge (CE) clients of one MC-LAG peer must be replicated as learned on the ICL interface of the other MC-LAG peer.
- MAC address learning on an ICL is disabled from the data path. It depends on software to install MAC addresses replicated through ICCP.

If you have a VLAN without an IRB or RVI configured, MAC address replication will synchronize the MAC addresses.

MAC Aging

MAC aging support in Junos OS extends aggregated Ethernet logic for a specified MC-LAG. A MAC address in software is not deleted until all Packet Forwarding Engines have deleted the MAC address.

Address Resolution Protocol Active-Active MC-LAG Support Methodology

The Address Resolution Protocol (ARP) maps IP addresses to MAC addresses. Junos OS uses ARP response packet snooping to support active-active MC-LAGs, providing easy synchronization without the need to maintain any specific state. Without synchronization, if one MC-LAG peer sends an ARP request, and the other MC-LAG peer receives the response, ARP resolution is not successful. With synchronization, the MC-LAG peers synchronize the ARP resolutions by sniffing the packet at the MC-LAG peer receiving the ARP response and replicating this to the other MC-LAG peer. This ensures that the entries in ARP tables on the MC-LAG peers are consistent.

When one of the MC-LAG peers restarts, the ARP destinations on its MC-LAG peer are synchronized. Because the ARP destinations are already resolved, its MC-LAG peer can forward Layer 3 packets out of the multichassis aggregated Ethernet interface.

NOTE:

- In some cases, ARP messages received by one MC-LAG peer are replicated to the other MC-LAG peer through ICCP. This optimization feature is applicable only for ARP replies, not ARP requests, received by the MC-LAG peers.
- Dynamic ARP resolution over the ICL interface is not supported. Consequently, incoming ARP replies on the ICL are discarded. However, ARP entries can be populated on the ICL interface through ICCP exchanges from a remote MC-LAG peer.
- During graceful Routing Engine switchover (GRES), ARP entries that were learned remotely are purged and then learned again.

DHCP Relay with Option 82

NOTE: DHCP relay is not supported with MAC address synchronization. If DHCP relay is required, configure VRRP over IRB or RVI for Layer 3 functionality.

BEST PRACTICE: In an MC-LAG active-active environment, we recommend that you use the bootp relay agent by configuring the DHCP relay agent with the **forwarding options helpers bootp** command to avoid stale session information issues that might arise for clients when the router is using the extended DHCP relay agent (jdhcp) process.

If your environment only supports IPv6 or you must use the extended DHCP relay agent (jdhcp) process for other reasons, then as a workaround, you can configure forward-only support by using the **forwarding-options dhcp-relay forward-only** command for IPv4 and

the **forwarding-options dhcpv6 forward-only** command for IPv6. You must also verify that your DHCP server in the network supports option 82.

DHCP relay with option 82 provides information about the network location of DHCP clients. The DHCP server uses this information to implement IP addresses or other parameters for the client. With DHCP relay enabled, DHCP request packets might take the path to the DHCP server through either of the MC-LAG peers. Because the MC-LAG peers have different hostnames, chassis MAC addresses, and interface names, you need to observe these requirements when you configure DHCP relay with option 82:

- Use the interface description instead of the interface name.
- Do not use the hostname as part of the circuit ID or remote ID string.
- Do not use the chassis MAC address as part of the remote ID string.
- Do not enable the vendor ID.
- If the ICL interface receives DHCP request packets, the packets are dropped to avoid duplicate packets in the network.

  A counter called *Due to received on ICL interface* has been added to the **show helper statistics** command, which tracks the packets that the ICL interface drops.

  An example of the CLI output follows:

  ```
  user@switch> show helper statistics
  BOOTP:
    Received packets: 6
    Forwarded packets: 0
    Dropped packets: 6
      Due to no interface in DHCP Relay database: 0
      Due to no matching routing instance: 0
      Due to an error during packet read: 0
      Due to an error during packet send: 0
      Due to invalid server address: 0
      Due to no valid local address: 0
      Due to no route to server/client: 0
      Due to received on ICL interface: 6
  ```

  The output shows that six packets received on the ICL interface have been dropped.

MC-LAG Packet Forwarding

To prevent the server from receiving multiple copies from both of the MC-LAG peers, a block mask is used to prevent forwarding of traffic received on the ICL toward the multichassis aggregated Ethernet interface. Preventing forwarding of traffic received on the ICL interface toward the multichassis aggregated Ethernet interface ensures that traffic received on MC-LAG links is not forwarded back to the same link on the other peer. The forwarding block mask for a given MC-LAG link is cleared if all of the local members of the MC-LAG link go down on the peer. To achieve faster convergence, if all local members of the MC-LAG link are down, outbound traffic on the MC-LAG is redirected to the ICL interface on the data plane.

Layer 3 Unicast Feature Support

Layer 3 unicast feature support includes the following:

- Address Resolution Protocol (ARP) synchronization enables ARP resolution on both of the MC-LAG peers.
- DHCP relay with option 82 enables option 82 on the MC-LAG peers. Option 82 provides information about the network location of DHCP clients. The DHCP server uses this information to implement IP addresses or other parameters for the client.

Virtual Router Redundancy Protocol (VRRP) over IRB and MAC Address Synchronization

There are two methods for enabling Layer 3 routing functionality across a multichassis link aggregation group (MC-LAG). You can choose either to configure the Virtual Router Redundancy Protocol (VRRP) over the integrated routing and bridging (IRB) interface or to synchronize the MAC addresses for the Layer 3 interfaces of the switches participating in the MC-LAG.

NOTE: On EX9200 and QFX Series switches, routing protocols are not supported on the downstream clients.

BEST PRACTICE: On EX9200 and QFX Series switches, we recommend that you use MAC address synchronization for the downstream clients. For the upstream routers, we recommend that you use VRRP over IRB or RVI method.

NOTE: On EX9200 and QFX Series switches, you cannot configure both VRRP over IRB and MAC synchronization, because processing MAC addresses might not work.

VRRP over IRB or RVI requires that you configure different IP addresses on IRB or RVI interfaces, and run VRRP over the IRB or RVI interfaces. The virtual IP address is the gateway IP address for the MC-LAG clients.

If you are using the VRRP over IRB method to enable Layer 3 functionality, you must configure static ARP entries for the IRB interface of the remote MC-LAG peer to allow routing protocols to run over the IRB interfaces. This step is required so you can issue the **ping** command to reach both the physical IP addresses and virtual IP addresses of the MC-LAG peers.

For example, you can issue the **set interfaces irb unit 18 family inet address 10.181.18.3/8 arp 10.181.18.2 mac 00:00:5E:00:2f:f0** command.

When you issue the **show interfaces irb** command after you have configured VRRP over IRB, you will see that the static ARP entries are pointing to the IRB MAC addresses of the remote MC-LAG peer:

```
user@switch> show interfaces irb
Physical interface: irb, Enabled, Physical link is Up
  Interface index: 180, SNMP ifIndex: 532
  Type: Ethernet, Link-level type: Ethernet, MTU: 1514
  Device flags   : Present Running
  Interface flags: SNMP-Traps
  Link type      : Full-Duplex
  Link flags     : None
  Current address: 00:00:5E:00:2f:f0, Hardware address: 00:00:5E:00:2f:f0
  Last flapped   : Never
    Input packets : 0
    Output packets: 0
```

NOTE: Use MAC synchronization if you require more than 1,000 VRRP instances.

MAC address synchronization enables MC-LAG peers to forward Layer 3 packets arriving on multichassis aggregated Ethernet interfaces with either their own IRB or RVI MAC address or their peer's IRB or RVI MAC address. Each MC-LAG peer installs its own IRB or RVI MAC address as well as the peer's IRB or RVI MAC address in the hardware. Each MC-LAG peer treats the packet as if it were its own packet. If MAC address synchronization is not enabled, the IRB or RVI MAC address is installed on the MC-LAG peer as if it were learned on the ICL.

NOTE: Here are some caveats with configuring MAC address synchronization:

- Use MAC address synchronization if you are not planning to run routing protocols on the IRB interfaces.

  MAC address synchronization does not support routing protocols on IRB interfaces, and routing protocols are not supported with downstream MC-LAG clients. If you need routing capability, configure both VRRP and routing protocols on each MC-LAG peer. Routing protocols are supported on upstream routers.

- DHCP relay is not supported with MAC address synchronization.

  If you need to configure DHCP relay, configure VRRP over IRB.

- Gratuitous ARP requests are not sent when the MAC address on the IRB interface changes.

MAC address synchronization requires you to configure the same IP address on the IRB interface in the VLAN on both MC-LAG peers. To enable the MAC address synchronization feature using the standard CLI, issue the **set vlan *vlan-name* mcae-mac-synchronize** command on each MC-LAG peer. If you are using the Enhanced Layer 2 CLI, issue the **set**

bridge-domains *name* **mcae-mac-synchronize** command on each MC-LAG peer. Configure the same IP address on both MC-LAG peers. This IP address is used as the default gateway for the MC-LAG servers or hosts.

Protocol Independent Multicast

Protocol Independent Multicast (PIM) and Internet Group Management Protocol (IGMP) provide support for Layer 3 multicast. In addition to the standard mode of PIM operation, there is a special mode called PIM dual designated router. PIM dual designated router minimizes multicast traffic loss in case of failures.

If you are using Layer 3 multicast, configure the IP address on the active MC-LAG peer with a high IP address or a high designated router priority.

NOTE: **PIM dual designated router is not supported on EX9200 and QFX10000 switches.**

PIM operation is discussed in the following sections:

PIM Operation with Normal Mode Designated Router Election

In normal mode with designated router election, the IRB or RVI interfaces on both of the MC-LAG peers are configured with PIM enabled. In this mode, one of the MC-LAG peers becomes the designated router through the PIM designated router election mechanism. The elected designated router maintains the rendezvous-point tree (RPT) and shortest-path tree (SPT) so it can receive data from the source device. The elected designated router participates in periodic PIM join and prune activities toward the rendezvous point or the source.

The trigger for initiating these join and prune activities is the IGMP membership reports that are received from interested receivers. IGMP reports received over multichassis aggregated Ethernet interfaces (potentially hashing on either of the MC-LAG peers) and single-homed links are synchronized to the MC-LAG peer through ICCP.

Both MC-LAG peers receive traffic on their incoming interface (IIF). The non-designated router receives traffic by way of the ICL interface, which acts as a multicast router (mrouter) interface.

If the designated router fails, the non-designated router has to build the entire forwarding tree (RPT and SPT), which can cause multicast traffic loss.

PIM Operation with Dual Designated Router Mode

In dual designated router mode, both of the MC-LAG peers act as designated routers (active and standby) and send periodic join and prune messages upstream toward the rendezvous point, or source, and eventually join the RPT or SPT.

The primary MC-LAG peer forwards the multicast traffic to the receiver devices even if the standby MC-LAG peer has a smaller preference metric.

The standby MC-LAG peer also joins the forwarding tree and receives the multicast data. The standby MC-LAG peer drops the data because it has an empty outgoing interface list (OIL). When the standby MC-LAG peer detects the primary MC-LAG peer failure, it adds the receiver VLAN to the OIL, and starts to forward the multicast traffic.

To enable a multicast dual designated router, issue the **set protocols pim interface interface-name dual-dr** command on the VLAN interfaces of each MC-LAG peer.

Failure Handling

To ensure faster convergence during failures, configure the IP address on the primary MC-LAG peer with a higher IP address or with a higher designated router priority. Doing this ensures that the primary MC-LAG peer retains the designated router membership if PIM peering goes down.

To ensure that traffic converges if an MC-AE interfaces goes down, the ICL-PL interface is always added as an mrouter port. Layer 3 traffic is flooded through the default entry or the snooping entry over the ICL-PL interface, and the traffic is forwarded on the MC-AE interface on the MC-LAG peer. If the ICL-PL interface goes down, PIM neighborship goes down. In this case, both MC-LAG peers become the designated router. The backup MC-LAG peer brings down its links and the routing peering is lost. If the ICCP connection goes down, the backup MC-LAG peer changes the LACP system ID and brings down the MC-AE interfaces. The state of PIM neighbors remains operational.

Miswiring Detection Guidelines

You can use STP to detect miswiring loops within the peer or across MC-LAG peers. An example of miswiring is when a port of a network element is accidentally connected to another port of the same network element. Using STP to detect loops on MC-LAG interfaces, however, is not supported.

NOTE: Do not use Multiple Spanning Tree Protocol (MSTP) or VLAN Spanning Tree Protocol (VSTP). There could be a loop if MSTP or VSTP is enabled in an MC-AE topology without enabling MSTP or VSTP on the MC-AE logical interfaces. Also, there could be a loop if an alternate path exists from access nodes to MC-AE nodes.

BEST PRACTICE:

To detect miswirings, we recommend that you do the following:

- Configure STP globally so that STP can detect local miswiring within and across MC-LAG peers.
- Disable STP on ICL links, however, because STP might block ICL interfaces and disable protection.
- Disable STP on interfaces that are connected to aggregation switches.

- Configure MC-LAG interfaces as edge ports.
- Enable bridge protocol data unit (BPDU) block on edge.
- Do not enable BPDU block on interfaces connected to aggregation switches.

For more information about BPDU block, see *Understanding BPDU Protection for STP, RSTP, and MSTP*.

Reverse Layer 2 Gateway Protocol (RL2GP) for Loop Prevention

With RL2GP, you can configure two edge MC-LAG nodes with the same STP virtual root ID. The virtual root ID must be superior to all bridges in the downstream network, and the downstream bridges must be capable of running STP. STP could block one of the interfaces in the downstream network and break any loop due to miswiring at the core or access layer, or due to a problem in the server software.

RL2GP must be configured on both MC-LAG nodes to prevent loops. Because both MC-LAG nodes would have the same virtual root ID, the MC-LAG interface would always be forwarding traffic. The downstream bridge would receive BPDUs from both nodes and thus receive twice the number of BPDUs on its aggregated Ethernet (AE) interface. If you do not want to receive twice the number of BPDUs, you can double the STP hello time on the virtual ID root. If both of the nodes use the same AE interface name, then the STP port number would be identical and would reduce the STP load on the downstream bridge.

MC-LAG Upgrade

Upgrade the MC-LAG peers according to the following guidelines.

NOTE: Upgrade both MC-LAG nodes to the same software version in order to achieve no loss during stable and failover conditions. The protocol states, data forwarding, and redundancy are guaranteed only after both nodes are upgraded to the same software version successfully.

NOTE: After a reboot, the multichassis aggregated Ethernet interfaces come up immediately and might start receiving packets from the server. If routing protocols are enabled, and the routing adjacencies have not been formed, packets might be dropped.

To prevent this scenario, issue the **set interfaces *interface-name* aggregated-ether-options mc-ae init-delay-time *time*** command to set a time by which the routing adjacencies are formed.

NOTE: On QFX and EX Series switches, the default session establishment hold time is 300 seconds. However, the session establishment time must be at least 100 seconds higher than the init delay time. You can optionally update the session establishment time to be 340 seconds and the init delay time to be 240 seconds.

1. Make sure that both of the MC-LAG peers (node1 and node2) are in the active-active state by using the following command on any one of the MC-LAG peers:

```
user@switch> show interfaces mc-ae id 1
Member Link                  : ae0
 Current State Machine's State: mcae active state
 Local Status                 : active<<<<<<<
 Local State                  : up
 Peer Status                  : active<<<<<<<
 Peer State                   : up
     Logical Interface        : ae0.0
     Topology Type            : bridge
     Local State              : up
     Peer State               : up
     Peer Ip/MCP/State        : 10.1.1.2 ae2.0 up
```

2. Upgrade node1 of the MC-LAG.

   When node1 is upgraded, it is rebooted, and all traffic is sent across the available LAG interfaces of node2, which is still up. The amount of traffic lost depends on how quickly the neighbor devices detect the link loss and rehash the flows of the LAG.

3. Verify that node1 is running the software you just installed by issuing the **show version** command.

4. Make sure that both nodes of the MC-LAG (node1 and node2) are in the active-active state after the reboot of node1.

5. Upgrade node2 of the MC-LAG.

   Repeat Step 1 through Step 3 to upgrade node2.

IGMP Report Synchronization

IGMP reports received over MC-AE interfaces and single-homed links are synchronized to the MC-LAG peers. The MCSNOOPD client application on the MC-LAG peer receives the synchronization packet over ICCP and then sends a copy of the packet to the kernel using the routing socket PKT_INJECT mechanism. When the kernel receives the packet, it sends the packet to the routing protocol process (rpd) enables Layer 3 multicast protocols, like PIM and IGMP, on routed VLAN interfaces (RVIs) configured on MC-LAG VLANs.

Release History Table

Release	Description
19.1R1	Starting with Junos OS Release 19.1R1, the number of number of ARP and ND entries has increased to 256,000 when enabling the **enhanced-convergence** and **arp-enhanced-scale** statements.
18.1R1	Starting with Junos OS Release 18.1R1, the number of vmembers has increased to 128k, and the number of ARP and ND entries has increased to 96k when enabling the **enhanced-convergence** statement.
15.1	Starting with Junos OS Release 15.1 on MX Series routers, configure the backup liveness detection feature to implement faster failover of data traffic during an MC-LAG peer reboot.
14.2R3	Starting with Junos OS Release 14.2R3 on MX Series routers, enhanced convergence improves Layer 2 and Layer 3 convergence time when a multichassis aggregated Ethernet (MC-AE) link goes down or comes up in a bridge domain or VLAN.

CHAPTER 2

Configuring MC-LAG for Providing Redundancy, Load Balancing, and Multihoming Support

Redundancy and Multihoming Using MC-LAG

Configuring Multichassis Link Aggregation on MX Series Routers

Multichassis link aggregation (MC-LAG) enables an MX Series 5G Universal Routing Platform to form a logical LAG interface with two or more other devices. MC-LAG provides additional benefits over traditional LAG in terms of node level redundancy, multihoming support, and a loop-free Layer 2 network without the need to run Spanning Tree Protocol (STP). MC-LAG can be configured for virtual private LAN service (VPLS) routing instances, circuit cross-connect (CCC) applications, and Layer 2 circuit encapsulation types.

The MC-LAG devices use Inter-Chassis Control Protocol (ICCP) to exchange the control information between two MC-LAG network devices.

On one end of the MC-LAG is an MC-LAG client device that has one or more physical links in a link aggregation group (LAG). This client device does not need to be aware of the MC-LAG configuration. On the other side of the MC-LAG are two MC-LAG network devices. Each of these network devices has one or more physical links connected to a single client device. The network devices coordinate with each other to ensure that data traffic is forwarded properly.

MC-LAG includes the following functionality:

- Only single-active MC-LAG mode with multi-homed VPLS instance is supported.
- MC-LAG operates only between two devices.
- Layer 2 circuit functions are supported with **ether-ccc** and **vlan-ccc** encapsulations.
- VPLS functions are supported with **ether-vpls** and **vlan-vpls** encapsulations.

NOTE: Ethernet connectivity fault management (CFM) specified in the IEEE 802.1ag standard for Operation, Administration, and Management (OAM) is *not* supported on MC-LAG interfaces.

To enable MC-LAG, include the **mc-ae** statement at the **[edit interfaces ae*X* aggregated-ether-options]** hierarchy level along with one of the following statements at the **[edit interfaces ae*X*]** hierarchy level: **encapsulation-ethernet-bridge**, **encapsulation ethernet-ccc**, **encapsulation ethernet-vpls**, or **encapsulation-flexible-ethernet-services**. You also need to configure the **lacp**, **admin-key**, and **system-id** statements at the **[edit interfaces ae*X* aggregated-ether-options]** hierarchy level:

NOTE: When you configure the **prefer-status-control-active** statement, you must also configure the **status-control active** statement. If you configure the **status-control standby** statement with the **prefer-status-control-active** statement, the system issues a warning.

To delete an MC-LAG interface from the configuration, issue the **delete interfaces ae*X* aggregated-ether-options mc-ae** command at the **[edit]** hierarchy level in configuration mode:

```
[edit]
user@host# delete interfaces aeX aggregated-ether-options mc-ae
```

Perform the following steps on each switch that is hosting an MC-LAG:

1. Specify the same multichassis aggregated Ethernet identification number for the MC-LAG that the aggregated Ethernet interface belongs to on each switch.

   ```
   [edit interfaces]
   user@host# set aeX aggregated-ether-options mc-ae mc-ae-id mc-ae-id
   ```

   For example:

   ```
   [edit interfaces]
   user@host# set ae1 aggregated-ether-options mc-ae mc-ae-id 3
   ```

2. Specify a unique chassis ID for the MC-LAG that the aggregated Ethernet interface belongs to on each switch.

```
[edit interfaces]
user@host# set aeX aggregated-ether-options chassis-id chassis-id
```

For example:

```
[edit interfaces]
user@host# set ae1 aggregated-ether-options mc-ae chassis-id 0
```

3. Specify the mode of the MC-LAG that the aggregated Ethernet interface belongs to.

NOTE: Only active/active mode is supported for Reverse Layer 2 Gateway Protocol (R-L2GP) at this time.

```
[edit interfaces]
user@host# set aeX aggregated-ether-options mc-ae mode mode
```

For example:

```
[edit interfaces]
user@host# set ae1 aggregated-ether-options mc-ae mode active-active
```

4. Specify whether the aggregated Ethernet interface participating in the MC-LAG is primary or secondary. Primary is **active**, and secondary is **standby**.

NOTE: You must configure status control on both switches hosting the MC-LAG. If one switch is in active mode, the other must be in standby mode.

```
[edit interfaces]
user@host# set aeX aggregated-ether-options mc-ae status-control (active | standby)
```

For example:

```
[edit interfaces]
user@host# set aeX aggregated-ether-options mc-ae status-control (active | standby)
```

5. Configure the MC-LAG interface to improve Layer 2 and Layer 3 convergence time when a multichassis aggregated Ethernet link goes down or comes up in a bridge domain.

```
[edit interfaces]
user@host# set aeX aggregated-ether-options mc-ae enhanced-convergence
```

6. Specify the same LACP system ID on each switch.

```
[edit interfaces]
user@host# set aeX aggregated-ether-options lacp system-id mac-address
```

For example:

```
[edit interfaces]
```

```
user@host# set ae1 aggregated-ether-options lacp system-id 00:01:02:03:04:05
```

7. Specify the same LACP administration key on each switch.

```
[edit interfaces]
user@host# set aeX aggregated-ether-options lacp admin-key number
```

For example:

```
[edit interfaces]
user@host# set ae1 aggregated-ether-options lacp admin-key 3
```

8. Configure ICCP by doing the following on each switch hosting the MC-LAG:

   a. Configure the local IP address to be used by all switches hosting the MC-LAG.

```
[edit protocols]
user@host# set iccp local-ip-addr local-ip-address
```

   For example:

```
[edit protocols]
user@host# set iccp local-ip-addr 10.3.3.1
```

   b. (Optional) Configure the IP address of the router and the time during which an ICCP connection must succeed between the routers hosting the MC-LAG.

NOTE: On QFX and EX Series switches, the default session establishment hold time is 300 seconds. However, the session establishment time must be at least 100 seconds higher than the init delay time. You can optionally update the session establishment time to be 340 seconds and the init delay time to be 240 seconds.

```
[edit protocols]
user@host# set iccp peer peer-ip-address session-establishment-hold-time seconds
```

   For example:

```
[edit protocols]
user@host# set iccp peer 10.3.3.2 session-establishment-hold-time 340
```

   c. (Optional) Configure the IP address to be used for backup liveness detection:

NOTE: By default, backup liveness detection is not enabled. Configure backup liveness detection if you require faster failover of data traffic loss during an MC-LAG peer reboot. Backup liveness detection helps achieve subsecond traffic loss during an MC-LAG peer reboot.

```
[edit protocols]
user@host# set iccp peer peer-ip-address backup-liveness-detection backup-peer-ip
ip-address
```

   For example:

```
[edit protocols]
user@host# set iccp peer 10.3.3.2 backup-liveness-detection backup-peer-ip 10.207.64.232
```

d. Configure the minimum interval at which the router must receive a reply from the other router with which it has established a Bidirectional Forwarding Detection (BFD) session.

NOTE: Configuring the minimum receive interval is required to enable BFD.

```
[edit protocols]
user@host# set iccp peer peer-ip-address liveness-detection minimum-receive-interval
milliseconds
```

For example:

```
[edit protocols]
user@host# set iccp peer 10.3.3.2 liveness-detection minimum-receive-interval 60
```

e. Configure the minimum transmit interval during which a router must receive a reply from a router with which it has established a BFD session.

```
[edit protocols]
user@host# set iccp peer peer-ip-address liveness-detection transmit-interval
minimum-interval milliseconds
```

For example:

```
[edit protocols]
user@host# set iccp peer 10.3.3.2 liveness-detection transmit-interval minimum-interval
60
```

f. Specify the switch service ID.

The switch service ID is used to synchronize applications, IGMP, ARP, and MAC learning across MC-LAG members.

```
[edit switch-options]
user@host# set service-id number
```

For example:

```
[edit switch-options]
user@host# set service-id 1
```

9. Configure a multichassis protection link between the routers.

```
[edit]
user@host# set multi-chassis multi-chassis-protection peer-ip-address interface
interface-name
```

For example:

```
[edit]
user@host# set multi-chassis multi-chassis-protection 10.3.3.1 interface ae0
```

10. Enable RSTP globally on all interfaces.

```
[edit]
user@host# set protocols rstp interface all mode point-to-point
```

11. Disable RSTP on the interchassis control link protection link (ICL-PL) interfaces on both routers.

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

    For example:

    ```
    [edit]
    user@host# set protocols rstp interface ae0.0 disable
    ```

12. Configure the MC-LAG interfaces as edge ports on both routers.

    ```
    user@host# set protocols rstp interface interface-name edge
    ```

    For example:

    ```
    [edit]
    user@host# set protocols rstp interface ae1 edge
    ```

13. Enable BPDU block on all interfaces except for the ICL-PL interfaces on both routers.

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

    For example:

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

Configuring Multichassis Link Aggregation on EX Series Switches

Multichassis link aggregation groups (MC-LAGs) enable a client device to form a logical LAG interface between two MC-LAG peers (for example, EX9200 switches). An MC-LAG provides redundancy and load balancing between the two MC-LAG peers, multihoming support, and a loop-free Layer 2 network without running Spanning Tree Protocol (STP).

On one end of an MC-LAG, there is an MC-LAG client device, such as a server, that has one or more physical links in a link aggregation group (LAG). This client device does not need to have an MC-LAG configured. On the other side of MC-LAG, there are two MC-LAG peers. Each of the MC-LAG peers has one or more physical links connected to a single client device.

The MC-LAG peers use Inter-Chassis Control Protocol (ICCP) to exchange control information and coordinate with each other to ensure that data traffic is forwarded properly.

NOTE: An interface with an already configured IP address cannot form part of the aggregated Ethernet interface or multichassis aggregated Ethernet interface group.

Perform the following steps on each switch that hosts an MC-LAG:

1. Specify the same multichassis aggregated Ethernet identification number for the MC-LAG that the aggregated Ethernet interface belongs to on each switch.

   ```
   [edit interfaces]
   user@switch# set aex aggregated-ether-options mc-ae mc-ae-id number
   ```

   For example:

   ```
   [edit interfaces]
   user@switch# set ae1 aggregated-ether-options mc-ae mc-ae-id 3
   ```

2. Specify a unique chassis ID for the MC-LAG that the aggregated Ethernet interface belongs to on each switch.

   ```
   [edit interfaces]
   user@switch# set aex aggregated-ether-options mc-ae chassis-id number
   ```

   For example:

   ```
   [edit interfaces]
   user@switch# set ae1 aggregated-ether-options mc-ae chassis-id 0
   ```

3. Specify the mode of the MC-LAG the aggregated Ethernet interface belongs to.

   ```
   [edit interfaces]
   user@switch# set aex aggregated-ether-options mc-ae mode mode
   ```

   For example:

```
[edit interfaces]
user@switch# set ae1 aggregated-ether-options mc-ae mode active-active
```

4. Specify whether the aggregated Ethernet interface participating in the MC-LAG is primary or secondary.

   Primary is **active**, and secondary is **standby**.

   NOTE: You must configure status control on both switches that host the MC-LAG. If one switch is in active mode, the other must be in standby mode.

   ```
   [edit interfaces]
   user@switch# set aex aggregated-ether-options mc-ae status-control (active | standby)
   ```

   For example:

   ```
   [edit interfaces]
   user@switch# set ae1 aggregated-ether-options mc-ae status-control active
   ```

   NOTE: If you configure both nodes as **prefer-status-control-active**, you must also configure ICCP peering using the peer's loopback address to make sure that the ICCP session does not go down because of physical link failures. Additionally, you must configure backup liveness detection on both of the MC-LAG nodes.

   NOTE: On EX9200 switches, the **prefer-status-control-active** statement was added in Junos OS Release 13.2R1.

5. Specify the init delay time.

   The init delay time specifies the number of seconds by which to delay bringing up the MC-LAG interface back to the up state when the MC-LAG peer is rebooted. By delaying the bring-up of the interface until after the protocol convergence, you can prevent packet loss during the recovery of failed links and devices.

   NOTE: On QFX and EX Series switches, the default session establishment hold time is 300 seconds. However, the session establishment time must be at least 100 seconds higher than the init delay time. You can optionally update the session establishment time to be 340 seconds and the init delay time to be 240 seconds.

   ```
   [edit interfaces]
   user@switch# set aex aggregated-ether-options mc-ae init-delay-time seconds
   ```

For example:

```
[edit interfaces]
user@switch# set ae0 aggregated-ether-options mc-ae init-delay-time 240
```

6. Specify the same LACP system ID on each switch.

```
[edit interfaces]
user@switch# set ae x aggregated-ether-options lacp system-id mac-address
```

For example:

```
[edit interfaces]
user@switch# set ae1 aggregated-ether-options lacp system-id 00:01:02:03:04:05
```

7. Specify the same LACP administration key on each switch.

```
[edit interfaces]
user@switch# set ae x aggregated-ether-options lacp admin-key number
```

For example:

```
[edit interfaces]
user@switch# set ae1 aggregated-ether-options lacp admin-key 3
```

8. Configure ICCP by performing the following steps on each switch that hosts the MC-LAG:

   a. Configure the local IP address to be used by the switches that host the MC-LAG.

   ```
   [edit protocols]
   user@switch# set iccp local-ip-addr local-ip-address
   ```

   For example:

   ```
   [edit protocols]
   user@switch# set iccp local-ip-addr 10.3.3.1
   ```

   b. (Optional) Configure the IP address of the switch and the time during which an ICCP connection must be established between the switches that host the MC-LAG.

   NOTE: On QFX and EX Series switches, the default session establishment hold time is 300 seconds. However, the session establishment time must be at least 100 seconds higher than the init delay time. You can optionally update the session establishment time to be 340 seconds and the init delay time to be 240 seconds.

   ```
   [edit protocols]
   user@switch# set iccp peer peer-ip-address session-establishment-hold-time seconds
   ```

   For example:

   ```
   [edit protocols]
   user@switch# set iccp peer 10.3.3.2 session-establishment-hold-time 340
   ```

c. (Optional) Configure the backup-liveness-detection statement on the management interface (fxp0) only.

We recommend that you configure the backup liveness detection feature to implement faster failover of data traffic during an MC-LAG peer reboot.

NOTE: On EX9200 switches, the backup-liveness-detection statement was added in Junos OS Release 13.2R1.

NOTE: By default, backup liveness detection is not enabled. Configure backup liveness detection if you require minimal traffic loss during a reboot. Backup liveness detection helps achieve sub-second traffic loss during an MC-LAG reboot.

```
[edit protocols]
user@switch# set iccp peer peer-ip-address backup-liveness-detection backup-peer-ip
ip-address
```

For example:

```
[edit protocols]
user@switch# set iccp peer 10.3.3.2 backup-liveness-detection backup-peer-ip
10.207.64.232
```

d. Configure the minimum interval at which the switch must receive a reply from the other switch with which it has established a Bidirectional Forwarding Detection (BFD) session.

NOTE: Configuring the minimum receive interval is required to enable BFD. We recommend a minimum receive interval value of 1000 seconds.

```
[edit protocols]
user@switch# set iccp peer peer-ip-address liveness-detection minimum-receive-interval
milliseconds
```

For example:

```
[edit protocols]
user@switch# set iccp peer 10.3.3.2 liveness-detection minimum-receive-interval 1000
```

e. Configure the minimum transmit interval during which a switch must receive a reply from a switch with which it has established a BFD session.

```
[edit protocols]
user@switch# set iccp peer peer-ip-address liveness-detection transmit-interval
minimum-interval milliseconds
```

For example:

```
[edit protocols]
```

```
user@switch# set iccp peer 10.3.3.2 liveness-detection transmit-interval minimum-interval
1000
```

9. Specify the switch service ID.

   The switch service ID is used to synchronize applications, IGMP, ARP, and MAC learning across MC-LAG members.

   ```
   [edit switch-options]
   user@switch# set service-id number
   ```

   For example:

   ```
   [edit switch-options]
   user@switch# set service-id 1
   ```

10. Configure a multichassis protection link between the switches.

    ```
    [edit multi-chassis]
    user@switch# set multi-chassis-protection peer-ip-address interface interface-name
    ```

    For example:

    ```
    [edit multi-chassis]
    user@switch# set multi-chassis-protection 10.3.3.1 interface ae0
    ```

See Also

- Configuring MC-LAG on EX9200 Switches in the Core for Campus Networks

Configuring Multichassis Link Aggregation

NOTE: Multichassis link aggregation (MC-LAG) is supported on QFX3500 and QFX3600 standalone switches running the original CLI, and on QFX5100, QFX5200, EX4600, QFX10002, QFX10008, and QFX10016 switches running Enhanced Layer 2 Software.

Multichassis link aggregation groups (MC-LAGs) enable a client device to form a logical LAG interface between two switches. An MC-LAG provides redundancy and load balancing between the two switches, multihoming support, and a loop-free Layer 2 network without running Spanning Tree Protocol (STP).

The MC-LAG switches use Inter-Chassis Control Protocol (ICCP) to exchange the control information between two MC-LAG switches.

NOTE: The ICCP link should be physically separate (out of band) from the data plane traffic.

On one end of an MC-LAG is an MC-LAG client device, such as a server, that has one or more physical links in a link aggregation group (LAG). This client device does not need to detect the MC-LAG. On the other side of MC-LAG are two MC-LAG switches. Each of the switches has one or more physical links connected to a single client device. The switches coordinate with each other to ensure that data traffic is forwarded properly.

NOTE: An interface with an already configured IP address cannot form part of the aggregated Ethernet interface or multichassis aggregated Ethernet interface group.

Perform the following steps on each switch that is hosting an MC-LAG:

1. Specify the same multichassis aggregated Ethernet identification number for the MC-LAG that the aggregated Ethernet interface belongs to on each switch.

   ```
   [edit interfaces]
   user@switch# set aeX aggregated-ether-options mc-ae mc-ae-id number
   ```

   For example:

   ```
   [edit interfaces]
   user@switch# set ae1 aggregated-ether-options mc-ae mc-ae-id 3
   ```

2. Specify a unique chassis ID for the MC-LAG that the aggregated Ethernet interface belongs to on each switch.

   ```
   [edit interfaces]
   user@switch# set aeX aggregated-ether-options mc-ae chassis-id number
   ```

For example:

```
[edit interfaces]
user@switch# set ae1 aggregated-ether-options mc-ae chassis-id 0
```

3. Specify the mode of the MC-LAG the aggregated Ethernet interface belongs to.

```
[edit interfaces]
user@switch# set aeX aggregated-ether-options mc-ae mode mode
```

For example:

```
[edit interfaces]
user@switch# set ae1 aggregated-ether-options mc-ae mode active-active
```

NOTE: By default, backup liveness detection is not enabled. Configure backup liveness detection if you require minimal traffic loss during a reboot. Backup liveness detection helps achieve sub-second traffic loss during an MC-LAG reboot.

4. Specify whether the aggregated Ethernet interface participating in the MC-LAG is primary or secondary.

Primary is **active**, and secondary is **standby**.

NOTE: You must configure status control on both switches hosting the MC-LAG. If one switch is in active mode, the other must be in standby mode.

```
[edit interfaces]
user@switch# set aeX aggregated-ether-options mc-ae status-control (active | standby)
```

For example:

```
[edit interfaces]
user@switch# set ae1 aggregated-ether-options mc-ae status-control active
```

NOTE: If you configure both nodes as **prefer-status-control-active**, you must also configure ICCP peering using the peer's loopback address to make sure that the ICCP session does not go down because of physical link failures. Additionally, you must configure backup liveness detection on both of the MC-LAG nodes.

5. Specify the init delay time.

The init delay time specifies the number of seconds by which to delay bringing up the MC-LAG interface back to the up state when the MC-LAG peer is rebooted. By delaying the bring-up of the interface until after the protocol convergence, you can prevent packet loss during the recovery of failed links and devices.

NOTE: On QFX and EX Series switches, the default session establishment hold time is 300 seconds. However, the session establishment time must be at least 100 seconds higher than the init delay time. You can optionally update the session establishment time to be 340 seconds and the init delay time to be 240 seconds.

```
[edit interfaces]
user@switch# set aex aggregated-ether-options mc-ae init-delay-time seconds
```

For example:

```
[edit interfaces]
user@switch# set ae0 aggregated-ether-options mc-ae init-delay-time 240
```

6. Specify the same LACP system ID on each switch.

```
[edit interfaces]
user@switch# set aeX aggregated-ether-options lacp system-id mac-address
```

For example:

```
[edit interfaces]
user@switch# set ae1 aggregated-ether-options lacp system-id 00:01:02:03:04:05
```

7. Specify the same LACP administration key on each switch.

```
[edit interfaces]
user@switch# set aeX aggregated-ether-options lacp admin-key number
```

For example:

```
[edit interfaces]
user@switch# set ae1 aggregated-ether-options lacp admin-key 3
```

8. Configure ICCP by doing the following on each switch hosting the MC-LAG:

NOTE: The ICCP link should be physically separate (out of band) from the data plane traffic.

a. Configure the local IP address to be used by all switches hosting the MC-LAG.

```
[edit protocols]
user@switch# set iccp local-ip-addr local-ip-address
```

For example:

```
[edit protocols]
user@switch# set iccp local-ip-addr 10.3.3.1
```

b. (Optional) Configure the IP address of the switch and the time during which an ICCP connection must succeed between the switches hosting the MC-LAG.

NOTE: On QFX and EX Series switches, the default session establishment hold time is 300 seconds. However, the session establishment time must be at least 100 seconds higher than the init delay time. You can optionally update the session establishment time to be 340 seconds and the init delay time to be 240 seconds.

```
[edit protocols]
user@switch# set iccp peer peer-ip-address session-establishment-hold-time seconds
```

For example:

```
[edit protocols]
user@switch# set iccp peer 10.3.3.2 session-establishment-hold-time 340
```

c. (Optional) Configure the backup-liveness-detection statement on the management interface (fxp0) only.

NOTE: By default, backup liveness detection is not enabled. Configure backup liveness detection if you require minimal traffic loss during a reboot. Backup liveness detection helps achieve sub-second traffic loss during an MC-LAG reboot.

```
[edit protocols]
user@switch# set iccp peer peer-ip-address backup-liveness-detection backup-peer-ip
ip-address
```

For example:

```
[edit protocols]
user@switch# set iccp peer 10.3.3.2 backup-liveness-detection backup-peer-ip
10.207.64.232
```

d. Configure the minimum interval at which the switch must receive a reply from the other switch with which it has established a Bidirectional Forwarding Detection (BFD) session.

NOTE: Configuring the minimum receive interval is required to enable BFD.

```
[edit protocols]
user@switch# set iccp peer peer-ip-address liveness-detection minimum-receive-interval
milliseconds
```

For example:

```
[edit protocols]
user@switch# set iccp peer 10.3.3.2 liveness-detection minimum-receive-interval 1000
```

e. Configure the minimum transmit interval during which a switch must receive a reply from a switch with which it has established a BFD session.

```
[edit protocols]
user@switch# set iccp peer peer-ip-address liveness-detection transmit-interval
minimum-interval milliseconds
```

For example:

```
[edit protocols]
user@switch# set iccp peer 10.3.3.2 liveness-detection transmit-interval minimum-interval
1000
```

9. Specify the switch service ID.

The switch service ID is used to synchronize applications, IGMP, ARP, and MAC learning across MC-LAG members.

```
[edit switch-options]
user@switch# set service-id number
```

For example:

```
[edit switch-options]
user@switch# set service-id 1
```

10. Configure a multichassis protection link between the switches.

```
[edit]
user@switch# set multi-chassis multi-chassis-protection peer-ip-address interface
interface-name
```

For example:

```
[edit]
user@switch# set multi-chassis multi-chassis-protection 10.3.3.1 interface ae0
```

11. If you are using ELS, configure the **service-id** on both switches.

The **service-id** must be the same number on both switches.

```
[edit switch-options]
user@switch# set service-id number
```

For example:

```
[edit switch-options]
user@switch# set service-id 10
```

12. Configure the MC-LAG interfaces as edge ports on both switches.

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

    For example:

    ```
    [edit protocols rstp]
    user@switch# set interface ae1 edge
    ```

13. Enable BPDU block on all interfaces except for the ICL-PL interfaces on both switches.

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

Forcing MC-LAG Links or Interfaces with Limited LACP Capability to Be Up

In an MC-LAG network, an MC-LAG client link without Link Access Control Protocol (LACP) configuration remains down and cannot be accessed by the MC-LAG switches.

To ensure that the client device with limited LACP capability is up and accessible on the MC-LAG network, configure one of the aggregated Ethernet links or interfaces on a MC-LAG switch to be up by using the **force-up** statement at the appropriate hierarchy level on your device:

- **[edit interfaces *interface-name* aggregated-ether-options lacp]**
- **[edit interfaces *interface-name* ether-options 802.3ad lacp]**

You can configure the *force-up* feature on the MC-LAG switches in either active mode or standby mode. However, in order to prevent duplicate traffic and packet drops, you configure the force-up feature only on one aggregated Ethernet link of the MC-LAG switches . If multiple aggregated Ethernet links are up on the MC-LAG switches with force-up feature configured, then the device selects the link based on the LACP port ID and port priority. The port with the lowest priority is given preference. In case of two ports with the same priority, the one with the lowest port ID is given preference.

NOTE: The force-up option is not supported on QFX10002 switches.

NOTE: On the QFX5100 switch, you can configure the force-up feature in Link Aggregation Control Protocol (LACP) on the MC-LAG switches starting with Junos OS Release 14.1X53-D10.

NOTE:

- If LACP comes up partially in the MC-LAG network—that is, it comes up on one of the MC-LAG switches and does not comes up on other MC-LAG switches—the force-up feature is disabled.

Example: Configuring Multichassis Link Aggregation

NOTE: Multichassis link aggregation (MC-LAG) is supported on QFX3500 and QFX3600 standalone switches running the original CLI and QFX5100 , EX4600, and QFX10002 standalone switches running Enhanced Layer 2 Software. (This example has not been tested on all devices that support MC-LAG. See Feature Explorer for a full listing of devices that support MC-LAG.

This example shows how multichassis link aggregation groups (MC-LAGs) enable a client device to form a logical LAG interface between two switches to provide redundancy and

load balancing between the two switches, multihoming support, and a loop-free Layer 2 network without running Spanning Tree Protocol (STP).

Requirements

This example uses the following hardware and software components:

- Junos OS Release 12.2 or later for the QFX3500 and QFX3600 standalone switches, Junos OS Release 13.2X51-D10 or later for the QFX5100 standalone switches, Junos OS Release 13.2X51-D25 or later for EX4600 switches, or Junos OS Release 15.1X53-D10 or later for QFX10002 standalone switches.
- Two QFX3500 or QFX3600 standalone switches, two QFX5100 standalone switches, two EX4600 switches, or two QFX10002 standalone switches.

Before you configure an MC-LAG, be sure that you understand how to:

- Configure aggregated Ethernet interfaces on a switch. See *Example: Configuring Link Aggregation Between a QFX Series Product and an Aggregation Switch*.
- Configure the Link Aggregation Control Protocol (LACP) on aggregated Ethernet interfaces on a switch. See *Example: Configuring Link Aggregation with LACP Between a QFX Series Product and an Aggregation Switch*.

Overview

In this example, you configure an MC-LAG across two switches, consisting of two aggregated Ethernet interfaces, an interchassis control link-protection link (ICL-PL), multichassis protection link for the ICL-PL, the Inter-Chassis Control Protocol for the peers hosting the MC-LAG, and Layer 3 connectivity between MC-LAG peers. Layer 3 connectivity is required for ICCP.

Topology

The topology used in this example consists of two switches hosting an MC-LAG. The two switches are connected to a server. Figure 2 on page 64 shows the topology used in this example.

Figure 2: Configuring a Multichassis LAG Between Switch A and Switch B

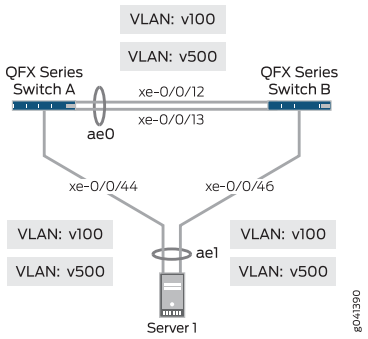

Table 5 on page 64 details the topology used in this configuration example.

Table 5: Components of the Topology for Configuring a Multichassis LAG Between Two Switches

Hostname	Base Hardware	Multichassis Link Aggregation Group
Switch A Switch B	QFX3500 or QFX3600 standalone switch, or QFX5100 standalone switch QFX3500 or QFX3600 standalone switch, or QFX5100 standalone switch QFX3500, QFX3600, EX4600, QFX5100, or QFX10002 standalone switch	ae0 is configured as an aggregated Ethernet interface, and is used as an ICL-PL. The following interfaces are part of ae0: xe-0/0/12 and xe-0/0/13 Switch A and xe-0/0/12 and xe-0/0/13 on Switch B. ae1 is configured as an MC-LAG, and the following two interfaces are part of ae1: xe-0/0/44 on Switch A and xe-0/0/46 on Switch B. .

Configuration

CLI Quick Configuration

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

NOTE: This example shows how to configure MC-LAG using both the original CLI and Enhanced Layer 2 Software (ELS).

In ELS, there are three statements and one additional statement that are different from the original CLI:

- The port-mode statement in the [edit interfaces *interface-name* unit *number* family ethernet-switching] hierarchy is not supported. Use the interface-mode statement instead.
- The vlan statement in the [edit interfaces *interface-name*] hierarchy is not supported. Use the irb statement instead.
- The vlan.*logical-interface-number* option in the [edit vlans *vlan-name* l3-interface] hierarchy is not supported. Use the irb.*logical-interface-number* option instead.
- The service-id statement in the [edit switch-options] hierarchy is required in the ELS CLI.

Switch A—Original CLI

```
set chassis aggregated-devices ethernet device-count 2
set interfaces xe-0/0/12 ether-options 802.3ad ae0
set interfaces xe-0/0/13 ether-options 802.3ad ae0
set interfaces xe-0/0/44 ether-options 802.3ad ae1
set interfaces ae0 unit 0 family ethernet-switching port-mode trunk
set interfaces ae0 unit 0 family ethernet-switching vlan members v500
set interfaces ae0 unit 0 family ethernet-switching vlan members v100
set interfaces ae1 aggregated-ether-options lacp active
set interfaces ae1 aggregated-ether-options lacp system-id 00:01:02:03:04:05
set interfaces ae1 aggregated-ether-options lacp admin-key 3
set interfaces ae1 aggregated-ether-options mc-ae mc-ae-id 3
set interfaces ae1 aggregated-ether-options mc-ae chassis-id 0
set interfaces ae1 aggregated-ether-options mc-ae mode active-active
set interfaces ae1 aggregated-ether-options mc-ae status-control active
set interfaces ae1 aggregated-ether-options mc-ae init-delay-time 240
set interfaces ae1 unit 0 family ethernet-switching port-mode trunk
set interfaces ae1 unit 0 family ethernet-switching vlan members v100
set interfaces ae1 unit 0 family ethernet-switching vlan members v500
set interfaces vlan unit 500 family inet address 10.3.3.2/8
set vlans v100 vlan-id 100
set vlans v500 vlan-id 500
set vlans v100 l3-interface vlan.100
set vlans v500 l3-interface vlan.500
set protocols iccp local-ip-addr 10.3.3.2
set protocols iccp peer 10.3.3.1 session-establishment-hold-time 340
set protocols iccp peer 10.3.3.1 backup-liveness-detection backup-peer-ip 10.207.64.233
set protocols iccp peer 10.3.3.1 liveness-detection minimum-receive-interval 1000
set protocols iccp peer 10.3.3.1 liveness-detection transmit-interval minimum-interval 1000
set protocol rstp system-identifier 00:01:02:03:04:05
set protocols rstp interface ae0 disable
set protocols rstp interface ae1 edge
set protocols rstp interface all mode point-to-point
```

```
set protocols rstp bpdu-block-on-edge
set multi-chassis multi-chassis-protection 10.3.31 interface ae0
set switch-options service-id 10
```

Switch A—ELS

```
set chassis aggregated-devices ethernet device-count 2
set interfaces xe-0/0/12 ether-options 802.3ad ae0
set interfaces xe-0/0/13 ether-options 802.3ad ae0
set interfaces xe-0/0/44 ether-options 802.3ad ae1
set interfaces ae0 unit 0 family ethernet-switching interface-mode trunk
set interfaces ae0 unit 0 family ethernet-switching vlan members v500
set interfaces ae0 unit 0 family ethernet-switching vlan members v100
set interfaces ae1 aggregated-ether-options lacp active
set interfaces ae1 aggregated-ether-options lacp system-id 00:01:02:03:04:05
set interfaces ae1 aggregated-ether-options lacp admin-key 3
set interfaces ae1 aggregated-ether-options mc-ae mc-ae-id 3
set interfaces ae1 aggregated-ether-options mc-ae chassis-id 0
set interfaces ae1 aggregated-ether-options mc-ae mode active-active
set interfaces ae1 aggregated-ether-options mc-ae status-control active
set interfaces ae1 aggregated-ether-options mc-ae init-delay-time 240
set interfaces ae1 unit 0 family ethernet-switching interface-mode trunk
set interfaces ae1 unit 0 family ethernet-switching vlan members v100
set interfaces ae1 unit 0 family ethernet-switching vlan members v500
set interfaces irb unit 500 family inet address 10.3.3.2/8
set vlans v100 vlan-id 100
set vlans v500 vlan-id 500
set vlans v100 l3-interface irb.100
set vlans v500 l3-interface irb.500
set protocols iccp local-ip-addr 10.3.3.2
set protocols iccp peer 10.3.3.1 session-establishment-hold-time 340
set protocols iccp peer 10.3.3.1 backup-liveness-detection backup-peer-ip 10.207.64.233
set protocols iccp peer 10.3.3.1 liveness-detection minimum-receive-interval 1000
set protocols iccp peer 10.3.3.1 liveness-detection transmit-interval minimum-interval 1000
set protocol rstp system-identifier 00:01:02:03:04:05
set protocols rstp interface ae1 edge
set protocols rstp interface ae1 mode point-to-point
set protocols rstp bpdu-block-on-edge
set multi-chassis multi-chassis-protection 10.3.3.1 interface ae0
set switch-options service-id 10
```

Switch B—Original CLI

```
set chassis aggregated-devices ethernet device-count 2
set interfaces xe-0/0/12 ether-options 802.3ad ae0
set interfaces xe-0/0/13 ether-options 802.3ad ae0
set interfaces xe-0/0/46 ether-options 802.3ad ae1
set interfaces ae0 unit 0 family ethernet-switching port-mode trunk
set interfaces ae0 unit 0 family ethernet-switching vlan members v500
set interfaces ae0 unit 0 family ethernet-switching vlan members v100
set interfaces ae1 aggregated-ether-options lacp active
set interfaces ae1 aggregated-ether-options lacp system-id 00:01:02:03:04:05
set interfaces ae1 aggregated-ether-options lacp admin-key 3
set interfaces ae1 aggregated-ether-options mc-ae mc-ae-id 3
set interfaces ae1 aggregated-ether-options mc-ae chassis-id 1
set interfaces ae1 aggregated-ether-options mc-ae mode active-active
set interfaces ae1 aggregated-ether-options mc-ae status-control standby
set interfaces ae1 aggregated-ether-options mc-ae init-delay-time 240
set interfaces ae1 unit 0 family ethernet-switching port-mode trunk
set interfaces ae1 unit 0 family ethernet-switching vlan members v100
```

```
set interfaces ae1 unit 0 family ethernet-switching vlan members v500
set interfaces vlan unit 500 family inet address 10.3.3.1/8
set vlans v100 vlan-id 100
set vlans v500 vlan-id 500
set vlans v100 l3-interface vlan.100
set vlans v500 l3-interface vlan.500
set protocols iccp local-ip-addr 10.3.3.1
set protocols iccp peer 10.3.3.2 session-establishment-hold-time 340
set protocols iccp peer 10.3.3.2 backup-liveness-detection backup-peer-ip 10.207.64.234
set protocols iccp peer 10.3.3.2 liveness-detection minimum-receive-interval 1000
set protocols iccp peer 10.3.3.2 liveness-detection transmit-interval minimum-interval 1000
set protocol rstp system-identifier 00:01:02:03:04:05
set protocols rstp interface ae0 disable
set protocols rstp interface ae1 edge
set protocols rstp interface all mode point-to-point
set protocols rstp bpdu-block-on-edge
set multi-chassis multi-chassis-protection 10.3.3.2 interface ae0
set switch-options service-id 10
```

Switch B—ELS

```
set chassis aggregated-devices ethernet device-count 2
set interfaces xe-0/0/12 ether-options 802.3ad ae0
set interfaces xe-0/0/13 ether-options 802.3ad ae0
set interfaces xe-0/0/46 ether-options 802.3ad ae1
set interfaces ae0 unit 0 family ethernet-switching interface-mode trunk
set interfaces ae0 unit 0 family ethernet-switching vlan members v500
set interfaces ae0 unit 0 family ethernet-switching vlan members v100
set interfaces ae1 aggregated-ether-options lacp active
set interfaces ae1 aggregated-ether-options lacp system-id 00:01:02:03:04:05
set interfaces ae1 aggregated-ether-options lacp admin-key 3
set interfaces ae1 aggregated-ether-options mc-ae mc-ae-id 3
set interfaces ae1 aggregated-ether-options mc-ae chassis-id 1
set interfaces ae1 aggregated-ether-options mc-ae mode active-active
set interfaces ae1 aggregated-ether-options mc-ae status-control standby
set interfaces ae1 aggregated-ether-options mc-ae init-delay-time 240
set interfaces ae1 unit 0 family ethernet-switching interface-mode trunk
set interfaces ae1 unit 0 family ethernet-switching vlan members v100
set interfaces ae1 unit 0 family ethernet-switching vlan members v500
set interfaces irb unit 500 family inet address 10.3.3.1/8
set vlans v100 vlan-id 100
set vlans v500 vlan-id 500
set vlans v100 l3-interface irb.100
set vlans v500 l3-interface irb.500
set protocols iccp local-ip-addr 10.3.3.1
set protocols iccp peer 10.3.3.2 session-establishment-hold-time 340
set protocols iccp peer 10.3.3.2 backup-liveness-detection backup-peer-ip 10.207.64.234
set protocols iccp peer 10.3.3.2 liveness-detection minimum-receive-interval 1000
set protocols iccp peer 10.3.3.2 liveness-detection transmit-interval minimum-interval 1000
set protocol rstp system-identifier 00:01:02:03:04:05
set protocols rstp interface ae1 edge
set protocols rstp interface ae1 mode point-to-point
set protocols rstp bpdu-block-on-edge
set multi-chassis multi-chassis-protection 10.3.3.2 interface ae0
set switch-options service-id 10
```

Configuring MC-LAG on Two Switches

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 enable multichassis protection link between MC-LAG peers:

1. Configure the number of LAGs on both Switch A and Switch B.

```
[edit chassis]
  user@switch# set aggregated-devices ethernet device-count 2
```

2. Add member interfaces to the aggregated Ethernet interfaces on both Switch A and Switch B.

```
[edit interfaces]
  user@switch# set xe-0/0/12 ether-options 802.3ad ae0
[edit interfaces]
  user@switch# set xe-0/0/13 ether-options 802.3ad ae0
```

```
[edit interfaces]
  user@switch# set xe-0/0/44 ether-options 802.3ad ae1
```

```
[edit interfaces]
  user@switch# set xe-0/0/46 ether-options 802.3ad ae1
```

3. Configure a trunk interface between Switch A and Switch B.

NOTE: The port-mode statement is not supported on Enhanced Layer 2 Software (ELS). If you are running ELS, use the interface-mode statement.

```
[edit interfaces]
  user@switch# set ae0 unit 0 family ethernet-switching port-mode trunk
```

or

```
[edit interfaces]
  user@switch# set ae0 unit 0 family ethernet-switching interface-mode trunk
```

4. Configure a multichassis protection link between Switch A and Switch B.

```
[edit]
  user@switch# set multichassis multi-chassis-protection 10.3.3.1 interface ae0
```

```
[edit]
  user@switch# set multichassis multi-chassis-protection 10.3.3.2 interface ae0
```

Step-by-Step Procedure

To enable ICCP:

1. Configure the local IP address to be in the ICCP connection on Switch A and Switch B.

```
[edit protocols]
  user@switch# set iccp local-ip-addr 10.3.3.2
```

```
[edit protocols]
  user@switch# set iccp local-ip-addr 10.3.3.1
```

2. Configure the peer IP address and minimum receive interval for a BFD session for ICCP on Switch A and Switch B.

NOTE: Configure at least 1000 ms as the minimum receive interval.

```
[edit protocols]
  user@switch# set iccp peer 10.3.3.1 liveness-detection minimum-receive-interval
  1000
```

```
[edit protocols]
  user@switch# set iccp peer 10.3.3.2 liveness-detection minimum-receive-interval
  1000
```

3. Configure the peer IP address and minimum transmit interval for BFD session for ICCP on Switch A and Switch B.

NOTE: Configure at least 1000 ms as the transmit interval minimum interval.

```
[edit protocols]
  user@switch# set iccp peer 10.3.3.1 liveness-detection transmit-interval
  minimum-interval 1000
```

```
[edit protocols]
  user@switch# set iccp peer 10.3.3.2 liveness-detection transmit-interval
  minimum-interval 1000
```

4. (Optional) Configure the time during which an ICCP connection must succeed between MC-LAG peers on Switch A and Switch B.

NOTE: On QFX and EX Series switches, the default session establishment hold time is 300 seconds. However, the session establishment time must be at least 100 seconds higher than the init delay time. You can optionally update the session establishment time to be 340 seconds and the init delay time to be 240 seconds.

```
[edit protocols]
  user@switch# set iccp peer 10.3.3.1 session-establishment-hold-time 340
```

```
[edit protocols]
  user@switch# set iccp peer 10.3.3.2 session-establishment-hold-time 340
```

5. (Optional) Configure the backup IP address to be used for backup liveness detection on both Switch A and Switch B.

NOTE: By default, backup liveness detection is not enabled. Configuring a backup IP address helps achieve sub-second traffic loss during an MC-LAG peer reboot.

```
[edit protocols]
  user@switch# set iccp peer 10.3.3.1 backup-liveness-detection backup-peer-ip
  10.207.64.233
```

```
[edit protocols]
  user@switch# set iccp peer 10.3.3.2 backup-liveness-detection backup-peer-ip
  10.207.64.234
```

6. Configure Layer 3 connectivity between the MC-LAG peers on both Switch A and Switch B.

```
[edit vlans]
  user@switch# set v100 l3-interface 100
[edit vlans]
  user@switch# set v500 l3-interface 500
```

```
[edit vlans]
  user@switch# set v100 l3-interface vlan.100
```

```
[edit vlans]
  user@switch# set v500 l3-interface vlan.500
```

```
[edit vlans]
  user@switch# set v100 l3-interface irb.100
```

```
[edit vlans]
  user@switch# set v500 l3-interface irb.500
```

NOTE: The port-mode statement is not supported on Enhanced Layer 2 Software (ELS). If you are running ELS, use the interface-mode statement.

Step-by-Step Procedure

To enable the MC-LAG interface:

1. Enable LACP on the MC-LAG interface on Switch A and Switch B.

NOTE: At least one end needs to be active. The other end can be either active or passive.

```
[edit interfaces]
  user@switch# set ae1 aggregated-ether-options lacp active
```

2. Specify the same multichassis aggregated Ethernet identification number on both MC-LAG peers on Switch A and Switch B.

```
[edit interfaces]
  user@switch# set ae1 aggregated-ether-options mc-ae mc-ae-id 3
```

3. Specify the same service ID on Switch A and Switch B.

```
[edit]
  user@switch# set switch-options service-id 10
```

4. Specify a unique chassis ID for the MC-LAG on the MC-LAG peers on Switch A and Switch B.

```
[edit interfaces]
  user@switch# set ae1 aggregated-ether-options mc-ae chassis-id 0
```

```
[edit interfaces]
  user@switch# set ae1 aggregated-ether-options mc-ae chassis-id 1
```

5. Specify the operating mode of the MC-LAG on both Switch A and Switch B.

NOTE: Only active-active mode is supported at this time.

```
[edit interfaces]
  user@switch# set ae1 aggregated-ether-options mc-ae mode active-active
```

6. Specify the status control for MC-LAG on Switch A and Switch B.

NOTE: You must configure status control on both Switch A and Switch B hosting the MC-LAG. If one peer is in active mode, the other must be in standby mode.

```
[edit interfaces]
  user@switch# set ae1 aggregated-ether-options mc-ae status-control active
```

```
[edit interfaces]
  user@switch# set ae1 aggregated-ether-options mc-ae status-control standby
```

7. Specify the number of seconds by which the bring-up of the multichassis aggregated Ethernet interface should be deferred after you reboot Switch A and Switch B.

NOTE: The recommended value for maximum VLAN configuration (for example, 4,000 VLANS) is 240 seconds. If IGMP snooping is enabled on all of the VLANs, the recommended value is 420 seconds.

```
[edit interfaces]
  user@switch# set ae1 aggregated-ether-options mc-ae init-delay-time 240
```

8. Specify the same LACP system ID for the MC-LAG on Switch A and Switch B.

```
[edit interfaces]
  user@switch# set ae1 aggregated-ether-options lacp system-ID 00:01:02:03:04:05
```

9. Specify the same LACP administration key on both Switch A and Switch B.

```
[edit interfaces]
  user@switch# set ae1 aggregated-ether-options lacp admin-key 3
```

10. Enable VLANs on the MC-LAG on Switch A and Switch B.

NOTE: The port-mode statement is not supported on Enhanced Layer 2 Software (ELS). If you are running ELS, use the interface-mode statement.

```
[edit interfaces]
  user@switch# set ae0 unit 0 family ethernet-switching port-mode trunk
```

```
[edit interfaces]
  user@switch# set ae0 unit 0 family ethernet-switching vlan members v100
```

```
[edit interfaces]
  user@switch# set ae0 unit 0 family ethernet-switching vlan members v500
```

or

```
[edit interfaces]
  user@switch# set ae0 unit 0 family ethernet-switching interface-mode trunk
```

```
[edit interfaces]
  user@switch# set ae0 unit 0 family ethernet-switching vlan members v100
```

```
[edit interfaces]
  user@switch# set ae0 unit 0 family ethernet-switching vlan members v500
```

Step-by-Step Procedure

To enable R2LGP:

NOTE: To process R2LGP requests, the mc-ae mode must be set to active-active.

1. Configure the RSTP system identifier on Switch A and Switch B:

   Configure the RSTP system identifier on Switch A and Switch B.

```
[edit protocols]
  user@switch# set rstp system-identifier 00:01:02:03:04:05
```

Step-by-Step Procedure

To enable RSTP:

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*.

1. Enable RSTP globally on all interfaces on Switch A and Switch B.

NOTE: The all option is not available on ELS, so you cannot issue this command on ELS.

```
[edit]
  user@switch# set protocols rstp interface all mode point-to-point
```

```
[edit]
  user@switch# set protocols rstp interface ae1 mode point-to-point
```

2. Disable RSTP on the ICL-PL interfaces on Switch A and Switch B:

NOTE: This command is not needed on ELS.

```
[edit]
  user@switch# set protocols rstp interface ae0 disable
```

3. Configure the MC-LAG interfaces as edge ports on Switch A and Switch B.

NOTE: The ae1 interface is a downstream interface. This is why RSTP and bpdu-block-on-edge need to be configured. MC LAG switches are usually configured as root bridge. When downstream switches send superior BPDUs to the MC LAG switches, the MC LAG interfaces will be set as blocked by the downstream switches. The default behavior for the original CLI is to drop superior BPDUs.

```
[edit]
  user@switch# set protocols rstp interface ae1 edge
```

4. Enable BPDU blocking on all interfaces except for the ICL-PL interfaces on Switch A and Switch B.

NOTE: The ae1 interface is a downstream interface. This is why RSTP and bpdu-block-on-edge need to be configured.

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

Results

From configuration mode, confirm your configuration by entering the **show chassis**, **show interfaces**, **show protocols**, **show multi-chassis**, **show switch-options**, and **show vlans** commands. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

Switch A—Original CLI

```
user@SwitchA# show chassis
aggregated-devices {
  ethernet {
    device-count 2;
  }
}
```

```
user@SwitchA# show interfaces
xe-0/0/12 {
  ether-options {
    802.3ad ae0;
  }
}
xe-0/0/13 {
  ether-options {
    802.3ad ae0;
  }
}
xe-0/0/44 {
  ether-options {
    802.3ad ae1;
  }
}
ae0 {
  unit 0 {
    family ethernet-switching {
      port-mode trunk;
      vlan {
        members v500;
      }
    }
  }
}
ae1 {
  aggregated-ether-options {
    lacp {
      active;
      system-id 00:01:02:03:04:05;
      admin-key 3;
    }
    mc-ae {
      mc-ae-id 3;
      chassis-id 0;
      mode active-active;
```

```
        status-control active;
        init-delay-time 240;
      }
    }
    unit 0 {
      family ethernet-switching {
        port-mode trunk;
        vlan {
          members v100;
          members v500;
        }
      }
    }
}
vlan {
    unit 500 {
      family inet {
        address 10.3.3.2/8;
      }
    }
}
```

```
user@SwitchA# show protocols
iccp {
    local-ip-addr 10.3.3.2;
    peer 10.3.3.1 {
      session-establishment-hold-time 340;
      backup-liveness-detection {
        backup-peer-ip 10.207.64.233;
      }
      liveness-detection {
        minimum-receive-interval 1000;
        transmit-interval {
          minimum-interval 1000;
        }
      }
    }
}
rstp {
    system-identifier 00:01:02:03:04:05;
    interface ae0 {
      disable;
    }
    interface ae1 {
      edge;
    }
    interface all {
      mode point-to-point;
    }
    bpdu-block-on-edge;
}
```

```
user@SwitchA# show multi-chassis
multi-chassis-protection 10.3.3.1 {
```

```
    interface ae0;
}
```

```
user@SwitchA# show switch-options
service-id 10;
```

```
user@SwitchA# show vlans
v100 {
    vlan-id 100;
    l3-interface vlan.100;
}
v500 {
    vlan-id 500;
    l3-interface vlan.500;
}
```

Switch A-ELS

```
user@SwitchA# show chassis
aggregated-devices {
    ethernet {
        device-count 2;
    }
}
```

```
user@SwitchA# show interfaces
xe-0/0/12 {
    ether-options {
        802.3ad ae0;
    }
}
xe-0/0/13 {
    ether-options {
        802.3ad ae0;
    }
}
xe-0/0/44 {
    ether-options {
        802.3ad ae1;
    }
}
ae0 {
    unit 0 {
        family ethernet-switching {
            interface-mode trunk;
            vlan {
                members v100;
                members v500;
            }
        }
    }
}
ae1 {
    aggregated-ether-options {
```

```
      lacp {
        active;
        system-id 00:01:02:03:04:05;
        admin-key 3;
      }
      mc-ae {
        mc-ae-id 3;
        chassis-id 0;
        mode active-active;
        status-control active;
        init-delay-time 240
      }
    }
    unit 0 {
      family ethernet-switching {
        interface-mode trunk;
        vlan {
          members v100;
          members v500;
        }
      }
    }
}
irb {
    unit 500 {
      family inet {
        address 10.3.3.2/8;
      }
    }
}
```

```
user@SwitchA# show protocols
iccp {
    local-ip-addr 10.3.3.2;
    peer 10.3.3.1 {
      session-establishment-hold-time 340;
      backup-liveness-detection {
        backup-peer-ip 10.207.64.233;
      }
      liveness-detection {
        minimum-receive-interval 1000;
        transmit-interval {
          minimum-interval 1000;
        }
      }
    }
}
rstp {
    system-identifier 00:01:02:03:04:05;
      interface ae1 {
      edge;
    }
    mode point-to-point;
    }
```

```
  bpdu-block-on-edge;
  }
```

```
user@SwitchA# show multi-chassis
multi-chassis-protection 10.3.3.1 {
  interface ae0;
}
```

```
user@SwitchA# show switch-options
service-id 10;
```

```
user@SwitchA# show vlans
v100 {
  vlan-id 100;
  l3-interface irb.100;
}
v500 {
  vlan-id 500;
  l3-interface irb.500;
}
```

Switch B—Original CLI

```
user@SwitchB# show chassis
aggregated-devices {
  ethernet {
    device-count 2;
  }
}
```

```
user@SwitchB# show interfaces
xe-0/0/12 {
  ether-options {
    802.3ad ae0;
  }
}
xe-0/0/13 {
  ether-options {
    802.3ad ae0;
  }
}
xe-0/0/46 {
  ether-options {
    802.3ad ae1;
  }
}
ae0 {
  unit 0 {
    family ethernet-switching {
      port-mode trunk;
      vlan {
        members v100;
        members v500;
```

```
        }
      }
    }
}
ae1 {
    aggregated-ether-options {
      lacp {
        active;
        system-id 00:01:02:03:04:05;
        admin-key 3;
      }
      mc-ae {
        mc-ae-id 3;
        chassis-id 1;
        mode active-active;
        status-control standby;
        init-delay-time 240;
      }
    }
    unit 0 {
      family ethernet-switching {
        port-mode trunk;
        vlan {
          members v100;
          members v500;
        }
      }
    }
}
vlan {
    unit 500 {
      family inet {
        address 10.3.3.1/8;
      }
    }
}
```

```
user@SwitchB# show protocols
iccp {
    local-ip-addr 10.3.3.1;
    peer 10.3.3.2 {
      session-establishment-hold-time 340;
      backup-liveness-detection {
        backup-peer-ip 10.207.64.234;
      }
      liveness-detection {
        minimum-receive-interval 1000;
        transmit-interval {
          minimum-interval 1000;
        }
      }
    }
}
rstp {
```

```
    system-identifier 00:01:02:03:04:05;
      interface ae0 {
      disable;
    }
    interface ae1 {
      edge;
    }
    interface all {
      mode point-to-point;
    }
    bpdu-block-on-edge;
}
```

```
user@SwitchB# show multi-chassis
multi-chassis-protection 10.3.3.2 {
    interface ae0;
}
```

```
user@SwitchB# show switch-options
service-id 10;
```

```
user@SwitchB# show vlans
v100 {
    vlan-id 100;
    l3-interface vlan.100;
}
v500 {
    vlan-id 500;
    l3-interface vlan.500;
}
```

Switch B—ELS

```
user@SwitchB# show chassis
aggregated-devices {
    ethernet {
        device-count 2;
    }
}
```

```
user@SwitchB# show interfaces
xe-0/0/12 {
    ether-options {
        802.3ad ae0;
    }
}
xe-0/0/13 {
    ether-options {
        802.3ad ae0;
    }
}
xe-0/0/46 {
    ether-options {
```

```
      802.3ad ae1;
    }
}
ae0 {
    unit 0 {
        family ethernet-switching {
            interface-mode trunk;
            vlan {
                members v100;
                members v500;
            }
        }
    }
}
ae1 {
    aggregated-ether-options {
        lacp {
            active;
            system-id 00:01:02:03:04:05;
            admin-key 3;
        }
        mc-ae {
            mc-ae-id 3;
            chassis-id 1;
            mode active-active;
            status-control standby;
            init-delay-time 240;
        }
    }
    unit 0 {
        family ethernet-switching {
            interface-mode trunk;
            vlan {
                members v100;
                members v500;
            }
        }
    }
}
irb {
    unit 500 {
        family inet {
            address 10.3.3.1/8;
        }
    }
}
```

```
user@SwitchB# show protocols
iccp {
    local-ip-addr 10.3.3.1;
    peer 10.3.3.2 {
        session-establishment-hold-time 340;
        backup-liveness-detection {
            backup-peer-ip 10.207.64.234;
```

```
    }
    liveness-detection {
      minimum-receive-interval 1000;
      transmit-interval {
        minimum-interval 1000;
      }
    }
  }
}
rstp {
  system-identifier 00:01:02:03:04:05;
    interface ae1 {
    edge;
  }
  mode point-to-point;
  }
  bpdu-block-on-edge;
  }
```

```
user@SwitchB# show multi-chassis
multi-chassis-protection 10.3.3.2 {
  interface ae0;
}
```

```
user@SwitchB# show switch-options
service-id 10;
```

```
user@SwitchB# show vlans
v100 {
  vlan-id 100;
  l3-interface irb.100;
}
v500 {
  vlan-id 500;
  l3-interface irb.500;
}
```

Verification

Verify that the configuration is working properly.

- Verifying That ICCP Is Working on Switch A on page 84
- Verifying That ICCP Is Working on Switch B on page 84
- Verifying That LACP Is Active on Switch A on page 84
- Verifying That LACP Is Active on Switch B on page 85
- Verifying That the multichassis aggregated Ethernet and ICL-PL Interfaces Are Up on Switch A on page 85
- Verifying That the multichassis aggregated Ethernet and ICL-PL Interfaces Are Up on Switch B on page 86

Verifying That ICCP Is Working on Switch A

Purpose Verify that ICCP is running on Switch A.

Action

```
[edit]
user@switch> show iccp
Redundancy Group Information for peer 10.3.3.1
  TCP Connection       : Established
  Liveliness Detection : Up

 Client Application: MCSNOOPD

 Client Application: eswd
```

Meaning This output shows that the TCP connection between the peers hosting the MC-LAG is up, liveness detection is up, and MCSNOOPD and ESWD client applications are running.

Verifying That ICCP Is Working on Switch B

Purpose Verify that ICCP is running on Switch B.

Action show iccp

```
[edit]
user@switch> show iccp
Redundancy Group Information for peer 10.3.3.2
  TCP Connection       : Established
  Liveliness Detection : Up

 Client Application: MCSNOOPD

 Client Application: eswd
```

Meaning This output shows that the TCP connection between the peers hosting the MC-LAG is up, liveness detection is up, and MCSNOOPD and ESWD client applications are running.

Verifying That LACP Is Active on Switch A

Purpose Verify that LACP is active on Switch A.

Action

```
[edit]
user@switch> show lacp interfaces
```

```
Aggregated interface: ae1
    LACP state:       Role   Exp   Def  Dist  Col  Syn  Aggr  Timeout  Activity
      xe-0/0/46      Actor    No    No   Yes  Yes  Yes   Yes     Fast    Active
      xe-0/0/46    Partner    No    No   Yes  Yes  Yes   Yes     Fast    Active
    LACP protocol:        Receive State  Transmit State          Mux State
      xe-0/0/46                 Current   Fast periodic Collecting distributing
```

Meaning This output shows that Switch A is participating in LACP negotiation.

Verifying That LACP Is Active on Switch B

Purpose Verify that LACP is active on Switch B

Action

```
[edit]
user@switch> show lacp interfaces
```

```
Aggregated interface: ae1
    LACP state:       Role   Exp   Def  Dist  Col  Syn  Aggr  Timeout  Activity
      xe-0/0/44      Actor    No    No   Yes  Yes  Yes   Yes     Fast    Active
      xe-0/0/44    Partner    No    No   Yes  Yes  Yes   Yes     Fast    Active
    LACP protocol:        Receive State  Transmit State          Mux State
      xe-0/0/44                 Current   Fast periodic Collecting distributing
```

Meaning This output shows that Switch B is participating in LACP negotiation.

Verifying That the multichassis aggregated Ethernet and ICL-PL Interfaces Are Up on Switch A

Purpose Verify that the multichassis aggregated Ethernet and ICL-PL interfaces are up on Switch A.

Action

```
[edit]
user@switch> show interfaces mc-ae
```

```
Member Link                  : ae1
 Current State Machine's State: mcae active state
 Local Status                : active
 Local State                 : up
 Peer Status                 : active
 Peer State                  : up
     Logical Interface       : ae1.0
     Topology Type           : bridge
     Local State             : up
     Peer State              : up
     Peer Ip/MCP/State       : 10.3.3.1 ae0.0 up
```

Meaning This output shows that the multichassis aggregated Ethernet interface on Switch A is up and active.

Verifying That the multichassis aggregated Ethernet and ICL-PL Interfaces Are Up on Switch B

Purpose Verify that the multichassis aggregated Ethernet and ICL-PL interfaces are up on Switch B.

Action

```
[edit]
user@switch> show interfaces mc-ae
```

```
Member Link                  : ae1
 Current State Machine's State: mcae active state
 Local Status                : active
 Local State                 : up
 Peer Status                 : active
 Peer State                  : up
     Logical Interface       : ae1.0
     Topology Type           : bridge
     Local State             : up
     Peer State              : up
     Peer Ip/MCP/State       : 10.3.3.2 ae0.0 up
```

Meaning This output shows that the multichassis aggregated Ethernet interface on Switch B is up and active.

Verifying That MAC Learning Is Occurring on Switch A

Purpose Verify that MAC learning is working on Switch A.

Action

```
[edit]
user@switch> show ethernet-switching table
```

```
Ethernet-switching table: 10 entries, 4 learned, 0 persistent entries
  VLAN              MAC address       Type         Age Interfaces
  V100              *                 Flood          - All-members
  V100              00:10:94:00:00:05 Learn(L)       33 ae0.0 (MCAE)
```

Meaning The output shows four learned MAC addresses entries.

Verifying That MAC Learning Is Occurring on Switch B

Purpose Verify that MAC learning is working on Switch B.

Action

```
[edit]
user@switch> show ethernet-switching table
```

```
Ethernet-switching table: 10 entries, 4 learned, 0 persistent entries
  VLAN              MAC address       Type         Age Interfaces
  V100              *                 Flood          - All-members
  V100              00:10:94:00:00:05 Learn(L)       33 ae0.0 (MCAE)
```

Meaning The output shows four learned MAC addresses entries.

Troubleshooting

Troubleshooting a LAG That Is Down

Problem The **show interfaces terse** command shows that the MC-LAG is **down**.

Solution Check the following:

1. Verify that there is no configuration mismatch.
2. Verify that all member ports are up.
3. Verify that the MC-LAG is part of family Ethernet switching (Layer 2 LAG).
4. Verify that the MC-LAG member is connected to the correct MC-LAG member at the other end.

CoS for FCoE Transit Switch Traffic Across an MC-LAG

Understanding MC-LAGs on an FCoE Transit Switch

Use an MC-LAG to provide a redundant aggregation layer for Fibre Channel over Ethernet (FCoE) traffic.

This topic describes:

Supported MC-LAG Topology

To support lossless transport of FCoE traffic across an MC-LAG, you must configure the appropriate class of service (CoS) on both of the switches with MC-LAG port members. The CoS configuration must be the same on both of the MC-LAG switches because MC-LAGs do not carry forwarding class and IEEE 802.1p priority information.

Switches that are not directly connected to FCoE hosts and that act as pass-through transit switches support MC-LAGs for FCoE traffic in an *inverted-U* network topology. Figure 3 on page 88 shows an inverted-U topology using QFX3500 switches.

Figure 3: Supported Topology for an MC-LAG on an FCoE Transit Switch

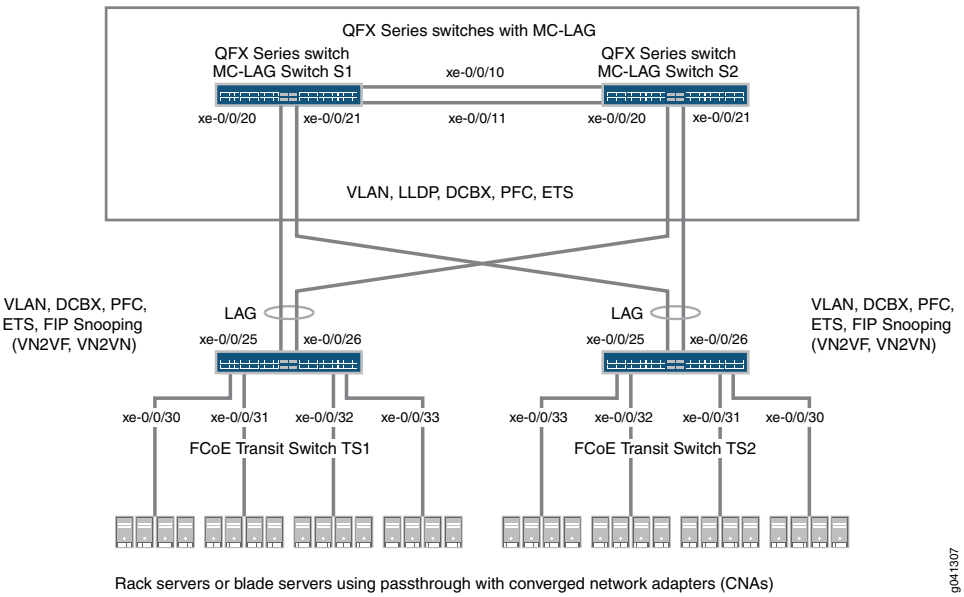

Standalone switches support MC-LAGs. QFabric system Node devices do not support MC-LAGs. Virtual Chassis and mixed-mode Virtual Chassis Fabric (VCF) configurations do not support FCoE. Only pure QFX5100 VCFs (consisting of only QFX5100 switches) support FCoE.

Ports that are part of an FCoE-FC gateway configuration (a virtual FCoE-FC gateway fabric) do not support MC-LAGs. Ports that are members of an MC-LAG act as pass-through transit switch ports.

The following rules and guidelines apply to MC-LAGs when used for FCoE traffic. The rules and guidelines help to ensure the proper handling and lossless transport characteristics required for FCoE traffic.

- The two switches that form the MC-LAG (Switches S1 and S2) cannot use ports that are part of an FCoE-FC gateway fabric. The MC-LAG switch ports must be pass-through transit switch ports (used as part of an intermediate transit switch that is not directly connected to FCoE hosts).
- MC-LAG Switches S1 and S2 cannot be directly connected to the FCoE hosts.
- The two switches that serve as access devices for FCoE hosts (FCoE Transit Switches TS1 and TS2) use standard LAGs to connect to MC-LAG Switches S1 and S2. FCoE Transit Switches TS1 and TS2 can be standalone switches or they can be Node devices in a QFabric system.
- Transit Switches TS1 and TS2 must use transit switch ports for the FCoE hosts and for the standard LAGs to MC-LAG Switches S1 and S2.
- Enable FIP snooping on the FCoE VLAN on Transit Switches TS1 and TS2. You can configure either VN_Port to VF_Port (VN2VF_Port) FIP snooping or VN_Port to VN_Port (VN2VN_Port) FIP snooping, depending on whether the FCoE hosts need to access targets in the FC SAN (VN2VF_Port FIP snooping) or targets in the Ethernet network (VN2VN_Port FIP snooping).

  FIP snooping should be performed at the access edge and is not supported on MC-LAG switches. Do not enable FIP snooping on MC-LAG Switches S1 and S2. (Do not enable FIP snooping on the MC-LAG ports that connect Switches S1 and S2 to Switches TS1 and TS2 or on the LAG ports that connect Switch S1 to S2.)

  

  NOTE: **Juniper Networks QFX10000 aggregation switches do not support FIP snooping, so they cannot be used as FIP snooping access switches (Transit Switches TS1 and TS2) in this topology.**

- The CoS configuration must be consistent on the MC-LAG switches. Because MC-LAGs carry no forwarding class or priority information, each MC-LAG switch needs to have the same CoS configuration to support lossless transport. (On each MC-LAG switch, the name, egress queue, and CoS provisioning of each forwarding class must be the same, and the priority-based flow control (PFC) configuration must be the same.)

Transit Switches (Server Access)

The role of FCoE Transit Switches TS1 and TS2 is to connect FCoE hosts in a multihomed fashion to the MC-LAG switches, so Transit Switches TS1 and TS2 act as access switches for the FCoE hosts. (FCoE hosts are directly connected to Transit Switches TS1 and TS2.)

The transit switch configuration depends on whether you want to do VN2VF_Port FIP snooping or VN2VN_Port FIP snooping, and whether the transit switches also have ports configured as part of an FCoE-FC gateway virtual fabric. Ports that a QFX3500 switch uses in an FCoE-FC gateway virtual fabric cannot be included in the transit switch LAG connection to the MC-LAG switches. (Ports cannot belong to both a transit switch and an FCoE-FC gateway; you must use different ports for each mode of operation.)

MC-LAG Switches (FCoE Aggregation)

The role of MC-LAG Switches S1 and S2 is to provide redundant, load-balanced connections between FCoE transit switches. The MC-LAG Switches S1 and S2 act as aggregation switches. FCoE hosts are not directly connected to the MC-LAG switches.

The MC-LAG switch configuration is the same regardless of which type of FIP snooping FCoE Transit Switches TS1 and TS2 perform.

FIP Snooping and FCoE Trusted Ports

To maintain secure access, enable VN2VF_Port FIP snooping or VN2VN_Port FIP snooping at the transit switch access ports connected directly to the FCoE hosts. FIP snooping should be performed at the access edge of the network to prevent unauthorized access. For example, in Figure 3 on page 88, you enable FIP snooping on the FCoE VLANs on Transit Switches TS1 and TS2 that include the access ports connected to the FCoE hosts.

Do not enable FIP snooping on the switches used to create the MC-LAG. For example, in Figure 3 on page 88, you would not enable FIP snooping on the FCoE VLANs on Switches S1 and S2.

Configure links between switches as FCoE trusted ports to reduce FIP snooping overhead and ensure that the system performs FIP snooping only at the access edge. In the sample topology, configure the Transit Switch TS1 and TS2 LAG ports connected to the MC-LAG switches as FCoE trusted ports, configure the Switch S1 and S2 MC-LAG ports connected to Switches TS1 and TS2 as FCoE trusted ports, and configure the ports in the LAG that connects Switches S1 to S2 as FCoE trusted ports.

CoS and Data Center Bridging (DCB)

The MC-LAG links do not carry forwarding class or priority information. The following CoS properties must have the same configuration on each MC-LAG switch or on each MC-LAG interface to support lossless transport:

- FCoE forwarding class name—For example, the forwarding class for FCoE traffic could use the default **fcoe** forwarding class on both MC-LAG switches.
- FCoE output queue—For example, the **fcoe** forwarding class could be mapped to queue 3 on both MC-LAG switches (queue 3 is the default mapping for the **fcoe** forwarding class).
- Classifier—The forwarding class for FCoE traffic must be mapped to the same IEEE 802.1p code point on each member interface of the MC-LAG on both MC-LAG switches. For example, the FCoE forwarding class **fcoe** could be mapped to IEEE 802.1p code point **011** (code point **011** is the default mapping for the **fcoe** forwarding class).
- Priority-based flow control (PFC)—PFC must be enabled on the FCoE code point on each MC-LAG switch and applied to each MC-LAG interface using a congestion notification profile.

You must also configure enhanced transmission selection (ETS) on the MC-LAG interfaces to provide sufficient scheduling resources (bandwidth, priority) for lossless transport.

The ETS configuration can be different on each MC-LAG switch, as long as enough resources are scheduled to support lossless transport for the expected FCoE traffic.

Link Layer Discovery Protocol (LLDP) and Data Center Bridging Capability Exchange Protocol (DCBX) must be enabled on each MC-LAG member interface (LLDP and DCBX are enabled by default on all interfaces).

NOTE: As with all other FCoE configurations, FCoE traffic requires a dedicated VLAN that carries only FCoE traffic, and IGMP snooping must be disabled on the FCoE VLAN.

Example: Configuring CoS for FCoE Transit Switch Traffic Across an MC-LAG

Multichassis link aggregation groups (MC-LAGs) provide redundancy and load balancing between two switches, multihoming support for client devices such as servers, and a loop-free Layer 2 network without running Spanning Tree Protocol (STP).

NOTE: This example uses Junos OS without support for the Enhanced Layer 2 Software (ELS) configuration style. If your switch runs software that does support ELS, see *Example: Configuring CoS Using ELS for FCoE Transit Switch Traffic Across an MC-LAG*. For ELS details, see *Using the Enhanced Layer 2 Software CLI*.

You can use an MC-LAG to provide a redundant aggregation layer for Fibre Channel over Ethernet (FCoE) traffic in an *inverted-U* topology. To support lossless transport of FCoE traffic across an MC-LAG, you must configure the appropriate class of service (CoS) on both of the switches with MC-LAG port members. The CoS configuration must be the same on both of the MC-LAG switches because an MC-LAG does not carry forwarding class and IEEE 802.1p priority information.

NOTE: This example describes how to configure CoS to provide lossless transport for FCoE traffic across an MC-LAG that connects two switches. It also describes how to configure CoS on the FCoE transit switches that connect FCoE hosts to the two switches that form the MC-LAG.

This example does *not* describe how to configure the MC-LAG itself. For a detailed example of MC-LAG configuration, see "Configuring Multichassis Link Aggregation" on page 56. However, this example includes a subset of MC-LAG configuration that only shows how to configure interface membership in the MC-LAG.

Ports that are part of an FCoE-FC gateway configuration (a virtual FCoE-FC gateway fabric) do not support MC-LAGs. Ports that are members of an MC-LAG act as FCoE pass-through transit switch ports.

QFX Series switches and EX4600 switches support MC-LAGs. QFabric system Node devices do not support MC-LAGs.

Requirements

This example uses the following hardware and software components:

- Two Juniper Networks QFX3500 switches that form an MC-LAG for FCoE traffic.
- Two Juniper Networks QFX3500 switches that provide FCoE server access in transit switch mode and that connect to the MC-LAG switches. These switches can be standalone QFX3500 switches or they can be Node devices in a QFabric system.
- FCoE servers (or other FCoE hosts) connected to the transit switches.
- Junos OS Release 12.2 or later for the QFX Series.

Overview

FCoE traffic requires lossless transport. This example shows you how to:

- Configure CoS for FCoE traffic on the two QFX3500 switches that form the MC-LAG, including priority-based flow control (PFC) and enhanced transmission selection (ETS; hierarchical scheduling of resources for the FCoE forwarding class priority and for the forwarding class set priority group).

  NOTE: Configuring or changing PFC on an interface blocks the entire port until the PFC change is completed. After a PFC change is completed, the port is unblocked and traffic resumes. Blocking the port stops ingress and egress traffic, and causes packet loss on all queues on the port until the port is unblocked.

- Configure CoS for FCoE on the two FCoE transit switches that connect FCoE hosts to the MC-LAG switches and enable FIP snooping on the FCoE VLAN at the FCoE transit switch access ports.
- Disable IGMP snooping on the FCoE VLAN.

  NOTE: This is only necessary if IGMP snooping is enabled on the VLAN. Before Junos OS Release 13.2, IGMP snooping was enabled by default on VLANs. Beginning with Junos OS Release 13.2, IGMP snooping is enabled by default only on the default VLAN.

- Configure the appropriate port mode, MTU, and FCoE trusted or untrusted state for each interface to support lossless FCoE transport.

Topology

Switches that act as transit switches support MC-LAGs for FCoE traffic in an inverted-U network topology, as shown in Figure 4 on page 93.

Figure 4: Supported Topology for an MC-LAG on an FCoE Transit Switch

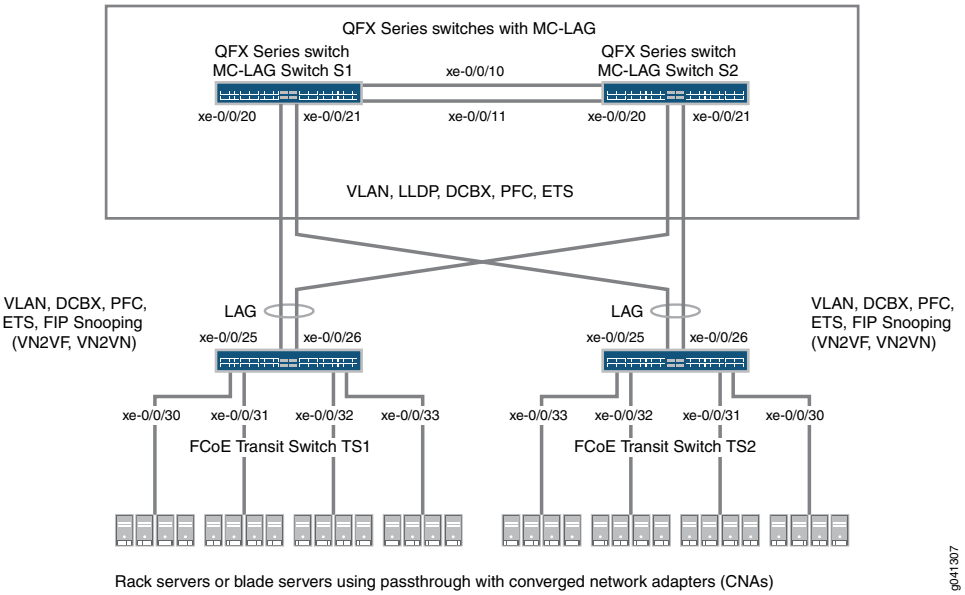

Table 6 on page 93 shows the configuration components for this example.

Table 6: Components of the CoS for FCoE Traffic Across an MC-LAG Configuration Topology

Component	Settings
Hardware	Four QFX3500 switches (two to form the MC-LAG as pass-through transit switches and two transit switches for FCoE access).
Forwarding class (all switches)	Default **fcoe** forwarding class.
Classifier (forwarding class mapping of incoming traffic to IEEE priority)	Default IEEE 802.1p trusted classifier on all FCoE interfaces.

