

Junos® OS

Multichassis Link Aggregation User Guide for Routing and Switching Devices





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

Juniper Networks, the Juniper Networks logo, Juniper, and Junos are registered trademarks of Juniper Networks, Inc. in the United States and other countries. All other trademarks, service marks, registered marks, or registered service marks are the property of their respective owners.

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

Junos® OS Multichassis Link Aggregation User Guide for Routing and Switching Devices Copyright © 2023 Juniper Networks, Inc. All rights reserved.

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

YEAR 2000 NOTICE

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

END USER LICENSE AGREEMENT

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

Table of Contents

About This Guide | ix

1

Understanding MC-LAG

Understanding Multichassis Link Aggregation Groups | 2

MC-LAG Concepts | 4

Advanced MC-LAG Concepts | 8

Understanding Configuration Synchronization | 8

Understanding Multichassis Link Aggregation Group Configuration Consistency Check | 14

Unknown Unicast and IGMP Snooping | 27

Layer 3 Unicast Feature Support | 28

MAC Address Management | 28

MAC Aging | 29

Address Resolution Protocol Active-Active MC-LAG Support Methodology | 29

DHCP Relay with Option 82 | 29

Layer 2 Unicast Features Supported | 31

Protocol Independent Multicast | 31

IGMP Report Synchronization | 33

IGMP Snooping in MC-LAG Active-Active Mode | 33

Understanding MC-LAGs on an FCoE Transit Switch | 40

Understanding EVPN-MPLS Interworking with Junos Fusion Enterprise and MC-LAG | 44

Understanding the Incremented Values of Statistical Counters for Loop-Free MC-LAG Networks | **50**

Enhanced Convergence | 54

IPv6 Neighbor Discovery Protocol | 55

Anycast Gateways | 56

2 Implementing MC-LAG

Getting Started with MC-LAG | 63

```
Configuring Multichassis Link Aggregation on MX Series Routers | 63

Configuring Multichassis Link Aggregation on EX Series Switches | 70

Configuring ICCP for MC-LAG | 76
```

MC-LAG Examples | 78

Example: Configuring Multichassis Link Aggregation on the QFX Series | 78

```
Requirements | 79

Overview | 79

Configuration | 80

Verification | 99

Troubleshooting | 103
```

Example: Configuring Multichassis Link Aggregation on the MX Series | 103

```
Requirements | 104

Overview | 104

Configuring the PE Routers | 105

Configuring the CE Device | 115

Configuring the Provider Router | 119

Verification | 122
```

Example: Configuring Multichassis Link Aggregation Between QFX Series Switches and MX Series