Table 6: Components of the CoS for FCoE Traffic Across an MC-LAG Configuration Topology (continued)

Component	Settings
LAGs and MC-LAG	S1—Ports xe-0/0/10 and x-0/0/11 are members of LAG ae0, which connects Switch S1 to Switch S2. Ports xe-0/0/20 and xe-0/0/21 are members of MC-LAG ae1. All ports are configured in **trunk** port mode, as **fcoe-trusted**, and with an MTU of **2180**. S2—Ports xe-0/0/10 and x-0/0/11 are members of LAG ae0, which connects Switch S2 to Switch S1. Ports xe-0/0/20 and xe-0/0/21 are members of MC-LAG ae1. All ports are configured in **trunk** port mode, as **fcoe-trusted**, and with an MTU of **2180**. NOTE: Ports xe-0/0/20 and xe-0/0/21 on Switches S1 and S2 are the members of the MC-LAG. TS1—Ports xe-0/0/25 and x-0/0/26 are members of LAG ae1, configured in **trunk** port mode, as **fcoe-trusted**, and with an MTU of **2180**. Ports xe-0/0/30, xe-0/0/31, xe-0/0/32, and xe-0/0/33 are configured in **tagged-access** port mode, with an MTU of **2180**. TS2—Ports xe-0/0/25 and x-0/0/26 are members of LAG ae1, configured in **trunk** port mode, as **fcoe-trusted**, and with an MTU of **2180**. Ports xe-0/0/30, xe-0/0/31, xe-0/0/32, and xe-0/0/33 are configured in **tagged-access** port mode, with an MTU of **2180**.
FCoE queue scheduler (all switches)	**fcoe-sched**: Minimum bandwidth **3g** Maximum bandwidth **100**% Priority **low**
Forwarding class-to-scheduler mapping (all switches)	Scheduler map **fcoe-map**: Forwarding class **fcoe** Scheduler **fcoe-sched**
Forwarding class set (FCoE priority group, all switches)	**fcoe-pg**: Forwarding class **fcoe** Egress interfaces: • S1—LAG ae0 and MC-LAG ae1 • S2—LAG ae0 and MC-LAG ae1 • TS1—LAG ae1, interfaces xe-0/0/30, xe-0/0/31, xe-0/0/32, and xe-0/0/33 • TS2—LAG ae1, interfaces xe-0/0/30, xe-0/0/31, xe-0/0/32, and xe-0/0/33
Traffic control profile (all switches)	**fcoe-tcp**: Scheduler map **fcoe-map** Minimum bandwidth **3g** Maximum bandwidth **100**%

Table 6: Components of the CoS for FCoE Traffic Across an MC-LAG Configuration Topology (continued)

Component	Settings
PFC congestion notification profile (all switches)	**fcoe-cnp**: Code point **011** Ingress interfaces: • S1—LAG ae0 and MC-LAG ae1 • S2—LAG ae0 and MC-LAG ae1 • TS1—LAG ae1, interfaces xe-0/0/30, xe-0/0/31, xe-0/0/32, and xe-0/0/33 • TS2—LAG ae1, interfaces xe-0/0/30, xe-0/0/31, xe-0/0/32, and xe-0/0/33
FCoE VLAN name and tag ID	Name—**fcoe_vlan** ID—**100** Include the FCoE VLAN on the interfaces that carry FCoE traffic on all four switches. Disable IGMP snooping on the interfaces that belong to the FCoE VLAN on all four switches.
FIP snooping	Enable FIP snooping on Transit Switches TS1 and TS2 on the FCoE VLAN. Configure the LAG interfaces that connect to the MC-LAG switches as FCoE trusted interfaces so that they do not perform FIP snooping. This example enables VN2VN_Port FIP snooping on the FCoE transit switch interfaces connected to the FCoE servers. The example is equally valid with VN2VF_Port FIP snooping enabled on the transit switch access ports. The method of FIP snooping you enable depends on your network configuration.

NOTE: This example uses the default IEEE 802.1p trusted BA classifier, which is automatically applied to trunk mode and tagged access mode ports if you do not apply an explicitly configured classifier.

To configure CoS for FCoE traffic across an MC-LAG:

- Use the default FCoE forwarding class and forwarding-class-to-queue mapping (do not explicitly configure the FCoE forwarding class or output queue). The default FCoE forwarding class is **fcoe**, and the default output queue is queue **3**.

  

  NOTE: In Junos OS Release 12.2, traffic mapped to explicitly configured forwarding classes, even lossless forwarding classes such as **fcoe**, is treated as lossy (**best-effort**) traffic and does *not* receive lossless treatment. To receive lossless treatment in Release 12.2, traffic must use one of the default lossless forwarding classes (**fcoe** or **no-loss**).

  In Junos OS Release 12.3 and later, you can include the *no-loss* packet drop attribute in the explicit forwarding class configuration to configure a lossless forwarding class.

- Use the default trusted BA classifier, which maps incoming packets to forwarding classes by the IEEE 802.1p code point (CoS priority) of the packet. The trusted classifier is the default classifier for interfaces in trunk and tagged-access port modes. The default trusted classifier maps incoming packets with the IEEE 802.1p code point 3 (**011**) to the FCoE forwarding class. If you choose to configure the BA classifier instead of using the default classifier, you must ensure that FCoE traffic is classified into forwarding classes in exactly the same way on both MC-LAG switches. Using the default classifier ensures consistent classifier configuration on the MC-LAG ports.
- Configure a congestion notification profile that enables PFC on the FCoE code point (code point **011** in this example). The congestion notification profile configuration must be the same on both MC-LAG switches.
- Apply the congestion notification profile to the interfaces.
- Configure enhanced transmission selection (ETS, also known as hierarchical scheduling) on the interfaces to provide the bandwidth required for lossless FCoE transport. Configuring ETS includes configuring bandwidth scheduling for the FCoE forwarding class, a forwarding class set (priority group) that includes the FCoE forwarding class, and a traffic control profile to assign bandwidth to the forwarding class set that includes FCoE traffic.
- Apply the ETS scheduling to the interfaces.
- Configure the port mode, MTU, and FCoE trusted or untrusted state for each interface to support lossless FCoE transport.

In addition, this example describes how to enable FIP snooping on the Transit Switch TS1 and TS2 ports that are connected to the FCoE servers and how to disable IGMP snooping on the FCoE VLAN. To provide secure access, FIP snooping must be enabled on the FCoE access ports.

This example focuses on the CoS configuration to support lossless FCoE transport across an MC-LAG. This example does not describe how to configure the properties of MC-LAGs and LAGs, although it does show you how to configure the port characteristics required

to support lossless transport and how to assign interfaces to the MC-LAG and to the LAGs.

Before you configure CoS, configure:

- The MC-LAGs that connect Switches S1 and S2 to Switches TS1 and TS2.
- The LAGs that connect the Transit Switches TS1 and TS2 to MC-LAG Switches S1 and S2.
- The LAG that connects Switch S1 to Switch S2.

Configuration

To configure CoS for lossless FCoE transport across an MC-LAG, perform these tasks:

- Configuring MC-LAG Switches S1 and S2 on page 99
- Configuring FCoE Transit Switches TS1 and TS2 on page 101
- Results on page 103

CLI Quick Configuration

To quickly configure CoS for lossless FCoE transport across an MC-LAG, copy the following commands, paste them in a text file, remove line breaks, change variables and details to match your network configuration, and then copy and paste the commands into the CLI for MC-LAG Switch S1 and MC-LAG Switch S2 at the **[edit]** hierarchy level. The configurations on Switches S1 and S2 are identical because the CoS configuration must be identical, and because this example uses the same ports on both switches.

Switch S1 and Switch S2

```
set class-of-service schedulers fcoe-sched priority low transmit-rate 3g
set class-of-service schedulers fcoe-sched shaping-rate percent 100
set class-of-service scheduler-maps fcoe-map forwarding-class fcoe scheduler fcoe-sched
set class-of-service forwarding-class-sets fcoe-pg class fcoe
set class-of-service traffic-control-profiles fcoe-tcp scheduler-map fcoe-map guaranteed-rate 3g
set class-of-service traffic-control-profiles fcoe-tcp shaping-rate percent 100
set class-of-service interfaces ae0 forwarding-class-set fcoe-pg output-traffic-control-profile fcoe-tcp
set class-of-service interfaces ae1 forwarding-class-set fcoe-pg output-traffic-control-profile fcoe-tcp
set class-of-service congestion-notification-profile fcoe-cnp input ieee-802.1 code-point 011 pfc
set class-of-service interfaces ae0 congestion-notification-profile fcoe-cnp
set class-of-service interfaces ae1 congestion-notification-profile fcoe-cnp
set vlans fcoe_vlan vlan-id 100
set protocols igmp-snooping vlan fcoe_vlan disable
set interfaces xe-0/0/10 ether-options 802.3ad ae0
set interfaces xe-0/0/11 ether-options 802.3ad ae0
set interfaces xe-0/0/20 ether-options 802.3ad ae1
set interfaces xe-0/0/21 ether-options 802.3ad ae1
set interfaces ae0 unit 0 family ethernet-switching port-mode trunk vlan members fcoe_vlan
set interfaces ae1 unit 0 family ethernet-switching port-mode trunk vlan members fcoe_vlan
set interfaces ae0 mtu 2180
set interfaces ae1 mtu 2180
set ethernet-switching-options secure-access-port interface ae0 fcoe-trusted
```

```
set ethernet-switching-options secure-access-port interface ae1 fcoe-trusted
```

To quickly configure CoS for lossless FCoE transport across an MC-LAG, copy the following commands, paste them in a text file, remove line breaks, change variables and details to match your network configuration, and then copy and paste the commands into the CLI for Transit Switch TS1 and Transit Switch TS2 at the **[edit]** hierarchy level. The configurations on Switches TS1 and TS2 are identical because the CoS configuration must be identical, and because this example uses the same ports on both switches.

Switch TS1 and Switch TS2

```
set class-of-service schedulers fcoe-sched priority low transmit-rate 3g
set class-of-service schedulers fcoe-sched shaping-rate percent 100
set class-of-service scheduler-maps fcoe-map forwarding-class fcoe scheduler fcoe-sched
set class-of-service forwarding-class-sets fcoe-pg class fcoe
set class-of-service traffic-control-profiles fcoe-tcp scheduler-map fcoe-map guaranteed-rate
3g
set class-of-service traffic-control-profiles fcoe-tcp shaping-rate percent 100
set class-of-service interfaces ae1 forwarding-class-set fcoe-pg output-traffic-control-profile
fcoe-tcp
set class-of-service interfaces xe-0/0/30 forwarding-class-set fcoe-pg
output-traffic-control-profile fcoe-tcp
set class-of-service interfaces xe-0/0/31 forwarding-class-set fcoe-pg
output-traffic-control-profile fcoe-tcp
set class-of-service interfaces xe-0/0/32 forwarding-class-set fcoe-pg
output-traffic-control-profile fcoe-tcp
set class-of-service interfaces xe-0/0/33 forwarding-class-set fcoe-pg
output-traffic-control-profile fcoe-tcp
set class-of-service congestion-notification-profile fcoe-cnp input ieee-802.1 code-point 011 pfc
set class-of-service interfaces ae1 congestion-notification-profile fcoe-cnp
set class-of-service interfaces xe-0/0/30 congestion-notification-profile fcoe-cnp
set class-of-service interfaces xe-0/0/31 congestion-notification-profile fcoe-cnp
set class-of-service interfaces xe-0/0/32 congestion-notification-profile fcoe-cnp
set class-of-service interfaces xe-0/0/33 congestion-notification-profile fcoe-cnp
set vlans fcoe_vlan vlan-id 100
set protocols igmp-snooping vlan fcoe_vlan disable
set interfaces xe-0/0/25 ether-options 802.3ad ae1
set interfaces xe-0/0/26 ether-options 802.3ad ae1
set interfaces ae1 unit 0 family ethernet-switching port-mode trunk vlan members fcoe_vlan
set interfaces xe-0/0/30 unit 0 family ethernet-switching port-mode tagged-access vlan members
fcoe_vlan
set interfaces xe-0/0/31 unit 0 family ethernet-switching port-mode tagged-access vlan members
fcoe_vlan
set interfaces xe-0/0/32 unit 0 family ethernet-switching port-mode tagged-access vlan members
fcoe_vlan
set interfaces xe-0/0/33 unit 0 family ethernet-switching port-mode tagged-access vlan members
fcoe_vlan
set interfaces ae1 mtu 2180
set interfaces xe-0/0/30 mtu 2180
set interfaces xe-0/0/31 mtu 2180
set interfaces xe-0/0/32 mtu 2180
set interfaces xe-0/0/33 mtu 2180
set ethernet-switching-options secure-access-port interface ae1 fcoe-trusted
set ethernet-switching-options secure-access-port vlan fcoe_vlan examine-fip examine-vn2v2
beacon-period 90000
```

Configuring MC-LAG Switches S1 and S2

Step-by-Step Procedure

To configure CoS resource scheduling (ETS), PFC, the FCoE VLAN, and the LAG and MC-LAG interface membership and characteristics to support lossless FCoE transport across an MC-LAG (this example uses the default **fcoe** forwarding class and the default classifier to map incoming FCoE traffic to the FCoE IEEE 802.1p code point **011**, so you do not configure them):

1. Configure output scheduling for the FCoE queue.

```
[edit class-of-service schedulers fcoe-sched]
user@switch# set priority low transmit-rate 3g
user@switch# set shaping-rate percent 100
```

2. Map the FCoE forwarding class to the FCoE scheduler (**fcoe-sched**).

```
[edit class-of-service]
user@switch# set scheduler-maps fcoe-map forwarding-class fcoe scheduler fcoe-sched
```

3. Configure the forwarding class set (**fcoe-pg**) for the FCoE traffic.

```
[edit class-of-service]
user@switch# set forwarding-class-sets fcoe-pg class fcoe
```

4. Define the traffic control profile (**fcoe-tcp**) to use on the FCoE forwarding class set.

```
[edit class-of-service traffic-control-profiles fcoe-tcp]
user@switch# set scheduler-map fcoe-map guaranteed-rate 3g
user@switch# set shaping-rate percent 100
```

5. Apply the FCoE forwarding class set and traffic control profile to the LAG and MC-LAG interfaces.

```
[edit class-of-service]
user@switch# set interfaces ae0 forwarding-class-set fcoe-pg output-traffic-control-profile
fcoe-tcp
user@switch# set interfaces ae1 forwarding-class-set fcoe-pg output-traffic-control-profile
fcoe-tcp
```

6. Enable PFC on the FCoE priority by creating a congestion notification profile (**fcoe-cnp**) that applies FCoE to the IEEE 802.1 code point **011**.

```
[edit class-of-service]
user@switch# set congestion-notification-profile fcoe-cnp input ieee-802.1 code-point
011 pfc
```

7. Apply the PFC configuration to the LAG and MC-LAG interfaces.

```
[edit class-of-service]
```

```
user@switch# set interfaces ae0 congestion-notification-profile fcoe-cnp
user@switch# set interfaces ae1 congestion-notification-profile fcoe-cnp
```

8. Configure the VLAN for FCoE traffic (**fcoe_vlan**).

```
[edit vlans]
user@switch# set fcoe_vlan vlan-id 100
```

9. Disable IGMP snooping on the FCoE VLAN.

```
[edit protocols]
user@switch# set igmp-snooping vlan fcoe_vlan disable
```

10. Add the member interfaces to the LAG between the two MC-LAG switches.

```
[edit interfaces]
user@switch# set xe-0/0/10 ether-options 802.3ad ae0
user@switch# set xe-0/0/11 ether-options 802.3ad ae0
```

11. Add the member interfaces to the MC-LAG.

```
[edit interfaces]
user@switch# set xe-0/0/20 ether-options 802.3ad ae1
user@switch# set xe-0/0/21 ether-options 802.3ad ae1
```

12. Configure the port mode as **trunk** and membership in the FCoE VLAN (**fcoe_vlan**) for the LAG (ae0) and for the MC-LAG (ae1).

```
[edit interfaces]
user@switch# set ae0 unit 0 family ethernet-switching port-mode trunk vlan members
fcoe_vlan
user@switch# set ae1 unit 0 family ethernet-switching port-mode trunk vlan members
fcoe_vlan
```

13. Set the MTU to **2180** for the LAG and MC-LAG interfaces.

    2180 bytes is the minimum size required to handle FCoE packets because of the payload and header sizes. You can configure the MTU to a higher number of bytes if desired, but not less than 2180 bytes.

```
[edit interfaces]
user@switch# set ae0 mtu 2180
user@switch# set ae1 mtu 2180
```

14. Set the LAG and MC-LAG interfaces as FCoE trusted ports.

    Ports that connect to other switches should be trusted and should not perform FIP snooping.


```
[edit ethernet-switching-options secure-access-port interface]
user@switch# set ae0 fcoe-trusted
user@switch# set ae1 fcoe-trusted
```

Configuring FCoE Transit Switches TS1 and TS2

Step-by-Step Procedure

The CoS configuration on FCoE Transit Switches TS1 and TS2 is similar to the CoS configuration on MC-LAG Switches S1 and S2. However, the port configurations differ, and you must enable FIP snooping on the Switch TS1 and Switch TS2 FCoE access ports.

To configure resource scheduling (ETS), PFC, the FCoE VLAN, and the LAG interface membership and characteristics to support lossless FCoE transport across the MC-LAG (this example uses the default **fcoe** forwarding class and the default classifier to map incoming FCoE traffic to the FCoE IEEE 802.1p code point **011**, so you do not configure them):

1. Configure output scheduling for the FCoE queue.

```
[edit class-of-service schedulers fcoe-sched]
user@switch# set priority low transmit-rate 3g
user@switch# set shaping-rate percent 100
```

2. Map the FCoE forwarding class to the FCoE scheduler (**fcoe-sched**).

```
[edit class-of-service]
user@switch# set scheduler-maps fcoe-map forwarding-class fcoe scheduler fcoe-sched
```

3. Configure the forwarding class set (**fcoe-pg**) for the FCoE traffic.

```
[edit class-of-service]
user@switch# set forwarding-class-sets fcoe-pg class fcoe
```

4. Define the traffic control profile (**fcoe-tcp**) to use on the FCoE forwarding class set.

```
[edit class-of-service]
user@switch# set traffic-control-profiles fcoe-tcp scheduler-map fcoe-map
guaranteed-rate 3g
user@switch# set traffic-control-profiles fcoe-tcp shaping-rate percent 100
```

5. Apply the FCoE forwarding class set and traffic control profile to the LAG interface and to the FCoE access interfaces.

```
[edit class-of-service]
user@switch# set interfaces ae1 forwarding-class-set fcoe-pg output-traffic-control-profile
fcoe-tcp
user@switch# set interfaces xe-0/0/30 forwarding-class-set fcoe-pg
output-traffic-control-profile fcoe-tcp
user@switch# set interfaces xe-0/0/31 forwarding-class-set fcoe-pg
output-traffic-control-profile fcoe-tcp
user@switch# set interfaces xe-0/0/32 forwarding-class-set fcoe-pg
output-traffic-control-profile fcoe-tcp
```

```
user@switch# set interfaces xe-0/0/33 forwarding-class-set fcoe-pg
output-traffic-control-profile fcoe-tcp
```

6. Enable PFC on the FCoE priority by creating a congestion notification profile (**fcoe-cnp**) that applies FCoE to the IEEE 802.1 code point **011**.

```
[edit class-of-service]
user@switch# set congestion-notification-profile fcoe-cnp input ieee-802.1 code-point
011 pfc
```

7. Apply the PFC configuration to the LAG interface and to the FCoE access interfaces.

```
[edit class-of-service]
user@switch# set interfaces ae1 congestion-notification-profile fcoe-cnp
user@switch# set interfaces xe-0/0/30 congestion-notification-profile fcoe-cnp
user@switch# set interfaces xe-0/0/31 congestion-notification-profile fcoe-cnp
user@switch# set interfaces xe-0/0/32 congestion-notification-profile fcoe-cnp
user@switch# set interfaces xe-0/0/33 congestion-notification-profile fcoe-cnp
```

8. Configure the VLAN for FCoE traffic (**fcoe_vlan**).

```
[edit vlans]
user@switch# set fcoe_vlan vlan-id 100
```

9. Disable IGMP snooping on the FCoE VLAN.

```
[edit protocols]
user@switch# set igmp-snooping vlan fcoe_vlan disable
```

10. Add the member interfaces to the LAG.

```
[edit interfaces]
user@switch# set xe-0/0/25 ether-options 802.3ad ae1
user@switch# set xe-0/0/26 ether-options 802.3ad ae1
```

11. On the LAG (**ae1**), configure the port mode as **trunk** and membership in the FCoE VLAN (**fcoe_vlan**).

```
[edit interfaces]
user@switch# set ae1 unit 0 family ethernet-switching port-mode trunk vlan members
fcoe_vlan
```

12. On the FCoE access interfaces (**xe-0/0/30**, **xe-0/0/31**, **xe-0/0/32**, **xe-0/0/33**), configure the port mode as **tagged-access** and membership in the FCoE VLAN (**fcoe_vlan**).

```
[edit interfaces]
user@switch# set xe-0/0/30 unit 0 family ethernet-switching port-mode tagged-access
vlan members fcoe_vlan
```



```
user@switch# set xe-0/0/31 unit 0 family ethernet-switching port-mode tagged-access
vlan members fcoe_vlan
user@switch# set xe-0/0/32 unit 0 family ethernet-switching port-mode tagged-access
vlan members fcoe_vlan
user@switch# set xe-0/0/33 unit 0 family ethernet-switching port-mode tagged-access
vlan members fcoe_vlan
```

13. Set the MTU to **2180** for the LAG and FCoE access interfaces.

    2180 bytes is the minimum size required to handle FCoE packets because of the payload and header sizes; you can configure the MTU to a higher number of bytes if desired, but not less than 2180 bytes.

```
[edit interfaces]
user@switch# set ae1 mtu 2180
user@switch# set xe-0/0/30 mtu 2180
user@switch# set xe-0/0/31 mtu 2180
user@switch# set xe-0/0/32 mtu 2180
user@switch# set xe-0/0/33 mtu 2180
```

14. Set the LAG interface as an FCoE trusted port. Ports that connect to other switches should be trusted and should not perform FIP snooping:

```
[edit ethernet-switching-options]
user@switch# set secure-access-port interface ae1 fcoe-trusted
```

NOTE: Access ports xe-0/0/30, xe-0/0/31, xe-0/0/32, and xe-0/0/33 are not configured as FCoE trusted ports. The access ports remain in the default state as untrusted ports because they connect directly to FCoE devices and must perform FIP snooping to ensure network security.

15. Enable FIP snooping on the FCoE VLAN to prevent unauthorized FCoE network access (this example uses VN2VN_Port FIP snooping; the example is equally valid if you use VN2VF_Port FIP snooping).

```
[edit ethernet-switching-options]
user@switch# set secure-access-port vlan fcoe_vlan examine-fip examine-vn2vn
beacon-period 90000
```

Results

Display the results of the CoS configuration on MC-LAG Switch S1 and on MC-LAG Switch S2 (the results on both switches are the same).

```
user@switch> show configuration class-of-service
traffic-control-profiles {
   fcoe-tcp {
```

```
      scheduler-map fcoe-map;
      shaping-rate percent 100;
      guaranteed-rate 3g;
    }
  }
  forwarding-class-sets {
    fcoe-pg {
      class fcoe;
    }
  }
  congestion-notification-profile {
    fcoe-cnp {
      input {
        ieee-802.1 {
          code-point 011 {
            pfc;
          }
        }
      }
    }
  }
  interfaces {
    ae0 {
      forwarding-class-set {
        fcoe-pg {
          output-traffic-control-profile fcoe-tcp;
        }
      }
      congestion-notification-profile fcoe-cnp;
    }
    ae1 {
      forwarding-class-set {
        fcoe-pg {
          output-traffic-control-profile fcoe-tcp;
        }
      }
      congestion-notification-profile fcoe-cnp;
    }
  }
  scheduler-maps {
    fcoe-map {
      forwarding-class fcoe scheduler fcoe-sched;
    }
  }
  schedulers {
    fcoe-sched {
      transmit-rate 3g;
      shaping-rate percent 100;
      priority low;
    }
  }
```

NOTE: The forwarding class and classifier configurations are not shown because the show command does not display default portions of the configuration.

Display the results of the CoS configuration on FCoE Transit Switch TS1 and on FCoE Transit Switch TS2 (the results on both transit switches are the same).

```
user@switch> show configuration class-of-service
traffic-control-profiles {
   fcoe-tcp {
      scheduler-map fcoe-map;
      shaping-rate percent 100;
      guaranteed-rate 3g;
   }
}
forwarding-class-sets {
   fcoe-pg {
      class fcoe;
   }
}
congestion-notification-profile {
   fcoe-cnp {
      input {
         ieee-802.1 {
            code-point 011 {
               pfc;
            }
         }
      }
   }
}
interfaces {
   xe-0/0/30 {
      forwarding-class-set {
         fcoe-pg {
            output-traffic-control-profile fcoe-tcp;
         }
      }
      congestion-notification-profile fcoe-cnp;
   }
   xe-0/0/31 {
      forwarding-class-set {
         fcoe-pg {
            output-traffic-control-profile fcoe-tcp;
         }
      }
      congestion-notification-profile fcoe-cnp;
   }
   xe-0/0/32 {
      forwarding-class-set {
         fcoe-pg {
```

```
            output-traffic-control-profile fcoe-tcp;
          }
        }
        congestion-notification-profile fcoe-cnp;
      }
      xe-0/0/33 {
        forwarding-class-set {
          fcoe-pg {
            output-traffic-control-profile fcoe-tcp;
          }
        }
        congestion-notification-profile fcoe-cnp;
      }
      ae1 {
        forwarding-class-set {
          fcoe-pg {
            output-traffic-control-profile fcoe-tcp;
          }
        }
        congestion-notification-profile fcoe-cnp;
      }
    }
    scheduler-maps {
      fcoe-map {
        forwarding-class fcoe scheduler fcoe-sched;
      }
    }
    schedulers {
      fcoe-sched {
        transmit-rate 3g;
        shaping-rate percent 100;
        priority low;
      }
    }
```

Verification

To verify that the CoS components and FIP snooping have been configured and are operating properly, perform these tasks. Because this example uses the default **fcoe** forwarding class and the default IEEE 802.1p trusted classifier, the verification of those configurations is not shown.

- Verifying That the Output Queue Schedulers Have Been Created on page 107
- Verifying That the Priority Group Output Scheduler (Traffic Control Profile) Has Been Created on page 107
- Verifying That the Forwarding Class Set (Priority Group) Has Been Created on page 108
- Verifying That Priority-Based Flow Control Has Been Enabled on page 108
- Verifying That the Interface Class of Service Configuration Has Been Created on page 109
- Verifying That the Interfaces Are Correctly Configured on page 111
- Verifying That FIP Snooping Is Enabled on the FCoE VLAN on FCoE Transit Switches TS1 and TS2 Access Interfaces on page 114

- Verifying That the FIP Snooping Mode Is Correct on FCoE Transit Switches TS1 and TS2 on page 115
- Verifying That IGMP Snooping Is Disabled on the FCoE VLAN on page 115

Verifying That the Output Queue Schedulers Have Been Created

Purpose Verify that the output queue scheduler for FCoE traffic has the correct bandwidth parameters and priorities, and is mapped to the correct forwarding class (output queue). Queue scheduler verification is the same on each of the four switches.

Action List the scheduler map using the operational mode command **show class-of-service scheduler-map fcoe-map**:

```
user@switch> show class-of-service scheduler-map fcoe-map
Scheduler map: fcoe-map, Index: 9023

  Scheduler: fcoe-sched, Forwarding class: fcoe, Index: 37289
    Transmit rate: 3000000000 bps, Rate Limit: none, Buffer size: remainder,
    Buffer Limit: none, Priority: low
    Excess Priority: unspecified
    Shaping rate: 100 percent,
    drop-profile-map-set-type: mark
    Drop profiles:
      Loss priority   Protocol    Index    Name
      Low             any             1    <default-drop-profile>
      Medium high     any             1    <default-drop-profile>
      High            any             1    <default-drop-profile>
```

Meaning The **show class-of-service scheduler-map fcoe-map** command lists the properties of the scheduler map **fcoe-map**. The command output includes:

- The name of the scheduler map (**fcoe-map**)
- The name of the scheduler (**fcoe-sched**)
- The forwarding classes mapped to the scheduler (**fcoe**)
- The minimum guaranteed queue bandwidth (transmit rate **3000000000 bps**)
- The scheduling priority (**low**)
- The maximum bandwidth in the priority group the queue can consume (shaping rate **100 percent**)
- The drop profile loss priority for each drop profile name. This example does not include drop profiles because you do not apply drop profiles to FCoE traffic.

Verifying That the Priority Group Output Scheduler (Traffic Control Profile) Has Been Created

Purpose Verify that the traffic control profile **fcoe-tcp** has been created with the correct bandwidth parameters and scheduler mapping. Priority group scheduler verification is the same on each of the four switches.

Action List the FCoE traffic control profile properties using the operational mode command **show class-of-service traffic-control-profile fcoe-tcp**:

```
user@switch> show class-of-service traffic-control-profile fcoe-tcp
Traffic control profile: fcoe-tcp, Index: 18303
  Shaping rate: 100 percent
  Scheduler map: fcoe-map
  Guaranteed rate: 3000000000
```

Meaning The **show class-of-service traffic-control-profile fcoe-tcp** command lists all of the configured traffic control profiles. For each traffic control profile, the command output includes:

- The name of the traffic control profile (**fcoe-tcp**)
- The maximum port bandwidth the priority group can consume (shaping rate **100 percent**)
- The scheduler map associated with the traffic control profile (**fcoe-map**)
- The minimum guaranteed priority group port bandwidth (guaranteed rate **3000000000** in bps)

Verifying That the Forwarding Class Set (Priority Group) Has Been Created

Purpose Verify that the FCoE priority group has been created and that the **fcoe** priority (forwarding class) belongs to the FCoE priority group. Forwarding class set verification is the same on each of the four switches.

Action List the forwarding class sets using the operational mode command **show class-of-service forwarding-class-set fcoe-pg**:

```
user@switch> show class-of-service forwarding-class-set fcoe-pg
Forwarding class set: fcoe-pg, Type: normal-type, Forwarding class set index:
31420
  Forwarding class                        Index
  fcoe                                      1
```

Meaning The **show class-of-service forwarding-class-set fcoe-pg** command lists all of the forwarding classes (priorities) that belong to the **fcoe-pg** priority group, and the internal index number of the priority group. The command output shows that the forwarding class set **fcoe-pg** includes the forwarding class **fcoe**.

Verifying That Priority-Based Flow Control Has Been Enabled

Purpose Verify that PFC is enabled on the FCoE code point. PFC verification is the same on each of the four switches.

Action List the FCoE congestion notification profile using the operational mode command **show class-of-service congestion-notification fcoe-cnp**:

```
user@switch> show class-of-service congestion-notification fcoe-cnp
Type: Input, Name: fcoe-cnp, Index: 6879
Cable Length: 100 m
  Priority     PFC          MRU
  000          Disabled
  001          Disabled
  010          Disabled
  011          Enabled      2500
  100          Disabled
  101          Disabled
  110          Disabled
  111          Disabled
Type: Output
  Priority     Flow-Control-Queues
  000
               0
  001
               1
  010
               2
  011
               3
  100
               4
  101
               5
  110
               6
  111
               7
```

Meaning The **show class-of-service congestion-notification fcoe-cnp** command lists all of the IEEE 802.1p code points in the congestion notification profile that have PFC enabled. The command output shows that PFC is enabled on code point **011** (**fcoe** queue) for the **fcoe-cnp** congestion notification profile.

The command also shows the default cable length (**100** meters), the default maximum receive unit (**2500** bytes), and the default mapping of priorities to output queues because this example does not include configuring these options.

Verifying That the Interface Class of Service Configuration Has Been Created

Purpose Verify that the CoS properties of the interfaces are correct. The verification output on MC-LAG Switches S1 and S2 differs from the output on FCoE Transit Switches TS1 and TS2.

Action List the interface CoS configuration on MC-LAG Switches S1 and S2 using the operational mode command **show configuration class-of-service interfaces**:

```
user@switch> show configuration class-of-service interfaces
```



```
ae0 {
    forwarding-class-set {
        fcoe-pg {
            output-traffic-control-profile fcoe-tcp;
        }
    }
    congestion-notification-profile fcoe-cnp;
}

ae1 {
    forwarding-class-set {
        fcoe-pg {
            output-traffic-control-profile fcoe-tcp;
        }
    }
    congestion-notification-profile fcoe-cnp;
}
```

List the interface CoS configuration on FCoE Transit Switches TS1 and TS2 using the operational mode command **show configuration class-of-service interfaces**:

```
user@switch> show configuration class-of-service interfaces
xe-0/0/30 {
    forwarding-class-set {
        fcoe-pg {
            output-traffic-control-profile fcoe-tcp;
        }
    }
    congestion-notification-profile fcoe-cnp;
}
xe-0/0/31 {
    forwarding-class-set {
        fcoe-pg {
            output-traffic-control-profile fcoe-tcp;
        }
    }
    congestion-notification-profile fcoe-cnp;
}
xe-0/0/32 {
    forwarding-class-set {
        fcoe-pg {
            output-traffic-control-profile fcoe-tcp;
        }
    }
    congestion-notification-profile fcoe-cnp;
}
xe-0/0/33 {
    forwarding-class-set {
        fcoe-pg {
            output-traffic-control-profile fcoe-tcp;
        }
    }
    congestion-notification-profile fcoe-cnp;
}
ae1 {
    forwarding-class-set {
        fcoe-pg {
```

```
            output-traffic-control-profile fcoe-tcp;
        }
    }
    congestion-notification-profile fcoe-cnp;
}
```

Meaning The **show configuration class-of-service interfaces** command lists the class of service configuration for all interfaces. For each interface, the command output includes:

- The name of the interface (for example, **ae0** or **xe-0/0/30**)
- The name of the forwarding class set associated with the interface (**fcoe-pg**)
- The name of the traffic control profile associated with the interface (output traffic control profile, **fcoe-tcp**)
- The name of the congestion notification profile associated with the interface (**fcoe-cnp**)

NOTE: Interfaces that are members of a LAG are not shown individually. The LAG or MC-LAG CoS configuration is applied to all interfaces that are members of the LAG or MC-LAG. For example, the interface CoS configuration output on MC-LAG Switches S1 and S2 shows the LAG CoS configuration but does not show the CoS configuration of the member interfaces separately. The interface CoS configuration output on FCoE Transit Switches TS1 and TS2 shows the LAG CoS configuration but also shows the configuration for interfaces xe-0/0/30, xe-0/0/31, xe-0/0/32, and xe-0/0/33, which are not members of a LAG.

Verifying That the Interfaces Are Correctly Configured

Purpose Verify that the LAG membership, MTU, VLAN membership, and port mode of the interfaces are correct. The verification output on MC-LAG Switches S1 and S2 differs from the output on FCoE Transit Switches TS1 and TS2.

Action List the interface configuration on MC-LAG Switches S1 and S2 using the operational mode command **show configuration interfaces**:

```
user@switch> show configuration interfaces
xe-0/0/10 {
    ether-options {
        802.3ad ae0;
    }
}
xe-0/0/11 {
    ether-options {
        802.3ad ae0;
    }
```

```
}
xe-0/0/20 {
    ether-options {
        802.3ad ae1;
    }
}
xe-0/0/21 {
    ether-options {
        802.3ad ae1;
    }
}
ae0 {
    mtu 2180;
    unit 0 {
        family ethernet-switching {
            port-mode trunk;
            vlan {
                members fcoe_vlan;
            }
        }
    }
}
ae1 {
    mtu 2180;
    unit 0 {
        family ethernet-switching {
            port-mode trunk;
            vlan {
                members fcoe_vlan;
            }
        }
    }
}
```

List the interface configuration on FCoE Transit Switches TS1 and TS2 using the operational mode command **show configuration interfaces**:

```
user@switch> show configuration interfaces
```

```
xe-0/0/25 {
    ether-options {
        802.3ad ae1;
    }
}
xe-0/0/26 {
    ether-options {
        802.3ad ae1;
    }
}
xe-0/0/30 {
    mtu 2180;
    unit 0 {
        family ethernet-switching {
            port-mode tagged-access;
            vlan {
                members fcoe_vlan;
            }
        }
```

```
    }
}
xe-0/0/31 {
    mtu 2180;
    unit 0 {
        family ethernet-switching {
            port-mode tagged-access;
            vlan {
                members fcoe_vlan;
            }
        }
    }
}
xe-0/0/32 {
    mtu 2180;
    unit 0 {
        family ethernet-switching {
            port-mode tagged-access;
            vlan {
                members fcoe_vlan;
            }
        }
    }
}
xe-0/0/33 {
    mtu 2180;
    unit 0 {
        family ethernet-switching {
            port-mode tagged-access;
            vlan {
                members fcoe_vlan;
            }
        }
    }
}

ae1 {
    mtu 2180;
    unit 0 {
        family ethernet-switching {
            port-mode trunk;
            vlan {
                members fcoe_vlan;
            }
        }
    }
}
```

Meaning The **show configuration interfaces** command lists the configuration of each interface by interface name.

For each interface that is a member of a LAG, the command lists only the name of the LAG to which the interface belongs.

For each LAG interface and for each interface that is not a member of a LAG, the command output includes:

- The MTU (**2180**)
- The unit number of the interface (**0**)
- The port mode (**trunk** mode for interfaces that connect two switches, **tagged-access** mode for interfaces that connect to FCoE hosts)
- The name of the VLAN in which the interface is a member (**fcoe_vlan**)

Verifying That FIP Snooping Is Enabled on the FCoE VLAN on FCoE Transit Switches TS1 and TS2 Access Interfaces

Purpose Verify that FIP snooping is enabled on the FCoE VLAN access interfaces. FIP snooping is enabled only on the FCoE access interfaces, so it is enabled only on FCoE Transit Switches TS1 and TS2. FIP snooping is not enabled on MC-LAG Switches S1 and S2 because FIP snooping is done at the Transit Switch TS1 and TS2 FCoE access ports.

Action List the port security configuration on FCoE Transit Switches TS1 and TS2 using the operational mode command **show configuration ethernet-switching-options secure-access-port**:

```
user@switch> show configuration ethernet-switching-options secure-access-port
interface ae1.0 {
    fcoe-trusted;
}
vlan fcoe_vlan {
    examine-fip {
        examine-vn2vn {
            beacon-period 90000;
        }
    }
}
```

Meaning The **show configuration ethernet-switching-options secure-access-port** command lists port security information, including whether a port is trusted. The command output shows that:

- LAG port **ae1.0**, which connects the FCoE transit switch to the MC-LAG switches, is configured as an FCoE trusted interface. FIP snooping is not performed on the member interfaces of the LAG (xe-0/0/25 and xe-0/0/26).
- FIP snooping is enabled (**examine-fip**) on the FCoE VLAN (**fcoe_vlan**), the type of FIP snooping is VN2VN_Port FIP snooping (**examine-vn2vn**), and the beacon period is set to **90000** milliseconds. On Transit Switches TS1 and TS2, all interface members of the FCoE VLAN perform FIP snooping unless the interface is configured as FCoE trusted. On Transit Switches TS1 and TS2, interfaces xe-0/0/30, xe-0/0/31, xe-0/0/32, and xe-0/0/33 perform FIP snooping because they are not configured as FCoE trusted. The interface members of LAG ae1 (xe-0/0/25 and xe-0/0/26) do not perform FIP snooping because the LAG is configured as FCoE trusted.

Verifying That the FIP Snooping Mode Is Correct on FCoE Transit Switches TS1 and TS2

Purpose Verify that the FIP snooping mode is correct on the FCoE VLAN. FIP snooping is enabled only on the FCoE access interfaces, so it is enabled only on FCoE Transit Switches TS1 and TS2. FIP snooping is not enabled on MC-LAG Switches S1 and S2 because FIP snooping is done at the Transit Switch TS1 and TS2 FCoE access ports.

Action List the FIP snooping configuration on FCoE Transit Switches TS1 and TS2 using the operational mode command **show fip snooping brief**:

```
user@switch> show fip snooping brief
VLAN: fcoe_vlan,    Mode: VN2VN Snooping
  FC-MAP: 0e:fd:00
…
```

NOTE: The output has been truncated to show only the relevant information.

Meaning The **show fip snooping brief** command lists FIP snooping information, including the FIP snooping VLAN and the FIP snooping mode. The command output shows that:

- The VLAN on which FIP snooping is enabled is **fcoe_vlan**
- The FIP snooping mode is VN2VN_Port FIP snooping (**VN2VN Snooping**)

Verifying That IGMP Snooping Is Disabled on the FCoE VLAN

Purpose Verify that IGMP snooping is disabled on the FCoE VLAN on all four switches.

Action List the IGMP snooping protocol information on each of the four switches using the **show configuration protocols igmp-snooping** command:

```
user@switch> show configuration protocols igmp-snooping
vlan fcoe_vlan {
    disable;
}
```

Meaning The **show configuration protocols igmp-snooping** command lists the IGMP snooping configuration for the VLANs configured on the switch. The command output shows that IGMP snooping is disabled on the FCoE VLAN (**fcoe_vlan**).

See Also
- *Configuring Link Aggregation*

- *Example: Configuring CoS PFC for FCoE Traffic*

Multichassis Link Aggregation for IPv6 Through NDP

Neighbor Discovery Protocol (NDP) is an IPv6 protocol that enables nodes on the same link to advertise their existence to their neighbors and to learn about the existence of their neighbors. NDP is built on top of Internet Control Message Protocol version 6 (ICMPv6). It replaces the following IPv4 protocols: Router Discovery (RDISC), Address Resolution Protocol (ARP), and ICMPv4 redirect.

You can use NDP in a multichassis link aggregation group (MC-LAG) active-active configuration on switches.

Neighbor Discovery Messages for MC-LAGs

NDP on MC-LAGs uses the following message types:

- Neighbor solicitation (NS)—Messages used for address resolution and to test reachability of neighbors.

  A host can verify that its address is unique by sending a neighbor solicitation message destined to the new address. If the host receives a neighbor advertisement in reply, the address is a duplicate.

- Neighbor advertisement (NA)—Messages used for address resolution and to test reachability of neighbors. Neighbor advertisements are sent in response to neighbor solicitation messages.

NDP Functions and Configuration Requirements on MC-LAGs

The functions provided through NDP on MC-LAGs on the switches are:

- Address resolution
- Neighbor unreachability detection

The requirements for an NDP configuration on MC-LAG on the switches include:

- An active-active configuration
- Virtual Router Redundancy Protocol (VRRP)

Related Documentation

- show l2-learning redundancy-groups on page 461

CHAPTER 3

Increasing ARP and Network Discovery Entries for Enhanced MC-LAG and Layer 3 VXLAN Topologies

Increasing ARP and Network Discovery Protocol Entries for Enhanced MC-LAG and Layer 3 VXLAN Topologies

Understanding the Need for an Increase in ARP and Network Discovery Protocol (NDP) Entries

The number of ARP and NDP entries has increased to 256,000 to improve enhanced MC-LAG and Layer 3 VXLAN scenarios.

Here are some enhanced MC-LAG and Layer 3 VXLAN scenarios in which an increase in ARP and NDP entries is needed:

- Enhanced MC-LAG topology with a large number of MC-AE interfaces that contain a large number of members per chassis.
- Non-collapsed spine-leaf topology, in which the leaf devices operate as Layer 2 gateways and handle traffic within the VXLAN, and the spine devices operate as Layer 3 gateways and handle traffic between the VXLANs using IRB interfaces.

  In this scenario, the increase in ARP and NDP entries is needed at the spine level.
- Leaf devices that operate as both Layer 2 and Layer 3 gateways.

  In this scenario, the transit spine devices provide Layer 3 routing functioning only, and the increased number of ARP and NDP entries in needed only at the leaf level.

Configuring Enhanced MC-LAG and Layer 3 VXLAN with Increased Number of ARP and NDP Entries Using IPv4 Transport

To increase the number of ARP and NDP entries using IPv4 transport, follow these steps. We recommend that you use the values provided in this procedure for optimal performance:

1. Enable the **arp-enhanced-scale** statement:

   ```
   [edit system]
   user@switch# set arp-enhanced-scale
   ```

2. Configure the maximum number of routes to be stored in the ARP cache.

   ```
   [edit system]
   user@switch# set arp-system-cache-limit number
   ```

   For example:

   ```
   [edit system]
   user@switch# set arp-system-cache-limit 2000000
   ```

3. Configure the amount of time between ARP updates.

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

   For example:

   ```
   [edit system]
   user@switch# set arp aging-timer 20
   ```

4. Enable enhanced convergence on the MC-AE interface:

   ```
   [edit interfaces]
   user@switch# set interface-name aggregated-ether-options mc-ae enhanced-convergence
   ```

5. Enable enhanced convergence on the IRB interface that you have configured as part of an MC-AE.

   ```
   [edit interfaces]
   user@switch# set irb unit number enhanced-convergence
   ```

6. Specify the amount of time that elapses before the MAC table entries are timed out and entries are deleted from the table.

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

   For example:

   ```
   [edit protocols l2-learning]
   user@switch# set global-mac-table-aging-time 3600
   ```

7. Specify the amount time that elapses before the entries in the MAC-IP bindings database are timed out and deleted.

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

For example:

```
[edit protocols l2-learning]
user@switch# set global-mac-ip-table-aging-time 1200
```

8. (Optional) If you have a Layer 3 VXLAN configuration, for each leaf device, specify the amount of time that elapses before the MAC table entries are timed out and entries are deleted from the table.

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

For example:

```
[edit protocols l2-learning]
user@switch# set global-mac-table-aging-time 3600
```

9. Reboot the device in order for these changes to take effect.

```
user@switch# request system reboot
```

Configuring Enhanced MC-LAG and Layer 3 VXLAN with Increased Number of ARP and NDP Entries Using IPv6 Transport

To increase the number of ARP and Network Discovery Protocol entries using IPv6 transport. We recommend that you use the values provided in this procedure for optimal performance:

1. Enable the **arp-enhanced-scale** statement:

```
[edit system]
user@switch# set arp-enhanced-scale
```

2. Specify the maximum system cache size for IPv6 next-hop addresses.

```
[edit system]
user@switch# set nd-system-cache-limitnumber
```

For example:

```
[edit system]
user@switch# set nd-system-cache-limit 2000000
```

3. Set the stale timer for IPv6 neighbor reachability confirmation.

```
[edit interfaces]
user@switch# set irb unit 1 family inet6 nd6-stale-timeseconds
```

For example:

```
[edit interfaces]
user@switch# set irb unit 1 family inet6 nd6-stale-time 1200
```

4. Enable enhanced convergence on the MC-AE interface:

```
[edit interfaces]
user@switch# set interface-name aggregated-ether-options mc-ae enhanced-convergence
```

5. Enable enhanced convergence on the IRB interface that you have configured as part of an MC-AE.

```
[edit interfaces]
user@switch# set irb unit number enhanced-convergence
```

6. Specify the amount of time that elapses before the MAC table entries are timed out and entries are deleted from the table.

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

For example:

```
[edit protocols l2-learning]
user@switch# set global-mac-table-aging-time 3600
```

7. Specify the amount time that elapses before the entries in the MAC-IP bindings database are timed out and deleted.

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

For example:

```
[edit protocols l2-learning]
user@switch# set global-mac-ip-table-aging-time 1200
```

8. (Optional) If you have a Layer 3 VXLAN configuration, for each leaf device, specify the amount of time that elapses before the MAC table entries are timed out and entries are deleted from the table.

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

For example:

```
[edit protocols l2-learning]
user@switch# set global-mac-table-aging-time 3600
```

9. Reboot the device in order for these changes to take effect.

```
user@switch# request system reboot
```

CHAPTER 4

Enabling High Availability in Layer 2 Networks Using Active-Active Bridging for MC-LAG

High Availability in Layer 2 Networks Using Active-Active Bridging for MC-LAG

Multichassis Link Aggregation on Logical Systems Overview

On MX Series routers, EX9200, and QFX10000 switches, multichassis link aggregation (MC-LAG) enables a device to form a logical LAG interface with two or more other devices. MC-LAG provides additional benefits over traditional LAG in terms of node-level redundancy, multihoming support, and a loop-free Layer 2 network without running Spanning Tree Protocol (STP). The MC-LAG devices use Inter-Chassis Control Protocol (ICCP) to exchange the control information between two MC-LAG network devices. Starting in Junos OS Release 14.1, you can configure MC-LAG interfaces on logical systems within a router. Starting with Junos OS Release 15.1, you can configure MC-LAG interfaces on logical systems on EX9200 switches.

NOTE: On QFX10008 switches, Layer 2 and Layer 3 IRB interfaces are not supported under the **[edit logical-systems]** hierarchy.

To configure ICCP for MC-LAG interfaces on logical systems, include the **iccp** statement at the **[edit logical-systems *logical-system-name* protocols]** hierarchy level. To view ICCP information for MC-LAG on logical systems, use the **show iccp logical-system *logical-system-name*** command. To view ARP statistics or remote MAC addresses for the multichassis aggregated Ethernet nodes for all or specified redundancy groups on a logical system, use the **show l2-learning redundancy-groups *group-name* logical-system *logical-system-name* (arp-statistics | remote-macs)** command. To view neighbor discovery (ND) statistical details for multichassis aggregated Ethernet nodes on redundancy groups of a logical group, use the **show l2-learning redundancy-groups *group-name* logical-system *logical-system-name* nd-statistics** command.

Logical systems enable effective, optimal segregation of a single router or switch into multiple virtual partitions, which can be configured and managed by diversified entities. Logical systems perform a subset of the actions of a physical router or switch and have their own unique routing tables, interfaces, policies, and routing instances. A set of logical systems within a single router or switch can handle the functions previously performed by several small routers or switches. As shown on the right side of Figure 5 on page 122, a set of logical systems within a single router can handle the functions previously performed by several small routers.

Figure 5: Comparison of Devices With and Without Logical Systems

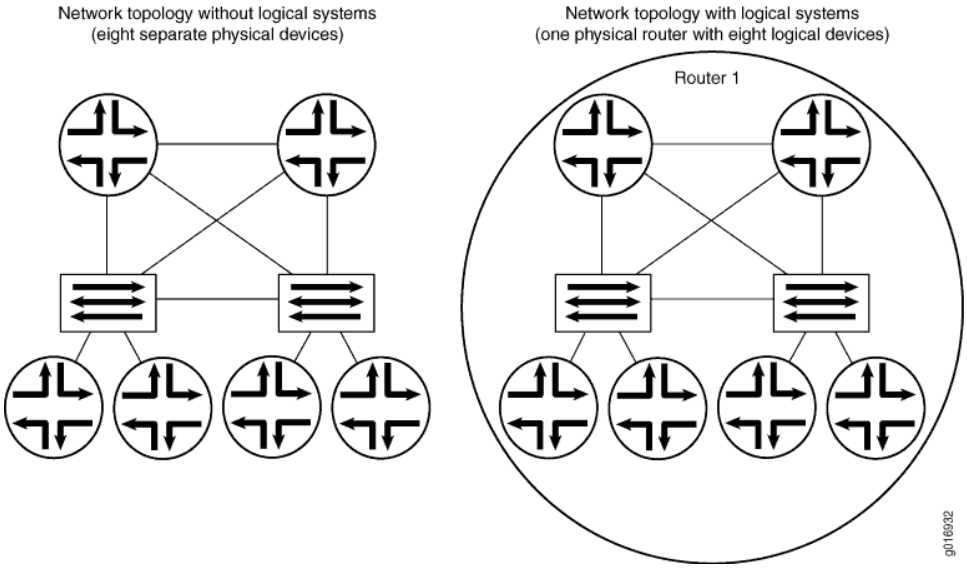

In a network deployment that contains MC-LAG interfaces, you can configure such interfaces on logical systems contained within a router or switch. When you configure multichassis aggregated Ethernet interfaces on a logical system, you must ensure that these interfaces are added with the same multichassis aggregated Ethernet identification number and redundancy group identifier for the MC-LAG on both the peers or devices that are connected by the multichassis aggregated Ethernet interfaces. It is not necessary

to specify the same logical system name on both the peers; however, you must ensure that ICCP to associate the routing or switching devices contained in a redundancy group is defined on both the peers within the logical systems of the devices. Such a configuration ensures that all the packets are transmitted using ICCP within the logical system network. The logical system information is added and removed by the ICCP process to prevent each packet from containing the logical system details. This behavior enables multiple disjoint users to employ MC-LAG capabilities within their networks transparently and seamlessly. A unique ICCP definition for a logical system is created, thereby enabling you to completely manage the ICCP parameters on one logical system without the need for access permissions to view other logical system networks on the same device. Configuration of MC-LAG interfaces on logical systems enables MC-LAG to be used across multiple routing tables and switch forwarding tables in active-active and active-standby modes of MC-LAG interfaces.

Because the Layer 2 address learning process supports logical systems, the ARP, neighbor discovery, and MAC synchronization packets that are traversing a multichassis aggregated Ethernet interface use the logical system:routing instance (LS:RI) combination to map the packets to the correct routing instance in a logical system. Link Aggregation Control Protocol (LACP) does not require the LS-RI combination to be identified because it operates on physical interfaces and is unique within a chassis. For a service, in the set of provider edge (PE) routers providing the service, the service ID distinguishes the routing instances in a logical system because it is unique for a logical system across a routing instance. MC-LAG is configured on the aggregated Ethernet (ae-) bundle interface. An ae- interface is a logical interface and is globally unique, which causes the MC-LAG configuration to be exclusive and separate for a router or switch. You can add ae- interfaces in an MC-LAG configuration to be part of a logical system and use it throughout that particular logical system.

Sample Configuration Scenario for MC-LAG on Logical Systems

Consider a sample scenario in which two MX Series routers, MX1 and MX2, are connected using an aggregated Ethernet interface that is enabled with MC-LAG. The peers in an MC-LAG use an interchassis link-protection link (ICL-PL) to replicate forwarding information across the peers. Additionally, ICCP propagates the operational state of MC-LAG members through the ICL-PL. The two PE devices, MX1 and MX2, each have a LAG connected to the CE devices, CE1 and CE2. Four logical systems are defined on each of the PE devices, MX1 and MX2. CE-1 and CE-2 can be part of the same VLAN with the same VLAN ID and located in the same IP subnet for MC-LAG in two different logical systems. All four logical system entities can work independently in MX1 and MX2.

The ICCP process can manage multiple client-server connections with its peer ICCP instances based on the ICCP configuration for the logical system:routing instance (LS-RI) combinations. Each ICCP connection is associated with an LS-RI combination. For example, with two routing instances, IP1 and IP2, on each of the logical systems, LS1 and LS2, the following mapping is performed for ICCP settings:

[ICCP] (LS1) (IP1) < = = > (IP2) (LS1) [ICCP] within LS1 network.

[ICCP] (LS2) (IP1) < = = > (IP2) (LS2) [ICCP] within LS2 network.

An ICCP instance in a logical system is linked with the ICCP instance of the peer logical system. The ICCP application transmits the relevant routing index depending on the LS:RI combination to the BFD process, when BFD is configured in your topology.

Figure 6 on page 124 shows the interconnection among logical systems on MX Series routers configured with MC-LAG.

Figure 6: Logical Systems with MC-LAG

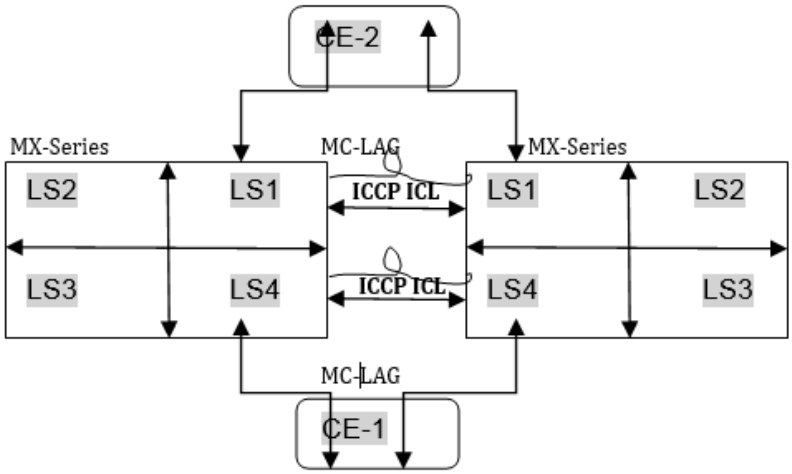

The Layer 2 address learning process (l2ald) transmits and receives Address Learning Protocol (ARP), neighbor discovery, and MAC synchronization packets with the LS-RI information. When the peer MAC synchronization packets are received, l2ald decodes the logical system details from the packet and determines whether an identical logical system has been previously created on the router. If a match is found for the logical system, the MAC forwarding entry for the corresponding bridge table for an interface bridge domain is created. If the logical system in the received packet does not match the defined logical system on the device, for the MAC synchronization packet, the default logical instance is used for processing. Similarly, upon receipt of the ARP and neighbor discovery packets, l2ald decapsulates the logical system information from the packets and determines if the corresponding logical instance has been previously created. If a match is found for the logical system, the ARP and neighbor discovery packets are processed according to the Layer 3 index that is unique in the system. The programming kernel entry might not require any logical system information since it is programmed on a Layer 3 index which is unique in the system. If the logical system in the received packet does not match the defined logical system on the device, for the ARP and neighbor discovery packets, the default logical instance is used for processing. The routing instance is determined using the service ID attribute. The logical system information is forwarded to ICCP, which in turn identifies the appropriate ICCP interface for the logical system and sends packets over it.

Guidelines for Configuring MC-LAG on Logical Systems

Keep the following points in mind while configuring MC-LAG interfaces on logical systems:

- You cannot use a single chassis to function as a provider edge (PE) device and a customer edge (CE) device in different logical systems.
- You cannot use a single chassis to function as two PE devices by configuring logical systems on the chassis and ICCP. ICL links between the two logical systems because the multichassis aggregated Ethernet ID is unique in a router or switch.
- Logical interfaces (IFLs) on the same **mc-ae** interface cannot be configured across multiple logical systems. In other words, in a multichassis link aggregation (MC-LAG) with both logical systems and logical interfaces (such as **mc-ae ae0 unit 0**), the same logical interface cannot be shared between logical systems.
- IGMP snooping in MC-LAG topologies with logical systems is not supported.
- VPLS and VPN protocols with MC-LAG in active-standby mode is not supported.
- Logical system information is not communicated to the peer chassis because this detail is derived from an ICCP instance.

Active-Active Bridging and VRRP over IRB Functionality Overview

Active-active bridging and VRRP over IRB support extends multichassis link aggregation group (MC-LAG) by adding the following functionality to MX Series routers and QFX Series switches:

- Interchassis link (ICL) pseudowire interface or Ethernet interface (ICL-PL field) for active-active bridging
- Active-active bridging
- VRRP over IRB for active-active bridging
- A single bridge domain not corresponding to two redundancy group IDs

How Active-Active Bridging over IRB Functionality Works

Active-Active bridging over IRB functionality uses the address resolution protocol (ARP) Active-Active MC-LAG.

Suppose one of the PE routers issues an ARP request and another PE router gets the response and, because of the aggregated Ethernet distribution logic, the ARP resolution is not successful. Junos OS uses ARP response packet snooping to perform active-active multichassis link aggregation group support, providing synchronization without the need to maintain any specific state.

Address Resolution Protocol Active-Active MC-LAG Support Methodology

Suppose one of the PE routers issues an ARP request and another PE router gets the response and, because of the aggregated Ethernet distribution logic, the ARP resolution is not successful. Junos OS uses ARP response packet snooping to perform active-active multichassis link aggregation group support, providing easy synchronization without the need to maintain any specific state.

Benefits of Active-Active Bridging and VRRP over IRB Functionality

Benefits of active-active bridging and VRRP over IRB functionality include:

- An MC-LAG reduces operational expenses by providing active-active links with a LAG, eliminates the need for Spanning Tree Protocol (STP), and provides faster Layer 2 convergence upon link and device failures.
- An MC-LAG adds node-level redundancy to the normal link-level redundancy that a LAG provides. An MC-LAG improves network resiliency, which reduces network down time as well as expenses.
- In data centers, it is desirable for servers to have redundant connections to the network. You probably want active-active connections along with links from any server to at least two separate routers.
- An MC-LAG allows you to bond two or more physical links into a logical link between two routers or between a server and a router, which improves network efficiency. An MC-LAG enables you to load-balance traffic on multiple physical links. If a link fails, the traffic can be forwarded through the other available link, and the logical aggregated link remains in the UP state.

Where Can I Use Active-Active Bridging and VRRP over IRB Functionality?

Active-active bridging and Virtual Router Redundancy Protocol (VRRP) over integrated routing and bridging (IRB) is supported on MX Series routers and QFX Series switches.

MC-LAG Functions in an Active-Active Bridging Domain

The following functions are supported for MC-LAG in an active-active bridging domain:

- MC-LAG is supported only between two chassis, using an interchassis link (ICL) pseudowire interface or Ethernet interface (ICL-PL field) for **active-active bridging**, and **active-active bridging** VRRP over IRB for **active-active bridging**.
- For VPLS networks, you can configure the aggregated Ethernet (aeX) interfaces on MC-LAG devices with the **encapsulation ethernet-vpls** statement to use Ethernet VPLS encapsulation on Ethernet interfaces that have VPLS enabled and that must accept packets carrying standard Tag Protocol ID (TPID) values or the **encapsulation vlan-vpls** statement to use Ethernet VLAN encapsulation on VPLS circuits.
- Layer 2 circuit functionalities are supported with **ethernet-ccc** as the encapsulation mode.
- Network topologies in a triangular and square pattern are supported. In a triangular network design, with equal-cost paths to all redundant nodes, slower, timer-based convergence can possibly be prevented. Instead of indirect neighbor or route loss detection using hellos and dead timers, you can identify the physical link loss and denote a path as unusable and reroute all traffic to the alternate equal-cost path. In a square network design, depending on the location of the failure, the routing protocol might converge to identify a new path to the subnet or the VLAN, causing the convergence of the network to be slower.
- Interoperation of Link Aggregation Control Protocol (LACP) for MC-LAG devices is supported. LACP is one method of bundling several physical interfaces to form one logical interface. When LACP is enabled, the local and remote sides of the aggregated Ethernet links exchange protocol data units (PDUs), which contain information about the state of the link. You can configure Ethernet links to actively transmit PDUs, or you can configure the links to passively transmit them, sending out LACP PDUs only when the links receive the PDUs from another link. One side of the link must be configured as active for the link to be up.
- Active-standby mode is supported using LACP. When an MC-LAG operates in the active-standby mode, one of the router's ports only becomes active when failure is detected in the active links. In this mode, the provider edge (PE) routers perform an election to determine the active and standby routers.
- Configuration of the pseudowire status type length variable (TLV) is supported. The pseudowire status TLV is used to communicate the status of a pseudowire back and forth between two PE routers. The pseudowire status negotiation process ensures that a PE router reverts back to the label withdraw method for pseudowire status if its remote PE router neighbor does not support the pseudowire status TLV.
- The MC-LAG devices use Inter-Chassis Control Protocol (ICCP) to exchange the control information between two MC-LAG network devices.

Points to Remember When Configuring MC-LAG Active-Active Bridge Domains

Keep the following points in mind when you configure MC-LAG in an active-active bridging domain:

- A single bridge domain cannot be associated with two redundancy groups. You cannot configure a bridge domain to contain logical interfaces from two different multichassis aggregated Ethernet interfaces and associate them with different redundancy group

IDs by using the **redundancy group group-id** statement at the **[edit interfaces ae*X* aggregated-ether-options]** hierarchy level.

- You must configure logical interfaces in a bridge domain from a single multichassis aggregated Ethernet interface and associate it with a redundancy group. You must configure a service ID by including the **service-id *vid*** statement at the **[edit bridge-domains *bd-name*]** hierarchy level for multichassis aggregated Ethernet interfaces if you configure logical interfaces on multichassis aggregated Ethernet interfaces that are part of the bridge domain.

More Data Traffic Forwarding Rules

In active-active bridging and VRRP over IRB topographies, network interfaces are categorized into three different interface types, as follows:

S-Links—Single-homed link (S-Link) terminating on MC-LAG-N device or MC-LAG in active-standby mode. In Figure 10 on page 133, interfaces ge-0/0/0.0 and ge-1/0/0.0 are S-Links.

MC-Links—MC-LAG links. In Figure 10 on page 133, interface ae0.0 is the MC-Link.

ICL—Interchassis link.

Based on incoming and outgoing interface types, some constraints are added to the Layer 2 forwarding rules for MC-LAG configurations, as described in the data traffic forwarding rules. Note that if only one of the MC-LAG member link is in the UP state, it is considered an S-Link.

The following data traffic forwarding rules apply:

1. When an MC-LAG network receives a packet from a local MC-Link or S-Link, the packet is forwarded to other local interfaces, including S-Links and MC-Links based on the normal Layer 2 forwarding rules and on the configuration of the **mesh-group** and **no-local-switching** statements. If MC-Links and S-Links are in the same mesh group and their **no-local-switching** statements are enabled, the received packets are only forwarded upstream and not sent to MC-Links and S-Links.

   NOTE: **The functionality described in Rule 2 is *not* supported.**

2. The following circumstances determine whether or not an ICL receives a packet from a local MC-Link or S-Link:

   a. If the peer MC-LAG network device has S-Links or MC-LAGs that do not reside on the local MC-LAG network device

   b. Whether or not interfaces on two peering MC-LAG network devices are allowed to talk to each other only if both a. and b. are true. Traffic is always forwarded to the ICL.

3. When an MC-LAG network receives a packet from the ICL, the packet is forwarded to all local S-Links and active MC-LAGs that do not exist in the MC-LAG network that the packet comes from.

4. NOTE: The topology shown in Figure 7 on page 129 is *not* supported.

   In certain cases, for example the topology shown in Figure 7 on page 129, there could be a loop caused by the ICL. To break the loop, one of the following mechanisms could be used:

   a. Run certain protocols, such as STP. In this case, whether packets received on one ICL are forwarded to other ICLs is determined by using Rule 3.

   b. Configure the ICL to be fully meshed among the MC-LAG network devices. In this case, traffic received on the ICL would not be forwarded to any other ICLs.

   In either case, duplicate packets could be forwarded to the MC-LAG clients. Consider the topology shown in Figure 7 on page 129, where if network routing instance N1 receives a packet from ge-0/0/0.0, it could be flooded to ICL1 and ICL3.

   When receiving from ICL1 and ICL3, network routing instances N3 and N2 could flood the same packet to MCL2, as shown in Figure 7 on page 129. To prevent this from happening, the ICL designated forwarder should be elected between MC-LAG peers, and traffic received on an ICL could be forwarded to the active-active MC-LAG client by the designated forwarder only.

Figure 7: Loop Caused by the ICL Links

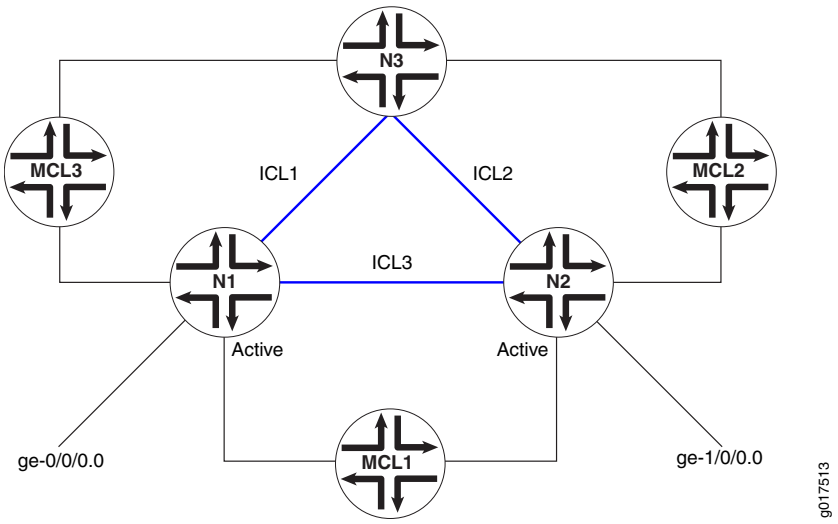

5. When received from an ICL, traffic should not be forwarded to the core-facing client link connection between two provider edge (PE) devices (MC-Link) if the peer chassis's (where the traffic is coming from) MC-Link is UP.

How to Configure MC-LAG Active-Active Bridge Domains

For a MC-LAG configured in an active-active bridge domain and with VRRP configured over an IRB interface, you must include the **accept-data** statement at the **[edit interfaces *interface-name* unit *logical-unit-number* family inet address *address* vrrp-group *group-id*]** hierarchy level to enable the router that functions as the master router to accept all packets destined for the virtual IP address.

On an MC-LAG, if you modify the source MAC address to be the virtual MAC address, you must specify the virtual IP address as the source IP address instead of the physical IP address. In such a case, the **accept-data** option is required for VRRP to prevent ARP from performing an incorrect mapping between IP and MAC addresses for customer edge (CE) devices. The **accept-data** attribute is needed for VRRP over IRB interfaces in MC-LAG to enable OSPF or other Layer 3 protocols and applications to work properly over multichassis aggregated Ethernet (mc-aeX) interfaces.

NOTE: On an MC-LAG, the unit number associated with aggregated Ethernet interfaces on provider edge router PE1 must match the unit number associated with aggregated Ethernet interfaces on provider edge router PE2. If the unit numbers differ, MAC address synchronization does not happen. As a result, the status of the MAC address on the remote provider edge router remains in a pending state.

If you are using the VRRP over IRB or RVI method to enable Layer 3 functionality, you must configure static ARP entries for the IRB or RVI interface of the remote MC-LAG peer to allow routing protocols to run over the IRB or RVI interfaces.

MAC Address Management

If an MC-LAG is configured to be active-active, upstream and downstream traffic could go through different MC-LAG network devices. Since the media access control (MAC) address is learned only on one of the MC-LAG network devices, the reverse direction's traffic could be going through the other MC-LAG network and be flooded unnecessarily. Also, a single-homed client's MAC address is only learned on the MC-LAG network device it is attached to. If a client attached to the peer MC-LAG network needs to communicate with that single-homed client, then traffic would be flooded on the peer MC-LAG network device. To avoid unnecessary flooding, whenever a MAC address is learned on one of the MC-LAG network devices, it gets replicated to the peer MC-LAG network device. The following conditions should be applied when MAC address replication is performed:

- MAC addresses learned on an MC-LAG of one MC-LAG network device should be replicated as learned on the same MC-LAG of the peer MC-LAG network device.
- MAC addresses learned on single-homed customer edge (CE) clients of one MC-LAG network device should be replicated as learned on the ICL-PL interface of the peer MC-LAG network device.

- MAC addresses learned on MC-LAG VE clients of one MC-LAG network device should be replicated as learned on the corresponding VE interface of the peer MC-LAG network device.
- MAC address learning on an ICL is disabled from the data path. It depends on software to install MAC addresses replicated through Inter-Chassis Control Protocol (ICCP).

MAC Aging

MAC aging support in Junos OS extends aggregated Ethernet logic for a specified MC-LAG. A MAC address in software is deleted until all Packet Forwarding Engines have deleted the MAC address. In the case of an MC-LAG, a remote provider edge is treated as a remote Packet Forwarding Engine and has a bit in the MAC data structure.

Layer 3 Routing

In general, when an MC-LAG is configured to provide Layer 3 routing functions to downstream clients, the MC-LAG network peers should be configured to provide the same gateway address to the downstream clients. To the upstream routers, the MC-LAG network peers could be viewed as either equal-cost multipath (ECMP) or two routes with different preference values.

Junos OS supports active-active MC-LAGs by using VRRP over IRB. Junos OS also supports active-active MC-LAGs by using IRB MAC address synchronization. You must configure IRB using the same IP address across MC-LAG peers. IRB MAC synchronization is supported on 32-bit interfaces and interoperates with earlier MPC and MIC releases.

To ensure that Layer 3 operates properly, instead of dropping the Layer 3 packet, the VRRP backup attempts to perform routing functions if the packet is received on an MC-LAG. A VRRP backup sends and responds to Address Resolution Protocol (ARP) requests.

For ARP, the same issue exists as with Layer 2 MAC addresses. Once ARP is learned, it must be replicated to the MC-LAG through ICCP. The peer must install an ARP route based on the ARP information received through ICCP.

For ARP aging, ARP requests on the MC-LAG peers can be aged out independently.

Topologies Supported for MC-LAG Active-Active Bridge Domains

The topologies shown in Figure 8 on page 132 and Figure 9 on page 132 are supported. These figures use the following abbreviations:

- Aggregated Ethernet (AE)
- Interchassis link (ICL)
- Multichassis link (MCL)

Figure 8: Single Multichassis Link

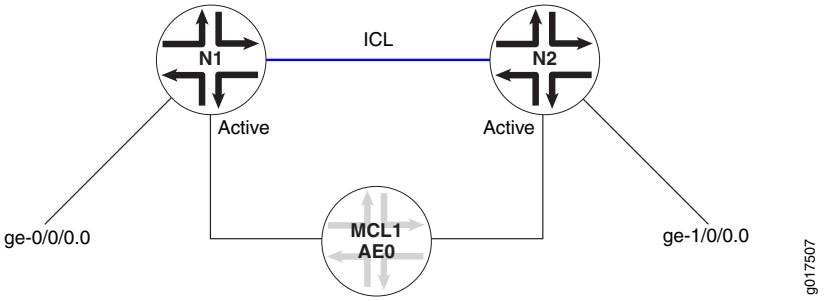

Figure 9: Dual Multichassis Link

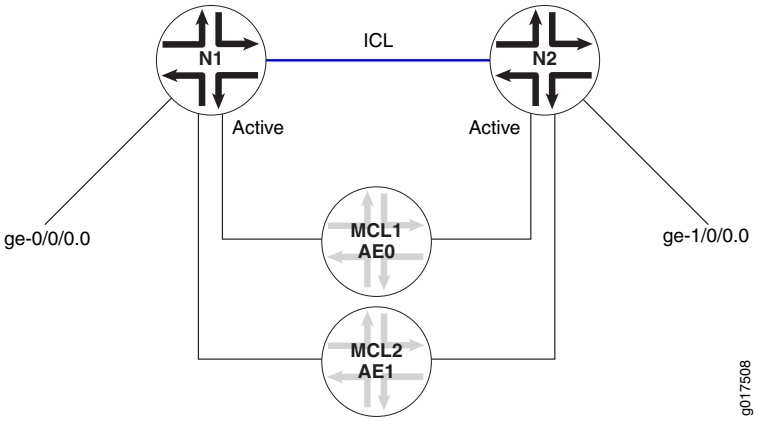

Potential Problems When Configuring MC-LAG Active-Active Bridge Domains

When configured to be active-active, the client device load-balances the traffic to the peering MC-LAG network devices. In a bridging environment, this could potentially cause the following problems:

- Traffic received on the MC-LAG from one MC-LAG network device could be looped back to the same MC-LAG on the other MC-LAG network device.
- Duplicated packets could be received by the MC-LAG client device.
- Traffic could be unnecessarily forwarded on the interchassis link.

To better illustrate the problems listed, consider Figure 10 on page 133, where an MC-LAG device MCL1 and single-homed clients ge-0/0/0.0 and ge-1/0/0.0 are allowed to talk to each other through an ICL. These problems could occur:

Figure 10: MC-LAG Device and Single-Homed Client

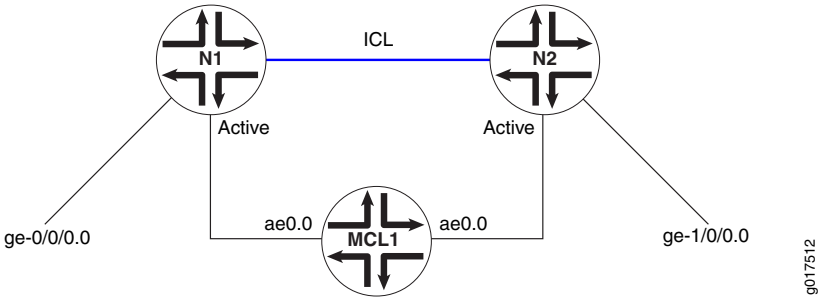

- Traffic received on network routing instance N1 from MCL1 could be flooded to ICL to reach network routing instance N2. Once it reaches network routing instance N2, it could flood again to MCL1.
- Traffic received on interface ge-0/0/0.0 could be flooded to MCL1 and ICL on network routing instance N1. Once network routing instance N2 receives such traffic from ICL, it could again be flooded to MCL1.
- If interface ge-1/0/0.0 does not exist on network routing instance N2, traffic received from interface ge-0/0/0.0 or MCL1 on network routing instance N1 could be flooded to network routing instance N2 through ICL unnecessarily since interface ge-0/0/0.0 and MCL1 could reach each other through network routing instance N1.

Restrictions When Configuring MC-LAG Active-Active Bridge Domains

In an IPv6 network, you cannot configure an MC-LAG in an active-active bridge domain if you specified the **vlan-id none** statement at the **[edit bridge-domain *bd-name*]** hierarchy level. The **vlan-id none** statement that enables the removal of the incoming VLAN tags identifying a Layer 2 logical interface when packets are sent over VPLS pseudowires is not supported for IPv6 packets in an MC-LAG.

The following functionality is *not* supported for MC-LAG active-active bridge domains:

- Virtual private LAN service (VPLS) within the core
- Bridged core
- Topology as described in Rule 4 of "More Data Traffic Forwarding Rules" on page 128
- Routed multichassis aggregated Ethernet interface, where the VRRP backup router is used in the edge of the network
- Track object, where in the case of an MC-LAG, the status of the uplinks from the provider edge can be monitored, and the MC-LAG can act on the status
- Mixed mode (active-active MC-LAG is supported on MX Series routers with MPC or MIC interfaces only)

  All interfaces in the bridge domain that are multichassis aggregated Ethernet active-active must be on MPCs or MICs.

The topologies shown in Figure 11 on page 134, Figure 12 on page 134, and Figure 13 on page 134 are *not* supported:

Figure 11: Interchassis Data Link Between Active-Active Nodes

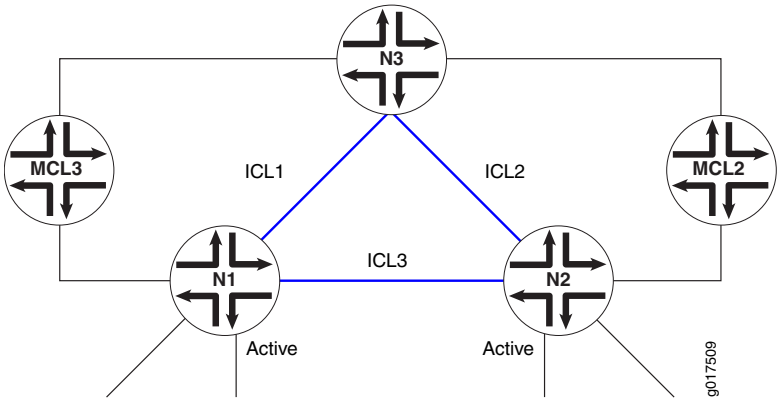

Figure 12: Active-Active MC-LAG with Single MC-LAG

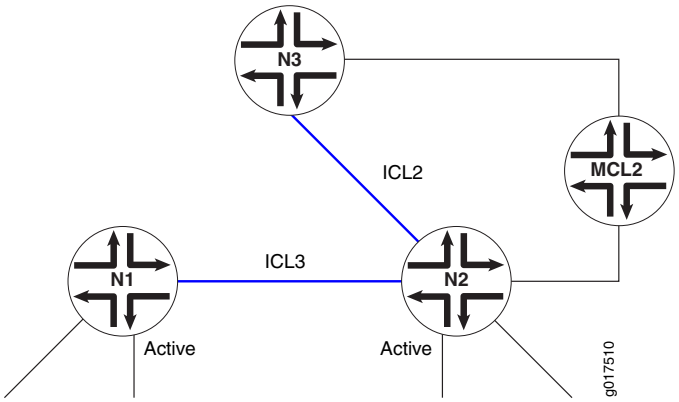

Figure 13: Active-Active MC-LAG with Multiple Nodes on a Single Multichassis Link

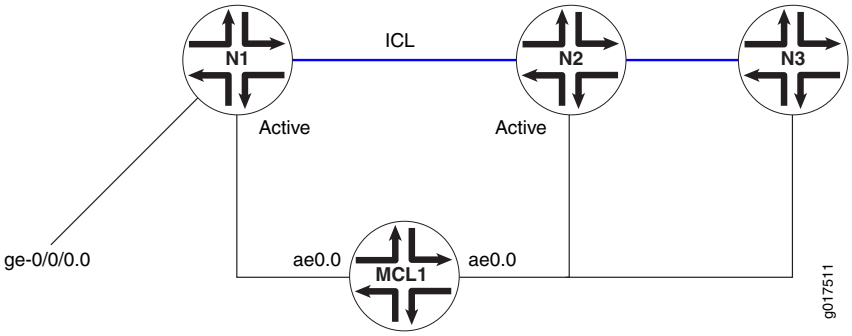

NOTE: A redundancy group cannot span more than two routers.

IGMP Snooping on Active-Active MC-LAG

IGMP Snooping on Active-Active MC-LAG

For multicast to work in an active-active MC-LAG scenario, the typical topology is as shown in Figure 14 on page 135 and Figure 15 on page 136 with interested receivers over

S-links and MC-Links. Starting in Junos OS Release 11.2, support is extended for sources connected over the Layer 2 interface.

If an MC-LAG is configured to be active-active, reports from MC-LAG clients could reach any of the MC-LAG network device peers. Therefore, the IGMP snooping module needs to replicate the states such that the Layer 2 multicast route state on both peers are the same. Additionally for S-Link clients, snooping needs to replicate these joins to its snooping peer, which in the case of Layer 3 connected source, passes this information to the PIM on IRB to enable the designated router to pull traffic for these groups,

The ICL should be configured as a router facing interface. For the scenario where traffic arrives through a Layer 3 interface, it is a requirement to have PIM and IGMP enabled on the IRB interface configured on the MC-LAG network device peers.

Figure 14: Multicast Topology with Source Connected Through Layer 3

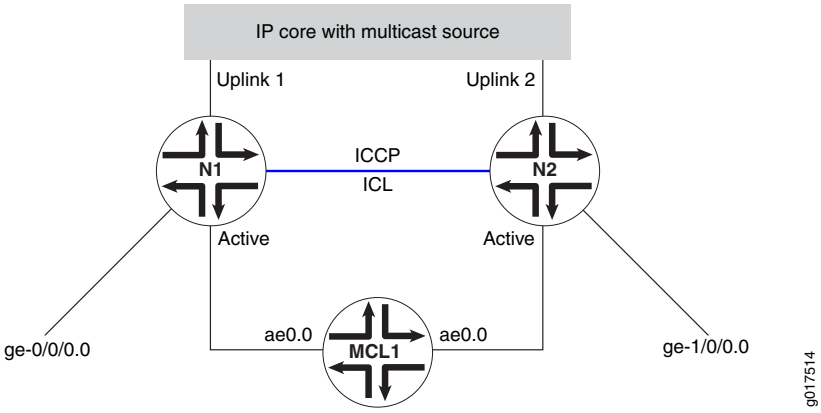

With reference to Figure 14 on page 135, either Device N1 or N2 becomes a designated router (for this example, N1 is the designated router). Router N1 therefore pulls the multicast traffic from the core. Once multicast data hits the network Device N1, the data is forwarded based on the snooping learned route.

For MC-Link clients, data is forwarded through N1. In the case of failover of the MC-Links, the data reaches the client through N2. For S-Link clients on N1, data would be forwarded through normal snooping routes.

For S-Link clients on N2, data is forwarded through the ICL interface. Layer 2 multicast routes on N1 do not show these groups unless there is interest for the same group over MC-Links or over S-Links on N1. For the IRB scenario, the IGMP membership and Layer 3 multicast route on N1 does however show these groups learned over the IRB interface.

Therefore, for a case where a specific group interest is only on the S-Link on N2, data arriving on N1 reaches N2 through the default route, and the Layer 2 multicast route on N2 has the S-Link in the outgoing interface list.

Figure 15: Multicast Topology with Source Connected Through MC-Link

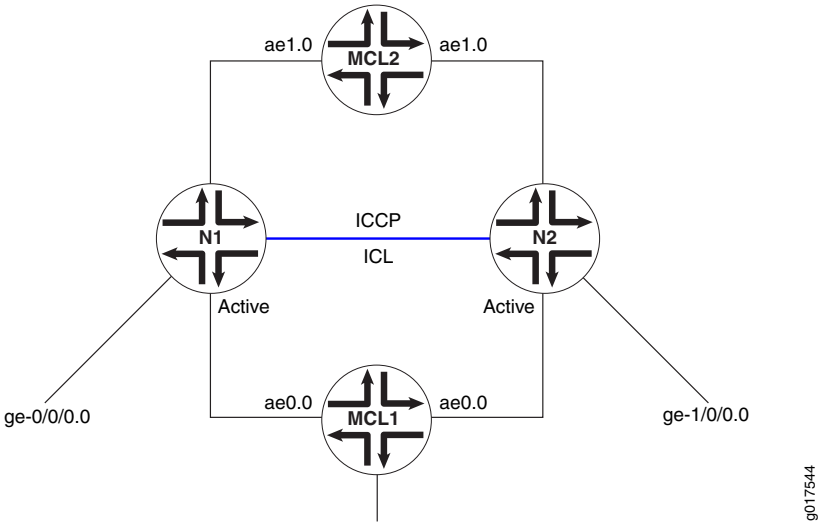

In Figure 15 on page 136, MCL1 and MCL2 are on different devices, and the multicast source or IGMP querier is connected through MCL2. The data forwarding behavior seen is similar to that explained for multicast topology with source connected through Layer 3.

NOTE: IGMP snooping should not be configured in proxy mode. There should be no IGMP hosts or IGMP or PIM routers sitting on the ICL interface.

Up and Down Event Handling

The following conditions apply to up and down event handling:

- If the Inter-Chassis Control Protocol (ICCP) connection is UP but the ICL interface goes DOWN, the router configured as the backup brings down all the multichassis aggregated Ethernet interfaces shared with the peer that is connected to ICL. This ensures that there are no loops in the network. Otherwise, both PEs become PIM-designated routers and, hence, forward multiple copies of the same packet to the customer edge.
- If the ICCP connection is UP and the ICL comes UP, the router configured as the backup brings up the multichassis aggregated Ethernet interfaces shared with the peer.
- If both the ICCP connection and the ICL are DOWN, the router configured as the backup brings up the multichassis aggregated Ethernet interfaces shared with the peer.
- The Layer 2 address learning process (l2ald) does not store the information about a MAC address learned from a peer in the kernel. If l2ald restarts, and if the MAC address was not learned from the local multichassis aggregated Ethernet interface, l2ald clears the MAC addresses, which causes the router to flood the packets destined to this MAC address. This behavior is similar to that in a Routing Engine switchover. (Note that currently l2ald runs on a Routing Engine only when it is a master). Also, during the time

l2ald is DOWN, ARP packets received from an ICCP peer are dropped. ARP retry takes care of this situation. This is the case with Routing Engine switchover, too.

- If ICCP restarts, l2ald does not identify that a MAC address was learned from a peer and, if the MAC address was learned only from the peer, that MAC address is deleted, and the packets destined to this MAC address are flooded.

Inter-Chassis Control Protocol

Inter-Chassis Control Protocol (ICCP) is used to synchronize configurations, states, and data.

ICCP supports the following types of state information:

- MC-LAG members and their operational states
- Single-homed members and their operational states

ICCP supports the following application database synchronization parameters:

- MAC addresses learned and to be aged
- ARP information learned over IRB

Inter-Chassis Control Protocol Message

ICCP messages and attribute-value pairs (AVPs) are used for synchronizing MAC address and ARP information.

Understanding the Incremented Values of Statistical Counters for Loop-Free MC-LAG Networks

In an MC-LAG in an active-active bridging domain, the output of the following command displays the MC-LAG color counters to be continuously increasing. This increase in the statistical count is an expected behavior because the MC-LAG color attribute or counter functions as a loop prevention mechanism.

```
request pfe execute target fpc0 command "show jnh 0 exceptions" |grep color
GOT: mc lag color                       DISC(88)   554712463 144488623417
request pfe execute target fpc0 command "show jnh 0 exceptions" |grep color
GOT: mc lag color                       DISC(88)   554712747 144488664296
```

The exception table stored in the Packet Forwarding Engine contains a list of counters as displayed in the following example output:

```
request pfe execute target fpc0 command "show jnh 0 exceptions"
SENT: Ukern command: show jnh 0 exceptions
GOT: Reason                             Type         Packets      Bytes
GOT: ============================================================================
GOT: Ucode Internal
GOT: ---------------------
GOT: mcast stack overflow               DISC(33)           0          0
GOT: sample stack error                 DISC(35)           0          0
GOT: undefined nexthop opcode           DISC(36)           0          0
GOT: internal ucode error               DISC(37)           0          0
```

```
GOT: invalid fabric hdr version         DISC(41)           0          0
GOT:
GOT: PFE State Invalid
GOT: ----------------------
GOT: sw error                           DISC(64)   803092438 59795128826
GOT: child ifl nonlocal to pfe          DISC(85)           0          0
GOT: invalid fabric token               DISC(75)         179      42346
GOT: unknown family                     DISC(73)           0          0
GOT: unknown vrf                        DISC(77)           0          0
GOT: iif down                           DISC(87)           0          0
GOT: unknown iif                        DISC( 1)
GOT: invalid stream                     DISC(72)           0          0
GOT: egress pfe unspecified             DISC(19)       10889    1536998
GOT: invalid L2 token                   DISC(86)          26       1224
GOT: mc lag color                       DISC(88)   554693648
144486028726<<<<<<<<<<<<<<<<<<<<<<<<<
GOT: dest interface non-local to pfe    DISC(27)  10939253797 2078273071209
GOT: invalid inline-svcs state          DISC(90)           0          0
GOT: nh id out of range                 DISC(93)           0          0
GOT: invalid encap                      DISC(96)           0          0
GOT: replication attempt on empty irb   DISC(97)           0          0
GOT: invalid selector entry             DISC(98)           0          0
GOT:
GOT:
GOT: Packet Exceptions
GOT: ----------------------
GOT: bad ipv4 hdr checksum              DISC( 2)
GOT: non-IPv4 layer3 tunnel             DISC( 4)           0          0
GOT: GRE unsupported flags              DISC( 5)           0          0
GOT: tunnel pkt too short               DISC( 6)           0          0
GOT: tunnel hdr too long                DISC( 7)           0          0
GOT: bad IPv6 options pkt               DISC( 9)           0          0
GOT: bad IP hdr                         DISC(11)           0          0
GOT: bad IP pkt len                     DISC(12)           0          0
GOT: L4 len too short                   DISC(13)
GOT: invalid TCP fragment               DISC(14)           0          0
GOT: mtu exceeded                       DISC(21)           0          0
GOT: frag needed but DF set             DISC(22)           0          0
GOT: ttl expired                        PUNT( 1)           9        769
GOT: IP options                         PUNT( 2)          16        512
GOT: xlated l2pt                        PUNT(14)           0          0
GOT: control pkt punt via ucode         PUNT( 4)           0          0
GOT: frame format error                 DISC( 0)
GOT: tunnel hdr needs reassembly        PUNT( 8)           0          0
GOT: GRE key mismatch                   DISC(76)           0          0
GOT: my-mac check failed                DISC(28)
GOT: frame relay type unsupported       DISC(38)           0          0
GOT: IGMP snooping control packet       PUNT(12)           0          0
GOT: bad CLNP hdr                       DISC(43)           0          0
GOT: bad CLNP hdr checksum              DISC(44)           0          0
GOT: Tunnel keepalives                  PUNT(58)           0          0
GOT:
GOT:
GOT: Bridging
GOT: ----------------------
GOT: lt unknown ucast                   DISC(84)           0          0
GOT: dmac miss                          DISC(15)           0          0
GOT: mac learn limit exceeded           DISC(17)           0          0
GOT: static mac on unexpected iif       DISC(18)           0          0
GOT: no local switching                 DISC(20)           0          0
```

```
GOT: bridge ucast split horizon         DISC(26)       39458   13232394
GOT: mcast smac on bridged iif          DISC(24)        1263     200152
GOT: bridge pkt punt                    PUNT( 7)           0          0
GOT: iif STP blocked                    DISC( 3)
GOT: oif STP blocked                    DISC(31)
GOT: vlan id out of oif's range         DISC(32)
GOT: mlp pkt                            PUNT(11)    15188054  440453569
GOT: input trunk vlan lookup failed     DISC(91)           0          0
GOT: output trunk vlan lookup failed    DISC(92)           0          0
GOT: LSI/VT vlan validation failed      DISC(94)           0          0
GOT:
GOT:
GOT: Firewall
GOT: ----------------------
GOT: mac firewall                       DISC(78)
GOT: firewall discard                   DISC(67)           0          0
GOT: tcam miss                          DISC(16)           0          0
GOT: firewall reject                    PUNT(36)      155559   59137563
GOT: firewall send to host              PUNT(54)           0          0
GOT: firewall send to host for NAT      PUNT(59)           0          0
GOT:
GOT:
GOT: Routing
GOT: ----------------------
GOT: discard route                      DISC(66)     1577352   82845749
GOT: dsc ifl discard route              DISC(95)           0          0
GOT: hold route                         DISC(70)       21130    1073961
GOT: mcast rpf mismatch                 DISC( 8)           0          0
GOT: resolve route                      PUNT(33)        2858     154202
GOT: control pkt punt via nh            PUNT(34)    51807272 5283911584
GOT: host route                         PUNT(32)    23473304 1370843994
GOT: ICMP redirect                      PUNT( 3)           0          0
GOT: mcast host copy                    PUNT( 6)           0          0
GOT: reject route                       PUNT(40)        1663     289278
GOT: link-layer-bcast-inet-check        DISC(99)           0          0
GOT:
```

Consider a sample deployment in which two provider edge (PE) routers, PE1 and PE2, are connected with an aggregated Ethernet interface, **ae0**. respectively. Multichassis link aggregation groups (MC-LAGs) are used between PE1 and PE2 to form a logical LAG interface between the two controllers. PE1 and PE2 in an MC-LAG use an interchassis control link-protection link (ICL-PL) to replicate forwarding information across the peers.

Inter-Chassis Control Protocol (ICCP) messages are sent between the two PE devices. In this example, you configure an MC-LAG across two routers, consisting of two aggregated Ethernet interfaces, an interchassis control link-protection link (ICL-PL), multichassis protection link for the ICL-PL, and ICCP for the peers hosting the MC-LAG.

The PE1 router is connected using another aggregated Ethernet interface, **ae3**, to a host, H1, and to another MC-LAG host called C1. MC-LAG is enabled on the **ae3** interface.

Traffic received on PE1 from MC-LAG C1 can be flooded over the ICL to reach PE2. When the packets arrive at PE2, they can be flooded back to MC- LAG C1. Traffic sent by the single-homed host H1 can be flooded to MC-LAG C1 and the ICL on PE1. When PE2 receives such traffic from ICL, it can be again flooded to MC-LAG C1. To protect the MC-LAG topology from such loops, the MC-LAG color capability is implemented. This functionality

is applied on the ingress of the ICL link. Therefore, when PE2 receives a packet from PE1, it sets the MC-LAG color as active or turns it on. When PE2 requires to flood the packet towards the MC-LAG link, it verifies whether the MC-LAG color bit is set or tagged as on. If the color is set, it drops the packet on the egress interface of MC-LAG **ae3** member link interfaces and the **mc-lag color** counter in the jnh exceptions is incremented.

Such a behavior of increase in counter value is an expected condition in an MC-LAG configured in an active/active bridging domain and when any form of traffic that needs to be flooded, such as ARP broadcast or multicast traffic, traverses the network.

Every VLAN might drop some packets to prevent loops and such a drop of packets might not be specific to a VLAN.

Sometimes, on both MC LAGs of the MX Series routers, you might notice that the counter increases on FPC0 and FPC2, but it does not increase on FPC2 as illustrated in the following sample output:

```
request pfe execute target fpc0 command "show jnh 0 exceptions" |grep color
GOT: mc lag color                       DISC(88)   558477875 144977739683
request pfe execute target fpc1 command "show jnh 0 exceptions" |grep color
GOT: mc lag color                       DISC(88)           0          0
request pfe execute target fpc2 command "show jnh 0 exceptions" |grep color
GOT: mc lag color                       DISC(88)   518499257 119130527834
```

This behavior occurs because on an MX Series router with a 16-port 10-Gigabit Ethernet MPC (16x10GE 3D MPC), there are four Packet Forwarding Engines for each MPC. If you examine one Packet Forwarding Engine in FPC 0, 1, and 2, PFE1 in FPC1 does not have any interfaces which are member of MC-LAG. It might contain interfaces in other aggregated Ethernet interfaces that are are not part of MC-LAG. Therefore, to obtain the correct counter statistics, you must examine the other Packet Forwarding Engines by entering the **request pfe execute target fpc0 command "show jnh X exceptions" |grep color** command where X can be 0, 1, 2, or 3.

When the counter named **dest interface non-local to pfe** is increasing, it is a desired behavior when aggregated Ethernet interfaces are split over more than one FPC. Consider an example in which an **ae5** interface contains the following member links: **xe-0/1/0** on (FPC0) and **xe-1/1/0** (FPC1) Based on the hash algorithm, traffic must be split between these two links. The hash algorithm is applied on the ingress FPC and performs an operation where it marks each packet through which FPC must be forwarded (FPC0 or FPC1). Then the packet is sent to the fabric. From the fabric, all of traffic is sent to both FPCs 0 and 1. On FPC0, the microkernel analyzes the packet and determines whether the packet needs to be forwarded by the local interface (local to pfe) or whether this packet has already been forwarded through FPC1 (non-local to pfe). If the packet has been already forwarded, the packet is dropped and the **non-local to pfe** counter is incremented.

Configuring Active-Active Bridging and VRRP over IRB in Multichassis Link Aggregation on MX Series Routers and QFX Series Switches

The following sections describe the configuration of active-active bridging and VRRP over IRB in a multichassis link aggregation (MC-LAG) :

Configuring MC-LAG

An MC-LAG is composed of logical link aggregation groups (LAGs) and is configured under the **[edit interfaces ae*X*]** hierarchy, as follows:

```
[edit]
interfaces {
   ae0 {
      encapsulation ethernet-bridge;
      multi-chassis-protection {
         peer 10.10.10.10 {
            interface ge-0/0/0;
            }
         }
      aggregated-ether-options {
         mc-ae {
            mode active-active; # see note below
         }
      }
   }
}
```

NOTE: The **mode active-active** statement is valid only if encapsulation is an **ethernet-bridge** or **extended-vlan-bridge**.

Use the **mode** statement to specify if an MC-LAG is **active-standby** or **active-active**. If the ICCP connection is UP and ICL comes UP, the router configured as standby brings up the multichassis aggregated Ethernet interfaces shared with the peer.

Using **multi-chassis-protection** at the physical interface level is a way to reduce the configuration at the logical interface level.

If there are n+1 logical interfaces under ae0, from ae0.0 through ae0.n, there are n+1 logical interfaces under ge-0/0/0 as well, from ge-0/0/0.0 through ge-0/0/0.n, each ge-0/0/0 logical interface is a protection link for the ae0 logical interface.

NOTE: A bridge domain cannot have multichassis aggregated Ethernet logical interfaces that belong to different redundancy groups.

Configuring the Interchassis Link-Protection Link

The interchassis link-protection link (ICL-PL) provides redundancy when a link failure (for example, an MC-LAG trunk failure) occurs on one of the active links. The ICL-PL is an aggregated Ethernet interface. You can configure only one ICL-PL between the two peers, although you can configure multiple MC-LAGs between them.

The ICL-PL assumes that interface ge-0/0/0.0 is used to protect interface ae0.0 of MC-LAG-1:

```
[edit]
interfaces {
   ae0 {
      ....
      unit 0 {
         multi-chassis-protection {
            peer 10.10.10.10 {
               interface ge-0/0/0.0;
            }
            ....
         }
         ...
      }
   }
}
```

The protection interface can be an Ethernet type interface such as ge or xe, or an aggregated Ethernet (ae) interface.

Configuring Multiple Chassis

A top-level hierarchy is used to specify a multichassis-related configuration, as follows:

```
[edit]
multi-chassis {
   multi-chassis-protection {
      peer 10.10.10.10 {
         interface ge-0/0/0;
      }
   }
}
```

This example specifies interface ge-0/0/0 as the multichassis protection interface for all the multichassis aggregated Ethernet interfaces which are also part of the peer. This can be overridden by specifying protection at the physical interface level and the logical interface level.

Configuring the Service ID

You must configure the same unique network-wide configuration for a service in the set of PE routers providing the service. You can configure the service IDs under the level of the hierarchies shown in the following examples:

Global Configuration (Default Logical System)

```
switch-options {
   service-id 10;
}
bridge-domains {
   bd0 {
      service-id 2;
   }
}
routing-instances {
   r1 {
      switch-options {
         service-id 10;
      }
      bridge-domains {
         bd0 {
            service-id 2;
         }
      }
   }
}
```

Logical Systems

```
ls1 {
   switch-options {
      service-id 10;
   }
   routing-instances {
      r1 {
         switch-options {
            service-id 10;
         }
      }
   }
}
```

NOTE: Using a service name per bridge domain is not supported.

The bridge-level service ID is required to link related bridge domains across peers, and should be configured with the same value. The **service-id** values share the name space across all bridging and routing instances, and across peers. Thus, duplicate values for service IDs are not permitted across these entities.

The service ID at the bridge domain level is mandatory for type non-single VLAN bridge domains. The service ID is optional for bridge domains with a VLAN identifier (VID) defined.

If no service ID is defined in the latter case, it is picked up from the service ID configuration for that routing instance.

NOTE: When this default routing instance (or any other routing instance) which contains a bridge domain containing a multichassis aggregated Ethernet interface is configured, you must configure a global-level **switch-options service-id** ***number***, irrespective of whether the contained bridge domains have specific service IDs configured.

In the sample illustration shown in Figure 16 on page 144, network routing instances N1 and N2, both for the same service ID, are configured with same service ID in both N1 and N2. Use of a name string instead of a number is not supported.

Figure 16: N1 and N2 for the Same Service with Same Service ID

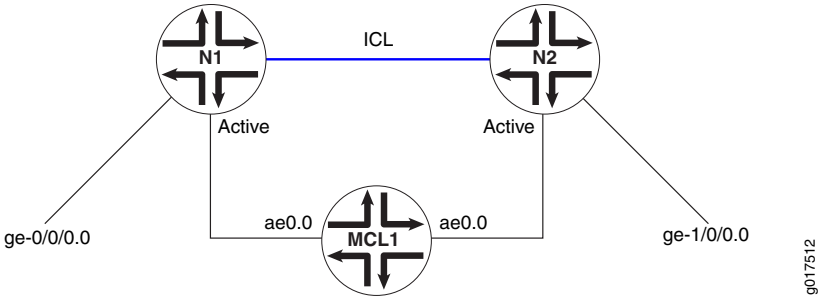

The following configuration restrictions apply:

- The service ID must be configured when the multichassis aggregated Ethernet interface is configured and a multichassis aggregated Ethernet logical interface is part of a bridge domain. This requirement is enforced.
- A single bridge domain cannot correspond to two redundancy group IDs.

  In Figure 17 on page 145, it is possible to configure a bridge domain consisting of logical interfaces from two multichassis aggregated Ethernet interfaces and map them to a separate redundancy group ID, which is not supported. A service must be mapped one-to-one with the redundancy group providing the service. This requirement is enforced.

Figure 17: Bridge Domain with Logical Interfaces from Two Multichassis Aggregated Ethernet Interfaces

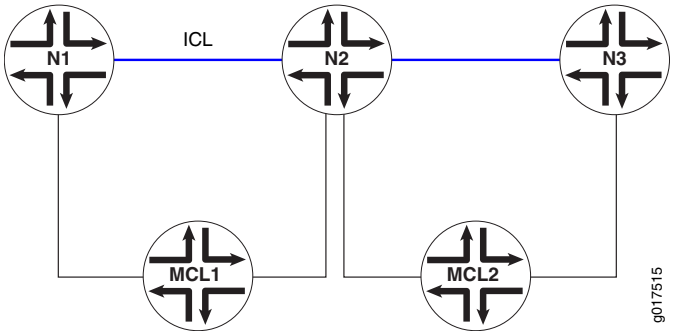

To display the multichassis aggregated Ethernet configuration, use the **show interfaces *mc-ae*** command. For more information, see the CLI Explorer.

Configuring IGMP Snooping for Active-Active MC-LAG

For the multicast solution to work, the following must be configured:

- The multichassis protection link must be configured as a router-facing interface.

```
[edit bridge-domain bd-name]
protocols {
  igmp-snooping {
    interface ge-0/0/0.0 {
      multicast-router-interface;
    }
  }
}
```

  In this example, ge-0/0/0.0 is an ICL interface.

- The **multichassis-lag-replicate-state** statement options must be configured under the **multicast-snooping-options** statement for that bridge domain.

NOTE: Snooping with active-active MC-LAG is only supported in non-proxy mode.

Configuring IGMP Snooping in MC-LAG Active-Active Mode

You can use the **bridge-domain** statement's **service-id** option to specify the multichassis aggregated Ethernet configuration on MX240 routers, MX480 routers, MX960 routers and QFX Series switches.

- The **service-id** statement is mandatory for non-single VLAN type bridge domains (**none**, **all**, or **vlan-id-tags:dual**).
- The statement is optional for bridge domains with a VID defined.

- If no **service-id** is defined in the latter case, it is picked up from the round-trip time's (RTT's) **service-id** configuration.
- The bridge-level **service-id** is required to link related bridge domains across peers, and should be configured with the same value.
- The **service-id** values share the name space across all bridging and routing instances, and across peers. Thus, duplicate **service-id** values are not permitted across these entities.
- A change of bridge **service-id** is considered catastrophic, and the bridge domain is changed.

This procedure allows you to enable or disable the replication feature.

To configure IGMP snooping in MC-LAG active-active mode :

1. Use the **multichassis-lag-replicate-state** statement at the **[edit multicast-snooping-options]** hierarchy level in the master instance.

```
multicast-snooping-options {
  ...
  multichassis-lag-replicate-state; # REQUIRED
}
```

2. Use the **interface *icl-intf-name*** statement at the **[edit protocols igmp-snooping]** hierarchy level, as shown in the following example:

```
protocols {
  igmp-snooping {
    interface icl-intf-name {
      multicast-router-interface;
    }
  }
}
```

NOTE: For QFX use the following configuration:

```
protocols {
  igmp-snooping {
    vlan vlan_name{
      }
      interface icl-intf-name {
      multicast-router-interface;
    }
  }
}
```

The interchassis link, **interface *icl-intf-name***, of the learning domain should be a router-facing interface.

Example: Configuring DHCP Relay on MC- LAG with VRRP on an EX9200 Switch

This example shows how to configure Dynamic Host Configuration Protocol (DHCP) relay on EX9200 switches with the multichassis link aggregation (MC-LAG) feature using Virtual Router Redundancy Protocol (VRRP).

Requirements

This example uses the following hardware and software components:

- Junos OS Release 12.3 or later for EX Series
- Two EX9200 switches

Before you configure DHCP relay, be sure that you understand how to:

- Configure MC-LAG and verify that MC-LAG and ICCP is up and running

To complete the configuration, enable VRRP by completing the following steps for each MC-LAG:

- Create an integrated routing and bridging (IRB) interface.
- Create a VRRP group and assign a virtual IP address that is shared between each switch in the VRRP group.
- Enable a member of a VRRP group to accept all packets destined for the virtual IP address if it is the master in the VRRP group.
- Configure Layer 3 connectivity between the VRRP groups.

Overview

In this example, you configure DHCP relay with MC-LAG across two switches consisting of two EX9200 switches, an interchassis link-protection link (ICL-PL), multichassis protection link for the ICL-PL, ICCP for the peers hosting the MC-LAG, and Layer 3 connectivity between MC-LAG peers. Layer 3 connectivity is required for ICCP.

NOTE: On EX9200 switches, dynamic ARP resolution is not supported over inter-chassis control links (ICLs). As a workaround, you can configure static ARP on both ends of the ICL.

Topology

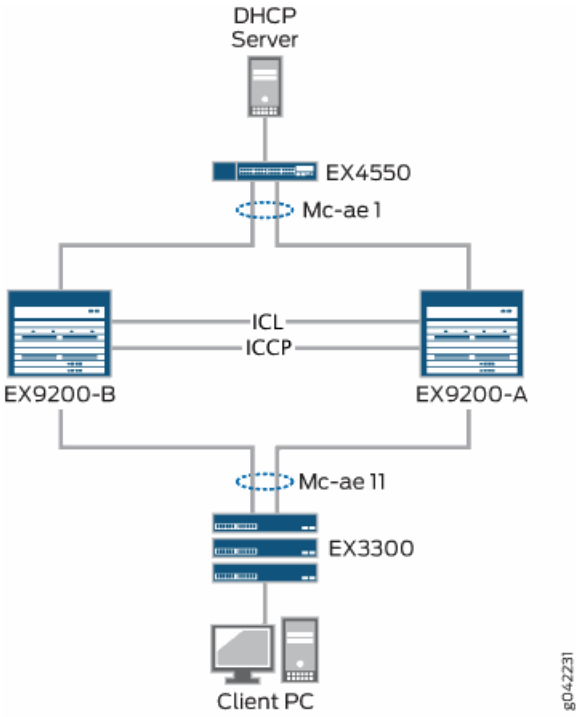

Table 7: Components of the Topology for Configuring DHCP Relay

Hostname	Hardware
Switch EX9200-A	EX9200 switch
Switch EX9200-B	EX9200 switch

Configuration

CLI Quick Configuration

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

Switch A and Switch B

```
set forwarding-options dhcp-relay forward-snooped-clients all-interfaces
set forwarding-options dhcp-relay server-group GVP-DHCP 10.168.61.5
set forwarding-options dhcp-relay overrides allow-snooped-clients
set forwarding-options dhcp-relay active-server-group GVP-DHCP
set forwarding-options dhcp-relay group Floor1 interface irb.2540
set forwarding-options dhcp-relay route-suppression destination
```

Step-by-Step Procedure

To configure DHCP relay on both Switch A and Switch B:

1. Configure forward snooped (unicast) packets on the interfaces.

```
[edit forwarding-options dhcp-relay]
  user@switch# set forward-snooped-clients all-interfaces
```

2. Create a DHCP server group. A DHCP server group can include 1 through 5 DHCP server IP addresses.

```
[edit forwarding-options dhcp-relay]
  user@switch# set server-group GVP-DHCP 10.168.61.5
```

3. Allow the creation of a binding entry using snooped (unicast) clients.

```
[edit forwarding-options dhcp-relay]
  user@switch# set overrides allow-snooped-clients
```

4. Apply a DHCP relay agent configuration to the named group of DHCP server addresses.

```
[edit forwarding-options dhcp-relay ]
  user@switch# set active-server-group GVP-DHCP
```

5. Create a DHCP relay group that includes at least one interface.

   DHCP relay runs on the interfaces defined in DHCP groups.

```
[edit forwarding-options dhcp-relay]
  user@switch# set group Floor1 interface irb.2540
```

6. Configure the relay agent to suppress the installation of ARP and route entries for corresponding client binding.

```
[edit forwarding-options dhcp-relay]
  user@switch# set route-suppression destination
```

Results

From configuration mode, confirm your configuration by entering the **show forwarding-options dhcp-relay** command on both Switch A and Switch B. If the output does not display the required configuration, repeat the instructions in this example to correct the configuration.

```
user@switch# show forwarding-options dhcp-relay
forward-snooped-clients {
  all-interfaces;
}
overrides {
  allow-snooped-clients;
```

```
}
server-group {
   GVP-DHCP {
      10.168.61.5;
   }
}
active-server-group {
   GVP-DHCP;
}
group Floor1 {
   interface {
      irb.2540;
   }
}
route-suppression {
   destination;
}
```

Overwriting Address Information

Step-by-Step Procedure

We recommend that you configure the DHCP relay agent to change the gateway IP address (giaddr) field in packets that it forwards between a DHCP client and a DHCP server.

To overwrite the address of every DHCP packet with the address of the DHCP relay agent before forwarding the packet to the DHCP server.

1. Specify that you want to configure override options.

```
[edit forwarding-options dhcp-relay]
   user@hots# set overrides
```

2. Specify that the address of DHCP packets is overwritten.

```
[edit forwarding-options dhcp-relay overrides]
   user@host# set always-write-giaddr
```

Verification

Confirm that the configuration is working properly.

- Verifying That DHCP Relay Binding Is Occurring on page 150
- Verifying That Relay Statistics Are Being Displayed on page 151

Verifying That DHCP Relay Binding Is Occurring

Purpose

Verify that address bindings in the DHCP client table are being displayed.

Action

```
root@switchA# show dhcp relay binding detail
Client IP Address:  10.168.103.20
Hardware Address:             84:18:88:a8:ca:80
State:            BOUND(RELAY_STATE_BOUND)
Lease Expires:                2013-10-03 12:17:43 CEST
Lease Expires in:             85829 seconds
Lease Start:                  2013-10-02 10:48:34 CEST
Last Packet Received:         2013-10-02 12:17:43 CEST
Incoming Client Interface:    ae0.0(irb.2540)
Server Ip Address:            10.168.61.5
Server Interface:             none
Bootp Relay Address:          10.168.103.2
Session Id:                   29
```

```
root@switchB# show dhcp relay binding detail
Client IP Address:  10.168.103.20
Hardware Address:             84:18:88:a8:ca:80
State:            BOUND(RELAY_STATE_BOUND)
Lease Expires:                2013-10-03 12:17:43 CEST
Lease Expires in:             86228 seconds
Lease Start:                  2013-10-02 10:48:34 CEST
Last Packet Received:         2013-10-02 10:48:34 CEST
Incoming Client Interface:    ae11.0(irb.2540)
Server Ip Address:            10.168.61.5
Server Interface:             none
Bootp Relay Address:          10.168.103.2
Session Id:                   16
```

Meaning The field State indicates the state of the DHCP relay address binding table on the DHCP client. The state BOUND indicates that the client has an active IP address lease.

Verifying That Relay Statistics Are Being Displayed

Purpose Verify that DHCP relay statistics are being displayed.

Action

```
root@switchA# show dhcp relay statistics
```

```
Packets dropped:
 Total                    9
 dhcp-service total       9
Messages received:
    BOOTREQUEST              4
    DHCPDECLINE              0
    DHCPDISCOVER             1
    DHCPINFORM               0
    DHCPRELEASE              0
    DHCPREQUEST              3

Messages sent:
    BOOTREPLY                0
    DHCPOFFER                0
    DHCPACK                  0
    DHCPNAK                  0
    DHCPFORCERENEW           0
```

Meaning The field Total displays the total number of packets discarded by the extended DHCP relay agent application.

Configuring Manual and Automatic Link Switchover for MC-LAG Interfaces on MX Series Routers

In a multichassis link aggregation (MC-LAG) topology with active-standby mode, a link switchover happens only if the active node goes down. You can override this default behavior by configuring an MC-LAG interface in active-standby mode to automatically revert to a preferred node. With this configuration, you can trigger a link switchover to a preferred node even when the active node is available. For example, consider two nodes, PE1 and PE2. PE1 is configured in active mode making it a preferred node, and PE2 is configured in active-standby mode. In case of any failure at PE1, PE2 becomes the active node. However, as soon as PE1 is available again, an automatic link switchover is triggered and the control is switched back to PE1 even though PE2 is active.

You can configure the link switchover in two modes: revertive and nonrevertive. In revertive mode, the link switchover is triggered automatically by using the **request interface mc-ae switchover** operational mode command. In nonrevertive mode, the link switchover must be triggered manually by the user. You can also configure a revert time that triggers an automatic or manual switchover when the specified timer expires.

NOTE:

- If two MC-LAG devices configured in an active-standby setup using Inter-Chassis Control Protocol (ICCP) and nonrevertive switchcover mode is configured on the aggregated Ethernet interfaces of both the MC-LAGs and when both mc-ae interfaces are linked together with a Layer 2 circuit local-switching configuration, we recommend that you perform switchover by entering the **request interface mc-ae switchover (immediate mcae-id mcae-id | mcae-id mcae-id)** operational mode command on only one of the aggregated Ethernet interfaces of an MC-LAG device. This command can be issued only on MC-LAG devices that are configured as active nodes (by using the **status-control active** statement at the **[edit interfaces ae*X* aggregated-ether-options mc-ae]** hierarchy level).
- In nonrevertive switchover mode, when an MC-LAG interface transitions to the standby state because of an MC-LAG member link failure and another MC-LAG interface moves to the active state, the MC-LAG in standby state remains in that state until the MC-LAG in active state encounters a failure and returns to the active state.
- If you perform a switchover on both the aggregated Ethernet interfaces in the MC-LAG, because of Layer 2 circuit local-switching configuration, a switchover on one aggregated Ethernet interface triggers a switchover on the other aggregated Ethernet interface. In such a scenario, both the aggregated Ethernet interfaces move to the standby state and then transition back to the active state. Therefore, you must not perform switchover on both the aggregated Ethernet interfaces in an MC-LAG at the same time.
- Layer 2 circuit configuration and VPLS functionalities are not supported if you configure an MC-LAG interface to be in revertive switchover mode. You can configure the revertive or nonrevertive switchover capability only if two MC-LAG devices are configured in an active-standby setup (one device set as an active node by using the **status-control standby** statement and the other device set as a standby node by using the **status-control active** statement at the **[edit interfaces ae*X* aggregated-ether-options mc-ae]** hierarchy level. You can perform a switchover by entering the **request interface mc-ae switchover (immediate mcae-id mcae-id | mcae-id mcae-id)** operational mode command only on MC-LAG devices configured as active nodes.

To configure the link switchover mechanism on an MC-LAG interface:

1. Configure the link switchover in revertive mode.

```
[edit interfaces aeX aggregated-ether-options mc-ae]
user@host# set switchover-mode revertive
```

2. (Optional) Configure the link switchover in nonrevertive mode.

```
[edit interfaces aeX aggregated-ether-options mc-ae]
user@host# set switchover-mode non-revertive
```

3. Configure the revert time.

```
[edit interfaces aeX aggregated-ether-options mc-ae]
user@host# set revert-time revert-time
```

4. Trigger manual switchover.

```
[edit request interface mc-ae]
user@host# set switchover < immediate> mcae-id mcae-id
```

You can use the **show interfaces mc-ae revertive-info** command to view the switchover configuration information.

Example: Configuring Multichassis Link Aggregation in Active-Active Mode

This example shows how to configure a multichassis link aggregation group (MC-LAG) in an active-active scenario, which load balances traffic across the PEs.

Requirements

This example uses the following hardware and software components:

NOTE: This example also applies to QFX10002 and QFX10008 switches.

- Four Juniper Networks MX Series routers (MX240, MX480, MX960)
- Junos OS Release 11.2 or later running on all four routers

Overview

Consider a sample topology in which a customer edge router, CE, is connected to two provider edge (PE) routers, PE1 and PE2, respectively. The two PE devices each have a link aggregation group (LAG) connected to the CE device. The configured mode is active -active, meaning that both PE routers' LAG ports are active and carrying traffic at the same time. PE1 and PE2 are connected to a single service provider router, P.

In this example, the CE router is not aware that its aggregated Ethernet links are connected to two separate PE devices. The two PE devices each have a LAG connected to the CE device. The configured mode is active-active, meaning that both PE routers' LAG ports are active and carrying traffic at the same time.

In Figure 18 on page 155, from the perspective of Router CE, all four ports belonging to a LAG are connected to a single service provider device. Because the configured mode is active-active, all four ports are active, and the CE device load-balances the traffic to the peering PE devices. On the PE routers, a regular LAG is configured facing the CE device.

On one end of an MC-LAG is an MC-LAG client device, such as a server, that has one or more physical links in a LAG. This client device does not need to detect the MC-LAG. On the other side of an MC-LAG are two MC-LAG routers. Each of the routers has one or more physical links connected to a single client device. The routers coordinate with each other to ensure that data traffic is forwarded properly.

ICCP messages are sent between the two PE devices. In this example, you configure an MC-LAG across two routers, consisting of two aggregated Ethernet interfaces, an interchassis link-protection link (ICL-PL), multichassis protection link for the ICL-PL, and ICCP for the peers hosting the MC-LAG.

Topology Diagram

Figure 18 on page 155 shows the topology used in this example.

Figure 18: MC-LAG Active-Active Mode on MX Series Routers

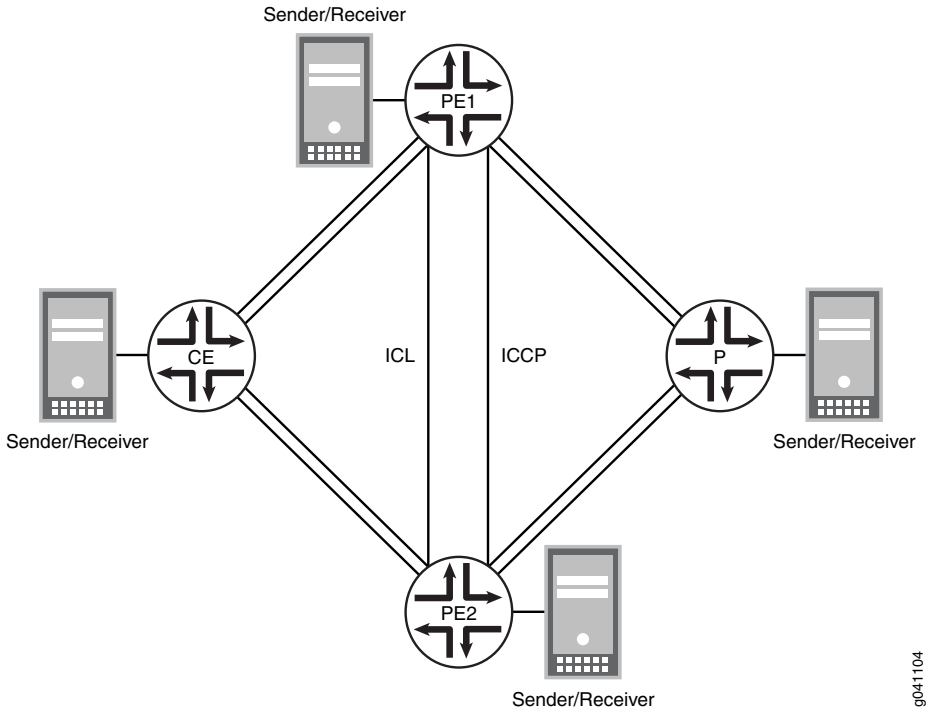

Configuring the PE Routers

CLI Quick Configuration

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

Router PE1

```
set chassis aggregated-devices ethernet device-count 5
set interfaces ge-1/0/1 gigether-options 802.3ad ae1
set interfaces ge-1/0/2 unit 0 family inet address 10.100.100.1/30
set interfaces ge-1/0/6 gigether-options 802.3ad ae0
set interfaces ge-1/1/1 flexible-vlan-tagging
set interfaces ge-1/1/1 encapsulation flexible-ethernet-services
set interfaces ge-1/1/1 unit 0 encapsulation vlan-bridge
set interfaces ge-1/1/1 unit 0 vlan-id-range 100-110
set interfaces ge-1/1/4 flexible-vlan-tagging
set interfaces ge-1/1/4 encapsulation flexible-ethernet-services
set interfaces ge-1/1/4 unit 0 encapsulation vlan-bridge
set interfaces ge-1/1/4 unit 0 vlan-id-range 100-110
set interfaces ae0 flexible-vlan-tagging
set interfaces ae0 encapsulation flexible-ethernet-services
set interfaces ae0 aggregated-ether-options lacp active
set interfaces ae0 aggregated-ether-options lacp system-priority 100
set interfaces ae0 aggregated-ether-options lacp system-id 00:00:00:00:00:05
set interfaces ae0 aggregated-ether-options lacp admin-key 1
set interfaces ae0 aggregated-ether-options mc-ae mc-ae-id 5
set interfaces ae0 aggregated-ether-options mc-ae redundancy-group 10
set interfaces ae0 aggregated-ether-options mc-ae chassis-id 1
set interfaces ae0 aggregated-ether-options mc-ae mode active-active
set interfaces ae0 aggregated-ether-options mc-ae status-control active
set interfaces ae0 unit 0 encapsulation vlan-bridge
set interfaces ae0 unit 0 vlan-id-range 100-110
set interfaces ae0 unit 0 multi-chassis-protection 10.100.100.2 interface ge-1/1/4.0
set interfaces ae1 flexible-vlan-tagging
set interfaces ae1 encapsulation flexible-ethernet-services
set interfaces ae1 aggregated-ether-options lacp active
set interfaces ae1 aggregated-ether-options lacp system-priority 100
set interfaces ae1 aggregated-ether-options lacp system-id 00:00:00:00:00:05
set interfaces ae1 aggregated-ether-options lacp admin-key 1
set interfaces ae1 aggregated-ether-options mc-ae mc-ae-id 10
set interfaces ae1 aggregated-ether-options mc-ae redundancy-group 10
set interfaces ae1 aggregated-ether-options mc-ae chassis-id 1
set interfaces ae1 aggregated-ether-options mc-ae mode active-active
set interfaces ae1 aggregated-ether-options mc-ae status-control active
set interfaces ae1 unit 0 encapsulation vlan-bridge
set interfaces ae1 unit 0 vlan-id-range 100-110
set interfaces ae1 unit 0 multi-chassis-protection 10.100.100.2 interface ge-1/1/4.0
set bridge-domains bd0 domain-type bridge
set bridge-domains bd0 vlan-id all
set bridge-domains bd0 service-id 20
set bridge-domains bd0 interface ae1.0
set bridge-domains bd0 interface ge-1/0/3.0
set bridge-domains bd0 interface ge-1/1/1.0
```

```
set bridge-domains bd0 interface ge-1/1/4.0
set bridge-domains bd0 interface ae0.0
set protocols iccp local-ip-addr 10.100.100.1
set protocols iccp peer 10.100.100.2 redundancy-group-id-list 10
set protocols iccp peer 10.100.100.2 liveness-detection minimum-interval 1000
set switch-options service-id 10
```

Router PE2

```
set chassis aggregated-devices ethernet device-count 5
set interfaces ge-1/0/2 unit 0 family inet address 10.100.100.2/30
set interfaces ge-1/0/3 flexible-vlan-tagging
set interfaces ge-1/0/3 encapsulation flexible-ethernet-services
set interfaces ge-1/0/3 unit 0 encapsulation vlan-bridge
set interfaces ge-1/0/3 unit 0 vlan-id-range 100-110
set interfaces ge-1/0/4 flexible-vlan-tagging
set interfaces ge-1/0/4 encapsulation flexible-ethernet-services
set interfaces ge-1/0/4 unit 0 encapsulation vlan-bridge
set interfaces ge-1/0/4 unit 0 vlan-id-range 100-110
set interfaces ge-1/0/5 gigether-options 802.3ad ae0
set interfaces ge-1/1/0 gigether-options 802.3ad ae1
set interfaces ae0 flexible-vlan-tagging
set interfaces ae0 encapsulation flexible-ethernet-services
set interfaces ae0 aggregated-ether-options lacp active
set interfaces ae0 aggregated-ether-options lacp system-priority 100
set interfaces ae0 aggregated-ether-options lacp system-id 00:00:00:00:00:05
set interfaces ae0 aggregated-ether-options lacp admin-key 1
set interfaces ae0 aggregated-ether-options mc-ae mc-ae-id 5
set interfaces ae0 aggregated-ether-options mc-ae redundancy-group 10
set interfaces ae0 aggregated-ether-options mc-ae chassis-id 0
set interfaces ae0 aggregated-ether-options mc-ae mode active-active
set interfaces ae0 aggregated-ether-options mc-ae status-control standby
set interfaces ae0 unit 0 encapsulation vlan-bridge
set interfaces ae0 unit 0 vlan-id-range 100-110
set interfaces ae0 unit 0 multi-chassis-protection 10.100.100.1 interface ge-1/0/4.0
set interfaces ae1 flexible-vlan-tagging
set interfaces ae1 encapsulation flexible-ethernet-services
set interfaces ae1 aggregated-ether-options lacp active
set interfaces ae1 aggregated-ether-options lacp system-priority 100
set interfaces ae1 aggregated-ether-options lacp system-id 00:00:00:00:00:05
set interfaces ae1 aggregated-ether-options lacp admin-key 1
set interfaces ae1 aggregated-ether-options mc-ae mc-ae-id 10
set interfaces ae1 aggregated-ether-options mc-ae redundancy-group 10
set interfaces ae1 aggregated-ether-options mc-ae chassis-id 0
set interfaces ae1 aggregated-ether-options mc-ae mode active-active
set interfaces ae1 aggregated-ether-options mc-ae status-control standby
set interfaces ae1 unit 0 encapsulation vlan-bridge
set interfaces ae1 unit 0 vlan-id-range 100-110
set interfaces ae1 unit 0 multi-chassis-protection 10.100.100.1 interface ge-1/0/4.0
set bridge-domains bd0 domain-type bridge
set bridge-domains bd0 vlan-id all
set bridge-domains bd0 service-id 20
set bridge-domains bd0 interface ae1.0
set bridge-domains bd0 interface ge-1/0/3.0
set bridge-domains bd0 interface ge-1/0/4.0
```

```
set bridge-domains bd0 interface ae0.0
set protocols iccp local-ip-addr 10.100.100.2
set protocols iccp peer 10.100.100.1 redundancy-group-id-list 10
set protocols iccp peer 10.100.100.1 liveness-detection minimum-interval 1000
set switch-options service-id 10
```

Configuring the PE1 Router

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 Router PE1:

1. Specify the number of aggregated Ethernet interfaces to be created.

```
[edit chassis]
user@PE1# set aggregated-devices ethernet device-count 5
```

2. Specify the members to be included within the aggregated Ethernet bundles.

```
[edit interfaces]
user@PE1# set ge-1/0/1 gigether-options 802.3ad ae1
user@PE1# set ge-1/0/6 gigether-options 802.3ad ae0
```

3. Configure the interfaces that connect to senders or receivers, the ICL interfaces, and the ICCP interfaces.

```
[edit interfaces]
user@PE1# set ge-1/1/1 flexible-vlan-tagging
user@PE1# set ge-1/1/1 encapsulation flexible-ethernet-services
user@PE1# set ge-1/1/1 unit 0 encapsulation vlan-bridge
user@PE1# set ge-1/1/1 unit 0 vlan-id-range 100-110

user@PE1# set ge-1/1/4 flexible-vlan-tagging
user@PE1# set ge-1/1/4 encapsulation flexible-ethernet-services
user@PE1# set ge-1/1/4 unit 0 encapsulation vlan-bridge
user@PE1# set ge-1/1/4 unit 0 vlan-id-range 100-110

user@PE1# set ge-1/0/2 unit 0 family inet address 10.100.100.1/30
```

4. Configure parameters on the aggregated Ethernet bundles.

```
[edit interfaces ae0]
user@PE1# set flexible-vlan-tagging
user@PE1# set encapsulation flexible-ethernet-services
user@PE1# set unit 0 encapsulation vlan-bridge
user@PE1# set unit 0 vlan-id-range 100-110
user@PE1# set unit 0 multi-chassis-protection 10.100.100.2 interface ge-1/1/4.0
```

```
[edit interfaces ae1]
user@PE1# set flexible-vlan-tagging
user@PE1# set encapsulation flexible-ethernet-services
user@PE1# set unit 0 encapsulation vlan-bridge
user@PE1# set unit 0 vlan-id-range 100-110
user@PE1# set unit 0 multi-chassis-protection 10.100.100.2 interface ge-1/1/4.0
```

5. Configure LACP on the aggregated Ethernet bundles.

```
[edit interfaces ae0 aggregated-ether-options]
user@PE1# set lacp active
user@PE1# set lacp system-priority 100
user@PE1# set lacp system-id 00:00:00:00:00:05
user@PE1# set lacp admin-key 1
[edit interfaces ae1 aggregated-ether-options]
user@PE1# set lacp active
user@PE1# set lacp system-priority 100
user@PE1# set lacp system-id 00:00:00:00:00:05
user@PE1# set lacp admin-key 1
```

6. Configure the MC-LAG interfaces.

```
[edit interfaces ae0 aggregated-ether-options]
user@PE1# set mc-ae mc-ae-id 5
user@PE1# set mc-ae redundancy-group 10
user@PE1# set mc-ae chassis-id 1
user@PE1# set mc-ae mode active-active
user@PE1# set mc-ae status-control active
[edit interfaces ae1 aggregated-ether-options]
user@PE1# set mc-ae mc-ae-id 10
user@PE1# set mc-ae redundancy-group 10
user@PE1# set mc-ae chassis-id 1
user@PE1# set mc-ae mode active-active
user@PE1# set mc-ae status-control active
```

The multichassis aggregated Ethernet identification number (**mc-ae-id**) specifies which link aggregation group the aggregated Ethernet interface belongs to. The ae0 interfaces on Router PE1 and Router PE2 are configured with **mc-ae-id 5**. The ae1 interfaces on Router PE1 and Router PE2 are configured with **mc-ae-id 10**.

The **redundancy-group 10** statement is used by ICCP to associate multiple chassis that perform similar redundancy functions and to establish a communication channel so that applications on peering chassis can send messages to each other. The ae0 and ae1 interfaces on Router PE1 and Router PE2 are configured with the same redundancy group, **redundancy-group 10**.

The **chassis-id** statement is used by LACP for calculating the port number of the MC-LAG's physical member links. Router PE1 uses **chassid-id 1** to identify both its ae0 and ae1 interfaces. Router PE2 uses **chassis-id 0** to identify both its ae0 and ae1 interfaces.

The **mode** statement indicates whether an MC-LAG is in active-standby mode or active-active mode. Chassis that are in the same group must be in the same mode.

7. Configure a domain that includes the set of logical ports.

```
[edit bridge-domains bd0]
user@PE1# set domain-type bridge
user@PE1# set vlan-id all
user@PE1# set service-id 20
user@PE1# set interface ae0.0
user@PE1# set interface ae1.0
user@PE1# set interface ge-1/0/3.0
user@PE1# set interface ge-1/1/1.0
user@PE1# set interface ge-1/1/4.0
```

The ports within a bridge domain share the same flooding or broadcast characteristics in order to perform Layer 2 bridging.

The bridge-level **service-id** statement is required to link related bridge domains across peers (in this case Router PE1 and Router PE2), and must be configured with the same value.

8. Configure ICCP parameters.

```
[edit protocols iccp]
user@PE1# set local-ip-addr 10.100.100.1
user@PE1# set peer 10.100.100.2 redundancy-group-id-list 10
user@PE1# set peer 10.100.100.2 liveness-detection minimum-interval 1000
```

9. Configure the service ID at the global level.

```
[edit switch-options]
user@PE1# set service-id 10
```

You must configure the same unique network-wide configuration for a service in the set of PE routers providing the service. This service ID is required if the multichassis aggregated Ethernet interfaces are part of a bridge domain.

Results

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

```
user@PE1#  show bridge-domains
bd0 {
   domain-type bridge;
   vlan-id all;
```

```
    service-id 20;
    interface ae1.0;
    interface ge-1/0/3.0;
    interface ge-1/1/1.0;
    interface ge-1/1/4.0;
    interface ae0.0;
}
```

```
user@PE1#  show chassis
aggregated-devices {
    ethernet {
        device-count 5;
    }
}
```

```
user@PE1#  show interfaces
ge-1/0/1 {
    gigether-options {
        802.3ad ae1;
    }
}
ge-1/0/6 {
    gigether-options {
        802.3ad ae0;
    }
}
ge-1/0/2 {
    unit 0 {
        family inet {
            address 10.100.100.1/30;
        }
    }
}
ge-1/1/1 {
    flexible-vlan-tagging;
    encapsulation flexible-ethernet-services;
    unit 0 {
        encapsulation vlan-bridge;
        vlan-id-range 100-110;
    }
}
ge-1/1/4 {
    flexible-vlan-tagging;
    encapsulation flexible-ethernet-services;
    unit 0 {
        encapsulation vlan-bridge;
        vlan-id-range 100-110;
    }
}
ae0 {
    flexible-vlan-tagging;
    encapsulation flexible-ethernet-services;
    aggregated-ether-options {
        lacp {
```

```
        active;
        system-priority 100;
        system-id 00:00:00:00:00:05;
        admin-key 1;
      }
      mc-ae {
        mc-ae-id 5;
        redundancy-group 10;
        chassis-id 1;
        mode active-active;
        status-control active;
      }
    }
    unit 0 {
      encapsulation vlan-bridge;
      vlan-id-range 100-110;
      multi-chassis-protection 10.100.100.2 {
        interface ge-1/1/4.0;
      }
    }
}
ae1 {
    flexible-vlan-tagging;
    encapsulation flexible-ethernet-services;
    aggregated-ether-options {
      lacp {
        active;
        system-priority 100;
        system-id 00:00:00:00:00:05;
        admin-key 1;
      }
      mc-ae {
        mc-ae-id 10;
        redundancy-group 10;
        chassis-id 1;
        mode active-active;
        status-control active;
      }
    }
    unit 0 {
      encapsulation vlan-bridge;
      vlan-id-range 100-110;
      multi-chassis-protection 10.100.100.2 {
        interface ge-1/1/4.0;
      }
    }
}
```

```
user@PE1#  show protocols
iccp {
    local-ip-addr 10.100.100.1;
    peer 10.100.100.2 {
      redundancy-group-id-list 10;
      liveness-detection {
```

```
        minimum-interval 1000;
    }
  }
}
```

```
user@PE1#  show switch-options
service-id 10;
```

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

Repeat the procedure for Router PE2, using the appropriate interface names and addresses.

Configuring the CE Device

CLI Quick Configuration

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

Device CE

```
set chassis aggregated-devices ethernet device-count 2
set interfaces ge-2/0/2 gigether-options 802.3ad ae0
set interfaces ge-2/0/3 gigether-options 802.3ad ae0
set interfaces ge-2/1/6 flexible-vlan-tagging
set interfaces ge-2/1/6 encapsulation flexible-ethernet-services
set interfaces ge-2/1/6 unit 0 encapsulation vlan-bridge
set interfaces ge-2/1/6 unit 0 vlan-id-range 100-110
set interfaces ae0 flexible-vlan-tagging
set interfaces ae0 encapsulation flexible-ethernet-services
set interfaces ae0 aggregated-ether-options lacp active
set interfaces ae0 aggregated-ether-options lacp system-priority 100
set interfaces ae0 unit 0 encapsulation vlan-bridge
set interfaces ae0 unit 0 vlan-id-range 100-500
set bridge-domains bd0 domain-type bridge
set bridge-domains bd0 vlan-id all
set bridge-domains bd0 interface ge-2/1/6.0
set bridge-domains bd0 interface ae0.0
```

Configuring the CE Device

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 the CE device:

1. Specify the number of aggregated Ethernet interfaces to be created.

```
[edit chassis]
user@CE# set aggregated-devices ethernet device-count 2
```

2. Specify the members to be included within the aggregated Ethernet bundle.

```
[edit interfaces]
user@CE# set ge-2/0/2 gigether-options 802.3ad ae0
user@CE# set ge-2/0/3 gigether-options 802.3ad ae0
```

3. Configure an interface that connects to senders or receivers.

```
[edit interfaces ge-2/1/6]
user@CE# set flexible-vlan-tagging
user@CE# set encapsulation flexible-ethernet-services
user@CE# set unit 0 encapsulation vlan-bridge
user@CE# set unit 0 vlan-id-range 100-110
```

4. Configure parameters on the aggregated Ethernet bundle.

```
[edit interfaces ae0]
user@CE# set flexible-vlan-tagging
user@CE# set encapsulation flexible-ethernet-services
user@CE# set unit 0 encapsulation vlan-bridge
user@CE# set unit 0 vlan-id-range 100-500
```

5. Configure LACP on the aggregated Ethernet bundle.

```
[edit interfaces ae0 aggregated-ether-options]
user@CE# set lacp active
user@CE# set lacp system-priority 100
```

The **active** statement initiates transmission of LACP packets.

For the **system-priority** statement, a smaller value indicates a higher priority. The device with the lower system priority value determines which links between LACP partner devices are active and which are in standby mode for each LACP group. The device on the controlling end of the link uses port priorities to determine which ports are bundled into the aggregated bundle and which ports are put in standby mode. Port priorities on the other device (the noncontrolling end of the link) are ignored.

6. Configure a domain that includes the set of logical ports.

```
[edit bridge-domains bd0]
user@CE# set domain-type bridge
user@CE# set vlan-id all
user@CE# set interface ge-2/1/6.0
user@CE# set interface ae0.0
```

The ports within a bridge domain share the same flooding or broadcast characteristics in order to perform Layer 2 bridging.

Results

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

```
user@CE# show bridge-domains
bd0 {
   domain-type bridge;
   vlan-id all;
   interface ge-2/1/6.0;
   interface ae0.0;
}
```

```
user@CE# show chassis
aggregated-devices {
   ethernet {
      device-count 2;
   }
}
```

```
user@CE# show interfaces
ge-2/0/2 {
   gigether-options {
      802.3ad ae0;
   }
}
ge-2/0/3 {
   gigether-options {
      802.3ad ae0;
   }
}
ge-2/1/6 {
   flexible-vlan-tagging;
   encapsulation flexible-ethernet-services;
   unit 0 {
      encapsulation vlan-bridge;
      vlan-id-range 100-110;
   }
}
ae0 {
   flexible-vlan-tagging;
   encapsulation flexible-ethernet-services;
   aggregated-ether-options {
      lacp {
         active;
         system-priority 100;
      }
   }
   unit 0 {
      encapsulation vlan-bridge;
      vlan-id-range 100-500;
   }
```

```
}
```

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

Configuring the Provider Router

CLI Quick Configuration

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

Router P

```
set chassis aggregated-devices ethernet device-count 2
set interfaces ge-1/0/5 gigether-options 802.3ad ae1
set interfaces ge-1/0/11 gigether-options 802.3ad ae1
set interfaces ge-1/1/3 flexible-vlan-tagging
set interfaces ge-1/1/3 encapsulation flexible-ethernet-services
set interfaces ge-1/1/3 unit 0 encapsulation vlan-bridge
set interfaces ge-1/1/3 unit 0 vlan-id-range 100-500
set interfaces ae1 flexible-vlan-tagging
set interfaces ae1 encapsulation flexible-ethernet-services
set interfaces ae1 aggregated-ether-options lacp active
set interfaces ae1 aggregated-ether-options lacp system-priority 100
set interfaces ae1 unit 0 encapsulation vlan-bridge
set interfaces ae1 unit 0 vlan-id-range 100-110
set bridge-domains bd0 vlan-id all
set bridge-domains bd0 domain-type bridge
set bridge-domains bd0 interface ge-1/1/3.0
set bridge-domains bd0 interface ae1.0
```

Configuring the P Router

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 the P router:

1. Specify the number of aggregated Ethernet interfaces to be created.

```
[edit chassis]
user@P# set aggregated-devices ethernet device-count 2
```

2. Specify the members to be included within the aggregated Ethernet bundle.

```
[edit interfaces]
user@P# set ge-1/0/5 gigether-options 802.3ad ae1
user@P# set ge-1/0/11 gigether-options 802.3ad ae1
```

3. Configure an interface that connects to senders or receivers.

```
[edit interfaces ge-1/1/3]
user@P# set flexible-vlan-tagging
user@P# set encapsulation flexible-ethernet-services
user@P# set unit 0 encapsulation vlan-bridge
user@P# set unit 0 vlan-id-range 100-500
```

4. Configure parameters on the aggregated Ethernet bundle.

```
[edit interfaces ae1]
user@P# set flexible-vlan-tagging
user@P# set encapsulation flexible-ethernet-services
user@P# set unit 0 encapsulation vlan-bridge
user@P# set unit 0 vlan-id-range 100-110
```

5. Configure LACP on the aggregated Ethernet bundle.

```
[edit interfaces ae1 aggregated-ether-options]
user@P# set lacp active
user@P# set lacp system-priority 100
```

6. Configure a domain that includes the set of logical ports.

```
[edit bridge-domains bd0]
user@P# set vlan-id all
user@P# set domain-type bridge
user@P# set interface ge-1/1/3.0
user@P# set interface ae1.0
```

Results

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

```
user@P#  show bridge-domains
bd0 {
   domain-type bridge;
   vlan-id all;
   interface ge-1/1/3.0;
   interface ae1.0;
}
```

```
user@P#  show chassis
aggregated-devices {
   ethernet {
      device-count 2;
```

```
    }
}
```

```
user@P# show interfaces
ge-1/0/5 {
    gigether-options {
        802.3ad ae1;
    }
}
ge-1/0/11 {
    gigether-options {
        802.3ad ae1;
    }
}
ge-1/1/3 {
    flexible-vlan-tagging;
    encapsulation flexible-ethernet-services;
    unit 0 {
        encapsulation vlan-bridge;
        vlan-id-range 100-500;
    }
}
ae1 {
    flexible-vlan-tagging;
    encapsulation flexible-ethernet-services;
    aggregated-ether-options {
        lacp {
            active;
            system-priority 100;
        }
    }
    unit 0 {
        encapsulation vlan-bridge;
        vlan-id-range 100-110;
    }
}
```

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

Verification

Confirm that the configuration is working properly by running the following commands:

- **show iccp**
- **show interfaces ae0**
- **show interfaces ae1**
- **show interfaces mc-ae**
- **show l2-learning instance**

Release History Table

Release	Description
15.1	Starting with Junos OS Release 15.1, you can configure MC-LAG interfaces on logical systems on EX9200 switches.
14.1	Starting in Junos OS Release 14.1, you can configure MC-LAG interfaces on logical systems within a router.

CHAPTER 5

Enabling High Availability in Layer 3 Networks Using VRRP and MAC Synchronization for MC-LAG

High Availability in Layer 3 Networks Using VRRP and MAC Address Synchronization for MC-LAG

Active-Active Bridging and VRRP over IRB Functionality Overview

Active-active bridging and VRRP over IRB support extends multichassis link aggregation group (MC-LAG) by adding the following functionality to MX Series routers and QFX Series switches:

- Interchassis link (ICL) pseudowire interface or Ethernet interface (ICL-PL field) for active-active bridging
- Active-active bridging
- VRRP over IRB for active-active bridging
- A single bridge domain not corresponding to two redundancy group IDs

How Active-Active Bridging over IRB Functionality Works

Active-Active bridging over IRB functionality uses the address resolution protocol (ARP) Active-Active MC-LAG.

Suppose one of the PE routers issues an ARP request and another PE router gets the response and, because of the aggregated Ethernet distribution logic, the ARP resolution is not successful. Junos OS uses ARP response packet snooping to perform active-active multichassis link aggregation group support, providing synchronization without the need to maintain any specific state.

Address Resolution Protocol Active-Active MC-LAG Support Methodology

Suppose one of the PE routers issues an ARP request and another PE router gets the response and, because of the aggregated Ethernet distribution logic, the ARP resolution is not successful. Junos OS uses ARP response packet snooping to perform active-active

multichassis link aggregation group support, providing easy synchronization without the need to maintain any specific state.

Benefits of Active-Active Bridging and VRRP over IRB Functionality

Benefits of active-active bridging and VRRP over IRB functionality include:

- An MC-LAG reduces operational expenses by providing active-active links with a LAG, eliminates the need for Spanning Tree Protocol (STP), and provides faster Layer 2 convergence upon link and device failures.
- An MC-LAG adds node-level redundancy to the normal link-level redundancy that a LAG provides. An MC-LAG improves network resiliency, which reduces network down time as well as expenses.
- In data centers, it is desirable for servers to have redundant connections to the network. You probably want active-active connections along with links from any server to at least two separate routers.
- An MC-LAG allows you to bond two or more physical links into a logical link between two routers or between a server and a router, which improves network efficiency. An MC-LAG enables you to load-balance traffic on multiple physical links. If a link fails, the traffic can be forwarded through the other available link, and the logical aggregated link remains in the UP state.

Where Can I Use Active-Active Bridging and VRRP over IRB Functionality?

Active-active bridging and Virtual Router Redundancy Protocol (VRRP) over integrated routing and bridging (IRB) is supported on MX Series routers and QFX Series switches.

MC-LAG Functions in an Active-Active Bridging Domain

The following functions are supported for MC-LAG in an active-active bridging domain:

- MC-LAG is supported only between two chassis, using an interchassis link (ICL) pseudowire interface or Ethernet interface (ICL-PL field) for **active-active bridging**, and **active-active bridging** VRRP over IRB for **active-active bridging**.
- For VPLS networks, you can configure the aggregated Ethernet (aeX) interfaces on MC-LAG devices with the **encapsulation ethernet-vpls** statement to use Ethernet VPLS encapsulation on Ethernet interfaces that have VPLS enabled and that must accept packets carrying standard Tag Protocol ID (TPID) values or the **encapsulation vlan-vpls** statement to use Ethernet VLAN encapsulation on VPLS circuits.
- Layer 2 circuit functionalities are supported with **ethernet-ccc** as the encapsulation mode.
- Network topologies in a triangular and square pattern are supported. In a triangular network design, with equal-cost paths to all redundant nodes, slower, timer-based convergence can possibly be prevented. Instead of indirect neighbor or route loss detection using hellos and dead timers, you can identify the physical link loss and denote a path as unusable and reroute all traffic to the alternate equal-cost path. In a square network design, depending on the location of the failure, the routing protocol

might converge to identify a new path to the subnet or the VLAN, causing the convergence of the network to be slower.

- Interoperation of Link Aggregation Control Protocol (LACP) for MC-LAG devices is supported. LACP is one method of bundling several physical interfaces to form one logical interface. When LACP is enabled, the local and remote sides of the aggregated Ethernet links exchange protocol data units (PDUs), which contain information about the state of the link. You can configure Ethernet links to actively transmit PDUs, or you can configure the links to passively transmit them, sending out LACP PDUs only when the links receive the PDUs from another link. One side of the link must be configured as active for the link to be up.
- Active-standby mode is supported using LACP. When an MC-LAG operates in the active-standby mode, one of the router's ports only becomes active when failure is detected in the active links. In this mode, the provider edge (PE) routers perform an election to determine the active and standby routers.
- Configuration of the pseudowire status type length variable (TLV) is supported. The pseudowire status TLV is used to communicate the status of a pseudowire back and forth between two PE routers. The pseudowire status negotiation process ensures that a PE router reverts back to the label withdraw method for pseudowire status if its remote PE router neighbor does not support the pseudowire status TLV.
- The MC-LAG devices use Inter-Chassis Control Protocol (ICCP) to exchange the control information between two MC-LAG network devices.

Points to Remember When Configuring MC-LAG Active-Active Bridge Domains

Keep the following points in mind when you configure MC-LAG in an active-active bridging domain:

- A single bridge domain cannot be associated with two redundancy groups. You cannot configure a bridge domain to contain logical interfaces from two different multichassis aggregated Ethernet interfaces and associate them with different redundancy group IDs by using the **redundancy group group-id** statement at the **[edit interfaces ae*X* aggregated-ether-options]** hierarchy level.
- You must configure logical interfaces in a bridge domain from a single multichassis aggregated Ethernet interface and associate it with a redundancy group. You must configure a service ID by including the **service-id *vid*** statement at the **[edit bridge-domains *bd-name*]** hierarchy level for multichassis aggregated Ethernet interfaces if you configure logical interfaces on multichassis aggregated Ethernet interfaces that are part of the bridge domain.

More Data Traffic Forwarding Rules

In active-active bridging and VRRP over IRB topographies, network interfaces are categorized into three different interface types, as follows:

S-Links—Single-homed link (S-Link) terminating on MC-LAG-N device or MC-LAG in active-standby mode. In Figure 10 on page 133, interfaces ge-0/0/0.0 and ge-1/0/0.0 are S-Links.

MC-Links—MC-LAG links. In Figure 10 on page 133, interface ae0.0 is the MC-Link.

ICL—Interchassis link.

Based on incoming and outgoing interface types, some constraints are added to the Layer 2 forwarding rules for MC-LAG configurations, as described in the data traffic forwarding rules. Note that if only one of the MC-LAG member link is in the UP state, it is considered an S-Link.

The following data traffic forwarding rules apply:

1. When an MC-LAG network receives a packet from a local MC-Link or S-Link, the packet is forwarded to other local interfaces, including S-Links and MC-Links based on the normal Layer 2 forwarding rules and on the configuration of the **mesh-group** and **no-local-switching** statements. If MC-Links and S-Links are in the same mesh group and their **no-local-switching** statements are enabled, the received packets are only forwarded upstream and not sent to MC-Links and S-Links.

   NOTE: The functionality described in Rule 2 is *not* supported.

2. The following circumstances determine whether or not an ICL receives a packet from a local MC-Link or S-Link:

   a. If the peer MC-LAG network device has S-Links or MC-LAGs that do not reside on the local MC-LAG network device

   b. Whether or not interfaces on two peering MC-LAG network devices are allowed to talk to each other only if both a. and b. are true. Traffic is always forwarded to the ICL.

3. When an MC-LAG network receives a packet from the ICL, the packet is forwarded to all local S-Links and active MC-LAGs that do not exist in the MC-LAG network that the packet comes from.

4. NOTE: The topology shown in Figure 7 on page 129 is *not* supported.

   In certain cases, for example the topology shown in Figure 7 on page 129, there could be a loop caused by the ICL. To break the loop, one of the following mechanisms could be used:

   a. Run certain protocols, such as STP. In this case, whether packets received on one ICL are forwarded to other ICLs is determined by using Rule 3.

   b. Configure the ICL to be fully meshed among the MC-LAG network devices. In this case, traffic received on the ICL would not be forwarded to any other ICLs.

   In either case, duplicate packets could be forwarded to the MC-LAG clients. Consider the topology shown in Figure 7 on page 129, where if network routing instance N1 receives a packet from ge-0/0/0.0, it could be flooded to ICL1 and ICL3.

When receiving from ICL1 and ICL3, network routing instances N3 and N2 could flood the same packet to MCL2, as shown in Figure 7 on page 129. To prevent this from happening, the ICL designated forwarder should be elected between MC-LAG peers, and traffic received on an ICL could be forwarded to the active-active MC-LAG client by the designated forwarder only.

Figure 19: Loop Caused by the ICL Links

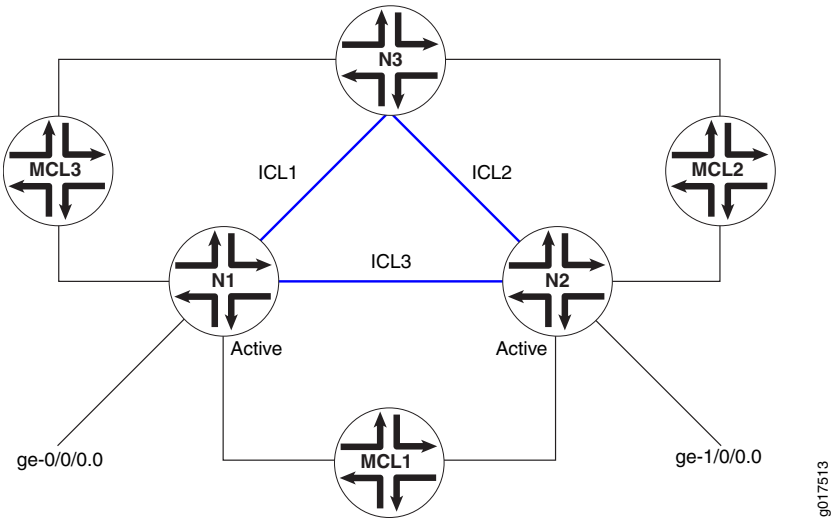

5. When received from an ICL, traffic should not be forwarded to the core-facing client link connection between two provider edge (PE) devices (MC-Link) if the peer chassis's (where the traffic is coming from) MC-Link is UP.

How to Configure MC-LAG Active-Active Bridge Domains

For a MC-LAG configured in an active-active bridge domain and with VRRP configured over an IRB interface, you must include the **accept-data** statement at the **[edit interfaces *interface-name* unit *logical-unit-number* family inet address *address* vrrp-group *group-id*]** hierarchy level to enable the router that functions as the master router to accept all packets destined for the virtual IP address.

On an MC-LAG, if you modify the source MAC address to be the virtual MAC address, you must specify the virtual IP address as the source IP address instead of the physical IP address. In such a case, the **accept-data** option is required for VRRP to prevent ARP from performing an incorrect mapping between IP and MAC addresses for customer edge (CE) devices. The **accept-data** attribute is needed for VRRP over IRB interfaces in MC-LAG to enable OSPF or other Layer 3 protocols and applications to work properly over multichassis aggregated Ethernet (mc-aeX) interfaces.

NOTE: On an MC-LAG, the unit number associated with aggregated Ethernet interfaces on provider edge router PE1 must match the unit number associated with aggregated Ethernet interfaces on provider edge router PE2. If the unit numbers differ, MAC address synchronization does not happen. As a result, the status of the MAC address on the remote provider edge router remains in a pending state.

If you are using the VRRP over IRB or RVI method to enable Layer 3 functionality, you must configure static ARP entries for the IRB or RVI interface of the remote MC-LAG peer to allow routing protocols to run over the IRB or RVI interfaces.

MAC Address Management

If an MC-LAG is configured to be active-active, upstream and downstream traffic could go through different MC-LAG network devices. Since the media access control (MAC) address is learned only on one of the MC-LAG network devices, the reverse direction's traffic could be going through the other MC-LAG network and be flooded unnecessarily. Also, a single-homed client's MAC address is only learned on the MC-LAG network device it is attached to. If a client attached to the peer MC-LAG network needs to communicate with that single-homed client, then traffic would be flooded on the peer MC-LAG network device. To avoid unnecessary flooding, whenever a MAC address is learned on one of the MC-LAG network devices, it gets replicated to the peer MC-LAG network device. The following conditions should be applied when MAC address replication is performed:

- MAC addresses learned on an MC-LAG of one MC-LAG network device should be replicated as learned on the same MC-LAG of the peer MC-LAG network device.
- MAC addresses learned on single-homed customer edge (CE) clients of one MC-LAG network device should be replicated as learned on the ICL-PL interface of the peer MC-LAG network device.
- MAC addresses learned on MC-LAG VE clients of one MC-LAG network device should be replicated as learned on the corresponding VE interface of the peer MC-LAG network device.
- MAC address learning on an ICL is disabled from the data path. It depends on software to install MAC addresses replicated through Inter-Chassis Control Protocol (ICCP).

MAC Aging

MAC aging support in Junos OS extends aggregated Ethernet logic for a specified MC-LAG. A MAC address in software is deleted until all Packet Forwarding Engines have deleted the MAC address. In the case of an MC-LAG, a remote provider edge is treated as a remote Packet Forwarding Engine and has a bit in the MAC data structure.

Layer 3 Routing

In general, when an MC-LAG is configured to provide Layer 3 routing functions to downstream clients, the MC-LAG network peers should be configured to provide the same gateway address to the downstream clients. To the upstream routers, the MC-LAG

network peers could be viewed as either equal-cost multipath (ECMP) or two routes with different preference values.

Junos OS supports active-active MC-LAGs by using VRRP over IRB. Junos OS also supports active-active MC-LAGs by using IRB MAC address synchronization. You must configure IRB using the same IP address across MC-LAG peers. IRB MAC synchronization is supported on 32-bit interfaces and interoperates with earlier MPC and MIC releases.

To ensure that Layer 3 operates properly, instead of dropping the Layer 3 packet, the VRRP backup attempts to perform routing functions if the packet is received on an MC-LAG. A VRRP backup sends and responds to Address Resolution Protocol (ARP) requests.

For ARP, the same issue exists as with Layer 2 MAC addresses. Once ARP is learned, it must be replicated to the MC-LAG through ICCP. The peer must install an ARP route based on the ARP information received through ICCP.

For ARP aging, ARP requests on the MC-LAG peers can be aged out independently.

Topologies Supported for MC-LAG Active-Active Bridge Domains

The topologies shown in Figure 8 on page 132 and Figure 9 on page 132 are supported. These figures use the following abbreviations:

- Aggregated Ethernet (AE)
- Interchassis link (ICL)
- Multichassis link (MCL)

Figure 20: Single Multichassis Link

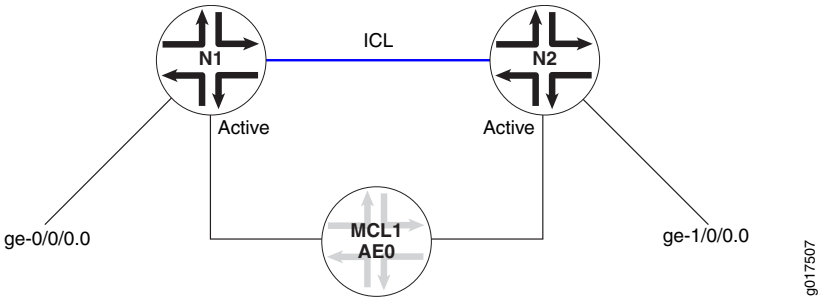

Figure 21: Dual Multichassis Link

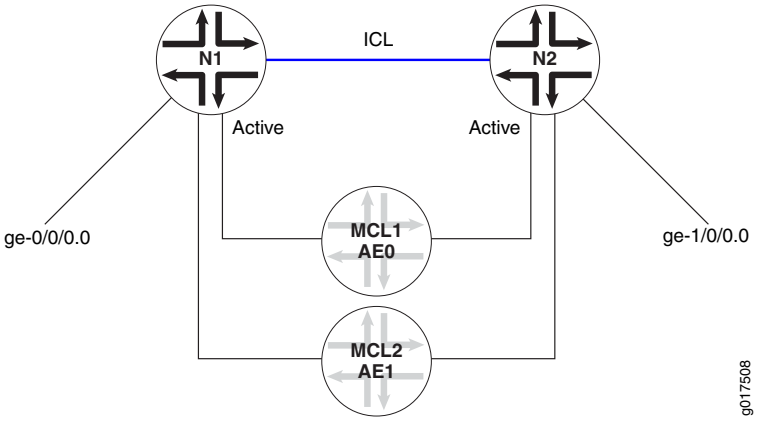

Potential Problems When Configuring MC-LAG Active-Active Bridge Domains

When configured to be active-active, the client device load-balances the traffic to the peering MC-LAG network devices. In a bridging environment, this could potentially cause the following problems:

- Traffic received on the MC-LAG from one MC-LAG network device could be looped back to the same MC-LAG on the other MC-LAG network device.
- Duplicated packets could be received by the MC-LAG client device.
- Traffic could be unnecessarily forwarded on the interchassis link.

To better illustrate the problems listed, consider Figure 10 on page 133, where an MC-LAG device MCL1 and single-homed clients ge-0/0/0.0 and ge-1/0/0.0 are allowed to talk to each other through an ICL. These problems could occur:

Figure 22: MC-LAG Device and Single-Homed Client

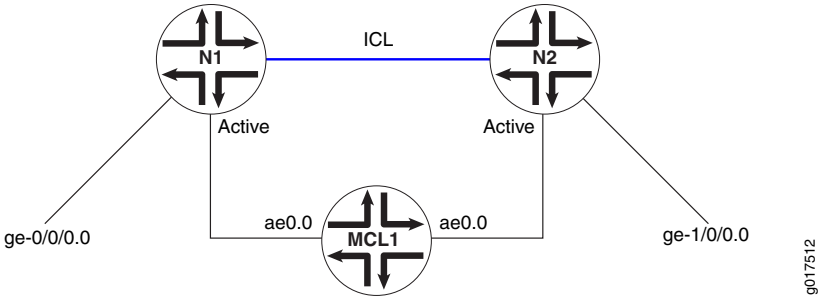

- Traffic received on network routing instance N1 from MCL1 could be flooded to ICL to reach network routing instance N2. Once it reaches network routing instance N2, it could flood again to MCL1.
- Traffic received on interface ge-0/0/0.0 could be flooded to MCL1 and ICL on network routing instance N1. Once network routing instance N2 receives such traffic from ICL, it could again be flooded to MCL1.
- If interface ge-1/0/0.0 does not exist on network routing instance N2, traffic received from interface ge-0/0/0.0 or MCL1 on network routing instance N1 could be flooded

to network routing instance N2 through ICL unnecessarily since interface ge-0/0/0.0 and MCL1 could reach each other through network routing instance N1.

Restrictions When Configuring MC-LAG Active-Active Bridge Domains

In an IPv6 network, you cannot configure an MC-LAG in an active-active bridge domain if you specified the **vlan-id none** statement at the **[edit bridge-domain *bd-name*]** hierarchy level. The **vlan-id none** statement that enables the removal of the incoming VLAN tags identifying a Layer 2 logical interface when packets are sent over VPLS pseudowires is not supported for IPv6 packets in an MC-LAG.

The following functionality is *not* supported for MC-LAG active-active bridge domains:

- Virtual private LAN service (VPLS) within the core
- Bridged core
- Topology as described in Rule 4 of "More Data Traffic Forwarding Rules" on page 128
- Routed multichassis aggregated Ethernet interface, where the VRRP backup router is used in the edge of the network
- Track object, where in the case of an MC-LAG, the status of the uplinks from the provider edge can be monitored, and the MC-LAG can act on the status
- Mixed mode (active-active MC-LAG is supported on MX Series routers with MPC or MIC interfaces only)

  All interfaces in the bridge domain that are multichassis aggregated Ethernet active-active must be on MPCs or MICs.

The topologies shown in Figure 11 on page 134, Figure 12 on page 134, and Figure 13 on page 134 are *not* supported:

Figure 23: Interchassis Data Link Between Active-Active Nodes

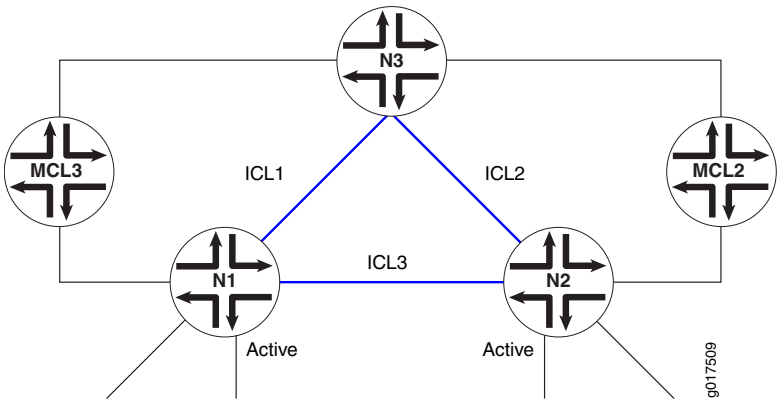

Figure 24: Active-Active MC-LAG with Single MC-LAG

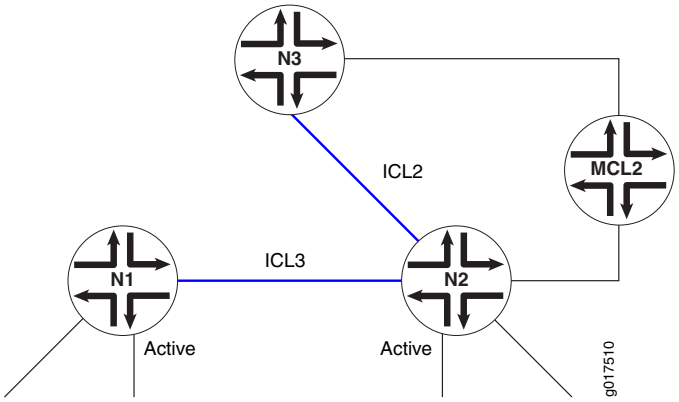

Figure 25: Active-Active MC-LAG with Multiple Nodes on a Single Multichassis Link

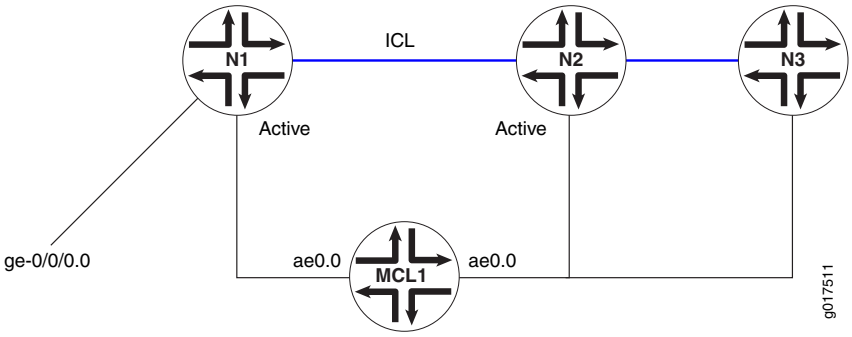

NOTE: A redundancy group cannot span more than two routers.

IGMP Snooping on Active-Active MC-LAG

IGMP Snooping on Active-Active MC-LAG

For multicast to work in an active-active MC-LAG scenario, the typical topology is as shown in Figure 14 on page 135 and Figure 15 on page 136 with interested receivers over S-links and MC-Links. Starting in Junos OS Release 11.2, support is extended for sources connected over the Layer 2 interface.

If an MC-LAG is configured to be active-active, reports from MC-LAG clients could reach any of the MC-LAG network device peers. Therefore, the IGMP snooping module needs to replicate the states such that the Layer 2 multicast route state on both peers are the same. Additionally for S-Link clients, snooping needs to replicate these joins to its snooping peer, which in the case of Layer 3 connected source, passes this information to the PIM on IRB to enable the designated router to pull traffic for these groups,

The ICL should be configured as a router facing interface. For the scenario where traffic arrives through a Layer 3 interface, it is a requirement to have PIM and IGMP enabled on the IRB interface configured on the MC-LAG network device peers.

Figure 26: Multicast Topology with Source Connected Through Layer 3

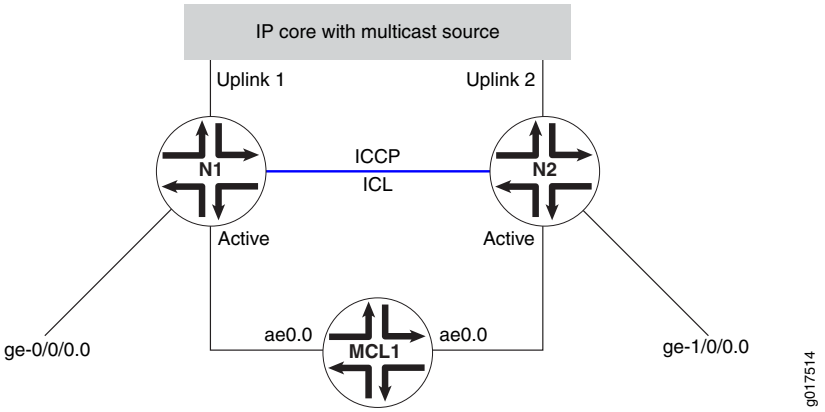

With reference to Figure 14 on page 135, either Device N1 or N2 becomes a designated router (for this example, N1 is the designated router). Router N1 therefore pulls the multicast traffic from the core. Once multicast data hits the network Device N1, the data is forwarded based on the snooping learned route.

For MC-Link clients, data is forwarded through N1. In the case of failover of the MC-Links, the data reaches the client through N2. For S-Link clients on N1, data would be forwarded through normal snooping routes.

For S-Link clients on N2, data is forwarded through the ICL interface. Layer 2 multicast routes on N1 do not show these groups unless there is interest for the same group over MC-Links or over S-Links on N1. For the IRB scenario, the IGMP membership and Layer 3 multicast route on N1 does however show these groups learned over the IRB interface.

Therefore, for a case where a specific group interest is only on the S-Link on N2, data arriving on N1 reaches N2 through the default route, and the Layer 2 multicast route on N2 has the S-Link in the outgoing interface list.

Figure 27: Multicast Topology with Source Connected Through MC-Link

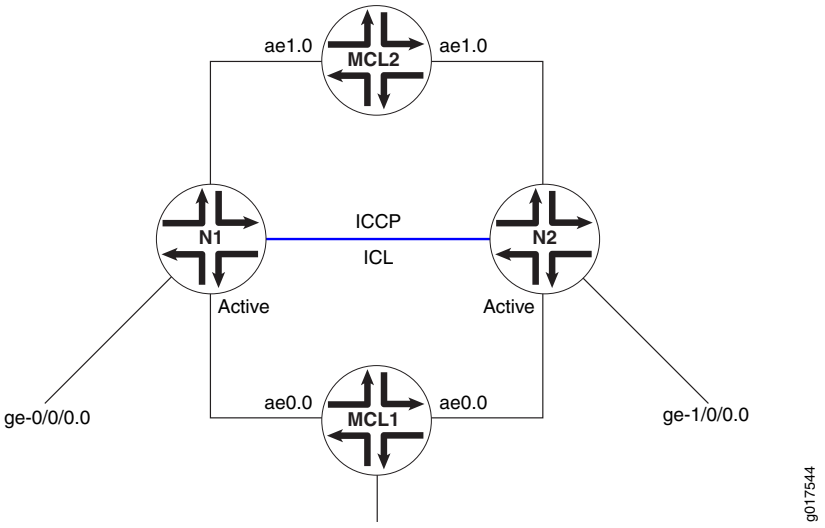

In Figure 15 on page 136, MCL1 and MCL2 are on different devices, and the multicast source or IGMP querier is connected through MCL2. The data forwarding behavior seen is similar to that explained for multicast topology with source connected through Layer 3.

NOTE: IGMP snooping should not be configured in proxy mode. There should be no IGMP hosts or IGMP or PIM routers sitting on the ICL interface.

Up and Down Event Handling

The following conditions apply to up and down event handling:

- If the Inter-Chassis Control Protocol (ICCP) connection is UP but the ICL interface goes DOWN, the router configured as the backup brings down all the multichassis aggregated Ethernet interfaces shared with the peer that is connected to ICL. This ensures that there are no loops in the network. Otherwise, both PEs become PIM-designated routers and, hence, forward multiple copies of the same packet to the customer edge.
- If the ICCP connection is UP and the ICL comes UP, the router configured as the backup brings up the multichassis aggregated Ethernet interfaces shared with the peer.
- If both the ICCP connection and the ICL are DOWN, the router configured as the backup brings up the multichassis aggregated Ethernet interfaces shared with the peer.
- The Layer 2 address learning process (l2ald) does not store the information about a MAC address learned from a peer in the kernel. If l2ald restarts, and if the MAC address was not learned from the local multichassis aggregated Ethernet interface, l2ald clears the MAC addresses, which causes the router to flood the packets destined to this MAC address. This behavior is similar to that in a Routing Engine switchover. (Note that currently l2ald runs on a Routing Engine only when it is a master). Also, during the time l2ald is DOWN, ARP packets received from an ICCP peer are dropped. ARP retry takes care of this situation. This is the case with Routing Engine switchover, too.
- If ICCP restarts, l2ald does not identify that a MAC address was learned from a peer and, if the MAC address was learned only from the peer, that MAC address is deleted, and the packets destined to this MAC address are flooded.

Inter-Chassis Control Protocol

Inter-Chassis Control Protocol (ICCP) is used to synchronize configurations, states, and data.

ICCP supports the following types of state information:

- MC-LAG members and their operational states
- Single-homed members and their operational states

ICCP supports the following application database synchronization parameters:

- MAC addresses learned and to be aged
- ARP information learned over IRB

Inter-Chassis Control Protocol Message

ICCP messages and attribute-value pairs (AVPs) are used for synchronizing MAC address and ARP information.

IGMP Snooping in MC-LAG Active-Active Mode

IGMP snooping in MC-LAG active-active mode is supported on MX240 routers, MX480 routers, MX960 routers and QFX Series switches.

The following topics are included:

IGMP Snooping in MC-LAG Active-Active Mode Functionality

Multichassis link aggregation group (MC-LAG) active-active mode and IGMP snooping in active-standby mode are supported. MC-LAG allows one device to form a logical LAG interface with two or more network devices. MC-LAG provides additional benefits including node level redundancy, multihoming, and a loop-free Layer 2 network without running Spanning Tree Protocol (STP). The following features are supported:

- State synchronization between peers for IGMP snooping in a bridge domain with only Layer 2 interfaces
- Qualified learning
- Router-facing multichassis links

The following enhancements to active-active bridging and Virtual Router Redundancy Protocol (VRRP) over integrated routing and bridging (IRB) are supported:

- MC-LAG support for IGMP snooping in a pure Layer 2 switch
- MC-LAG support for IGMP snooping in bridge domains doing qualified learning
- Support for MC-Links being router-facing interfaces

The following functions are *not* supported:

- MC-LAG for VPLS instances
- MC-Links trunk ports
- Proxy mode for active-active
- Adding interchassis links to outgoing interfaces on an as needed basis

  Interchassis links can be added to the outgoing interface list as router-facing interfaces.

Typically Supported Network Topology for IGMP Snooping with MC-LAG Active-Active Bridging

Figure 28 on page 186 depicts a typical network topology over which IGMP snooping with MC-LAG active-active bridging is supported.

Figure 28: Typical Network Over Which Active-Active Is Supported

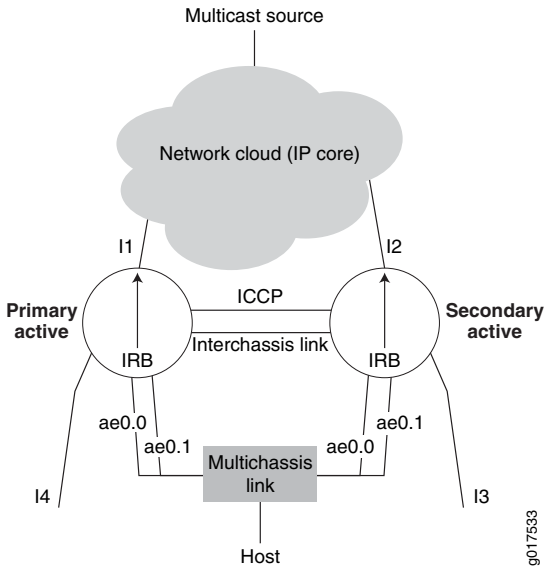

Interfaces I3 and I4 are single-homed interfaces. The multichassis links ae0.0 and ae0.1 belong to the same bridge domain in both the chassis. Interfaces I3, ae0.0, and ae0.1 are in the same bridge domain in the secondary active (S-A) router. Interfaces I4, ae0.0, and ae0.1 are in the same bridge domain in the primary active (P-A) router. Interfaces I3, I4, ae0.0, and ae0.1 are in the same learning domain as is the interchassis link (ICL) connecting the two chassis.

The primary active router is the chassis in which integrated routing and bridging has become PIM-DR. The secondary active router is the chassis in which integrated routing and bridging is not PIM-DR. Router P-A is the chassis responsible for pulling traffic from the IP core. Hence, PIM-DR election is used to avoid duplication of data traffic.

Learning domains are described in "Qualified Learning" on page 188.

For the IGMP speakers (hosts and routers) in the learning domain, P-A and S-A together should appear as one device with interfaces I4, I3, ae0.0, and ae0.1.

No duplicate control packets should be sent on multichassis links, meaning the control packet should be sent through only one link.

Control Plane State Updates Triggered by Packets Received on Remote Chassis

Following are the control plane state updates that are triggered by the packets received on remote chassis:

- The membership state in Layer 3 multicast routing is updated as a result of reports learned on remote legs of multichassis links and S-links attached to the remote chassis.
- The membership state and routing entry in snooping are updated when reports are received on the remote legs of a multichassis link.

NOTE:

- When reports are received on S-links attached to the remote chassis, the membership state or routing entry in snooping is not updated.
- When synchronizing multicast snooping state between PE routers, timers, such as the Group Membership Timeout timer, are not synchronized. When the synch notification is received, the remote PE router receiving the notification starts or restarts the relevant timer.
- The list of <s,g>s for which the state is maintained is the same in both the chassis under snooping as long as the outgoing interface lists involve only multichassis links.

Data Forwarding

This discussion assumes integrated routing and bridging on Router P-A is the PIM-DR. It pulls the traffic from sources in the core. Traffic might also come on Layer 2 interfaces in the bridge domain. For hosts directly connected to the P-A chassis, there is no change in the way data is delivered.

For delivering traffic to hosts connected to S-A (which is the non-DR) on the single-homed link like I3, we rely on the interchassis link. The traffic that hits P-A is sent over ICL to S-A to be delivered to the links that have reported interests in s,g and the links that are router-facing.

When the ae0 leg in P-A goes down, the hosts connected to the multichassis link receive traffic through ICL. In S-A, traffic received on ICL is sent to multichassis links in the outgoing interface list for which the ae counterpart in P-A is down.

Pure Layer 2 Topology Without Integrated Routing and Bridging

Figure 29 on page 188 shows that the chassis connecting to the PIM-DR is the primary active (P-A) router and the other is the secondary active (S-A) router.

Figure 29: Layer 2 Configuration Without Integrated Routing and Bridging

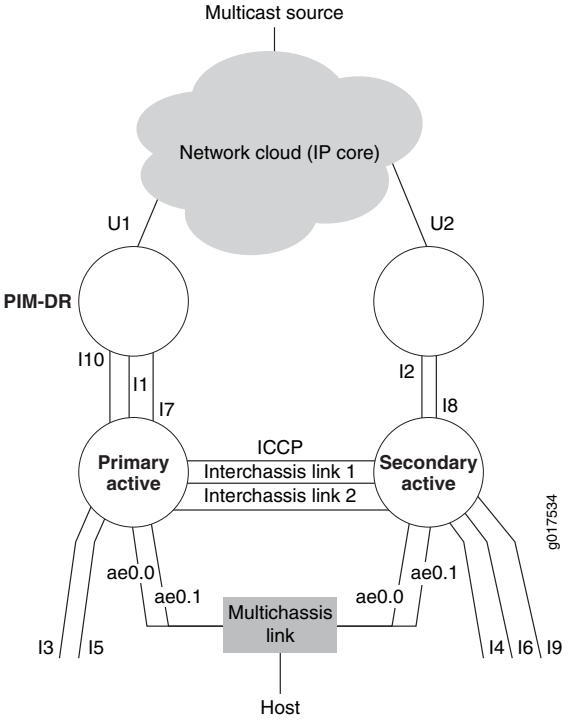

Qualified Learning

In this topology, interfaces I1, I2, I3, I4, I5, I6, I7, I8, I9, and I10 are single-homed interfaces. The multichassis links ae0.0 and ae0.1 belong to the same bridge domain in both the chassis. Interfaces I10, I1,I7,I3,I5, ae0.0 and ae0.1 are in same bridge domain, bd1 in P-A. Interfaces I9, I2, I8, I4, I6, ae0.0, and ae0.1 are in same bridge domain, bd1 in S-A.

This discussion assumes the following configuration:

- In P-A and S-A, qualified learning is ON in bd1.
- Interfaces I1, I2, I3, ae0.0, and I4 belong to vlan1, learning domain ld1.
- Interfaces I7, I8, I5, ae0.1, and I6 belong to vlan2, learning domain ld2.
- Interfaces I9 and I10 belong to vlan3, learning domain ld3.

For the IGMP speakers (hosts and routers) in the same learning domain ld1, P-A and S-A linked should appear to be one switch.

For the IGMP speakers (hosts and routers) in the same learning domain ld2, P-A and S-A linked should appear to be one switch.

Since there are no multichassis links in learning domain ld3, for the IGMP speakers (hosts and routers) in learning domain ld3, P-A and S-A will not appear to be one switch.

This discussion assumes interchassis link ICL1 corresponds to learning domain ld1 and interchassis link ICL2 corresponds to learning domain ld2.

Control packet flow is supported, with the exception of passing information to IRB.

Data Forwarding with Qualified Learning

This discussion assumes one learning domain (LD), ld1, and further assumes that interface I1 on Router P-A is connected to the PIM-DR in the learning domain and pulls the traffic from sources in the core.

For delivering traffic to hosts connected to Router S-A (which is the non-DR) on the single-homed link like I2, I4 (belonging to ld1), we rely on ICL1. The traffic that hits Router P-A on interface I1 is sent over interchassis link ICL1 to Router S-A to be delivered to the links that have reported interests in s,g or the links that are router-facing in learning domain ld1.

When the interface ae0 leg in Router P-A goes down, the hosts connected to the multichassis link receive traffic from interface I1 using the interchassis link ICL1. In Router S-A, traffic received on interchassis link ICL1 is sent to multichassis links in the outgoing interface list for which the aggregated Ethernet counterpart in Router P-A is down.

It is further assumed that interface I9 in Router S-A belongs to the learning domain ld3 with interests in s,g, and that interface I10 in learning domain ld3 in Router P-A receives traffic for s,g. Interface I9 does not receive data in this topology because there are no multichassis links (in a-a mode) and hence no interchassis link in learning domain ld3.

Static Groups on Single-Homed Interfaces

For multichassis links, the static group configuration should exist on both legs, and synchronization with the other chassis is not required.

Synchronization of the static groups on single-homed interfaces between the chassis is not supported. However, the addition of logical interfaces to the default outgoing interface list supports traffic delivery to the interface within a static configuration.

Router-Facing Interfaces as Multichassis Links

IGMP queries could arrive on either leg of the multichassis links, but in both peers, the multichassis link should be considered as router-facing.

Reports should exit only once from the multichassis link, that is, from only one leg.

The following MC-LAG support for IGMP snooping in IRB is provided:

- Non-proxy snooping
- Logical interfaces must be outgoing interfaces for all routes including the default route
- IGMP snooping in a pure Layer 2 switch

- IGMP snooping in bridge domains doing qualified learning
- Router-facing interface MC-Links

The following features are *not* supported:

- Proxy mode for active-active
- MC-LAG support for VPLS instances
- Trunk ports as multichassis links
- Adding logical interfaces to outgoing interfaces on an as need basis.

  However, logical interfaces are always added as a router-facing interface to the outgoing interface list.

See Also
- *Example: Configuring IGMP Snooping*

Example: Configuring IGMP Snooping in MC-LAG Active-Active Mode

This example shows how to configure Internet Group Management Protocol (IGMP) snooping for uninterrupted traffic flow with a multichassis link aggregation group (MC-LAG) in an active-active scenario.

Requirements

This example uses the following hardware and software components:

NOTE: **This example also applies to QFX10002 and QFX10008 switches.**

- Four Juniper Networks MX Series routers
- Junos OS Release 11.2 or later running on all four routers

Before you begin, make sure that Protocol Independent Multicast (PIM) and Internet Group Management Protocol (IGMP) are running on all interfaces that will receive multicast packets. IGMP is automatically enabled on all IPv4 interfaces on which you configure PIM.

Overview

When links are aggregated, the links can be treated as if they were a single link. Link aggregation increases bandwidth, provides graceful degradation as failure occurs, and

increases availability. MC-LAG provides redundant Layer 2 access connectivity at the node level. This enables two or more systems to share a common LAG endpoint. The multiple endpoints present a single logical chassis to the start point, and the start node does not need to be aware that MC-LAG is being used.

In this example, the CE router is not aware that its aggregated Ethernet links are connected to two separate PE devices. The two PE devices each have a LAG connected to the CE device. The configured mode is active-active, meaning that both PE routers' LAG ports are active and carrying traffic at the same time.

NOTE: The other possible mode is active-standby, in which one of the router's ports only becomes active when failure is detected in the active links. In active-standby mode, the PE routers perform an election to determine the active and standby routers.

From the perspective of the CE device, all four ports belonging to a LAG are connected to a single service provider device. Because the configured mode is active-active, all four ports are active, and the CE device load-balances the traffic to the peering PE devices. On the PE routers, a regular LAG is configured facing the CE device.

Inter-Chassis Control Protocol (ICCP) messages are sent between the two PE devices. These messages exchange MC-LAG configuration parameters and ensure that both chassis use the correct Link Aggregation Control Protocol (LACP) parameters when talking to the CE device.

The interchassis link-protection link (ICL) provides redundancy when a link failure occurs on one of the active links. The ICL-PL between the MC-LAG peering devices relays traffic that would otherwise be dropped due to a link failure.

Topology Diagram

Figure 30 on page 192 shows the topology used in this example.

Figure 30: IGMP Snooping in MC-LAG Active-Active Mode on MX Series Routers

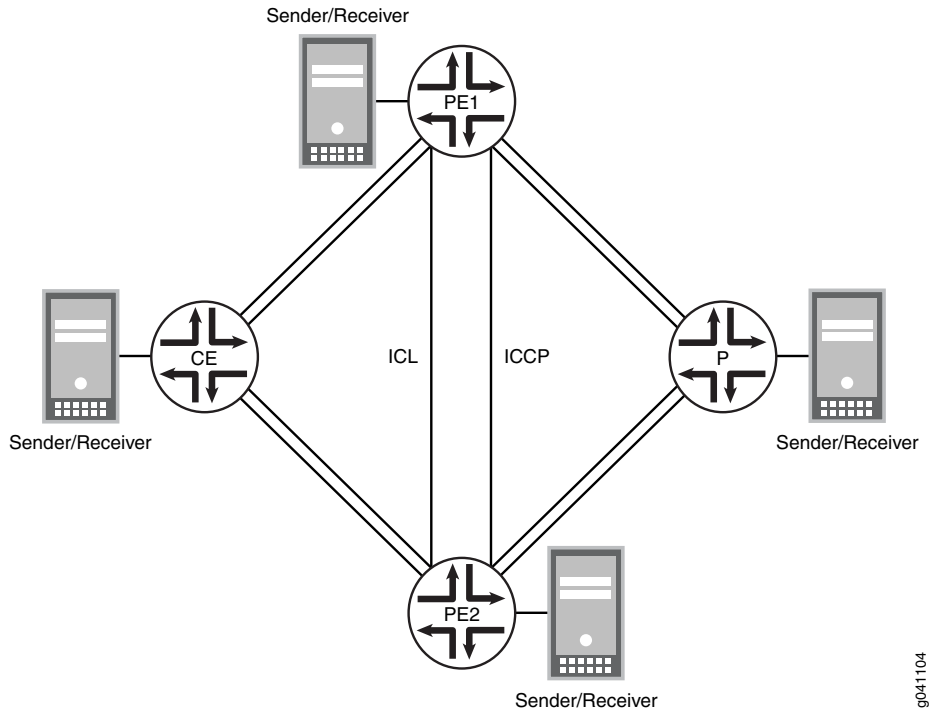

Configuring the PE Routers

CLI Quick Configuration

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

Router PE1

```
set chassis aggregated-devices ethernet device-count 5
set interfaces ge-1/0/1 gigether-options 802.3ad ae1
set interfaces ge-1/0/2 unit 0 family inet address 10.100.100.1/30
set interfaces ge-1/0/6 gigether-options 802.3ad ae0
set interfaces ge-1/1/1 flexible-vlan-tagging
set interfaces ge-1/1/1 encapsulation flexible-ethernet-services
set interfaces ge-1/1/1 unit 0 encapsulation vlan-bridge
set interfaces ge-1/1/1 unit 0 vlan-id-range 100-110
set interfaces ge-1/1/4 flexible-vlan-tagging
set interfaces ge-1/1/4 encapsulation flexible-ethernet-services
set interfaces ge-1/1/4 unit 0 encapsulation vlan-bridge
set interfaces ge-1/1/4 unit 0 vlan-id-range 100-110
set interfaces ae0 flexible-vlan-tagging
set interfaces ae0 encapsulation flexible-ethernet-services
set interfaces ae0 aggregated-ether-options lacp active
set interfaces ae0 aggregated-ether-options lacp system-priority 100
set interfaces ae0 aggregated-ether-options lacp system-id 00:00:00:00:00:05
set interfaces ae0 aggregated-ether-options lacp admin-key 1
set interfaces ae0 aggregated-ether-options mc-ae mc-ae-id 5
set interfaces ae0 aggregated-ether-options mc-ae redundancy-group 10
```

```
set interfaces ae0 aggregated-ether-options mc-ae chassis-id 1
set interfaces ae0 aggregated-ether-options mc-ae mode active-active
set interfaces ae0 aggregated-ether-options mc-ae status-control active
set interfaces ae0 unit 0 encapsulation vlan-bridge
set interfaces ae0 unit 0 vlan-id-range 100-110
set interfaces ae0 unit 0 multi-chassis-protection 10.100.100.2 interface ge-1/1/4.0
set interfaces ae1 flexible-vlan-tagging
set interfaces ae1 encapsulation flexible-ethernet-services
set interfaces ae1 aggregated-ether-options lacp active
set interfaces ae1 aggregated-ether-options lacp system-priority 100
set interfaces ae1 aggregated-ether-options lacp system-id 00:00:00:00:00:05
set interfaces ae1 aggregated-ether-options lacp admin-key 1
set interfaces ae1 aggregated-ether-options mc-ae mc-ae-id 10
set interfaces ae1 aggregated-ether-options mc-ae redundancy-group 10
set interfaces ae1 aggregated-ether-options mc-ae chassis-id 1
set interfaces ae1 aggregated-ether-options mc-ae mode active-active
set interfaces ae1 aggregated-ether-options mc-ae status-control active
set interfaces ae1 unit 0 encapsulation vlan-bridge
set interfaces ae1 unit 0 vlan-id-range 100-110
set interfaces ae1 unit 0 multi-chassis-protection 10.100.100.2 interface ge-1/1/4.0
set bridge-domains bd0 domain-type bridge
set bridge-domains bd0 vlan-id all
set bridge-domains bd0 service-id 20
set bridge-domains bd0 interface ae1.0
set bridge-domains bd0 interface ge-1/0/3.0
set bridge-domains bd0 interface ge-1/1/1.0
set bridge-domains bd0 interface ge-1/1/4.0
set bridge-domains bd0 interface ae0.0
set bridge-domains bd0 multicast-snooping-options multichassis-lag-replicate-state
set bridge-domains bd0 protocols igmp-snooping vlan 100 interface ge-1/1/4.0
  multicast-router-interface
set bridge-domains bd0 protocols igmp-snooping vlan 101 interface ge-1/1/4.0
  multicast-router-interface
set bridge-domains bd0 protocols igmp-snooping vlan 200 interface ge-1/1/4.0
  multicast-router-interface
set multicast-snooping-options multichassis-lag-replicate-state
set protocols iccp local-ip-addr 10.100.100.1
set protocols iccp peer 10.100.100.2 redundancy-group-id-list 10
set protocols iccp peer 10.100.100.2 liveness-detection minimum-interval 1000
set switch-options service-id 10
```

Router PE2

```
set chassis aggregated-devices ethernet device-count 5
set interfaces ge-1/0/2 unit 0 family inet address 10.100.100.2/30
set interfaces ge-1/0/3 flexible-vlan-tagging
set interfaces ge-1/0/3 encapsulation flexible-ethernet-services
set interfaces ge-1/0/3 unit 0 encapsulation vlan-bridge
set interfaces ge-1/0/3 unit 0 vlan-id-range 100-110
set interfaces ge-1/0/4 flexible-vlan-tagging
set interfaces ge-1/0/4 encapsulation flexible-ethernet-services
set interfaces ge-1/0/4 unit 0 encapsulation vlan-bridge
set interfaces ge-1/0/4 unit 0 vlan-id-range 100-110
set interfaces ge-1/0/5 gigether-options 802.3ad ae0
set interfaces ge-1/1/0 gigether-options 802.3ad ae1
```

```
set interfaces ae0 flexible-vlan-tagging
set interfaces ae0 encapsulation flexible-ethernet-services
set interfaces ae0 aggregated-ether-options lacp active
set interfaces ae0 aggregated-ether-options lacp system-priority 100
set interfaces ae0 aggregated-ether-options lacp system-id 00:00:00:00:00:05
set interfaces ae0 aggregated-ether-options lacp admin-key 1
set interfaces ae0 aggregated-ether-options mc-ae mc-ae-id 5
set interfaces ae0 aggregated-ether-options mc-ae redundancy-group 10
set interfaces ae0 aggregated-ether-options mc-ae chassis-id 0
set interfaces ae0 aggregated-ether-options mc-ae mode active-active
set interfaces ae0 aggregated-ether-options mc-ae status-control active
set interfaces ae0 unit 0 encapsulation vlan-bridge
set interfaces ae0 unit 0 vlan-id-range 100-110
set interfaces ae0 unit 0 multi-chassis-protection 10.100.100.1 interface ge-1/0/4.0
set interfaces ae1 flexible-vlan-tagging
set interfaces ae1 encapsulation flexible-ethernet-services
set interfaces ae1 aggregated-ether-options lacp active
set interfaces ae1 aggregated-ether-options lacp system-priority 100
set interfaces ae1 aggregated-ether-options lacp system-id 00:00:00:00:00:05
set interfaces ae1 aggregated-ether-options lacp admin-key 1
set interfaces ae1 aggregated-ether-options mc-ae mc-ae-id 10
set interfaces ae1 aggregated-ether-options mc-ae redundancy-group 10
set interfaces ae1 aggregated-ether-options mc-ae chassis-id 0
set interfaces ae1 aggregated-ether-options mc-ae mode active-active
set interfaces ae1 aggregated-ether-options mc-ae status-control active
set interfaces ae1 unit 0 encapsulation vlan-bridge
set interfaces ae1 unit 0 vlan-id-range 100-110
set interfaces ae1 unit 0 multi-chassis-protection 10.100.100.1 interface ge-1/0/4.0
set bridge-domains bd0 domain-type bridge
set bridge-domains bd0 vlan-id all
set bridge-domains bd0 service-id 20
set bridge-domains bd0 interface ae1.0
set bridge-domains bd0 interface ge-1/0/3.0
set bridge-domains bd0 interface ge-1/0/4.0
set bridge-domains bd0 interface ae0.0
set bridge-domains bd0 multicast-snooping-options multichassis-lag-replicate-state
set bridge-domains bd0 protocols igmp-snooping vlan 100 interface ge-1/0/4.0
    multicast-router-interface
set bridge-domains bd0 protocols igmp-snooping vlan 101 interface ge-1/0/4.0
    multicast-router-interface
set bridge-domains bd0 protocols igmp-snooping vlan 200 interface ge-1/0/4.0
    multicast-router-interface
set multicast-snooping-options multichassis-lag-replicate-state
set protocols iccp local-ip-addr 10.100.100.2
set protocols iccp peer 10.100.100.1 redundancy-group-id-list 10
set protocols iccp peer 10.100.100.1 liveness-detection minimum-interval 1000
set switch-options service-id 10
```

Configuring the PE1 Router

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 Router PE1:

1. Specify the number of aggregated Ethernet interfaces to be created.

```
[edit chassis]
user@PE1# set aggregated-devices ethernet device-count 5
```

2. Specify the members to be included within the aggregated Ethernet bundles.

```
[edit interfaces]
user@PE1# set ge-1/0/1 gigether-options 802.3ad ae1
user@PE1# set ge-1/0/6 gigether-options 802.3ad ae0
```

3. Configure the interfaces that connect to multicast senders or receivers, the ICL interfaces, and the ICCP interfaces.

```
[edit interfaces]
user@PE1# set ge-1/1/1 flexible-vlan-tagging
user@PE1# set ge-1/1/1 encapsulation flexible-ethernet-services
user@PE1# set ge-1/1/1 unit 0 encapsulation vlan-bridge
user@PE1# set ge-1/1/1 unit 0 vlan-id-range 100-110

user@PE1# set ge-1/1/4 flexible-vlan-tagging
user@PE1# set ge-1/1/4 encapsulation flexible-ethernet-services
user@PE1# set ge-1/1/4 unit 0 encapsulation vlan-bridge
user@PE1# set ge-1/1/4 unit 0 vlan-id-range 100-110

user@PE1# set ge-1/0/2 unit 0 family inet address 10.100.100.1/30
```

4. Configure parameters on the aggregated Ethernet bundles.

```
[edit interfaces ae0]
user@PE1# set flexible-vlan-tagging
user@PE1# set encapsulation flexible-ethernet-services
user@PE1# set unit 0 encapsulation vlan-bridge
user@PE1# set unit 0 vlan-id-range 100-110
user@PE1# set unit 0 multi-chassis-protection 10.100.100.2 interface ge-1/1/4.0
[edit interfaces ae1]
user@PE1# set flexible-vlan-tagging
user@PE1# set encapsulation flexible-ethernet-services
user@PE1# set unit 0 encapsulation vlan-bridge
user@PE1# set unit 0 vlan-id-range 100-110
user@PE1# set unit 0 multi-chassis-protection 10.100.100.2 interface ge-1/1/4.0
```

5. Configure LACP on the aggregated Ethernet bundles.

```
[edit interfaces ae0 aggregated-ether-options]
user@PE1# set lacp active
user@PE1# set lacp system-priority 100
user@PE1# set lacp system-id 00:00:00:00:00:05
user@PE1# set lacp admin-key 1
[edit interfaces ae1 aggregated-ether-options]
user@PE1# set lacp active
user@PE1# set lacp system-priority 100
user@PE1# set lacp system-id 00:00:00:00:00:05
user@PE1# set lacp admin-key 1
```

6. Configure the MC-LAG interfaces.

```
[edit interfaces ae0 aggregated-ether-options]
user@PE1# set mc-ae mc-ae-id 5
user@PE1# set mc-ae redundancy-group 10
user@PE1# set mc-ae chassis-id 1
user@PE1# set mc-ae mode active-active
user@PE1# set mc-ae status-control active
[edit interfaces ae1 aggregated-ether-options]
user@PE1# set mc-ae mc-ae-id 10
user@PE1# set mc-ae redundancy-group 10
user@PE1# set mc-ae chassis-id 1
user@PE1# set mc-ae mode active-active
user@PE1# set mc-ae status-control active
```

The multichassis aggregated Ethernet identification number (**mc-ae-id**) specifies which link aggregation group the aggregated Ethernet interface belongs to. The ae0 interfaces on Router PE1 and Router PE2 are configured with **mc-ae-id 5**. The ae1 interfaces on Router PE1 and Router PE2 are configured with **mc-ae-id 10**.

The **redundancy-group 10** statement is used by ICCP to associate multiple chassis that perform similar redundancy functions and to establish a communication channel so that applications on peering chassis can send messages to each other. The ae0 and ae1 interfaces on Router PE1 and Router PE2 are configured with the same redundancy group, **redundancy-group 10**.

The **chassis-id** statement is used by LACP for calculating the port number of the MC-LAG's physical member links. Router PE1 uses **chassid-id 1** to identify both its ae0 and ae1 interfaces. Router PE2 uses **chassis-id 0** to identify both its ae0 and ae1 interfaces.

The **mode** statement indicates whether an MC-LAG is in active-standby mode or active-active mode. Chassis that are in the same group must be in the same mode.

7. Configure a domain that includes the set of logical ports.

```
[edit bridge-domains bd0]
user@PE1# set domain-type bridge
user@PE1# set vlan-id all
```

```
user@PE1# set service-id 20
user@PE1# set interface ae0.0
user@PE1# set interface ae1.0
user@PE1# set interface ge-1/0/3.0
user@PE1# set interface ge-1/1/1.0
user@PE1# set interface ge-1/1/4.0
```

The ports within a bridge domain share the same flooding or broadcast characteristics in order to perform Layer 2 bridging.

The bridge-level **service-id** statement is required to link related bridge domains across peers (in this case Router PE1 and Router PE2), and should be configured with the same value.

8. At the global level and also in the bridge domain, replicate IGMP join and leave messages from the active link to the standby link of a dual-link MC-LAG interface, to enable faster recovery of membership information after failover.

```
[edit multicast-snooping-options]
user@PE1# set multichassis-lag-replicate-state
[edit bridge-domains bd0 multicast-snooping-options]
user@PE1# set multichassis-lag-replicate-state
```

9. (Optional) Suppress MC-LAG reports to optimize the syncing of the ICCP messages. By default, every IGMP packet received on the MC-AE interface is replicated to the peer. If multiple hosts behind the CE router send reports for the same group, all the packets are synced even though only a single report is used for building the IGMP snooping state on the peer. Also all subsequent refreshes sent in response to the IGMP queries are also synced to this peer. This requires significant CPU cycles on both peers which send and receive these reports over ICCP. Starting with Junos OS Release 16.1, you can configure the **suppress-report** statement at the [**edit multicast-snooping-options multichassis-lag-replicate-state**] hierarchy level to optimize the syncing of the ICCP messages.

```
[edit multicast-snooping-options]
user@PE1# set multichassis-lag-replicate-state suppress-report
```

Optimizing the syncing of ICCP messages ensures that the message exchanges using ICCP between the peers is more efficient. This also improves scaling by ensuring that the membership state is present only at the receiving PE.

NOTE:

- Because the IGMP reports/leaves sent between the MC-LAG peers are suppressed, IGMP snooping statistics will not be the same on both peers. Total statistics will be the sum of the IGMP reports received on both MCLAG peers.
- When MC-LAG reports are suppressed, the MCSNOOPD client application will not receive the source IP address (host information).

10. Configure multicast snooping for the MC-LAG interfaces.

```
[edit bridge-domains bd0]
user@PE1# set protocols igmp-snooping vlan 100 interface ge-1/1/4.0
  multicast-router-interface
user@PE1# set protocols igmp-snooping vlan 101 interface ge-1/1/4.0
  multicast-router-interface
user@PE1# set protocols igmp-snooping vlan 200 interface ge-1/1/4.0
  multicast-router-interface
```

NOTE: Starting with Junos OS Release 16.1, you can selectively add ICL to preserve ICL bandwidth. To do this, you must *not* configure the ICL as an **multicast-router-interface** as specified in this step. Instead, you must configure the **enhanced-ip** statement.

When you configure to selectively add ICL, control packets are directly sent from PFE to RPD. Therefore, if an IRB interface is attached to a bridge-domain, the proxy functionality in the L2 domain will not be effective because MCSNOOPD only proxies to external routers connected to the physical interfaces. In such scenarios, you can enable proxy to IRB. To do this, configure the **irb** statement at the [**edit protocols igmp-snooping proxy**] hierarchy level.

11. Configure ICCP parameters.

```
[edit protocols iccp]
user@PE1# set local-ip-addr 10.100.100.1
user@PE1# set peer 10.100.100.2 redundancy-group-id-list 10
user@PE1# set peer 10.100.100.2 liveness-detection minimum-interval 1000
```

12. Configure the service ID at the global level.

```
[edit switch-options]
user@PE1# set service-id 10
```

You must configure the same unique network-wide configuration for a service in the set of PE routers providing the service. This service ID is required if the multichassis aggregated Ethernet interfaces are part of a bridge domain.

Results

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

```
user@PE1# show bridge-domains
bd0 {
    domain-type bridge;
    vlan-id all;
    service-id 20;
    interface ae1.0;
    interface ge-1/0/3.0;
    interface ge-1/1/1.0;
    interface ge-1/1/4.0;
    interface ae0.0;
    multicast-snooping-options {
        multichassis-lag-replicate-state;
    }
    protocols {
        igmp-snooping {
            vlan 100 {
                interface ge-1/1/4.0 {
                    multicast-router-interface;
                }
            }
            vlan 101 {
                interface ge-1/1/4.0 {
                    multicast-router-interface;
                }
            }
            vlan 200 {
                interface ge-1/1/4.0 {
                    multicast-router-interface;
                }
            }
        }
    }
}
```

```
user@PE1# show chassis
aggregated-devices {
    ethernet {
        device-count 5;
    }
}
```

```
user@PE1# show interfaces
ge-1/0/1 {
    gigether-options {
        802.3ad ae1;
    }
}
ge-1/0/6 {
    gigether-options {
        802.3ad ae0;
    }
}
ge-1/0/2 {
    unit 0 {
```

```
      family inet {
        address 10.100.100.1/30;
      }
    }
}
ge-1/1/1 {
    flexible-vlan-tagging;
    encapsulation flexible-ethernet-services;
    unit 0 {
      encapsulation vlan-bridge;
      vlan-id-range 100-110;
    }
}
ge-1/1/4 {
    flexible-vlan-tagging;
    encapsulation flexible-ethernet-services;
    unit 0 {
      encapsulation vlan-bridge;
      vlan-id-range 100-110;
    }
}
ae0 {
    flexible-vlan-tagging;
    encapsulation flexible-ethernet-services;
    aggregated-ether-options {
      lacp {
        active;
        system-priority 100;
        system-id 00:00:00:00:00:05;
        admin-key 1;
      }
      mc-ae {
        mc-ae-id 5;
        redundancy-group 10;
        chassis-id 1;
        mode active-active;
        status-control active;
      }
    }
    unit 0 {
      encapsulation vlan-bridge;
      vlan-id-range 100-110;
      multi-chassis-protection 10.100.100.2 {
        interface ge-1/1/4.0;
      }
    }
}
ae1 {
    flexible-vlan-tagging;
    encapsulation flexible-ethernet-services;
    aggregated-ether-options {
      lacp {
        active;
        system-priority 100;
        system-id 00:00:00:00:00:05;
```

```
        admin-key 1;
      }
      mc-ae {
        mc-ae-id 10;
        redundancy-group 10;
        chassis-id 1;
        mode active-active;
        status-control active;
      }
    }
    unit 0 {
      encapsulation vlan-bridge;
      vlan-id-range 100-110;
      multi-chassis-protection 10.100.100.2 {
        interface ge-1/1/4.0;
      }
    }
}
```

```
user@PE1#  show multicast-snooping-options
multichassis-lag-replicate-state;
```

```
user@PE1#  show protocols
iccp {
  local-ip-addr 10.100.100.1;
  peer 10.100.100.2 {
    redundancy-group-id-list 10;
    liveness-detection {
      minimum-interval 1000;
    }
  }
}
```

```
user@PE1# run show switch-options
service-id 10;
```

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

Repeat the procedure for Router PE2, using the appropriate interface names and addresses.

Configuring the CE Device

CLI Quick Configuration

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

Device CE

```
set chassis aggregated-devices ethernet device-count 2
set interfaces ge-2/0/2 gigether-options 802.3ad ae0
set interfaces ge-2/0/3 gigether-options 802.3ad ae0
```

```
set interfaces ge-2/1/6 flexible-vlan-tagging
set interfaces ge-2/1/6 encapsulation flexible-ethernet-services
set interfaces ge-2/1/6 unit 0 encapsulation vlan-bridge
set interfaces ge-2/1/6 unit 0 vlan-id-range 100-110
set interfaces ae0 flexible-vlan-tagging
set interfaces ae0 encapsulation flexible-ethernet-services
set interfaces ae0 aggregated-ether-options lacp active
set interfaces ae0 aggregated-ether-options lacp system-priority 100
set interfaces ae0 unit 0 encapsulation vlan-bridge
set interfaces ae0 unit 0 vlan-id-range 100-500
set bridge-domains bd0 domain-type bridge
set bridge-domains bd0 vlan-id all
set bridge-domains bd0 interface ge-2/1/6.0
set bridge-domains bd0 interface ae0.0
```

Configuring the CE Device

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 the CE device:

1. Specify the number of aggregated Ethernet interfaces to be created.

```
[edit chassis]
user@CE# set aggregated-devices ethernet device-count 2
```

2. Specify the members to be included within the aggregated Ethernet bundle.

```
[edit interfaces]
user@CE# set ge-2/0/2 gigether-options 802.3ad ae0
user@CE# set ge-2/0/3 gigether-options 802.3ad ae0
```

3. Configure an interface that connects to multicast senders or receivers.

```
[edit interfaces ge-2/1/6]
user@CE# set flexible-vlan-tagging
user@CE# set encapsulation flexible-ethernet-services
user@CE# set unit 0 encapsulation vlan-bridge
user@CE# set unit 0 vlan-id-range 100-110
```

4. Configure parameters on the aggregated Ethernet bundle.

```
[edit interfaces ae0]
user@CE# set flexible-vlan-tagging
user@CE# set encapsulation flexible-ethernet-services
user@CE# set unit 0 encapsulation vlan-bridge
user@CE# set unit 0 vlan-id-range 100-500
```

5. Configure LACP on the aggregated Ethernet bundle.

```
[edit interfaces ae0 aggregated-ether-options]
user@CE# set lacp active
user@CE# set lacp system-priority 100
```

The **active** statement initiates transmission of LACP packets.

For the **system-priority** statement, a smaller value indicates a higher priority. The device with the lower system priority value determines which links between LACP partner devices are active and which are in standby mode for each LACP group. The device on the controlling end of the link uses port priorities to determine which ports are bundled into the aggregated bundle and which ports are put in standby mode. Port priorities on the other device (the noncontrolling end of the link) are ignored.

6. Configure a domain that includes the set of logical ports.

```
[edit bridge-domains bd0]
user@CE# set domain-type bridge
user@CE# set vlan-id all
user@CE# set interface ge-2/1/6.0
user@CE# set interface ae0.0
```

The ports within a bridge domain share the same flooding or broadcast characteristics in order to perform Layer 2 bridging.

Results

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

```
user@CE#  show bridge-domains
bd0 {
   domain-type bridge;
   vlan-id all;
   interface ge-2/1/6.0;
   interface ae0.0;
}
```

```
user@CE#  show chassis
aggregated-devices {
   ethernet {
      device-count 2;
   }
}
```

```
user@CE#  show interfaces
ge-2/0/2 {
   gigether-options {
```

```
      802.3ad ae0;
    }
}
ge-2/0/3 {
    gigether-options {
      802.3ad ae0;
    }
}
ge-2/1/6 {
    flexible-vlan-tagging;
    encapsulation flexible-ethernet-services;
    unit 0 {
      encapsulation vlan-bridge;
      vlan-id-range 100-110;
    }
}
ae0 {
    flexible-vlan-tagging;
    encapsulation flexible-ethernet-services;
    aggregated-ether-options {
      lacp {
        active;
        system-priority 100;
      }
    }
    unit 0 {
      encapsulation vlan-bridge;
      vlan-id-range 100-500;
    }
}
```

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

Configuring the Provider Router

CLI Quick Configuration

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

Router P

```
set chassis aggregated-devices ethernet device-count 2
set interfaces ge-1/0/5 gigether-options 802.3ad ae1
set interfaces ge-1/0/11 gigether-options 802.3ad ae1
set interfaces ge-1/1/3 flexible-vlan-tagging
set interfaces ge-1/1/3 encapsulation flexible-ethernet-services
set interfaces ge-1/1/3 unit 0 encapsulation vlan-bridge
set interfaces ge-1/1/3 unit 0 vlan-id-range 100-500
set interfaces ae1 flexible-vlan-tagging
set interfaces ae1 encapsulation flexible-ethernet-services
set interfaces ae1 aggregated-ether-options lacp active
set interfaces ae1 aggregated-ether-options lacp system-priority 100
set interfaces ae1 unit 0 encapsulation vlan-bridge
set interfaces ae1 unit 0 vlan-id-range 100-110
```

```
set bridge-domains bd0 vlan-id all
set bridge-domains bd0 domain-type bridge
set bridge-domains bd0 interface ge-1/1/3.0
set bridge-domains bd0 interface ae1.0
```

Configuring the P Router

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 Router P:

1. Specify the number of aggregated Ethernet interfaces to be created.

```
[edit chassis]
user@P# set aggregated-devices ethernet device-count 2
```

2. Specify the members to be included within the aggregated Ethernet bundle.

```
[edit interfaces]
user@P# set ge-1/0/5 gigether-options 802.3ad ae1
user@P# set ge-1/0/11 gigether-options 802.3ad ae1
```

3. Configure an interface that connects to multicast senders or receivers.

```
[edit interfaces ge-1/1/3]
user@P# set flexible-vlan-tagging
user@P# set encapsulation flexible-ethernet-services
user@P# set unit 0 encapsulation vlan-bridge
user@P# set unit 0 vlan-id-range 100-500
```

4. Configure parameters on the aggregated Ethernet bundle.

```
[edit interfaces ae1]
user@P# set flexible-vlan-tagging
user@P# set encapsulation flexible-ethernet-services
user@P# set unit 0 encapsulation vlan-bridge
user@P# set unit 0 vlan-id-range 100-110
```

5. Configure LACP on the aggregated Ethernet bundle.

```
[edit interfaces ae1 aggregated-ether-options]
user@P# set lacp active
user@P# set lacp system-priority 100
```

6. Configure a domain that includes the set of logical ports.

```
[edit bridge-domains bd0]
user@P# set vlan-id all
user@P# set domain-type bridge
user@P# set interface ge-1/1/3.0
user@P# set interface ae1.0
```

Results

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

```
user@P#  show bridge-domains
bd0 {
   domain-type bridge;
   vlan-id all;
   interface ge-1/1/3.0;
   interface ae1.0;
}
```

```
user@P#  show chassis
aggregated-devices {
   ethernet {
      device-count 2;
   }
}
```

```
user@P#  show interfaces
ge-1/0/5 {
   gigether-options {
      802.3ad ae1;
   }
}
ge-1/0/11 {
   gigether-options {
      802.3ad ae1;
   }
}
ge-1/1/3 {
   flexible-vlan-tagging;
   encapsulation flexible-ethernet-services;
   unit 0 {
      encapsulation vlan-bridge;
      vlan-id-range 100-500;
   }
}
ae1 {
   flexible-vlan-tagging;
   encapsulation flexible-ethernet-services;
   aggregated-ether-options {
```

```
      lacp {
        active;
        system-priority 100;
      }
    }
    unit 0 {
      encapsulation vlan-bridge;
      vlan-id-range 100-110;
    }
}
```

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

Verification

Verify that the configuration is working properly by running the following commands:

- **show iccp**
- ***show igmp snooping interface***
- ***show igmp snooping membership***
- **show interfaces ae0**
- **show interfaces ae1**
- **show interfaces mc-ae**
- ***show l2-learning instance*** **extensive**
- ***show multicast snooping route*** **extensive**

Example: Configuring Multichassis Link Aggregation for Layer 3 Multicast Using VRRP on EX9200 Switches

There are two methods for enabling Layer 3 multicast functionality across a multichassis link aggregation group (MC-LAG). You can choose either to configure Virtual Router Redundancy Protocol (VRRP) over the integrated routing and bridging (IRB) interface or to synchronize the MAC addresses for the Layer 3 interfaces of the switches participating in the MC-LAG. This provides redundancy and load balancing between the two MC-LAG peers. The procedure to configure VRRP for use in a Layer 3 multicast MC-LAG is included in this example.

Requirements

This example uses the following hardware and software components:

- Two EX9200 switches

- Junos OS Release 13.2R1 or later

Before you configure an MC-LAG for Layer 3 multicast using VRRP, be sure that you understand how to:

- Configure aggregated Ethernet interfaces on a switch. See *Configuring an Aggregated Ethernet Interface*.
- Configure Link Aggregation Control Protocol (LACP) on aggregated Ethernet interfaces on a switch. See *Configuring Aggregated Ethernet LACP*.
- Configure Virtual Router Redundancy Protocol (VRRP) on a switch. See *Configuring Basic VRRP Support*.

Overview

In this example, you configure two MC-LAGs between two switches, consisting of two aggregated Ethernet interfaces (ae1 and ae2). To support the MC-LAG, you create a third aggregated Ethernet interface (ae0) for the interchassis link (ICL). You also configure a multichassis protection link for the ICL, Inter-Chassis Control Protocol (ICCP) for the peers hosting the MC-LAG, and Layer 3 connectivity between MC-LAG peers.

NOTE: **Layer 3 connectivity is required for ICCP.**

To complete the MC-LAG configuration, enable VRRP by completing the following tasks for each MC-LAG:

1. Create an IRB interface.
2. Create a VRRP group and assign a virtual IP address that is shared between each switch in the VRRP group.
3. Enable a member of a VRRP group to accept all packets destined for the virtual IP address if it is the master in the VRRP group.
4. Configure Layer 3 connectivity between the VRRP groups.

Topology

The topology used in this example consists of two switches hosting two MC-LAGs—ae1 and ae2. The two switches are connected to a multicast source (Server 1) over MC-LAG ae1, and a multicast receiver (Server 2) over MC-LAG ae2. Figure 31 on page 209 shows the topology for this example.

Figure 31: Configuring Two MC-LAGs Between Switch A and Switch B

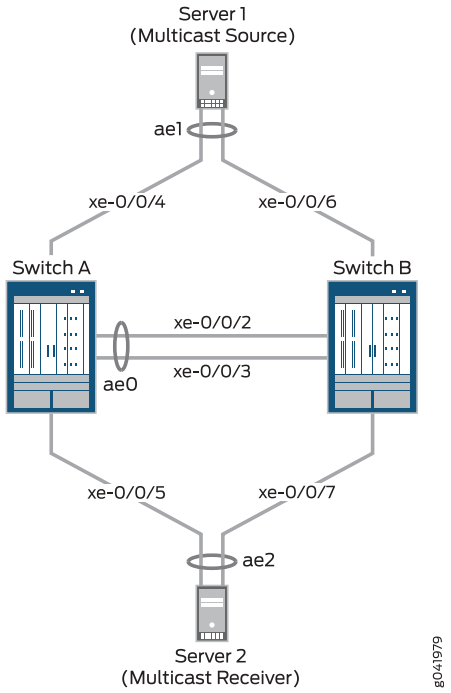

Table 8 on page 209 details the topology used in this configuration example.

Table 8: Components of the Topology for Configuring Two MC-LAGs Between Switch A and Switch B

Hostname	Base Hardware	Link Aggregation Groups
Switch A Switch B	EX9200 switch	• ae0 is configured as an aggregated Ethernet interface, and is used as an ICL. The following two interfaces are part of ae0: xe-0/0/2 and xe-0/0/3 on Switch A and xe-0/0/2 and xe-0/0/3 on Switch B. • ae1 is configured as an MC-LAG for the multicast source (Server 1), and the following two interfaces are part of ae1: xe-0/0/4 on Switch A and xe-0/0/6 on Switch B. • ae2 is configured as an MC-LAG for the multicast receiver (Server 2), and the following two interfaces are part of ae2: xe-0/0/5 on Switch A and xe-0/0/7 on Switch B.

Configuration

CLI Quick Configuration

To quickly configure this example, copy the following commands, paste them in a text file, remove any line breaks, change any details necessary to match your network

configuration, copy and paste the commands into the CLI at the **[edit]** hierarchy level, and then enter **commit** from configuration mode.

Switch A

```
set chassis aggregated-devices ethernet device-count 3
set interfaces xe-0/0/2 ether-options 802.3ad ae0
set interfaces xe-0/0/3 ether-options 802.3ad ae0
set interfaces xe-0/0/4 ether-options 802.3ad ae1
set interfaces xe-0/0/5 ether-options 802.3ad ae2
set interfaces ae0 unit 0 family ethernet-switching interface-mode trunk
set interfaces ae0 unit 0 family ethernet-switching vlan members v500
set interfaces ae1 aggregated-ether-options lacp active
set interfaces ae1 aggregated-ether-options lacp system-id 00:01:02:03:04:05
set interfaces ae1 aggregated-ether-options lacp admin-key 3
set interfaces ae1 aggregated-ether-options mc-ae mc-ae-id 3
set interfaces ae1 aggregated-ether-options mc-ae chassis-id 0
set interfaces ae1 aggregated-ether-options mc-ae mode active-active
set interfaces ae1 aggregated-ether-options mc-ae status-control active
set interfaces ae1 aggregated-ether-options mc-ae init-delay-time 240
set interfaces ae1 unit 0 family ethernet-switching interface-mode trunk
set interfaces ae1 unit 0 family ethernet-switching vlan members v100
set interfaces ae2 aggregated-ether-options lacp active
set interfaces ae2 aggregated-ether-options lacp system-id 00:01:02:03:04:06
set interfaces ae2 aggregated-ether-options lacp admin-key 3
set interfaces ae2 aggregated-ether-options mc-ae mc-ae-id 4
set interfaces ae2 aggregated-ether-options mc-ae chassis-id 0
set interfaces ae2 aggregated-ether-options mc-ae mode active-active
set interfaces ae2 aggregated-ether-options mc-ae status-control active
set interfaces ae2 aggregated-ether-options mc-ae init-delay-time 240
set interfaces ae2 unit 0 family ethernet-switching interface-mode trunk
set interfaces ae2 unit 0 family ethernet-switching vlan members v200
set interfaces irb unit 100 family inet address 10.1.1.11/8 vrrp-group 1 virtual-address 10.1.1.1
set interfaces irb unit 100 family inet address 10.1.1.11/8 vrrp-group 1 priority 200
set interfaces irb unit 100 family inet address 10.1.1.11/8 vrrp-group 1 accept-data
set interfaces irb unit 200 family inet address 10.1.1.21/8 vrrp-group 2 virtual-address
  10.1.1.2
set interfaces irb unit 200 family inet address 10.1.1.21/8 vrrp-group 2 priority 200
set interfaces irb unit 200 family inet address 10.1.1.21/8 vrrp-group 2 accept-data
set interfaces irb unit 500 family inet address 10.3.3.2/8
set vlans v100 vlan-id 100
set vlans v100 l3-interface irb.100
set vlans v200 vlan-id 200
set vlans v200 l3-interface irb.200
set vlans v500 vlan-id 500
set vlans v500 l3-interface irb.500
set protocols iccp local-ip-addr 10.3.3.2
set protocols iccp peer 10.3.3.1 session-establishment-hold-time 340
set protocols iccp peer 10.3.3.1 backup-liveness-detection backup-peer-ip 10.207.64.233
set protocols iccp peer 10.3.3.1 liveness-detection minimum-receive-interval 60
set protocols iccp peer 10.3.3.1 liveness-detection transmit-interval minimum-interval
  60
set protocols igmp-snooping vlan v100
set protocols igmp-snooping vlan v200
set protocols igmp-snooping vlan v500
```

```
set protocols ospf area 0.0.0.0 interface irb.100 bfd-liveness-detection
  minimum-receive-interval 700
set protocols ospf area 0.0.0.0 interface irb.100 bfd-liveness-detection transmit-interval
  minimum-interval 350
set protocols ospf area 0.0.0.0 interface irb.100 bfd-liveness-detection transmit-interval
  threshold 500
set protocols ospf area 0.0.0.0 interface irb.200 bfd-liveness-detection
  minimum-receive-interval 700
set protocols ospf area 0.0.0.0 interface irb.200 bfd-liveness-detection transmit-interval
  minimum-interval 350
set protocols ospf area 0.0.0.0 interface irb.200 bfd-liveness-detection transmit-interval
  threshold 500
set protocols pim rp static address 10.0.0.3 group-ranges 233.252.0.0/8
set protocols pim interface irb.100 priority 200
set protocols pim interface irb.200 priority 600
set multi-chassis multi-chassis-protection 10.3.3.1 interface ae0
```

Switch B

```
set chassis aggregated-devices ethernet device-count 2
set interfaces xe-0/0/2 ether-options 802.3ad ae0
set interfaces xe-0/0/3 ether-options 802.3ad ae0
set interfaces xe-0/0/6 ether-options 802.3ad ae1
set interfaces xe-0/0/7 ether-options 802.3ad ae2
set interfaces ae0 unit 0 family ethernet-switching interface-mode trunk
set interfaces ae0 unit 0 family ethernet-switching vlan members v500
set interfaces ae1 aggregated-ether-options lacp active
set interfaces ae1 aggregated-ether-options lacp system-id 00:01:02:03:04:05
set interfaces ae1 aggregated-ether-options lacp admin-key 3
set interfaces ae1 aggregated-ether-options mc-ae mc-ae-id 3
set interfaces ae1 aggregated-ether-options mc-ae chassis-id 1
set interfaces ae1 aggregated-ether-options mc-ae mode active-active
set interfaces ae1 aggregated-ether-options mc-ae status-control standby
set interfaces ae1 aggregated-ether-options mc-ae init-delay-time 240
set interfaces ae1 unit 0 family ethernet-switching interface-mode trunk
set interfaces ae1 unit 0 family ethernet-switching vlan members v100
set interfaces ae2 aggregated-ether-options lacp active
set interfaces ae2 aggregated-ether-options lacp system-id 00:01:02:03:04:06
set interfaces ae2 aggregated-ether-options lacp admin-key 3
set interfaces ae2 aggregated-ether-options mc-ae mc-ae-id 4
set interfaces ae2 aggregated-ether-options mc-ae chassis-id 1
set interfaces ae2 aggregated-ether-options mc-ae mode active-active
set interfaces ae2 aggregated-ether-options mc-ae status-control standby
set interfaces ae2 aggregated-ether-options mc-ae init-delay-time 240
set interfaces ae2 unit 0 family ethernet-switching interface-mode trunk
set interfaces ae2 unit 0 family ethernet-switching vlan members v200
set interfaces irb unit 100 family inet address 10.1.1.10/8 vrrp-group 1 virtual-address 10.1.1.1
set interfaces irb unit 100 family inet address 10.1.1.10/8 vrrp-group 1 priority 150
set interfaces irb unit 100 family inet address 10.1.1.10/8 vrrp-group 1 accept-data
set interfaces irb unit 200 family inet address 10.1.1.20/8 vrrp-group 2 virtual-address
  10.1.1.2
set interfaces irb unit 200 family inet address 10.1.1.20/8 vrrp-group 2 priority 150
set interfaces irb unit 200 family inet address 10.1.1.20/8 vrrp-group 2 accept-data
set interfaces irb unit 500 family inet address 10.3.3.1/8
set vlans v100 vlan-id 100
```

```
set vlans v100 l3-interface irb.100
set vlans v200 vlan-id 200
set vlans v200 l3-interface irb.200
set vlans v500 vlan-id 500
set vlans v500 l3-interface irb.500
set protocols iccp local-ip-addr 10.3.3.1
set protocols iccp peer 10.3.3.2 session-establishment-hold-time 340
set protocols iccp peer 10.3.3.2 backup-liveness-detection backup-peer-ip 10.207.64.234
set protocols iccp peer 10.3.3.2 liveness-detection minimum-receive-interval 60
set protocols iccp peer 10.3.3.2 liveness-detection transmit-interval minimum-interval
  60
set protocols igmp-snooping vlan v100
set protocols igmp-snooping vlan v200
set protocols igmp-snooping vlan v500
set protocols ospf area 0.0.0.0 interface irb.100 bfd-liveness-detection
  minimum-receive-interval 700
set protocols ospf area 0.0.0.0 interface irb.100 bfd-liveness-detection transmit-interval
  minimum-interval 350
set protocols ospf area 0.0.0.0 interface irb.100 bfd-liveness-detection transmit-interval
  threshold 500
set protocols ospf area 0.0.0.0 interface irb.200 bfd-liveness-detection
  minimum-receive-interval 700
set protocols ospf area 0.0.0.0 interface irb.200 bfd-liveness-detection transmit-interval
  minimum-interval 350
set protocols ospf area 0.0.0.0 interface irb.200 bfd-liveness-detection transmit-interval
  threshold 500
set protocols pim rp static address 10.0.0.3 group-ranges 233.252.0.0/8
set protocols pim interface irb.100 priority 100
set protocols pim interface irb.200 priority 500
set multi-chassis multi-chassis-protection 10.3.3.2 interface ae0
```

Configuring MC-LAG for Layer 3 Multicast Using VRRP on Two Switches

Step-by-Step Procedure

The following procedure 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 enable a multichassis protection link between MC-LAG peers:

1. Configure the number of LAGs on both Switch A and Switch B.

```
[edit chassis]
user@switch# set aggregated-devices ethernet device-count 3
```

2. Add member interfaces to the aggregated Ethernet interfaces on both Switch A and Switch B.

```
[edit interfaces]
user@switch# set xe-0/0/2 ether-options 802.3ad ae0
user@switch# set xe-0/0/3 ether-options 802.3ad ae0
```

```
[edit interfaces]
```

```
user@switch# set xe-0/0/4 ether-options 802.3ad ae1
user@switch# set xe-0/0/5 ether-options 802.3ad ae2
```

```
[edit interfaces]
user@switch# set xe-0/0/6 ether-options 802.3ad ae1
user@switch# set xe-0/0/7 ether-options 802.3ad ae2
```

3. Configure ae0 as the trunk interface between Switch A and Switch B.

```
[edit interfaces]
user@switch# set ae0 unit 0 family ethernet-switching interface-mode trunk
```

4. Configure ae0 as the multichassis protection link between Switch A and Switch B.

```
[edit]
user@switch# set multi-chassis multi-chassis-protection 10.3.3.1 interface ae0
```

```
[edit]
user@switch# set multi-chassis multi-chassis-protection 10.3.3.2 interface ae0
```

Step-by-Step Procedure

To enable ICCP:

1. Configure the local IP address to be in the ICCP connection on Switch A and Switch B.

```
[edit protocols]
user@switch# set iccp local-ip-addr 10.3.3.2
```

```
[edit protocols]
user@switch# set iccp local-ip-addr 10.3.3.1
```

2. Configure the peer IP address, minimum receive interval, and minimum transmit interval for a Bidirectional Forwarding Detection (BFD) session for ICCP on Switch A and Switch B.

NOTE: Configuring the minimum receive interval is required to enable BFD. We recommend a minimum receive interval value of 60 seconds.

```
[edit protocols]
user@switch# set iccp peer 10.3.3.1 liveness-detection minimum-receive-interval
  60
user@switch# set iccp peer 10.3.3.1 liveness-detection transmit-interval
  minimum-interval 60
```

```
[edit protocols]
```

```
user@switch# set iccp peer 10.3.3.2 liveness-detection minimum-receive-interval
  60
user@switch# set iccp peer 10.3.3.2 liveness-detection transmit-interval
  minimum-interval 60
```

3. (Optional) Configure the time during which an ICCP connection must be established between MC-LAG peers on Switch A and Switch B.

NOTE: On QFX and EX Series switches, the default session establishment hold time is 300 seconds. However, the session establishment time must be at least 100 seconds higher than the init delay time. You can optionally update the session establishment time to be 340 seconds and the init delay time to be 240 seconds.

```
[edit protocols]
user@switch# set iccp peer 10.3.3.1 session-establishment-hold-time 340
```

```
[edit protocols]
user@switch# set iccp peer 10.3.3.2 session-establishment-hold-time 340
```

4. (Optional) Configure the backup-liveness-detection statement on the management interface (fxp0) only.

   We recommend that you configure the backup liveness detection feature to implement faster failover of data traffic during an MC-LAG peer reboot.

NOTE: The backup-liveness-detection statement is supported starting in Junos OS Release 13.2R1.

```
[edit protocols]
user@switch# set iccp peer 10.3.3.1 backup-liveness-detection backup-peer-ip
  10.207.64.233
```

```
[edit protocols]
user@switch# set iccp peer 10.3.3.2 backup-liveness-detection backup-peer-ip
  10.207.64.234
```

5. Configure Layer 3 connectivity between the MC-LAG peers on both Switch A and Switch B.

```
[edit vlans]
user@switch# set v500 vlan-id 500
user@switch# set v500 l3-interface irb.500
```

```
[edit interfaces]
user@switch# set ae0 unit 0 family ethernet-switching vlan members v500
```

```
[edit interfaces]
user@switch# set irb unit 500 family inet address 10.3.3.2/8
```

```
[edit interfaces]
user@switch# set irb unit 500 family inet address 10.3.3.1/8
```

Step-by-Step Procedure

To enable the ae1 and ae2 MC-LAG interfaces:

1. Enable LACP on the MC-LAG interfaces on Switch A and Switch B.

NOTE: At least one end needs to be active. The other end can be either active or passive.

```
[edit interfaces]
user@switch# set ae1 aggregated-ether-options lacp active
user@switch# set ae2 aggregated-ether-options lacp active
```

2. Specify the same multichassis aggregated Ethernet identification number for each MC-LAG peer on Switch A and Switch B.

```
[edit interfaces]
user@switch# set ae1 aggregated-ether-options mc-ae mc-ae-id 3
user@switch# set ae2 aggregated-ether-options mc-ae mc-ae-id 4
```

3. Specify a unique chassis ID for the MC-LAG on the MC-LAG peers on Switch A and Switch B.

```
[edit interfaces]
user@switch# set ae1 aggregated-ether-options mc-ae chassis-id 0
user@switch# set ae2 aggregated-ether-options mc-ae chassis-id 0
```

```
[edit interfaces]
user@switch# set ae1 aggregated-ether-options mc-ae chassis-id 1
user@switch# set ae2 aggregated-ether-options mc-ae chassis-id 1
```

4. Specify the operating mode of the MC-LAGs on both Switch A and Switch B.

```
[edit interfaces]
user@switch# set ae1 aggregated-ether-options mc-ae mode active-active
user@switch# set ae2 aggregated-ether-options mc-ae mode active-active
```

5. Specify the status control for the MC-LAGs on Switch A and Switch B.

NOTE: You must configure status control on both Switch A and Switch B that host the MC-LAGs. If one peer is in active mode, the other must be in standby mode.

```
[edit interfaces]
user@switch# set ae1 aggregated-ether-options mc-ae status-control active
user@switch# set ae2 aggregated-ether-options mc-ae status-control active
```

```
[edit interfaces]
user@switch# set ae1 aggregated-ether-options mc-ae status-control standby
user@switch# set ae2 aggregated-ether-options mc-ae status-control standby
```

NOTE: If you configure both nodes as prefer-status-control-active, you must also configure ICCP peering using the peer's loopback address to make sure that the ICCP session does not go down because of physical link failures. Additionally, you must configure backup liveness detection on both of the MC-LAG nodes.

6. To minimize traffic loss, specify the number of seconds by which to delay bringing the multichassis aggregated Ethernet interface back to the up state when you reboot Switch A or Switch B.

```
[edit interfaces]
user@switch# set ae1 aggregated-ether-options mc-ae init-delay-time 240
user@switch# set ae2 aggregated-ether-options mc-ae init-delay-time 240
```

7. Specify the same LACP system ID for each MC-LAG on Switch A and Switch B.

```
[edit interfaces]
user@switch# set ae1 aggregated-ether-options lacp system-ID 00:01:02:03:04:05
user@switch# set ae2 aggregated-ether-options lacp system-ID 00:01:02:03:04:06
```

8. Specify the same LACP administration key on both Switch A and Switch B.

```
[edit interfaces]
user@switch# set ae1 aggregated-ether-options lacp admin-key 3
user@switch# set ae2 aggregated-ether-options lacp admin-key 3
```

9. Enable a VLAN for each MC-LAG on Switch A and Switch B.

```
[edit vlans]
```

```
user@switch# set v100 vlan-id 100
user@switch# set v200 vlan-id 200
```

```
[edit interfaces]
user@switch# set ae1 unit 0 family ethernet-switching vlan members v100
user@switch# set ae2 unit 0 family ethernet-switching vlan members v200
```

10. Configure ae1 and ae2 as trunk interfaces between Switch A and Switch B.

```
[edit interfaces]
user@switch# set ae1 unit 0 family ethernet-switching interface-mode trunk
user@switch# set ae2 unit 0 family ethernet-switching interface-mode trunk
```

Step-by-Step Procedure

To enable VRRP on the MC-LAGs:

1. Create an integrated routing and bridging (IRB) interface for each MC-LAG, assign a virtual IP address that is shared between each switch in the VRRP groups, and assign an individual IP address for each switch in the VRRP groups.

```
[edit interfaces]
user@switch# set irb unit 100 family inet address 10.1.1.11/8 vrrp-group 1
  virtual-address 10.1.1.1
user@switch# set irb unit 200 family inet address 10.1.1.21/8 vrrp-group 2
  virtual-address 10.1.1.2
```

```
[edit interfaces]
user@switch# set irb unit 100 family inet address 10.1.1.10/8 vrrp-group 1
  virtual-address 10.1.1.1
user@switch# set irb unit 200 family inet address 10.1.1.20/8 vrrp-group 2
  virtual-address 10.1.1.2
```

2. Assign the priority for each switch in the VRRP groups.

NOTE: The switch configured with the highest priority is the master.

```
[edit interfaces]
user@switch# set irb unit 100 family inet address 10.1.1.11/8 vrrp-group 1 priority 200
user@switch# set irb unit 200 family inet address 10.1.1.21/8 vrrp-group 2 priority
  200
```

```
[edit interfaces]
user@switch# set irb unit 100 family inet address 10.1.1.10/8 vrrp-group 1 priority
  150
user@switch# set irb unit 200 family inet address 10.1.1.20/8 vrrp-group 2 priority
  150
```

3. Enable the switch to accept all packets destined for the virtual IP address if it is the master in a VRRP group.

```
[edit interfaces]
user@switch# set irb unit 100 family inet address 10.1.1.11/8 vrrp-group 1 accept-data
user@switch# set irb unit 200 family inet address 10.1.1.21/8 vrrp-group 2 accept-data
```

```
[edit interfaces]
user@switch# set irb unit 100 family inet address 10.1.1.10/8 vrrp-group 1 accept-data
user@switch# set irb unit 200 family inet address 10.1.1.20/8 vrrp-group 2
  accept-data
```

4. Configure Layer 3 connectivity between Switch A and Switch B.

```
[edit]
user@switch# set v100 l3-interface irb.100
user@switch# set v200 l3-interface irb.200
```

Step-by-Step Procedure

To enable IGMP snooping:

- Enable IGMP snooping for all VLANs on Switch A and Switch B.

```
[edit protocols]
user@switch# set igmp-snooping vlan v100
user@switch# set igmp-snooping vlan v200
user@switch# set igmp-snooping vlan v500
```

NOTE: You must configure the multichassis-lag-replicate-state statement for IGMP snooping to work properly in an MC-LAG environment.

Step-by-Step Procedure

To configure OSPF as the Layer 3 protocol:

1. Configure an OSPF area on Switch A and Switch B.

```
[edit protocols ospf]
user@switch# set area 0.0.0.0
```

2. Assign the VLAN interfaces for the MC-LAGs as interfaces to the OSPF area on Switch A and Switch B.

```
[edit protocols ospf area 0.0.0.0]
user@switch# set interface irb.100
user@switch# set interface irb.200
```

3. Configure the minimum receive interval, minimum transmit interval, and transmit interval threshold for a Bidirectional Forwarding Detection (BFD) session for the OSPF interfaces on Switch A and Switch B.

```
[edit protocols ospf area 0.0.0.0]
user@switch# set interface irb.100 bfd-liveness-detection minimum-receive-interval
  700
user@switch# set interface irb.100 bfd-liveness-detection transmit-interval
  minimum-interval 350
user@switch# set interface irb.100 bfd-liveness-detection transmit-interval threshold
  500
user@switch# set interface irb.200 bfd-liveness-detection minimum-receive-interval
  700
user@switch# set interface irb.200 bfd-liveness-detection transmit-interval
  minimum-interval 350
user@switch# set interface irb.200 bfd-liveness-detection transmit-interval
  threshold 500
```

Step-by-Step Procedure

To configure Protocol Independent Multicast (PIM) as the multicast protocol:

1. Configure a static rendezvous point (RP) address on Switch A and Switch B.

```
[edit protocols pim]
user@switch# set rp static address 10.0.0.3
```

2. Configure the address ranges of the multicast groups for which Switch A and Switch B can be an RP.

```
[edit protocols pim rp static address 10.0.0.3]
user@switch# set group-ranges 233.252.0.0/8
```

3. Configure the priority of each PIM interface for being selected as the designated router.

   An interface with a higher priority value has a higher probability of being selected as the designated router.

NOTE: Configure the IP address on the active MC-LAG peer with a high IP address or a high designated router priority.

```
[edit protocols pim]
user@switch# set interface irb.100 priority 200
user@switch# set interface irb.200 priority 600
```

```
[edit protocols pim]
user@switch# set interface irb.100 priority 100
user@switch# set interface irb.200 priority 500
```

Results

From configuration mode on Switch A, confirm your configuration by entering the **show chassis**, **show interfaces**, **show multi-chassis**, **show protocols**, and **show vlans** commands. If the output does not display the required configuration, repeat the instructions in this example to correct the configuration.

Switch A

```
user@SwitchA# show chassis
aggregated-devices {
   ethernet {
      device-count 3;
   }
}
```

```
user@SwitchA# show interfaces
xe-0/0/2 {
   ether-options {
      802.3ad ae0;
   }
}
xe-0/0/3 {
   ether-options {
      802.3ad ae0;
   }
}
xe-0/0/4 {
   ether-options {
      802.3ad ae1;
   }
}
xe-0/0/5 {
   ether-options {
      802.3ad ae2;
   }
}
ae0 {
   unit 0 {
      family ethernet-switching {
         interface-mode trunk;
         vlan {
            members v500;
         }
      }
   }
}
ae1 {
   aggregated-ether-options {
      lacp {
         active;
         system-id 00:01:02:03:04:05;
         admin-key 3;
```

```
    }
    mc-ae {
      mc-ae-id 3;
      chassis-id 0;
      mode active-active;
      status-control active;
      init-delay-time 240;
    }
  }
  unit 0 {
    family ethernet-switching {
      interface-mode trunk;
      vlan {
        members v100;
      }
    }
  }
}
ae2 {
  aggregated-ether-options {
    lacp {
      active;
      system-id 00:01:02:03:04:06;
      admin-key 3;
    }
    mc-ae {
      mc-ae-id 4;
      chassis-id 0;
      mode active-active;
      status-control active;
      init-delay-time 240;
    }
  }
  unit 0 {
    family ethernet-switching {
      interface-mode trunk;
      vlan {
        members v200;
      }
    }
  }
}
irb {
  unit 100 {
    family inet {
      address 10.1.1.11/8 {
        vrrp-group 1 {
          virtual-address 10.1.1.1;
          priority 200;
          accept-data;
        }
      }
    }
  }
  unit 200 {
```

```
    family inet {
      address 10.1.1.21/8 {
        vrrp-group 2 {
          virtual-address 10.1.1.2;
          priority 200;
          accept-data;
        }
      }
    }
  }
  unit 500 {
    family inet {
      address 10.3.3.2/8;
    }
  }
}
```

```
user@SwitchA# show protocols
ospf {
  area 0.0.0.0 {
    interface irb.100 {
      bfd-liveness-detection {
        minimum-receive-interval 700;
        transmit-interval {
          minimum-interval 350;
          threshold 500;
        }
      }
    }
    interface irb.200 {
      bfd-liveness-detection {
        minimum-receive-interval 700;
        transmit-interval {
          minimum-interval 350;
          threshold 500;
        }
      }
    }
  }
}
pim {
  rp {
    static {
      address 10.0.0.3 {
        group-ranges {
          233.252.0.0/8;
        }
      }
    }
  }
  interface irb.100 {
    priority 200;
  }
  interface irb.200 {
```

```
      priority 600;
    }
  }
  iccp {
    local-ip-addr 10.3.3.2;
    peer 10.3.3.1 {
      session-establishment-hold-time 340;
      backup-liveness-detection {
        backup-peer-ip 10.207.64.233;
      }
      liveness-detection {
        minimum-receive-interval 60;
        transmit-interval {
          minimum-interval 60;
        }
      }
    }
  }
  igmp-snooping {
    vlan v100;
    vlan v200;
    vlan v500;
  }
```

```
user@SwitchA# show multi-chassis
multi-chassis-protection 10.3.3.1 {
  interface ae0;
}
```

```
user@SwitchA# show vlans
v100 {
  vlan-id 100;
  l3-interface irb.100;
}
v200 {
  vlan-id 200;
  l3-interface irb.200;
}
v500 {
  vlan-id 500;
  l3-interface irb.500;
}
```

Switch B

```
user@SwitchB# show chassis
aggregated-devices {
  ethernet {
    device-count 2;
  }
}
```

```
user@SwitchB# show interfaces
```

```
xe-0/0/2 {
  ether-options {
    802.3ad ae0;
  }
}
xe-0/0/3 {
  ether-options {
    802.3ad ae0;
  }
}
xe-0/0/6 {
  ether-options {
    802.3ad ae1;
  }
}
xe-0/0/7 {
  ether-options {
    802.3ad ae2;
  }
}
ae0 {
  unit 0 {
    family ethernet-switching {
      interface-mode trunk;
      vlan {
        members v500;
      }
    }
  }
}
ae1 {
  aggregated-ether-options {
    lacp {
      active;
      system-id 00:01:02:03:04:05;
      admin-key 3;
    }
    mc-ae {
      mc-ae-id 3;
      chassis-id 1;
      mode active-active;
      status-control standby;
      init-delay-time 240;
    }
  }
  unit 0 {
    family ethernet-switching {
      interface-mode trunk;
      vlan {
        members v100;
      }
    }
  }
}
ae2 {
```

```
    aggregated-ether-options {
      lacp {
        active;
        system-id 00:01:02:03:04:06;
        admin-key 3;
      }
      mc-ae {
        mc-ae-id 4;
        chassis-id 1;
        mode active-active;
        status-control standby;
        init-delay-time 240;
      }
    }
    unit 0 {
      family ethernet-switching {
        interface-mode trunk;
        vlan {
          members v200;
        }
      }
    }
}
irb {
    unit 100 {
      family inet {
        address 10.1.1.10/8 {
          vrrp-group 1 {
            virtual-address 10.1.1.1;
            priority 150;
            accept-data;
          }
        }
      }
    }
    unit 200 {
      family inet {
        address 10.1.1.20/8 {
          vrrp-group 2 {
            virtual-address 10.1.1.2;
            priority 150;
            accept-data;
          }
        }
      }
    }
    unit 500 {
      family inet {
        address 10.3.3.1/8;
      }
    }
}
```

```
user@SwitchB# show protocols
```

```
ospf {
  area 0.0.0.0 {
    interface irb.100 {
      bfd-liveness-detection {
        minimum-receive-interval 700;
        transmit-interval {
          minimum-interval 350;
          threshold 500;
        }
      }
    }
    interface irb.200 {
      bfd-liveness-detection {
        minimum-receive-interval 700;
        transmit-interval {
          minimum-interval 350;
          threshold 500;
        }
      }
    }
  }
}
pim {
  rp {
    static {
      address 10.0.0.3 {
        group-ranges {
          233.252.0.0/8;
        }
      }
    }
  }
  interface irb.100 {
    priority 100;
  }
  interface irb.200 {
    priority 500;
  }
}
iccp {
  local-ip-addr 10.3.3.1;
  peer 10.3.3.2 {
    session-establishment-hold-time 340;
    backup-liveness-detection {
      backup-peer-ip 10.207.64.234;
    }
    liveness-detection {
      minimum-receive-interval 60;
      transmit-interval {
        minimum-interval 60;
      }
    }
  }
}
igmp-snooping {
```

```
    vlan v100;
    vlan v200;
    vlan v500;
}
```

```
user@SwitchB# show multi-chassis
multi-chassis-protection 10.3.3.2 {
    interface ae0;
}
```

```
user@SwitchB# show vlans
v100 {
    vlan-id 100;
    l3-interface irb.100;
}
v200 {
    vlan-id 200;
    l3-interface irb.200;
}
v500 {
    vlan-id 500;
    l3-interface irb.500;
}
```

Verification

Verify that the configuration is working properly.

- Confirm That Switch A Is the Master Designated Router on page 227
- Verifying That Switch B Is the Backup Designated Router on page 228

Confirm That Switch A Is the Master Designated Router

Purpose Verify that Switch A is the master designated router (DR).

Action `user@switch>` **show pim interfaces**

```
Stat = Status, V = Version, NbrCnt = Neighbor Count,
S = Sparse, D = Dense, B = Bidirectional,
DR = Designated Router, P2P = Point-to-point link,
Active = Bidirectional is active, NotCap = Not Bidirectional Capable

Name               Stat Mode IP V State        NbrCnt JoinCnt(sg/*g) DR address
pime.32769         Down S     4 2   P2P,NotCap         0 0/0
irb.100            Up   S     4 2   DDR-DR,NotCap      1 0/0           10.1.1.11
irb.200            Up   S     4 2   DDR-DR,NotCap      2 0/0           10.1.1.21
```

Meaning The **DDR-DR** state of the VLAN interfaces in the output shows that Switch A is the master designated router.

Verifying That Switch B Is the Backup Designated Router

Purpose Confirm that Switch B is the backup designated router.

Action

```
user@switch> show pim interfaces
Stat = Status, V = Version, NbrCnt = Neighbor Count,
S = Sparse, D = Dense, B = Bidirectional,
DR = Designated Router, P2P = Point-to-point link,
Active = Bidirectional is active, NotCap = Not Bidirectional Capable

Name               Stat Mode IP V State        NbrCnt JoinCnt(sg/*g) DR address
pime.32769         Down S     4 2   P2P,NotCap         0 0/0
irb.100            Up   S     4 2   DDR-BDR,NotCap     1 0/0           10.1.1.11
irb.200            Up   S     4 2   DDR-BDR,NotCap     2 0/0           10.1.1.21
```

Meaning The **DDR-BDR** state of the VLAN interfaces in the output shows that Switch B is the backup designated router.

Example: Configuring Multichassis Link Aggregation for Layer 3 Unicast using MAC Address Synchronization

There are two methods for enabling Layer 3 unicast functionality across a multichassis link aggregation group (MC-LAG). You can choose either to synchronize the MAC addresses for the Layer 3 interfaces of the switches participating in the MC-LAG, or you can configure Virtual Router Redundancy Protocol (VRRP), but you cannot configure both at the same time. Because RVI interfaces share the same MAC address, if you enable MAC address synchronization, packets received on an MC-LAG peer with a destination MAC address that is the same as that of the peer's IRB MAC address will not be forwarded. The procedure to configure MAC address synchronization is included in this example. For more information on configuring VRRP for use in a Layer 3 unicast MC-LAG, see "Example: Configuring Multichassis Link Aggregation for Layer 3 Unicast Using VRRP" on page 246.

Requirements

This example uses the following hardware and software components:

- Junos OS Release 12.3 or later for the QFX Series
- Two QFX3500, QFX3600, EX4600, QFX5100, or QFX10000 switches

Before you configure an MC-LAG for Layer 3 unicast, be sure that you understand how to:

- Configure aggregated Ethernet interfaces on a switch. See *Example: Configuring Link Aggregation Between a QFX Series Product and an Aggregation Switch*.
- Configure the Link Aggregation Control Protocol (LACP) on aggregated Ethernet interfaces on a switch. See *Example: Configuring Link Aggregation with LACP Between a QFX Series Product and an Aggregation Switch*.
- Configure a standard MC-LAG between switches. See "Example: Configuring Multichassis Link Aggregation" on page 62.

Overview

In this example, you configure an MC-LAG across two switches by including interfaces from both switches in an aggregated Ethernet interface (ae1). To support the MC-LAG, you create a second aggregated Ethernet interface (ae0) for the interchassis link-protection link (ICL-PL). You also configure a multichassis protection link for the ICL-PL, Inter-Chassis Control Protocol (ICCP) for the peers hosting the MC-LAG, and Layer 3 connectivity between MC-LAG peers.

NOTE: Layer 3 connectivity is required for ICCP.

NOTE: On QFX5100 and QFX10000 switches, if you try to configure both VRRP over IRB and MAC synchronization, you will receive a commit error.

To complete the configuration, configure MAC address synchronization between the peers and specify the same IP address on both Layer 3 interface members (also known as the routed VLAN interface [RVI] or the integrated routing and bridging (IRB) interface) in the MC-LAG VLAN.

Topology

The topology used in this example consists of two switches hosting an MC-LAG. The two switches are connected to a server. Figure 32 on page 230 shows the topology of this example.

Figure 32: Configuring a Multichassis LAG Between Switch A and Switch B

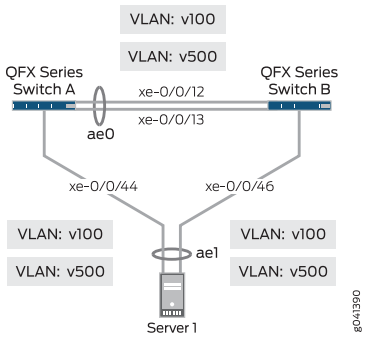

Table 9 on page 230 details the topology used in this configuration example.

Table 9: Components of the Topology for Configuring a Multichassis LAG Between Two Switches

Hostname	Base Hardware	Multichassis Link Aggregation Group
Switch A Switch B	QFX3500, QFX3600, EX4600, QFX5100, or QFX10000 switch QFX3500, QFX3600, EX4600, QFX5100, or QFX10000 switch	ae0 is configured as an aggregated Ethernet interface, and is used as an ICL-PL. The following interfaces are part of ae0: xe-0/0/12and xe-0/0/13 on Switch A and xe-0/0/12 and xe-0/0/13 on Switch B. These interfaces are included in VLAN v500. ae1 is configured as an MC-LAG, and the following two interfaces are part of ae1: xe-0/0/44 on Switch A and xe-0/0/46 on Switch B. These interfaces are included in VLAN v100.

Configuration

CLI Quick Configuration

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

Switch A

```
set chassis aggregated-devices ethernet device-count 2
set interfaces xe-0/0/12 ether-options 802.3ad ae0
set interfaces xe-0/0/13 ether-options 802.3ad ae0
set interfaces xe-0/0/44 ether-options 802.3ad ae1
set interfaces ae0 unit 0 family ethernet-switching port-mode trunk
set interfaces ae0 unit 0 family ethernet-switching vlan members v500
set interfaces ae0 unit 0 family ethernet-switching vlan members v100
set interfaces ae1 aggregated-ether-options lacp active
set interfaces ae1 aggregated-ether-options lacp system-id 00:01:02:03:04:05
set interfaces ae1 aggregated-ether-options lacp admin-key 3
set interfaces ae1 aggregated-ether-options mc-ae mc-ae-id 3
```

```
set interfaces ae1 aggregated-ether-options mc-ae chassis-id 0
set interfaces ae1 aggregated-ether-options mc-ae mode active-active
set interfaces ae1 aggregated-ether-options mc-ae status-control active
set interfaces ae1 aggregated-ether-options mc-ae init-delay-time 240
set interfaces ae1 unit 0 family ethernet-switching port-mode trunk
set interfaces ae1 unit 0 family ethernet-switching vlan members v100
set interfaces vlan unit 100 family inet address 10.1.1.10
set interfaces vlan unit 500 family inet address 10.3.3.2/8
set vlans v100 vlan-id 100
set vlans v100 l3-interface vlan.100
set vlans v100 mcae-mac-synchronize
set vlans v500 vlan-id 500
set vlans v500 l3-interface vlan.500
set protocols iccp local-ip-addr 10.3.3.2
set protocols iccp peer 10.3.3.1 session-establishment-hold-time 340
set protocols iccp peer 10.3.3.1 backup-liveness-detection backup-peer-ip 10.207.64.233
set protocols iccp peer 10.3.3.1 liveness-detection minimum-receive-interval 1000
set protocols iccp peer 10.3.3.1 liveness-detection transmit-interval minimum-interval
  1000
set protocols rstp interface ae0.0 disable
set protocols rstp interface ae1.0 edge
set protocols rstp interface all mode point-to-point
set protocols rstp bpdu-block-on-edge
set multi-chassis multi-chassis-protection 10.3.3.1 interface ae0
```

Switch B

```
set chassis aggregated-devices ethernet device-count 2
set interfaces xe-0/0/12 ether-options 802.3ad ae0
set interfaces xe-0/0/13 ether-options 802.3ad ae0
set interfaces xe-0/0/46 ether-options 802.3ad ae1
set interfaces ae0 unit 0 family ethernet-switching port-mode trunk
set interfaces ae0 unit 0 family ethernet-switching vlan members v500
set interfaces ae0 unit 0 family ethernet-switching vlan members v100
set interfaces ae1 aggregated-ether-options lacp active
set interfaces ae1 aggregated-ether-options lacp system-id 00:01:02:03:04:05
set interfaces ae1 aggregated-ether-options lacp admin-key 3
set interfaces ae1 aggregated-ether-options mc-ae mc-ae-id 3
set interfaces ae1 aggregated-ether-options mc-ae chassis-id 1
set interfaces ae1 aggregated-ether-options mc-ae mode active-active
set interfaces ae1 aggregated-ether-options mc-ae status-control standby
set interfaces ae1 aggregated-ether-options mc-ae init-delay-time 240
set interfaces ae1 unit 0 family ethernet-switching port-mode trunk
set interfaces ae1 unit 0 family ethernet-switching vlan members v100
set interfaces vlan unit 100 family inet address 10.1.1.10
set interfaces vlan unit 500 family inet address 10.3.3.1/8
set vlans v100 vlan-id 100
set vlans v100 l3-interface vlan.100
set vlans v100 mcae-mac-synchronize
set vlans v500 vlan-id 500
set vlans v500 l3-interface vlan.500
set protocols iccp local-ip-addr 10.3.3.1
set protocols iccp peer 10.3.3.2 session-establishment-hold-time 340
set protocols iccp peer 10.3.3.2 backup-liveness-detection backup-peer-ip 10.207.64.234
set protocols iccp peer 10.3.3.2 liveness-detection minimum-receive-interval 1000
```

```
set protocols iccp peer 10.3.3.2 liveness-detection transmit-interval minimum-interval
  1000
set protocols rstp interface ae0.0 disable
set protocols rstp interface ae1.0 edge
set protocols rstp interface all mode point-to-point
set protocols rstp bpdu-block-on-edge
set multi-chassis multi-chassis-protection 10.3.3.2 interface ae0
```

Configuring MC-LAG on Two Switches

Step-by-Step Procedure

To enable multichassis protection link between MC-LAG peers:

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 enable multichassis protection link between MC-LAG peers:

1. Configure the number of LAGs on both Switch A and Switch B.

```
[edit chassis]
  user@switch# set aggregated-devices ethernet device-count 2
```

2. Add member interfaces to the aggregated Ethernet interfaces on both Switch A and Switch B.

```
[edit interfaces]
  user@switch# set xe-0/0/12 ether-options 802.3ad ae0
    [edit interfaces]
      user@switch# set xe-0/0/12 ether-options 802.3ad ae0
    [edit interfaces]
      user@switch# set xe-0/0/13 ether-options 802.3ad ae0
    user@switch# set xe-0/0/13 ether-options 802.3ad ae0

[edit interfaces]
  user@switch# set xe-0/0/44 ether-options 802.3ad ae1
```

```
[edit interfaces]
  user@switch# set xe-0/0/46 ether-options 802.3ad ae1
```

3. Configure a trunk interface between Switch A and Switch B.

```
[edit interfaces]
  user@switch# set ae0 unit 0 family ethernet-switching port-mode trunk
```

4. Configure a multichassis protection link between Switch A and Switch B.

Switch A:

```
[edit]
  user@switch# set multi-chassis multi-chassis-protection 10.3.3.2 interface ae0
```

Switch B:

```
[edit]
  user@switch# set multi-chassis multi-chassis-protection 10.3.3.1 interface ae0
```

Step-by-Step Procedure

To enable ICCP:

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* .

1. Configure the local IP address to be in the ICCP connection on Switch A and Switch B.

   Switch A:

   ```
   [edit protocols]
   user@switch# set iccp local-ip-addr 10.3.3.2
   ```

   Switch B:

   ```
   [edit protocols]
   user@switch# set iccp local-ip-addr 10.3.3.1
   ```

2. Configure the peer IP address and minimum receive interval for a Bidirectional Forwarding Detection (BFD) session for ICCP on Switch A and Switch B.

   ```
   [edit protocols]
     user@switch# set iccp peer 10.3.3.1 liveness-detection minimum-receive-interval
     1000
   ```

   ```
   [edit protocols]
     user@switch# set iccp peer 10.3.3.2 liveness-detection minimum-receive-interval
     1000
   ```

3. Configure the peer IP address and minimum transmit interval for a BFD session for ICCP on Switch A and Switch B.

   NOTE: Configure at least 1000 milliseconds as the transmit interval minimum interval.

   ```
   [edit protocols]
   ```

```
user@switch# set iccp peer 10.3.3.1 liveness-detection transmit-interval
minimum-interval 1000
```

```
[edit protocols]
  user@switch# set iccp peer 10.3.3.2 liveness-detection transmit-interval
  minimum-interval 1000
```

4. (Optional) Configure the time during which an ICCP connection must succeed between MC-LAG peers on Switch A and Switch B.

NOTE: On QFX and EX Series switches, the default session establishment hold time is 300 seconds. However, the session establishment time must be at least 100 seconds higher than the init delay time. You can optionally update the session establishment time to be 340 seconds and the init delay time to be 240 seconds.

```
[edit protocols]
  user@switch# set iccp peer 10.3.3.1 session-establishment-hold-time 340
```

```
[edit protocols]
  user@switch# set iccp peer 10.3.3.2 session-establishment-hold-time 340
```

5. (Optional) Configure the backup IP address to be used for backup liveness detection on both Switch A and Switch B.

NOTE: By default, backup liveness detection is not enabled. Configuring a backup IP address helps achieve sub-second traffic loss during an MC-LAG peer reboot.

```
[edit protocols]
  user@switch# set iccp peer 10.3.3.1 backup-liveness-detection backup-peer-ip
  10.207.64.233
```

```
[edit protocols]
  user@switch# set iccp peer 10.3.3.2 backup-liveness-detection backup-peer-ip
  10.207.64.234
```

6. Configure Layer 3 connectivity across the ae0 ICCP link by adding a Layer 3 interface on both Switch A and Switch B.

```
[edit vlans v500]
  user@switch# set vlan-id 500
user@switch# set l3-interface vlan.500
```

```
[edit interfaces ae0 unit 0 family ethernet-switching]
   user@switch# set port-mode trunk
user@switch# set vlan members v500
```

Step-by-Step Procedure

To enable the MC-LAG interface:

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* .

1. Enable LACP on the MC-LAG interface on Switch A and Switch B.

   

   NOTE: At least one end needs to be active. The other end can be either active or passive.

   ```
   [edit interfaces]
   user@switch# set ae1 aggregated-ether-options lacp active
   ```

2. Specify the same multichassis aggregated Ethernet identification number on both MC-LAG peers on Switch A and Switch B.

   ```
   [edit interfaces]
   user@switch# set ae1 aggregated-ether-options mc-ae mc-ae-id 3
   ```

3. Specify a unique chassis ID for the MC-LAG on the MC-LAG peers on Switch A and Switch B.

   Switch A:

   ```
   [edit interfaces]
   user@switch# set ae1 aggregated-ether-options mc-ae chassis-id 0
   ```

   Switch B:

   ```
   [edit interfaces]
   user@switch# set ae1 aggregated-ether-options mc-ae chassis-id 1
   ```

4. Specify the operating mode of the MC-LAG on both Switch A and Switch B.

   

   NOTE: Only active-active mode is supported at this time.

   ```
   [edit interfaces]
   user@switch# set ae1 aggregated-ether-options mc-ae mode active-active
   ```

5. Specify the status control for MC-LAG on Switch A and Switch B.

NOTE: You must configure status control on both Switch A and Switch B hosting the MC-LAG. If one peer is in active mode, the other must be in standby mode.

Switch A:

```
[edit interfaces]
user@switch# set ae1 aggregated-ether-options mc-ae status-control active
```

Switch B:

```
[edit interfaces]
user@switch# set ae1 aggregated-ether-options mc-ae status-control standby
```

NOTE: If you configure both nodes as prefer-status-control-active, you must also configure ICCP peering using the peer's loopback address to make sure that the ICCP session does not go down because of physical link failures. Additionally, you must configure backup liveness detection on both of the MC-LAG nodes.

6. Specify the number of seconds by which the bring-up of the MC-AE interface should be deferred after you reboot Switch A and Switch B.

NOTE: On QFX and EX Series switches, the default session establishment hold time is 300 seconds. However, the session establishment time must be at least 100 seconds higher than the init delay time. You can optionally update the session establishment time to be 340 seconds and the init delay time to be 240 seconds.

```
[edit interfaces]
user@switch# set ae1 aggregated-ether-options mc-ae init-delay-time 240
```

7. Specify the same LACP system ID for the MC-LAG on Switch A and Switch B.

```
[edit interfaces]
user@switch# set ae1 aggregated-ether-options lacp system-ID 00:01:02:03:04:05
```

8. Specify the same LACP administration key on both Switch A and Switch B.

```
[edit interfaces]
user@switch# set ae1 aggregated-ether-options lacp admin-key 3
```

9. Enable a VLAN on the MC-LAG on Switch A and Switch B.

```
[edit interfaces]
user@switch# set ae1 unit 0 family ethernet-switching port-mode trunk

[edit]
user@switch# set vlans v100 vlan-id 100

[edit interfaces]
user@switch# set ae1 unit 0 family ethernet-switching vlan members v100
```

10. Create a Layer 3 interface for the MC-LAG VLAN and assign the same IP address on both Switch A and Switch B.

```
[edit]
user@switch# set vlans v100 l3-interface vlan.100

[edit interfaces]
user@switch# set vlan unit 100 family inet address 10.1.1.10
```

11. Configure MAC address synchronization in the MC-LAG VLAN on both Switch A and Switch B.

```
[edit]
user@switch# set vlans v100 mcae-mac-synchronize
```

Step-by-Step Procedure

To enable RSTP:

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* .

1. Enable RSTP globally on all interfaces on Switch A and Switch B.

```
[edit]
user@switch# set protocols rstp interface all mode point-to-point
```

2. Disable RSTP on the ICL-PL interfaces on Switch A and Switch B:

```
[edit]
user@switch# set protocols rstp interface ae0.0 disable
```

3. Configure the MC-LAG interfaces as edge ports on Switch A and Switch B.

NOTE: The ae1 interface is a downstream interface. This is why RSTP and bpdu-block-on-edge need to be configured.

```
[edit]
user@switch# set protocols rstp interface ae1.0 edge
```

4. Enable BPDU blocking on all interfaces except for the ICL-PL interfaces on Switch A and Switch B.

NOTE: The ae1 interface is a downstream interface. This is why RSTP and bpdu-block-on-edge need to be configured.

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

Results

Display the results of the configuration on Switch A.

```
chassis {
    aggregated-devices {
        ethernet {
            device-count 2;
        }
    }
}
interfaces {
    xe-0/0/12 {
        ether-options {
            802.3ad ae0;
        }
    }
    xe-0/0/13 {
        ether-options {
            802.3ad ae0;
        }
    }
    xe-0/0/44 {
        ether-options {
            802.3ad ae1;
        }
    }
    ae0 {
        unit 0 {
            family ethernet-switching {
                port-mode trunk;
                vlan {
                    members v500;
                }
            }
        }
    }
    ae1 {
        aggregated-ether-options {
            lacp {
                active;
                system-id 00:01:02:03:04:05;
                admin-key 3;
            }
```

```
        mc-ae {
          mc-ae-id 3;
          chassis-id 0;
          mode active-active;
          status-control active;
          init-delay-time 240
        }
      }
      unit 0 {
        family ethernet-switching {
          port-mode trunk;
          vlan {
            members v100;
          }
        }
      }
    }
    vlan {
      unit 100 {
        family inet {
          address 10.1.1.10;
        }
      }
      unit 500 {
        family inet {
          address 10.3.3.2/8;
        }
      }
    }
```

```
user@SwitchA# show protocols
iccp {
    local-ip-addr 10.3.3.2;
    peer 10.3.3.1 {
        session-establishment-hold-time 340;
        backup-liveness-detection {
            backup-peer-ip 10.207.64.233;
        }
        liveness-detection {
            minimum-receive-interval 1000;
            transmit-interval {
                minimum-interval 1000;
            }
        }
    }
}
rstp {
    interface ae0.0 {
        disable;
    }
    interface ae1.0 {
        edge;
    }
    interface all {
```

```
    mode point-to-point;
  }
  bpdu-block-on-edge;
}
```

```
user@SwitchA# show multi-chassis
multi-chassis-protection 10.3.3.1 {
  interface ae0;
}
```

```
user@SwitchA# show vlans
vlans {
  v100 {
    vlan-id 100;
    l3-interface vlan.100;
    mcae-mac-synchronize;
  }
  v500 {
    vlan-id 500;
    l3-interface vlan.500;
  }
}
```

Display the results of the configuration on Switch B.

```
chassis {
  aggregated-devices {
    ethernet {
      device-count 2;
    }
  }
}
interfaces {
  xe-0/0/12 {
    ether-options {
      802.3ad ae0;
    }
  }
  xe-0/0/13 {
    ether-options {
      802.3ad ae0;
    }
  }
  xe-0/0/46 {
    ether-options {
      802.3ad ae1;
    }
  }
  ae0 {
    unit 0 {
      family ethernet-switching {
        port-mode trunk;
        vlan {
```

```
            members v500;
          }
        }
      }
    }
    ae1 {
      aggregated-ether-options {
        lacp {
          active;
          system-id 00:01:02:03:04:05;
          admin-key 3;
        }
        mc-ae {
          mc-ae-id 3;
          chassis-id 1;
          mode active-active;
          status-control standby;
          init-delay-time 240
        }
      }
      unit 0 {
        family ethernet-switching {
          port-mode trunk;
          vlan {
            members v100;
          }
        }
      }
    }
    vlan {
      unit 100 {
        family inet {
          address 10.1.1.10;
        }
      }
      unit 500 {
        family inet {
          address 10.3.3.1/8;
        }
      }
    }
}
```

```
user@SwitchB# show protocols
iccp {
    local-ip-addr 10.3.3.1;
    peer 10.3.3.2 {
        session-establishment-hold-time 340;
        backup-liveness-detection {
            backup-peer-ip 10.207.64.234;
        }
        liveness-detection {
            minimum-receive-interval 1000;
            transmit-interval {
```

```
            minimum-interval 1000;
          }
        }
      }
    }
    rstp {
      interface ae0.0 {
        disable;
      }
      interface ae1.0 {
        edge;
      }
      interface all {
        mode point-to-point;
      }
      bpdu-block-on-edge;
    }
    }
    multi-chassis {
      multi-chassis-protection 10.3.3.2 {
        interface ae0;
      }
    }
    vlans {
      v100 {
        vlan-id 100;
        l3-interface vlan.100;
        mcae-mac-synchronize;
      }
      v500 {
        vlan-id 500;
        l3-interface vlan.500;
      }
    }
```

Verification

To verify that the MC-LAG group has been created and is working properly, perform these tasks:

- Verifying That ICCP Is Working on Switch A on page 242
- Verifying That ICCP Is Working on Switch B on page 243
- Verifying That LACP Is Active on Switch A on page 243
- Verifying That LACP Is Active on Switch B on page 244
- Verifying That the MC-AE and ICL-PL Interfaces Are Up on Switch A on page 244
- Verifying That the MC-AE and ICL-PL Interfaces Are Up on Switch B on page 245
- Verifying MAC Address Synchronization on Switch A and Switch B on page 245

Verifying That ICCP Is Working on Switch A

Purpose Verify that ICCP is running on Switch A.

Action

```
[edit]
user@switch# show iccp
```

```
Redundancy Group Information for peer 10.3.3.1
  TCP Connection       : Established
  Liveliness Detection : Up

 Client Application: MCSNOOPD

 Client Application: eswd
```

Meaning This output shows that the TCP connection between the peers hosting the MC-LAG is up, liveness detection is up, and MCSNOOPD and ESWD client applications are running.

Verifying That ICCP Is Working on Switch B

Purpose Verify that ICCP is running on Switch B.

Action show iccp

```
[edit]
user@switch# show iccp
```

```
Redundancy Group Information for peer 10.3.3.2
  TCP Connection       : Established
  Liveliness Detection : Up

 Client Application: MCSNOOPD

 Client Application: eswd
```

Meaning This output shows that the TCP connection between the peers hosting the MC-LAG is up, liveness detection is up, and MCSNOOPD and ESWD client applications are running.

Verifying That LACP Is Active on Switch A

Purpose Verify that LACP is active on Switch A.

Action

```
[edit]
user@switch# show lacp interfaces
```

```
Aggregated interface: ae1
    LACP state:       Role   Exp   Def  Dist  Col  Syn  Aggr  Timeout  Activity
      xe-0/0/46      Actor    No    No   Yes  Yes  Yes   Yes     Fast    Active
      xe-0/0/46    Partner    No    No   Yes  Yes  Yes   Yes     Fast    Active
    LACP protocol:        Receive State  Transmit State          Mux State
      xe-0/0/46                 Current   Fast periodic Collecting distributing
```

Meaning This output shows that Switch A is participating in LACP negotiation.

Verifying That LACP Is Active on Switch B

Purpose Verify that LACP is active on Switch B

Action

```
[edit]
user@switch# show lacp interfaces
```

```
Aggregated interface: ae1
    LACP state:       Role   Exp   Def  Dist  Col  Syn  Aggr  Timeout  Activity
      xe-0/0/44      Actor    No    No   Yes  Yes  Yes   Yes     Fast    Active
      xe-0/0/44    Partner    No    No   Yes  Yes  Yes   Yes     Fast    Active
    LACP protocol:        Receive State  Transmit State          Mux State
      xe-0/0/44                 Current   Fast periodic Collecting distributing
```

Meaning This output shows that Switch B is participating in LACP negotiation.

Verifying That the MC-AE and ICL-PL Interfaces Are Up on Switch A

Purpose Verify that the MC-AE and ICL-PL interfaces are up on Switch A.

Action

```
[edit]
user@switch# show interfaces mc-ae
```

```
Member Link                  : ae1
 Current State Machine's State: mcae active state
 Local Status                 : active
 Local State                  : up
 Peer Status                  : active
 Peer State                   : up
     Logical Interface        : ae1.0
     Topology Type            : bridge
     Local State              : up
     Peer State               : up
     Peer Ip/MCP/State        : 10.3.3.1 ae0.0 up
```

Meaning This output shows that the MC-AE interface on Switch A is up and active.

Verifying That the MC-AE and ICL-PL Interfaces Are Up on Switch B

Purpose Verify that the MC-AE and ICL-PL interfaces are up on Switch B.

Action

```
[edit]
user@switch# show interfaces mc-ae
```

```
Member Link                  : ae1
 Current State Machine's State: mcae active state
 Local Status                 : active
 Local State                  : up
 Peer Status                  : active
 Peer State                   : up
     Logical Interface        : ae1.0
     Topology Type            : bridge
     Local State              : up
     Peer State               : up
     Peer Ip/MCP/State        : 10.3.3.2 ae0.0 up
```

Meaning This output shows that the MC-AE interface on Switch B is up and active.

Verifying MAC Address Synchronization on Switch A and Switch B

Purpose Confirm that MAC address synchronization is working on both Switch A and Switch B.

Action

```
[edit]
user@switch# show ethernet-switching table vlan v100
```

```
Ethernet-switching table: 3 unicast entries
  VLAN              MAC address       Type         Age Interfaces
  v100               *                 Flood          - All-members
  v100               84:18:88:df:35:36 Static         - Router
  v100               84:18:88:df:83:0a Static         - Router
```

Meaning The output shows two static MAC addresses in VLAN v100. The addresses belong to the Layer 3 IRB/RVI interfaces of both Switch A and Switch B that you configured in the MC-LAG VLAN. Appearance of both addresses indicates that MAC address synchronization is working.

Troubleshooting

Troubleshooting a LAG That Is Down

Problem The **show interfaces terse** command shows that the MC-LAG is **down**

Solution Check the following:

- Verify that there is no configuration mismatch.
- Verify that all member ports are up.
- Verify that the MC-LAG is part of family Ethernet switching (Layer 2 LAG).
- Verify that the MC-LAG member is connected to the correct MC-LAG member at the other end.

Example: Configuring Multichassis Link Aggregation for Layer 3 Unicast Using VRRP

NOTE: Multichassis link aggregation (MC-LAG) is supported on QFX3500 and QFX3600 standalone switches running the original CLI, and QFX3500, QFX3600, QFX5100, EX4600, and QFX10000 standalone switches running Enhanced Layer 2 Software.

There are two methods for enabling Layer 3 unicast functionality across a multichassis link aggregation group (MC-LAG) to control traffic flow. You can choose either to synchronize the MAC addresses for the Layer 3 interfaces of the switches participating in the MC-LAG , or you can configure Virtual Router Redundancy Protocol (VRRP), but you cannot configure both at the same time. Because RVI interfaces share the same MAC address, if you enable MAC address synchronization, packets received on an MC-LAG peer with a destination MAC address that is the same as that of the peer's IRB MAC address will not be forwarded. The procedure to configure VRRP for use in a Layer 3 unicast MC-LAG is included in this example. For more information about configuring MAC

address synchronization, see "Example: Configuring Multichassis Link Aggregation for Layer 3 Unicast using MAC Address Synchronization" on page 228.

Requirements

This example uses the following hardware and software components:

- Junos OS Release 12.3 or later for the QFX Series
- Two QFX3500, QFX3600, EX4600, QFX5100, or QFX10000 switches

Before you configure an MC-LAG, be sure that you understand how to:

- Configure aggregated Ethernet interfaces on a switch. See *Example: Configuring Link Aggregation Between a QFX Series Product and an Aggregation Switch*.
- Configure the Link Aggregation Control Protocol (LACP) on aggregated Ethernet interfaces on a switch. See *Example: Configuring Link Aggregation with LACP Between a QFX Series Product and an Aggregation Switch*.
- Configure Virtual Router Redundancy Protocol (VRRP) on a switch. See *Configuring Basic VRRP Support for QFX*.

Overview

In this example, you configure an MC-LAG across two switches by including interfaces from both switches in an aggregated Ethernet interface (ae1). To support the MC-LAG, create a second aggregated Ethernet interface (ae0) for the interchassis control link-protection link (ICL-PL). Configure a multichassis protection link for the ICL-PL, the Inter-Chassis Control Protocol (ICCP) for the peers hosting the MC-LAG, and Layer 3 connectivity between MC-LAG peers.

NOTE: Layer 3 connectivity is required for ICCP.

NOTE: On QFX5100 and QFX10000 switches, if you try to configure both VRRP over IRB and MAC synchronization, you will receive a commit error.

To complete the configuration, enable VRRP by completing the following steps:

1. Create a routed VLAN interface (RVI).
2. Create a VRRP group and assign a virtual IP address that is shared between each switch in the VRRP group.

3. Enable a member of a VRRP group to accept all packets destined for the virtual IP address if it is the master in the VRRP group.
4. Configure Layer 3 connectivity between the VRRP groups.

Topology

The topology used in this example consists of two switches hosting MC-LAGs. The two switches are connected to a server. Figure 33 on page 248 shows the topology of this example.

Figure 33: Configuring a Multichassis LAG Between Switch A and Switch B

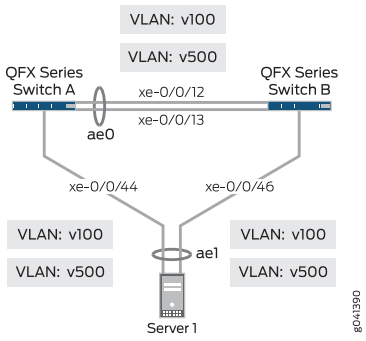

Table 10 on page 248 details the topology used in this configuration example.

Table 10: Components of the Topology for Configuring a Multichassis LAG Between Two Switches

Hostname	Base Hardware	Multichassis Link Aggregation Group
Switch A Switch B	QFX3500, QFX3600, EX4600, QFX5100, QFX10000 switch QFX3500, QFX3600, EX4600, QFX5100, QFX10000 switch	ae0 is configured as an aggregated Ethernet interface, and is used as an ICL-PL. The following interfaces are part of ae0: xe-0/0/12 and xe-0/0/13 on Switch A and xe-0/0/12 and xe-0/0/13 on Switch B. ae1 is configured as an MC-LAG, and the following two interfaces are part of ae1: xe-0/0/44 on Switch A and xe-0/0/46 on Switch B. .

Configuration

CLI Quick Configuration

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

NOTE: This example shows how to configure MC-LAG using both the original CLI and Enhanced Layer 2 Software (ELS).

In ELS, there are three statements and one additional statement that are different from the original CLI:

- The port-mode statement in the [edit interfaces *interface-name* unit *number* family ethernet-switching] hierarchy is not supported. Use the interface-mode statement instead.
- The vlan statement in the [edit interfaces *interface-name*] hierarchy is not supported. Use the irb statement instead.
- The vlan.*logical-interface-number* option in the [edit vlans *vlan-name* l3-interface] option is not supported. Use the irb.*logical-interface-number* option instead.
- The service-id statement in the [edit switch-options] hierarchy is required in the ELS CLI.

Switch A—Original CLI:

```
set chassis aggregated-devices ethernet device-count 2
set interfaces xe-0/0/12 ether-options 802.3ad ae0
set interfaces xe-0/0/13 ether-options 802.3ad ae0
set interfaces xe-0/0/44 ether-options 802.3ad ae1
set interfaces ae0 unit 0 family ethernet-switching port-mode trunk
set interfaces ae0 unit 0 family ethernet-switching vlan members v500 v100
set interfaces ae1 aggregated-ether-options lacp active
set interfaces ae1 aggregated-ether-options lacp system-id 00:01:02:03:04:05
set interfaces ae1 aggregated-ether-options lacp admin-key 3
set interfaces ae1 aggregated-ether-options mc-ae mc-ae-id 3
set interfaces ae1 aggregated-ether-options mc-ae chassis-id 0
set interfaces ae1 aggregated-ether-options mc-ae mode active-active
set interfaces ae1 aggregated-ether-options mc-ae status-control active
set interfaces ae1 aggregated-ether-options mc-ae init-delay-time 240
set interfaces ae1 unit 0 family ethernet-switching port-mode trunk
set interfaces ae1 unit 0 family ethernet-switching vlan members v100
set interfaces vlan unit 100 family inet address 10.1.1.11/8 vrrp-group 1 virtual-address 10.1.1.1
set interfaces vlan unit 100 family inet address 10.1.1.11/8 vrrp-group 1 priority 200
set interfaces vlan unit 100 family inet address 10.1.1.11/8 vrrp-group 1 accept-data
set interfaces vlan unit 500 family inet address 10.3.3.2/8
set vlans v100 vlan-id 100
set vlans v100 l3-interface vlan.100
set vlans v500 vlan-id 500
set vlans v500 l3-interface vlan.500
set protocols iccp local-ip-addr 10.3.3.2
set protocols iccp peer 10.3.3.1 session-establishment-hold-time 340
set protocols iccp peer 10.3.3.1 backup-liveness-detection backup-peer-ip 10.207.64.233
set protocols iccp peer 10.3.3.1 liveness-detection minimum-receive-interval 1000
set protocols iccp peer 10.3.3.1 liveness-detection transmit-interval minimum-interval 1000
set protocols rstp interface ae0 disable
set protocols rstp interface ae1 edge
set protocols rstp interface all mode point-to-point
```

```
set protocols rstp bpdu-block-on-edge
set multi-chassis multi-chassis-protection 10.3.3.1 interface ae0
```

Switch A—ELS

```
set chassis aggregated-devices ethernet device-count 2
set interfaces xe-0/0/12 ether-options 802.3ad ae0
set interfaces xe-0/0/13 ether-options 802.3ad ae0
set interfaces xe-0/0/44 ether-options 802.3ad ae1
set interfaces ae0 unit 0 family ethernet-switching interface-mode trunk
set interfaces ae0 unit 0 family ethernet-switching vlan members v500 v100
set interfaces ae1 aggregated-ether-options lacp active
set interfaces ae1 aggregated-ether-options lacp system-id 00:01:02:03:04:05
set interfaces ae1 aggregated-ether-options lacp admin-key 3
set interfaces ae1 aggregated-ether-options mc-ae mc-ae-id 3
set interfaces ae1 aggregated-ether-options mc-ae chassis-id 0
set interfaces ae1 aggregated-ether-options mc-ae mode active-active
set interfaces ae1 aggregated-ether-options mc-ae status-control active
set interfaces ae1 aggregated-ether-options mc-ae init-delay-time 240
set interfaces ae1 unit 0 family ethernet-switching interface-mode trunk
set interfaces ae1 unit 0 family ethernet-switching vlan members v100
set interfaces irb unit 100 family inet address 10.1.1.11/8 vrrp-group 1 virtual-address 10.1.1.1
set interfaces irb unit 100 family inet address 10.1.1.11/8 vrrp-group 1 priority 200
set interfaces irb unit 100 family inet address 10.1.1.11/8 vrrp-group 1 accept-data
set interfaces irb unit 500 family inet address 10.3.3.2/8
set vlans v100 vlan-id 100
set vlans v100 l3-interface irb.100
set vlans v500 vlan-id 500
set vlans v500 l3-interface irb.500
set protocols iccp local-ip-addr 10.3.3.2
set protocols iccp peer 10.3.3.1 session-establishment-hold-time 340
set protocols iccp peer 10.3.3.1 backup-liveness-detection backup-peer-ip 10.207.64.233
set protocols iccp peer 10.3.3.1 liveness-detection minimum-receive-interval 1000
set protocols iccp peer 10.3.3.1 liveness-detection transmit-interval minimum-interval 1000
set protocols rstp interface ae1 edge
set protocols rstp interface ae1 mode point-to-point
set protocols rstp bpdu-block-on-edge
set multi-chassis multi-chassis-protection 10.3.3.1 interface ae0
set switch-options service-id 10
```

Switch B—Original CLI:

```
set chassis aggregated-devices ethernet device-count 2
set interfaces xe-0/0/12 ether-options 802.3ad ae0
set interfaces xe-0/0/13 ether-options 802.3ad ae0
set interfaces xe-0/0/46 ether-options 802.3ad ae1
set interfaces ae0 unit 0 family ethernet-switching port-mode trunk
set interfaces ae0 unit 0 family ethernet-switching vlan members v500 v100
set interfaces ae1 aggregated-ether-options lacp active
set interfaces ae1 aggregated-ether-options lacp system-id 00:01:02:03:04:05
set interfaces ae1 aggregated-ether-options lacp admin-key 3
set interfaces ae1 aggregated-ether-options mc-ae mc-ae-id 3
set interfaces ae1 aggregated-ether-options mc-ae chassis-id 1
set interfaces ae1 aggregated-ether-options mc-ae mode active-active
set interfaces ae1 aggregated-ether-options mc-ae status-control standby
set interfaces ae1 aggregated-ether-options mc-ae init-delay-time 240
set interfaces ae1 unit 0 family ethernet-switching port-mode trunk
set interfaces ae1 unit 0 family ethernet-switching vlan members v100
set interfaces vlan unit 100 family inet address 10.1.1.10/8 vrrp-group 1 virtual-address 10.1.1.1
set interfaces vlan unit 100 family inet address 10.1.1.10/8 vrrp-group 1 priority 150
```

```
set interfaces vlan unit 100 family inet address 10.1.1.10/8 vrrp-group 1 accept-data
set interfaces vlan unit 500 family inet address 10.3.3.1/8
set vlans v100 vlan-id 100
set vlans v100 l3-interface vlan.100
set vlans v500 vlan-id 500
set vlans v500 l3-interface vlan.500
set protocols iccp local-ip-addr 10.3.3.1
set protocols iccp peer 10.3.3.2 session-establishment-hold-time 340
set protocols iccp peer 10.3.3.2 backup-liveness-detection backup-peer-ip 10.207.64.234
set protocols iccp peer 10.3.3.2 liveness-detection minimum-receive-interval 1000
set protocols iccp peer 10.3.3.2 liveness-detection transmit-interval minimum-interval 1000
set protocols rstp interface ae0 disable
set protocols rstp interface ae1 edge
set protocols rstp interface all mode point-to-point
set protocols rstp bpdu-block-on-edge
set multi-chassis multi-chassis-protection 10.3.3.2 interface ae0
```

Switch B—ELS:

```
set chassis aggregated-devices ethernet device-count 2
set interfaces xe-0/0/12 ether-options 802.3ad ae0
set interfaces xe-0/0/13 ether-options 802.3ad ae0
set interfaces xe-0/0/46 ether-options 802.3ad ae1
set interfaces ae0 unit 0 family ethernet-switching interface-mode trunk
set interfaces ae0 unit 0 family ethernet-switching vlan members v500 v100
set interfaces ae1 aggregated-ether-options lacp active
set interfaces ae1 aggregated-ether-options lacp system-id 00:01:02:03:04:05
set interfaces ae1 aggregated-ether-options lacp admin-key 3
set interfaces ae1 aggregated-ether-options mc-ae mc-ae-id 3
set interfaces ae1 aggregated-ether-options mc-ae chassis-id 1
set interfaces ae1 aggregated-ether-options mc-ae mode active-active
set interfaces ae1 aggregated-ether-options mc-ae status-control standby
set interfaces ae1 aggregated-ether-options mc-ae init-delay-time 240
set interfaces ae1 unit 0 family ethernet-switching interface-mode trunk
set interfaces ae1 unit 0 family ethernet-switching vlan members v100
set interfaces irb unit 100 family inet address 10.1.1.10/8 vrrp-group 1 virtual-address 10.1.1.1
set interfaces irb unit 100 family inet address 10.1.1.10/8 vrrp-group 1 priority 150
set interfaces irb unit 100 family inet address 10.1.1.10/8 vrrp-group 1 accept-data
set interfaces irb unit 500 family inet address 10.3.3.1/8
set vlans v100 vlan-id 100
set vlans v100 l3-interface irb.100
set vlans v500 vlan-id 500
set vlans v500 l3-interface irb.500
set protocols iccp local-ip-addr 10.3.3.1
set protocols iccp peer 10.3.3.2 session-establishment-hold-time 340
set protocols iccp peer 10.3.3.2 backup-liveness-detection backup-peer-ip 10.207.64.234
set protocols iccp peer 10.3.3.2 liveness-detection minimum-receive-interval 1000
set protocols iccp peer 10.3.3.2 liveness-detection transmit-interval minimum-interval 1000
set protocols rstp interface ae1 edge
set protocols rstp interface ae1 mode point-to-point
set protocols rstp bpdu-block-on-edge
set multi-chassis multi-chassis-protection 10.3.3.2 interface ae0
set switch-options service-id 10
```

Configuring MC-LAG on Two Switches

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 enable multichassis protection link between MC-LAG peers:

1. Configure the number of LAGs on both Switch A and Switch B.

```
[edit chassis]
  user@switch# set aggregated-devices ethernet device-count 2
```

2. Add member interfaces to the aggregated Ethernet interfaces on both Switch A and Switch B.

```
[edit interfaces]
  user@switch# set xe-0/0/12 ether-options 802.3ad ae0
[edit interfaces]
  user@switch# set xe-0/0/13 ether-options 802.3ad ae0 ae0
```

```
[edit interfaces]
  user@switch# set xe-0/0/44 ether-options 802.3ad ae0 ae1
```

```
[edit interfaces]
  user@switch# set xe-0/0/46 ether-options 802.3ad ae0 ae1
```

3. Configure a trunk interface between Switch A and Switch B using the original CLI.

```
[edit interfaces]
  user@switch# set ae0 unit 0 family ethernet-switching port-mode trunk
```

4. Configure a trunk interface between Switch A and Switch B using ELS.

```
[edit interfaces]
  user@switch# set ae0 unit 0 family ethernet-switching interface-mode trunk
```

5. Configure a multichassis protection link between Switch A and Switch B.

```
[edit]
  user@switch# set multi-chassis multi-chassis-protection 10.3.3.1 interface ae0
```

```
[edit]
  user@switch# set multi-chassis multi-chassis-protection 10.3.3.2 interface ae0
```

Step-by-Step Procedure

To enable ICCP:

1. Configure the local IP address to be in the ICCP connection on Switch A and Switch B.

```
[edit protocols]
  user@switch# set iccp local-ip-addr 10.3.3.2
```

```
[edit protocols]
  user@switch# set iccp local-ip-addr 10.3.3.1
```

2. Configure the peer IP address and minimum receive interval for a Bidirectional Forwarding Detection (BFD) session for ICCP on Switch A and Switch B.

```
[edit protocols]
  user@switch# set iccp peer 10.3.3.1 liveness-detection minimum-receive-interval
  1000
```

```
[edit protocols]
  user@switch# set iccp peer 10.3.3.2 liveness-detection minimum-receive-interval
  1000
```

3. Configure the peer IP address and minimum transmit interval for a BFD session for ICCP on Switch A and Switch B.

```
[edit protocols]
  user@switch# set iccp peer 10.3.3.1 liveness-detection transmit-interval
  minimum-interval 1000
```

```
[edit protocols]
  user@switch# set iccp peer 10.3.3.2 liveness-detection transmit-interval
  minimum-interval 1000
```

4. (Optional) Configure the time during which an ICCP connection must succeed between MC-LAG peers on Switch A and Switch B.

NOTE: On QFX and EX Series switches, the default session establishment hold time is 300 seconds. However, the session establishment time must be at least 100 seconds higher than the init delay time. You can optionally update the session establishment time to be 340 seconds and the init delay time to be 240 seconds.

```
[edit protocols]
  user@switch# set iccp peer10.3.3.1 session-establishment-hold-time 340
```

```
[edit protocols]
```

```
user@switch# set iccp peer 10.3.3.2 session-establishment-hold-time 340
```

5. (Optional) Configure the backup IP address to be used for backup liveness detection on both Switch A and Switch B.

NOTE: By default, backup liveness detection is not enabled. Configuring a backup IP address helps achieve sub-second traffic loss during an MC-LAG peer reboot.

```
[edit protocols]
  user@switch# set iccp peer 10.3.3.1 backup-liveness-detection backup-peer-ip
  10.207.64.233
```

```
[edit protocols]
  user@switch# set iccp peer 10.3.3.2 backup-liveness-detection backup-peer-ip
  10.207.64.234
```

6. Configure Layer 3 connectivity between the MC-LAG peers on both Switch A and Switch B using the original CLI.

```
[edit vlans]
  user@switch# set v500 vlan-id 500
```

```
[edit vlans]
  user@switch# set v500 l3-interface vlan.500
```

```
[edit interfaces]
  user@switch# set ae0 unit 0 family ethernet-switching port-mode trunk vlan
  members v500 v100
```

7. Configure Layer 3 connectivity between the MC-LAG peers on both Switch A and Switch B using ELS.

```
edit vlans]
  user@switch# set v500 vlan-id 500
```

```
[edit vlans]
  user@switch# set v500 l3-interface irb.500
```

```
[edit interfaces]
  user@switch# set ae0 unit 0 family ethernet-switching interface-mode trunk vlan
  members v500 v100
```

Step-by-Step Procedure

To enable the MC-LAG interface:

1. Enable LACP on the MC-LAG interface on Switch A and Switch B.

   

   NOTE: At least one end needs to be active. The other end can be either active or passive.

   ```
   [edit interfaces]
     user@switch# set ae1 aggregated-ether-options lacp active
   ```

2. Specify the same multichassis aggregated Ethernet identification number on both MC-LAG peers on Switch A and Switch B.

   ```
   [edit interfaces]
     user@switch# set ae1 aggregated-ether-options mc-ae mc-ae-id 3
   ```

3. Specify the same service ID on Switch A and Switch B.

   ```
   [edit]
     user@switch# set switch-options service-id 10
   ```

4. Specify a unique chassis ID for the MC-LAG on the MC-LAG peers on Switch A and Switch B.

   ```
   [edit interfaces]
     user@switch# set ae1 aggregated-ether-options mc-ae chassis-id 0
   ```

   ```
   [edit interfaces]
     user@switch# set ae1 aggregated-ether-options mc-ae chassis-id 1
   ```

5. Specify the operating mode of the MC-LAG on both Switch A and Switch B.

   

   NOTE: Only active-active mode is supported at this time.

   ```
   [edit interfaces]
     user@switch# set ae1 aggregated-ether-options mc-ae mode active-active
   ```

6. Specify the status control for MC-LAG on Switch A and Switch B.

NOTE: You must configure status control on both Switch A and Switch B hosting the MC-LAG. If one peer is in active mode, the other must be in standby mode.

```
[edit interfaces]
  user@switch# set ae1 aggregated-ether-options mc-ae status-control active
```

```
[edit interfaces]
  user@switch# set ae1 aggregated-ether-options mc-ae status-control standby
```

NOTE: f you configure both nodes as prefer-status-control-active, you must also configure ICCP peering using the peer's loopback address to make sure that the ICCP session does not go down because of physical link failures. Additionally, you must configure backup liveness detection on both of the MC-LAG nodes.

7. Specify the number of seconds by which the bring-up of the multichassis aggregated Ethernet interface should be deferred after you reboot Switch A and Switch B.

NOTE: On QFX and EX Series switches, the default session establishment hold time is 300 seconds. However, the session establishment time must be at least 100 seconds higher than the init delay time. You can optionally update the session establishment time to be 340 seconds and the init delay time to be 240 seconds.

```
[edit interfaces]
  user@switch# set ae1 aggregated-ether-options mc-ae init-delay-time 240
```

8. Specify the same LACP system ID for the MC-LAG on Switch A and Switch B.

```
[edit interfaces]
  user@switch# set ae1 aggregated-ether-options lacp system-ID 00:01:02:03:04:05
```

9. Specify the same LACP administration key on both Switch A and Switch B.

```
[edit interfaces]
  user@switch# set ae1 aggregated-ether-options lacp admin-key 3
```

10. Enable a VLAN on the MC-LAG on Switch A and Switch B using the original CLI.

```
[edit interfaces]
  user@switch# set ae1 unit 0 family ethernet-switching port-mode trunk
```

```
[edit]
  user@switch# set vlans v100 vlan-id 100
```

```
[edit interfaces]
  user@switch# set ae1 unit 0 family ethernet-switching vlan members v100
```

11. Enable a VLAN on the MC-LAG on Switch A and Switch B using ELS.

```
[edit interfaces]
  user@switch# set ae1 unit 0 family ethernet-switching interface-mode trunk
```

```
[edit]
  user@switch# set vlans v100 vlan-id 100
```

```
[edit interfaces]
  user@switch# set ae1 unit 0 family ethernet-switching vlan members v100
```

12. Enable VRRP on the MC-LAG on Switch A and Switch B:

- Create a routed VLAN interface (RVI), assign a virtual IP address that is shared between each switch in the VRRP group, and assign an individual IP address for each switch in the VRRP group:

```
[edit interfaces]
  user@switch# set vlan unit 100 family inet address 10.1.1.11/8 vrrp-group 1
  virtual-address 10.1.1.1
```

```
[edit interfaces]
  user@switch# set vlan unit 100 family inet address 10.1.1.10/8 vrrp-group 1
  virtual-address 10.1.1.1
```

- Assign the priority for each switch in the VRRP group:

NOTE: The switch configured with the highest priority is the master.

```
[edit interfaces]
  user@switch# set vlan unit 100 family inet address 10.1.1.11/8 vrrp-group 1 priority
  200
```

```
[edit interfaces]
  user@switch# set vlan unit 100 family inet address 10.1.1.10/8 vrrp-group 1
  priority 150
```

- Enable the switch to accept all packets destined for the virtual IP address if it is the master in the VRRP group:

```
[edit interfaces]
  user@switch# set vlan unit 100 family inet address 10.1.1.11/8 vrrp-group 1
  accept-data
```

```
[edit interfaces]
  user@switch# set vlan unit 100 family inet address 10.1.1.10/8 vrrp-group 1
  accept-data
```

- Configure Layer 3 connectivity between Switch A and Switch B.

```
[edit interfaces]
  user@switch# set vlans v100 l3-interface vlan.100
```

13. Enable VRRP on the MC-LAG on Switch A and Switch B using ELS:

- Create a routed VLAN interface (RVI), assign a virtual IP address that is shared between each switch in the VRRP group, and assign an individual IP address for each switch in the VRRP group:

```
[edit interfaces]
  user@switch# set irb unit 100 family inet address 10.1.1.11/8 vrrp-group 1
  virtual-address 10.1.1.1
```

```
[edit interfaces]
  user@switch# set irb unit 100 family inet address 10.1.1.10/8 vrrp-group 1
  virtual-address 10.1.1.1
```

- Assign the priority for each switch in the VRRP group:

NOTE: The switch configured with the highest priority is the master.

```
[edit interfaces]
  user@switch# set irb unit 100 family inet address 10.1.1.11/8 vrrp-group 1 priority
  200
```

```
[edit interfaces]
  user@switch# set irb unit 100 family inet address 10.1.1.10/8 vrrp-group 1 priority
  150
```

- Enable the switch to accept all packets destined for the virtual IP address if it is the master in the VRRP group:

```
[edit interfaces]
  user@switch# set irb unit 100 family inet address 10.1.1.11/8 vrrp-group 1
  accept-data
```

```
[edit interfaces]
  user@switch# set irb unit 100 family inet address 10.1.1.10/8 vrrp-group 1
  accept-data
```

- Configure Layer 3 connectivity between Switch A and Switch B.

```
[edit interfaces]
  user@switch# set irb v100 l3-interface irb.100
```

Step-by-Step Procedure

To enabled RSTP:

1. Enable RSTP globally on all interfaces on Switch A and Switch B.

NOTE: The all option is not available on ELS, so you cannot issue this command on ELS.

```
[edit]
  user@switch# set protocols rstp interface all mode point-to-point
```

```
[edit]
  user@switch# set protocols rstp interface ae1 mode point-to-point
```

2. Disable RSTP on the ICL-PL interfaces on Switch A and Switch B.

NOTE: This command is not needed on ELS.

```
[edit]
  user@switch# set protocols rstp interface ae0 disable
```

3. Configure the MC-LAG interfaces as edge ports on Switch A and Switch B.

NOTE: The ae1 interface is a downstream interface. This is why RSTP and bpdu-block-on-edge need to be configured.

```
[edit]
  user@switch# set protocols rstp interface ae1 edge
```

4. Enable BPDU blocking on all interfaces except for the ICL-PL interfaces on Switch A and Switch B.

NOTE: The ae1 interface is a downstream interface. This is why RSTP and bpdu-block-on-edge need to be configured.

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

Results

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

Switch A—Original CLI

```
user@SwitchA# show chassis
aggregated-devices {
  ethernet {
    device-count 2;
  }
}
```

```
user@SwitchA# show interfaces
  interfaces {
    xe-0/0/12 {
      ether-options {
        802.3ad ae0;
      }
    }
    xe-0/0/13 {
      ether-options {
        802.3ad ae0;
      }
    }
    xe-0/0/44 {
      ether-options {
        802.3ad ae1;
      }
    }
    ae0 {
      unit 0 {
        family ethernet-switching {
          port-mode trunk;
          vlan {
            members v500;
          }
        }
      }
    }
```

```
    ae1 {
      aggregated-ether-options {
        lacp {
          active;
          system-id 00:01:02:03:04:05;
          admin-key 3;
        }
        mc-ae {
          mc-ae-id 3;
          chassis-id 0;
          mode active-active;
          status-control active;
          init-delay-time 240;
        }
      }
      unit 0 {
        family ethernet-switching {
          port-mode trunk;
          vlan {
            members v100;
          }
        }
      }
    }
    vlan {
      unit 100 {
        family inet {
          address 10.1.1.11/8 {
            vrrp-group 1 {
              virtual-address 10.1.1.1;
              priority 200;
              accept-data;
            }
          }
        }
      }
      unit 500 {
        family inet {
          address 10.3.3.2/8;
        }
      }
    }
  }
```

```
user@SwitchA# show protocols
iccp {
  local-ip-addr 10.3.3.2;
  peer 10.3.3.1 {
    session-establishment-hold-time 340;
    backup-liveness-detection {
      backup-peer-ip 10.207.64.233;
    }
    liveness-detection {
      minimum-receive-interval 1000;
```

```
        transmit-interval {
          minimum-interval 1000;
        }
      }
    }
}
rstp {
    interface ae0.0 {
      disable;
    }
    interface ae1.0 {
      edge;
    }
    interface all {
      mode point-to-point;
    }
    bpdu-block-on-edge;
}
```

```
user@SwitchA# show multi-chassis
multi-chassis-protection 10.3.3.1 {
    interface ae0;
}
```

```
user@SwitchA# show vlans
v100 {
    vlan-id 100;
    l3-interface vlan.100;
}
v500 {
    vlan-id 500;
    l3-interface vlan.500;
}
```

Switch A—ELS

```
user@SwitchA# show chassis
aggregated-devices {
    ethernet {
      device-count 2;
    }
}
```

```
user@SwitchA# show interfaces
interfaces {
    xe-0/0/12 {
      ether-options {
        802.3ad ae0;
      }
    }
    xe-0/0/13 {
      ether-options {
        802.3ad ae0;
```

```
    }
  }
  xe-0/0/44 {
    ether-options {
      802.3ad ae1;
    }
  }
  ae0 {
    unit 0 {
      family ethernet-switching {
        interface-mode trunk;
        vlan {
          members v500;
        }
      }
    }
  }
  ae1 {
    aggregated-ether-options {
      lacp {
        active;
        system-id 00:01:02:03:04:05;
        admin-key 3;
      }
      mc-ae {
        mc-ae-id 3;
        chassis-id 0;
        mode active-active;
        status-control active;
        init-delay-time 240;
      }
    }
    unit 0 {
      family ethernet-switching {
        interface-mode trunk;
        vlan {
          members v100;
        }
      }
    }
  }
  vlan {
    unit 100 {
      family inet {
        address 10.1.1.11/8 {
          vrrp-group 1 {
            virtual-address 10.1.1.1;
            priority 200;
            accept-data;
          }
        }
      }
    }
    unit 500 {
      family inet {
```

```
            address 10.3.3.2/8;
        }
    }
  }
}
```

```
user@SwitchA# show protocols
iccp {
  local-ip-addr 10.3.3.2;
  peer 10.3.3.1 {
    session-establishment-hold-time 340;
    backup-liveness-detection {
      backup-peer-ip 10.207.64.233;
    }
    liveness-detection {
      minimum-receive-interval 1000;
      transmit-interval {
        minimum-interval 1000;
      }
    }
  }
}
rstp {
  interface ae1.0 {
    edge;
  }
  mode point-to-point;
  }
  bpdu-block-on-edge;
  }
```

```
user@SwitchA# show multi-chassis
multi-chassis-protection 10.3.3.1 {
  interface ae0;
}
```

```
user@SwitchA# show switch-options
service-id 10;
```

```
user@SwitchA# show vlans
v100 {
  vlan-id 100;
  l3-interface irb.100;
}
v500 {
  vlan-id 500;
  l3-interface irb.500;
}
```

Switch B—Original CLI

```
user@SwitchB# show chassis
```

```
aggregated-devices {
  ethernet {
    device-count 2;
  }
}
```

```
user@SwitchB# show interfaces
xe-0/0/12 {
  ether-options {
    802.3ad ae0;
  }
}
xe-0/0/13 {
  ether-options {
    802.3ad ae0;
  }
}
xe-0/0/44 {
  ether-options {
    802.3ad ae1;
  }
}
ae0 {
  unit 0 {
    family ethernet-switching {
      port-mode trunk;
      vlan {
        members v500;
      }
    }
  }
}
ae1 {
  aggregated-ether-options {
    lacp {
      active;
      system-id 00:01:02:03:04:05;
      admin-key 3;
    }
    mc-ae {
      mc-ae-id 3;
      chassis-id 1;
      mode active-active;
      status-control active;
      init-delay-time 240;
    }
  }
  unit 0 {
    family ethernet-switching {
      port-mode trunk;
      vlan {
        members v100;
      }
    }
```

```
    }
}
vlan {
    unit 100 {
        family inet {
            address 10.1.1.10/8 {
                vrrp-group 1 {
                    virtual-address 10.1.1.1;
                    priority 200;
                    accept-data;
                }
            }
        }
    }
    unit 500 {
        family inet {
            address 10.3.1/8;
        }
    }
}
```

```
user@SwitchB# show protocols
iccp {
    local-ip-addr 10.3.3.1;
    peer 10.3.3.2 {
        session-establishment-hold-time 340;
        backup-liveness-detection {
            backup-peer-ip 10.207.64.234;
        }
        liveness-detection {
            minimum-receive-interval 1000;
            transmit-interval {
                minimum-interval 1000;
            }
        }
    }
}
rstp {
    interface ae0.0 {
        disable;
    }
    interface ae1.0 {
        edge;
    }
    interface all {
        mode point-to-point;
    }
    bpdu-block-on-edge;
}
```

```
user@SwitchB# show multi-chassis
multi-chassis-protection 10.3.3.2 {
    interface ae0;
}
```

```
user@SwitchB# show vlans
v100 {
   vlan-id 100;
   l3-interface vlan.100;
}
v500 {
   vlan-id 500;
   l3-interface vlan.500;
}
```

Switch B—ELS

```
user@SwitchB# show chassis
aggregated-devices {
   ethernet {
      device-count 2;
   }
}
```

```
user@SwitchB# show interfaces
xe-0/0/12 {
   ether-options {
      802.3ad ae0;
   }
}
xe-0/0/13 {
   ether-options {
      802.3ad ae0;
   }
}
xe-0/0/44 {
   ether-options {
      802.3ad ae1;
   }
}
ae0 {
   unit 0 {
      family ethernet-switching {
         interface-mode trunk;
         vlan {
            members v500;
         }
      }
   }
}
ae1 {
   aggregated-ether-options {
      lacp {
         active;
         system-id 00:01:02:03:04:05;
         admin-key 3;
      }
      mc-ae {
         mc-ae-id 3;
```

```
        chassis-id 1;
        mode active-active;
        status-control active;
        init-delay-time 240;
      }
    }
    unit 0 {
      family ethernet-switching {
        interface-mode trunk;
        vlan {
          members v100;
        }
      }
    }
}
vlan {
    unit 100 {
      family inet {
        address 10.1.1.10/8 {
          vrrp-group 1 {
            virtual-address 10.1.1.1;
            priority 200;
            accept-data;
          }
        }
      }
    }
    unit 500 {
      family inet {
        address 10.3.3.1/8;
      }
    }
}
```

```
user@SwitchB# show protocols
iccp {
    local-ip-addr 10.3.3.1;
    peer 10.3.3.2 {
      session-establishment-hold-time 340;
      backup-liveness-detection {
        backup-peer-ip 10.207.64.234;
      }
      liveness-detection {
        minimum-receive-interval 1000;
        transmit-interval {
          minimum-interval 1000;
        }
      }
    }
}
rstp {
    interface ae1.0 {
      edge;
    }
```

```
    mode point-to-point;
    }
    bpdu-block-on-edge;
    }
```

```
user@SwitchB# show multi-chassis
multi-chassis-protection 10.3.3.2 {
    interface ae0;
}
```

```
user@SwitchB# show switch-options
service-id 10;
```

```
user@SwitchB# show vlans
v100 {
    vlan-id 100;
    l3-interface irb.100;
}
v500 {
    vlan-id 500;
    l3-interface irb.500;
}
```

Verification

Verify that the configuration is working properly.

Verifying That ICCP Is Working on Switch A

Purpose Verify that ICCP is running on Switch A.

Action

```
[edit]
user@switch> show iccp
```

```
Redundancy Group Information for peer 10.3.3.1
  TCP Connection       : Established
  Liveliness Detection : Up

 Client Application: MCSNOOPD

 Client Application: eswd
```

Meaning This output shows that the TCP connection between the peers hosting the MC-LAG is up, liveness detection is up, and MCSNOOPD and ESWD client applications are running.

Verifying That ICCP Is Working on Switch B

Purpose Verify that ICCP is running on Switch B.

Action show iccp

```
[edit]
user@switch> show iccp
```

```
Redundancy Group Information for peer 10.3.3.2
  TCP Connection       : Established
  Liveliness Detection : Up

 Client Application: MCSNOOPD

 Client Application: eswd
```

Meaning This output shows that the TCP connection between the peers hosting the MC-LAG is up, liveness detection is up, and MCSNOOPD and ESWD client applications are running.

Verifying That LACP Is Active on Switch A

Purpose Verify that LACP is active on Switch A.

Action

```
[edit]
user@switch> show lacp interfaces
```

```
Aggregated interface: ae1
    LACP state:       Role   Exp   Def  Dist  Col  Syn  Aggr  Timeout  Activity
      xe-0/0/46      Actor    No    No   Yes  Yes  Yes   Yes     Fast    Active
      xe-0/0/46    Partner    No    No   Yes  Yes  Yes   Yes     Fast    Active
    LACP protocol:        Receive State  Transmit State          Mux State
      xe-0/0/46                 Current   Fast periodic Collecting distributing
```

Meaning This output shows that Switch A is participating in LACP negotiation.

Verifying That LACP Is Active on Switch B

Purpose Verify that LACP is active on Switch B.

Action

```
[edit]
user@switch> show lacp interfaces
```

```
Aggregated interface: ae1
    LACP state:       Role   Exp   Def  Dist  Col  Syn  Aggr  Timeout  Activity
      xe-0/0/44      Actor    No    No   Yes  Yes  Yes   Yes     Fast    Active
      xe-0/0/44    Partner    No    No   Yes  Yes  Yes   Yes     Fast    Active
    LACP protocol:        Receive State  Transmit State          Mux State
      xe-0/0/44                 Current   Fast periodic Collecting distributing
```

Meaning This output shows that Switch B is participating in LACP negotiation.

Verifying That the multichassis aggregated Ethernet and ICL-PL Interfaces Are Up on Switch A

Purpose Verify that the multichassis aggregated Ethernet and Inter-chassis Link Protection (ICL-PL) interfaces are up on Switch A.

Action

```
[edit]
user@switch> show interfaces mc-ae
```

```
Member Link                  : ae1
 Current State Machine's State: mcae active state
 Local Status                 : active
 Local State                  : up
 Peer Status                  : active
 Peer State                   : up
     Logical Interface        : ae1.0
     Topology Type            : bridge
     Local State              : up
     Peer State               : up
     Peer Ip/MCP/State        : 10.3.3.1 ae0.0 up
```

Meaning This output shows that the multichassis aggregated Ethernet and ICL-PL on Switch A is up and active.

Verifying That the Multichassis Aggregated Ethernet and ICL-PL Interfaces Are Up on Switch B

Purpose Verify that the multichassis aggregated Ethernet and ICL-PL interfaces are up on Switch B.

Action

```
[edit]
user@switch> show interfaces mc-ae
```

```
Member Link                  : ae1
 Current State Machine's State: mcae active state
 Local Status                 : active
 Local State                  : up
 Peer Status                  : active
 Peer State                   : up
     Logical Interface        : ae1.0
     Topology Type            : bridge
     Local State              : up
     Peer State               : up
     Peer Ip/MCP/State        : 10.3.3.2 ae0.0 up
```

Meaning This output shows that the multichassis aggregated Ethernet and ICL-PL interface on Switch B is up and active.

Verifying that MAC Learning Is Occurring on Switch A

Purpose Verify that MAC learning is working on Switch A.

Action

```
[edit]
user@switch> show ethernet-switching table
```

```
Ethernet-switching table: 6 entries, 1 learned, 0 persistent entriesC
  VLAN              MAC address       Type         Age Interfaces
  v100              *                 Flood          - All-members
  v100              00:00:5e:00:01:01 Static         - Router
  v100              78:fe:3d:5a:07:42 Static         - Router
  v100              78:fe:3d:5b:ad:c2 Learn(R)        0 ae0.0
  v500              *                 Flood          - All-members
  v500              78:fe:3d:5a:07:42 Static         - Router
```

Meaning The output shows two static MAC address in VLAN v100 and one static MAC address in VLAN v500. These addresses belong to the Layer 3 RVI addresses on both Switch A and Switch B that you configured in the MC-LAG. The ICL-PL interface configured on the VRRP master member learned the VLAN v100 Learn (R) MAC address of the VRRP backup member.

Verifying that MAC Learning Is Occurring on Switch B

Purpose Verify that MAC learning is working on Switch B.

Action

```
[edit]
user@switch> show ethernet-switching table
```

```
Ethernet-switching table: 7 entries, 1 learned, 0 persistent entries
  VLAN              MAC address       Type         Age Interfaces
  v100              *                 Flood          - All-members
  v100              00:00:5e:00:01:01 Static         - Router
  v100              78:fe:3d:5a:07:42 Learn(R)        0 ae0.0
  v100              78:fe:3d:5b:ad:c2 Static         - Router
  v200              78:fe:3d:5b:ad:c2 Static         - Router
  v500              *                 Flood          - All-members
  v500              78:fe:3d:5b:ad:c2 Static         - Router
```

Meaning The output shows two static MAC address in VLAN v100 and one static MAC address in VLAN v500. These addresses belong to the Layer 3 RVI addresses on both Switch A and Switch B that you configured in the MC-LAG. The ICL-PL interface configured on the VRRP backup member learned the VLAN v100 Learn (R) MAC address of the VRRP master member.

Verifying that Switch A is the Master in the VRRP Group

Purpose Verify that Switch A is the master member in the VRRP group.

Action

```
[edit]
user@switch> show vrrp
```

```
Interface     State       Group   VR state VR Mode   Timer    Type   Address
vlan.100      up              1   master   Active      A  0.605 lcl    10.1.1.11

                                                                vip    10.1.1.1
```

Meaning The output shows that Switch A is the master member in the VRRP group.

Verifying that Switch B is the Backup Member in the VRRP Group

Purpose Verify that Switch B is the backup member in the VRRP group.

Action

```
[edit]
user@switch> show vrrp
```

```
Interface     State       Group   VR state VR Mode   Timer    Type   Address
vlan.100      up              1   backup   Active      A  0.605 lcl    10.1.1.10

                                                                vip    10.1.1.1
```

Meaning The output shows that Switch B is the backup member in the VRRP group.

Verifying that the Virtual IP Address is Attached to an Individual Address on Switch A

Action

```
[edit]
user@switch# run show interfaces terse vlan
```

```
Interface               Admin Link Proto    Local                 Remote
vlan                    up    up
vlan.100                up    up   inet     10.1.1.1/8
                                            10.1.1.11/8
vlan.500                up    up   inet     10.3.3.2/8
```

Meaning The output shows that the virtual IP address (10.1.1.1/8) is bound to the individual IP address (10.1.1.11/8) on Switch A.

Verifying that the Virtual IP Address is Attached to an Individual Address on Switch B

Action

```
[edit]
user@switch# run show interfaces terse vlan
```

```
Interface               Admin Link Proto    Local                 Remote
vlan                    up    up
vlan.100                up    up   inet     10.1.1.1/8
                                            10.1.1.10/8
vlan.500                up    up   inet     10.3.3.1/8
```

Meaning The output shows that the virtual IP address (10.1.1.1/8) is bound to the individual IP address (10.1.1.10/8) on Switch B.

Troubleshooting

Troubleshooting a LAG That Is Down

Problem The **show interfaces terse** command shows that the MC-LAG is **down**.

Solution Check the following:

1. Verify that there is no configuration mismatch.
2. Verify that all member ports are up.
3. Verify that the MC-LAG is part of family Ethernet switching (Layer 2 LAG).
4. Verify that the MC-LAG member is connected to the correct MC-LAG member at the other end.

Example: Configuring Multichassis Link Aggregation for Layer 3 Multicast Using VRRP

NOTE: Multichassis link aggregation (MC-LAG) is supported on QFX3500 and QFX3600 standalone switches running the original CLI and QFX3500, QFX3600, QFX5100, EX4600, and QFX10000 standalone switches running Enhanced Layer 2 Software (ELS).

There are two methods for enabling Layer 3 multicast functionality across a multichassis link aggregation group (MC-LAG) to control traffic. You can choose either to synchronize the MAC addresses for the Layer 3 interfaces of the switches participating in the MC-LAG , or you can configure Virtual Router Redundancy Protocol (VRRP), but you cannot configure both at the same time. Because RVI interfaces share the same MAC address, if you enable MAC address synchronization, packets received on an MC-LAG peer with a destination MAC address that is the same as that of the peer's IRB MAC address will

not be forwarded. The procedure to configure VRRP for use in a Layer 3 multicast MC-LAG is included in this example.

Requirements

This example uses the following hardware and software components:

- Junos OS Release 12.3 or later for the QFX3500 and QFX3600 standalone switches and Junos OS Release 13.2X51-D10 or later for the QFX5100 and EX4600 standalone switches, and Junos OS Release 15.1X53-D10 or later for the standalone QFX10000 switches.
- Two QFX3500 or QFX3600 standalone switches, two QFX5100 standalone switches, two EX4600, or two QFX10002 standalone switches.

Before you configure an MC-LAG for Layer 3 multicast using VRRP, be sure that you understand how to:

- Configure aggregated Ethernet interfaces on a switch. See *Example: Configuring Link Aggregation Between a QFX Series Product and an Aggregation Switch*.
- Configure the Link Aggregation Control Protocol (LACP) on aggregated Ethernet interfaces on a switch. See *Example: Configuring Link Aggregation with LACP Between a QFX Series Product and an Aggregation Switch*.

Overview

In this example, you configure two MC-LAGs across two switches, consisting of two aggregated Ethernet interfaces (ae1 and ae2). To support the MC-LAG, create a third aggregated Ethernet interface (ae0) for the interchassis control link-protection link (ICL-PL). Configure a multichassis protection link for the ICL-PL, the Inter-Chassis Control Protocol (ICCP) for the peers hosting the MC-LAG, and Layer 3 connectivity between MC-LAG peers.

NOTE: Layer 3 connectivity is required for ICCP.

NOTE: On QFX5100 and QFX10000 switches, if you try to configure both VRRP over IRB and MAC synchronization, you will receive a commit error.

To complete the configuration, enable VRRP by completing the following steps for each MC-LAG:

1. Create a routed VLAN interface (RVI).
2. Create a VRRP group and assign a virtual IP address that is shared between each switch in the VRRP group.
3. Enable a member of a VRRP group to accept all packets destined for the virtual IP address if it is the master in the VRRP group.
4. Configure Layer 3 connectivity between the VRRP groups.

Topology

The topology used in this example consists of two switches hosting two MC-LAGs—ae1 and ae2. The two switches are connected to a multicast source (Server 1) over the MC-LAG ae1, and a multicast receiver (Server 2) over the MC-LAG ae2. Figure 34 on page 277 shows the topology of this example.

Figure 34: Configuring a Multichassis LAG for Layer 3 Multicast Using VRRP

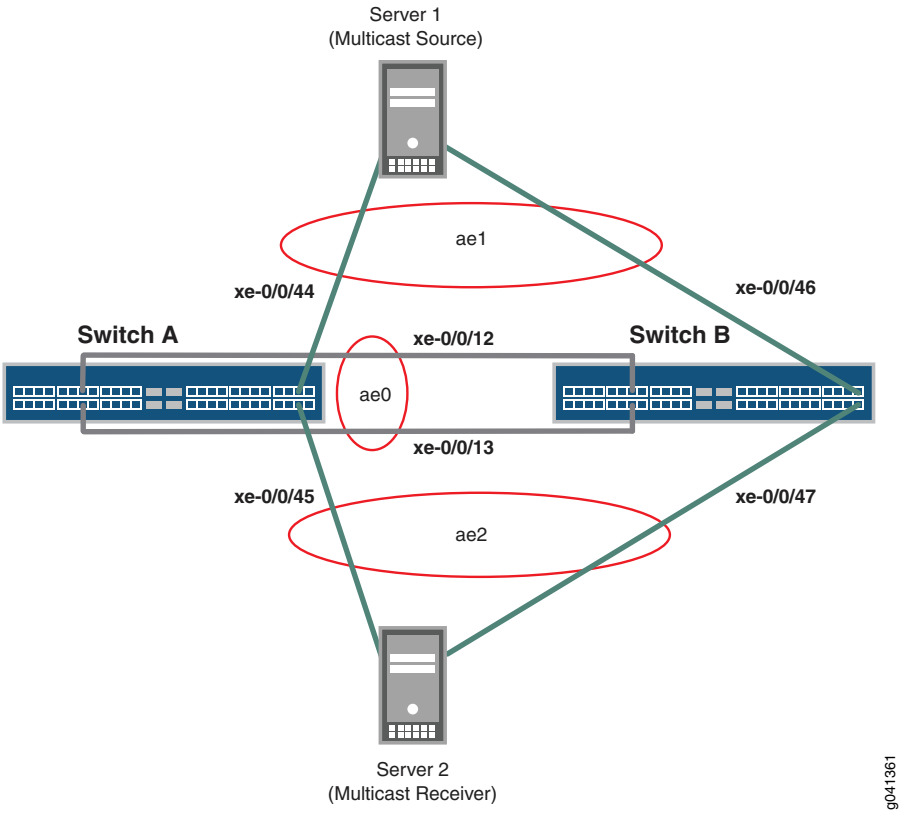

Table 11 on page 278 details the topology used in this configuration example.

Table 11: Components of the Topology for Configuring a Multichassis LAG for Layer 3 Multicast Using VRRP

Hostname	Base Hardware	Multichassis Link Aggregation Group
Switch A Switch B	QFX3500, QFX3600, EX4600, QFX5100, or QFX10000 standalone switch QFX3500, QFX3600, EX4600, QFX5100, or QFX10000 standalone switch	• ae0 is configured as an aggregated Ethernet interface, and is used as an ICL-PL. The following two interfaces are part of ae0: xe-0/0/12 and xe-0/0/13 on Switch A and xe-0/0/12 and xe-0/0/13 on Switch B. • ae1 is configured as an MC-LAG for the multicast source (Server 1), and the following two interfaces are part of ae1: xe-0/0/44 on Switch A and xe-0/0/46 on Switch B. • ae2 is configured as an MC-LAG for the multicast receiver (Server 2), and the following two interfaces are part of ae2: xe-0/0/45 on Switch A and xe-0/0/47 on Switch B.

Configuration

CLI Quick Configuration

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

NOTE: This example shows how to configure MC-LAG using both the original CLI and Enhanced Layer 2 Software (ELS).

In ELS, there are three different statements and one different option from the original CLI:

- The **port-mode** statement in the **[edit interfaces** *interface-name* **unit** *number* **family ethernet-switching]** hierarchy is not supported. Use the **interface-mode** statement instead.
- The **vlan** statement in the **[edit interfaces** *interface-name***]** hierarchy is not supported. Use the **irb** statement instead.
- The **vlan.***logical-interface-number* hierarchy in the **[edit vlans** *vlan-name* **l3-interface]** option is not supported. Use the **irb.***logical-interface-number* option instead.
- The **service-id** statement in the **[edit switch-options]** hierarchy is required in the ELS CLI.

Switch A—Original CLI

```
set chassis aggregated-devices ethernet device-count 3
set interfaces xe-0/0/12 ether-options 802.3ad ae0
```

```
set interfaces xe-0/0/13 ether-options 802.3ad ae0
set interfaces xe-0/0/44 ether-options 802.3ad ae1
set interfaces xe-0/0/45 ether-options 802.3ad ae2
set interfaces ae0 unit 0 family ethernet-switching port-mode trunk
set interfaces ae0 unit 0 family ethernet-switching vlan members v500
set interfaces ae1 aggregated-ether-options lacp active
set interfaces ae1 aggregated-ether-options lacp system-id 00:01:02:03:04:05
set interfaces ae1 aggregated-ether-options lacp admin-key 3
set interfaces ae1 aggregated-ether-options mc-ae mc-ae-id 3
set interfaces ae1 aggregated-ether-options mc-ae chassis-id 0
set interfaces ae1 aggregated-ether-options mc-ae mode active-active
set interfaces ae1 aggregated-ether-options mc-ae status-control active
set interfaces ae1 aggregated-ether-options mc-ae init-delay-time 240
set interfaces ae1 unit 0 family ethernet-switching port-mode trunk
set interfaces ae1 unit 0 family ethernet-switching vlan members v100
set interfaces ae2 aggregated-ether-options lacp active
set interfaces ae2 aggregated-ether-options lacp system-id 00:01:02:03:04:06
set interfaces ae2 aggregated-ether-options lacp admin-key 3
set interfaces ae2 aggregated-ether-options mc-ae mc-ae-id 4
set interfaces ae2 aggregated-ether-options mc-ae chassis-id 0
set interfaces ae2 aggregated-ether-options mc-ae mode active-active
set interfaces ae2 aggregated-ether-options mc-ae status-control active
set interfaces ae2 aggregated-ether-options mc-ae init-delay-time 240
set interfaces ae2 unit 0 family ethernet-switching port-mode trunk
set interfaces ae2 unit 0 family ethernet-switching vlan members v200
set interfaces vlan unit 100 family inet address 10.1.1.11/8 vrrp-group 1 virtual-address
  10.1.1.1
set interfaces vlan unit 100 family inet address 10.1.1.11/8 vrrp-group 1 priority 200
set interfaces vlan unit 100 family inet address 10.1.1.11/8 vrrp-group 1 accept-data
set interfaces vlan unit 200 family inet address 10.1.1.21/8 vrrp-group 2 virtual-address
  10.1.1.2
set interfaces vlan unit 200 family inet address 10.1.1.21/8 vrrp-group 2 priority 200
set interfaces vlan unit 200 family inet address 10.1.1.21/8 vrrp-group 2 accept-data
set interfaces vlan unit 500 family inet address 10.3.3.2/8
set vlans v100 vlan-id 100
set vlans v100 l3-interface vlan.100
set vlans v200 vlan-id 200
set vlans v200 l3-interface vlan.200
set vlans v500 vlan-id 500
set vlans v500 l3-interface vlan.500
set protocols iccp local-ip-addr 10.3.3.2
set protocols iccp peer 10.3.3.1 session-establishment-hold-time 340
set protocols iccp peer 10.3.3.1 backup-liveness-detection backup-peer-ip 10.207.64.233
set protocols iccp peer 10.3.3.1 liveness-detection minimum-receive-interval 1000
set protocols iccp peer 10.3.3.1 liveness-detection transmit-interval minimum-interval
  1000
set protocols igmp-snooping vlan all
set protocols ospf area 0.0.0.0 interface vlan.100 bfd-liveness-detection
  minimum-receive-interval 700
set protocols ospf area 0.0.0.0 interface vlan.100 bfd-liveness-detection transmit-interval
  minimum-interval 350
set protocols ospf area 0.0.0.0 interface vlan.100 bfd-liveness-detection transmit-interval
  threshold 500
set protocols ospf area 0.0.0.0 interface vlan.200 bfd-liveness-detection
  minimum-receive-interval 700
```

```
set protocols ospf area 0.0.0.0 interface vlan.200 bfd-liveness-detection transmit-interval
  minimum-interval 350
set protocols ospf area 0.0.0.0 interface vlan.200 bfd-liveness-detection transmit-interval
  threshold 500
set protocols pim rp static address 10.0.0.3 group-ranges 233.252.0.0/8
set protocols pim interface vlan.100 priority 200
set protocols pim interface vlan.100 dual-dr
set protocols pim interface vlan.100 bfd-liveness-detection minimum-receive-interval
  700
set protocols pim interface vlan.100 bfd-liveness-detection transmit-interval
  minimum-interval 350
set protocols pim interface vlan.100 bfd-liveness-detection transmit-interval threshold
  500
set protocols pim interface vlan.200 priority 600
set protocols pim interface vlan.200 dual-dr
set protocols pim interface vlan.200 bfd-liveness-detection minimum-receive-interval
  700
set protocols pim interface vlan.200 bfd-liveness-detection transmit-interval
  minimum-interval 350
set protocols pim interface vlan.200 bfd-liveness-detection transmit-interval threshold
  500
set protocols rstp interface ae0.0 disable
set protocols rstp interface ae1.0 edge
set protocols rstp interface ae2.0 edge
set protocols rstp interface all mode point-to-point
set protocols rstp bpdu-block-on-edge
set multi-chassis multi-chassis-protection 10.3.3.1 interface ae0
```

Switch A—ELS

```
set chassis aggregated-devices ethernet device-count 3
set interfaces xe-0/0/12 ether-options 802.3ad ae0
set interfaces xe-0/0/13 ether-options 802.3ad ae0
set interfaces xe-0/0/44 ether-options 802.3ad ae1
set interfaces xe-0/0/45 ether-options 802.3ad ae2
set interfaces ae0 unit 0 family ethernet-switching interface-mode trunk
set interfaces ae0 unit 0 family ethernet-switching vlan members v500
set interfaces ae1 aggregated-ether-options lacp active
set interfaces ae1 aggregated-ether-options lacp system-id 00:01:02:03:04:05
set interfaces ae1 aggregated-ether-options lacp admin-key 3
set interfaces ae1 aggregated-ether-options mc-ae mc-ae-id 3
set interfaces ae1 aggregated-ether-options mc-ae chassis-id 0
set interfaces ae1 aggregated-ether-options mc-ae mode active-active
set interfaces ae1 aggregated-ether-options mc-ae status-control active
set interfaces ae1 aggregated-ether-options mc-ae init-delay-time 240
set interfaces ae1 unit 0 family ethernet-switching interface-mode trunk
set interfaces ae1 unit 0 family ethernet-switching vlan members v100
set interfaces ae2 aggregated-ether-options lacp active
set interfaces ae2 aggregated-ether-options lacp system-id 00:01:02:03:04:06
set interfaces ae2 aggregated-ether-options lacp admin-key 3
set interfaces ae2 aggregated-ether-options mc-ae mc-ae-id 4
set interfaces ae2 aggregated-ether-options mc-ae chassis-id 0
set interfaces ae2 aggregated-ether-options mc-ae mode active-active
set interfaces ae2 aggregated-ether-options mc-ae status-control active
set interfaces ae2 aggregated-ether-options mc-ae init-delay-time 240
```

```
set interfaces ae2 unit 0 family ethernet-switching interface-mode trunk
set interfaces ae2 unit 0 family ethernet-switching vlan members v200
set interfaces irb unit 100 family inet address 10.1.1.11/8 vrrp-group 1 virtual-address 10.1.1.1
set interfaces irb unit 100 family inet address 10.1.1.11/8 vrrp-group 1 priority 200
set interfaces irb unit 100 family inet address 10.1.1.11/8 vrrp-group 1 accept-data
set interfaces irb unit 200 family inet address 10.1.1.21/8 vrrp-group 2 virtual-address
  10.1.1.2
set interfaces irb unit 200 family inet address 10.1.1.21/8 vrrp-group 2 priority 200
set interfaces irb unit 200 family inet address 10.1.1.21/8 vrrp-group 2 accept-data
set interfaces irb unit 500 family inet address 10.3.3.2/8
set vlans v100 vlan-id 100
set vlans v100 l3-interface irb.100
set vlans v200 vlan-id 200
set vlans v200 l3-interface irb.200
set vlans v500 vlan-id 500
set vlans v500 l3-interface irb.500
set protocols iccp local-ip-addr 10.3.3.2
set protocols iccp peer 10.3.3.1 session-establishment-hold-time 340
set protocols iccp peer 10.3.3.1 backup-liveness-detection backup-peer-ip 10.207.64.233
set protocols iccp peer 10.3.3.1 liveness-detection minimum-receive-interval 1000
set protocols iccp peer 10.3.3.1 liveness-detection transmit-interval minimum-interval
  1000
set protocols igmp-snooping vlan all
set protocols ospf area 0.0.0.0 interface vlan.100 bfd-liveness-detection
  minimum-receive-interval 700
set protocols ospf area 0.0.0.0 interface vlan.100 bfd-liveness-detection transmit-interval
  minimum-interval 350
set protocols ospf area 0.0.0.0 interface vlan.100 bfd-liveness-detection transmit-interval
  threshold 500
set protocols ospf area 0.0.0.0 interface vlan.200 bfd-liveness-detection
  minimum-receive-interval 700
set protocols ospf area 0.0.0.0 interface vlan.200 bfd-liveness-detection transmit-interval
  minimum-interval 350
set protocols ospf area 0.0.0.0 interface vlan.200 bfd-liveness-detection transmit-interval
  threshold 500
set protocols pim rp static address 10.0.0.3 group-ranges 233.252.0.0/8
set protocols pim interface vlan.100 priority 200
set protocols pim interface vlan.100 dual-dr
set protocols pim interface vlan.100 bfd-liveness-detection minimum-receive-interval
  700
set protocols pim interface vlan.100 bfd-liveness-detection transmit-interval
  minimum-interval 350
set protocols pim interface vlan.100 bfd-liveness-detection transmit-interval threshold
  500
set protocols pim interface vlan.200 priority 600
set protocols pim interface vlan.200 dual-dr
set protocols pim interface vlan.200 bfd-liveness-detection minimum-receive-interval
  700
set protocols pim interface vlan.200 bfd-liveness-detection transmit-interval
  minimum-interval 350
set protocols pim interface vlan.200 bfd-liveness-detection transmit-interval threshold
  500
set protocols rstp interface ae1.0 edge
set protocols rstp interface ae2.0 edge
set protocols rstp interface ae1.0 mode point-to-point
```

```
set protocols rstp bpdu-block-on-edge
set multi-chassis multi-chassis-protection 10.3.3.1 interface ae0
set switch-options service-id 10
```

Switch B—Original CLI:

```
set chassis aggregated-devices ethernet device-count 2
set interfaces xe-0/0/12 ether-options 802.3ad ae0
set interfaces xe-0/0/13 ether-options 802.3ad ae0
set interfaces xe-0/0/46 ether-options 802.3ad ae1
set interfaces xe-0/0/47 ether-options 802.3ad ae2
set interfaces ae0 unit 0 family ethernet-switching port-mode trunk
set interfaces ae0 unit 0 family ethernet-switching vlan members v500
set interfaces ae1 aggregated-ether-options lacp active
set interfaces ae1 aggregated-ether-options lacp system-id 00:01:02:03:04:05
set interfaces ae1 aggregated-ether-options lacp admin-key 3
set interfaces ae1 aggregated-ether-options mc-ae mc-ae-id 3
set interfaces ae1 aggregated-ether-options mc-ae chassis-id 1
set interfaces ae1 aggregated-ether-options mc-ae mode active-active
set interfaces ae1 aggregated-ether-options mc-ae status-control standby
set interfaces ae1 aggregated-ether-options mc-ae init-delay-time 240
set interfaces ae1 unit 0 family ethernet-switching port-mode trunk
set interfaces ae1 unit 0 family ethernet-switching vlan members v100
set interfaces ae2 aggregated-ether-options lacp active
set interfaces ae2 aggregated-ether-options lacp system-id 00:01:02:03:04:06
set interfaces ae2 aggregated-ether-options lacp admin-key 3
set interfaces ae2 aggregated-ether-options mc-ae mc-ae-id 4
set interfaces ae2 aggregated-ether-options mc-ae chassis-id 1
set interfaces ae2 aggregated-ether-options mc-ae mode active-active
set interfaces ae2 aggregated-ether-options mc-ae status-control standby
set interfaces ae2 aggregated-ether-options mc-ae init-delay-time 240
set interfaces ae2 unit 0 family ethernet-switching port-mode trunk
set interfaces ae2 unit 0 family ethernet-switching vlan members v200
set interfaces vlan unit 100 family inet address 10.1.1.10/8 vrrp-group 1 virtual-address
  10.1.1.1
set interfaces vlan unit 100 family inet address 10.1.1.10/8 vrrp-group 1 priority 150
set interfaces vlan unit 100 family inet address 10.1.1.10/8 vrrp-group 1 accept-data
set interfaces vlan unit 200 family inet address 10.1.1.20/8 vrrp-group 2 virtual-address
  10.1.1.2
set interfaces vlan unit 200 family inet address 10.1.1.20/8 vrrp-group 2 priority 150
set interfaces vlan unit 200 family inet address 10.1.1.20/8 vrrp-group 2 accept-data
set interfaces vlan unit 500 family inet address 10.3.3.1/8
set vlans v100 vlan-id 100
set vlans v100 l3-interface vlan.100
set vlans v200 vlan-id 200
set vlans v200 l3-interface vlan.200
set vlans v500 vlan-id 500
set vlans v500 l3-interface vlan.500
set protocols iccp local-ip-addr 10.3.3.1
set protocols iccp peer 10.3.3.2 session-establishment-hold-time 340
set protocols iccp peer 10.3.3.2 backup-liveness-detection backup-peer-ip 10.207.64.234
set protocols iccp peer 10.3.3.2 liveness-detection minimum-receive-interval 1000
set protocols iccp peer 10.3.3.2 liveness-detection transmit-interval minimum-interval
  1000
set protocols igmp-snooping vlan all
```

```
set protocols ospf area 0.0.0.0 interface vlan.100 bfd-liveness-detection
  minimum-receive-interval 700
set protocols ospf area 0.0.0.0 interface vlan.100 bfd-liveness-detection transmit-interval
  minimum-interval 350
set protocols ospf area 0.0.0.0 interface vlan.100 bfd-liveness-detection transmit-interval
  threshold 500
set protocols ospf area 0.0.0.0 interface vlan.200 bfd-liveness-detection
  minimum-receive-interval 700
set protocols ospf area 0.0.0.0 interface vlan.200 bfd-liveness-detection transmit-interval
  minimum-interval 350
set protocols ospf area 0.0.0.0 interface vlan.200 bfd-liveness-detection transmit-interval
  threshold 500
set protocols pim rp static address 10.0.0.3 group-ranges 233.252.0.0/8
set protocols pim interface vlan.100 priority 100
set protocols pim interface vlan.100 dual-dr
set protocols pim interface vlan.100 bfd-liveness-detection minimum-receive-interval
  700
set protocols pim interface vlan.100 bfd-liveness-detection transmit-interval
  minimum-interval 350
set protocols pim interface vlan.100 bfd-liveness-detection transmit-interval threshold
  500
set protocols pim interface vlan.200 priority 500
set protocols pim interface vlan.200 dual-dr
set protocols pim interface vlan.200 bfd-liveness-detection minimum-receive-interval
  700
set protocols pim interface vlan.200 bfd-liveness-detection transmit-interval
  minimum-interval 350
set protocols pim interface vlan.200 bfd-liveness-detection transmit-interval threshold
  500
set protocols rstp interface ae0.0 disable
set protocols rstp interface ae1.0 edge
set protocols rstp interface ae2.0 edge
set protocols rstp interface all mode point-to-point
set protocols rstp bpdu-block-on-edge
set multi-chassis multi-chassis-protection 10.3.3.2 interface ae0
```

Switch B—ELS

```
set chassis aggregated-devices ethernet device-count 2
set interfaces xe-0/0/12 ether-options 802.3ad ae0
set interfaces xe-0/0/13 ether-options 802.3ad ae0
set interfaces xe-0/0/46 ether-options 802.3ad ae1
set interfaces xe-0/0/47 ether-options 802.3ad ae2
set interfaces ae0 unit 0 family ethernet-switching interface-mode trunk
set interfaces ae0 unit 0 family ethernet-switching vlan members v500
set interfaces ae1 aggregated-ether-options lacp active
set interfaces ae1 aggregated-ether-options lacp system-id 00:01:02:03:04:05
set interfaces ae1 aggregated-ether-options lacp admin-key 3
set interfaces ae1 aggregated-ether-options mc-ae mc-ae-id 3
set interfaces ae1 aggregated-ether-options mc-ae chassis-id 1
set interfaces ae1 aggregated-ether-options mc-ae mode active-active
set interfaces ae1 aggregated-ether-options mc-ae status-control standby
set interfaces ae1 aggregated-ether-options mc-ae init-delay-time 240
set interfaces ae1 unit 0 family ethernet-switching interface-mode trunk
set interfaces ae1 unit 0 family ethernet-switching vlan members v100
```

```
set interfaces ae2 aggregated-ether-options lacp active
set interfaces ae2 aggregated-ether-options lacp system-id 00:01:02:03:04:06
set interfaces ae2 aggregated-ether-options lacp admin-key 3
set interfaces ae2 aggregated-ether-options mc-ae mc-ae-id 4
set interfaces ae2 aggregated-ether-options mc-ae chassis-id 1
set interfaces ae2 aggregated-ether-options mc-ae mode active-active
set interfaces ae2 aggregated-ether-options mc-ae status-control standby
set interfaces ae2 aggregated-ether-options mc-ae init-delay-time 240
set interfaces ae2 unit 0 family ethernet-switching interface-mode trunk
set interfaces ae2 unit 0 family ethernet-switching vlan members v200
set interfaces irb unit 100 family inet address 10.1.1.10/8 vrrp-group 1 virtual-address 10.1.1.1
set interfaces irb unit 100 family inet address 10.1.1.10/8 vrrp-group 1 priority 150
set interfaces irb unit 100 family inet address 10.1.1.10/8 vrrp-group 1 accept-data
set interfaces irb unit 200 family inet address 10.1.1.20/8 vrrp-group 2 virtual-address
  10.1.1.2
set interfaces irb unit 200 family inet address 10.1.1.20/8 vrrp-group 2 priority 150
set interfaces irb unit 200 family inet address 10.1.1.20/8 vrrp-group 2 accept-data
set interfaces irb unit 500 family inet address 10.3.3.1/8
set vlans v100 vlan-id 100
set vlans v100 l3-interface irb.100
set vlans v200 vlan-id 200
set vlans v200 l3-interface irb.200
set vlans v500 vlan-id 500
set vlans v500 l3-interface irb.500
set protocols iccp local-ip-addr 10.3.3.1
set protocols iccp peer 10.3.3.2 session-establishment-hold-time 340
set protocols iccp peer 10.3.3.2 backup-liveness-detection backup-peer-ip 10.207.64.234
set protocols iccp peer 10.3.3.2 liveness-detection minimum-receive-interval 1000
set protocols iccp peer 10.3.3.2 liveness-detection transmit-interval minimum-interval
  1000
set protocols igmp-snooping vlan all
set protocols ospf area 0.0.0.0 interface vlan.100 bfd-liveness-detection
  minimum-receive-interval 700
set protocols ospf area 0.0.0.0 interface vlan.100 bfd-liveness-detection transmit-interval
  minimum-interval 350
set protocols ospf area 0.0.0.0 interface vlan.100 bfd-liveness-detection transmit-interval
  threshold 500
set protocols ospf area 0.0.0.0 interface vlan.200 bfd-liveness-detection
  minimum-receive-interval 700
set protocols ospf area 0.0.0.0 interface vlan.200 bfd-liveness-detection transmit-interval
  minimum-interval 350
set protocols ospf area 0.0.0.0 interface vlan.200 bfd-liveness-detection transmit-interval
  threshold 500
set protocols pim rp static address 10.0.0.3 group-ranges 233.252.0.0/8
set protocols pim interface vlan.100 priority 100
set protocols pim interface vlan.100 dual-dr
set protocols pim interface vlan.100 bfd-liveness-detection minimum-receive-interval
  700
set protocols pim interface vlan.100 bfd-liveness-detection transmit-interval
  minimum-interval 350
set protocols pim interface vlan.100 bfd-liveness-detection transmit-interval threshold
  500
set protocols pim interface vlan.200 priority 500
set protocols pim interface vlan.200 dual-dr
```

```
set protocols pim interface vlan.200 bfd-liveness-detection minimum-receive-interval
  700
set protocols pim interface vlan.200 bfd-liveness-detection transmit-interval
  minimum-interval 350
set protocols pim interface vlan.200 bfd-liveness-detection transmit-interval threshold
  500
set protocols rstp interface ae1.0 edge
set protocols rstp interface ae2.0 edge
set protocols rstp interface ae1.0 mode point-to-point
set protocols rstp bpdu-block-on-edge
set multi-chassis multi-chassis-protection 10.3.3.2 interface ae0
set switch-options service-id 10
```

Configuring MC-LAG for Layer 3 Multicast Using VRRP on Two Switches

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 enable multichassis protection link between MC-LAG peers:

1. Configure the number of LAGs on both Switch A and Switch B.

```
[edit chassis]
user@switch# set aggregated-devices ethernet device-count 3
```

2. Add member interfaces to the aggregated Ethernet interfaces on both Switch A and Switch B.

```
[edit interfaces]
user@switch# set xe-0/0/12 ether-options 802.3ad ae0
user@switch# set xe-0/0/13 ether-options 802.3ad ae0
```

```
[edit interfaces]
user@switch# set xe-0/0/44 ether-options 802.3ad ae1
user@switch# set xe-0/0/45 ether-options 802.3ad ae2
```

```
[edit interfaces]
user@switch# set xe-0/0/46 ether-options 802.3ad ae1
user@switch# set xe-0/0/47 ether-options 802.3ad ae2
```

3. Configure ae0 as the trunk interface between Switch A and Switch B.

```
[edit interfaces]
user@switch# set ae0 unit 0 family ethernet-switching port-mode trunk
```

```
[edit interfaces]
user@switch# set ae0 unit 0 family ethernet-switching interface-mode trunk
```

4. Configure ae0 as the multichassis protection link between Switch A and Switch B.

```
[edit]
user@switch# set multi-chassis multi-chassis-protection 10.3.3.1 interface ae0
```

```
[edit]
user@switch# set multi-chassis multi-chassis-protection 10.3.3.2 interface ae0
```

Step-by-Step Procedure

To enable ICCP:

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* .

1. Configure the local IP address to be in the ICCP connection on Switch A and Switch B.

```
[edit protocols]
user@switch# set iccp local-ip-addr 10.3.3.2
```

```
[edit protocols]
user@switch# set iccp local-ip-addr 10.3.3.1
```

2. Configure the peer IP address, minimum receive interval, and minimum transmit interval for a Bidirectional Forwarding Detection (BFD) session for ICCP on Switch A and Switch B.

```
[edit protocols]
user@switch# set iccp peer 10.3.3.1 liveness-detection minimum-receive-interval
  1000
user@switch# set iccp peer 10.3.3.1 liveness-detection transmit-interval
  minimum-interval 1000
```

```
[edit protocols]
user@switch# set iccp peer 10.3.3.2 liveness-detection minimum-receive-interval
  1000
user@switch# set iccp peer 10.3.3.2 liveness-detection transmit-interval
  minimum-interval 1000
```

3. (Optional) Configure the time during which an ICCP connection must succeed between MC-LAG peers on Switch A and Switch B.

NOTE: On QFX and EX Series switches, the default session establishment hold time is 300 seconds. However, the session establishment time must be at least 100 seconds higher than the init delay time. You can optionally update the session establishment time to be 340 seconds and the init delay time to be 240 seconds.

```
[edit protocols]
user@switch# set iccp peer 10.3.3.1 session-establishment-hold-time 340
```

```
[edit protocols]
user@switch# set iccp peer 10.3.3.2 session-establishment-hold-time 340
```

4. (Optional) Configure the backup IP address to be used for backup liveness detection on both Switch A and Switch B.

NOTE: By default, backup liveness detection is not enabled. Configuring a backup IP address helps achieve sub-second traffic loss during an MC-LAG peer reboot.

```
[edit protocols]
user@switch# set iccp peer 10.3.3.1 backup-liveness-detection backup-peer-ip
  10.207.64.233
```

```
[edit protocols]
user@switch# set iccp peer 10.3.3.2 backup-liveness-detection backup-peer-ip
  10.207.64.233
```

5. Configure Layer 3 connectivity between the MC-LAG peers on both Switch A and Switch B.

NOTE: In ELS, use the irb.*logical-interface-number* instead.

```
[edit vlans]
user@switch# set v500 vlan-id 500
user@switch# set v500 l3-interface vlan.500
```

```
[edit vlans]
user@switch# set v500 vlan-id 500
user@switch# set v500 l3-interface irb.500
```

```
[edit interfaces]
user@switch# set ae0 unit 0 family ethernet-switching vlan members v500
```

```
[edit interfaces]
user@switch# set vlan unit 500 family inet address 10.3.3.2/8
```

```
[edit interfaces]
user@switch# set irb unit 500 family inet address 10.3.3.2/8
```

```
[edit interfaces]
user@switch# set vlan unit 500 family inet address 10.3.3.1/8
```

```
[edit interfaces]
user@switch# set irb unit 500 family inet address 10.3.3.1/8
```

Step-by-Step Procedure

To enable the ae1 and ae2 MC-LAG interfaces:

1. Enable LACP on the MC-LAG interfaces on Switch A and Switch B.

NOTE: At least one end needs to be active. The other end can be either active or passive.

```
[edit interfaces]
user@switch# set ae1 aggregated-ether-options lacp active
user@switch# set ae2 aggregated-ether-options lacp active
```

2. Specify the same multichassis aggregated Ethernet identification number for each MC-LAG peer on Switch A and Switch B.

```
[edit interfaces]
user@switch# set ae1 aggregated-ether-options mc-ae mc-ae-id 3
user@switch# set ae2 aggregated-ether-options mc-ae mc-ae-id 4
```

3. Specify the same service ID on Switch A and Switch B.

```
[edit]
set switch-option service-id 10
```

4. Specify a unique chassis ID for the MC-LAG on the MC-LAG peers on Switch A and Switch B.

```
[edit interfaces]
user@switch# set ae1 aggregated-ether-options mc-ae chassis-i 0
user@switch# set ae2 aggregated-ether-options mc-ae chassis-i 0
```

```
[edit interfaces]
user@switch# set ae1 aggregated-ether-options mc-ae chassis-id 1
user@switch# set ae2 aggregated-ether-options mc-ae chassis-id 1
```

5. Specify the operating mode of the MC-LAGs on both Switch A and Switch B.

NOTE: Only active-active mode is supported at this time.

```
[edit interfaces]
user@switch# set ae1 aggregated-ether-options mc-ae mod active-active
user@switch# set ae2 aggregated-ether-options mc-ae mod active-active
```

6. Specify the status control for the MC-LAGs on Switch A and Switch B.

NOTE: You must configure status control on both Switch A and Switch B hosting the MC-LAGs. If one peer is in active mode, the other must be in standby mode.

```
[edit interfaces]
user@switch# set ae1 aggregated-ether-options mc-ae status-control active
user@switch# set ae2 aggregated-ether-options mc-ae status-control active
```

```
[edit interfaces]
user@switch# set ae1 aggregated-ether-options mc-ae status-control standby
user@switch# set ae2 aggregated-ether-options mc-ae status-control standby
```

NOTE: If you configure both nodes as prefer-status-control-active, you must also configure ICCP peering using the peer's loopback address to make sure that the ICCP session does not go down because of physical link failures. Additionally, you must configure backup liveness detection on both of the MC-LAG nodes.

7. Specify the number of seconds by which the bring-up of the MC-LAG interfaces should be deferred after you reboot Switch A or Switch B.

NOTE: The recommended value for maximum VLAN configuration (for example, 4,000 VLANS) is 240 seconds. If IGMP snooping is enabled on all of the VLANs, the recommended value is 420 seconds.

```
[edit interfaces]
user@switch# set ae1 aggregated-ether-options mc-ae init-delay-time 240
user@switch# set ae2 aggregated-ether-options mc-ae init-delay-time 240
```

8. Specify the same LACP system ID for each MC-LAG on Switch A and Switch B.

```
[edit interfaces]
user@switch# set ae1 aggregated-ether-options lacp system-ID 00:01:02:03:04:05
user@switch# set ae2 aggregated-ether-options lacp system-ID 00:01:02:03:04:06
```

9. Specify the same LACP administration key on both Switch A and Switch B.

```
[edit interfaces]
user@switch# set ae1 aggregated-ether-options lacp admin-key 3
user@switch# set ae2 aggregated-ether-options lacp admin-key 3
```

10. Enable a VLAN for each MC-LAG on Switch A and Switch B.

```
[edit vlans]
user@switch# set v100 vlan-id 100
user@switch# set v200 vlan-id 200
```

```
[edit interfaces]
user@switch# set ae1 unit 0 family ethernet-switching vlan members v100
user@switch# set ae2 unit 0 family ethernet-switching vlan members v200
```

11. Configure ae1 and ae2 as trunk interfaces between Switch A and Switch B.

```
[edit interfaces]
user@switch# set ae1 unit 0 family ethernet-switching port-mode trunk
user@switch# set ae2 unit 0 family ethernet-switching port-mode trunk
```

```
[edit interfaces]
user@switch# set ae1 unit 0 family ethernet-switching interface-mode trunk
user@switch# set ae2 unit 0 family ethernet-switching interface-mode trunk
```

Step-by-Step Procedure

To enable VRRP on the MC-LAGs on Switch A and Switch B:

1. Create a routed VLAN interface (RVI) for each MC-LAG, assign a virtual IP address that is shared between each switch in the VRRP groups, and assign an individual IP address for each switch in the VRRP groups.

```
[edit interfaces]
user@switch# set vlan unit 100 family inet address 10.1.1.11/8 vrrp-group 1
  virtual-address 10.1.1.1
user@switch# set vlan unit 200 family inet address 10.1.1.21/8 vrrp-group 2
  virtual-address 10.1.1.2
```

```
[edit interfaces]
user@switch# set irb unit 100 family inet address 10.1.1.11/8 vrrp-group 1
  virtual-address 10.1.1.1
user@switch# set irb unit 200 family inet address 10.1.1.21/8 vrrp-group 2
  virtual-address 10.1.1.2
```

```
[edit interfaces]
user@switch# set vlan unit 100 family inet address 10.1.1.10/8 vrrp-group 1
  virtual-address 10.1.1.1
user@switch# set vlan unit 200 family inet address 10.1.1.20/8 vrrp-group 2
  virtual-address 10.1.1.2
```

```
[edit interfaces]
user@switch# set irb unit 100 family inet address 10.1.1.10/8 vrrp-group 1
  virtual-address 10.1.1.1
user@switch# set irb unit 200 family inet address 10.1.1.20/8 vrrp-group 2
  virtual-address 10.1.1.2
```

2. Assign the priority for each switch in the VRRP groups:

NOTE: The switch configured with the highest priority is the master.

```
[edit interfaces]
user@switch# set vlan unit 100 family inet address 10.1.1.11/8 vrrp-group 1 priority
  200
user@switch# set vlan unit 200 family inet address 10.1.1.21/8 vrrp-group 2 priority
  200
```

```
[edit interfaces]
user@switch# set irb unit 100 family inet address 10.1.1.11/8 vrrp-group 1 priority 200
user@switch# set irb unit 200 family inet address 10.1.1.21/8 vrrp-group 2 priority
  200
```

```
[edit interfaces]
user@switch# set vlan unit 100 family inet address 10.1.1.10/8 vrrp-group 1 priority
  150
user@switch# set vlan unit 200 family inet address 10.1.1.20/8 vrrp-group 2 priority
  150
```

```
[edit interfaces]
user@switch# set irb unit 100 family inet address 10.1.1.10/8 vrrp-group 1 priority
  150
user@switch# set irb unit 200 family inet address 10.1.1.20/8 vrrp-group 2 priority
  150
```

3. Enable the switch to accept all packets destined for the virtual IP address if it is the master in a VRRP group:

```
[edit interfaces]
user@switch# set vlan unit 100 family inet address 10.1.1.11/8 vrrp-group 1
  accept-data
user@switch# set vlan unit 200 family inet address 10.1.1.21/8 vrrp-group 2
  accept-data
```

```
[edit interfaces]
user@switch# set irb unit 100 family inet address 10.1.1.11/8 vrrp-group 1 accept-data
user@switch# set irb unit 200 family inet address 10.1.1.21/8 vrrp-group 2 accept-data
```

```
[edit interfaces]
user@switch# set vlan unit 100 family inet address 10.1.1.10/8 vrrp-group 1
  accept-data
user@switch# set vlan unit 200 family inet address 10.1.1.20/8 vrrp-group 2
  accept-data
```

```
[edit interfaces]
user@switch# set irb unit 100 family inet address 10.1.1.10/8 vrrp-group 1 accept-data
user@switch# set irb unit 200 family inet address 10.1.1.20/8 vrrp-group 2
  accept-data
```

4. Configure Layer 3 connectivity between Switch A and Switch B.

```
[edit interfaces]
user@switch# set v100 l3-interface vlan.100
user@switch# set v200 l3-interface vlan.200
```

```
[edit interfaces]
user@switch# set v100 l3-interface irb.100
user@switch# set v200 l3-interface irb.200
```

Step-by-Step Procedure

To enable IGMP snooping:

1. Enable IGMP snooping for all VLANs on Switch A and Switch B.

```
[edit protocols]
user@switch# set igmp-snooping vlan all
```

Step-by-Step Procedure

To configure OSPF as the Layer 3 protocol:

1. Configure an OSPF area on Switch A and Switch B.

```
[edit protocols ospf]
user@switch# set area 0.0.0.0
```

2. Assign the VLAN interfaces for the MC-LAGs as interfaces to the OSPF area on Switch A and Switch B.

```
[edit protocols ospf area 0.0.0.0]
user@switch# set interface vlan.100
user@switch# set interface vlan.200
```

3. Configure the minimum receive interval, minimum transmit interval, and transmit interval threshold for a Bidirectional Forwarding Detection (BFD) session for the OSPF interfaces on Switch A and Switch B.

NOTE: On a QFX5100 switch, the minimum transmit interval must be 1000 milliseconds or greater. Sub-second timers are not supported in Junos OS 13.2X51-D10 and later. If you configure the minimum transmit interval timer lower than 1000 milliseconds, the state of the MC-LAG can be affected.

```
[edit protocols ospf area 0.0.0.0]
user@switch# set interface vlan.100 bfd-liveness-detection
  minimum-receive-interval 700
user@switch# set interface vlan.100 bfd-liveness-detection transmit-interval
  minimum-interval 350
user@switch# set interface vlan.100 bfd-liveness-detection transmit-interval
  threshold 500
user@switch# set interface vlan.200 bfd-liveness-detection
  minimum-receive-interval 700
user@switch# set interface vlan.200 bfd-liveness-detection transmit-interval
  minimum-interval 350
user@switch# set interface vlan.200 bfd-liveness-detection transmit-interval
  threshold 500
```

Step-by-Step Procedure

To configure PIM as the multicast protocol:

1. Configure a static rendezvous point (RP) address on Switch A and Switch B.

```
[edit protocols pim]
user@switch# set rp static address 10.0.0.3
```

2. Configure the address ranges of the multicast groups for which Switch A and Switch B can be a rendezvous point (RP).

```
[edit protocols pim rp static address 10.0.0.3]
user@switch# set group-ranges 233.252.0.0/8
```

3. Enable PIM on the VLAN interfaces for the MC-LAGs on Switch A and Switch B.

```
[edit protocols pim]
user@switch# set interface vlan.100 dual-dr
user@switch# set interface vlan.200 dual-dr
```

4. Configure each PIM interface's priority for being selected as the designated router (DR) on Switch A and Switch B.

An interface with a higher priority value has a higher probability of being selected as the DR.

```
[edit protocols pim]
user@switch# set interface vlan.100 priority 200
user@switch# set interface vlan.200 priority 600
```

```
[edit protocols pim]
user@switch# set interface vlan.100 priority 100
user@switch# set interface vlan.200 priority 500
```

5. Configure the minimum receive interval, minimum transmit interval, and transmit interval threshold for a Bidirectional Forwarding Detection (BFD) session for the PIM interfaces on Switch A and Switch B.

```
[edit protocols pim]
user@switch# set interface vlan.100 bfd-liveness-detection
  minimum-receive-interval 700
user@switch# set interface vlan.100 bfd-liveness-detection transmit-interval
  minimum-interval 350
user@switch# set interface vlan.100 bfd-liveness-detection transmit-interval
  threshold 500
user@switch# set interface vlan.200 bfd-liveness-detection
  minimum-receive-interval 700
user@switch# set interface vlan.200 bfd-liveness-detection transmit-interval
  minimum-interval 350
user@switch# set interface vlan.200 bfd-liveness-detection transmit-interval
  threshold 500
```

Step-by-Step Procedure

To enable RSTP:

1. Enable RSTP on Switch A and Switch B.

```
[edit protocols rstp]
user@switch# set interface ae1.0 mode point-to-point
```

2. Enable RSTP on Switch B.

```
[edit protocols rstp]
user@switch# set interface ae1.0 mode point-to-point
```

3. Disable RSTP on the ICL-PL interfaces on Switch A and Switch B.

NOTE: This command does not apply on ELS.

```
[edit protocols rstp]
```

```
user@switch# set interface ae0.0 disable
```

4. Configure the MC-LAG interfaces as edge ports on Switch A and Switch B.

NOTE: The ae1 and ae2 interfaces are downstream interfaces. This is why RSTP and bpdu-block-on-edge need to be configured.

```
[edit protocols rstp]
user@switch# set interface ae1.0 edge
user@switch# set interface ae2.0 edge
```

5. Enable BPDU blocking on all interfaces except for the ICL-PL interfaces on Switch A and Switch B.

NOTE: The ae1 and ae2 interfaces are downstream interfaces. This is why RSTP and bpdu-block-on-edge need to be configured.

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

Results

From configuration mode on Switch A, confirm your configuration by entering the **show chassis**, **show interfaces**, **show multi-chassis**, **show protocols**, and **show vlans** commands. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

Switch A—Original CLI:

```
user@SwitchA# show chassis
aggregated-devices {
   ethernet {
      device-count 3;
   }
}
```

```
user@SwitchA# show interfaces
xe-0/0/12 {
   ether-options {
      802.3ad ae0;
   }
}
xe-0/0/13 {
```

```
    ether-options {
      802.3ad ae0;
    }
  }
  xe-0/0/44 {
    ether-options {
      802.3ad ae1;
    }
  }
  xe-0/0/45 {
    ether-options {
      802.3ad ae2;
    }
  }
  ae0 {
    unit 0 {
      family ethernet-switching {
        port-mode trunk;
        vlan {
          members v500;
        }
      }
    }
  }
  ae1 {
    aggregated-ether-options {
      lacp {
        active;
        system-id 00:01:02:03:04:05;
        admin-key 3;
      }
      mc-ae {
        mc-ae-id 3;
        chassis-id 0;
        mode active-active;
        status-control active;
        init-delay-time 240;
      }
    }
    unit 0 {
      family ethernet-switching {
        port-mode trunk;
        vlan {
          members v100;
        }
      }
    }
  }
  ae2 {
    aggregated-ether-options {
      lacp {
        active;
        system-id 00:01:02:03:04:06;
        admin-key 3;
      }
```

```
    mc-ae {
      mc-ae-id 4;
      chassis-id 0;
      mode active-active;
      status-control active;
      init-delay-time 240;
    }
  }
  unit 0 {
    family ethernet-switching {
      port-mode trunk;
      vlan {
        members v200;
      }
    }
  }
}
vlan {
  unit 100 {
    family inet {
      address 10.1.1.11/8 {
        vrrp-group 1 {
          virtual-address 10.1.1.1;
          priority 200;
          accept-data;
        }
      }
    }
  }
  unit 200 {
    family inet {
      address 10.1.1.21/8 {
        vrrp-group 2 {
          virtual-address 10.1.1.2;
          priority 200;
          accept-data;
        }
      }
    }
  }
  unit 500 {
    family inet {
      address 10.3.3.2/8;
    }
  }
}
}
```

```
user@SwitchA# show protocols
ospf {
  area 0.0.0.0 {
    interface vlan.100 {
      bfd-liveness-detection {
        minimum-receive-interval 700;
```

```
            transmit-interval {
              minimum-interval 350;
              threshold 500;
            }
          }
        }
        interface vlan.200 {
          bfd-liveness-detection {
            minimum-receive-interval 700;
            transmit-interval {
              minimum-interval 350;
              threshold 500;
            }
          }
        }
      }
    }
    pim {
      rp {
        static {
          address 10.0.0.3 {
            group-ranges {
              233.252.0.0/8;
            }
          }
        }
      }
      interface vlan.100 {
        priority 200;
        dual-dr;
        bfd-liveness-detection { ## Warning: 'bfd-liveness-detection' is deprecated
          minimum-receive-interval 700;
          transmit-interval {
            minimum-interval 350;
            threshold 500;
          }
        }
      }
      interface vlan.200 {
        priority 600;
        dual-dr;
        bfd-liveness-detection { ## Warning: 'bfd-liveness-detection' is deprecated
          minimum-receive-interval 700;
          transmit-interval {
            minimum-interval 350;
            threshold 500;
          }
        }
      }
    }
    iccp {
      local-ip-addr 10.3.3.2;
      peer 10.3.3.1 {
        session-establishment-hold-time 340;
        backup-liveness-detection {
```

```
        backup-peer-ip 10.207.64.233;
      }
      liveness-detection {
        minimum-receive-interval 1000;
        transmit-interval {
          minimum-interval 1000;
        }
      }
    }
}
igmp-snooping {
    vlan all;
}
rstp {
    interface ae0.0 {
      disable;
    }
    interface ae1.0 {
      edge;
    }
    interface ae2.0 {
      edge;
    }
    interface all {
      mode point-to-point;
    }
    bpdu-block-on-edge;
}
```

```
user@SwitchA# show multi-chassis
multi-chassis-protection 10.3.3.1 {
    interface ae0;
}
```

```
user@SwitchA# show vlans
v100 {
    vlan-id 100;
    l3-interface vlan.100;
}
v200 {
    vlan-id 200;
    l3-interface vlan.200;
}
v500 {
    vlan-id 500;
    l3-interface vlan.500;
}
```

Switch A—ELS

```
user@SwitchA# show chassis
aggregated-devices {
    ethernet {
```

```
    device-count 3;
  }
}
```

```
user@SwitchA# show interfaces
xe-0/0/12 {
  ether-options {
    802.3ad ae0;
  }
}
xe-0/0/13 {
  ether-options {
    802.3ad ae0;
  }
}
xe-0/0/44 {
  ether-options {
    802.3ad ae1;
  }
}
xe-0/0/45 {
  ether-options {
    802.3ad ae2;
  }
}
ae0 {
  unit 0 {
    family ethernet-switching {
      interface-mode trunk;
      vlan {
        members v100;
      }
    }
  }
}
ae1 {
  aggregated-ether-options {
    lacp {
      active;
      system-id 00:01:02:03:04:05;
      admin-key 3;
    }
    mc-ae {
      mc-ae-id 3;
      chassis-id 0;
      mode active-active;
      status-control active;
      init-delay-time 240;
    }
  }
  unit 0 {
    family ethernet-switching {
      interface-mode trunk;
      vlan {
```

```
            members v100;
        }
      }
    }
}
ae2 {
    aggregated-ether-options {
      lacp {
        active;
        system-id 00:01:02:03:04:06;
        admin-key 3;
      }
      mc-ae {
        mc-ae-id 4;
        chassis-id 0;
        mode active-active;
        status-control active;
        init-delay-time 240;
      }
    }
    unit 0 {
      family ethernet-switching {
        interface-mode trunk;
        vlan {
          members v200;
        }
      }
    }
}
irb {
    unit 100 {
      family inet {
        address 10.1.1.11/8 {
          vrrp-group 1 {
            virtual-address 10.1.1.1;
            priority 200;
            accept-data;
          }
        }
      }
    }
    unit 200 {
      family inet {
        address 10.1.1.21/8 {
          vrrp-group 2 {
            virtual-address 10.1.1.2;
            priority 200;
            accept-data;
          }
        }
      }
    }
    unit 500 {
      family inet {
        address 10.3.3.2/8;
```

```
    }
  }
}
```

```
user@SwitchA# show protocols
ospf {
  area 0.0.0.0 {
    interface vlan.100 {
      bfd-liveness-detection {
        minimum-receive-interval 700;
        transmit-interval {
          minimum-interval 350;
          threshold 500;
        }
      }
    }
    interface vlan.200 {
      bfd-liveness-detection {
        minimum-receive-interval 700;
        transmit-interval {
          minimum-interval 350;
          threshold 500;
        }
      }
    }
  }
}
pim {
  rp {
    static {
      address 10.0.0.3 {
        group-ranges {
          233.252.0.0/8;
        }
      }
    }
  }
  interface vlan.100 {
    priority 200;
    dual-dr;
    bfd-liveness-detection {
      minimum-receive-interval 700;
      transmit-interval {
        minimum-interval 350;
        threshold 500;
      }
    }
  }
  interface vlan.200 {
    priority 600;
    dual-dr;
    bfd-liveness-detection {
      minimum-receive-interval 700;
      transmit-interval {
```

```
            minimum-interval 350;
            threshold 500;
          }
        }
      }
    }
    iccp {
      local-ip-addr 10.3.3.2;
      peer 10.3.3.1 {
        session-establishment-hold-time 340;
        backup-liveness-detection {
          backup-peer-ip 10.207.64.233;
        }
        liveness-detection {
          minimum-receive-interval 1000;
          transmit-interval {
            minimum-interval 1000;
          }
        }
      }
    }
    igmp-snooping {
      vlan all;
    }
    rstp {
      interface ae1.0 {
        edge;
      }
      interface ae2.0 {
        edge;
      }
      interface ae1.0{
        mode point-to-point;
      }
      bpdu-block-on-edge;
    }
```

```
user@SwitchA# show multi-chassis
multi-chassis-protection 10.3.3.1 {
    interface ae0;
}
```

```
user@SwitchA# show switch-options
service-id 10;
```

```
user@SwitchA# show vlans
v100 {
    vlan-id 100;
    l3-interface irb.100;
}
v200 {
    vlan-id 200;
    l3-interface irb.200;
```

```
}
v500 {
   vlan-id 500;
   l3-interface irb.500;
}
```

From configuration mode on Switch B, confirm your configuration by entering the **show chassis**, **show interfaces**, **show multi-chassis**, **show protocols**, and **show vlans** commands. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

Switch B--Original CLI

```
user@SwitchB# show chassis
aggregated-devices {
   ethernet {
      device-count 3;
   }
}
```

```
user@SwitchB# show interfaces
xe-0/0/12 {
   ether-options {
      802.3ad ae0;
   }
}
xe-0/0/13 {
   ether-options {
      802.3ad ae0;
   }
}
xe-0/0/46 {
   ether-options {
      802.3ad ae1;
   }
}
xe-0/0/47 {
   ether-options {
      802.3ad ae2;
   }
}
ae0 {
   unit 0 {
      family ethernet-switching {
         port-mode trunk;
         vlan {
            members v500;
         }
      }
   }
}
ae1 {
   aggregated-ether-options {
```

```
      lacp {
        active;
        system-id 00:01:02:03:04:05;
        admin-key 3;
      }
      mc-ae {
        mc-ae-id 3;
        chassis-id 1;
        mode active-active;
        status-control standby;
        init-delay-time 240;
      }
    }
    unit 0 {
      family ethernet-switching {
        port-mode trunk;
        vlan {
          members v100;
        }
      }
    }
  }
  ae2 {
    aggregated-ether-options {
      lacp {
        active;
        system-id 00:01:02:03:04:06;
        admin-key 3;
      }
      mc-ae {
        mc-ae-id 4;
        chassis-id 1;
        mode active-active;
        status-control standby;
        init-delay-time 240;
      }
    }
    unit 0 {
      family ethernet-switching {
        port-mode trunk;
        vlan {
          members v200;
        }
      }
    }
  }
  irb {
    unit 100 {
      family inet {
        address 10.1.1.10/8 {
          vrrp-group 1 {
            virtual-address 10.1.1.1;
            priority 150;
            accept-data;
          }
```

```
        }
      }
    }
    unit 200 {
      family inet {
        address 10.1.1.20/8 {
          vrrp-group 2 {
            virtual-address 10.1.1.2;
            priority 150;
            accept-data;
          }
        }
      }
    }
    unit 500 {
      family inet {
        address 10.3.3.1/8;
      }
    }
}
}
```

```
user@SwitchB# show protocols
ospf {
    area 0.0.0.0 {
        interface vlan.100 {
            bfd-liveness-detection {
                minimum-receive-interval 700;
                transmit-interval {
                    minimum-interval 350;
                    threshold 500;
                }
            }
        }
        interface vlan.200 {
            bfd-liveness-detection {
                minimum-receive-interval 700;
                transmit-interval {
                    minimum-interval 350;
                    threshold 500;
                }
            }
        }
    }
}
pim {
    rp {
        static {
            address 10.0.0.3 {
                group-ranges {
                    233.252.0.0/8;
                }
            }
        }
```

```
    }
    interface vlan.100 {
      priority 100;
      dual-dr;
      bfd-liveness-detection { ## Warning: 'bfd-liveness-detection' is deprecated
        minimum-receive-interval 700;
        transmit-interval {
          minimum-interval 350;
          threshold 500;
        }
      }
    }
    interface vlan.200 {
      priority 500;
      dual-dr;
      bfd-liveness-detection { ## Warning: 'bfd-liveness-detection' is deprecated
        minimum-receive-interval 700;
        transmit-interval {
          minimum-interval 350;
          threshold 500;
        }
      }
    }
}
iccp {
    local-ip-addr 10.3.3.1;
    peer 10.3.3.2 {
      session-establishment-hold-time 340;
      backup-liveness-detection {
        backup-peer-ip 10.207.64.234;
      }
      liveness-detection {
        minimum-receive-interval 1000;
        transmit-interval {
          minimum-interval 1000;
        }
      }
    }
}
igmp-snooping {
    vlan all;
}
rstp {
    interface ae0.0 {
      disable;
    }
    interface ae1.0 {
      edge;
    }
    interface ae2.0 {
      edge;
    }
    interface all {
      mode point-to-point;
    }
```

```
    bpdu-block-on-edge;
}
```

```
user@SwitchB# show multi-chassis
multi-chassis-protection 10.3.3.2 {
    interface ae0;
}
```

```
user@SwitchB# show vlans
v100 {
    vlan-id 100;
    l3-interface vlan.100;
}
v200 {
    vlan-id 200;
    l3-interface vlan.200;
}
v500 {
    vlan-id 500;
    l3-interface vlan.500;
}
```

Switch B—ELS

```
user@SwitchB# show chassis
aggregated-devices {
    ethernet {
        device-count 3;
    }
}
```

```
user@SwitchB# show interfaces
xe-0/0/12 {
    ether-options {
        802.3ad ae0;
    }
}
xe-0/0/13 {
    ether-options {
        802.3ad ae0;
    }
}
xe-0/0/46 {
    ether-options {
        802.3ad ae1;
    }
}
xe-0/0/47 {
    ether-options {
        802.3ad ae2;
    }
}
```

```
ae0 {
  unit 0 {
    family ethernet-switching {
      interface-mode trunk;
      vlan {
        members v100;
      }
    }
  }
}
ae1 {
  aggregated-ether-options {
    lacp {
      active;
      system-id 00:01:02:03:04:05;
      admin-key 3;
    }
    mc-ae {
      mc-ae-id 3;
      chassis-id 1;
      mode active-active;
      status-control standby;
      init-delay-time 240;
    }
  }
  unit 0 {
    family ethernet-switching {
      interface-mode trunk;
      vlan {
        members v100;
      }
    }
  }
}
ae2 {
  aggregated-ether-options {
    lacp {
      active;
      system-id 00:01:02:03:04:06;
      admin-key 3;
    }
    mc-ae {
      mc-ae-id 4;
      chassis-id 1;
      mode active-active;
      status-control standby;
      init-delay-time 240;
    }
  }
  unit 0 {
    family ethernet-switching {
      interface-mode trunk;
      vlan {
        members v200;
      }
```

```
    }
  }
}
irb {
  unit 100 {
    family inet {
      address 10.1.1.10/8 {
        vrrp-group 1 {
          virtual-address 10.1.1.1;
          priority 150;
          accept-data;
        }
      }
    }
  }
  unit 200 {
    family inet {
      address 10.1.1.20/8 {
        vrrp-group 2 {
          virtual-address 10.1.1.2;
          priority 150;
          accept-data;
        }
      }
    }
  }
  unit 500 {
    family inet {
      address 10.3.3.1/8;
    }
  }
}
```

```
user@SwitchB# show protocols
ospf {
  area 0.0.0.0 {
    interface vlan.100 {
      bfd-liveness-detection {
        minimum-receive-interval 700;
        transmit-interval {
          minimum-interval 350;
          threshold 500;
        }
      }
    }
    interface vlan.200 {
      bfd-liveness-detection {
        minimum-receive-interval 700;
        transmit-interval {
          minimum-interval 350;
          threshold 500;
        }
      }
    }
```

```
    }
}
pim {
    rp {
        static {
            address 10.0.0.3 {
                group-ranges {
                    233.252.0.0/8;
                }
            }
        }
    }
    interface vlan.100 {
        priority 100;
        dual-dr;
        bfd-liveness-detection {
            minimum-receive-interval 700;
            transmit-interval {
                minimum-interval 350;
                threshold 500;
            }
        }
    }
    interface vlan.200 {
        priority 500;
        dual-dr;
        bfd-liveness-detection {
            minimum-receive-interval 700;
            transmit-interval {
                minimum-interval 350;
                threshold 500;
            }
        }
    }
}
iccp {
    local-ip-addr 10.3.3.1;
    peer 10.3.3.2 {
        session-establishment-hold-time 340;
        backup-liveness-detection {
            backup-peer-ip 10.207.64.234;
        }
        liveness-detection {
            minimum-receive-interval 1000;
            transmit-interval {
                minimum-interval 1000;
            }
        }
    }
}
igmp-snooping {
    vlan all;
}
rstp {
    interface ae1.0 {
```

```
    edge;
  }
  interface ae2.0 {
    edge;
  }
  interface ae1.0{
    mode point-to-point;
  }
  bpdu-block-on-edge;
}
```

```
user@SwitchB# show multi-chassis
multi-chassis-protection 10.3.3.2 {
  interface ae0;
}
```

```
user@SwitchB# show switch-options
service-id 10;
```

```
user@SwitchB# show vlans
v100 {
  vlan-id 100;
  l3-interface irb.100;
}
v200 {
  vlan-id 200;
  l3-interface irb.200;
}
v500 {
  vlan-id 500;
  l3-interface irb.500;
}
```

Verification

Verify that the configuration is working properly.

- Verifying That Switch A is the Master Designated Router on page 312
- Verifying That Switch B is the Backup Designated Router on page 313

Verifying That Switch A is the Master Designated Router

Purpose Verify that Switch A is the master designated router (DR).

Action From operational mode, enter the **show pim interfaces** command.

```
user@switch> show pim interfaces
```

```
Stat = Status, V = Version, NbrCnt = Neighbor Count,
S = Sparse, D = Dense, B = Bidirectional,
DR = Designated Router, P2P = Point-to-point link,
Active = Bidirectional is active, NotCap = Not Bidirectional Capable
```

```
Name               Stat Mode IP V State        NbrCnt JoinCnt(sg/*g) DR address
pime.32769         Down S     4 2   P2P,NotCap         0 0/0
vlan.100           Up   S     4 2   DDR-DR,NotCap      1 0/0           10.1.1.11
vlan.200           Up   S     4 2   DDR-DR,NotCap      2 0/0           10.1.1.21
```

Meaning The DDR-DR state of the VLAN interfaces in the output shows that Switch A is the master designated router.

Verifying That Switch B is the Backup Designated Router

Purpose Verify that Switch B is the backup designated router (BDR).

Action From operational mode, enter the **show pim interfaces** command.

```
user@switch> show pim interfaces
Stat = Status, V = Version, NbrCnt = Neighbor Count,
S = Sparse, D = Dense, B = Bidirectional,
DR = Designated Router, P2P = Point-to-point link,
Active = Bidirectional is active, NotCap = Not Bidirectional Capable

Name               Stat Mode IP V State        NbrCnt JoinCnt(sg/*g) DR address
pime.32769         Down S     4 2   P2P,NotCap         0 0/0
vlan.100           Up   S     4 2   DDR-BDR,NotCap     1 0/0           10.1.1.11
vlan.200           Up   S     4 2   DDR-BDR,NotCap     2 0/0           10.1.1.21
```

Meaning The DDR-BDR state of the VLAN interfaces in the output shows that Switch B is the backup designated router.

See Also
- Configuring Multichassis Link Aggregation on page 56

Example: Configuring Multichassis Link Aggregation for Layer 3 Multicast Using VRRP on MX Series Routers

[Warning: element unresolved in stylesheets: <author> (in <title>). This is probably a new element that is not yet supported in the stylesheets.]

There are two methods for enabling Layer 3 multicast functionality across a multichassis link aggregation group (MC-LAG). You can choose either to configure the Virtual Router Redundancy Protocol (VRRP) or synchronize the MAC addresses for the Layer 3 interfaces of the routers participating in the MC-LAG to load balance the traffic. The procedure to configure VRRP for use in a Layer 3 multicast MC-LAG is included in this example.

Requirements

This example uses the following hardware and software components:

- Four Juniper Networks MX Series routers
- Junos OS Release 11.2 or later running on all four routers

Before you configure an MC-LAG for Layer 3 multicast using VRRP, be sure that you understand how to:

- Configure aggregated Ethernet interfaces on a router.
- Configure the Link Aggregation Control Protocol (LACP) on aggregated Ethernet interfaces on a router.
- Configure the Virtual Router Redundancy Protocol (VRRP) on a router.

Overview

In this example, you configure an MC-LAG across two routers by including interfaces from both routers in an aggregated Ethernet interface (ae1). To support the MC-LAG, create a second aggregated Ethernet interface (ae0) for the interchassis control link-protection link (ICL-PL). Configure a multichassis protection link for the ICL-PL, Inter-Chassis Control Protocol (ICCP) for the peers hosting the MC-LAG, and Layer 3 connectivity between MC-LAG peers.

NOTE: **Layer 3 connectivity is required for ICCP.**

To complete the configuration, enable VRRP by completing the following steps:

- Create a routed VLAN interface (RVI).
- Create a VRRP group and assign a virtual IP address that is shared between each router in the VRRP group.
- Enable a member of a VRRP group to accept all packets destined for the virtual IP address if it is the master in the VRRP group.

Consider a sample topology in which a customer edge router, CE, is connected to two provider edge (PE) routers, PE1 and PE2, respectively. The two PE devices each have a link aggregation group (LAG) connected to the CE device. The configured mode is active-active, meaning that both PE routers' LAG ports are active and carrying traffic at the same time. PE1 and PE2 are connected to a single service provider router, P.

From the perspective of the CE device, all four ports belonging to a LAG are connected to a single service provider device. Because the configured mode is active-active, all four

ports are active, and the CE device load-balances the traffic to the peering PE devices. On the PE routers, a regular LAG is configured facing the CE device.

On one end of an MC-LAG is an MC-LAG client device, such as a server, that has one or more physical links in a LAG. This client device does not need to detect the MC-LAG. On the other side of an MC-LAG are two MC-LAG routers. Each of the routers has one or more physical links connected to a single client device. The routers coordinate with each other to ensure that data traffic is forwarded properly.

Topology Diagram

Figure 35 on page 315 shows the topology used in this example.

Figure 35: MC-LAG Active-Active on MX Series Routers

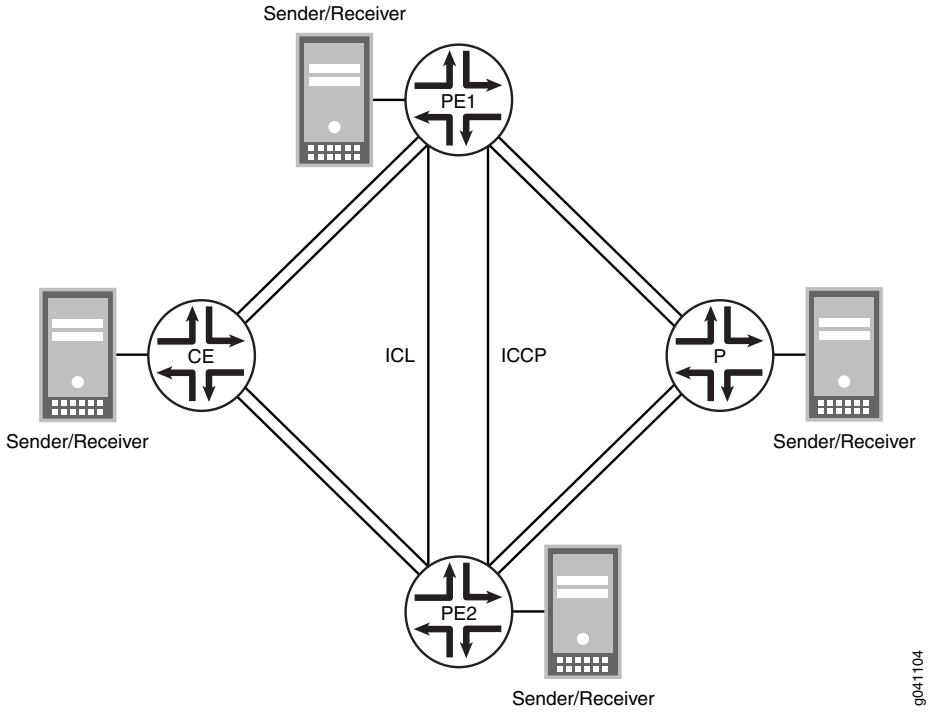

Configuring the PE Routers

CLI Quick Configuration

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

Router PE1

```
set chassis aggregated-devices ethernet device-count 5
set interfaces ge-1/0/1 gigether-options 802.3ad ae1
set interfaces ge-1/0/2 unit 0 family inet address 10.100.100.1/30
set interfaces ge-1/0/6 gigether-options 802.3ad ae0
set interfaces ge-1/1/1 flexible-vlan-tagging
set interfaces ge-1/1/1 encapsulation flexible-ethernet-services
set interfaces ge-1/1/1 unit 0 encapsulation vlan-bridge
```

```
set interfaces ge-1/1/1 unit 0 vlan-id-range 100-110
set interfaces ge-1/1/4 flexible-vlan-tagging
set interfaces ge-1/1/4 encapsulation flexible-ethernet-services
set interfaces ge-1/1/4 unit 0 encapsulation vlan-bridge
set interfaces ge-1/1/4 unit 0 vlan-id-range 100-110
set interfaces ae0 flexible-vlan-tagging
set interfaces ae0 encapsulation flexible-ethernet-services
set interfaces ae0 aggregated-ether-options lacp active
set interfaces ae0 aggregated-ether-options lacp system-priority 100
set interfaces ae0 aggregated-ether-options lacp system-id 00:00:00:00:00:05
set interfaces ae0 aggregated-ether-options lacp admin-key 1
set interfaces ae0 aggregated-ether-options mc-ae mc-ae-id 5
set interfaces ae0 aggregated-ether-options mc-ae redundancy-group 10
set interfaces ae0 aggregated-ether-options mc-ae chassis-id 1
set interfaces ae0 aggregated-ether-options mc-ae mode active-active
set interfaces ae0 aggregated-ether-options mc-ae status-control active
set interfaces ae0 unit 0 encapsulation vlan-bridge
set interfaces ae0 unit 0 vlan-id-range 100-110
set interfaces ae0 unit 0 multi-chassis-protection 10.100.100.2 interface ge-1/1/4.0
set interfaces ae1 flexible-vlan-tagging
set interfaces ae1 encapsulation flexible-ethernet-services
set interfaces ae1 aggregated-ether-options lacp active
set interfaces ae1 aggregated-ether-options lacp system-priority 100
set interfaces ae1 aggregated-ether-options lacp system-id 00:00:00:00:00:05
set interfaces ae1 aggregated-ether-options lacp admin-key 1
set interfaces ae1 aggregated-ether-options mc-ae mc-ae-id 10
set interfaces ae1 aggregated-ether-options mc-ae redundancy-group 10
set interfaces ae1 aggregated-ether-options mc-ae chassis-id 1
set interfaces ae1 aggregated-ether-options mc-ae mode active-active
set interfaces ae1 aggregated-ether-options mc-ae status-control active
set interfaces ae1 unit 0 encapsulation vlan-bridge
set interfaces ae1 unit 0 vlan-id-range 100-110
set interfaces ae1 unit 0 multi-chassis-protection 10.100.100.2 interface ge-1/1/4.0
set bridge-domains bd0 domain-type bridge
set bridge-domains bd0 vlan-id all
set bridge-domains bd0 service-id 20
set bridge-domains bd0 interface ae1.0
set bridge-domains bd0 interface ge-1/0/3.0
set bridge-domains bd0 interface ge-1/1/1.0
set bridge-domains bd0 interface ge-1/1/4.0
set bridge-domains bd0 interface ae0.0
set protocols iccp local-ip-addr 10.100.100.1
set protocols iccp peer 10.100.100.2 redundancy-group-id-list 10
set protocols iccp peer 10.100.100.2 liveness-detection minimum-interval 1000
set protocols ospf area 0.0.0.0 interface ge-1/1/1.0 bfd-liveness-detection
  minimum-receive-interval 700
set protocols ospf area 0.0.0.0 interface ge-1/1/1.0 bfd-liveness-detection
  transmit-interval minimum-interval 350
set protocols ospf area 0.0.0.0 interface ge-1/1/1.0 bfd-liveness-detection
  transmit-interval threshold 500
set protocols ospf area 0.0.0.0 interface ge-1/1/4.0 bfd-liveness-detection
  minimum-receive-interval 700
set protocols ospf area 0.0.0.0 interface ge-1/1/4.0 bfd-liveness-detection
  transmit-interval minimum-interval 350
```

```
set protocols ospf area 0.0.0.0 interface ge-1/1/4.0 bfd-liveness-detection
  transmit-interval threshold 500
set protocols pim rp static address 1.0.0.3 group-ranges 233.252.0.0/8
set protocols pim interface ge-1/1/4.0 priority 200
set protocols pim interface ge-1/1/4.0 version 2
set protocols pim interface ge-1/1/4.0 bfd-liveness-detection minimum-receive-interval
  700
set protocols pim interface ge-1/1/4.0 bfd-liveness-detection transmit-interval
  minimum-interval 350
set protocols pim interface ge-1/1/4.0 bfd-liveness-detection transmit-interval threshold
  500
set protocols pim interface ge-1/1/1.0 priority 600
set protocols pim interface ge-1/1/1.0 version 2
set protocols pim interface ge-1/1/1.0 bfd-liveness-detection minimum-receive-interval
  700
set protocols pim interface ge-1/1/1.0 bfd-liveness-detection transmit-interval
  minimum-interval 350
set protocols pim interface ge-1/1/1.0 bfd-liveness-detection transmit-interval threshold
  500
set protocols rstp interface ae1.0 edge
set protocols rstp interface ae1.0 mode point-to-point
set protocols rstp bpdu-block-on-edge
set switch-options service-id 10
```

Router PE2

```
set chassis aggregated-devices ethernet device-count 5
set interfaces ge-1/0/2 unit 0 family inet address 10.100.100.2/30
set interfaces ge-1/0/3 flexible-vlan-tagging
set interfaces ge-1/0/3 encapsulation flexible-ethernet-services
set interfaces ge-1/0/3 unit 0 encapsulation vlan-bridge
set interfaces ge-1/0/3 unit 0 vlan-id-range 100-110
set interfaces ge-1/0/4 flexible-vlan-tagging
set interfaces ge-1/0/4 encapsulation flexible-ethernet-services
set interfaces ge-1/0/4 unit 0 encapsulation vlan-bridge
set interfaces ge-1/0/4 unit 0 vlan-id-range 100-110
set interfaces ge-1/0/5 gigether-options 802.3ad ae0
set interfaces ge-1/1/0 gigether-options 802.3ad ae1
set interfaces ae0 flexible-vlan-tagging
set interfaces ae0 encapsulation flexible-ethernet-services
set interfaces ae0 aggregated-ether-options lacp active
set interfaces ae0 aggregated-ether-options lacp system-priority 100
set interfaces ae0 aggregated-ether-options lacp system-id 00:00:00:00:00:05
set interfaces ae0 aggregated-ether-options lacp admin-key 1
set interfaces ae0 aggregated-ether-options mc-ae mc-ae-id 5
set interfaces ae0 aggregated-ether-options mc-ae redundancy-group 10
set interfaces ae0 aggregated-ether-options mc-ae chassis-id 0
set interfaces ae0 aggregated-ether-options mc-ae mode active-active
set interfaces ae0 aggregated-ether-options mc-ae status-control active
set interfaces ae0 unit 0 encapsulation vlan-bridge
set interfaces ae0 unit 0 vlan-id-range 100-110
set interfaces ae0 unit 0 multi-chassis-protection 10.100.100.1 interface ge-1/0/4.0
set interfaces ae1 flexible-vlan-tagging
set interfaces ae1 encapsulation flexible-ethernet-services
set interfaces ae1 aggregated-ether-options lacp active
```

```
set interfaces ae1 aggregated-ether-options lacp system-priority 100
set interfaces ae1 aggregated-ether-options lacp system-id 00:00:00:00:00:05
set interfaces ae1 aggregated-ether-options lacp admin-key 1
set interfaces ae1 aggregated-ether-options mc-ae mc-ae-id 10
set interfaces ae1 aggregated-ether-options mc-ae redundancy-group 10
set interfaces ae1 aggregated-ether-options mc-ae chassis-id 0
set interfaces ae1 aggregated-ether-options mc-ae mode active-active
set interfaces ae1 aggregated-ether-options mc-ae status-control active
set interfaces ae1 unit 0 encapsulation vlan-bridge
set interfaces ae1 unit 0 vlan-id-range 100-110
set interfaces ae1 unit 0 multi-chassis-protection 10.100.100.1 interface ge-1/0/4.0
set bridge-domains bd0 domain-type bridge
set bridge-domains bd0 vlan-id all
set bridge-domains bd0 service-id 20
set bridge-domains bd0 interface ae1.0
set bridge-domains bd0 interface ge-1/0/3.0
set bridge-domains bd0 interface ge-1/0/4.0
set bridge-domains bd0 interface ae0.0
set protocols iccp local-ip-addr 10.100.100.2
set protocols iccp peer 10.100.100.1 redundancy-group-id-list 10
set protocols iccp peer 10.100.100.1 liveness-detection minimum-interval 1000
set protocols ospf area 0.0.0.0 interface ge-1/0/4.0 bfd-liveness-detection
  minimum-receive-interval 700
set protocols ospf area 0.0.0.0 interface ge-1/0/4.0 bfd-liveness-detection
  transmit-interval minimum-interval 350
set protocols ospf area 0.0.0.0 interface ge-1/0/4.0 bfd-liveness-detection
  transmit-interval threshold 500
set protocols pim rp static address 1.0.0.3 group-ranges 233.252.0.0/8
set protocols pim interface ge-1/0/4.0 priority 200
set protocols pim interface ge-1/0/4.0 version 2
set protocols pim interface ge-1/0/4.0 bfd-liveness-detection minimum-receive-interval
  700
set protocols pim interface ge-1/0/4.0 bfd-liveness-detection transmit-interval
  minimum-interval 350
set protocols pim interface ge-1/0/4.0 bfd-liveness-detection transmit-interval threshold
  500
set protocols rstp interface ae1.0 edge
set protocols rstp interface ae1.0 mode point-to-point
set protocols rstp bpdu-block-on-edge
set switch-options service-id 10
```

Configuring the PE1 Router

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* .

NOTE: Repeat the procedure for Router PE2, after modifying the appropriate interface names, addresses, and other parameters.

To configure Router PE1:

1. Specify the number of aggregated Ethernet interfaces to be created.

```
[edit chassis]
user@PE1# set aggregated-devices ethernet device-count 5
```

2. Specify the members to be included within the aggregated Ethernet bundles.

```
[edit interfaces]
user@PE1# set ge-1/0/1 gigether-options 802.3ad ae1
user@PE1# set ge-1/0/6 gigether-options 802.3ad ae0
```

3. Configure the interfaces that connect to senders or receivers, the interchassis link(ICL) interfaces, and the ICCP interfaces.

```
[edit interfaces]
user@PE1# set ge-1/1/1 flexible-vlan-tagging
user@PE1# set ge-1/1/1 encapsulation flexible-ethernet-services
user@PE1# set ge-1/1/1 unit 0 encapsulation vlan-bridge
user@PE1# set ge-1/1/1 unit 0 vlan-id-range 100-110

user@PE1# set ge-1/1/4 flexible-vlan-tagging
user@PE1# set ge-1/1/4 encapsulation flexible-ethernet-services
user@PE1# set ge-1/1/4 unit 0 encapsulation vlan-bridge
user@PE1# set ge-1/1/4 unit 0 vlan-id-range 100-110

user@PE1# set ge-1/0/2 unit 0 family inet address 10.100.100.1/30
```

4. Configure parameters on the aggregated Ethernet bundles.

```
[edit interfaces ae0]
user@PE1# set flexible-vlan-tagging
user@PE1# set encapsulation flexible-ethernet-services
user@PE1# set unit 0 encapsulation vlan-bridge
user@PE1# set unit 0 vlan-id-range 100-110
user@PE1# set unit 0 multi-chassis-protection 10.100.100.2 interface ge-1/1/4.0
```

```
[edit interfaces ae1]
user@PE1# set flexible-vlan-tagging
user@PE1# set encapsulation flexible-ethernet-services
user@PE1# set unit 0 encapsulation vlan-bridge
user@PE1# set unit 0 vlan-id-range 100-110
user@PE1# set unit 0 multi-chassis-protection 10.100.100.2 interface ge-1/1/4.0
```

5. Configure LACP on the aggregated Ethernet bundles.

```
[edit interfaces ae0 aggregated-ether-options]
user@PE1# set lacp active
user@PE1# set lacp system-priority 100
user@PE1# set lacp system-id 00:00:00:00:00:05
user@PE1# set lacp admin-key 1
[edit interfaces ae1 aggregated-ether-options]
user@PE1# set lacp active
user@PE1# set lacp system-priority 100
user@PE1# set lacp system-id 00:00:00:00:00:05
user@PE1# set lacp admin-key 1
```

6. Configure the MC-LAG interfaces.

```
[edit interfaces ae0 aggregated-ether-options]
user@PE1# set mc-ae mc-ae-id 5
user@PE1# set mc-ae redundancy-group 10
user@PE1# set mc-ae chassis-id 1
user@PE1# set mc-ae mode active-active
user@PE1# set mc-ae status-control active
[edit interfaces ae1 aggregated-ether-options]
user@PE1# set mc-ae mc-ae-id 10
user@PE1# set mc-ae redundancy-group 10
user@PE1# set mc-ae chassis-id 1
user@PE1# set mc-ae mode active-active
user@PE1# set mc-ae status-control active
```

The multichassis aggregated Ethernet identification number (**mc-ae-id**) specifies which link aggregation group the aggregated Ethernet interface belongs to. The ae0 interfaces on Router PE1 and Router PE2 are configured with **mc-ae-id 5**. The ae1 interfaces on Router PE1 and Router PE2 are configured with **mc-ae-id 10**.

The **redundancy-group 10** statement is used by ICCP to associate multiple chassis that perform similar redundancy functions and to establish a communication channel so that applications on peering chassis can send messages to each other. The **ae0** and **ae1** interfaces on Router PE1 and Router PE2 are configured with the same redundancy group **redundancy-group 10**.

The **chassis-id** statement is used by LACP for calculating the port number of the MC-LAG's physical member links. Router PE1 uses **chassid-id 1** to identify both its ae0 and ae1 interfaces. Router PE2 uses **chassis-id 0** to identify both its ae0 and ae1 interfaces.

The **mode** statement indicates whether an MC-LAG is in active-standby mode or active-active mode. Chassis that are in the same group must be in the same mode.

7. Configure a domain that includes the set of logical ports.

```
[edit bridge-domains bd0]
user@PE1# set domain-type bridge
user@PE1# set vlan-id all
user@PE1# set service-id 20
user@PE1# set interface ae0.0
user@PE1# set interface ae1.0
user@PE1# set interface ge-1/0/3.0
user@PE1# set interface ge-1/1/1.0
user@PE1# set interface ge-1/1/4.0
```

The ports within a bridge domain share the same flooding or broadcast characteristics in order to perform Layer 2 bridging.

The bridge-level **service-id** statement is required to link related bridge domains across peers (in this case Router PE1 and Router PE2), and should be configured with the same value.

8. Configure ICCP parameters.

```
[edit protocols iccp]
user@PE1# set local-ip-addr 10.100.100.1
user@PE1# set peer 10.100.100.2 redundancy-group-id-list 10
user@PE1# set peer 10.100.100.2 liveness-detection minimum-interval 1000
```

9. Configure the service ID at the global level.

```
[edit switch-options]
user@PE1# set service-id 10
```

You must configure the same unique network-wide configuration for a service in the set of PE routers providing the service. This service ID is required if the multichassis aggregated Ethernet interfaces are part of a bridge domain.

Step-by-Step Procedure

To enable VRRP on the MC-LAGs :

1. Assign the priority for each router in the VRRP groups.

NOTE: The router configured with the highest priority is the master.

```
[edit interfaces]
user@PE1# set vlan unit 100 family inet address 10.1.1.11/8 vrrp-group 1 priority 200
```

```
user@PE1 #set vlan unit 200 family inet address 10.1.1.21/8 vrrp-group 2 priority 200
```

```
[edit interfaces]
user@PE2# set vlan unit 100 family inet address 10.1.1.10/8 vrrp-group 1 priority 150
user@PE2#  set vlan unit 200 family inet address 10.1.1.20/8 vrrp-group 2 priority
  150
```

2. Enable the router to accept all packets destined for the virtual IP address if it is the master in a VRRP group.

```
[edit interfaces]
user@PE1# set vlan unit 100 family inet address 10.1.1.11/8 vrrp-group 1 accept-data
```

```
[edit interfaces]
user@PE2# set vlan unit 100 family inet address 10.1.1.10/8 vrrp-group 1 accept-data
user@PE2# set vlan unit 200 family inet address 10.1.1.20/8 vrrp-group 2 accept-data
```

Step-by-Step Procedure

To configure OSPF as the Layer 3 protocol:

1. Configure an OSPF area .

```
[edit protocols ospf]
user@PE1# set area 0.0.0.0
```

2. Assign the VLAN interfaces for the MC-LAGs as interfaces to the OSPF area .

```
[edit protocols ospf area 0.0.0.0]
user@PE1# set interface ge-1/1/1.0
user@PE1# set interface ge-1/4/1.0
```

3. Configure the minimum receive interval, minimum transmit interval, and transmit interval threshold for a Bidirectional Forwarding Detection (BFD) session for the OSPF interfaces .

```
[edit protocols ospf area 0.0.0.0]
user@PE1# set interface ge-1/1/1.0 bfd-liveness-detection minimum-receive-interval
  700
user@PE1# set interface ge-1/1/1.0 bfd-liveness-detection transmit-interval
  minimum-interval 350
user@PE1# set interface ge-1/1/1.0 bfd-liveness-detection transmit-interval threshold
  500
user@PE1# set interface ge-1/4/1.0 bfd-liveness-detection minimum-receive-interval
  700
user@PE1# set interface ge-1/4/1.0 bfd-liveness-detection transmit-interval
  minimum-interval 350
user@PE1# set interface ge-1/4/1.0 bfd-liveness-detection transmit-interval
  threshold 500
```

Step-by-Step Procedure

To configure PIM as the multicast protocol:

1. Configure a static rendezvous point (RP) address .

```
[edit protocols pim]
user@PE1# set rp static address 10.0.0.3
```

2. Configure the address ranges of the multicast groups for which PE1 and PE2 can be a rendezvous point (RP).

```
[edit protocols pim rp static address 10.0.0.3]
user@PE1# set group-ranges 233.252.0.0/8
```

3. Enable PIM on the VLAN interfaces for the MC-LAGs on PE1 and PE2.

```
[edit protocols pim]
user@PE1# set interface ge-1/1/1.0 version 2
user@PE1# set interface ge-1/4/1.0 version 2
```

4. Configure each PIM interface's priority for being selected as the designated router (DR).

   An interface with a higher priority value has a higher probability of being selected as the DR.

```
[edit protocols pim]
user@PE1# set interface ge-1/1/1.0 priority 600
user@PE1# set interface ge-1/4/1.0 priority 200
```

5. Configure the minimum receive interval, minimum transmit interval, and transmit interval threshold for a Bidirectional Forwarding Detection (BFD) session for the PIM interfaces on PE1 and PE2.

```
[edit protocols pim]
user@PE1# set interface ge-1/1/1.0 bfd-liveness-detection minimum-receive-interval
  700
user@PE1# set interface ge-1/1/1.0 bfd-liveness-detection transmit-interval
  minimum-interval 350
user@PE1# set interface ge-1/1/1.0 bfd-liveness-detection transmit-interval threshold
  500
user@PE1# set interface ge-1/4/1.0 bfd-liveness-detection minimum-receive-interval
  700
user@PE1# set interface ge-1/4/1.0 bfd-liveness-detection transmit-interval
  minimum-interval 350
user@PE1# set interface ge-1/4/1.0 bfd-liveness-detection transmit-interval
  threshold 500
```

Step-by-Step Procedure

To enable RSTP:

1. Enable RSTP globally on all interfaces.

```
[edit]
user@PE1#  set protocols rstp interface ae1.0 mode point-to-point
```

2. Configure the MC-LAG interfaces as edge ports on PE1 and PE2.

NOTE: The ae1 interface is a downstream interface. This is why RSTP and bpdu-block-on-edge need to be configured.

```
[edit]
user@PE1#  set protocols rstp interface ae1.0 edge
```

3. Enable BPDU blocking on all interfaces except for the ICL-PL interfaces on PE1 and PE2.

NOTE: The ae1 interface is a downstream interface. This is why RSTP and bpdu-block-on-edge need to be configured.

```
[edit]
user@PE1#  set protocols rstp bpdu-block-on-edge
```

Results

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

```
user@PE1# show bridge-domains
bd0 {
   domain-type bridge;
   vlan-id all;
   service-id 20;
   interface ae1.0;
   interface ge-1/0/3.0
   interface ge-1/1/1.0;
   interface ge-1/1/4.0;
   interface ae0.0;
}
```

```
user@PE1# show chassis
aggregated-devices {
   ethernet {
```

```
    device-count 5;
  }
}
```

```
user@PE1# show interfaces
ge-1/0/1 {
  gigether-options {
    802.3ad ae1;
  }
}
ge-1/0/6 {
  gigether-options {
    802.3ad ae0;
  }
}
ge-1/0/2 {
  unit 0 {
    family inet {
      address 10.100.100.1/30;
    }
  }
}
ge-1/1/1 {
  flexible-vlan-tagging;
  encapsulation flexible-ethernet-services;
  unit 0 {
    encapsulation vlan-bridge;
    vlan-id-range 100-110;
  }
}
ge-1/1/4 {
  flexible-vlan-tagging;
  encapsulation flexible-ethernet-services;
  unit 0 {
    encapsulation vlan-bridge;
    vlan-id-range 100-110;
  }
}
ae0 {
  flexible-vlan-tagging;
  encapsulation flexible-ethernet-services;
  aggregated-ether-options {
    lacp {
      active;
      system-priority 100;
      system-id 00:00:00:00:00:05;
      admin-key 1;
    }
    mc-ae {
      mc-ae-id 5;
      redundancy-group 10;
      chassis-id 1;
      mode active-active;
      status-control active;
```

```
    }
  }
  unit 0 {
    encapsulation vlan-bridge;
    vlan-id-range 100-110;
    multi-chassis-protection 10.100.100.2 {
      interface ge-1/1/4.0;
    }
  }
}
ae1 {
  flexible-vlan-tagging;
  encapsulation flexible-ethernet-services;
  aggregated-ether-options {
    lacp {
      active;
      system-priority 100;
      system-id 00:00:00:00:00:05;
      admin-key 1;
    }
    mc-ae {
      mc-ae-id 10;
      redundancy-group 10;
      chassis-id 1;
      mode active-active;
      status-control active;
    }
  }
  unit 0 {
    encapsulation vlan-bridge;
    vlan-id-range 100-110;
    multi-chassis-protection 10.100.100.2 {
      interface ge-1/1/4.0;
    }
  }
}
  vlan {
    unit 100 {
      family inet {
        address 10.1.1.11/8 {
          vrrp-group 1 {
            priority 200;
            accept-data;
          }
        }
```

```
user@PE1# show vrrp
vlan {
  unit 100 {
    family inet {
      address 10.1.1.11/8 {
        vrrp-group 1 {
          virtual-address 10.1.1.1;
          priority 200;
```

```
            accept-data;
          }
        }
      }
    }
    unit 200 {
      family inet {
        address 10.1.1.21/8 {
          vrrp-group 2 {
            virtual-address 10.1.1.2;
            priority 200;
            accept-data;
          }
        }
      }
    }
}
```

```
user@PE1# show protocols
iccp {
    local-ip-addr 10.100.100.1;
    peer 10.100.100.2 {
        redundancy-group-id-list 10;
        liveness-detection {
            minimum-interval 1000;
        }
        }
        }
    }
}
rstp {
    interface ae1.0 {
        edge;
    }
    interface ae1.0 {
        mode point-to-point;
    }
    bpdu-block-on-edge;
}
ospf {
    area 0.0.0.0 {
        interface ge-1/1/1.0 {
            bfd-liveness-detection {
                minimum-receive-interval 700;
                transmit-interval {
                    minimum-interval 350;
                    threshold 500;
                }
            }
        }
        interface ge-1/4/1.0 {
            bfd-liveness-detection {
                minimum-receive-interval 700;
                transmit-interval {
```

```
            minimum-interval 350;
            threshold 500;
          }
        }
      }
    }
}
pim {
  rp {
    static {
      address 10.0.0.3 {
        group-ranges {
          239.0.0.0/8;
        }
      }
    }
  }
  interface ge-1/1/1.0 {
    priority 600;
    version 2;
    bfd-liveness-detection { ## Warning: 'bfd-liveness-detection' is deprecated
      minimum-receive-interval 700;
      transmit-interval {
        minimum-interval 350;
        threshold 500;
      }
    }
  }
  interface ge-1/4/1.0 {
    priority 200;
    version 2;
    bfd-liveness-detection { ## Warning: 'bfd-liveness-detection' is deprecated
      minimum-receive-interval 700;
      transmit-interval {
        minimum-interval 350;
        threshold 500;
      }
    }
  }
}
```

```
user@PE1> show switch-options
service-id 10;
```

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

Repeat the procedure for Router PE2, using the appropriate interface names and addresses.

Configuring the CE Device

CLI Quick Configuration

To quickly configure this example, copy the following commands, paste them in a text file, remove any line breaks, change any details necessary to match your network

configuration, copy and paste the commands into the CLI at the **[edit]** hierarchy level, and then enter **commit** from configuration mode.

Device CE

```
set chassis aggregated-devices ethernet device-count 2
set interfaces ge-2/0/2 gigether-options 802.3ad ae0
set interfaces ge-2/0/3 gigether-options 802.3ad ae0
set interfaces ge-2/1/6 flexible-vlan-tagging
set interfaces ge-2/1/6 encapsulation flexible-ethernet-services
set interfaces ge-2/1/6 unit 0 encapsulation vlan-bridge
set interfaces ge-2/1/6 unit 0 vlan-id-range 100-110
set interfaces ae0 flexible-vlan-tagging
set interfaces ae0 encapsulation flexible-ethernet-services
set interfaces ae0 aggregated-ether-options lacp active
set interfaces ae0 aggregated-ether-options lacp system-priority 100
set interfaces ae0 unit 0 encapsulation vlan-bridge
set interfaces ae0 unit 0 vlan-id-range 100-500
set bridge-domains bd0 domain-type bridge
set bridge-domains bd0 vlan-id all
set bridge-domains bd0 interface ge-2/1/6.0
set bridge-domains bd0 interface ae0.0
```

Configuring the CE Device

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 Device CE:

1. Specify the number of aggregated Ethernet interfaces to be created.

```
[edit chassis]
user@CE# set aggregated-devices ethernet device-count 2
```

2. Specify the members to be included within the aggregated Ethernet bundle.

```
[edit interfaces]
user@CE# set ge-2/0/2 gigether-options 802.3ad ae0
user@CE# set ge-2/0/3 gigether-options 802.3ad ae0
```

3. Configure an interface that connects to senders or receivers.

```
[edit interfaces ge-2/1/6]
user@CE# set flexible-vlan-tagging
user@CE# set encapsulation flexible-ethernet-services
user@CE# set unit 0 encapsulation vlan-bridge
user@CE# set unit 0 vlan-id-range 100-110
```

4. Configure parameters on the aggregated Ethernet bundle.

```
[edit interfaces ae0]
user@CE# set flexible-vlan-tagging
user@CE# set encapsulation flexible-ethernet-services
user@CE# set unit 0 encapsulation vlan-bridge
user@CE# set unit 0 vlan-id-range 100-500
```

5. Configure LACP on the aggregated Ethernet bundle.

```
[edit interfaces ae0 aggregated-ether-options]
user@CE# set lacp active
user@CE# set lacp system-priority 100
```

The **active** statement initiates transmission of LACP packets.

For the **system-priority** statement, a smaller value indicates a higher priority. The device with the lower system priority value determines which links between LACP partner devices are active and which are in standby mode for each LACP group. The device on the controlling end of the link uses port priorities to determine which ports are bundled into the aggregated bundle and which ports are put in standby mode. Port priorities on the other device (the noncontrolling end of the link) are ignored.

6. Configure a domain that includes the set of logical ports.

```
[edit bridge-domains bd0]
user@CE# set domain-type bridge
user@CE# set vlan-id all
user@CE# set interface ge-2/1/6.0
user@CE# set interface ae0.0
```

The ports within a bridge domain share the same flooding or broadcast characteristics in order to perform Layer 2 bridging.

Results

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

```
user@CE# show bridge-domains
bd0 {
   domain-type bridge;
   vlan-id all;
   interface ge-2/1/6.0;
   interface ae0.0;
}
```

```
user@CE# show chassis
```

```
aggregated-devices {
  ethernet {
    device-count 2;
  }
}
```

```
user@CE# show interfaces
ge-2/0/2 {
  gigether-options {
    802.3ad ae0;
  }
}
ge-2/0/3 {
  gigether-options {
    802.3ad ae0;
  }
}
ge-2/1/6 {
  flexible-vlan-tagging;
  encapsulation flexible-ethernet-services;
  unit 0 {
    encapsulation vlan-bridge;
    vlan-id-range 100-110;
  }
}
ae0 {
  flexible-vlan-tagging;
  encapsulation flexible-ethernet-services;
  aggregated-ether-options {
    lacp {
      active;
      system-priority 100;
    }
  }
  unit 0 {
    encapsulation vlan-bridge;
    vlan-id-range 100-500;
  }
}
```

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

Configuring the Provider Router

CLI Quick Configuration

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

Router P

```
set chassis aggregated-devices ethernet device-count 2
set interfaces ge-1/0/5 gigether-options 802.3ad ae1
set interfaces ge-1/0/11 gigether-options 802.3ad ae1
```

```
set interfaces ge-1/1/3 flexible-vlan-tagging
set interfaces ge-1/1/3 encapsulation flexible-ethernet-services
set interfaces ge-1/1/3 unit 0 encapsulation vlan-bridge
set interfaces ge-1/1/3 unit 0 vlan-id-range 100-500
set interfaces ae1 flexible-vlan-tagging
set interfaces ae1 encapsulation flexible-ethernet-services
set interfaces ae1 aggregated-ether-options lacp active
set interfaces ae1 aggregated-ether-options lacp system-priority 100
set interfaces ae1 unit 0 encapsulation vlan-bridge
set interfaces ae1 unit 0 vlan-id-range 100-110
set bridge-domains bd0 vlan-id all
set bridge-domains bd0 domain-type bridge
set bridge-domains bd0 interface ge-1/1/3.0
set bridge-domains bd0 interface ae1.0
```

Configuring the P Router

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 Router P:

1. Specify the number of aggregated Ethernet interfaces to be created.

```
[edit chassis]
user@P# set aggregated-devices ethernet device-count 2
```

2. Specify the members to be included within the aggregated Ethernet bundle.

```
[edit interfaces]
user@P# set ge-1/0/5 gigether-options 802.3ad ae1
user@P# set ge-1/0/11 gigether-options 802.3ad ae1
```

3. Configure an interface that connects to senders or receivers.

```
[edit interfaces ge-1/1/3]
user@P# set flexible-vlan-tagging
user@P# set encapsulation flexible-ethernet-services
user@P# set unit 0 encapsulation vlan-bridge
user@P# set unit 0 vlan-id-range 100-500
```

4. Configure parameters on the aggregated Ethernet bundle.

```
[edit interfaces ae1]
user@P# set flexible-vlan-tagging
user@P# set encapsulation flexible-ethernet-services
user@P# set unit 0 encapsulation vlan-bridge
user@P# set unit 0 vlan-id-range 100-110
```

5. Configure LACP on the aggregated Ethernet bundle.

```
[edit interfaces ae1 aggregated-ether-options]
user@P# set lacp active
user@P# set lacp system-priority 100
```

6. Configure a domain that includes the set of logical ports.

```
[edit bridge-domains bd0]
user@P# set vlan-id all
user@P# set domain-type bridge
user@P# set interface ge-1/1/3.0
user@P# set interface ae1.0
```

Results

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

```
user@P# show bridge-domains
bd0 {
   domain-type bridge;
   vlan-id all;
   interface ge-1/1/3.0;
   interface ae1.0;
}
```

```
user@P# show chassis
aggregated-devices {
   ethernet {
      device-count 2;
   }
}
```

```
user@P# show interfaces
ge-1/0/5 {
   gigether-options {
      802.3ad ae1;
   }
}
ge-1/0/11 {
   gigether-options {
      802.3ad ae1;
   }
}
ge-1/1/3 {
   flexible-vlan-tagging;
   encapsulation flexible-ethernet-services;
   unit 0 {
```

```
    encapsulation vlan-bridge;
    vlan-id-range 100-500;
  }
}
ae1 {
  flexible-vlan-tagging;
  encapsulation flexible-ethernet-services;
  aggregated-ether-options {
    lacp {
      active;
      system-priority 100;
    }
  }
  unit 0 {
    encapsulation vlan-bridge;
    vlan-id-range 100-110;
  }
}
```

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

Verification

Confirm that the configuration is working properly by running the following commands:

- **show iccp**
- **show interfaces ae0**
- **show interfaces ae1**
- **show interfaces mc-ae**
- **show pim interfaces**
- **show vrrp**
- ***show l2-learning instance*** **extensive**

Troubleshooting

Troubleshooting a LAG That Is Down

Problem The **show interfaces terse** command shows that the MC-LAG is **down**.

Solution Check the following:

- Verify that there is no configuration mismatch.
- Verify that all member ports are up.
- Verify that the MC-LAG is part of family Ethernet switching (Layer 2 LAG).
- Verify that the MC-LAG member is connected to the correct MC-LAG member at the other end.

Example: Configuring Multichassis Link Aggregation for Layer 3 Unicast Using VRRP on MX Series Routers

There are two methods for enabling Layer 3 unicast functionality across a multichassis link aggregation group (MC-LAG) to control traffic flow. You can choose either to configure Virtual Router Redundancy Protocol (VRRP) or synchronize the MAC addresses for the Layer 3 interfaces of the routers participating in the MC-LAG. The procedure to configure VRRP for use in a Layer 3 unicast MC-LAG is included in this example.

Requirements

This example uses the following hardware and software components:

- Four Juniper Networks MX Series routers
- Junos OS Release 11.2 or later running on all four routers

Before you configure an MC-LAG, be sure that you understand how to:

- Configure aggregated Ethernet interfaces on a router.
- Configure Link Aggregation Control Protocol (LACP) on aggregated Ethernet interfaces on a router.
- Configure Virtual Router Redundancy Protocol (VRRP) on a router.

Overview

In this example, you configure an MC-LAG across two routers by including interfaces from both routers in an aggregated Ethernet interface (ae0). Configure a multichassis protection link for the ICL-PL, Inter-Chassis Control Protocol (ICCP) for the peers hosting the MC-LAG, and Layer 3 connectivity between MC-LAG peers.

NOTE: **Layer 3 connectivity is required for ICCP.**

To complete the configuration, enable VRRP by completing the following steps:

1. Create a routed VLAN interface (RVI).
2. Create a VRRP group and assign a virtual IP address that is shared between each router in the VRRP group.

3. Enable a member of a VRRP group to accept all packets destined for the virtual IP address.

Consider a sample topology in which a customer edge router, CE, is connected to two provider edge (PE) routers, PE1 and PE2, respectively. The two PE devices each have a LAG connected to the CE device. The configured mode is active-active, meaning that both PE routers' LAG ports are active and carrying traffic at the same time. PE1 and PE2 are connected to a single service provider router, the P router.

In this example, the CE router is not aware that its aggregated Ethernet links are connected to two separate PE devices. The two PE devices each have a LAG connected to the CE device. The configured mode is active-active, meaning that both PE routers' LAG ports are active and carrying traffic at the same time.

From the perspective of Router CE, the two ports belonging to a LAG are connected to a single service provider device. Because the configured mode is active-active, the two ports are active, and the CE device load-balances the traffic to the peering PE devices. On the PE routers, a regular LAG is configured facing the CE device.

On one end of an MC-LAG is an MC-LAG client device, such as a server, that has one or more physical links in a link aggregation group (LAG). This client device does not need to detect the MC-LAG. On the other side of an MC-LAG are two MC-LAG routers. Each of the routers has one or more physical links connected to a single client device. The routers coordinate with each other to ensure that data traffic is forwarded properly.

ICCP messages are sent between the two PE devices. In this example, you configure an MC-LAG across two routers, consisting of two aggregated Ethernet interfaces, an interchassis link-protection link (ICL-PL), multichassis protection link for the ICL-PL, and ICCP for the peers hosting the MC-LAG.

Topology Diagram

Figure 36 on page 337 shows the topology used in this example.

Figure 36: MC-LAG Active-Active on MX Series Routers

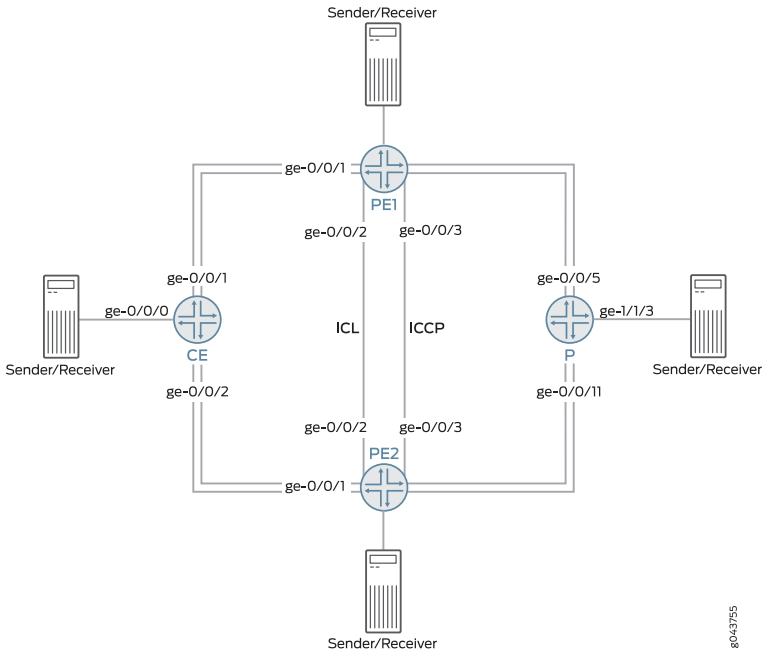

Configuring the PE Routers

CLI Quick Configuration

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

Router PE1

```
set chassis aggregated-devices ethernet device-count 4
set interfaces ge-0/0/1 gigether-options 802.3ad ae0
set interfaces ge-0/0/2 description "icl link"
set interfaces ge-0/0/2 flexible-vlan-tagging
set interfaces ge-0/0/2 encapsulation flexible-ethernet-services
set interfaces ge-0/0/2 unit 1 encapsulation vlan-bridge
set interfaces ge-0/0/2 unit 1 vlan-id 1
set interfaces ge-0/0/3 description "ICCP Link"
set interfaces ge-0/0/3 unit 0 family inet address 192.168.143.17/24
set interfaces ae0 flexible-vlan-tagging
set interfaces ae0 encapsulation flexible-ethernet-services
set interfaces ae0 aggregated-ether-options lacp active
set interfaces ae0 aggregated-ether-options lacp periodic fast
set interfaces ae0 aggregated-ether-options lacp system-id 00:00:00:00:00:02
set interfaces ae0 aggregated-ether-options lacp admin-key 10
set interfaces ae0 aggregated-ether-options mc-ae mc-ae-id 1
set interfaces ae0 aggregated-ether-options mc-ae redundancy-group 1
set interfaces ae0 aggregated-ether-options mc-ae chassis-id 1
set interfaces ae0 aggregated-ether-options mc-ae mode active-active
set interfaces ae0 aggregated-ether-options mc-ae status-control standby
set interfaces ae0 unit 1 encapsulation vlan-bridge
```

```
set interfaces ae0 unit 1 vlan-id 1
set interfaces ae0 unit 1 multi-chassis-protection 192.168.143.16 interface ge-0/0/2.1
set interfaces irb unit 1 family inet address 10.1.1.2/24 vrrp-group 1 virtual-address 10.1.1.5
set interfaces irb unit 1 family inet address 10.1.1.2/24 vrrp-group 1 priority 200
set interfaces irb unit 1 family inet address 10.1.1.2/24 vrrp-group 1 accept-data
set protocols iccp local-ip-addr 192.168.143.17
set protocols iccp peer 192.168.143.16 redundancy-group-id-list 1
set protocols iccp peer 192.168.143.16 liveness-detection minimum-interval 2500
set protocols iccp peer 192.168.143.16 liveness-detection multiplier 3
set bridge-domains bd1 vlan-id 1
set bridge-domains bd1 interface ae0.1
set bridge-domains bd1 interface ge-0/0/2.1
set bridge-domains bd1 routing-interface irb.1
set switch-options service-id 1
```

Router PE2

```
set chassis aggregated-devices ethernet device-count 4
set interfaces ge-0/0/1 gigether-options 802.3ad ae0
set interfaces ge-0/0/2 description "icl link"
set interfaces ge-0/0/2 flexible-vlan-tagging
set interfaces ge-0/0/2 encapsulation flexible-ethernet-services
set interfaces ge-0/0/2 unit 1 encapsulation vlan-bridge
set interfaces ge-0/0/2 unit 1 vlan-id 1
set interfaces ge-0/0/3 description "ICCP Link"
set interfaces ge-0/0/3 unit 0 family inet address 192.168.143.16/24
set interfaces ae0 flexible-vlan-tagging
set interfaces ae0 encapsulation flexible-ethernet-services
set interfaces ae0 aggregated-ether-options lacp active
set interfaces ae0 aggregated-ether-options lacp periodic fast
set interfaces ae0 aggregated-ether-options lacp system-id 00:00:00:00:00:02
set interfaces ae0 aggregated-ether-options lacp admin-key 10
set interfaces ae0 aggregated-ether-options mc-ae mc-ae-id 1
set interfaces ae0 aggregated-ether-options mc-ae redundancy-group 1
set interfaces ae0 aggregated-ether-options mc-ae chassis-id 0
set interfaces ae0 aggregated-ether-options mc-ae mode active-active
set interfaces ae0 aggregated-ether-options mc-ae status-control active
set interfaces ae0 unit 1 encapsulation vlan-bridge
set interfaces ae0 unit 1 vlan-id 1
set interfaces ae0 unit 1 multi-chassis-protection 192.168.143.17 interface ge-0/0/2.1
set interfaces irb unit 1 family inet address 10.1.1.3/24 vrrp-group 1 virtual-address 10.1.1.5
set interfaces irb unit 1 family inet address 10.1.1.3/24 vrrp-group 1 priority 200
set interfaces irb unit 1 family inet address 10.1.1.3/24 vrrp-group 1 accept-data
set protocols iccp local-ip-addr 192.168.143.16
set protocols iccp peer 192.168.143.17 redundancy-group-id-list 1
set protocols iccp peer 192.168.143.17 liveness-detection minimum-interval 2500
set protocols iccp peer 192.168.143.17 liveness-detection multiplier 3
set bridge-domains bd1 vlan-id 1
set bridge-domains bd1 interface ae0.1
set bridge-domains bd1 interface ge-0/0/2.1
set bridge-domains bd1 routing-interface irb.1
set switch-options service-id 1
```

Configuring the PE1 Router

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* .

NOTE: Repeat this procedure for Router PE2, after modifying the appropriate interface names, addresses, and other parameters.

To configure Router PE1:

1. Specify the number of aggregated Ethernet interfaces to be created.

```
[edit chassis]
user@PE1# set chassis aggregated-devices ethernet device-count 4
```

2. Specify the members to be included within the aggregated Ethernet bundles.

```
[edit interfaces]
user@PE1# set ge-0/0/1 gigether-options 802.3ad ae0
```

3. Configure the interfaces that connect to senders or receivers, the ICL interfaces, and the ICCP interfaces.

```
[edit interfaces]
user@PE1# set ge-0/0/2 description "icl link"
user@PE1# set ge-0/0/2 flexible-vlan-tagging
user@PE1# set ge-0/0/2 encapsulation flexible-ethernet-services
user@PE1# set ge-0/0/2 unit 1 encapsulation vlan-bridge
user@PE1# set ge-0/0/2 unit 1 vlan-id 1
user@PE1# set ge-0/0/3 description "ICCP Link"
user@PE1# set ge-0/0/3 unit 0 family inet address 192.168.143.17/24
```

4. Configure parameters on the aggregated Ethernet bundles.

```
[edit interfaces ae0]
user@PE1# set flexible-vlan-tagging
user@PE1# set encapsulation flexible-ethernet-services
```

5. Configure LACP on the aggregated Ethernet bundles.

```
[edit interfaces ae0 aggregated-ether-options]
user@PE1# set lacp active
user@PE1# set lacp periodic fast
```

```
user@PE1# set lacp system-id 00:00:00:00:00:02
user@PE1# set lacp admin-key 10
```

6. Configure the MC-LAG interfaces.

```
[edit interfaces ae0 aggregated-ether-options]
user@PE1# set mc-ae mc-ae-id 1
user@PE1# set mc-ae redundancy-group 1
user@PE1# set mc-ae chassis-id 1
user@PE1# set mc-ae mode active-active
user@PE1# set interfaces ae0 aggregated-ether-options mc-ae status-control
  standby
```

The multichassis aggregated Ethernet identification number (**mc-ae-id**) specifies which link aggregation group the aggregated Ethernet interface belongs to. The ae0 interface on Router PE1 is configured with **mc-ae-id 1**.

The **redundancy-group 1** statement is used by ICCP to associate multiple chassis that perform similar redundancy functions and to establish a communication channel so that applications on peering chassis can send messages to each other. The ae0 interfaces on Router PE1 and Router PE2 are configured with the same redundancy group, **redundancy-group 1**.

The **chassis-id** statement is used by LACP for calculating the port number of the MC-LAG's physical member links. Router PE1 uses **chassid-id 1** to identify ae0 interface.

The **mode** statement indicates whether an MC-LAG is in active-standby mode or active-active mode. Chassis that are in the same group must be in the same mode.

7. Configure a domain that includes the set of logical ports.

```
[edit bridge-domains bd0]
user@PE1# set domain-type bridge
user@PE1# set vlan-id 1
user@PE1# set service-id 1
user@PE1# set interface ae0.1
user@PE1# set interface ge-0/0/2.1
```

The ports within a bridge domain share the same flooding or broadcast characteristics in order to perform Layer 2 bridging.

The bridge-level **service-id** statement is required to link related bridge domains across peers (in this case Router PE1 and Router PE2), and should be configured with the same value.

8. Configure ICCP parameters.

```
[edit protocols iccp]
user@PE1# set local-ip-addr 192.168.143.17
user@PE1# set peer 192.168.143.16 redundancy-group-id-list 1
```

```
user@PE1# set peer 1192.168.143.16 liveness-detection minimum-interval 2500
user@PE1# set peer 1192.168.143.16 liveness-detection multiplier 3
```

9. Configure the service ID at the global level.

```
[edit switch-options]
user@PE1# set service-id 1
```

You must configure the same unique network-wide configuration for a service in the set of PE routers providing the service. This service ID is required if the multichassis aggregated Ethernet interfaces are part of a bridge domain.

Step-by-Step Procedure

To enable VRRP on the MC-LAGs :

- Enable VRRP on the MC-LAG.
  - Create a Integrated Routing and Bridging (IRB), assign a virtual IP address that is shared between each router in the VRRP group, and assign an individual IP address for each router in the VRRP group.

    ```
    [edit]
    user@PE1# set interfaces irb unit 1 family inet address 10.1.1.3/24 vrrp-group 1
      virtual-address 10.1.1.5
    ```

  - Assign the priority for each router in the VRRP group.

    NOTE: The router configured with the highest priority is the master.

    ```
    [edit interfaces irb]
    user@PE1# unit 1 family inet address 10.1.1.2/24 vrrp-group 1 virtual-address 10.1.1.5
    user@PE1# unit 1 family inet address 10.1.1.2/24 vrrp-group 1 priority 200
    ```

  - Enable the router to accept all packets destined for the virtual IP address.

    ```
    [edit interfaces irb]
    user@PE1# unit 1 family inet address 10.1.1.2/24 vrrp-group 1 accept-data
    ```

Configuring the PE2 Router

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 Router PE2:

1. Specify the number of aggregated Ethernet interfaces to be created.

```
[edit chassis]
user@PE2# set chassis aggregated-devices ethernet device-count 4
```

2. Specify the members to be included within the aggregated Ethernet bundles.

```
[edit interfaces]
user@PE2# set ge-0/0/1 gigether-options 802.3ad ae0
```

3. Configure the interfaces that connect to senders or receivers, the ICL interfaces, and the ICCP interfaces.

```
[edit interfaces]
user@PE2# set ge-0/0/2 description "icl link"
user@PE2# set ge-0/0/2 flexible-vlan-tagging
user@PE2# set ge-0/0/2 encapsulation flexible-ethernet-services
user@PE2# set ge-0/0/2 unit 1 encapsulation vlan-bridge
user@PE2# set ge-0/0/2 unit 1 vlan-id 1
user@PE2# set ge-0/0/3 description "ICCP Link"
user@PE2# set ge-0/0/3 unit 0 family inet address 192.168.143.16/24
```

4. Configure parameters on the aggregated Ethernet bundles.

```
[edit interfaces ae0]
user@PE2# set flexible-vlan-tagging
user@PE2# set encapsulation flexible-ethernet-services
```

5. Configure LACP on the aggregated Ethernet bundles.

```
[edit interfaces ae0 aggregated-ether-options]
user@PE2# set lacp active
user@PE2# set lacp periodic fast
user@PE2# set lacp system-id 00:00:00:00:00:02
user@PE2# set lacp admin-key 10
```

6. Configure the MC-LAG interfaces.

```
[edit interfaces ae0 aggregated-ether-options]
user@PE2# set mc-ae mc-ae-id 1
user@PE2# set mc-ae redundancy-group 1
user@PE2# set mc-ae chassis-id 0
user@PE2# set mc-ae mode active-active
user@PE2# set interfaces ae0 aggregated-ether-options mc-ae status-control
  active
```

The multichassis aggregated Ethernet identification number (**mc-ae-id**) specifies which link aggregation group the aggregated Ethernet interface belongs to. The ae0 interface on Router PE2 is configured with **mc-ae-id 1**.

The **redundancy-group 1** statement is used by ICCP to associate multiple chassis that perform similar redundancy functions and to establish a communication channel so that applications on peering chassis can send messages to each other. The ae0 interfaces on Router PE1 and Router PE2 are configured with the same redundancy group, **redundancy-group 1**.

The **chassis-id** statement is used by LACP for calculating the port number of the MC-LAG's physical member links. Router PE2 uses **chassid-id 0** to identify its ae0.

The **mode** statement indicates whether an MC-LAG is in active-standby mode or active-active mode. Chassis that are in the same group must be in the same mode.

7. Configure a domain that includes the set of logical ports.

```
[edit bridge-domains bd0]
user@PE2# set domain-type bridge
user@PE2# set vlan-id 1
user@PE2# set service-id 1
user@PE2# set interface ae0.1
user@PE2# set interface ge-0/0/2.1
```

The ports within a bridge domain share the same flooding or broadcast characteristics in order to perform Layer 2 bridging.

The bridge-level **service-id** statement is required to link related bridge domains across peers (in this case Router PE1 and Router PE2), and should be configured with the same value.

8. Configure ICCP parameters.

```
[edit protocols iccp]
user@PE2# set local-ip-addr 192.168.143.16
user@PE2# set peer 192.168.143.17 redundancy-group-id-list 1
user@PE2# set peer 1192.168.143.17 liveness-detection minimum-interval 2500
user@PE2# set peer 1192.168.143.17 liveness-detection multiplier 3
```

9. Configure the service ID at the global level.

```
[edit switch-options]
user@PE1# set service-id 1
```

You must configure the same unique network-wide configuration for a service in the set of PE routers providing the service. This service ID is required if the multichassis aggregated Ethernet interfaces are part of a bridge domain.

Step-by-Step Procedure

To enable VRRP on the MC-LAGs :

- Enable VRRP on the MC-LAG.
  - Create a Integrated and Bridging Interface (IRB), assign a virtual IP address that is shared between each router in the VRRP group, and assign an individual IP address for each router in the VRRP group.

```
[edit]
user@PE2set interfaces irb unit 1 family inet address 10.1.1.3/24 vrrp-group 1
   virtual-address 10.1.1.5
```

  - Assign the priority for each router in the VRRP group.

NOTE: The router configured with the highest priority is the master.

```
[edit interfaces irb]
user@PE2#unit 1 family inet address 10.1.1.3/24 vrrp-group 1 virtual-address 10.1.1.5
user@PE2#unit 1 family inet address 10.1.1.3/24 vrrp-group 1 priority 200
```

  - Enable the router to accept all packets destined for the virtual IP address.

```
[edit interfaces irb]
user@PE1# unit 1 family inet address 10.1.1.3/24 vrrp-group 1 accept-data
```

Results

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

```
user@PE1#  show bridge-domains
bd1 {
   domain-type bridge;
   vlan-id 1;
   interface ge-0/0/2.1;
   interface ae0.1;
```

```
    interface irb.1;
}
```

```
user@PE1#  show interfaces
ge-0/0/1 {
    gigether-options {
        802.3ad ae0;
    }
}
ge-0/0/2 {
    description "icl link";
        flexible-vlan-tagging;
        encapsulation flexible-ethernet-services;
        unit 1 {
        encapsulation vlan-bridge;
        vlan-id 1;
    }
}
ge-0/0/3 {
    description "ICCP Link";
        unit 0 {
        family inet {
            address 192.168.143.17/24;
        }
    }
}
ae0 {
    flexible-vlan-tagging;
        encapsulation flexible-ethernet-services;
        aggregated-ether-options {
        lacp {
        active;
            periodic fast;
            system-id 00:00:00:00:00:02;
        admin-key 10;
        }
        mc-ae {
        mc-ae-id 1;
            redundancy-group 1;
            chassis-id 1;
            mode active-active;
        status-control standby;
    }
}
unit 1 {
    encapsulation vlan-bridge;
    vlan-id 1;
    multi-chassis-protection 192.168.143.16 {
        interface ge-0/0/2.1;
    }
}
irb {
    unit 1 {
        family inet {
```

```
        address 10.1.1.2/24 {
          vrrp-group 1 {
            virtual-address 10.1.1.5;
            priority 200;
            accept-data;
          }
        }
      }
    }
}
```

```
user@PE1#  show chassis
aggregated-devices {
   ethernet {
      device-count 4;
   }
}
```

```
user@PE1#  show protocols
iccp {
   local-ip-addr 192.168.143.17;
   peer 192.168.143.16 {
      redundancy-group-id-list 1;
      liveness-detection {
         minimum-interval 2500;
         multiplier 3;
      }
   }
}
```

```
user@PE1#  show switch-options
service-id 1;
```

```
user@PE2#  show bridge-domains
bd1 {
   vlan-id 1;
   interface ae0.1;
   interface ge-0/0/2.1;
   routing-interface irb.1;
}
```

```
user@PE2#  show interfaces
ge-0/0/1 {
   gigether-options {
      802.3ad ae0;
   }
}
ge-0/0/2 {
   description "icl link";
   flexible-vlan-tagging;
   encapsulation flexible-ethernet-services;
   unit 1 {
```

```
      encapsulation vlan-bridge;
      vlan-id 1;
    }
}
ge-0/0/3 {
    description "ICCP Link";
    unit 0 {
      family inet {
        address 192.168.143.16/24;
      }
    }
}
ae0 {
    flexible-vlan-tagging;
    encapsulation flexible-ethernet-services;
    aggregated-ether-options {
      lacp {
        active;
        periodic fast;
        system-id 00:00:00:00:00:02;
        admin-key 10;
      }
      mc-ae {
        mc-ae-id 1;
        redundancy-group 1;
        chassis-id 0;
        mode active-active;
        status-control active;
      }
    }
    unit 1 {
      encapsulation vlan-bridge;
      vlan-id 1;
      multi-chassis-protection 192.168.143.17 {
        interface ge-0/0/2.1;
      }
    }
}
irb {
    unit 1 {
      family inet {
        address 10.1.1.3/24 {
          vrrp-group 1 {
            virtual-address 10.1.1.5;
            priority 200;
            accept-data;
          }
        }
      }
    }
}
```

```
user@PE2#  show protocols
iccp {
```

```
    local-ip-addr 192.168.143.16;
    peer 192.168.143.17 {
      redundancy-group-id-list 1;
      liveness-detection {
        minimum-interval 2500;
        multiplier 3;
      }
    }
}
```

```
user@PE2#  show switch-options
service-id 1;
```

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

Configuring the CE Device

CLI Quick Configuration

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

Device CE

```
set chassis aggregated-devices ethernet device-count 4
set interfaces ge-0/0/1 gigether-options 802.3ad ae0
set interfaces ge-0/0/2 gigether-options 802.3ad ae0
set interfaces ge-0/0/0 flexible-vlan-tagging
set interfaces ge-0/0/0 encapsulation flexible-ethernet-services
set interfaces ge-0/0/0 unit 1 encapsulation vlan-bridge
set interfaces ge-0/0/0 unit 1 vlan-id 1
set interfaces ae0 flexible-vlan-tagging
set interfaces ae0 encapsulation flexible-ethernet-services
set interfaces ae0 aggregated-ether-options lacp active
set interfaces ae0 aggregated-ether-options lacp periodic fast
set interfaces ae0 aggregated-ether-options lacp system-priority 127
set interfaces ae0 aggregated-ether-options lacp system-id 00:00:00:00:00:01
set interfaces ae0 unit 1 encapsulation vlan-bridge
set interfaces ae0 unit 1 vlan-id 1
set bridge-domains bd1 vlan-id 1
set bridge-domains bd1 interface ae0.1
set bridge-domains bd1 interface ge-0/0/0.1
```

Configuring the CE Device

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 Device CE:

1. Specify the number of aggregated Ethernet interfaces to be created.

```
[edit chassis]
user@CE# set aggregated-devices ethernet device-count 4
```

2. Specify the members to be included within the aggregated Ethernet bundle.

```
[edit interfaces]
user@CE# set ge-0/0/1 gigether-options 802.3ad ae0
user@CE# set ge-0/0/2 gigether-options 802.3ad ae0
```

3. Configure an interface that connects to senders or receivers.

```
[edit interfaces ge-0/0/0]
user@CE# set flexible-vlan-tagging
user@CE# set encapsulation flexible-ethernet-services
user@CE# set unit 1 encapsulation vlan-bridge
user@CE# set unit 1 vlan-id 1
```

4. Configure parameters on the aggregated Ethernet bundle.

```
[edit interfaces ae0]
user@CE# set flexible-vlan-tagging
user@CE# set encapsulation flexible-ethernet-services
```

5. Configure LACP on the aggregated Ethernet bundle.

```
[edit interfaces ae0 aggregated-ether-options]
user@CE# set lacp active
user@CE# set lacp periodic fast
user@CE# set lacp system-priority 127
user@CE# set lacp system-id 00:00:00:00:00:01
```

The **active** statement initiates transmission of LACP packets.

For the **system-priority** statement, a smaller value indicates a higher priority. The device with the lower system priority value determines which links between LACP partner devices are active and which are in standby mode for each LACP group. The device on the controlling end of the link uses port priorities to determine which ports are bundled into the aggregated bundle and which ports are put in standby mode. Port priorities on the other device (the noncontrolling end of the link) are ignored.

6. Configure a domain that includes the set of logical ports.

```
[edit bridge-domains bd0]
user@CE# set domain-type bridge
user@CE# set vlan-id 1
user@CE# set interface ge-0/0/0.1
user@CE# set interface ae0.1
```

The ports within a bridge domain share the same flooding or broadcast characteristics in order to perform Layer 2 bridging.

Results

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

```
user@CE#  show bridge-domains
bd1 {
   vlan-id 1;
   interface ae0.1;
   interface ge-0/0/0.1;
}
```

```
user@CE#  show chassis
aggregated-devices {
   ethernet {
      device-count 4;
   }
}
```

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

Configuring the Provider Router, P

- Configuring the P Router on page 351
- Results on page 352

CLI Quick Configuration

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

Router P

```
set chassis aggregated-devices ethernet device-count 4
set interfaces ge-1/0/5 gigether-options 802.3ad ae1
set interfaces ge-1/0/11 gigether-options 802.3ad ae1
set interfaces ge-1/1/3 flexible-vlan-tagging
set interfaces ge-1/1/3 encapsulation flexible-ethernet-services
set interfaces ge-1/1/3 unit 0 encapsulation vlan-bridge
set interfaces ge-1/1/3 unit 0 vlan-id-range 100-500
```

```
set interfaces ae1 flexible-vlan-tagging
set interfaces ae1 encapsulation flexible-ethernet-services
set interfaces ae1 aggregated-ether-options lacp active
set interfaces ae1 aggregated-ether-options lacp system-priority 100
set interfaces ae1 unit 0 encapsulation vlan-bridge
set interfaces ae1 unit 0 vlan-id-range 100-110
set bridge-domains bd0 vlan-id all
set bridge-domains bd0 domain-type bridge
set bridge-domains bd0 interface ge-1/1/3.0
set bridge-domains bd0 interface ae1.0
```

Configuring the P Router

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 Router P:

1. Specify the number of aggregated Ethernet interfaces to be created.

```
[edit chassis]
user@P# set aggregated-devices ethernet device-count 4
```

2. Specify the members to be included within the aggregated Ethernet bundle.

```
[edit interfaces]
user@P# setge-0/0/1 gigether-options 802.3ad ae0
user@P# set ge-1/0/11 gigether-options 802.3ad ae1
```

3. Configure an interface that connects to senders or receivers.

```
[edit interfaces ge-1/1/3]
user@P# set flexible-vlan-tagging
user@P# set encapsulation flexible-ethernet-services
user@P# set unit 0 encapsulation vlan-bridge
user@P# set unit 0 vlan-id-range 100-500
```

4. Configure parameters on the aggregated Ethernet bundle.

```
[edit interfaces ae1]
user@P# set flexible-vlan-tagging
user@P# set encapsulation flexible-ethernet-services
user@P# set unit 0 encapsulation vlan-bridge
user@P# set unit 0 vlan-id-range 100-110
```

5. Configure LACP on the aggregated Ethernet bundle.

```
[edit interfaces ae1 aggregated-ether-options]
user@P# set lacp active
user@P# set lacp system-priority 100
```

6. Configure a domain that includes the set of logical ports.

```
[edit bridge-domains bd0]
user@P# set vlan-id all
user@P# set domain-type bridge
user@P# set interface ge-1/1/3.0
user@P# set interface ae1.0
```

Results

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

```
user@P#  show bridge-domains
bd0 {
   domain-type bridge;
   vlan-id all;
   interface ge-1/1/3.0;
   interface ae1.0;
}
```

```
user@P#  show chassis
aggregated-devices {
   ethernet {
      device-count 2;
   }
}
```

```
user@P#  show interfaces
ge-1/0/5 {
   gigether-options {
      802.3ad ae1;
   }
}
ge-1/0/11 {
   gigether-options {
      802.3ad ae1;
   }
}
ge-1/1/3 {
   flexible-vlan-tagging;
   encapsulation flexible-ethernet-services;
   unit 0 {
```



```
      encapsulation vlan-bridge;
      vlan-id-range 100-500;
    }
  }
  ae1 {
    flexible-vlan-tagging;
    encapsulation flexible-ethernet-services;
    aggregated-ether-options {
      lacp {
        active;
        system-priority 100;
      }
    }
    unit 0 {
      encapsulation vlan-bridge;
      vlan-id-range 100-110;
    }
  }
```

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

Verification

Confirm that the configuration is working properly by running the following commands:

- **show iccp**
- **show interfaces ae0**
- **show interfaces ae1**
- **show interfaces mc-ae**
- **show vrrp**
- **show l2-learning instance extensive**

Troubleshooting

- Troubleshooting a LAG That Is Down on page 353

Troubleshooting a LAG That Is Down

Problem The **show interfaces terse** command shows that the MC-LAG is **down**.

Solution Check the following:

1. Verify that there is no configuration mismatch.
2. Verify that all member ports are up.
3. Verify that the MC-LAG is part of family Ethernet switching (Layer 2 LAG).
4. Verify that the MC-LAG member is connected to the correct MC-LAG member at the other end.

CHAPTER 6

Managing MC-LAG Configurations

Synchronizing the Configuration Across an MC-LAG

Understanding Configuration Synchronization

Configuration synchronization works on QFX Series switches, Junos Fusion Provider Edge, Junos Fusion Enterprise, EX Series switches, and MX Series routers.

This topic describes:

Benefits of Configuration Synchronization

Configuration synchronization enables you to propagate, synchronize, and commit configurations from one device to another. You can log into any one of those devices to manage all devices, thus having a single point of management.

How Configuration Synchronization Works

Use configuration groups to simplify the configuration process. For example, you can create one configuration group for the local device, one or more for the remote devices, and one for the global configuration, which is essentially a configuration that is common to all devices.

In addition, you can create conditional groups to specify when a configuration is synchronized with another device. You can enable the **peers-synchronize** statement at the **[edit system commit]** hierarchy to synchronize the configurations and commits across the devices by default. NETCONF over SSH provides a secure connection between the devices, and Secure Copy Protocol (SCP) copies the configurations securely between them.

How to Enable Configuration Synchronization

To enable configuration synchronization, perform the following steps:

1. Statically map the local device to the remote devices.
2. Create configuration groups for local, remote, and global configurations.
3. Create conditional groups.
4. Create apply groups.
5. Enable NETCONF over SSH.
6. Configure the device details and user authentication details for configuration synchronization.
7. Enable the **peers-synchronize** statement or issue the **commit peers-synchronize** command to synchronize and commit the configurations between local and remote devices.

How Configuration Synchronization is Supported

On MX Series routers and Junos Fusion, support for configuration synchronization started with Junos OS Release 14.2R6. On QFX Series switches, support for configuration synchronization started with Junos OS Release 15.1X53-D60. On Junos Fusion Data Center, support for configuration synchronization started with Junos OS Release 17.4R1.

Configuration Groups for Local, Remote and Global Configurations

You can create configuration groups for local, remote and global configurations. A local configuration group is used by the local device, a remote configuration group is used by the remote device, and a global configuration group is shared between the local and remote devices.

For example, you could create a local configuration group called Group A, which would include the configuration used by the local device (Switch A), a remote configuration group called Group B, which would include the configuration used by remote devices (Switch B, Switch C, and Switch D), and a global configuration group called Group C, which would include the configuration that is common to all devices.

Create configuration groups at the **[edit groups]** hierarchy level.

NOTE: Configuration synchronization does not support nested groups.

Creating Conditional Groups for Certain Devices

You can create conditional groups to specify when a particular configuration should be applied to a device. If you want to apply the global configuration to all devices in a four-device configuration, for example, enable the **when peers [<*name of local peer*> <*name of remote peer*> <*name of remote peer*> <*name of remote peer*>]** statement at the **[edit groups]** hierarchy level. If, for example, you want to apply the global configuration (Group C) to the local and remote devices (Switch A, Switch B, Switch C, and Switch D), you could issue the **set groups Group C when peers [Switch A Switch B Switch C Switch D]** command.

Applying Configuration Groups

To apply configuration groups, enable the **apply-groups** statement at the **[edit]** hierarchy level. For example, to apply the local configuration group (Group A, for example), remote configuration group (Group B, for example), and global configuration group (Group C, for example), issue the **set apply-groups [GroupA GroupB GroupC]** command.

Device Configuration Details for Configuration Synchronization

To synchronize configurations between devices, you need to configure the hostname or IP address, username, and password for the remote devices. To do this, issue the **set peers <*hostname-of-remote-peer*> user <*name-of-user*> authentication <*plain-text-password-string*>** command at the **[edit system commit]** hierarchy on the local device.

For example, to synchronize a configuration from Switch A to Switch B, issue the **set peers SwitchB user administrator authentication test123** command on Switch A.

You also need to statically map the local device to the remote devices. To this, issue the **set system commit peers**

For example, to synchronize a configuration from Switch A to Switch B, Switch C, and Switch D, configure the following on Switch A:

Switch A

```
[edit system commit]
peers {
   switchB {
      user admin-swB;
      authentication "$ABC123";
   }
   switchC {
      user admin-swC;
      authentication ""$ABC123";
   }
   switchD {
      user admin-swD;
```

```
    authentication "$ABC123";
  }
}
  [edit system]
    static-host-mapping [
      SwitchA{
        inet [ 10.92.76.2 ];
      }
      SwitchB{
        inet [ 10.92.76.4 ];
      }
      SwitchC{
        inet [ 10.92.76.6 ];
      }
      SwitchD{
        inet [ 10.92.76.8 ];
      }
    }
  }
```

If you only want to synchronize configurations from Switch A to Switch B, Switch C, and Switch D, you do not need to configure the **peers** statement on Switch B, Switch C, and Switch D.

The configuration details from the peers statements are also used to establish a NETCONF over SSH connection between the devices. To enable NETCONF over SSH, issue the **set system services netconf ssh** command on all devices.

How Configurations and Commits Are Synchronized Between Devices

The local (or requesting) device on which you enable the **peers-synchronize** statement or issue the **commit peers-synchronize** command copies and loads its configuration to the remote (or responding) device. Each device then performs a syntax check on the configuration file being committed. If no errors are found, the configuration is activated and becomes the current operational configuration on all devices. The commits are propagated using a remote procedural call (RPC).

The following events occur during configuration synchronization:

1. The local device sends the sync-peers.conf file (the configuration that will be shared with the devices specified in the conditional group) to the remote devices.
2. The remote devices load the configuration, send the results of the load to the local device, export their configuration to the local device, and reply that the commit is complete.
3. The local device reads the replies from the remote devices.
4. If successful, the configuration is committed.

Configuration synchronization is not successful if either a) the remote device is unavailable or b) the remote device is reachable, but there are failures due to the following reasons:

- SSH connection fails because of user and authentication issues.
- Junos OS RPC fails because a lock cannot be obtained on the remote database.
- Loading the configuration fails because of syntax problems.
- Commit check fails.

The **peers-synchronize** statement uses the hostname or IP address, username, and password for the devices you configured in the **peers** statement. With the **peers-synchronize** statement enabled, you can simply issue the **commit** command to synchronize the configuration from one device to another. For example, if you configured the **peers** statement on the local device, and want to synchronize the configuration with the remote device, you can simply issue the **commit** command on the local device. However, if you issue the **commit** command on the local device and the remote device is not reachable, you will receive a warning message saying that the remote device is not reachable and only the configuration on the local device is committed:

Here is an example warning message:

```
error: netconf: could not read hello
error: did not receive hello packet from server
error: Setting up sessions for peer: 'peer1' failed
warning: Cannot connect to remote peers, ignoring it
commit complete
```

If you do not have the **peers** statement configured with the remote device information and you issue the **commit** command, only the configuration on the local device is committed. If the remote device is unreachable and there are other failures, the commit is unsuccessful on both the local and remote devices.

NOTE: When you enable the peers-synchronize statement and issue the commit command, the commit might take longer than a normal commit. Even if the configuration is the same across the devices and does not require synchronization, the system still attempts to synchronize the configurations.

The **commit peers-synchronize** command also uses the hostname or IP address, username, and password for the devices configured in the **peers** statement. If you issue the **commit peers-synchronize** command on the local device to synchronize the configuration with the remote device and the remote device is reachable but there are other failures, the commit fails on both the local and remote devices.

Synchronizing and Committing Configurations

To propagate, synchronize, and commit configuration changes from one device (Junos Fusion Provider Edge, Junos Fusion Enterprise, EX Series switches, and MX Series routers) to another, perform following tasks:

- Configure Devices for Configuration Synchronization on page 360
- Create a Global Configuration Group on page 361

Configure Devices for Configuration Synchronization

Configure the hostnames or IP addresses for the devices that will be synchronizing their configurations as well as the usernames and authentication details for the users administering configuration synchronization. Additionally, enable a NETCONF connection so that the devices can synchronize their configurations. Secure Copy Protocol (SCP) copies the configurations securely between the devices.

For example, if you have a local device named Switch A and want to synchronize a configuration with remote devices named Switch B, Switch C, and Switch D, you need to configure the details for Switch B, Switch C, and Switch D on Switch A.

To specify the configuration details:

1. On the local device, specify the configuration details for the remote device.

   ```
   [edit system commit]
   user@switch# set peers hostname user username authentication password string
   ```

   For example, if the local device is Switch A, and the remote devices are Switch B, Switch C, and Switch D:

   ```
   [edit system commit]
   user@Switch A# set peers Switch B user admin-SwitchB authentication "$ABC123"
   user@Switch A# set peers Switch C user admin-SwitchC authentication "$ABC123"
   user@Switch A# set peers Switch D user admin-SwitchD authentication "$ABC123"
   ```

   The password string is stored as an authenticated password string.

   The output for Switch A is as follows:

   ```
   [edit system commit]
     peers {
     Switch B{
       user admin-SwitchB;
       authentication "$ABC123";
     }
     Switch C{
       user admin-SwitchC;
       authentication "$ABC123";
     }
     Switch D{
       user admin-SwitchD;
       authentication "$ABC123";
     }
   }
   ```

2. Statically map Switch A to Switch B, Switch C, and Switch D.

   For example:

```
[edit system ]
user@Switch A# set static-host-mapping Switch A inet 10.92.76.2
user@Switch A# set static-host-mapping Switch B inet 10.92.76.4
user@Switch A# set static-host-mapping Switch C inet 10.92.76.6
user@Switch A# set static-host-mapping Switch D inet 10.92.76.8
```

   The output is as follows:

```
[edit system]
  static-host-mapping [
    SwitchA{
      inet [ 10.92.76.2 ];
    }
    SwitchB{
      inet [ 10.92.76.4 ];
    }
    SwitchC{
      inet [ 10.92.76.6 ];
    }
    SwitchD{
      inet [ 10.92.76.8 ];
    }
  }
}
```

3. Enable a NETCONF connection using SSH between all devices (Switch A, Switch B, Switch C, and Switch D).

   For example:

```
[edit]
user@Switch A# set system services netconf ssh

[edit]
user@Switch B# set system services netconf ssh

[edit]
user@Switch C# set system services netconf ssh

[edit]
user@Switch D# set system services netconf ssh
```

Create a Global Configuration Group

Create a global configuration group the local and remote devices.

To create a global configuration group:

1. Specify the devices that will receive the configuration:

```
[edit]
```

```
user@switch# set groups <name of group> when peers [<name of local peer> <name of remote peer>]
```

For example:

```
[edit]
user@switch# set groups global when peers [Switch A Switch B Switch C Switch D]
```

2. Create the global configuration that will be shared between the devices.

   For example:

```
interfaces {
  ge-0/0/0 {
    unit 0 {
      family inet {
        address 10.1.1.1/8;
      }
    }
  }
  ge-0/0/1 {
    ether-options {
      802.3ad ae0;
    }
  }
  ge-0/0/2 {
    ether-options {
      802.3ad ae1;
    }
  }
  ae0 {
    aggregated-ether-options {
      lacp {
        active;
      }
    }
    unit 0 {
      family ethernet-switching {
        interface-mode trunk;
        vlan {
          members v1;
        }
      }
    }
  }
  ae1 {
    aggregated-ether-options {
      lacp {
        active;
        system-id 00:01:02:03:04:05;
        admin-key 3;
      }
      mc-ae {
        mc-ae-id 1;
        redundancy-group 1;
```

```
            mode active-active;
          }
        }
        unit 0 {
          family ethernet-switching {
            interface-mode access;
            vlan {
              members v1;
            }
          }
        }
      }
    }
    switch-options {
      service-id 1;
    }
    vlans {
      v1 {
        vlan-id 100;
        l3-interface irb.100;
      }
    }
```

The output for the configuration is as follows:

```
groups {
  global {
    when {
      peers [ Switch A Switch B Switch C Switch D ];
    }
    interfaces {
      ge-0/0/0 {
        unit 0 {
          family inet {
            address 10.1.1.1/8;
          }
        }
      }
      ge-0/0/1 {
        ether-options {
          802.3ad ae0;
        }
      }
      ge-0/0/2 {
        ether-options {
          802.3ad ae1;
        }
      }
      ae0 {
        aggregated-ether-options {
          lacp {
            active;
          }
```

```
        }
        unit 0 {
          family ethernet-switching {
            interface-mode trunk;
            vlan {
              members v1;
            }
          }
        }
      }
      ae1 {
        aggregated-ether-options {
          lacp {
            active;
            system-id 00:01:02:03:04:05;
            admin-key 3;
          }
          mc-ae {
            mc-ae-id 1;
            redundancy-group 1;
            mode active-active;
          }
        }
        unit 0 {
          family ethernet-switching {
            interface-mode access;
            vlan {
              members v1;
            }
          }
        }
      }
    }
    switch-options {
      service-id 1;
    }
    vlans {
      v1 {
        vlan-id 100;
        l3-interface irb.100;
      }
    }
  }
}
```

Create a Local Configuration Group

Create a local configuration group for the local device.

To create a local configuration group:

1. Specify the local configuration group name.

```
[edit]
user@switch# set groups name of group when peers [name of local peer]
```

For example:

```
[edit]
user@switch# set groups local when peers [Switch A]
```

2. Include the local configuration that will be used by the local device.

   For example:

```
interfaces {
   ae1 {
      aggregated-ether-options {
         mc-ae {
            chassis-id 0;
            status-control active;
            events {
               iccp-peer-down {
                  prefer-status-control-active;
               }
            }
         }
      }
   }
   irb {
      unit 100 {
         family inet {
            address 10.10.10.3/8 {
               arp 10.10.10.2 l2-interface ae0.0 mac 00:00:5E:00:53:00;
            }
         }
      }
   }
}
      multi-chassis {
         multi-chassis-protection 10.1.1.1 {
            interface ae0;
         }
      }
   }
}
```

The output for the configuration is as follows:

```
groups {
   local {
      when {
         peers Switch A;
      }
      interfaces {
         ae1 {
            aggregated-ether-options {
               mc-ae {
                  chassis-id 0;
                  status-control active;
```

```
            events {
              iccp-peer-down {
                prefer-status-control-active;
              }
            }
          }
        }
      }
      irb {
        unit 100 {
          family inet {
            address 10.10.10.3/8 {
              arp 10.10.10.2 l2-interface ae0.0 mac 00:00:5E:00:53:00;
            }
          }
        }
      }
    }
    multi-chassis {
      multi-chassis-protection 10.1.1.1 {
        interface ae0;
      }
    }
  }
}
```

Create a Remote Configuration Group

Create a remote configuration group for remote devices.

To create a remote configuration group:

1. Specify the remote configuration group name.

   ```
   [edit]
   user@switch# set groups name of group when peers [names of remote peers]
   ```

   For example:

   ```
   [edit]
   user@switch# set groups remote when peers [Switch B Switch C Switch D]
   ```

2. Include the remote configuration that will be used by the remote devices.

   For example:

   ```
   interfaces {
     ae1 {
       aggregated-ether-options {
         mc-ae {
           chassis-id 1;
           status-control standby;
           events {
             iccp-peer-down {
               prefer-status-control-active;
   ```

```
            }
          }
        }
      }
    }
    irb {
      unit 100 {
        family inet {
          address 10.10.10.3/8 {
            arp 10.10.10.2 l2-interface ae0.0 mac 00:00:5E:00:53:00;
          }
        }
      }
    }
  }
    multi-chassis {
      multi-chassis-protection 10.1.1.1 {
        interface ae0;
      }
    }
  }
```

The output for the configuration is as follows:

```
groups {
  remote {
    when {
      peers Switch B Switch C Switch D
    }
    interfaces {
      ae1 {
        aggregated-ether-options {
          mc-ae {
            chassis-id 1;
            status-control standby;
            events {
              iccp-peer-down {
                prefer-status-control-active;
              }
            }
          }
        }
      }
      irb {
        unit 100 {
          family inet {
            address 10.10.10.3/8 {
              arp 10.10.10.2 l2-interface ae0.0 mac 00:00:5E:00:53:00;
            }
          }
        }
      }
    }
```

```
    multi-chassis {
      multi-chassis-protection 10.1.1.1 {
        interface ae0;
      }
    }
  }
}
```

Create Apply Groups for the Local, Remote, and Global Configurations

Create apply groups so changes in the configuration are inherited by local, remote, and global configuration groups. List the configuration groups in order of inheritance, where the configuration data in the first configuration group takes priority over the data in subsequent configuration groups.

When you apply the configuration groups and issue the **commit peers-synchronize** command, changes are committed on both the local and remote devices. If there is an error on any of the devices, an error message is issued, and the commit is aborted.

To apply the configuration groups:

1. Specify the names of the configuration groups.

   ```
   [edit]
   user@switch# set apply-groups [<name of global configuration group> <name of local configuration group> <name of remote configuration group>]
   ```

   For example:

   ```
   [edit]
   user@switch# set apply-groups [ global local remote ]
   ```

   The output for the configuration is as follows:

   ```
   apply-groups [ global local remote ];
   ```

Synchronizing and Committing Configurations

NOTE: The commit at <"string"> command is not supported when performing configuration synchronization.

You can enable the **peers-synchronize** statement on the local (or requesting) device to copy and load its configuration to the remote (or responding) device by default. You can alternatively issue the **commit peers-synchronize** command.

- Configure the **commit** command on the local (or requesting) to automatically perform a **peers-synchronize** action between devices.

  ```
  [edit]
  user@switch# set system commit peers-synchronize
  ```

The output for the configuration is as follows:

```
system {
  commit {
    peers-synchronize;
  }
}
```

- Issue the **commit peers-synchronize** command on the local (or requesting) device.

```
[edit]
user@switch# commit peers-synchronize
```

Troubleshooting Remote Device Connections

Problem **Description:**

When you issue the **commit** command, the system issues the following error message:

```
root@Switch A# commit
error: netconf: could not read hello error: did not receive hello packet from server error: Setting
up sessions for peer: 'Switch B' failed warning: Cannot connect to remote peers, ignoring it
```

The error message shows that there is a NETCONF connection issue between the local device and remote device.

Resolution

1. Verify that the SSH connection to the remote device (Switch B) is working.

   ```
   root@Switch A# ssh root@Switch B
   @@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@
   @ WARNING: REMOTE HOST IDENTIFICATION HAS CHANGED! @
   @@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@
   IT IS POSSIBLE THAT SOMEONE IS DOING SOMETHING NASTY! Someone could be
   eavesdropping on you right now (man-in-the-middle attack)! It is also possible that a host
   key has just been changed. The fingerprint for the ECDSA key sent by the remote host is
   21:e8:5a:58:bb:29:8b:96:a4:eb:cc:8a:32:95:53:c0. Please contact your system administrator.
   Add correct host key in /root/.ssh/known_hosts to get rid of this message. Offending ECDSA
   key in /root/.ssh/known_hosts:1 ECDSA host key for Switch A has changed and you have
   requested strict checking. Host key verification failed.
   ```

   The error message shows that the SSH connection is not working.

2. Delete the key entry in the **/root/.ssh/known_hosts:1** directory and try to connect to Switch B again.

   ```
   root@Switch A# ssh root@Switch B
   The authenticity of host 'Switch B (10.92.76.235)' can't be established. ECDSA key fingerprint
   is 21:e8:5a:58:bb:29:8b:96:a4:eb:cc:8a:32:95:53:c0. Are you sure you want to continue
   connecting (yes/no)? yes Warning: Permanently added 'Switch A,10.92.76.235' (ECDSA) to
   the list of known hosts. Password: Last login: Wed Apr 13 15:29:58 2016 from 192.168.61.129 -
   JUNOS 15.1I20160412_0929_dc-builder Kernel 64-bit FLEX
   JNPR-10.1-20160217.114153_fbsd-builder_stable_10 At least one package installed on this
   device has limited support. Run 'file show /etc/notices/unsupported.txt' for details.
   ```

   Connection to Switch B was successful.

3. Log out of Switch B.

   ```
   root@Switch B# exit
   logout Connection to Switch B closed.
   ```

4. Verify that NETCONF over SSH is working.

   ```
   root@Switch A# ssh root@Switch B -s netconf
   logout Connection to st-72q-01 closed.
   Password:
   <hello xmlns="urn:ietf:params:xml:ns:netconf:base:1.0">
   <capabilities>
   <capability>urn:ietf:params:netconf:base:1.0</capability>
   <capability>urn:ietf:params:netconf:capability:candidate:1.0</capability>
   ```

   The log message shows that the NETCONF over SSH was successful.

   If the error message showed that NETCONF over SSH was not successful, enable NETCONF over SSH by issuing the **set system services netconf ssh** command.

5. Create configuration groups to synchronize if you have not done so already.

   You can issue the **show | compare** command to see if any configuration groups have

been created.

```
root@Switch A# show | compare
```

6. Issue the **commit** command.

```
root@Switch A# commit
[edit chassis]
configuration check succeeds
commit complete
{master:0}[edit]
```

The log message shows that the commit was successful.

Release History Table

Release	Description
17.4R1	On Junos Fusion Data Center, support for configuration synchronization started with Junos OS Release 17.4R1.
15.1X53-D60	On QFX Series switches, support for configuration synchronization started with Junos OS Release 15.1X53-D60.
14.2R6	On MX Series routers and Junos Fusion, support for configuration synchronization started with Junos OS Release 14.2R6.

Understanding Multichassis Link Aggregation Group Configuration Consistency Check

When there is a Multichassis Link Aggregation Group (MC-LAG) inconsistency, you are notified and can take action to resolve it. An example of an inconsistency is configuring identical chassis IDs on both peers instead of configuring unique chassis IDs on both peers. Only committed MC-LAG parameters are checked for consistency.

Benefits of Using MC-LAG Consistency Check

- This feature helps you find configuration-parameter inconsistencies between multichassis link aggregation group (MC-LAG) peers.

How MC-LAG Consistency Checks Work

The following events take place during configuration consistency check after you issue a commit on the local MC-LAG peer:

1. Commit an MC-LAG configuration on the local MC-LAG peer.
2. ICCP parses the MC-LAG configuration and then sends the configuration to the remote MC-LAG peer.
3. The remote MC-LAG peer receives the MC-LAG configuration from the local MC-LAG peer and compares it with its own MC-LAG configuration.

   If the there is a severe inconsistency between the two MC-LAG configurations, the MC-LAG interface is brought down, and syslog messages are issued.

   If there is a moderate inconsistency between the two configurations, syslog messages are issued.

The following events take place during configuration consistency check after you issue a commit on the remote MC-LAG peer:

- Commit an MC-LAG configuration on the remote MC-LAG peer.
- ICCP parses the MC-LAG configuration and then sends the configuration to the local MC-LAG peer.
- The local MC-LAG peer receives the configuration from the remote MC-LAG peer and compares it with its own configuration.

  If the there is a severe inconsistency between the two configurations, the MC-LAG interface is brought down, and syslog messages are issued.

  If there is a moderate inconsistency between the two configurations, syslog messages are issued.

Configuration Consistency Requirements

There are different configuration consistency requirements depending on the MC-LAG parameters. The consistency requirements are either identical or unique, which means that some parameters must be configured identically and some must be configured uniquely on the MC-LAG peers. For example, the chassis ID must be unique on both peers, whereas the LACP mode must be identical on both peers.

The enforcement level of the consistency requirements (identical or unique) is either mandatory or desired. When the enforcement level is mandatory, and you configure the MC-LAG parameter incorrectly, the system brings down the interface and issues a syslog message.

For example, you receive a syslog message that says, **"Some of the Multichassis Link Aggregation (MC-LAG) configuration parameters between the peer devices are not consistent. The concerned MC-LAG interfaces were explictly brought down to prevent unwanted behavior."** When you correct the inconsistency, and issue a successful commit, the system will bring up the interface. When the enforcement is desired, and you configure the MC-LAG parameter incorrectly, you receive a syslog message that says, **"Some of the Multichassis Link Aggregation(MC-LAG) configuration parameters between the peer devices are not consistent. This may lead to sub-optimal performance of the feature."** As noted in the syslog message, performance will be sub-optimal in this situation. You can also issue the **show interfaces mc-ae** command to display the configuration consistency check status of the multichassis aggregated Ethernet interface.

If there are multiple inconsistencies, only the first inconsistency is shown. If the enforcement level for an MC-LAG parameter is mandatory, and you did not configure that parameter correctly, the command shows that the MC-LAG interface is down.

When Remote Peers are Not Reachable

When you issue a commit on the local peer, and the remote peer is not reachable, configuration consistency check will pass so that the local peer can come up in standalone mode. When the remote peer comes up, ICCP exchanges the configurations between the peers. If the consistency check fails, the MC-LAG interface goes down, and the system notifies you of the parameter that caused the inconsistency. When you correct the inconsistency, and issue a successful commit, the system brings up the interface.

Enabling MC-LAG Configuration Consistency Checking

Consistency check is not enabled by default. To enable consistency check, issue the **set multi-chassis mc-lag consistency-check** command.

Table 12 on page 373 provides a sample list of committed MC-LAG parameters that are checked for consistency, along with their consistency requirements (identical or unique), hierarchies in which the parameters are configured, and the consistency check enforcement levels (mandatory or desired).

Table 12: MC-LAG Parameters Checked for Configuration Consistency

Configuration Knob	Hierarchy	Consistency Requirement	Enforcement
session-establishment-hold-time Specify the time during which an Inter-Chassis Control Protocol (ICCP) connection must be established between peers.	**Global, ICCP Peer**	**Identical**	**Mandatory**
mac-limit Specify the maximum number of MAC addresses to be associated with a VLAN—the default is unlimited, which can leave the network vulnerable to flooding.	**Global**	**Identical**	**Desired**

Table 12: MC-LAG Parameters Checked for Configuration Consistency (continued)

Configuration Knob	Hierarchy	Consistency Requirement	Enforcement
mac-aging-timer Specify how long MAC addresses remain in the Ethernet switching table.	**Global**	**Identical**	**Desired**
arp-aging-timer Specify the ARP aging timer in minutes for a logical interface of **inet**.	**Global**	**Identical**	**Desired**
rstp-system-identifier Specify different bridge identifiers for different RSTP routing instances.	**Global**	**Identical**	**Desired**
mstp-system-identifier Specify different bridge identifiers for different MSTP routing instances.	**Global**	**Identical**	**Desired**
rstp-bridge-priority Determine which bridge is elected as the root bridge for RSTP. If two bridges have the same path cost to the root bridge, the bridge priority determines which bridge becomes the designated bridge for a LAN segment.	**Global**	**Identical**	**Desired**
mstp-bridge-priority Determine which bridge is elected as the root bridge for MSTP. If two bridges have the same path cost to the root bridge, the bridge priority determines which bridge becomes the designated bridge for a LAN segment.	**Global**	**Identical**	**Desired**
rstp-bpdu-block-on-edge Configure bridge protocol data unit (BPDU) protection on all edge ports of a switch for RSTP.	**Global**	**Identical**	**Desired**
vstp-bpdu-block-on-edge Configure bridge protocol data unit (BPDU) protection on all edge ports of a switch for VSTP.	**Global**	**Identical**	**Desired**

Table 12: MC-LAG Parameters Checked for Configuration Consistency (continued)

Configuration Knob	Hierarchy	Consistency Requirement	Enforcement
mstp-bpdu-block-on-edge Configure bridge protocol data unit (BPDU) protection on all edge ports of a switch for MSTP.	**Global**	**Identical**	**Desired**
service-id Specify a service identifier for each multichassis aggregated Ethernet interface that belongs to a link aggregation group (LAG).	**Global**	**Identical**	**Mandatory**
bfd minimum-interval Configure the minimum interval after which the local routing device transmits hello packets and then expects to receive a reply from a neighbor with which it has established a BFD session.	**ICCP Peer**	**Identical**	**Mandatory**
iccp/minimum-transmit-interval Specify the minimum interval at which the local routing device transmits hello packets to a neighbor with which it has established a BFD session.	**ICCP Peer**	**Identical**	**Mandatory**
iccp/minimum-receive-interval Specify the minimum interval after which the local routing device must receive a reply from a neighbor with which it has established a BFD session.	**ICCP Peer**	**Identical**	**Mandatory**
iccp/bfd multiplier Configure the number of hello packets not received by a neighbor that causes the originating interface to be declared down.	**ICCP Peer**	**Identical**	**Mandatory**
iccp single-hop Configure single hop BFD sessions.	**ICCP Peer**	**Identical**	**Mandatory**
iccp/authentication-key Specify the authentication key password to verify the authenticity of packets sent from the peers hosting an MC-LAG.	**ICCP Peer**	**Identical**	**Mandatory**

Table 12: MC-LAG Parameters Checked for Configuration Consistency (continued)

Configuration Knob	Hierarchy	Consistency Requirement	Enforcement
redundancy-group-id-list Specify the redundancy group identification number. The Inter-Chassis Control Protocol (ICCP) uses the redundancy group ID to associate multiple chassis that perform similar redundancy functions.	ICCP Peer	Identical	Mandatory
backup-liveness-detection Determine whether a peer is up or down by exchanging keepalive messages over the management link between the two Inter-Chassis Control Protocol (ICCP) peers.	ICCP Peer	Unique	Mandatory
mc-ae-id Specify the identification number of the MC-LAG device.	MCAE ifd	Identical	Mandatory
mcae redundancy-group Used by ICCP to associate multiple chassis that perform similar redundancy functions and to establish a communication channel so that applications on peering chassis can send messages to each other.	MCAE ifd	Identical	Mandatory
mcae chassis-id Used by LACP for calculating the port number of the MC-LAG's physical member links.	MCAE ifd	Unique	Mandatory
mcae deployment mode Indicates whether an MC-LAG is in active-standby mode or active-active mode.	MCAE ifd	Identical	Mandatory
mcae status-control Specify whether the chassis becomes active or remains in standby mode when an interchassis link failure occurs.	MCAE ifd	Unique	Mandatory

Table 12: MC-LAG Parameters Checked for Configuration Consistency (continued)

Configuration Knob	Hierarchy	Consistency Requirement	Enforcement
force-icl-down Specify that if the node's ICCP peer goes down, the system brings down the interchassis-link logical interface.	**MCAE ifd**	**Unique**	**Mandatory**
prefer-status-control-active Specify that the node configured as **status-control active** becomes the active node if the peer of this node goes down.	**MCAE ifd**	**Unique**	**Desired**
lacp mode Specify LACP is active or passive.	**MCAE ifd**	**Identical**	**Mandatory**
lacp periodic Specify the interval for periodic transmission of LACP packets.	**MCAE ifd**	**Identical**	**Mandatory**
lacp system-id Define the LACP system identifier at the aggregated Ethernet interface level.	**MCAE ifd**	**Identical**	**Mandatory**
lacp admin-key Specify an administrative key for the router or switch.	**MCAE ifd**	**Identical**	**Mandatory**
native-vlan-id Configure mixed tagging support for untagged packets on a port.	**MCAE ifd**	**Identical**	**Mandatory**
mcae-mac-synchronize Synchronize the MAC addresses for the Layer 3 interfaces of the switches participating in the MC-LAG.	**VLAN**	**Identical**	**Mandatory**
Interface mac Limit 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.	**VLAN**	**Identical**	**Desired**

Table 12: MC-LAG Parameters Checked for Configuration Consistency (continued)

Configuration Knob	Hierarchy	Consistency Requirement	Enforcement
l3-interface Associate a Layer 3 interface with the VLAN.	**VLAN**	**Identical**	**Desired**
igmp-snooping Enable IGMP snooping. A Layer 2 device monitors the IGMP join and leave messages sent from each connected host to a multicast router. This enables the Layer 2 device to keep track of the multicast groups and associated member ports. The Layer 2 device uses this information to make intelligent decisions and to forward multicast traffic to only the intended destination hosts.	**VLAN**	**Identical**	**Mandatory**
family Specify the protocol family configured on the logical interface.	**IRB Interface**	**Identical**	**Mandatory**
ipv4 address Specify an IPv4 address for the IRB interface.	**IRB Interface**	**Unique**	**Mandatory**
ipv6 address Specify an IPv6 address for the IRB interface.	**IRB Interface**	**Unique**	**Mandatory**
vrrp-group id Specify a VRRP group identifier.	**IRB Interface**	**Identical**	**Mandatory**
proxy-arp-type For Ethernet interfaces only, configure the router or switch to respond to any ARP request, as long as the router or switch has an active route to the ARP request's target address.	**IRB Interface**	**Identical**	**Mandatory**
vrrp-group priority Configure a Virtual Router Redundancy Protocol (VRRP) router's priority for becoming the master default router. The router with the highest priority within the group becomes the master.	**VRRP Group**	**Unique**	**Mandatory**

Table 12: MC-LAG Parameters Checked for Configuration Consistency (continued)

Configuration Knob	Hierarchy	Consistency Requirement	Enforcement
vrrp-group authentication-key Configure a Virtual Router Redundancy Protocol (VRRP) IPv4 authentication key. You also must specify a VRRP authentication scheme by including the authentication-type statement.	**VRRP Group**	**Identical**	**Mandatory**
vrrp-group authentication-type Enable Virtual Router Redundancy Protocol (VRRP) IPv4 authentication and specify the authentication scheme for the VRRP group.	**VRRP Group**	**Identical**	**Mandatory**
vrrp-group virtual-address Configure the addresses of the virtual routers in a Virtual Router Redundancy Protocol (VRRP) IPv4 or IPv6 group.	**VRRP Group**	**Identical**	**Mandatory**
encapsulation Configure a logical link-layer encapsulation type.	**MCAE ifd**	**Identical**	**Mandatory**
flexible-vlan-tagging Support simultaneous transmission of 802.1Q VLAN single-tag and dual-tag frames on logical interfaces on the same Ethernet port, and on pseudowire logical interfaces.	**MCAE ifd**	**Identical**	**Mandatory**
vlan-tagging For Fast Ethernet and Gigabit Ethernet interfaces, aggregated Ethernet interfaces configured for VPLS, and pseudowire subscriber interfaces, enable the reception and transmission of 802.1Q VLAN-tagged frames on the interface.	**MCAE ifd**	**Identical**	**Mandatory**
mtu Specify the maximum transmission unit (MTU) size for the media or protocol.	**MCAE ifd, ICL ifd**	**Identical**	**Mandatory**

Table 12: MC-LAG Parameters Checked for Configuration Consistency (continued)

Configuration Knob	Hierarchy	Consistency Requirement	Enforcement
interface-mode Determine whether the logical interface accepts or discards packets based on VLAN tags.	**MCAE ifl**	**Identical**	**Mandatory**
vlan membership Specify the name of the VLAN that belongs to an interface.	**MCAE ifl**	**Identical**	**Mandatory**

Learning the Status of a Configuration Consistency Check

The following commands provide information regarding the status of configuration consistency check:

- Issue the show multi-chassis mc-lag configuration-consistency list-of-parameters command to view the list of committed MC-LAG parameters that are checked for inconsistencies, the consistency requirement (identical or unique), and the enforcement level (mandatory or desired).
- Issue the show multi-chassis mc-lag configuration-consistency command to view the list of committed MC-LAG parameters that are checked for inconsistencies, the consistency requirement (identical or unique), the enforcement level (mandatory or desired), and the result of the configuration consistency check. The results are either pass or fail.
- Issue the show multi-chassis mc-lag configuration-consistency global-config command to view configuration consistency check status for all global configuration related to MC-LAG functionality, the consistency requirement (identical or unique), the enforcement level (mandatory or desired), and the result of the configuration consistency check. The results are either pass or fail..
- Issue the show multi-chassis mc-lag configuration-consistency icl-config command to view configuration consistency check status for parameters related to the interchassis control link, the consistency requirement (identical or unique), the enforcement level (mandatory or desired), and the result of the configuration consistency check. The results are either pass or fail.
- Issue the show multi-chassis mc-lag configuration-consistency mcae-config command to view configuration consistency check status for parameters related to the multichassis aggregated Ethernet interface, the consistency requirement (identical or unique), the enforcement level (mandatory or desired), and the result of the configuration consistency check. The results are either pass or fail.
- Issue the show multi-chassis mc-lag configuration-consistency vlan-config command to view configuration consistency check status for parameters related to VLAN configuration, the consistency requirement (identical or unique), the enforcement level

(mandatory or desired), and the result of the configuration consistency check. The results are either pass or fail..

- Issue the show multi-chassis mc-lag configuration-consistency vrrp-config command to view configuration consistency check status for parameters related to VRRP configuration, the consistency requirement (identical or unique), the enforcement level (mandatory or desired), and the result of the configuration consistency check. The results are either pass or fail.
- Issue the show interfaces mc-ae command to view configuration consistency check status of the multichassis aggregated Ethernet interface. If there are multiple inconsistencies, only the first inconsistency is shown. If the enforcement level for the MC-LAG parameter is mandatory, and you did not configure that parameter correctly, the command will show that the MC-LAG interface is down.

Support for MC-LAG Configuration Consistency Checking

Both EX Series switches and QFX Series switches support MC-LAG configuration consistency checking.

Starting with Junos OS Release 15.1X53-D60 on QFX10000 switches, configuration consistency check uses the Inter-Chassis Control Protocol (ICCP) to exchange MC-LAG configuration parameters (chassis ID, service ID, and so on) and checks for any configuration inconsistencies across MC-LAG peers.

Release History Table

Release	Description
15.1X53-D60	Starting with Junos OS Release 15.1X53-D60 on QFX10000 switches, configuration consistency check uses the Inter-Chassis Control Protocol (ICCP) to exchange MC-LAG configuration parameters (chassis ID, service ID, and so on) and checks for any configuration inconsistencies across MC-LAG peers.

Related Documentation

- Configuring MC-LAG on EX9200 Switches in the Core for Campus Networks

Extending an MC-LAG Topology Using EVPN-MPLS

- Understanding EVPN-MPLS Interworking with Junos Fusion Enterprise and MC-LAG on page 382
- Example: EVPN-MPLS Interworking With an MC-LAG Topology on page 386

Understanding EVPN-MPLS Interworking with Junos Fusion Enterprise and MC-LAG

Starting with Junos OS Release 17.4R1, you can use Ethernet VPN (EVPN) to extend a Junos Fusion Enterprise or multichassis link aggregation group (MC-LAG) network over an MPLS network to a data center or campus network. With the introduction of this feature, you can now interconnect dispersed campus and data center sites to form a single Layer 2 virtual bridge.

Figure 37 on page 382 shows a Junos Fusion Enterprise topology with two EX9200 switches that serve as aggregation devices (PE2 and PE3) to which the satellite devices are multihomed. The two aggregation devices use an interchassis link (ICL) and the Inter-Chassis Control Protocol (ICCP) protocol from MC-LAG to connect and maintain the Junos Fusion Enterprise topology. PE1 in the EVPN-MPLS environment interworks with PE2 and PE3 in the Junos Fusion Enterprise with MC-LAG.

Figure 37: EVPN-MPLS Interworking with Junos Fusion Enterprise

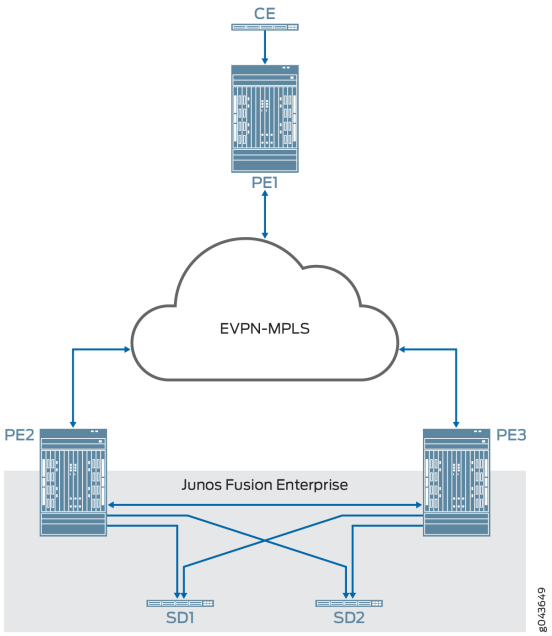

Figure 38 on page 383 shows an MC-LAG topology in which customer edge (CE) device CE1 is multihomed to PE2 and PE3. PE2 and PE3 use an ICL and the ICCP protocol from MC-LAG to connect and maintain the topology. PE1 in the EVPN-MPLS environment interworks with PE2 and PE3 in the MC-LAG environment.

Figure 38: EVPN-MPLS Interworking with MC-LAG

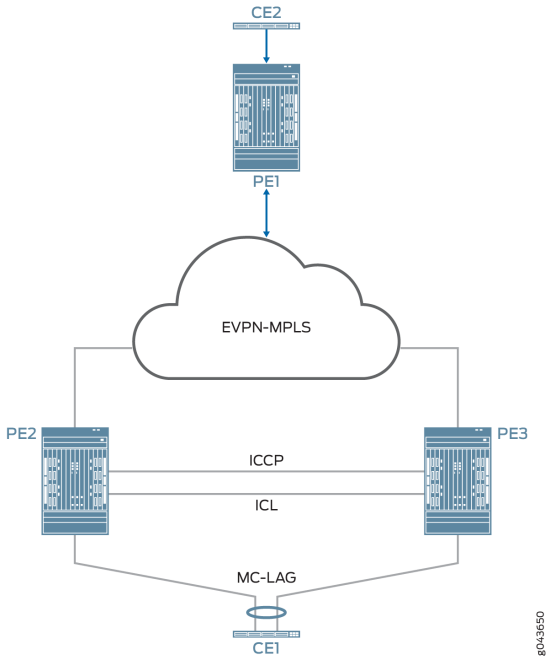

Throughout this topic, Figure 37 on page 382 and Figure 38 on page 383 serve as references to illustrate various scenarios and points.

The use cases depicted in Figure 37 on page 382 and Figure 38 on page 383 require the configuration of both EVPN multihoming in active-active mode and MC-LAG on PE2 and PE3. EVPN with multihoming active-active and MC-LAG have their own forwarding logic for handling traffic, in particular, broadcast, unknown unicast, and multicast (BUM) traffic. At times, the forwarding logic for EVPN with multihoming active-active and MC-LAG contradict each other and causes issues. This topic describes the issues and how the EVPN-MPLS interworking feature resolves these issues.

NOTE:

Other than the EVPN-MPLS interworking-specific implementations described in this topic, EVPN-MPLS, Junos Fusion Enterprise, and MC-LAG offer the same functionality and function the same as the standalone features.

Benefits of Using EVPN-MPLS with Junos Fusion Enterprise and MC-LAG

Use EVPN-MPLS with Junos Fusion Enterprise and MC-LAG to interconnect dispersed campus and data center sites to form a single Layer 2 virtual bridge.

BUM Traffic Handling

In the use cases shown in Figure 37 on page 382 and Figure 38 on page 383, PE1, PE2, and PE3 are EVPN peers, and PE2 and PE3 are MC-LAG peers. Both sets of peers exchange control information and forward traffic to each other, which causes issues. Table 13 on page 384 outlines the issues that arise, and how Juniper Networks resolves the issues in its implementation of the EVPN-MPLS interworking feature.

Table 13: BUM Traffic: Issues and Resolutions

BUM Traffic Direction	EVPN Interworking with Junos Fusion Enterprise and MC-LAG Logic	Issue	Juniper Networks Implementation Approach
North bound (PE2 receives BUM packet from a locally attached single- or dual-homed interfaces).	PE2 floods BUM packet to the following: • All locally attached interfaces, including the ICL, for a particular broadcast domain. • All remote EVPN peers for which PE2 has received inclusive multicast routes.	Between PE2 and PE3, there are two BUM forwarding paths—the MC-LAG ICL and an EVPN-MPLS path. The multiple forwarding paths result in packet duplication and loops.	• BUM traffic is forwarded on the ICL only. • Incoming traffic from the EVPN core is not forwarded on the ICL. • Incoming traffic from the ICL is not forwarded to the EVPN core.
South bound (PE1 forwards BUM packet to PE2 and PE3).	PE2 and PE3 both receive a copy of the BUM packet and flood the packet out of all of their local interfaces, including the ICL.	PE2 and PE3 both forward the BUM packet out of the ICL, which results in packet duplication and loops.	

Split Horizon

In the use cases shown in Figure 37 on page 382 and Figure 38 on page 383, split horizon prevents multiple copies of a BUM packet from being forwarded to a CE device (satellite device). However, the EVPN-MPLS and MC-LAG split horizon implementations contradict each other, which causes an issue. Table 14 on page 385 explains the issue and how Juniper Networks resolves it in its implementation of the EVPN-MPLS interworking feature.

Table 14: BUM Traffic: Split Horizon-Related Issue and Resolution

BUM Traffic Direction	EVPN Interworking with Junos Fusion Enterprise and MC-LAG Logic	Issue	Juniper Networks Implementation Approach

North bound (PE2 receives BUM packet from a locally attached dual-homed interface).	• Per EVPN-MPLS forwarding logic:
	  ◦ Only the designated forwarder (DF) for the Ethernet segment (ES) can forward BUM traffic.
	  ◦ The local bias rule, in which the local peer forwards the BUM packet and the remote peer drops it, is not supported.
	• Per MC-LAG forwarding logic, local bias is supported.
		The EVPN-MPLS and MC-LAG forwarding logic contradicts each other and can prevent BUM traffic from being forwarded to the ES.	Support local bias, thereby ignoring the DF and non-DF status of the port for locally switched traffic.

South bound (PE1 forwards BUM packet to PE2 and PE3).	Traffic received from PE1 follows the EVPN DF and non-DF forwarding rules for a mulithomed ES.	None.	Not applicable.

MAC Learning

EVPN and MC-LAG use the same method for learning MAC addresses—namely, a PE device learns MAC addresses from its local interfaces and synchronizes the addresses to its peers. However, given that both EVPN and MC-LAG are synchronizing the addresses, an issue arises.

Table 15 on page 385 describes the issue and how the EVPN-MPLS interworking implementation prevents the issue. The use cases shown in Figure 37 on page 382 and Figure 38 on page 383 illustrate the issue. In both use cases, PE1, PE2, and PE3 are EVPN peers, and PE2 and PE3 are MC-LAG peers.

Table 15: MAC Learning: EVPN and MC-LAG Synchronization Issue and Implementation Details

MAC Synchronization Use Case	EVPN Interworking with Junos Fusion Enterprise and MC-LAG Logic	Issue	Juniper Networks Implementation Approach
MAC addresses learned locally on single- or dual-homed interfaces on PE2 and PE3.	• Between the EVPN peers, MAC addresses are synchronized using the EVPN BGP control plane. • Between the MC-LAG peers, MAC addresses are synchronized using the MC-LAG ICCP control plane.	PE2 and PE3 function as both EVPN peers and MC-LAG peers, which result in these devices having multiple MAC synchronization paths.	• For PE1: use MAC addresses synchronized by EVPN BGP control plane. • For PE2 and PE3: use MAC addresses synchronized by MC-LAG ICCP control plane.

Table 15: MAC Learning: EVPN and MC-LAG Synchronization Issue and Implementation Details (continued)

MAC Synchronization Use Case	EVPN Interworking with Junos Fusion Enterprise and MC-LAG Logic	Issue	Juniper Networks Implementation Approach
MAC addresses learned locally on single- or dual-homed interfaces on PE1.	Between the EVPN peers, MAC addresses are synchronized using the EVPN BGP control plane.	None.	Not applicable.

Handling Down Link Between Cascade and Uplink Ports in Junos Fusion Enterprise

NOTE: **This section applies only to EVPN-MPLS interworking with a Junos Fusion Enterprise.**

In the Junos Fusion Enterprise shown in Figure 37 on page 382, assume that aggregation device PE2 receives a BUM packet from PE1 and that the link between the cascade port on PE2 and the corresponding uplink port on satellite device SD1 is down. Regardless of whether the BUM packet is handled by MC-LAG or EVPN multihoming active-active, the result is the same—the packet is forwarded via the ICL interface to PE3, which forwards it to dual-homed SD1.

To further illustrate how EVPN with multihoming active-active handles this situation with dual-homed SD1, assume that the DF interface resides on PE2 and is associated with the down link and that the non-DF interface resides on PE3. Typically, per EVPN with multihoming active-active forwarding logic, the non-DF interface drops the packet. However, because of the down link associated with the DF interface, PE2 forwards the BUM packet via the ICL to PE3, and the non-DF interface on PE3 forwards the packet to SD1.

Layer 3 Gateway Support

The EVPN-MPLS interworking feature supports the following Layer 3 gateway functionality for extended bridge domains and VLANs:

- Integrated routing and bridging (IRB) interfaces to forward traffic between the extended bridge domains or VLANs.
- Default Layer 3 gateways to forward traffic from a physical (bare-metal) server in an extended bridge domain or VLAN to a physical server or virtual machine in another extended bridge domain or VLAN.

Example: EVPN-MPLS Interworking With an MC-LAG Topology

This example shows how to use Ethernet VPN (EVPN) to extend a multichassis link aggregation (MC-LAG) network over an MPLS network to a data center network or geographically distributed campus network.

EVPN-MPLS interworking is supported with an MC-LAG topology in which two MX Series routers, two EX9200 switches, or a mix of the two Juniper Networks devices function as MC-LAG peers, which use the Inter-Chassis Control Protocol (ICCP) and an interchassis

link (ICL) to connect and maintain the topology. The MC-LAG peers are connected to a provider edge (PE) device in an MPLS network. The PE device can be either an MX Series router or an EX9200 switch.

This example shows how to configure the MC-LAG peers and PE device in the MPLS network to interwork with each other.

- Starting with Junos OS Release 19.1R1, the **no-arp-suppression** configuration statement is no longer supported on any device.

Requirements

This example uses the following hardware and software components:

- Three EX9200 switches:
  - PE1 and PE2, which both function as MC-LAG peers in the MC-LAG topology and EVPN BGP peers in the EVPN-MPLS overlay network.
  - PE3, which functions as an EVPN BGP peer in the EVPN-MPLS overlay network.
- The EX9200 switches are running Junos OS Release 17.4R1 or later software.

NOTE: **Although the MC-LAG topology includes two customer edge (CE) devices, this example focuses on the configuration of the PE1, PE2, and PE3.**

Overview and Topology

Figure 39 on page 388 shows an MC-LAG topology with provider edge devices PE1 and PE2 that are configured as MC-LAG peers. The MC-LAG peers exchange control information over an ICCP link and data traffic over an ICL. In this example, the ICL is an aggregated Ethernet interface that is comprised of two interfaces.

Figure 39: EVPN-MPLS Interworking With an MC-LAG Topology

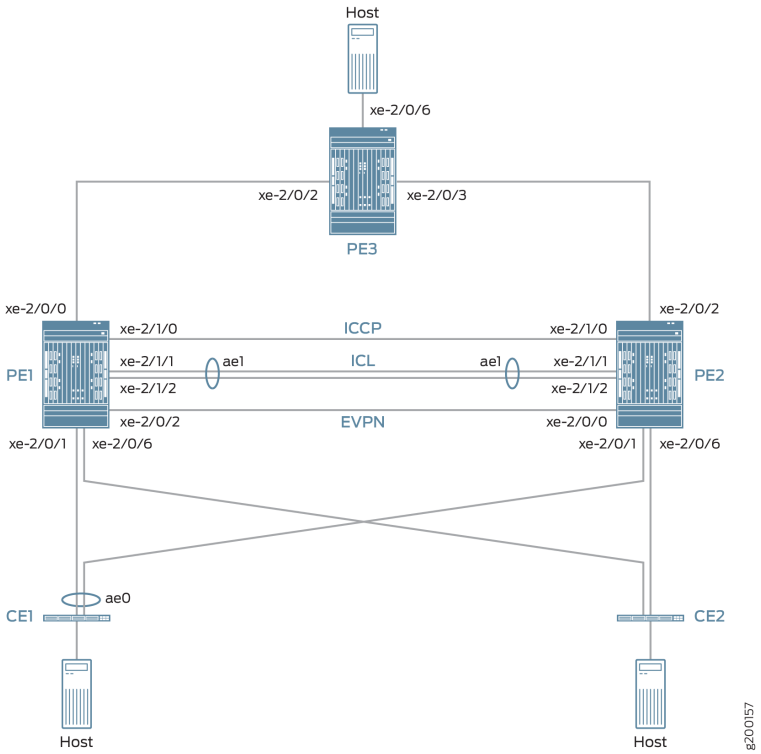

The topology in Figure 39 on page 388 also includes CE devices CE1 and CE2, which are both multihomed to each PE device. The links between CE1 and the two PE devices are bundled as an aggregated Ethernet interface on which MC-LAG in active-active mode is configured.

The topology in Figure 39 on page 388 also includes PE3 at the edge of an MPLS network. PE3 functions as the gateway between the MC-LAG network and either a data center or a geographically distributed campus network. PE1, PE2, and PE3 run EVPN, which enables hosts in the MC-LAG network to communicate with hosts in the data center or other campus network by way of an intervening MPLS network.

From the perspective of the EVPN-MPLS interworking feature, PE3 functions solely as an EVPN BGP peer, and PE1 and PE2 in the MC-LAG topology have dual roles:

- MC-LAG peers in the MC-LAG network.
- EVPN BGP peers in the EVPN-MPLS network.

Because of the dual roles, PE1 and PE2 are configured with MC-LAG, EVPN, BGP, and MPLS attributes.

Table 16 on page 389 outlines key MC-LAG and EVPN (BGP and MPLS) attributes configured on PE1, PE2, and PE3.

Table 16: Key MC-LAG and EVPN (BGP and MPLS) Attributes Configured on PE1, PE2, and PE3

Key Attributes	PE1	PE2	PE3
MC-LAG Attributes			
Interfaces	ICL: aggregated Ethernet interface ae1, which is comprised of xe-2/1/1 and xe-2/1/2 ICCP: xe-2/1/0	ICL: aggregated Ethernet interface ae1, which is comprised of xe-2/1/1 and xe-2/1/2 ICCP: xe-2/1/0	Not applicable
EVPN-MPLS			
Interfaces	Connection to PE3: xe-2/0/0 Connection to PE2: xe-2/0/2	Connection to PE3: xe-2/0/2 Connection to PE1: xe-2/0/0	Connection to PE1: xe-2/0/2 Connection to PE2: xe-2/0/3
IP addresses	BGP peer address: 198.51.100.1	BGP peer address: 198.51.100.2	BGP peer address: 198.51.100.3
Autonomous system	65000	65000	65000
Virtual switch routing instances	evpn1, evpn2, evpn3	evpn1, evpn2, evpn3	evpn1, evpn2, evpn3

Note the following about the EVPN-MPLS interworking feature and its configuration:

- You must configure Ethernet segment identifiers (ESIs) on the dual-homed interfaces in the MC-LAG topology. The ESIs enable EVPN to identify the dual-homed interfaces.
- The only type of routing instance that is supported is the virtual switch instance (**set routing-instances *name* instance-type virtual-switch**).
- On the MC-LAG peers, you must include the **bgp-peer** configuration statement in the **[edit routing-instances *name* protocols evpn mclag]** hierarchy level. This configuration statement enables the interworking of EVPN-MPLS with MC-LAG on the MC-LAG peers.
- Address Resolution Protocol (ARP) suppression is not supported.

PE1 and PE2 Configuration

To configure PE1 and PE2, perform these tasks:

- PE1: Configuring MC-LAG on page 394
- PE1: Configuring EVPN-MPLS on page 396
- PE2: Configuring MC-LAG on page 398
- PE2: Configuring EVPN-MPLS on page 400

CLI Quick Configuration

PE1: MC-LAG Configuration

```
set chassis aggregated-devices ethernet device-count 3
set interfaces xe-2/0/1 gigether-options 802.3ad ae0
set interfaces ae0 flexible-vlan-tagging
set interfaces ae0 encapsulation flexible-ethernet-services
set interfaces ae0 aggregated-ether-options lacp active
set interfaces ae0 aggregated-ether-options lacp periodic fast
set interfaces ae0 aggregated-ether-options lacp system-id 00:00:11:11:11:11
set interfaces ae0 aggregated-ether-options lacp admin-key 1
set interfaces ae0 aggregated-ether-options mc-ae mc-ae-id 1
set interfaces ae0 aggregated-ether-options mc-ae redundancy-group 2
set interfaces ae0 aggregated-ether-options mc-ae chassis-id 0
set interfaces ae0 aggregated-ether-options mc-ae mode active-active
set interfaces ae0 aggregated-ether-options mc-ae status-control active
set interfaces ae0 unit 1 esi 00:11:22:33:44:55:66:77:88:99
set interfaces ae0 unit 1 esi all-active
set interfaces ae0 unit 1 family ethernet-switching interface-mode trunk
set interfaces ae0 unit 1 family ethernet-switching vlan members 1
set interfaces ae0 unit 2 esi 00:11:11:11:11:11:11:11:11:11
set interfaces ae0 unit 2 esi all-active
set interfaces ae0 unit 2 family ethernet-switching interface-mode trunk
set interfaces ae0 unit 2 family ethernet-switching vlan members 2
set interfaces ae0 unit 3 esi 00:11:22:22:22:22:22:22:22:22
set interfaces ae0 unit 3 esi all-active
set interfaces ae0 unit 3 family ethernet-switching interface-mode trunk
set interfaces ae0 unit 3 family ethernet-switching vlan members 3
set interfaces xe-2/0/6 enable
set interfaces xe-2/0/6 flexible-vlan-tagging
set interfaces xe-2/0/6 encapsulation flexible-ethernet-services
set interfaces xe-2/0/6 unit 1 family ethernet-switching interface-mode trunk
set interfaces xe-2/0/6 unit 1 family ethernet-switching vlan members 1
set interfaces xe-2/0/6 unit 2 family ethernet-switching interface-mode trunk
set interfaces xe-2/0/6 unit 2 family ethernet-switching vlan members 2
set interfaces xe-2/0/6 unit 3 family ethernet-switching interface-mode trunk
set interfaces xe-2/0/6 unit 3 family ethernet-switching vlan members 3
set interfaces xe-2/1/0 unit 0 family inet address 203.0.113.1/24
set interfaces xe-2/1/1 gigether-options 802.3ad ae1
set interfaces xe-2/1/2 gigether-options 802.3ad ae1
set interfaces ae1 flexible-vlan-tagging
set interfaces ae1 encapsulation flexible-ethernet-services
set interfaces ae1 aggregated-ether-options lacp active
set interfaces ae1 unit 1 family ethernet-switching interface-mode trunk
set interfaces ae1 unit 1 family ethernet-switching vlan members 1
set interfaces ae1 unit 2 family ethernet-switching interface-mode trunk
set interfaces ae1 unit 2 family ethernet-switching vlan members 2
set interfaces ae1 unit 3 family ethernet-switching interface-mode trunk
set interfaces ae1 unit 3 family ethernet-switching vlan members 3
set multi-chassis multi-chassis-protection 203.0.113.2 interface ae1
set protocols iccp local-ip-addr 203.0.113.1
set protocols iccp peer 203.0.113.2 session-establishment-hold-time 600
set protocols iccp peer 203.0.113.2 redundancy-group-id-list 2
set protocols iccp peer 203.0.113.2 liveness-detection minimum-interval 10000
set protocols iccp peer 203.0.113.2 liveness-detection multiplier 3
```

PE1: EVPN-MPLS Configuration

```
set interfaces lo0 unit 0 family inet address 198.51.100.1/32 primary
set interfaces xe-2/0/0 unit 0 family inet address 192.0.2.2/24
set interfaces xe-2/0/0 unit 0 family mpls
set interfaces xe-2/0/2 unit 0 family inet address 192.0.2.111/24
set interfaces xe-2/0/2 unit 0 family mpls
set interfaces irb unit 1 family inet address 10.2.1.1/24 virtual-gateway-address 10.2.1.254
set interfaces irb unit 2 family inet address 10.2.2.1/24 virtual-gateway-address 10.2.2.254
set interfaces irb unit 3 family inet address 10.2.3.1/24 virtual-gateway-address 10.2.3.254
set routing-options router-id 198.51.100.1
set routing-options autonomous-system 65000
set routing-options forwarding-table export evpn-pplb
set protocols mpls interface xe-2/0/0.0
set protocols mpls interface xe-2/0/2.0
set protocols bgp group evpn type internal
set protocols bgp group evpn local-address 198.51.100.1
set protocols bgp group evpn family evpn signaling
set protocols bgp group evpn local-as 65000
set protocols bgp group evpn neighbor 198.51.100.2
set protocols bgp group evpn neighbor 198.51.100.3
set protocols ospf area 0.0.0.0 interface lo0.0
set protocols ospf area 0.0.0.0 interface xe-2/0/0.0
set protocols ospf area 0.0.0.0 interface xe-2/0/2.0
set protocols ldp interface xe-2/0/0.0
set protocols ldp interface xe-2/0/2.0
set protocols ldp interface lo0.0
set policy-options policy-statement evpn-pplb from protocol evpn
set policy-options policy-statement evpn-pplb then load-balance per-packet
set routing-instances evpn1 instance-type virtual-switch
set routing-instances evpn1 interface xe-2/0/6.1
set routing-instances evpn1 interface ae0.1
set routing-instances evpn1 interface ae1.1
set routing-instances evpn1 route-distinguisher 1:10
set routing-instances evpn1 vrf-target target:1:5
set routing-instances evpn1 protocols evpn extended-vlan-list 1
set routing-instances evpn1 protocols evpn mclag bgp-peer 198.51.100.2
set routing-instances evpn1 switch-options service-id 1
set routing-instances evpn1 vlans v1 vlan-id 1
set routing-instances evpn1 vlans v1 l3-interface irb.1
set routing-instances evpn1 vlans v1 no-arp-suppression
set routing-instances evpn2 instance-type virtual-switch
set routing-instances evpn2 interface xe-2/0/6.2
set routing-instances evpn2 interface ae0.2
set routing-instances evpn2 interface ae1.2
set routing-instances evpn2 route-distinguisher 1:20
set routing-instances evpn2 vrf-target target:1:6
set routing-instances evpn2 protocols evpn extended-vlan-list 2
set routing-instances evpn2 protocols evpn mclag bgp-peer 198.51.100.2
set routing-instances evpn2 switch-options service-id 2
set routing-instances evpn2 vlans v1 vlan-id 2
set routing-instances evpn2 vlans v1 l3-interface irb.2
set routing-instances evpn2 vlans v1 no-arp-suppression
set routing-instances evpn3 instance-type virtual-switch
set routing-instances evpn3 interface xe-2/0/6.3
```

```
set routing-instances evpn3 interface ae0.3
set routing-instances evpn3 interface ae1.3
set routing-instances evpn3 route-distinguisher 1:30
set routing-instances evpn3 vrf-target target:1:7
set routing-instances evpn3 protocols evpn extended-vlan-list 3
set routing-instances evpn3 protocols evpn mclag bgp-peer 198.51.100.2
set routing-instances evpn3 switch-options service-id 3
set routing-instances evpn3 vlans v1 vlan-id 3
set routing-instances evpn3 vlans v1 l3-interface irb.3
set routing-instances evpn3 vlans v1 no-arp-suppression
```

PE2: MC-LAG Configuration

```
set chassis aggregated-devices ethernet device-count 3
set interfaces xe-2/0/1 gigether-options 802.3ad ae0
set interfaces xe-2/0/6 enable
set interfaces xe-2/0/6 flexible-vlan-tagging
set interfaces xe-2/0/6 encapsulation flexible-ethernet-services
set interfaces xe-2/0/6 unit 1 family ethernet-switching interface-mode trunk
set interfaces xe-2/0/6 unit 1 family ethernet-switching vlan members 1
set interfaces xe-2/0/6 unit 2 family ethernet-switching interface-mode trunk
set interfaces xe-2/0/6 unit 2 family ethernet-switching vlan members 2
set interfaces xe-2/0/6 unit 3 family ethernet-switching interface-mode trunk
set interfaces xe-2/0/6 unit 3 family ethernet-switching vlan members 3
set interfaces xe-2/1/0 unit 0 family inet address 203.0.113.2/24
set interfaces xe-2/1/1 gigether-options 802.3ad ae1
set interfaces xe-2/1/2 gigether-options 802.3ad ae1
set interfaces ae0 flexible-vlan-tagging
set interfaces ae0 encapsulation flexible-ethernet-services
set interfaces ae0 aggregated-ether-options lacp active
set interfaces ae0 aggregated-ether-options lacp periodic fast
set interfaces ae0 aggregated-ether-options lacp system-id 00:00:11:11:11:11
set interfaces ae0 aggregated-ether-options lacp admin-key 1
set interfaces ae0 aggregated-ether-options mc-ae mc-ae-id 1
set interfaces ae0 aggregated-ether-options mc-ae redundancy-group 2
set interfaces ae0 aggregated-ether-options mc-ae chassis-id 1
set interfaces ae0 aggregated-ether-options mc-ae mode active-active
set interfaces ae0 aggregated-ether-options mc-ae status-control standby
set interfaces ae0 unit 1 esi 00:11:22:33:44:55:66:77:88:99
set interfaces ae0 unit 1 esi all-active
set interfaces ae0 unit 1 family ethernet-switching interface-mode trunk
set interfaces ae0 unit 1 family ethernet-switching vlan members 1
set interfaces ae0 unit 2 esi 00:11:11:11:11:11:11:11:11:11
set interfaces ae0 unit 2 esi all-active
set interfaces ae0 unit 2 family ethernet-switching interface-mode trunk
set interfaces ae0 unit 2 family ethernet-switching vlan members 2
set interfaces ae0 unit 3 esi 00:11:22:22:22:22:22:22:22:22
set interfaces ae0 unit 3 esi all-active
set interfaces ae0 unit 3 family ethernet-switching interface-mode trunk
set interfaces ae0 unit 3 family ethernet-switching vlan members 3
set interfaces ae1 flexible-vlan-tagging
set interfaces ae1 encapsulation flexible-ethernet-services
set interfaces ae1 aggregated-ether-options lacp active
set interfaces ae1 unit 1 family ethernet-switching interface-mode trunk
```

```
set interfaces ae1 unit 1 family ethernet-switching vlan members 1
set interfaces ae1 unit 2 family ethernet-switching interface-mode trunk
set interfaces ae1 unit 2 family ethernet-switching vlan members 2
set interfaces ae1 unit 3 family ethernet-switching interface-mode trunk
set interfaces ae1 unit 3 family ethernet-switching vlan members 3
set multi-chassis multi-chassis-protection 203.0.113.1 interface ae1
set protocols iccp local-ip-addr 203.0.113.2
set protocols iccp peer 203.0.113.1 session-establishment-hold-time 600
set protocols iccp peer 203.0.113.1 redundancy-group-id-list 2
set protocols iccp peer 203.0.113.1 liveness-detection minimum-interval 10000
set protocols iccp peer 203.0.113.1 liveness-detection multiplier 3
```

PE2: EVPN-MPLS Configuration

```
set interfaces xe-2/0/0 unit 0 family inet address 192.0.2.222/24
set interfaces xe-2/0/0 unit 0 family mpls
set interfaces xe-2/0/2 unit 0 family inet address 192.0.2.22/24
set interfaces xe-2/0/2 unit 0 family mpls
set interfaces lo0 unit 0 family inet address 198.51.100.2/32 primary
set interfaces irb unit 1 family inet address 10.2.1.2/24 virtual-gateway-address 10.2.1.254
set interfaces irb unit 2 family inet address 10.2.2.2/24 virtual-gateway-address 10.2.2.254
set interfaces irb unit 3 family inet address 10.2.3.2/24 virtual-gateway-address 10.2.3.254
set routing-options router-id 198.51.100.2
set routing-options autonomous-system 65000
set routing-options forwarding-table export evpn-pplb
set protocols mpls interface xe-2/0/2.0
set protocols mpls interface xe-2/0/0.0
set protocols bgp group evpn type internal
set protocols bgp group evpn local-address 198.51.100.2
set protocols bgp group evpn family evpn signaling
set protocols bgp group evpn local-as 65000
set protocols bgp group evpn neighbor 198.51.100.1
set protocols bgp group evpn neighbor 198.51.100.3
set protocols ospf area 0.0.0.0 interface lo0.0
set protocols ospf area 0.0.0.0 interface xe-2/0/0.0
set protocols ospf area 0.0.0.0 interface xe-2/0/2.0
set protocols ldp interface xe-2/0/0.0
set protocols ldp interface xe-2/0/2.0
set protocols ldp interface lo0.0
set policy-options policy-statement evpn-pplb from protocol evpn
set policy-options policy-statement evpn-pplb then load-balance per-packet
set routing-instances evpn1 instance-type virtual-switch
set routing-instances evpn1 interface xe-2/0/6.1
set routing-instances evpn1 interface ae0.1
set routing-instances evpn1 interface ae1.1
set routing-instances evpn1 route-distinguisher 1:11
set routing-instances evpn1 vrf-target target:1:5
set routing-instances evpn1 protocols evpn extended-vlan-list 1
set routing-instances evpn1 protocols evpn mclag bgp-peer 198.51.100.1
set routing-instances evpn1 switch-options service-id 1
set routing-instances evpn1 vlans v1 vlan-id 1
set routing-instances evpn1 vlans v1 l3-interface irb.1
set routing-instances evpn1 vlans v1 no-arp-suppression
set routing-instances evpn2 instance-type virtual-switch
```

```
set routing-instances evpn2 interface xe-2/0/6.2
set routing-instances evpn2 interface ae0.2
set routing-instances evpn2 interface ae1.2
set routing-instances evpn2 route-distinguisher 1:21
set routing-instances evpn2 vrf-target target:1:6
set routing-instances evpn2 protocols evpn extended-vlan-list 2
set routing-instances evpn2 protocols evpn mclag bgp-peer 198.51.100.1
set routing-instances evpn2 switch-options service-id 2
set routing-instances evpn2 vlans v1 vlan-id 2
set routing-instances evpn2 vlans v1 l3-interface irb.2
set routing-instances evpn2 vlans v1 no-arp-suppression
set routing-instances evpn3 instance-type virtual-switch
set routing-instances evpn3 interface xe-2/0/6.3
set routing-instances evpn3 interface ae0.3
set routing-instances evpn3 interface ae1.3
set routing-instances evpn3 route-distinguisher 1:31
set routing-instances evpn3 vrf-target target:1:7
set routing-instances evpn3 protocols evpn extended-vlan-list 3
set routing-instances evpn3 protocols evpn mclag bgp-peer 198.51.100.1
set routing-instances evpn3 switch-options service-id 3
set routing-instances evpn3 vlans v1 vlan-id 3
set routing-instances evpn3 vlans v1 l3-interface irb.3
set routing-instances evpn3 vlans v1 no-arp-suppression
```

PE1: Configuring MC-LAG

Step-by-Step Procedure

1. Set the number of aggregated Ethernet interfaces on PE1.

```
[edit]
user@switch# set chassis aggregated-devices ethernet device-count 3
```

2. Configure aggregated Ethernet interface ae0 on interface xe-2/0/1, and configure LACP and MC-LAG on ae0. Divide aggregated Ethernet interface ae0 into three logical interfaces (ae0.1, ae0.2, and ae0.3). For each logical interface, specify an ESI, place the logical interface is in MC-LAG active-active mode, and map the logical interface to a VLAN.

```
[edit]
user@switch# set interfaces xe-2/0/1 gigether-options 802.3ad ae0
user@switch# set interfaces ae0 flexible-vlan-tagging
user@switch# set interfaces ae0 encapsulation flexible-ethernet-services
user@switch# set interfaces ae0 aggregated-ether-options lacp active
user@switch# set interfaces ae0 aggregated-ether-options lacp periodic fast
user@switch# set interfaces ae0 aggregated-ether-options lacp system-id
  00:00:11:11:11:11
user@switch# set interfaces ae0 aggregated-ether-options lacp admin-key 1
user@switch# set interfaces ae0 aggregated-ether-options mc-ae mc-ae-id 1
user@switch# set interfaces ae0 aggregated-ether-options mc-ae redundancy-group
  2
user@switch# set interfaces ae0 aggregated-ether-options mc-ae chassis-id 0
user@switch# set interfaces ae0 aggregated-ether-options mc-ae mode
  active-active
```

```
user@switch# set interfaces ae0 aggregated-ether-options mc-ae status-control
  active
user@switch# set interfaces ae0 unit 1 esi 00:11:22:33:44:55:66:77:88:99
user@switch# set interfaces ae0 unit 1 esi all-active
user@switch# set interfaces ae0 unit 1 family ethernet-switching interface-mode
  trunk
user@switch# set interfaces ae0 unit 1 family ethernet-switching vlan members 1
user@switch# set interfaces ae0 unit 2 esi 00:11:11:11:11:11:11:11:11:11
user@switch# set interfaces ae0 unit 2 esi all-active
user@switch# set interfaces ae0 unit 2 family ethernet-switching interface-mode
  trunk
user@switch# set interfaces ae0 unit 2 family ethernet-switching vlan members 2
user@switch# set interfaces ae0 unit 3 esi 00:11:22:22:22:22:22:22:22:22
user@switch# set interfaces ae0 unit 3 esi all-active
user@switch# set interfaces ae0 unit 3 family ethernet-switching interface-mode
  trunk
user@switch# set interfaces ae0 unit 3 family ethernet-switching vlan members 3
```

3. Configure physical interface xe-2/0/6, and divide it into three logical interfaces (xe-2/0/6.1, xe-2/0/6.2, and xe-2/0/6.3). Map each logical interface to a VLAN.

```
[edit]
user@switch# set interfaces xe-2/0/6 enable
user@switch# set interfaces xe-2/0/6 flexible-vlan-tagging
user@switch# set interfaces xe-2/0/6 encapsulation flexible-ethernet-services
user@switch# set interfaces xe-2/0/6 unit 1 family ethernet-switching
  interface-mode trunk
user@switch# set interfaces xe-2/0/6 unit 1 family ethernet-switching vlan members
  1
user@switch# set interfaces xe-2/0/6 unit 2 family ethernet-switching
  interface-mode trunk
user@switch# set interfaces xe-2/0/6 unit 2 family ethernet-switching vlan members
  2
user@switch# set interfaces xe-2/0/6 unit 3 family ethernet-switching
  interface-mode trunk
user@switch# set interfaces xe-2/0/6 unit 3 family ethernet-switching vlan members
  3
```

4. Configure physical interface xe-2/1/0 as a Layer 3 interface, on which you configure ICCP. Specify the interface with the IP address of 203.0.113.2 on PE2 as the ICCP peer to PE1.

```
[edit]
user@switch# set interfaces xe-2/1/0 unit 0 family inet address 203.0.113.1/24
user@switch# set protocols iccp local-ip-addr 203.0.113.1
user@switch# set protocols iccp peer 203.0.113.2 session-establishment-hold-time
  600
user@switch# set protocols iccp peer 203.0.113.2 redundancy-group-id-list 2
user@switch# set protocols iccp peer 203.0.113.2 liveness-detection
  minimum-interval 10000
user@switch# set protocols iccp peer 203.0.113.2 liveness-detection multiplier 3
```

5. Configure aggregated Ethernet interface ae1 on interfaces xe-2/1/1 and xe-2/1/2, and configure LACP on ae1. Divide aggregated Ethernet interface ae1 into three logical interfaces (ae1.1, ae1.2, and ae1.3), and map each logical interface to a VLAN. Specify ae1 as the multichassis protection link between PE1 and PE2.

```
[edit]
user@switch# set interfaces xe-2/1/1 gigether-options 802.3ad ae1
user@switch# set interfaces xe-2/1/2 gigether-options 802.3ad ae1
user@switch# set interfaces ae1 flexible-vlan-tagging
user@switch# set interfaces ae1 encapsulation flexible-ethernet-services
user@switch# set interfaces ae1 aggregated-ether-options lacp active
user@switch# set interfaces ae1 unit 1 family ethernet-switching interface-mode
  trunk
user@switch# set interfaces ae1 unit 1 family ethernet-switching vlan members 1
user@switch# set interfaces ae1 unit 2 family ethernet-switching interface-mode
  trunk
user@switch# set interfaces ae1 unit 2 family ethernet-switching vlan members 2
user@switch# set interfaces ae1 unit 3 family ethernet-switching interface-mode
  trunk
user@switch# set interfaces ae1 unit 3 family ethernet-switching vlan members 3
user@switch# set multi-chassis multi-chassis-protection 203.0.113.2 interface ae1
```

PE1: Configuring EVPN-MPLS

Step-by-Step Procedure

1. Configure the loopback interface, and the interfaces connected to the other PE devices.

```
[edit]
user@switch# set interfaces lo0 unit 0 family inet address 198.51.100.1/32 primary
user@switch# set interfaces xe-2/0/0 unit 0 family inet address 192.0.2.2/24
user@switch# set interfaces xe-2/0/0 unit 0 family mpls
user@switch# set interfaces xe-2/0/2 unit 0 family inet address 192.0.2.111/24
user@switch# set interfaces xe-2/0/2 unit 0 family mpls
```

2. Configure IRB interfaces irb.1, irb.2, and irb.3.

```
[edit]
user@switch# set interfaces irb unit 1 family inet address 10.2.1.1/24
  virtual-gateway-address 10.2.1.254
user@switch# set interfaces irb unit 2 family inet address 10.2.2.1/24
  virtual-gateway-address 10.2.2.254
user@switch# set interfaces irb unit 3 family inet address 10.2.3.1/24
  virtual-gateway-address 10.2.3.254
```

3. Assign a router ID and the autonomous system in which PE1, PE2, and PE3 reside.

```
[edit]
user@switch# set routing-options router-id 198.51.100.1
user@switch# set routing-options autonomous-system 65000
```

4. Enable per-packet load-balancing for EVPN routes when EVPN multihoming active-active mode is used.

```
[edit]
user@switch# set routing-options forwarding-table export evpn-pplb
user@switch# set policy-options policy-statement evpn-pplb from protocol evpn
user@switch# set policy-options policy-statement evpn-pplb then load-balance
  per-packet
```

5. Enable MPLS on interfaces xe-2/0/0.0 and xe-2/0/2.0.

```
[edit]
user@switch# set protocols mpls interface xe-2/0/0.0
user@switch# set protocols mpls interface xe-2/0/2.0
```

6. Configure an IBGP overlay that includes PE1, PE2, and PE3.

```
[edit]
user@switch# set protocols bgp group evpn type internal
user@switch# set protocols bgp group evpn local-address 198.51.100.1
user@switch# set protocols bgp group evpn family evpn signaling
user@switch# set protocols bgp group evpn local-as 65000
user@switch# set protocols bgp group evpn neighbor 198.51.100.2
user@switch# set protocols bgp group evpn neighbor 198.51.100.3
```

7. Configure OSPF as the internal routing protocol for EVPN by specifying an area ID and interfaces on which EVPN-MPLS is enabled.

```
[edit]
user@switch# set protocols ospf area 0.0.0.0 interface lo0.0
user@switch# set protocols ospf area 0.0.0.0 interface xe-2/0/0.0
user@switch# set protocols ospf area 0.0.0.0 interface xe-2/0/2.0
```

8. Configure the Label Distribution Protocol (LDP) on the loopback interface and the interfaces on which EVPN-MPLS is enabled.

```
[edit]
user@switch# set protocols ldp interface lo0.0
user@switch# set protocols ldp interface xe-2/0/0.0
user@switch# set protocols ldp interface xe-2/0/2.0
```

9. Configure virtual switch routing instances for VLAN v1, which is assigned VLAN IDs of 1, 2, and 3, and include the interfaces and other entities associated with the VLAN.

```
[edit]
user@switch# set routing-instances evpn1 instance-type virtual-switch
user@switch# set routing-instances evpn1 interface xe-2/0/6.1
user@switch# set routing-instances evpn1 interface ae0.1
```

```
user@switch# set routing-instances evpn1 interface ae1.1
user@switch# set routing-instances evpn1 route-distinguisher 1:10
user@switch# set routing-instances evpn1 vrf-target target:1:5
user@switch# set routing-instances evpn1 protocols evpn extended-vlan-list 1
user@switch# set routing-instances evpn1 protocols evpn mclag bgp-peer
  198.51.100.2
user@switch# set routing-instances evpn1 switch-options service-id 1
user@switch# set routing-instances evpn1 vlans v1 vlan-id 1
user@switch# set routing-instances evpn1 vlans v1 l3-interface irb.1
user@switch# set routing-instances evpn1 vlans v1 no-arp-suppression
user@switch# set routing-instances evpn2 instance-type virtual-switch
user@switch# set routing-instances evpn2 interface xe-2/0/6.2
user@switch# set routing-instances evpn2 interface ae0.2
user@switch# set routing-instances evpn2 interface ae1.2
user@switch# set routing-instances evpn2 route-distinguisher 1:20
user@switch# set routing-instances evpn2 vrf-target target:1:6
user@switch# set routing-instances evpn2 protocols evpn extended-vlan-list 2
user@switch# set routing-instances evpn2 protocols evpn mclag bgp-peer
  198.51.100.2
user@switch# set routing-instances evpn2 switch-options service-id 2
user@switch# set routing-instances evpn2 vlans v1 vlan-id 2
user@switch# set routing-instances evpn2 vlans v1 l3-interface irb.2
user@switch# set routing-instances evpn2 vlans v1 no-arp-suppression
user@switch# set routing-instances evpn3 instance-type virtual-switch
user@switch# set routing-instances evpn3 interface xe-2/0/6.3
user@switch# set routing-instances evpn3 interface ae0.3
user@switch# set routing-instances evpn3 interface ae1.3
user@switch# set routing-instances evpn3 route-distinguisher 1:30
user@switch# set routing-instances evpn3 vrf-target target:1:7
user@switch# set routing-instances evpn3 protocols evpn extended-vlan-list 3
user@switch# set routing-instances evpn3 protocols evpn mclag bgp-peer
  198.51.100.2
user@switch# set routing-instances evpn3 switch-options service-id 3
user@switch# set routing-instances evpn3 vlans v1 vlan-id 3
user@switch# set routing-instances evpn3 vlans v1 l3-interface irb.3
user@switch# set routing-instances evpn3 vlans v1 no-arp-suppression
```

PE2: Configuring MC-LAG

Step-by-Step Procedure

1. Set the number of aggregated Ethernet interfaces on PE2.

```
[edit]
user@switch# set chassis aggregated-devices ethernet device-count 3
```

2. Configure aggregated Ethernet interface ae0 on interface xe-2/0/1, and configure LACP and MC-LAG on ae0. Divide aggregated Ethernet interface ae0 into three logical interfaces (ae0.1, ae0.2, and ae0.3). For each logical interface, specify an ESI, place the logical interface is in MC-LAG active-active mode, and map the logical interface to a VLAN.

```
[edit]
```

```
user@switch# set interfaces xe-2/0/1 gigether-options 802.3ad ae0
user@switch# set interfaces ae0 flexible-vlan-tagging
user@switch# set interfaces ae0 encapsulation flexible-ethernet-services
user@switch# set interfaces ae0 aggregated-ether-options lacp active
user@switch# set interfaces ae0 aggregated-ether-options lacp periodic fast
user@switch# set interfaces ae0 aggregated-ether-options lacp system-id
  00:00:11:11:11:11
user@switch# set interfaces ae0 aggregated-ether-options lacp admin-key 1
user@switch# set interfaces ae0 aggregated-ether-options mc-ae mc-ae-id 1
user@switch# set interfaces ae0 aggregated-ether-options mc-ae redundancy-group
  2
user@switch# set interfaces ae0 aggregated-ether-options mc-ae chassis-id 1
user@switch# set interfaces ae0 aggregated-ether-options mc-ae mode
  active-active
user@switch# set interfaces ae0 aggregated-ether-options mc-ae status-control
  standby
user@switch# set interfaces ae0 unit 1 esi 00:11:22:33:44:55:66:77:88:99
user@switch# set interfaces ae0 unit 1 esi all-active
user@switch# set interfaces ae0 unit 1 family ethernet-switching interface-mode
  trunk
user@switch# set interfaces ae0 unit 1 family ethernet-switching vlan members 1
user@switch# set interfaces ae0 unit 2 esi 00:11:11:11:11:11:11:11:11:11
user@switch# set interfaces ae0 unit 2 esi all-active
user@switch# set interfaces ae0 unit 2 family ethernet-switching interface-mode
  trunk
user@switch# set interfaces ae0 unit 2 family ethernet-switching vlan members 2
user@switch# set interfaces ae0 unit 3 esi 00:11:22:22:22:22:22:22:22:22
user@switch# set interfaces ae0 unit 3 esi all-active
user@switch# set interfaces ae0 unit 3 family ethernet-switching interface-mode
  trunk
user@switch# set interfaces ae0 unit 3 family ethernet-switching vlan members 3
```

3. Configure physical interface xe-2/0/6, and divide it into three logical interfaces (xe-2/0/6.1, xe-2/0/6.2, and xe-2/0/6.3). Map each logical interface to a VLAN.

```
[edit]
set interfaces xe-2/0/6 enable
set interfaces xe-2/0/6 flexible-vlan-tagging
set interfaces xe-2/0/6 encapsulation flexible-ethernet-services
set interfaces xe-2/0/6 unit 1 family ethernet-switching interface-mode trunk
set interfaces xe-2/0/6 unit 1 family ethernet-switching vlan members 1
set interfaces xe-2/0/6 unit 2 family ethernet-switching interface-mode trunk
set interfaces xe-2/0/6 unit 2 family ethernet-switching vlan members 2
set interfaces xe-2/0/6 unit 3 family ethernet-switching interface-mode trunk
set interfaces xe-2/0/6 unit 3 family ethernet-switching vlan members 3
```

4. Configure physical interface xe-2/1/0 as a Layer 3 interface, on which you configure ICCP. Specify the interface with the IP address of 203.0.113.1 on PE1 as the ICCP peer to PE2.

```
[edit]
set interfaces xe-2/1/0 unit 0 family inet address 203.0.113.2/24
```

```
set protocols iccp local-ip-addr 203.0.113.2
set protocols iccp peer 203.0.113.1 session-establishment-hold-time 600
set protocols iccp peer 203.0.113.1 redundancy-group-id-list 2
set protocols iccp peer 203.0.113.1 liveness-detection minimum-interval 10000
set protocols iccp peer 203.0.113.1 liveness-detection multiplier 3
```

5. Configure aggregated Ethernet interface ae1 on interfaces xe-2/1/1 and xe-2/1/2, and configure LACP on ae1. Divide aggregated Ethernet interface ae1 into three logical interfaces (ae1.1, ae1.2, and ae1.3), and map each logical interface to a VLAN. Specify ae1 as the multichassis protection link between PE1 and PE2.

```
[edit]
set interfaces xe-2/1/1 gigether-options 802.3ad ae1
set interfaces xe-2/1/2 gigether-options 802.3ad ae1
set interfaces ae1 flexible-vlan-tagging
set interfaces ae1 encapsulation flexible-ethernet-services
set interfaces ae1 aggregated-ether-options lacp active
set interfaces ae1 unit 1 family ethernet-switching interface-mode trunk
set interfaces ae1 unit 1 family ethernet-switching vlan members 1
set interfaces ae1 unit 2 family ethernet-switching interface-mode trunk
set interfaces ae1 unit 2 family ethernet-switching vlan members 2
set interfaces ae1 unit 3 family ethernet-switching interface-mode trunk
set interfaces ae1 unit 3 family ethernet-switching vlan members 3
set multi-chassis multi-chassis-protection 203.0.113.1 interface ae1
```

PE2: Configuring EVPN-MPLS

Step-by-Step Procedure

1. Configure the loopback interface, and the interfaces connected to the other PE devices.

```
[edit]
user@switch# set interfaces lo0 unit 0 family inet address 198.51.100.2/32 primary
user@switch# set interfaces xe-2/0/0 unit 0 family inet address 192.0.2.222/24
user@switch# set interfaces xe-2/0/0 unit 0 family mpls
user@switch# set interfaces xe-2/0/2 unit 0 family inet address 192.0.2.22/24
user@switch# set interfaces xe-2/0/2 unit 0 family mpls
```

2. Configure IRB interfaces irb.1, irb.2, and irb.3.

```
[edit]
user@switch# set interfaces irb unit 1 family inet address 10.2.1.2/24
  virtual-gateway-address 10.2.1.254
user@switch# set interfaces irb unit 2 family inet address 10.2.2.2/24
  virtual-gateway-address 10.2.2.254
user@switch# set interfaces irb unit 3 family inet address 10.2.3.2/24
  virtual-gateway-address 10.2.3.254
```

3. Assign a router ID and the autonomous system in which PE1, PE2, and PE3 reside.

```
[edit]
user@switch# set routing-options router-id 198.51.100.2
user@switch# set routing-options autonomous-system 65000
```

4. Enable per-packet load-balancing for EVPN routes when EVPN multihoming active-active mode is used.

```
[edit]
user@switch# set routing-options forwarding-table export evpn-pplb
user@switch# set policy-options policy-statement evpn-pplb from protocol evpn
user@switch# set policy-options policy-statement evpn-pplb then load-balance
  per-packet
```

5. Enable MPLS on interfaces xe-2/0/0.0 and xe-2/0/2.0.

```
[edit]
user@switch# set protocols mpls interface xe-2/0/0.0
user@switch# set protocols mpls interface xe-2/0/2.0
```

6. Configure an IBGP overlay that includes PE1, PE2, and PE3.

```
[edit]
user@switch# set protocols bgp group evpn type internal
user@switch# set protocols bgp group evpn local-address 198.51.100.2
user@switch# set protocols bgp group evpn family evpn signaling
user@switch# set protocols bgp group evpn local-as 65000
user@switch# set protocols bgp group evpn neighbor 198.51.100.1
user@switch# set protocols bgp group evpn neighbor 198.51.100.3
```

7. Configure OSPF as the internal routing protocol for EVPN by specifying an area ID and interfaces on which EVPN-MPLS is enabled.

```
[edit]
user@switch# set protocols ospf area 0.0.0.0 interface lo0.0
user@switch# set protocols ospf area 0.0.0.0 interface xe-2/0/0.0
user@switch# set protocols ospf area 0.0.0.0 interface xe-2/0/2.0
```

8. Configure the Label Distribution Protocol (LDP) on the loopback interface and the interfaces on which EVPN-MPLS is enabled.

```
[edit]
user@switch# set protocols ldp interface lo0.0
user@switch# set protocols ldp interface xe-2/0/0.0
user@switch# set protocols ldp interface xe-2/0/2.0
```

9. Configure virtual switch routing instances for VLAN v1, which is assigned VLAN IDs of 1, 2, and 3, and include the interfaces and other entities associated with the VLAN.

```
[edit]
user@switch# set routing-instances evpn1 instance-type virtual-switch
user@switch# set routing-instances evpn1 interface xe-2/0/6.1
user@switch# set routing-instances evpn1 interface ae0.1
user@switch# set routing-instances evpn1 interface ae1.1
user@switch# set routing-instances evpn1 route-distinguisher 1:11
user@switch# set routing-instances evpn1 vrf-target target:1:5
user@switch# set routing-instances evpn1 protocols evpn extended-vlan-list 1
user@switch# set routing-instances evpn1 protocols evpn mclag bgp-peer
  198.51.100.1
user@switch# set routing-instances evpn1 switch-options service-id 1
user@switch# set routing-instances evpn1 vlans v1 vlan-id 1
user@switch# set routing-instances evpn1 vlans v1 l3-interface irb.1
user@switch# set routing-instances evpn1 vlans v1 no-arp-suppression
user@switch# set routing-instances evpn2 instance-type virtual-switch
user@switch# set routing-instances evpn2 interface xe-2/0/6.2
user@switch# set routing-instances evpn2 interface ae0.2
user@switch# set routing-instances evpn2 interface ae1.2
user@switch# set routing-instances evpn2 route-distinguisher 1:21
user@switch# set routing-instances evpn2 vrf-target target:1:6
user@switch# set routing-instances evpn2 protocols evpn extended-vlan-list 2
user@switch# set routing-instances evpn2 protocols evpn mclag bgp-peer
  198.51.100.1
user@switch# set routing-instances evpn2 switch-options service-id 2
user@switch# set routing-instances evpn2 vlans v1 vlan-id 2
user@switch# set routing-instances evpn2 vlans v1 l3-interface irb.2
user@switch# set routing-instances evpn2 vlans v1 no-arp-suppression
user@switch# set routing-instances evpn3 instance-type virtual-switch
user@switch# set routing-instances evpn3 interface xe-2/0/6.3
user@switch# set routing-instances evpn3 interface ae0.3
user@switch# set routing-instances evpn3 interface ae1.3
user@switch# set routing-instances evpn3 route-distinguisher 1:31
user@switch# set routing-instances evpn3 vrf-target target:1:7
user@switch# set routing-instances evpn3 protocols evpn extended-vlan-list 3
user@switch# set routing-instances evpn3 protocols evpn mclag bgp-peer
  198.51.100.1
user@switch# set routing-instances evpn3 switch-options service-id 3
user@switch# set routing-instances evpn3 vlans v1 vlan-id 3
user@switch# set routing-instances evpn3 vlans v1 l3-interface irb.3
user@switch# set routing-instances evpn3 vlans v1 no-arp-suppression
```

PE3 Configuration

CLI Quick Configuration

PE3: EVPN-MPLS Configuration

```
set interfaces lo0 unit 0 family inet address 198.51.100.3/32 primary
set interfaces xe-2/0/2 unit 0 family inet address 192.0.2.1/24
set interfaces xe-2/0/2 unit 0 family mpls
set interfaces xe-2/0/3 unit 0 family inet address 192.0.2.11/24
```

```
set interfaces xe-2/0/3 unit 0 family mpls
set interfaces xe-2/0/6 enable
set interfaces xe-2/0/6 flexible-vlan-tagging
set interfaces xe-2/0/6 encapsulation flexible-ethernet-services
set interfaces xe-2/0/6 unit 1 family ethernet-switching interface-mode trunk
set interfaces xe-2/0/6 unit 1 family ethernet-switching vlan members 1
set interfaces xe-2/0/6 unit 2 family ethernet-switching interface-mode trunk
set interfaces xe-2/0/6 unit 2 family ethernet-switching vlan members 2
set interfaces xe-2/0/6 unit 3 family ethernet-switching interface-mode trunk
set interfaces xe-2/0/6 unit 3 family ethernet-switching vlan members 3
set interfaces irb unit 1 family inet address 10.2.1.3/24 virtual-gateway-address 10.2.1.254
set interfaces irb unit 2 family inet address 10.2.2.3/24 virtual-gateway-address 10.2.2.254
set interfaces irb unit 3 family inet address 10.2.3.3/24 virtual-gateway-address 10.2.3.254
set routing-options router-id 198.51.100.3
set routing-options autonomous-system 65000
set routing-options forwarding-table export evpn-pplb
set protocols mpls interface xe-2/0/2.0
set protocols mpls interface xe-2/0/3.0
set protocols bgp group evpn type internal
set protocols bgp group evpn local-address 198.51.100.3
set protocols bgp group evpn family evpn signaling
set protocols bgp group evpn local-as 65000
set protocols bgp group evpn neighbor 198.51.100.1
set protocols bgp group evpn neighbor 198.51.100.2
set protocols ospf area 0.0.0.0 interface lo0.0
set protocols ospf area 0.0.0.0 interface xe-2/0/2.0
set protocols ospf area 0.0.0.0 interface xe-2/0/3.0
set protocols ldp interface lo0.0
set protocols ldp interface xe-2/0/2.0
set protocols ldp interface xe-2/0/3.0
set policy-options policy-statement evpn-pplb from protocol evpn
set policy-options policy-statement evpn-pplb then load-balance per-packet
set routing-instances evpn1 instance-type virtual-switch
set routing-instances evpn1 interface xe-2/0/6.1
set routing-instances evpn1 route-distinguisher 1:12
set routing-instances evpn1 vrf-target target:1:5
set routing-instances evpn1 protocols evpn extended-vlan-list 1
set routing-instances evpn1 switch-options service-id 1
set routing-instances evpn1 vlans v1 vlan-id 1
set routing-instances evpn1 vlans v1 l3-interface irb.1
set routing-instances evpn1 vlans v1 no-arp-suppression
set routing-instances evpn2 instance-type virtual-switch
set routing-instances evpn2 interface xe-2/0/6.2
set routing-instances evpn2 route-distinguisher 1:22
set routing-instances evpn2 vrf-target target:1:6
set routing-instances evpn2 protocols evpn extended-vlan-list 2
set routing-instances evpn2 switch-options service-id 2
set routing-instances evpn2 vlans v1 vlan-id 2
set routing-instances evpn2 vlans v1 l3-interface irb.2
set routing-instances evpn2 vlans v1 no-arp-suppression
set routing-instances evpn3 instance-type virtual-switch
set routing-instances evpn3 interface xe-2/0/6.3
set routing-instances evpn3 route-distinguisher 1:32
set routing-instances evpn3 vrf-target target:1:7
set routing-instances evpn3 protocols evpn extended-vlan-list 3
```

```
set routing-instances evpn3 switch-options service-id 3
set routing-instances evpn3 vlans v1 vlan-id 3
set routing-instances evpn3 vlans v1 l3-interface irb.3
set routing-instances evpn3 vlans v1 no-arp-suppression
```

PE3: Configuring EVPN-MPLS

Step-by-Step Procedure

1. Configure the loopback interface, and the interfaces connected to the other PE devices.

```
[edit]
user@switch# set interfaces lo0 unit 0 family inet address 198.51.100.3/32 primary
user@switch# set interfaces xe-2/0/2 unit 0 family inet address 192.0.2.1/24
user@switch# set interfaces xe-2/0/2 unit 0 family mpls
user@switch# set interfaces xe-2/0/3 unit 0 family inet address 192.0.2.11/24
user@switch# set interfaces xe-2/0/3 unit 0 family mpls
```

2. Configure interface xe-2/0/6, which is connected to the host.

```
[edit]
user@switch# set interfaces xe-2/0/6 enable
user@switch# set interfaces xe-2/0/6 flexible-vlan-tagging
user@switch# set interfaces xe-2/0/6 encapsulation flexible-ethernet-services
user@switch# set interfaces xe-2/0/6 unit 1 family ethernet-switching
  interface-mode trunk
user@switch# set interfaces xe-2/0/6 unit 1 family ethernet-switching vlan members
  1
user@switch# set interfaces xe-2/0/6 unit 2 family ethernet-switching
  interface-mode trunk
user@switch# set interfaces xe-2/0/6 unit 2 family ethernet-switching vlan members
  2
user@switch# set interfaces xe-2/0/6 unit 3 family ethernet-switching
  interface-mode trunk
user@switch# set interfaces xe-2/0/6 unit 3 family ethernet-switching vlan members
  3
```

3. Configure IRB interfaces irb.1, irb.2, and irb.3.

```
[edit]
user@switch# set interfaces irb unit 1 family inet address 10.2.1.3/24
  virtual-gateway-address 10.2.1.254
user@switch# set interfaces irb unit 2 family inet address 10.2.2.3/24
  virtual-gateway-address 10.2.2.254
user@switch# set interfaces irb unit 3 family inet address 10.2.3.3/24
  virtual-gateway-address 10.2.3.254
```

4. Assign a router ID and the autonomous system in which PE1, PE2, and PE3 reside.

```
[edit]
user@switch# set routing-options router-id 198.51.100.3
```

```
user@switch# set routing-options autonomous-system 65000
```

5. Enable per-packet load-balancing for EVPN routes when EVPN multihoming active-active mode is used.

```
[edit]
user@switch# set routing-options forwarding-table export evpn-pplb
user@switch# set policy-options policy-statement evpn-pplb from protocol evpn
user@switch# set policy-options policy-statement evpn-pplb then load-balance
  per-packet
```

6. Enable MPLS on interfaces xe-2/0/2.0 and xe-2/0/3.0.

```
[edit]
user@switch# set protocols mpls interface xe-2/0/2.0
user@switch# set protocols mpls interface xe-2/0/3.0
```

7. Configure an IBGP overlay that includes PE1, PE2, and PE3.

```
[edit]
user@switch# set protocols bgp group evpn type internal
user@switch# set protocols bgp group evpn local-address 198.51.100.3
user@switch# set protocols bgp group evpn family evpn signaling
user@switch# set protocols bgp group evpn local-as 65000
user@switch# set protocols bgp group evpn neighbor 198.51.100.1
user@switch# set protocols bgp group evpn neighbor 198.51.100.2
```

8. Configure OSPF as the internal routing protocol for EVPN by specifying an area ID and interfaces on which EVPN-MPLS is enabled.

```
[edit]
user@switch# set protocols ospf area 0.0.0.0 interface lo0.0
user@switch# set protocols ospf area 0.0.0.0 interface xe-2/0/2.0
user@switch# set protocols ospf area 0.0.0.0 interface xe-2/0/3.0
```

9. Configure the LDP on the loopback interface and the interfaces on which EVPN-MPLS is enabled.

```
[edit]
user@switch# set protocols ldp interface lo0.0
user@switch# set protocols ldp interface xe-2/0/2.0
user@switch# set protocols ldp interface xe-2/0/3.0
```

10. Configure virtual switch routing instances for VLAN v1, which is assigned VLAN IDs of 1, 2, and 3, and include the interfaces and other entities associated with the VLAN.

```
[edit]
```

```
user@switch# set routing-instances evpn1 instance-type virtual-switch
user@switch# set routing-instances evpn1 interface xe-2/0/6.1
user@switch# set routing-instances evpn1 route-distinguisher 1:12
user@switch# set routing-instances evpn1 vrf-target target:1:5
user@switch# set routing-instances evpn1 protocols evpn extended-vlan-list 1
user@switch# set routing-instances evpn1 switch-options service-id 1
user@switch# set routing-instances evpn1 vlans v1 vlan-id 1
user@switch# set routing-instances evpn1 vlans v1 l3-interface irb.1
user@switch# set routing-instances evpn1 vlans v1 no-arp-suppression
user@switch# set routing-instances evpn2 instance-type virtual-switch
user@switch# set routing-instances evpn2 interface xe-2/0/6.2
user@switch# set routing-instances evpn2 route-distinguisher 1:22
user@switch# set routing-instances evpn2 vrf-target target:1:6
user@switch# set routing-instances evpn2 protocols evpn extended-vlan-list 2
user@switch# set routing-instances evpn2 switch-options service-id 2
user@switch# set routing-instances evpn2 vlans v1 vlan-id 2
user@switch# set routing-instances evpn2 vlans v1 l3-interface irb.2
user@switch# set routing-instances evpn2 vlans v1 no-arp-suppression
user@switch# set routing-instances evpn3 instance-type virtual-switch
user@switch# set routing-instances evpn3 interface xe-2/0/6.3
user@switch# set routing-instances evpn3 route-distinguisher 1:32
user@switch# set routing-instances evpn3 vrf-target target:1:7
user@switch# set routing-instances evpn3 protocols evpn extended-vlan-list 3
user@switch# set routing-instances evpn3 switch-options service-id 3
user@switch# set routing-instances evpn3 vlans v1 vlan-id 3
user@switch# set routing-instances evpn3 vlans v1 l3-interface irb.3
user@switch# set routing-instances evpn3 vlans v1 no-arp-suppression
```

Release History Table

Release	Description
19.1R1	Starting with Junos OS Release 19.1R1, the **no-arp-suppression** configuration statement is no longer supported on any device.
17.4R1	Starting with Junos OS Release 17.4R1, you can use Ethernet VPN (EVPN) to extend a Junos Fusion Enterprise or multichassis link aggregation group (MC-LAG) network over an MPLS network to a data center or campus network.

CHAPTER 7

Troubleshooting Multichassis Link Aggregation

Troubleshooting Multichassis Link Aggregation

Use the following information to troubleshoot multichassis link aggregation configuration issues:

MAC Addresses Learned on Multichassis Aggregated Ethernet Interfaces Are Not Removed from the MAC Address Table

Problem **Description:** When both of the mulitchassis aggregated Ethernet interfaces on both connected multichassis link aggregation group (MC-LAG) peers are down, the MAC addresses learned on the multichassis aggregated Ethernet interfaces are not removed from the MAC address table.

For example, if you disable the multichassis aggregated Ethernet interface (ae0) on both MC-LAG peers by issuing the **set interfaces ae0 disable** command and commit the configuration, the MAC table still shows the MAC addresses as being learned on the multichassis aggregated Ethernet interfaces of both MC-LAG peers.

```
user@switchA> show ethernet-switching table
Ethernet-switching table: 6 entries, 2 learned, 0 persistent entries
  VLAN             MAC address       Type         Age Interfaces
  v10               *                 Flood          - All-members
  v10               00:00:5E:00:53:00 Learn(L)     3:55 ae0.0 (MCAE)
  v10               00:00:5E:00:53:01 Learn(R)        0 xe-0/0/9.0
  v20               *                 Flood          - All-members
  v30               *                 Flood          - All-members
  v30               00:00:5E:00:53:03 Static         - Router
```

```
user@switchB> show ethernet-switching table
Ethernet-switching table: 6 entries, 2 learned, 0 persistent entries
  VLAN             MAC address       Type         Age Interfaces
  v10               *                 Flood          - All-members
  v10               00:00:5E:00:53:04 Learn(R)        0 ae0.0 (MCAE)
  v10               00:00:5E:00:53:05 Learn         40 xe-0/0/10.0
  v20               *                 Flood          - All-members
  v30               *                 Flood          - All-members
  v30               00:00:5E:00:53:06 Static         - Router
```

Solution This is expected behavior.

MC-LAG Peer Does Not Go into Standby Mode

Problem **Description:** A multichassis link aggregation group (MC-LAG) peer does not go into standby mode if the MC-LAG peer IP address specified in the Inter-Chassis Control Protocol (ICCP) configuration and the IP address specified in the multichassis protection configuration are different.

Solution To prevent failure to enter standby mode, make sure that the peer IP address in the ICCP configurations and the IP address in multichassis protection configurations are the same.

Secondary MC-LAG Peer with Status Control Set to Standby Becomes Inactive

Problem **Description:** When the interchassis control link-protection link (ICL-PL) and multichassis aggregated Ethernet interfaces go down on the primary multichassis link aggregation group (MC-LAG) peer, the secondary MC-LAG peer's multichassis aggregated Ethernet interfaces with status control set to standby become inactive instead of active.

Solution This is expected behavior.

Redirect Filters Take Priority over User-Defined Filters

Problem **Description:** Multichassis link aggregation group (MC-LAG) implicit failover redirection filters take precedence over user-configured explicit filters.

Solution This is expected behavior.

Operational Command Output Is Wrong

Problem **Description:** After you deactivate Inter-Chassis Control Protocol (ICCP), the **show iccp** operational command output still shows registered client daemons, such as mcsnoopd, lacpd, and eswd.

For example:

```
user@switch> show iccp
Client Application: MCSNOOPD
  Redundancy Group IDs Joined: None

Client Application: lacpd
  Redundancy Group IDs Joined: 1

Client Application: eswd
  Redundancy Group IDs Joined: 1
```

The **show iccp** command output always shows registered modules regardless of whether or not ICCP peers are configured.

Solution This is expected behavior.

ICCP Connection Might Take Up to 60 Seconds to Become Active

Problem **Description:** When the Inter-Chassis Control Protocol (ICCP) configuration and the routed VLAN interface (RVI) configuration are committed together, the ICCP connection might take up to 60 seconds to become active.

Solution This is expected behavior.

MAC Address Age Learned on a Multichassis Aggregated Ethernet Interface Is Reset to Zero

Problem **Description:** When you activate and then deactivate an interchassis link-protection link (ICL-PL), the MAC address age learned on the multichassis aggregated Ethernet interface is reset to zero. The next-hop interface changes trigger MAC address updates in the hardware, which then triggers aging updates in the Packet Forwarding Engine. The result is that the MAC address age is updated to zero.

For example, the ICL-PL has been deactivated, and the **show ethernet-switching table** command output shows that the MAC addresses have an age of 0.

user@switch> show ethernet-switching table

```
Ethernet-switching table: 3 entries, 2 learned, 0 persistent entries
  VLAN              MAC address       Type         Age Interfaces
  v100              *                 Flood          - All-members
  v100              00:10:00:00:00:01 Learn(L)     0 ae0.0 (MCAE)
  v100              00:10:00:00:00:02 Learn(L)     0 ae0.0 (MCAE)
```

Solution This is expected behavior.

MAC Address Is Not Learned Remotely in a Default VLAN

Problem **Description:** On a QFX3500 switch running Junos OS Release 12.3 or earlier, if a multichassis link aggregation group (MC-LAG) peer learns a MAC address in the default VLAN, Inter-Chassis Control Protocol (ICCP) does not synchronize the MAC address with the MAC address of the other MC-LAG peer.

Solution This is expected behavior.

Snooping Entries Learned on Multichassis Aggregated Ethernet Interfaces Are Not Removed

Problem **Description:** When multichassis aggregated Ethernet interfaces are configured on a VLAN that is enabled for multicast snooping, the membership entries learned on the multichassis aggregated Ethernet interfaces on the VLAN are not cleared when the multichassis aggregated Ethernet interfaces go down. This is done to speed up convergence time when the interfaces come up, or come up and go down.

Solution This is expected behavior.

ICCP Does Not Come Up After You Add or Delete an Authentication Key

Problem **Description:** The Inter-Chassis Control Protocol (ICCP) connection is not established when you add an authentication key and then delete it only at the global ICCP level. However, authentication works correctly at the ICCP peer level.

Solution Delete the ICCP configuration, and then add the ICCP configuration.

Local Status Is Standby When It Should Be Active

Problem **Description:** If the multichassis aggregated Ethernet interface is down when the state machine is in a synchronized state, the multichassis link aggregation group (MC-LAG) peer local status is standby. If the multichassis aggregated Ethernet interface goes down after the state machine is in an active state, then the local status remains active, and the local state indicates that the interface is down.

Solution This is expected behavior.

Packets Loop on the Server When ICCP Fails

Problem **Description:** When you enable backup liveness detection for a multichassis link aggregation group (MC-LAG), and the backup liveness detection packets are lost because of a temporary failure on the MC-LAG, then both of the peers in the MC-LAG remain active. If this happens, both of the MC-LAG peers send packets to the connected server.

Solution This is expected behavior.

Both MC-LAG Peers Use the Default System ID After a Reboot or an ICCP Configuration Change

Problem **Description:** After a reboot or after a new Inter-Chassis Control Protocol (ICCP) configuration has been committed, and the ICCP connection does not become active, the Link Aggregation Control Protocol (LACP) messages transmitted over the multichassis

aggregated Ethernet interfaces use the default system ID. The configured system ID is used instead of the default system ID only after the MC-LAG peers synchronize with each other.

Solution This is expected behavior.

No Commit Checks Are Done for ICL-PL Interfaces

Problem **Description:** There are no commit checks on the interface being configured as an interchassis link-protection link (ICL-PL), so you must provide a valid interface name for the ICL-PL.

Solution This is expected behavior.

Double Failover Scenario

Problem **Description:** If the following events happen in this exact order—Inter-Chassis Control Protocol (ICCP) goes down, and the multichassis aggregated Ethernet interface on the multichassis link aggregation group (MC-LAG) peer in active mode goes down—a double failover occurs. In this scenario, the MC-LAG peer in standby mode does not detect what happens on the active MC-LAG peer. The MC-LAG peer in standby mode operates as if the multichassis aggregated Ethernet interface on the MC-LAG in active mode were up and blocks the interchassis link-protection link (ICL-PL) traffic. The ICL-PL traffic is not forwarded.

Solution This is expected behavior.

Multicast Traffic Floods the VLAN When the ICL-PL Interface Goes Down and Up

Problem **Description:** When interchassis link-protection link (ICL-PL) goes down and comes up, multicast traffic is flooded to all of the interfaces in the VLAN. The Packet Forwarding Engine flag Ip4McastFloodMode for the VLAN is changed to MCAST_FLOOD_ALL. This problem only occurs when a multichassis link aggregation group (MC-LAG) is configured for Layer 2.

Solution This is expected behavior.

Layer 3 Traffic Sent to the Standby MC-LAG Peer Is Not Redirected to Active MC-LAG Peer

Problem **Description:** When Inter-chassis Control Protocol (ICCP) is down, the status of a remote MC-LAG peer is unknown. Even if the MC-LAG peer is configured as standby, the traffic is not redirected to this peer because it is assumed that this peer is down.

Solution This is expected behavior.

Aggregated Ethernet Interfaces Go Down

Problem **Description:** When a multichassis aggregated Ethernet interface is converted to an aggregated Ethernet interface, it retains some multichassis aggregated Ethernet interface properties. For example, the aggregated Ethernet interface might retain the administrative key of the multichassis aggregated Ethernet interface. When this happens, the aggregated Ethernet interface goes down.

Solution Restart Link Aggregation Control Protocol (LACP) on the multichassis link aggregation group (MC-LAG) peer hosting the aggregated Ethernet interface to bring up the aggregated Ethernet interface. Restarting LACP removes the multichassis aggregated Ethernet properties of the aggregated Ethernet interface.

Flooding of Upstream Traffic

Problem **Description:** When MAC synchronization is enabled, the multichassis link aggregation group (MC-LAG) peer can resolve Address Resolution Protocol (ARP) entries for the MC-LAG routed VLAN interface (RVI) with either of the MC-LAG peer MAC addresses. If the downstream traffic is sent with one MAC address (MAC1) but the peer has resolved the MAC address with a different MAC address (MAC2), the MAC2 address might not be learned by any of the access layer switches. Flooding of the upstream traffic for the MAC2 address might then occur.

Solution Make sure that downstream traffic is sent from the MC-LAG peers periodically to prevent the MAC addresses from aging out.

ARP and MAC Table Entries Become Out of Sync in an MC-LAG Configuration

Problem **Description:** The ARP and MAC address tables in an MC-LAG configuration normally stay synchronized, but updates might be lost in extreme situations when table updates are happening very frequently, such as when link flapping occurs in an MC-LAG group.

Solution To avoid ARP and MAC entries becoming out of sync in an MC-LAG configuration, you can configure the arp-l2-validate option on the switch's IRB interface, as follows:

```
user@switch> set interfaces irb arp-l2-validate
```

The **arp-l2-validate** option is available only on QFX Series switches and EX4300 switches starting with Junos OS Release 15.1R4, and EX9200 switches starting with Junos OS Release 13.2R4.

This option turns on validation of ARP and MAC table entries, automatically applying updates if they become out of sync. You might want to enable this option as a workaround

when the network is experiencing other issues that also cause loss of ARP and MAC synchronization, but disable it during normal operation because this option might impact performance in scale configurations.

Configuring Interface Diagnostics Tools to Test the Physical Layer Connections

Configuring Loopback Testing

Loopback testing allows you to verify the connectivity of a circuit. You can configure any of the following interfaces to execute a loopback test: aggregated Ethernet, Fast Ethernet, Gigabit Ethernet, E1, E3, *N*xDS0, serial, SONET/SDH, T1, and T3.

The physical path of a network data circuit usually consists of segments interconnected by devices that repeat and regenerate the transmission signal. The transmit path on one device connects to the receive path on the next device. If a circuit fault occurs in the form of a line break or a signal corruption, you can isolate the problem by using a loopback test. Loopback tests allow you to isolate segments of the circuit and test them separately.

To do this, configure a *line loopback* on one of the routers. Instead of transmitting the signal toward the far-end device, the line loopback sends the signal back to the originating router. If the originating router receives back its own Data Link Layer packets, you have verified that the problem is beyond the originating router. Next, configure a line loopback farther away from the local router. If this originating router does not receive its own Data Link Layer packets, you can assume that the problem is on one of the segments between the local router and the remote router's interface card. In this case, the next troubleshooting step is to configure a line loopback closer to the local router to find the source of the problem.

The following types of loopback testing are supported by Junos OS:

- DCE local—Loops packets back on the local data circuit-terminating equipment (DCE).
- DCE remote—Loops packets back on the remote DCE.
- Local—Useful for troubleshooting physical PIC errors. Configuring local loopback on an interface allows transmission of packets to the channel service unit (CSU) and then to the circuit toward the far-end device. The interface receives its own transmission, which includes data and timing information, on the local router's PIC. The data received from the CSU is ignored. To test a local loopback, issue the **show interfaces *interface-name*** command. If PPP keepalives transmitted on the interface are received by the PIC, the **Device Flags** field contains the output **Loop-Detected**.
- Payload—Useful for troubleshooting the physical circuit problems between the local router and the remote router. A payload loopback loops data only (without clocking

information) on the remote router's PIC. With payload loopback, overhead is recalculated.

- Remote—Useful for troubleshooting the physical circuit problems between the local router and the remote router. A remote loopback loops packets, including both data and timing information, back on the remote router's interface card. A router at one end of the circuit initiates a remote loopback toward its remote partner. When you configure a remote loopback, the packets received from the physical circuit and CSU are received by the interface. Those packets are then retransmitted by the PIC back toward the CSU and the circuit. This loopback tests all the intermediate transmission segments.

Table 17 on page 415 shows the loopback modes supported on the various interface types.

Table 17: Loopback Modes by Interface Type

Interface	Loopback Modes	Usage Guidelines
Aggregated Ethernet, Fast Ethernet, Gigabit Ethernet	Local	*Configuring Ethernet Loopback Capability*
Circuit Emulation E1	Local and remote	*Configuring E1 Loopback Capability*
Circuit Emulation T1	Local and remote	*Configuring T1 Loopback Capability*
E1 and E3	Local and remote	*Configuring E1 Loopback Capability* and *Configuring E3 Loopback Capability*
*N*xDS0	Payload	*Configuring NxDS0 IQ and IQE Interfaces*, *Configuring T1 and NxDS0 Interfaces*, *Configuring Channelized OC12/STM4 IQ and IQE Interfaces (SONET Mode)*, *Configuring Fractional E1 IQ and IQE Interfaces*, and *Configuring Channelized T3 IQ Interfaces*
Serial (V.35 and X.21)	Local and remote	*Configuring Serial Loopback Capability*
Serial (EIA-530)	DCE local, DCE remote, local, and remote	*Configuring Serial Loopback Capability*
SONET/SDH	Local and remote	*Configuring SONET/SDH Loopback Capability to Identify a Problem as Internal or External*
T1 and T3	Local, payload, and remote	*Configuring T1 Loopback Capability* and *Configuring T3 Loopback Capability* See also *Configuring the T1 Remote Loopback Response*

To configure loopback testing, include the **loopback** statement:

```
user@host# loopback mode;
```

You can include this statement at the following hierarchy levels:

- **[edit interfaces *interface-name* aggregated-ether-options]**
- **[edit interfaces *interface-name* ds0-options]**
- **[edit interfaces *interface-name* e1-options]**
- **[edit interfaces *interface-name* e3-options]**
- **[edit interfaces *interface-name* fastether-options]**
- **[edit interfaces *interface-name* gigether-options]**
- **[edit interfaces *interface-name* serial-options]**
- **[edit interfaces *interface-name* sonet-options]**
- **[edit interfaces *interface-name* t1-options]**
- **[edit interfaces *interface-name* t3-options]**

Configuring BERT Testing

To configure BERT:

- Configure the duration of the test.

  ```
  [edit interfaces interface-name interface-type-options]
  user@host#bert-period seconds;
  ```

  You can configure the BERT period to last from 1 through 239 seconds on some PICs and from 1 through 240 seconds on other PICs. By default, the BERT period is 10 seconds.

- Configure the error rate to monitor when the inbound pattern is received.

  ```
  [edit interfaces interface-name interface-type-options]
  user@host#bert-error-rate rate;
  ```

  rate is the bit error rate. This can be an integer from 0 through 7, which corresponds to a bit error rate from 10^{-0} (1 error per bit) to 10^{-7} (1 error per 10 million bits).

- Configure the bit pattern to send on the transmit path.

  ```
  [edit interfaces interface-name interface-type-options]
  user@host#bert-algorithm algorithm;
  ```

  algorithm is the pattern to send in the bit stream. For a list of supported algorithms, enter a **?** after the **bert-algorithm** statement; for example:

  ```
  [edit interfaces t1-0/0/0 t1-options]
  ```

  ```
  user@host# set bert-algorithm ?
  ```

  ```
  Possible completions:
    pseudo-2e11-o152     Pattern is 2^11 -1 (per O.152 standard)
    pseudo-2e15-o151     Pattern is 2^15 - 1 (per O.152 standard)
    pseudo-2e20-o151     Pattern is 2^20 - 1 (per O.151 standard)
  ```

```
pseudo-2e20-o153     Pattern is 2^20 - 1 (per O.153 standard)
      ...
```

For specific hierarchy information, see the individual interface types.

NOTE: The four-port E1 PIC supports only the following algorithms:

```
pseudo-2e11-o152     Pattern is 2^11 -1 (per O.152 standard)
pseudo-2e15-o151     Pattern is 2^15 - 1 (per O.151 standard)
pseudo-2e20-o151     Pattern is 2^20 - 1 (per O.151 standard)
pseudo-2e23-o151     Pattern is 2^23 (per O.151 standard)
```

When you issue the help command from the CLI, all BERT algorithm options are displayed, regardless of the PIC type, and no commit check is available. Unsupported patterns for a PIC type can be viewed in system log messages.

NOTE: The 12-port T1/E1 Circuit Emulation (CE) PIC supports only the following algorithms:

```
all-ones-repeating   Repeating one bits
all-zeros-repeating  Repeating zero bits
alternating-double-ones-zeros  Alternating pairs of ones and zeros
alternating-ones-zeros  Alternating ones and zeros
pseudo-2e11-o152     Pattern is 2^11 -1 (per O.152 standard)
pseudo-2e15-o151     Pattern is 2^15 - 1 (per O.151 standard)
pseudo-2e20-o151     Pattern is 2^20 - 1 (per O.151 standard)
pseudo-2e7           Pattern is 2^7 - 1
pseudo-2e9-o153      Pattern is 2^9 - 1 (per O.153 standard)
repeating-1-in-4     1 bit in 4 is set
repeating-1-in-8     1 bit in 8 is set
repeating-3-in-24    3 bits in 24 are set
```

When you issue the help command from the CLI, all BERT algorithm options are displayed, regardless of the PIC type, and no commit check is available. Unsupported patterns for a PIC type can be viewed in system log messages.

NOTE: The IQE PICs support only the following algorithms:

```
all-ones-repeating   Repeating one bits
all-zeros-repeating  Repeating zero bits
alternating-double-ones-zeros  Alternating pairs of ones and zeros
alternating-ones-zeros  Alternating ones and zeros
pseudo-2e9-o153      Pattern is 2^9  -1 (per 0.153 (511 type) standard)
pseudo-2e11-o152    Pattern is 2^11 -1 (per 0.152 and 0.153 (2047 type)
 standards)
pseudo-2e15-o151     Pattern is 2^15 -1 (per 0.151 standard)
pseudo-2e20-o151     Pattern is 2^20 -1 (per 0.151 standard)
pseudo-2e20-o153     Pattern is 2^20 -1 (per 0.153 standard)
pseudo-2e23-o151     Pattern is 2^23 -1 (per 0.151 standard)
repeating-1-in-4     1 bit in 4 is set
repeating-1-in-8     1 bit in 8 is set
repeating-3-in-24    3 bits in 24 are set
```

When you issue the help command from the CLI, all BERT algorithm options are displayed, regardless of the PIC type, and no commit check is available. Unsupported patterns for a PIC type can be viewed in system log messages.

NOTE: BERT is supported on the PDH interfaces of the Channelized SONET/SDH OC3/STM1 (Multi-Rate) MIC with SFP and the DS3/E3 MIC. The following BERT algorithms are supported:

```
all-ones-repeating            Repeating one bits
all-zeros-repeating           Repeating zero bits
alternating-double-ones-zeros Alternating pairs of ones and zeros
alternating-ones-zeros        Alternating ones and zeros
repeating-1-in-4              1 bit in 4 is set
repeating-1-in-8              1 bit in 8 is set
repeating-3-in-24             3 bits in 24 are set
pseudo-2e9-o153               Pattern is 2^9 - 1 (per 0.153 standard)
pseudo-2e11-o152              Pattern is 2^11 - 1 (per 0.152 standard)
pseudo-2e15-o151              Pattern is 2^15 - 1 (per 0.151 standard)
pseudo-2e20-o151              Pattern is 2^20 - 1 (per 0.151 standard)
pseudo-2e20-o153              Pattern is 2^20 - 1 (per 0.153 standard)
pseudo-2e23-o151              Pattern is 2^23 (per 0.151 standard)
```

Table 18 on page 418 shows the BERT capabilities for various interface types.

Table 18: BERT Capabilities by Interface Type

Interface	T1 BERT	T3 BERT	Comments
12-port T1/E1 Circuit Emulation	Yes (ports 0–11)	—	• Limited algorithms

Table 18: BERT Capabilities by Interface Type (continued)

Interface	T1 BERT	T3 BERT	Comments
4-port Channelized OC3/STM1 Circuit Emulation	Yes (port 0–3)	—	• Limited algorithms
E1 or T1	Yes (port 0–3)	Yes (port 0–3)	• Single port at a time • Limited algorithms
E3 or T3	Yes (port 0–3)	Yes (port 0–3)	• Single port at a time
Channelized OC12	—	Yes (channel 0–11)	• Single channel at a time • Limited algorithms • No bit count
Channelized STM1	Yes (channel 0–62)	—	• Multiple channels • Only one algorithm • No error insert • No bit count
Channelized T3 and Multichannel T3	Yes (channel 0–27)	Yes (port 0–3 on channel 0)	• Multiple ports and channels • Limited algorithms for T1 • No error insert for T1 • No bit count for T1

These limitations do not apply to channelized IQ interfaces. For information about BERT capabilities on channelized IQ interfaces, see *Channelized IQ and IQE Interfaces Properties*.

Starting and Stopping a BERT Test

Before you can start the BERT test, you must disable the interface. To do this, include the **disable** statement at the **[edit interfaces *interface-name*]** hierarchy level:

```
[edit interfaces interface-name]
disable;
```

After you configure the BERT properties and commit the configuration, begin the test by issuing the **test interface *interface-name interface-type*-bert-start** operational mode command:

```
user@host> test interface interface-name interface-type-bert-start
```

The test runs for the duration you specify with the **bert-period** statement. If you want to terminate the test sooner, issue the **test interface *interface-name interface-type*-bert-stop** command:

```
user@host> test interface interface-name interface-type-bert-stop
```

For example:

```
user@host> test interface t3-1/2/0 t3-bert-start
user@host> test interface t3-1/2/0 t3-bert-stop
```

To view the results of the BERT test, issue the **show interfaces extensive | find BERT** command:

```
user@host> show interfaces interface-name extensive | find BERT
```

For more information about running and evaluating the results of the BERT procedure, see the CLI Explorer.

NOTE: To exchange BERT patterns between a local router and a remote router, include the **loopback remote** statement in the interface configuration at the remote end of the link. From the local router, issue the **test interface** command.

Related Documentation

- *show interfaces diagnostics optics (Gigabit Ethernet, 10-Gigabit Ethernet, 40-Gigabit Ethernet, 100-Gigabit Ethernet, and Virtual Chassis Port)*

CHAPTER 8

Configuration Statements

- status-control on page 446
- session-establishment-hold-time on page 447
- threshold (Detection Time) on page 448
- transmit-interval (Liveness Detection) on page 449
- version (Liveness Detection) on page 449

apply-groups

Syntax apply-groups [*group-names*];

Hierarchy Level All hierarchy levels

Release Information Statement introduced before Junos OS Release 7.4.

Description Apply a configuration group to a specific hierarchy level in a configuration, to have a configuration inherit the statements in the configuration group.

You can specify more than one group name. You must list them in order of inheritance priority. The configuration data in the first group takes priority over the data in subsequent groups.

Options ***group-names***—One or more names specified in the **groups** statement.

Required Privilege Level configure—To enter configuration mode, but other required privilege levels depend on where the statement is located in the configuration hierarchy.

Related Documentation

- *Applying a Junos OS Configuration Group*
- groups on page 430

arp-enhanced-scale

Syntax arp-enhanced-scale;

Hierarchy Level [edit system]

Release Information Statement introduced in Junos OS Release 19.1R1 for QFX10008 and QFX10016 switches.

Description Increases the number of ARP and neighbor discovery entries for MC-LAG configured with enhanced convergence and Layer 3 VXLAN deployments.

NOTE: To increase the ARP and discovery entries for MC-LAG with enhanced convergence, you also need to enable the enhanced-convergence statement at the [edit system] hierarchy. For information on how to configure enhanced convergence, see "Understanding Multichassis Link Aggregation Groups" on page 21 and "Increasing ARP and Network Discovery Protocol Entries for Enhanced MC-LAG and Layer 3 VXLAN Topologies" on page 117.

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

Related Documentation

- enhanced-convergence on page 429
- Understanding Multichassis Link Aggregation Groups on page 21

arp-l2-validate

Syntax arp-l2-validate

Hierarchy Level [edit interfaces irb]

Release Information Statement introduced in Junos OS Release 13.2R4 for EX9200 switches.
Statement introduced in Junos OS Release 15.1R4 for QFX Series switches and EX4300 switches.

Description Enables periodic checking of ARP Layer 3 addressing and MAC Layer 2 addressing tables, and fixes entries if they become out of sync.

Normally, the ARP and MAC address tables stay synchronized. However, you can configure this option on the **irb** interface of the switch to help avoid traffic loss in network conditions that might cause unresolved inconsistencies to occur between the ARP and MAC address tables, such as:

- When link flapping occurs in a multichassis link aggregation (MC-LAG) group, and the network is attempting to achieve convergence. In this case, frequent MAC table updates are happening, and occasionally a corresponding ARP table update might be lost.

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

Related Documentation

-

authentication-key (ICCP)

Syntax

```
authentication-key key;
```

Hierarchy Level

[edit protocols iccp],
[edit protocols iccp peer *peer-IP-address*]

Release Information

Statement introduced in Junos OS Release 10.0 for MX Series routers.
Statement introduced in Junos OS Release 12.2 for the QFX Series.

Description

Specify the authentication key (password). The QFX3500 and MX Series devices use this password to verify the authenticity of packets sent from the peers hosting a multichassis link aggregation group (MC-LAG). Peer-level authentication takes precedence over global-level authentication.

Inter-Chassis Control Protocol (ICCP) uses MD5 authentication.

Options

key—Authentication password. It can be 1 through 16 contiguous digits or letters. Separate decimal digits with periods. Separate hexadecimal digits with periods and precede the string with 0x. If you include spaces in the password, enclose the entire password in quotation marks (" ").

Required Privilege Level

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

backup-liveness-detection

Syntax

```
backup-liveness-detection {
  backup-peer-ip ipv4-address;
}
```

Hierarchy Level

[edit protocols iccp peer]

Release Information

Statement introduced in Junos OS Release 12.2 for the QFX Series.
Statement introduced in Junos OS Release 13.2R1 for EX Series switches.
Statement introduced in Junos OS Release 15.2R1 for MX Series routers.

Description

Determine whether a peer is up or down by exchanging keepalive messages over the management link between the two Inter-Chassis Control Protocol (ICCP) peers.

When an ICCP connection is operationally down, the status of the peers hosting a multichassis link aggregation group (MC-LAG) is detected by sending liveness detection requests to each other. Peers must respond to liveness detection requests within a specified amount of time. If the responses are not received within that time for a given number of consecutive attempts, the liveness detection check fails, and a failure action is implemented. Backup liveness detection must be configured on both peers hosting the MC-LAG.

The remaining statement is explained separately. See CLI Explorer.

NOTE: If backup liveness detection is configured, the peer status is always up when either the ICCP TCP Connection is established, or Bidirectional Forwarding Protocol (BFD) is configured and the peer is up. The backup liveness check is only performed when the ICCP connection is down.

Required Privilege Level

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

Related Documentation

- Configuring Multichassis Link Aggregation on MX Series Routers on page 45

backup-peer-ip

Syntax backup-peer-ip *ipv4-address*;

Hierarchy Level [edit protocols iccp peer backup-liveness-detection]

Release Information Statement introduced in Junos OS Release 12.2 for the QFX Series.
Statement introduced in Junos OS Release 13.2R1 for EX Series switches.

Description Specify the IP address of the peer being used as a backup peer in the Bidirectional Forwarding Detection (BFD) configuration.

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

bgp-peer

Syntax bgp-peer *ip-address*;

Hierarchy Level [edit routing-instances *name* protocols evpn mclag]

Release Information Statement introduced in Junos OS Release 17.4R1 on MX Series routers, EX Series switches, and Junos Fusion Enterprise.

Description Configure an aggregation device in a Junos Fusion Enterprise or a multichassis link aggregation group (MC-LAG) topology to interwork with an Ethernet VPN-MPLS (EVPN-MPLS) device.

Options ***ip-address***—IP address of the BGP peer. Typically, a BGP peer is identified by the IP address of the device's loopback interface.

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

Related Documentation
- Understanding EVPN-MPLS Interworking with Junos Fusion Enterprise and MC-LAG on page 382

chassis-id

Syntax

```
chassis-id chassis-id;
```

Hierarchy Level

```
[editinterfaces aggregated-ether-options mc-ae]
```

Release Information

Statement introduced in Junos OS Release 12.2 for the QFX Series.
Statement introduced in Junos OS Release 12.3R2 for EX Series switches.

Description

Specify the chassis ID of the multichassis aggregated Ethernet interface device. LACP uses the chassis ID to calculate the port number of the multichassis link aggregation group (MC-LAG) physical member links.

Required Privilege Level

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

detection-time (Liveness Detection)

Syntax

```
detection-time {
  threshold milliseconds;
}
```

Hierarchy Level

```
[edit protocols iccp peer liveness-detection]
```

Release Information

Statement introduced in Junos OS Release 10.0 for MX Series routers.
Statement introduced in Junos OS Release 12.2 for the QFX Series.

Description

The Bidirectional Forwarding Detection (BFD) timers are adaptive and can be adjusted to be faster or slower.

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.

enhanced-convergence

Syntax

```
enhanced-convergence;
```

Hierarchy Level

[edit interfaces ae*X* aggregated-ether-options mc-ae]
[edit interfaces irb unit *unit-number*]

Release Information

Statement introduced in Junos OS Release 15.1R1.
Statement introduced in Junos OS Release 15.1X53-D60 for the QFX Series.

Description

NOTE: On EX9200 and QFX10000 switches, enhanced convergence is applicable for unicast traffic only—for example, when a MAC address is learned over an MC-AE interface, or an ARP entry is resolved over an MC-AE interface.

NOTE: Enhanced convergence is not supported on QFX5100, QFX5110, QFX5120, QFX5200-48Y, QFX5200-32C, and QFX5210-64C switches.

Improves Layer 2 and Layer 3 convergence time when a multichassis aggregated Ethernet (MC-AE) link goes down or comes up in a bridge domain or VLAN. Convergence time is improved because the traffic on the MC-AE interface is switched to the interchassis link (ICL) without waiting for a MAC address update.

If you have configured an IRB interface over an MC-AE interface that has enhanced convergences enabled, then you must configure enhanced convergence on the IRB interface as well. Enhanced convergence must be enabled for both Layer 2 and Layer 3 interfaces.

Required Privilege Level

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

Related Documentation

- Configuring Multichassis Link Aggregation on MX Series Routers on page 45
- Configuring Multichassis Link Aggregation on page 56

groups

Syntax

```
groups {
  group-name {
    configuration-data;
    when {
      chassis chassis-id;
      member member-id;
      model model-id;
      node node-id;
      peers [ names-of-peers ]
      routing-engine routing-engine-id;
      time <start-time> [to <end-time>];
    }
      conditional-data;
  }
  lccn-re0 {
    configuration-data;
  }
  lccn-re1 {
    configuration-data;
  }
}
```

Hierarchy Level [edit]

Release Information Statement introduced before Junos OS Release 7.4.

Description Create a configuration group.

Options ***group-name***—Name of the configuration group. To configure multiple groups, specify more than one group name.

configuration-data—The configuration statements that are to be applied elsewhere in the configuration with the **apply-groups** statement, to have the target configuration inherit the statements in the group.

when—Define conditions under which the configuration group should be applied. Conditions include the type of chassis, model, or Routing Engine, virtual chassis member, cluster node, and start and optional end time of day. If you specify multiple conditions in a single configuration group, all conditions must be met before the configuration group is applied.

- **chassis** *chassis-id*—Specify the chassis type of the router. Valid types include SCC0, SCC1, LCC0, LCC1 ... LCC3.
- **member** *member-id*—Specify the name of the member of the virtual chassis.
- **model** *model-id*—Specify the model name of the router, such as m7i or tx100.

- **node** *node-id*—Specify the cluster node.
- **peers** *names-of-peers*—Specify the names of the MC-LAG peers participating in commit synchronization.
- **routing-engine** *routing-engine-id*—Specify the type of Routing Engine, re0 or re1.
- **time** *start-time* **[to** *end-time***]**—Specify the start time or time duration for this configuration group to be applied. If only the start time is specified, the configuration group is applied at the specified time and remains in effect until the time is changed. If the end time is specified, then on each day, the applied configuration group is started and stopped at the specified times. The syntax for specifying the time uses the format yyyy-mm-dd.hh:mm, hh:mm, or hh.

conditional-data—Option introduced in Junos 11.3. The conditional statements that are to be applied when this configuration group is applied. On routers that support multiple Routing Engines, you can also specify two special group names:

- **re0**—Configuration statements that are to be applied to the Routing Engine in slot 0.
- **re1**—Configuration statements that are to be applied to the Routing Engine in slot 1.

On routers that support multiple Routing Engines, you can also specify two special group names:

The configuration specified in group **re0** is applied only if the current Routing Engine is in slot 0; likewise, the configuration specified in group **re1** is applied only if the current Routing Engine is in slot 1. Therefore, both Routing Engines can use the same configuration file, each using only the configuration statements that apply to it. Each **re0** or **re1** group contains at a minimum the configuration for the hostname and the management interface (**fxp0**). If each Routing Engine uses a different management interface, the group also should contain the configuration for the backup router and static routes.

(Routing matrix only) The TX Matrix router supports group names for the Routing Engines in each connected T640 router in the following formats:

NOTE: The management Ethernet interface used for the TX Matrix Plus router, T1600 routers in a routing matrix, and PTX Series Packet Transport Routers, is em0. Junos OS automatically creates the router's management Ethernet interface, em0.

- **lcc*n*-re0**—Configuration statements applied to the Routing Engine in slot 0 of the specified T640 router that is connected to a TX Matrix router.
- **lcc*n*-re1**—Configuration statements applied to the specified to the Routing Engine in slot 1 of the specified T640 router that is connected to a TX Matrix router.

  n identifies the T640 router and can be from 0 through 3.

Required Privilege Level configure—To enter configuration mode.

Related Documentation

- *Creating a Junos OS Configuration Group*
- apply-groups on page 422
- *apply-groups-except*

iccp

Syntax

```
iccp {
  authentication-key string;
  local-ip-addr local-ip-addr;
  peer ip-address{
    authentication-key string;
    backup-liveness-detection {
      backup-peer-ip ip-address;
    }
    liveness-detection {
      detection-time {
        threshold milliseconds;
      }
      minimum-interval milliseconds;
      minimum-receive-interval milliseconds;
      multiplier number;
      no-adaptation;
      transmit-interval {
        minimum-interval milliseconds;
        threshold milliseconds;
      }
      version (1 | automatic);
    }
    local-ip-addr ipv4-address;
    session-establishment-hold-time seconds;
  }
  session-establishment-hold-time seconds;
  traceoptions {
    file <filename> <files number> <match regular-expression> <microsecond-stamp>
      <size size> <world-readable | no-world-readable>;
    flag flag;
    no-remote-trace;
  }
}
```

Hierarchy Level

```
[edit protocols]
```

Release Information

Statement introduced in Junos OS Release 10.0 for MX Series routers.
Statement introduced in Junos OS Release 12.2 for the QFX Series.
Statement introduced in Junos OS Release 12.3R2 for EX Series switches.

Description

Configure Inter-Chassis Control Protocol (ICCP) between the multichassis link aggregation group (MC-LAG) peers. ICCP replicates forwarding information, validates configurations, and propagates the operational state of the MC-LAG members.

NOTE: Backup liveness detection is not supported on MX Series routers.

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.

interface (Multichassis Protection)

Syntax
interface *interface-name*;

Hierarchy Level
[edit multi-chassis multi-chassis-protection peer]

Release Information
Statement introduced in Junos OS Release 9.6 for MX Series routers.
Statement introduced in Junos OS Release 12.2 for the QFX Series.
Statement introduced in Junos OS Release 12.3R2 for EX Series switches.

Description
Specify the name of the interface that is being used as an interchassis link-protection link (ICL-PL). The two switches hosting a multichassis link aggregation group (MC-LAG) use this link to pass Inter-Chassis Control Protocol (ICCP) and data traffic.

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

local-ip-addr (ICCP)

Syntax
local-ip-addr *local-ip-address*;

Hierarchy Level
[edit protocols iccp],
[edit protocols iccp peer *peer-IP-address*]

Release Information
Statement introduced in Junos OS Release 10.0 for MX Series routers.
Statement introduced in Junos OS Release 12.2 for the QFX Series.
Statement introduced in Junos OS Release 12.3R2 for EX Series switches.

Description
Specify the local IP address of the interchassis link (ICL) interface that Inter-Chassis Control Protocol (ICCP) uses to communicate to the peers that host a multichassis link aggregation group (MC-LAG).

Options
local-ip-address—Default local IP address to be used by all peers.

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

mc-ae

Syntax

```
mc-ae {
  chassis-id chassis-id;
    events {
      iccp-peer-down;
        force-icl-down;
      prefer-status-control-active;
      }
      init-delay-time seconds;
      mc-ae-id mc-ae-id;
      mode (active-active | active-standby);
      redundancy-group group-id;
      revert-time revert-time;
      status-control (active | standby);
      switchover-mode (non-revertive |revertive);
    }
```

Hierarchy Level

[edit interfaces ae*X* aggregated-ether-options],
[edit logical-systems *logical-system-name interfaces* ae*X* aggregated-ether-options]

Release Information

Statement introduced in Junos OS Release 9.6 for MX Series routers.
events statement introduced in Junos OS Release 11.4R4 for MX Series routers.
Statement introduced in Junos OS Release 12.2 for the QFX Series. Only the **chassis-id**, **mc-ae-id**, **mode active-active**, and **status-control (active | standby)** options are supported on QFX Series devices.
Statement introduced in Junos OS Release 12.3R2 for EX Series switches.
prefer-status-control-active statement introduced in Junos OS Release 13.2R1 for EX Series switches.
init-delay-time *seconds* statement introduced in Junos OS Release 13.2R3 for EX Series switches.
switchover-mode and **revert-time** statements introduced in Junos OS Release 13.3.
Support for logical systems introduced in Junos OS Release 14.1.

Description

Enable multichassis link aggregation groups (MC-LAG), which enables one device to form a logical LAG interface with two or more other devices.

Options

chassis-id—Specify the chassis ID for Link Aggregation Control Protocol (LACP) to calculate the port number of MC-LAG physical member links.

Values: 0 or 1

events—Specify an action if a specific MC-LAG event occurs.

iccp-peer-down—Specify an action if the ICCP peer of this node goes down.

force-icl-down—If the node's ICCP peer goes down, bring down the interchassis-link logical interface.

prefer-status-control-active—Specify that the node configured as **status-control active** become the active node if the peer of this node goes down.

When ICCP goes down, you can use this keyword to make a mc-lag PE to become the active PE. For example, if you want mc-lag PE1 to be Active on ICCP down, then configure this keyword in PE1. It is not recommended to configure this keyword in both the mc-lag PEs.

NOTE: The **prefer-status-control-active** statement can be configured with the **status-control standby** configuration to prevent the LACP MC-LAG system ID from reverting to the default LACP system ID on ICCP failure. Use this configuration only if you can ensure that ICCP will not go down unless the router or switch is down. You must also configure the **hold-time down** value (at the **[edit interfaces *interface-name*]** hierarchy level) for the interchassis link with the **status-control standby** configuration to be higher than the ICCP BFD timeout. This configuration prevents data traffic loss by ensuring that when the router or switch with the **status-control active** configuration goes down, the router or switch with the **status-control standby** configuration does not go into standby mode.

To make the **prefer-status-control-active** configuration work with the **status-control standby** configuration when an interchassis-link logical interface is configured on aggregate Ethernet interface, you must either configure the **lacp periodic *interval*** statement at the **[edit interface *interface-name* aggregated-ether-options]** hierarchy level as **slow** or configure the **detection-time threshold** statement at the **[edit protocols iccp peer liveness-detection]** hierarchy level as less than 3 seconds.

init-delay-time *seconds*—To minimize traffic loss, specify the number of seconds in which to delay bringing the multichassis aggregated Ethernet interface back to the up state when you reboot an MC-LAG peer.

NOTE: On QFX and EX Series switches, the default session establishment hold time is 300 seconds. However, the session establishment time must be at least 100 seconds higher than the init delay time. You can optionally update the session establishment time to be 340 seconds and the init delay time to be 240 seconds.

mc-ae-id *mc-ae-id*—Specify the identification number of the MC-LAG device. The two MC-LAG network devices that manage a given MC-LAG must have the same identification number.

Range: 1 through 65,535

mode (active-active | active-standby)—Specify whether the MC-LAG is in active-active or active-standby mode.

NOTE: You can configure IPv4 (inet) and IPv6 (inet6) addresses on mc-ae interfaces when the active-standby mode is configured.

redundancy-group *group-id*—Specify the redundancy group identification number. The Inter-Chassis Control Protocol (ICCP) uses the redundancy group ID to associate multiple chassis that perform similar redundancy functions.

Range: 1 through 4,294,967,294

revert-time—Wait interval (in minutes) before the switchover to the preferred node is performed when the **switchover-mode** is configured as revertive.

Range: 1 through 10

status-control (active | standby)—Specify whether the chassis becomes active or remains in standby mode when an interchassis link failure occurs.

switchover-mode (non-revertive | revertive)—Specify whether Junos OS should trigger a link switchover to the preferred node when the active node is available.

NOTE: For revertive mode to automatically switch over to the preferred node, the status-control statement should be configured as active.

Required Privilege Level

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

mc-ae-id

Syntax

```
mc-ae-id mc-ae-id;
```

Hierarchy Level

```
[edit interfaces aggregated-ether-options mc-ae]
```

Release Information

Statement introduced in Junos OS Release 12.2 for the QFX Series.
Statement introduced in Junos OS Release 12.3R2 for EX Series switches.

Description

Specify the multichassis aggregated Ethernet (MC-AE) identification number of the MC-AE that a given aggregated Ethernet interface belongs to. The two peers that host a given multichassis link aggregation group (MC-LAG) must have the same multichassis aggregated Ethernet ID.

Options

Range: 1 through 65535.

Required Privilege Level

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

mclag

Syntax

```
mclag {
  bgp-peer ip-address;
}
```

Hierarchy Level

```
[edit routing-instances name protocols evpn]
```

Release Information

Statement introduced in Junos OS Release 17.4R1 on MX Series routers, EX Series switches, and Junos Fusion Enterprise.

Description

Configure parameters that enable the interworking of Ethernet VPN-MPLS (EVPN-MPLS) with a Junos Fusion Enterprise or a multichassis link aggregation group (MC-LAG) topology.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

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

Related Documentation

- Understanding EVPN-MPLS Interworking with Junos Fusion Enterprise and MC-LAG on page 382

minimum-interval (Liveness Detection)

Syntax minimum-interval *milliseconds*;

Hierarchy Level [edit protocols iccp peer liveness-detection]

Release Information Statement introduced in Junos OS Release 10.0 for MX Series routers.
Statement introduced in Junos OS Release 12.2 for the QFX Series.
Statement introduced in Junos OS Release 12.3R2 for EX Series switches.

Description Configure simultaneously the minimum interval at which the peer transmits liveness detection requests and the minimum interval at which the peer expects to receive a reply from a peer with which it has established a Bidirectional Forwarding Detection (BFD) session. Optionally, instead of using this statement, you can specify the minimum transmit and receive intervals separately by using the **transmit-interval minimal-interval** and **minimum-receive-interval** statements, respectively.

Options *milliseconds*—Specify the minimum interval value for Bidirectional Forwarding Detection (BFD).

Range: 1 through 255,000

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

minimum-receive-interval (Liveness Detection)

Syntax minimum-receive-interval *milliseconds*;

Hierarchy Level [edit protocols iccp peer liveness-detection]

Release Information Statement introduced in Junos OS Release 10.0 for MX Series routers.
Statement introduced in Junos OS Release 12.2 for the QFX Series.
Statement introduced in Junos OS Release 12.3R2 for EX Series switches.

Description Configure the minimum interval at which the peer must receive a reply from a peer with which it has established a Bidirectional Forwarding Detection (BFD) session.

Options ***milliseconds***—Specify the minimum interval value.

Range: 1 through 255,000

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

mode (QFX Series)

Syntax mode active-active ;

Hierarchy Level [edit interfaces aggregated-ether-options mc-ae]

Release Information Statement introduced in Junos OS Release 12.2 for the QFX Series.
Statement introduced in Junos OS Release 12.3R2 for EX Series switches.

Description Configure the multichassis link aggregation group (MC-LAG) to be in active-active mode. In active-active mode, all of the members of the MC-LAG will be active on both routing or switching devices. Only active-active mode is supported at this time.

Options **active-active**—Specify that all of the members of the MC-LAG will be active on both routing or switching devices.

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

multiplier (Liveness Detection)

Syntax

```
multiplier number;
```

Hierarchy Level

```
[edit protocols iccp peer liveness-detection]
```

Release Information Statement introduced in Junos OS Release 12.2 for the QFX Series.

Description Configure the number of liveness detection requests not received by the peer before Bidirectional Forwarding Detection (BFD) declares the peer is down.

Options ***number***—Maximum allowable number of liveness detection requests missed by the peer.

Range: 1 through 255

Default: 3

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

multi-chassis

Syntax

```
multi-chassis {
  multi-chassis-protection peer-ip-address {
    interface interface-name;
  }
  mc-lag consistency-check;
}
```

Hierarchy Level

```
[edit]
```

Release Information Statement introduced in Junos OS Release 9.6 for MX Series routers.
Statement introduced in Junos OS Release 12.2 for the QFX Series.
Statement introduced in Junos OS Release 12.3R2 for EX Series switches.

Description Configure an interchassis link-protection link (ICL-PL) between the two peers that host a multichassis link aggregation group (MC-LAG). You can configure either an aggregated Ethernet interface or a 10-Gigabit Ethernet interface to be an ICL-PL.

The remaining statements are explained separately. See CLI Explorer.

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

multi-chassis-protection

Syntax

```
multi-chassis-protection peer-ip-address {
    interface interface-name;
}
```

Hierarchy Level

[edit multi-chassis]

Release Information

Statement introduced in Junos OS Release 9.6 for MX Series routers.
Statement introduced in Junos OS Release 12.2 for the QFX Series.
Statement introduced in Junos OS Release 12.3R2 for EX Series switches.

Description

Configure multichassis link protection between the two peers that host a multichassis link aggregation group (MC-LAG). If the Interchassis Control Protocol (ICCP) connection is up and the interchassis link (ICL) comes up, the peer configured as standby brings up the multichassis aggregated Ethernet interfaces shared with the peer. Multichassis protection must be configured on one interface for each peer.

The remaining statement is explained separately. See CLI Explorer.

Required Privilege Level

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

no-adaptation (Liveness Detection)

Syntax

```
no-adaptation;
```

Hierarchy Level

[edit protocols iccp peer liveness-detection]

Release Information

Statement introduced in Junos OS Release 10.0 for MX Series routers.
Statement introduced in Junos OS Release 12.2 for the QFX Series.

Description

Configure Bidirectional Forwarding Detection (BFD) sessions to not adapt to changing network conditions.

Required Privilege Level

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

peer (ICCP)

Syntax

```
peer ip-address {
  authentication-key string;
  backup-liveness-detection {
    backup-peer-ip ip-address;
  }
  liveness-detection {
    detection-time {
      threshold milliseconds;
    }
    minimum-intervalmilliseconds;
    minimum-receive-interval milliseconds;
    multiplier number;
    no-adaptation;
    transmit-interval {
      minimum-interval milliseconds;
      threshold milliseconds;
    }
    version (1 | automatic);
  }
  local-ip-addr ipv4-address;
  session-establishment-hold-time seconds;
}
```

Hierarchy Level

```
[edit protocols iccp]
```

Release Information

Statement introduced in Junos OS Release 10.0 for MX Series routers.
Statement introduced in Junos OS Release 12.2 for the QFX Series.
Statement introduced in Junos OS Release 12.3R2 for EX Series switches.

Description

Configure the peers that host a multichassis link aggregation group (MC-LAG). You must configure Inter-Chassis Control Protocol (ICCP) for both peers that host the MC-LAG.

NOTE: Backup liveness detection is not supported on MX Series routers.

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.

peer (Multichassis)

Syntax

```
peer ip-address {
    interface interface-name;
}
```

Hierarchy Level

[edit multi-chassis]

Release Information

Statement introduced in Junos OS Release 9.6 for MX Series routers.
Statement introduced in Junos OS Release 12.2 for the QFX Series.

Description

Configure the IP address of the peer that is part of the interchassis link-protection link (ICL-PL). If Inter-Chassis Control Protocol (ICCP) is up and the interchassis link (ICL) comes up, the peer configured as standby will bring up the multichassis aggregated Ethernet interfaces shared with the active peer specified by the **peer** statement. You must specify the physical interface of the peer.

The remaining statement is explained separately. See CLI Explorer.

Required Privilege Level

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

peers (Commit)

Syntax

```
peers {
  name of peer {
    user name of user;
    authentication string;
  }
}
```

Hierarchy Level

```
[edit system commit]
```

Release Information

Statement introduced in Junos OS Release 14.2R6 for the MX Series and Junos Fusion.
Statement introduced in Junos OS Release 15.1X53-D60 for the QFX Series.
Statement introduced in Junos OS Release 16.1R1 for the EX Series.

Description

Configure options for the peers participating in commit synchronization.

Options

name of peer—Hostname or IP address of the peer participating in commit synchronization.

user—Name of administrator configuring commit synchronization.

authentication—Plain-text password string that is stored as an encrypted password string.

Required Privilege Level

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

peers-synchronize

Syntax
```
peers-synchronize;
```

Hierarchy Level
```
[edit system commit]
```

Release Information
Statement introduced in Junos OS Release 14.2R6 for the MX Series and Junos Fusion.
Statement introduced in Junos OS Release 15.1X53-D60 for the QFX Series.
Statement introduced in Junos OS Release 16.1R1 for the EX Series.

Description
Configure the **commit** command to automatically perform a **peers-synchronize** action between peers. The local peer (or requesting peer) on which you enable the **peers-synchronize** statement copies and loads its configuration to the remote (or responding) peer. Each peer then performs a syntax check on the configuration file being committed. If no errors are found, the configuration is activated and becomes the current operational configuration on both peers.

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

Related Documentation
- *server*
- *synchronize*

status-control

Syntax
```
status-control (active | standby);
```

Hierarchy Level
```
[edit interfaces aggregated-ether-options mc-ae]
```

Release Information
Statement introduced in Junos OS Release 12.2 for the QFX Series.
Statement introduced in Junos OS Release 12.3R2 for EX Series switches.

Description
Specify whether a peer hosting a multichassis link aggregation group (MC-LAG) is primary or secondary. Primary is considered active, and secondary is considered standby.

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

session-establishment-hold-time

Syntax
session-establishment-hold-time *seconds*;

Hierarchy Level
[edit protocols iccp],
[edit protocols iccp peer]

Release Information
Statement introduced in Junos OS Release 10.0 for MX Series routers.
Statement introduced in Junos OS Release 12.2 for the QFX Series.
Statement introduced in Junos OS Release 12.3R2 for EX Series switches.

Description
Specify the time during which an Inter-Chassis Control Protocol (ICCP) connection must be established between peers.

NOTE: On QFX and EX Series switches, the default session establishment hold time is 300 seconds. However, the session establishment time must be at least 100 seconds higher than the init delay time. You can optionally update the session establishment time to be 340 seconds and the init delay time to be 240 seconds.

Options
seconds—Time (in seconds) within which a successful ICCP connection must be established.

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

threshold (Detection Time)

Syntax threshold *milliseconds*;

Hierarchy Level [edit protocols iccp peer liveness-detection detection-time]

Release Information Statement introduced in Junos OS Release 10.0 for MX Series routers.
Statement introduced in Junos OS Release 12.2 for the QFX Series.

Description Specify the threshold for the adaptation of the detection time for a Bidirectional Forwarding Detection (BFD) session. When the detection time adapts to a value equal to or greater than the threshold, a single trap and a single system log message are sent.

NOTE: The threshold time must be greater than or equal to the minimum-interval or the minimum-receive-interval values.

Options *milliseconds*— Value for the detection time adaptation threshold.
Range: 1 through 255,000

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

transmit-interval (Liveness Detection)

Syntax

```
transmit-interval {
  minimum-interval milliseconds;
  threshold milliseconds;
}
```

Hierarchy Level

[edit protocols iccp peer liveness-detection]

Release Information

Statement introduced in Junos OS Release 10.0 for MX Series routers.
Statement introduced in Junos OS Release 12.2 for the QFX Series.
Statement introduced in Junos OS Release 12.3R2 for EX Series switches.

Description

Configure the Bidirectional Forwarding Detection (BFD) transmit interval. The negotiated transmit interval for a peer is the interval between the sending of BFD liveness detection requests to peers. The receive interval for a peer is the minimum interval between receiving packets sent from its peer; the receive interval is not negotiated between peers. To determine the transmit interval, each peer compares its configured minimum transmit interval with its peer's minimum receive interval. The larger of the two numbers is accepted as the transmit interval for that peer.

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.

version (Liveness Detection)

Syntax

```
version (1 | automatic);
```

Hierarchy Level

[edit protocols iccp peer liveness-detection]

Release Information

Statement introduced in Junos OS Release 10.0 for MX Series routers.
Statement introduced in Junos OS Release 12.2 for the QFX Series.

Description

Configure the Bidirectional Forwarding Detection (BFD) protocol version to detect.

Options

1—Use BFD protocol version 1.

automatic—Autodetect the BFD protocol version.

Required Privilege Level

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

CHAPTER 9

Operational Commands

- request interface mc-ae switchover (Multichassis Link Aggregation)
- request interface (revert | switchover) (Aggregated Ethernet Link Protection)
- request lacp link-switchover
- show iccp
- show interfaces mc-ae
- show l2-learning redundancy-groups
- show multi-chassis mc-lag configuration-consistency list-of-parameters
- show multi-chassis mc-lag configuration-consistency
- show multi-chassis mc-lag configuration-consistency global-config
- show multi-chassis mc-lag configuration-consistency icl-config
- show multi-chassis mc-lag configuration-consistency mcae-config
- show multi-chassis mc-lag configuration-consistency vlan-config
- show multi-chassis mc-lag configuration-consistency vrrp-config

request interface mc-ae switchover (Multichassis Link Aggregation)

Syntax

```
request interface mc-ae switchover
<immediate> mcae-id mcae-id;
mcae-id mcae-id;
```

Release Information Command introduced in Junos OS Release 13.3.

Description Manually revert egress traffic from the active node to the designated preferred node of a multichassis aggregated Ethernet interface. You can use this command to manually switch over traffic to the preferred node when the **switchover-mode** statement for the multichassis aggregated Ethernet interface is configured as **non-revertive** at the **[edit interfaces ae*X* mc-ae]** hierarchy level.

NOTE: To run this command successfully, the **status-control** statement should be configured as **active** at the **[edit interfaces ae*X* mc-ae]** hierarchy level.

Options

immediate—(Optional) Trigger immediate switchover to the preferred node. If this option is not configured, Junos OS waits for the timer configured using the **revert-time** statement at the **[edit interfaces ae*X* mc-ae]** hierarchy level to expire before it triggers the switchover.

mcae-id *mcae-id*—Triggers switchover for the specified mc-ae interface.

Required Privilege Level view

Related Documentation

- Configuring Multichassis Link Aggregation on MX Series Routers on page 45
- Configuring Manual and Automatic Link Switchover for MC-LAG Interfaces on MX Series Routers on page 152

List of Sample Output

request interface mc-ae switchover immediate mcae-id on page 452
request interface mc-ae switchover mcae-id on page 453

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

Sample Output

request interface mc-ae switchover immediate mcae-id

```
user@host >request interface mc-ae switchover immediate mcae-id 2
MCAE: Switchover Done
```

Sample Output

request interface mc-ae switchover mcae-id

```
user@host >request interface mc-ae switchover mcae-id 2
Switchover In Progress: Please check after 1 minutes,

Use “show interfaces mc-ae revertive-info” to check for the status
```

request interface (revert | switchover) (Aggregated Ethernet Link Protection)

Syntax request interface (revert | switchover) ae*x*

Release Information Command introduced in Junos OS Release 8.3.

Description Manually revert egress traffic from the designated backup link to the designated primary link of an aggregated Ethernet interface for which link protection is enabled, or manually switch egress traffic from the primary link to the backup link. This traffic includes transit traffic and local traffic originated on the router itself.

NOTE: When link protection is enabled on an aggregated Ethernet interface, if the primary link fails, the router automatically routes egress traffic to the backup link. However, the router does not automatically route egress traffic back to the primary link when the primary link is subsequently reestablished. Instead, you manually control when to have traffic diverted back to the primary link by issuing the **request interface (revert | switchover)** (Aggregated Ethernet Link Protection) operational command and specifying the **revert** keyword.

On M Series and T Series routers, use the **request interface (revert | switchover)** (Adaptive Services) operational command to manually revert to the primary adaptive services interface or link services interface, or to switch from the primary to the secondary interface. For information about this command, see *request interface (revert | switchover) (Adaptive Services)*.

Options **revert**—Restores egress traffic processing to the primary link.

switchover—Transfers egress traffic processing to the secondary (backup) link.

ae*x*—Aggregated Ethernet logical interface number: 0 through 15.

Required Privilege Level view

List of Sample Output request interface revert on page 454

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

Sample Output

request interface revert

```
user@host >request interface revert ae1
```


request lacp link-switchover

Syntax request lacp link-switchover ae*x*

Release Information Command introduced in Junos OS Release 9.3.

Description Manually switch aggregated Ethernet active or standby LACP links.

NOTE: Because this command overrides LACP priority calculations, we strongly recommend that you use this command only when the actor (in this case, the Juniper Networks router) is controlling the active or standby link and the partner (peer) is following. This scenario occurs when you configure only the actor for link protection.

Options **ae*x***—Aggregated Ethernet logical interface number: 0 through 15.

Required Privilege Level view

List of Sample Output request lacp link-switchover aeX on page 455

Output Fields When you enter this command, you are provided feedback on the status of your request. To view the switchover, use the **show lacp interfaces** command.

Sample Output

request lacp link-switchover aeX

```
user@host >request lacp link-switchover ae0ae0: Request succeeded
```

show iccp

Syntax

show iccp <brief | detail>
logical-system [*system-name* | all]

Release Information

Command introduced in Junos OS Release 10.0 for the MX Series.
Support for logical systems added in Junos OS Release 14.1 for MX Series routers.
Command introduced in Junos OS Release 12.2 for the QFX Series.

Description

Display Inter-Chassis Control Protocol (ICCP) information about the multichassis link aggregation group (MC-LAG) peers, including the state of the TCP connection, Bidirectional Forwarding Detection (BFD) protocol, backup liveness peer status, and MCSNOOPD, LACPD, and ESWD applications.

Options

none—Display ICCP information about the MC-LAG peers, including the state of the TCP connection and BFD protocol, and MCSNOOPD, LACP, and ESWD applications.

brief—Display brief ICCP information about the MC-LAG peers, including the state of the TCP connection and BFD protocol, and MCSNOOPD, LACP, and ESWD applications.

detail—Display detailed ICCP information about the MC-LAG peers, including the state of the TCP connection and BFD protocol, and MCSNOOPD, LACP, and ESWD applications.

logical-system [*system-name* | all]—(Optional) Display information for a specified logical system or all systems.

Required Privilege Level

view

Related Documentation

- iccp on page 433

List of Sample Output

show iccp (QFX Series) on page 457
show iccp (MX Series) on page 457

Output Fields

Table 19 on page 456 lists the output fields for the **show iccp** command. Output fields are listed in the approximate order in which they appear.

Table 19: show iccp Output Fields

Field Name	Field Description
Redundancy Group Information for peer	Aggregated Ethernet interface name.
TCP Connection	Specifies if the TCP connection between the peers hosting the MC-LAG is up or down.
Liveness Detection	Specifies if liveness detection, also known as Bidirectional Forwarding Detection (BFD), is up or down.

Table 19: show iccp Output Fields (continued)

Field Name	Field Description
Client Application	Specifies information regarding the state of the MCSNOOPD and client applications.

Sample Output

show iccp (QFX Series)

```
user@switch> show iccp
```

```
Redundancy Group Information for peer 10.3.3.2
  TCP Connection       : Established
  Liveliness Detection : Up

Client Application: MCSNOOPD

Client Application: eswd
```

show iccp (MX Series)

```
user@host> show iccp
```

```
Logical system :LS1
    Redundancy Group Information for peer 10.1.1.1
       TCP Connection       : Established
       Liveliness Detection : Up
       Redundancy Group ID          Status
         1                           Up
         2                           Up

Client Application: lacpd
Redundancy Group IDs Joined:  1
Redundancy Group IDs Joined:  2

Client Application: l2ald_iccpd_client
 Redundancy Group IDs Joined:  1
 Redundancy Group IDs Joined:  2
```

show interfaces mc-ae

Syntax **show interfaces mc-ae id** *identifier* **unit** *number*

Release Information Command introduced in Junos OS Release 9.6 for the MX Series.
Command introduced in Junos OS Release 12.2 for the QFX Series.
Statement introduced in Junos OS Release 12.3R2 for EX Series switches.
Configuration Consistency Check output field added in Junos OS Release 15.1X53-D60 for the QFX Series.

Description On peers with multichassis aggregated Ethernet (**mc-ae***X*) interfaces, use this command to display information about the multichassis aggregated Ethernet interfaces.

NOTE: In Junos OS Release 17.4R1, this command is not supported on EX4300, EX9200, PTX10000, QFX10002, and QFX10008 devices.

Options **id** *identifier*—(Optional) Specify the name of the multichassis aggregated Ethernet interface.

unit *number*—(Optional) Specify the logical interface by unit number.

Required Privilege Level view

List of Sample Output
show interfaces mc-ae (EX Series) on page 459
show interfaces mc-ae (MX Series) on page 459
show interfaces mc-ae (Active/Active Bridging and VRRP over IRB on MX Series) on page 460

Output Fields Table 20 on page 458 lists the output fields for the **show interfaces mc-ae** command. Output fields are listed in the approximate order in which they appear.

Table 20: show interfaces mc-ae Output Fields

Output Field Name	Field Description
Current State Machine's State	Specifies the state of the MC-LAG initialization state machine.
Configuration Consistency Check	Specifies the status of the MC-LAG configuration consistency check feature. The status is either **Passed** or **Failed**. If the status is **Failed**, the system will display the name of the parameter that failed consistency check. If there are multiple inconsistencies, only the first inconsistency is shown. If the enforcement level for the MC-LAG parameter was mandatory, and you did not configure that parameter correctly, the command will show that the MC-LAG interface is down.

Table 20: show interfaces mc-ae Output Fields (continued)

Output Field Name	Field Description
Member Link	Specifies the identifiers of the configured multichassis link aggregated interface members.
Local Status	Specifies the status of the local link: **active** or **standby**.
Peer Status	Specifies the status of the peer link: **active** or **standby**.
Peer State	Specifies the status of the local and peer links in an **active/active** MC-LAG configuration.
Logical Interface	Specifies the identifier and unit of the AE interface.
Topology Type	Specifies the bridge configured on the AE.
Local State	Specifies if the local device is up or down.
Peer State	Specifies if the peer device is up or down.
Peer Ip/MCP/State	Specifies the multichassis protection (MCP) link or the interchassis link-protection link (ICL-PL) for all of the multichassis aggregated Ethernet interfaces that are part of the peer.

Sample Output

show interfaces mc-ae (EX Series)

```
user@switch> show interfaces mc-ae ae1 512
Member Link                 : ae1
 Current State Machine's State: mcae active state
 Configuration Consistency Check    : Failed (redundancy group id mismatch)
 Local Status                : active
 Local State                 : up
 Peer Status                 : standby
 Peer State                  : up
     Logical Interface       : ae1.0
     Topology Type           : bridge
     Local State             : up
     Peer State              : up
     Peer Ip/MCP/State       : 10.1.1.1 ae0.0 up
```

show interfaces mc-ae (MX Series)

```
user@host> show interfaces mc-ae ae0 unit 512
 Member Links   : ae0
 Local Status   : active
 Peer Status    : active
 Logical Interface       : ae0.512
```

```
Core Facing Interface : Label Ethernet Interface
ICL-PL         : Label Ethernet Interface
```

show interfaces mc-ae (Active/Active Bridging and VRRP over IRB on MX Series)

```
user@host# show interfaces mc-ae ge-0/0/0.0
```

```
Member Link                  : ae0
Current State Machine's State: active
Local Status                 : active
Local State                  : up
Peer Status                  : active
Peer State                   : up
    Logical Interface        : ae0.0
    Topology Type            : bridge
    Local State              : up
    Peer State               : up
    Peer Ip/ICL-PL/State     : 192.168.100.10 ge-0/0/0.0 up
```

show l2-learning redundancy-groups

Syntax

```
show l2-learning redundancy-groups
logical-system [system-name | all]
<redundancy-group-id [0 to 4294967294]>
arp-statistics
nd-statistics
remote-macs
```

Release Information

Command introduced in Junos OS Release 13.2.
Support for logical systems added in Junos OS Release 14.1.
Command introduced in Junos OS Release 15.1R1 for EX Series switches

Description

(MX Series routers only) Display ARP statistics, Neighbor Discovery statistics, or remote MAC addresses for the Multi-Chassis Aggregated Ethernet (MC-AE) nodes for all or specified redundancy groups on a router or switch or logical systems on a router or switch. Note that the Redundancy Group ID is inherited by the bridging domain or VLAN from member AE interfaces.

Options

logical-system [*system-name* | all]—(Optional) Display information for a specified logical system or all systems.

redundancy-group-id—(Optional) The redundancy group identification number. The Inter-Chassis Control Protocol (ICCP) uses the redundancy group ID to associate the routing or switching devices contained in a redundancy group.

arp-statistics—(Optional) Count of ARP packets sent and received by the two MC-AE nodes.

nd-statistics—(Optional) Count of Neighbor Discovery packets sent and received by the two MC-AE nodes.

remote-macs —(Optional) List of remote MAC addresses in the “Installed” state, as learned from the remote MC-AE node.

Required Privilege Level

view

Related Documentation

- Configuring Multichassis Link Aggregation on MX Series Routers on page 45
- *show interfaces mc-ae*
- Configuring Active-Active Bridging and VRRP over IRB in Multichassis Link Aggregation on MX Series Routers and QFX Series Switches on page 141
- Configuring Multichassis Link Aggregation on EX Series Switches on page 51

List of Sample Output

show l2-learning redundancy-groups arp-statistics on page 463

show l2-learning redundancy-groups nd-statistics on page 463
show l2-learning redundancy-groups remote-macs on page 464
show l2-learning redundancy-groups logical-system arp-statistics (for Logical Systems) on page 464
show l2-learning redundancy-groups logical-system nd-statistics (for Logical Systems) on page 464
show l2-learning redundancy-groups group-id on page 464
show l2-learning redundancy-groups logical-system on page 465

Output Fields Output fields are listed in the approximate order in which they appear.

Table 21: show l2-learning redundancy-groups arp-statistics Output Fields

Field Name	Field Description
Redundancy Group ID	Redundancy Group to which the following details apply.
MCLAG ARP Statistics Group ID	ARP statistics for this Multichassis Link Aggregation Group (MC-LAG) instance.
ARP Rx Count From Line	Total number of ARPs received from the Line.
ARP Tx Count To Peer	Total number of ARPs sent to the peer.
ARP Rx Count From Peer	Total number of ARPs received from the peer.
ARP Drop Count received from line	Total number of ARPs sent by the peer that were received.
ARP Drop Count received from peer	Total number of ARPs sent by the peer that were dropped
Service-id	Service ID (configured at the routing instance level).

Table 22: show l2-learning redundancy-groups nd-statistics Output Fields

Field Name	Field Description
Redundancy Group ID	Redundancy Group to which the following details apply.
MCLAG ND Statistics Group ID	Neighbor Discovery statistics for this Multichassis Link Aggregation Group (MC-LAG) instance.
ND Rx Count From Line	Total number of Neighbor Discovery packets received from the Line.
ND Tx Count To Peer	Total number of Neighbor Discovery packets sent to the peer.
ND Rx Count From Peer	Total number of Neighbor Discovery packets received from the peer.

Table 22: show l2-learning redundancy-groups nd-statistics Output Fields (continued)

Field Name	Field Description
ND Drop Count received from line	Total number of Neighbor Discovery packets sent by the peer that were received.
ND Drop Count received from peer	Total number of Neighbor Discovery packets sent by the peer that were dropped
Service-id	Service ID (configured at the routing instance level).

Table 23: show l2-learning redundancy-groups remote-macs Output Fields

Field Name	Field Description
Redundancy Group ID	Redundancy Group to which the following details apply.
Peer-Addr	IP address of the remote peer.
VLAN	Virtual LAN identifier associated with the redundancy group.
MAC	Hardware media access control address associated with the redundancy group.
MCAE-ID	ID number of the MC-AE used by the redundancy group.
Flags	Connection state: local connect or Remote connect. If no flag is shown, the redundancy group may not be connected.
Status	Installation state: Installed or Not Installed.

Sample Output

show l2-learning redundancy-groups arp-statistics

```
user@host> show l2-learning redundancy-groups arp-statistics
Logical System : default
 Redundancy Group ID : 1     Flags : Local Connect, Remote Connect

MCLAG ARP Statistics
Group ID                                : 1
ARP Rx Count From Line                  : 52
ARP Tx Count To Peer                    : 15
ARP Rx Count From Peer                  : 39
ARP Install Count                       : 34
ARP Drop Count received from line       : 37
ARP Drop Count received from peer       : 5
```

show l2-learning redundancy-groups nd-statistics

```
user@host> show l2-learning redundancy-groups nd-statistics
```

```
Logical System : default
 Redundancy Group ID : 1     Flags : Local Connect, Remote Connect

MCLAG ND Statistics
Group ID                                : 1
ND Rx Count From Line                   : 52
ND Tx Count To Peer                     : 15
ND Rx Count From Peer                   : 39
ND Install Count                        : 34
ND Drop Count received from line        : 37
ND Drop Count received from peer        : 5
```

show l2-learning redundancy-groups remote-macs

```
user@host> show l2-learning redundancy-groups <redundancy-group-id> remote-macs
```

```
  Redundancy Group ID : 1     Flags : Local Connect, Remote Connect

Service-id  Peer-Addr    VLAN        MAC            MCAE-ID  Subunit  Opcode
Flags   Status
10          10.1.1.2      100   64:87:88:6a:df:f0      1        0       1
 0     Installed
```

show l2-learning redundancy-groups logical-system arp-statistics (for Logical Systems)

```
user@host> show l2-learning redundancy-groups logical-system LS1 arp-statistics
```

```
 Redundancy Group ID : 1     Flags : Local Connect, Remote Connect

MCLAG ARP Statistics
Group ID                                : 1
ARP Rx Count From Line                  : 52
ARP Tx Count To Peer                    : 15
ARP Rx Count From Peer                  : 39
ARP Install Count                       : 34
ARP Drop Count received from line       : 37
ARP Drop Count received from peer       : 5
```

show l2-learning redundancy-groups logical-system nd-statistics (for Logical Systems)

```
user@host> show l2-learning redundancy-groups logical-system LS1 nd-statistics
```

```
 Redundancy Group ID : 1     Flags : Local Connect, Remote Connect

MCLAG ND Statistics
Group ID                                : 1
ND Rx Count From Line                   : 52
ND Tx Count To Peer                     : 15
ND Rx Count From Peer                   : 39
ND Install Count                        : 34
ND Drop Count received from line        : 37
ND Drop Count received from peer        : 5
```

show l2-learning redundancy-groups group-id

```
user@host> show l2-learning redundancy-groups 1
```

```
Redundancy Group ID : 1       Flags : Local Connect,Remote Connect
```

show l2-learning redundancy-groups logical-system

```
user@host> show l2-learning redundancy-groups logical-system ls1
```

```
Redundancy Group ID : 2       Flags : Local Connect,Remote Connect
```

show multi-chassis mc-lag configuration-consistency list-of-parameters

Syntax show multichassis configuration-check list-of-parameters

Release Information Command introduced in Junos OS Release 16.1R1 for the EX Series.

Description Displays the list of MC-LAG parameters (referred to as configuration knobs in the CLI) that are checked for consistency across MC-LAG peers. There are certain parameters that must identical and others that must be unique on both peers. Enforcement of the consistency check for the parameters is either mandatory or desired. If the enforcement is mandatory, and you have not configured the parameter correctly, the multichassis aggregated Ethernet interface (MC-AE) interface will not come up. If the enforcement is desired, and you have not configured the parameter correctly, the MC-AE interface will come up, but performance might be sub-optimal. In this situation, the system will issue a syslog message.

The following list provides the hierarchies in which the MC-LAG parameters are configured:

- **ICL ifd**

  Specifies configuration parameters related to the interchassis control link

- **ICCP Peer**

  Specifies configuration parameters related to ICCP functionality

- **IRB Interface**

  Specifies configuration parameters related to the integrated routing and bridging interface

- **MCAE IFBD**

  Specifies configuration parameters related to the VLAN membership of a given MC-AE interface

- **MCAE ifd**

  Specifies configuration parameters related to a given MC-AE interface

- **MCAE ifl**

  Specifies configuration parameters related to a given MC-AE logical interface

- **VLAN**

  Specifies configuration parameters related to a given VLAN

- **VRRP Group**

  Specifies configuration parameters related to a VRRP session

Options There are no options for this command.

Required Privilege Level view

List of Sample Output show multi-chassis mc-lag configuration-consistency list-of-parameters on page 472

Output Fields Table 24 on page 467 lists the output fields for the **show multi-chassis mc-lag configuration-consistency list-of-parameters** command.

Table 24: show multi-chassis mc-lag configuration-consistency list-of-parameters Output Fields

Configuration Knob	Hierarchy
session-establishment-hold-time Specify the time during which an Inter-Chassis Control Protocol (ICCP) connection must be established between peers.	**Global**, **ICCP Peer**
mac-limit Specify the maximum number of MAC addresses to be associated with a VLAN—the default is unlimited, which can leave the network vulnerable to flooding.	**Global**
mac-aging-timer Specify how long MAC addresses remain in the Ethernet switching table.	**Global**
arp-aging-timer Specify the ARP aging timer in minutes for a logical interface of **inet**.	**Global**
rstp-system-identifier Specify different bridge identifiers for different RSTP routing instances.	**Global**
mstp-system-identifier Specify different bridge identifiers for different MSTP routing instances.	**Global**
rstp-bridge-priority Determine which bridge is elected as the root bridge for RSTP. If two bridges have the same path cost to the root bridge, the bridge priority determines which bridge becomes the designated bridge for a LAN segment.	**Global**
mstp-bridge-priority Determine which bridge is elected as the root bridge for MSTP. If two bridges have the same path cost to the root bridge, the bridge priority determines which bridge becomes the designated bridge for a LAN segment.	**Global**

Table 24: show multi-chassis mc-lag configuration-consistency list-of-parameters Output Fields (contin

Configuration Knob	Hierarchy
rstp-bpdu-block-on-edge Configure bridge protocol data unit (BPDU) protection on all edge ports of a switch for RSTP.	**Global**
vstp-bpdu-block-on-edge Configure bridge protocol data unit (BPDU) protection on all edge ports of a switch for VSTP.	**Global**
mstp-bpdu-block-on-edge Configure bridge protocol data unit (BPDU) protection on all edge ports of a switch for MSTP.	**Global**
service-id Specify a service identifier for each multichassis aggregated Ethernet interface that belongs to a link aggregation group (LAG).	**Global**
bfd minimum-interval Configure the minimum interval after which the local routing device transmits hello packets and then expects to receive a reply from a neighbor with which it has established a BFD session.	**ICCP Peer**
iccp/minimum-transmit-interval Specify the minimum interval at which the local routing device transmits hello packets to a neighbor with which it has established a BFD session.	**ICCP Peer**
iccp/minimum-receive-interval Specify the minimum interval after which the local routing device must receive a reply from a neighbor with which it has established a BFD session.	**ICCP Peer**
iccp/bfd multiplier Configure the number of hello packets not received by a neighbor that causes the originating interface to be declared down.	**ICCP Peer**
iccp single-hop Configure single hop BFD sessions.	**ICCP Peer**
iccp/authentication-key Specify the authentication key password to verify the authenticity of packets sent from the peers hosting an MC-LAG.	**ICCP Peer**

Table 24: show multi-chassis mc-lag configuration-consistency list-of-parameters Output Fields (contin

Configuration Knob	Hierarchy
redundancy-group-id-list Specify the redundancy group identification number. The Inter-Chassis Control Protocol (ICCP) uses the redundancy group ID to associate multiple chassis that perform similar redundancy functions.	**ICCP Peer**
backup-liveness-detection Determine whether a peer is up or down by exchanging keepalive messages over the management link between the two Inter-Chassis Control Protocol (ICCP) peers.	**ICCP Peer**
mc-ae-id Specify the identification number of the MC-LAG device.	**MCAE ifd**
mcae redundancy-group Used by ICCP to associate multiple chassis that perform similar redundancy functions and to establish a communication channel so that applications on peering chassis can send messages to each other.	**MCAE ifd**
mcae chassis-id Used by LACP for calculating the port number of the MC-LAG's physical member links.	**MCAE ifd**
mcae deployment mode Indicates whether an MC-LAG is in active-standby mode or active-active mode.	**MCAE ifd**
mcae status-control Specify whether the chassis becomes active or remains in standby mode when an interchassis link failure occurs.	**MCAE ifd**
force-icl-down Specify that if the node's ICCP peer goes down, the system brings down the interchassis-link logical interface.	**MCAE ifd**
prefer-status-control-active Specify that the node configured as **status-control active** becomes the active node if the peer of this node goes down.	**MCAE ifd**
lacp mode Specify LACP is active or passive.	**MCAE ifd**

Table 24: show multi-chassis mc-lag configuration-consistency list-of-parameters Output Fields (contin

Configuration Knob	Hierarchy
lacp periodic Specify the interval for periodic transmission of LACP packets.	**MCAE ifd**
lacp system-id Define the LACP system identifier at the aggregated Ethernet interface level.	**MCAE ifd**
lacp admin-key Specify an administrative key for the router or switch.	**MCAE ifd**
native-vlan-id Configure mixed tagging support for untagged packets on a port.	**MCAE ifd**
mcae-mac-synchronize Synchronize the MAC addresses for the Layer 3 interfaces of the switches participating in the MC-LAG.	**VLAN**
Interface mac Limit 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.	**VLAN**
l3-interface Associate a Layer 3 interface with the VLAN.	**VLAN**
igmp-snooping Enable IGMP snooping. A Layer 2 device monitors the IGMP join and leave messages sent from each connected host to a multicast router. This enables the Layer 2 device to keep track of the multicast groups and associated member ports. The Layer 2 device uses this information to make intelligent decisions and to forward multicast traffic to only the intended destination hosts.	**VLAN**
family Specify the protocol family configured on the logical interface.	**IRB Interface**
ipv4 address Specify an IPv4 address for the IRB interface.	**IRB Interface**

Table 24: show multi-chassis mc-lag configuration-consistency list-of-parameters Output Fields (contin

Configuration Knob	Hierarchy
ipv6 address Specify an IPv6 address for the IRB interface.	**IRB Interface**
vrrp-group id Specify a VRRP group identifier.	**IRB Interface**
proxy-arp-type For Ethernet interfaces only, configure the router or switch to respond to any ARP request, as long as the router or switch has an active route to the ARP request's target address.	**IRB Interface**
vrrp-group priority Configure a Virtual Router Redundancy Protocol (VRRP) router's priority for becoming the master default router. The router with the highest priority within the group becomes the master.	**VRRP Group**
vrrp-group authentication-key Configure a Virtual Router Redundancy Protocol (VRRP) IPv4 authentication key. You also must specify a VRRP authentication scheme by including the authentication-type statement.	**VRRP Group**
vrrp-group authentication-type Enable Virtual Router Redundancy Protocol (VRRP) IPv4 authentication and specify the authentication scheme for the VRRP group.	**VRRP Group**
vrrp-group virtual-address Configure the addresses of the virtual routers in a Virtual Router Redundancy Protocol (VRRP) IPv4 or IPv6 group.	**VRRP Group**
encapsulation Configure a logical link-layer encapsulation type.	**MCAE ifd**
flexible-vlan-tagging Support simultaneous transmission of 802.1Q VLAN single-tag and dual-tag frames on logical interfaces on the same Ethernet port, and on pseudowire logical interfaces.	**MCAE ifd**

Table 24: show multi-chassis mc-lag configuration-consistency list-of-parameters Output Fields (contin

Configuration Knob	Hierarchy
vlan-tagging For Fast Ethernet and Gigabit Ethernet interfaces, aggregated Ethernet interfaces configured for VPLS, and pseudowire subscriber interfaces, enable the reception and transmission of 802.1Q VLAN-tagged frames on the interface.	**MCAE ifd**
mtu Specify the maximum transmission unit (MTU) size for the media or protocol.	**MCAE ifd, ICL ifd**
interface-mode Determine whether the logical interface accepts or discards packets based on VLAN tags.	**MCAE ifl**
vlan membership Specify the name of the VLAN that belongs to an interface.	**MCAE ifl**

Sample Output

show multi-chassis mc-lag configuration-consistency list-of-parameters

```
user@host> show multi-chassis mc-lag configuration-consistency list-of-parameters
Possible completions:
  configuration-consistency  Show all configuration consistency information
regress@liki-pe2_re0> show multi-chassis mc-lag configuration-consistency
list-of-parameters
#Item   Configuration Knob                  Hierarchy         Consistency
Requirement Enforcement
------  -----------------                  ---------
----------------------  -----------
0       local-ip-addr                      Global               Unique
       Mandatory
1       session-establishment-hold-time    Global               Identical
       Mandatory
2       local-ip-addr                      ICCP Peer            Unique
       Mandatory
3       session-establishment-hold-time    ICCP Peer            Identical
       Mandatory
5       bfd minimum-interval               ICCP Peer            Identical
       Mandatory
6       iccp/minimum-transmit-interval     ICCP Peer            Identical
       Mandatory
7       iccp/minimum-receive-interval      ICCP Peer            Identical
       Mandatory
8       iccp/bfd multiplier                ICCP Peer            Identical
       Mandatory
9       iccp single-hop                    ICCP Peer            Identical
       Mandatory
11      iccp/authentication-key            ICCP Peer            Identical
```

```
      Mandatory
4      redundancy-group-id-list            ICCP Peer            Identical
      Mandatory
12     backup-liveness-detection           ICCP Peer            Unique
      Mandatory
13     service-id                          Global               Identical
      Mandatory
14     mac-limit                           Global               Identical
      Desired
15     mac-ageing-timer                    Global               Identical
      Desired
16     arp-ageing-timer                    Global               Identical
      Desired
17     rstp-bpdu-block-on-edge             Global               Identical
      Desired
18     rstp-bridge-priority                Global               Identical
      Desired
19     rstp-system-identifier              Global               Identical
      Desired
20     vstp-bpdu-block-on-edge             Global               Identical
      Desired
21     mstp-bpdu-block-on-edge             Global               Identical
      Desired
22     mstp-bridge-priority                Global               Identical
      Desired
23     mstp-system-identifier              Global               Identical
      Desired
24     mc-ae-id                            MCAE ifd             Identical
      Mandatory
25     mcae redundancy-group               MCAE ifd             Identical
      Mandatory
26     mcae chassis-id                     MCAE ifd             Unique
      Mandatory
27     mcae deployment mode                MCAE ifd             Identical
      Mandatory
28     mcae status-control                 MCAE ifd             Unique
      Mandatory
29     force-icl-down                      MCAE ifd             Unique
      Mandatory
30     prefer-status-control-active        MCAE ifd             Unique
      Desired
31     lacp mode                           MCAE ifd             Identical
      Mandatory
32     lacp periodic                       MCAE ifd             Identical
      Mandatory
33     lacp system-id                      MCAE ifd             Identical
      Mandatory
34     lacp admin-key                      MCAE ifd             Identical
      Mandatory
59     vlan id list                        VLAN                 Identical
      Mandatory
60     vlan-ids                            VLAN                 Identical
      Mandatory
62     Interface mac Limit                 VLAN                 Identical
      Desired
58     service-id                          VLAN                 Identical
      Mandatory
64     igmp-snooping-enabled               VLAN                 Identical
      Mandatory
61     mcae-mac-synchronize                VLAN                 Identical
```

```
      Mandatory
45     l3-interface                        VLAN                Identical
      Desired
47     ipv4 address                        IRB Interface       Unique
      Mandatory
48     ipv6 address                        IRB Interface       Unique
      Mandatory
49     vrrp-group id                       IRB Interface       Identical
      Mandatory
53     vrrp-group priority                 VRRP Group          Unique
      Mandatory
51     vrrp-group authentication-key       VRRP Group          Identical
      Mandatory
52     vrrp-group authentication-type      VRRP Group          Identical
      Mandatory
50     vrrp-group virtual-address          VRRP Group          Identical
      Mandatory
54     proxy-arp-type                      IRB Interface       Identical
      Mandatory
35     encapsulation                       MCAE ifd            Identical
      Mandatory
36     flexible-vlan-tagging               MCAE ifd            Identical
      Mandatory
37     vlan-tagging                        MCAE ifd            Identical
      Mandatory
38     mtu                                 MCAE ifd            Identical
      Mandatory
39     native-vlan-id                      MCAE ifd            Identical
      Mandatory
40     family                              MCAE ifl            Identical
      Mandatory
42     interface-mode                      MCAE ifl            Identical
      Mandatory
43     vlans                               MCAE ifl            Identical
      Mandatory
44     vlan membership                     MCAE ifl            Identical
      Mandatory
65     ICL interface                       ICL ifd             Identical
      Mandatory
65     ICL interface                       ICL ifd             Identical
      Mandatory
65     ICL interface                       ICL ifd             Identical
      Mandatory
68     encapsulation                       ICL ifd             Identical
      Mandatory
69     flexible-vlan-tagging               ICL ifd             Identical
      Mandatory
71     vlan-tagging                        ICL ifd             Identical
      Mandatory
70     mtu                                 ICL ifd             Identical
      Mandatory
72     native-vlan-id                      ICL ifd             Identical
      Mandatory
73     family                              ICL ifl             Identical
      Mandatory
75     interface-mode                      ICL ifl             Identical
      Mandatory
76     vlans                               ICL ifl             Identical
      Mandatory
```

```
77      vlan membership                     ICL ifl              Identical
       Mandatory
```

show multi-chassis mc-lag configuration-consistency

Syntax show multi-chassis mc-lag configuration-consistency (brief | detail)

Release Information Command introduced in Junos OS Release 16.1R1 for the EX Series.

Description Displays configuration consistency check status for various MC-LAG parameters .

NOTE: This command only displays MC-LAG parameters that are committed.

Options **none**—Display configuration consistency check status for various MC-LAG parameters.

Required Privilege Level view

List of Sample Output show multi-chassis mc-lag configuration-consistency on page 477

Output Fields Table 25 on page 476 lists the output fields for the **show multi-chassis mc-lag configuration-consistency** command. Output fields are listed in the approximate order in which they appear.

Table 25: show multi-chassis mc-lag configuration-consistency Output Fields

Output Field Name	Field Description
Configuration Item	Name of the committed MC-LAG parameter.
Local Physical Interface	Name of the physical interface configured on the local MC-LAG peer.
Peer Physical Interface	Name of the physical interface configured on the remote MC-LAG peer.
Local Logical Interface	Name of the logical interface configured on the local MC-LAG peer.
Peer Logical Interface	Name of the logical interface configured on the remote MC-LAG peer.
Local IRB	Name of the integrated routing and bridging (IRB) interface configured on the local MC-LAG peer.
Peer IRB	Name of the integrated routing and bridging (IRB) interface configured on the remote MC-LAG peer.
Local VLAN	Name of the VLAN configured on the local MC-LAG peer.

Table 25: show multi-chassis mc-lag configuration-consistency Output Fields (continued)

Output Field Name	Field Description
Peer VLAN	Name of the VLAN configured on the remote MC-LAG peer.
Enforcement Level	Enforcement level for the MC-LAG parameter is **Mandatory** or **Desirable**.
Local Value	Value of the committed MC-LAG parameter for the local MC-LAG peer.
Peer Value	Value of the committed MC-LAG parameter for the remote MC-LAG peer.
Result	Result of the configuration consistency check of the MC-LAG parameter is **PASS** or **FAIL**.

Sample Output

show multi-chassis mc-lag configuration-consistency

```
user@host> show multi-chassis mc-lag configuration-consistency
Configuration Item                  Enforcement Level   Local Value
Peer Value            Result
------------------                  -----------------   -----------
----------            -------
service-id                          Mandatory           1                    1
                     PASS
session-establishment-hold-time     Mandatory           300
300                   PASS
local-ip-addr                       Mandatory           10.1.1.1
10.1.1.2              PASS
iccp/bfd multiplier                 Mandatory           3                    3
                     PASS
iccp/minimum-transmit-interval      Mandatory           60
60                    PASS
bfd minimum-interval                Mandatory           1000
1000                  PASS

Local Physical Interface:ae0
Peer Physical Interface :ae0
Configuration Item                  Enforcement Level   Local Value
Peer Value            Result
------------------                  -----------------   -----------
----------            -------
lacp admin-key                      Mandatory           21
21                    PASS
lacp system-id                      Mandatory           00:00:00:00:00:01
00:00:00:00:00:01     PASS
lacp mode                           Mandatory           0                    0
                     PASS
prefer-status-control-active        Desirable           TRUE
TRUE                  FAIL
mcae status-control                 Mandatory           standby
active                PASS
```

```
mcae deployment mode                Mandatory           active-active
active-active         PASS
mcae chassis-id                     Mandatory           0                   1
                     PASS
mcae redundancy-group               Mandatory           101
101                   PASS
force-icl-down                      Mandatory           --
TRUE                  PASS

Local Logical Interface:ae0.0
Peer Logical Interface :ae0.0
Configuration Item                  Enforcement Level   Local Value
Peer Value            Result
------------------                  -----------------   -----------
----------            -------
vlan membership                     Mandatory           501-502
501-502               PASS
interface-mode                      Mandatory           trunk
trunk                 PASS

Local Physical Interface:ae1
Peer Physical Interface :ae1
Configuration Item                  Enforcement Level   Local Value
Peer Value            Result
------------------                  -----------------   -----------
----------            -------
lacp admin-key                      Mandatory           22
22                    PASS
lacp system-id                      Mandatory           00:00:00:00:00:03
00:00:00:00:00:03     PASS
lacp mode                           Mandatory           0                   0
                     PASS
prefer-status-control-active        Desirable           TRUE
TRUE                  FAIL
mcae status-control                 Mandatory           standby
active                PASS
mcae deployment mode                Mandatory           active-active
active-active         PASS
mcae chassis-id                     Mandatory           0                   1
                     PASS
mcae redundancy-group               Mandatory           101
101                   PASS
force-icl-down                      Mandatory           --
TRUE                  PASS

Local Logical Interface:ae1.0
Peer Logical Interface :ae1.0
Configuration Item                  Enforcement Level   Local Value
Peer Value            Result
------------------                  -----------------   -----------
----------            -------
vlan membership                     Mandatory           601-602
601-602               PASS
interface-mode                      Mandatory           trunk
trunk                 PASS

Local Physical Interface:ae5
Peer Physical Interface :ae5

Local Logical Interface:ae5.0
```

```
Peer Logical Interface :ae5.0
Configuration Item                  Enforcement Level   Local Value
Peer Value            Result
------------------                  -----------------   -----------
----------            -------
vlan membership                     Mandatory           501-502,601-602
501-502,601-602       PASS
interface-mode                      Mandatory           trunk
trunk                 PASS

Local VLAN:client-vlan-1
Peer VLAN :client-vlan-1
Configuration Item                  Enforcement Level   Local Value
Peer Value            Result
------------------                  -----------------   -----------
----------            -------
mcae-mac-synchronize                Mandatory           TRUE
TRUE                  PASS

Local IRB:irb.501
Peer IRB :irb.501
Configuration Item                  Enforcement Level   Local Value
Peer Value            Result
------------------                  -----------------   -----------
----------            -------
IPv4 Addresses                      Mandatory           10.1.1.1/24
10.1.1.2/24           FAIL

Local VLAN:client-vlan-2
Peer VLAN :client-vlan-2
Configuration Item                  Enforcement Level   Local Value
Peer Value            Result
------------------                  -----------------   -----------
----------            -------
mcae-mac-synchronize                Mandatory           TRUE
TRUE                  PASS

Local IRB:irb.502
Peer IRB :irb.502
Configuration Item                  Enforcement Level   Local Value
Peer Value            Result
------------------                  -----------------   -----------
----------            -------
IPv4 Addresses                      Mandatory           10.1.1.1/24
10.1.1.2/24           FAIL

Local VLAN:server-vlan-1
Peer VLAN :server-vlan-1
Configuration Item                  Enforcement Level   Local Value
Peer Value            Result
------------------                  -----------------   -----------
----------            -------
mcae-mac-synchronize                Mandatory           TRUE
TRUE                  PASS

Local IRB:irb.601
Peer IRB :irb.601
Configuration Item                  Enforcement Level   Local Value
Peer Value            Result
------------------                  -----------------   -----------
```

```
----------            -------
IPv4 Addresses                      Mandatory           10.2.1.1/24
10.2.1.2/24           FAIL

Local VLAN:server-vlan-2
Peer VLAN :server-vlan-2
Configuration Item                  Enforcement Level   Local Value
Peer Value            Result
------------------                  -----------------   -----------
----------            -------
mcae-mac-synchronize                Mandatory           TRUE
TRUE                  PASS

Local IRB:irb.602
Peer IRB :irb.602
Configuration Item                  Enforcement Level   Local Value
Peer Value            Result
------------------                  -----------------   -----------
----------            -------
IPv4 Addresses                      Mandatory           10.3.1.1/24
10.3.1.2/24           FAIL
```

show multi-chassis mc-lag configuration-consistency global-config

Syntax show multi-chassis mc-lag configuration-consistency global-config (brief | detail)

Release Information Command introduced in Junos OS Release 15.1X53-D60 for the QFX Series.
Command introduced in Junos OS Release 16.1R1 for the EX Series.

Description View configuration consistency check status for all committed global configuration related to MC-LAG functionality, the consistency requirement (identical or unique), the enforcement level (mandatory or desired), and the result of the configuration consistency check. The results are either pass or fail.

This command shows only a subset of what is shown in the **show multi-chassis mc-lag configuration-consistency** command. The following parameters related to the global configuration are checked for consistency.

- ICL interface
- RSTP bridge priority
- service ID
- session establishment hold time
- local IP address of the ICCP interface
- backup liveness detection peer IP address
- ICCP/BFD multiplier

NOTE: **This command only displays MC-LAG parameters that are committed.**

Options **none**—Display configuration consistency check status for all global configuration related to MC-LAG functionality.

Required Privilege Level view

List of Sample Output show multi-chassis mc-lag configuration-consistency global-config on page 482

Output Fields Table 26 on page 482 lists the output fields for the **show multi-chassis mc-lag configuration-consistency global-config** command. Output fields are listed in the approximate order in which they appear.

Table 26: show multi-chassis mc-lag configuration-consistency global-config Output Fields

Output Field Name	Field Description
Configuration Item	Name of the committed MC-LAG parameter.
Enforcement Level	Enforcement level for the MC-LAG parameter is **Mandatory** or **Desirable**.
Local Value	Value of the committed MC-LAG parameter on the local peer.
Peer Value	Value of the committed MC-LAG parameter on the remote peer.
Result	Result of the configuration consistency check of the MC-LAG parameter is **PASS** or **FAIL**.

Sample Output

show multi-chassis mc-lag configuration-consistency global-config

```
user@host> show multi-chassis mc-lag configuration-consistency global-config
Configuration Item                  Enforcement Level   Local Value
Peer Value            Result
------------------                  -----------------   -----------
----------            -------
service-id                          Mandatory           1                    1
                     PASS
session-establishment-hold-time     Mandatory           300
300                   PASS
local-ip-addr                       Mandatory           10.1.1.1
10.1.1.2              PASS
iccp/bfd multiplier                 Mandatory           3                    3
                     PASS
iccp/minimum-transmit-interval      Mandatory           60
60                    PASS
bfd minimum-interval                Mandatory           1000
1000                  PASS
```

show multi-chassis mc-lag configuration-consistency icl-config

Syntax show multi-chassis mc-lag configuration-consistency icl-config (brief | detail)

Release Information Command introduced in Junos OS Release 15.1X53-D60 for the QFX Series.
Command introduced in Junos OS Release 16.1R1 for the EX Series.

Description View configuration consistency check status for parameters related to the ICL, the consistency requirement (identical or unique), the enforcement level (mandatory or desired), and the result of the configuration consistency check. The results are either pass or fail. Some example of parameters related to the ICL interface are the interface mode and which VLAN the interface belongs to.

This command shows only a subset of what is shown in the **show multi-chassis mc-lag configuration-consistency** command. The following parameters related to the ICL configuration are checked for consistency check:

- VLAN membership
- interface mode

NOTE: This command only displays MC-LAG parameters that are committed.

Options **none**—Display configuration consistency check status for MC-LAG parameters related to the interchassis control link.

Required Privilege Level view

List of Sample Output show multi-chassis mc-lag configuration-consistency icl-config on page 484

Output Fields Table 27 on page 483 lists the output fields for the **show multi-chassis mc-lag configuration-consistency icl-config** command. Output fields are listed in the approximate order in which they appear.

Table 27: show multi-chassis mc-lag configuration-consistency icl-config Output Fields

Output Field Name	Field Description
Configuration Item	Name of the committed MC-LAG parameter.
Local Physical Interface	Name of the physical interface configured on the local MC-LAG peer.
Peer Physical Interface	Name of the physical interface configured on the remote MC-LAG peer.

Table 27: show multi-chassis mc-lag configuration-consistency icl-config Output Fields (continued)

Output Field Name	Field Description
Local Logical Interface	Name of the logical interface configured on the local MC-LAG peer.
Peer Logical Interface	Name of the logical interface configured on the remote MC-LAG peer.
Enforcement Level	Enforcement level for the MC-LAG parameter is **Mandatory** or **Desirable**.
Local Value	Value of the committed MC-LAG parameter on the local peer.
Peer Value	Value of the committed MC-LAG parameter on the remote peer.
Result	Result of the configuration consistency check of the MC-LAG parameter is **PASS** or **FAIL**.

Sample Output

show multi-chassis mc-lag configuration-consistency icl-config

```
user@host> show multi-chassis mc-lag configuration-consistency icl-config

Local Physical Interface:ae5
Peer Physical Interface :ae5

Local Logical Interface:ae5.0
Peer Logical Interface :ae5.0
Configuration Item                  Enforcement Level   Local Value
Peer Value            Result
------------------                  -----------------   -----------
----------            -------
vlan membership                     Mandatory           501-502,601-602
501-502,601-602       PASS
interface-mode                      Mandatory           trunk
trunk                 PASS
```

show multi-chassis mc-lag configuration-consistency mcae-config

Syntax show multi-chassis mc-lag configuration-consistency mcae-config (brief | detail)

Release Information Command introduced in Junos OS Release 15.1X53-D60 for the QFX Series.
Command introduced in Junos OS Release 16.1R1 for the EX Series.

Description View configuration consistency check status for committed parameters related to the multichassis aggregated Ethernet interfaces, the consistency requirement (identical or unique), the enforcement level (mandatory or desired), and the result of the configuration consistency check. The results are either pass or fail.

This command shows only a subset of what is shown in the **show multi-chassis mc-lag configuration-consistency** command. The following parameters related to the MC-AE interfaces are checked for consistency:

- LACP administrative key
- LACP system ID
- LACP periodic interval
- prefer status control setting
- status control setting
- mode
- chassis ID
- redundancy group ID
- VLAN membership of the ICL
- interface mode of the ICL

NOTE: **This command only displays MC-LAG parameters that are committed.**

Options **none**—Display configuration consistency check status for MC-LAG parameters related to the multichassis aggregated Ethernet interface.

Required Privilege Level view

List of Sample Output show multi-chassis mc-lag configuration-consistency mcae-config on page 486

Output Fields Table 28 on page 486 lists the output fields for the **show multi-chassis mc-lag configuration-consistency mcae-config** command. Output fields are listed in the approximate order in which they appear.

Table 28: show multi-chassis mc-lag configuration-consistency mcae-config Output Fields

Output Field Name	Field Description
Configuration Item	Name of the committed MC-LAG parameter.
Local Physical Interface	Name of the physical interface configured on the local MC-LAG peer.
Peer Physical Interface	Name of the physical interface configured on the remote MC-LAG peer.
Local Logical Interface	Name of the logical interface configured on the local MC-LAG peer.
Peer Logical Interface	Name of the logical interface configured on the remote MC-LAG peer.
Enforcement Level	Enforcement level for the MC-LAG parameter is **Mandatory** or **Desirable**.
Local Value	Value of the committed MC-LAG parameter on the local peer.
Peer Value	Value of the committed MC-LAG parameter on the remote peer.
Result	Result of the configuration consistency check of the MC-LAG parameter is **PASS** or **FAIL**.

Sample Output

show multi-chassis mc-lag configuration-consistency mcae-config

```
user@host> show multi-chassis mc-lag configuration-consistency mcae-config
Local Physical Interface:ae0
Peer Physical Interface :ae0
Configuration Item                  Enforcement Level   Local Value
Peer Value            Result
------------------                  -----------------   -----------
----------            -------
lacp admin-key                      Mandatory           21
21                    PASS
lacp system-id                      Mandatory           00:00:00:00:00:01
00:00:00:00:00:01     PASS
lacp mode                           Mandatory           0                    0
                     PASS
prefer-status-control-active        Desirable           TRUE
TRUE                  FAIL
mcae status-control                 Mandatory           standby
active                PASS
mcae deployment mode                Mandatory           active-active
active-active         PASS
mcae chassis-id                     Mandatory           0                    1
                     PASS
mcae redundancy-group               Mandatory           101
```

```
101                   PASS
force-icl-down                        Mandatory           --
TRUE                  PASS

Local Logical Interface:ae0.0
Peer Logical Interface :ae0.0
Configuration Item                    Enforcement Level   Local Value
Peer Value            Result
------------------                    -----------------   -----------
----------            -------
vlan membership                       Mandatory           501-502
501-502               PASS
interface-mode                        Mandatory           trunk
trunk                 PASS

Local Physical Interface:ae1
Peer Physical Interface :ae1
Configuration Item                    Enforcement Level   Local Value
Peer Value            Result
------------------                    -----------------   -----------
----------            -------
lacp admin-key                        Mandatory           22
22                    PASS
lacp system-id                        Mandatory           00:00:00:00:00:03
00:00:00:00:00:03     PASS
lacp mode                             Mandatory           0                     0
                     PASS
prefer-status-control-active          Desirable           TRUE
TRUE                  FAIL
mcae status-control                   Mandatory           standby
active                PASS
mcae deployment mode                  Mandatory           active-active
active-active         PASS
mcae chassis-id                       Mandatory           0                     1
                     PASS
mcae redundancy-group                 Mandatory           101
101                   PASS
force-icl-down                        Mandatory           --
TRUE                  PASS

Local Logical Interface:ae1.0
Peer Logical Interface :ae1.0
Configuration Item                    Enforcement Level   Local Value
Peer Value            Result
------------------                    -----------------   -----------
----------            -------
vlan membership                       Mandatory           601-602
601-602               PASS
interface-mode                        Mandatory           trunk
trunk                 PASS
```

show multi-chassis mc-lag configuration-consistency vlan-config

Syntax show multi-chassis mc-lag configuration-consistency vlan-config (brief | detail)

Release Information Command introduced in Junos OS Release 15.1X53-D60 for the QFX Series.
Command introduced in Junos OS Release 16.1R1 for the EX Series.

Description View configuration consistency check status for committed parameters related to MC-LAG VLAN configuration, the consistency requirement (identical or unique), the enforcement level (mandatory or desired), and the result of the configuration consistency check. The results are either pass or fail.

This command shows only a subset of what is shown in the **show multi-chassis mc-lag configuration-consistency** command. The following parameters related to the VLAN and IRB configuration are checked for consistency:

- VRRP group ID
- IP address of IRB interface

NOTE: This command only displays MC-LAG parameters that are committed.

Options **none**—Display configuration consistency check status for MC-LAG parameters related to VLAN configuration.

Required Privilege Level view

List of Sample Output show multi-chassis mc-lag configuration-consistency vlan-config on page 489

Output Fields Table 29 on page 488 lists the output fields for the **show multi-chassis mc-lag configuration-consistency vlan-config** command. Output fields are listed in the approximate order in which they appear.

Table 29: show multi-chassis mc-lag configuration-consistency vlan-config Output Fields

Output Field Name	Field Description
Configuration Item	Name of the committed MC-LAG parameter.
Local Physical Interface	Name of the physical interface configured on the local MC-LAG peer.
Peer Physical Interface	Name of the physical interface configured on the remote MC-LAG peer.

Table 29: show multi-chassis mc-lag configuration-consistency vlan-config Output Fields (continued)

Output Field Name	Field Description
Local Logical Interface	Name of the logical interface configured on the local MC-LAG peer.
Peer Logical Interface	Name of the logical interface configured on the remote MC-LAG peer.
Local IRB	Name of the integrated routing and bridging (IRB) interface configured on the local MC-LAG peer.
Peer IRB	Name of the integrated routing and bridging (IRB) interface configured on the remote MC-LAG peer.
Local VLAN	Name of the VLAN configured on the local MC-LAG peer.
Peer VLAN	Name of the VLAN configured on the remote MC-LAG peer.
Enforcement Level	Enforcement level for the MC-LAG parameter is **Mandatory** or **Desirable**.
Local Value	Value of the committed MC-LAG parameter on the local peer.
Peer Value	Value of the committed MC-LAG parameter on the remote peer.
Result	Result of the configuration consistency check of the MC-LAG parameter is **PASS** or **FAIL**.

Sample Output

show multi-chassis mc-lag configuration-consistency vlan-config

```
user@host> show multi-chassis mc-lag configuration-consistency vlan-config
Local VLAN:client-vlan-1
Peer VLAN :client-vlan-1

Local IRB:irb.501
Peer IRB :irb.501
Configuration Item                  Enforcement Level   Local Value
Peer Value           Result
------------------                  -----------------   -----------
----------           -------
vrrp-group id                       Mandatory           11
11                    PASS
IPv4 Addresses                      Mandatory           10.1.1.2/8
10.1.1.1/8           PASS

Local VLAN:client-vlan-2
Peer VLAN :client-vlan-2
```

```
Local IRB:irb.502
Peer IRB :irb.502
Configuration Item                  Enforcement Level   Local Value
Peer Value           Result
------------------                  -----------------   -----------
----------           -------
vrrp-group id                       Mandatory           12
12                    PASS
IPv4 Addresses                      Mandatory           10.0.1.2/8
10.0.1.1/8           PASS
```

show multi-chassis mc-lag configuration-consistency vrrp-config

Syntax

```
show multi-chassis mc-lag configuration-consistency vrrp-config
```

Release Information

Command introduced in Junos OS Release 15.1X53-D60 for the QFX Series.
Command introduced in Junos OS Release 16.1R1 for the EX Series.

Description

View configuration consistency check status for committed parameters related to VRRP configuration, the consistency requirement (identical or unique), the enforcement level (mandatory or desired), and the result of the configuration consistency check. The results are either pass or fail.

This command shows only a subset of what is shown in the **show multi-chassis mc-lag configuration-consistency** command. The following parameters related to the VRRP configuration are checked for consistency: VRRP group virtual IP address and VRRP group priority value.

NOTE: This command only displays MC-LAG parameters that are committed.

Options

none—Displays configuration consistency check status for MC-LAG parameters related to Virtual Router Redundancy Protocol (VRRP) configuration.

Required Privilege Level

view

List of Sample Output

show multi-chassis mc-lag configuration-consistency vrrp-config on page 492

Output Fields

Table 30 on page 491 lists the output fields for the **show multi-chassis mc-lag configuration-consistency vrrp-config** command. Output fields are listed in the approximate order in which they appear.

Table 30: show multi-chassis mc-lag configuration-consistency vrrp-config Output Fields

Output Field Name	Field Description
Configuration Item	Name of the committed MC-LAG parameter.
Local Physical Interface	Name of the physical interface configured on the local MC-LAG peer.
Peer Physical Interface	Name of the physical interface configured on the remote MC-LAG peer.
Local Logical Interface	Name of the logical interface configured on the local MC-LAG peer.

Table 30: show multi-chassis mc-lag configuration-consistency vrrp-config Output Fields (continued)

Output Field Name	Field Description
Peer Logical Interface	Name of the logical interface configured on the remote MC-LAG peer.
Local IRB	Name of the integrated routing and bridging (IRB) interface configured on the local MC-LAG peer.
Peer IRB	Name of the integrated routing and bridging (IRB) interface configured on the remote MC-LAG peer.
Local VLAN	Name of the VLAN configured on the local MC-LAG peer.
Peer VLAN	Name of the VLAN configured on the remote MC-LAG peer.
Enforcement Level	Enforcement level for the MC-LAG parameter is **Mandatory** or **Desirable**.
Local Value	Value of the committed MC-LAG parameter on the local peer.
Peer Value	Value of the committed MC-LAG parameter on the remote peer.
Result	Result of the configuration consistency check of the MC-LAG parameter is **PASS** or **FAIL**.

Sample Output

show multi-chassis mc-lag configuration-consistency vrrp-config

```
user@host> show multi-chassis mc-lag configuration-consistency vrrp-config
Lcal VRRP Group:11

Peer VRRP Group :11
Configuration Item                  Enforcement Level   Local Value
Peer Value           Result
------------------                  -----------------   -----------
----------           -------
vrrp-group virtual-address          Mandatory           010.001.001.010
010.001.001.010      PASS
vrrp-group priority                 Mandatory           202
201                  PASS

Local VRRP Group:12

Peer VRRP Group :12
Configuration Item                  Enforcement Level   Local Value
Peer Value           Result
------------------                  -----------------   -----------
----------           -------
vrrp-group virtual-address          Mandatory           011.001.001.010
```

```
011.001.001.010       PASS
vrrp-group priority                 Mandatory          202
201                   PASS
```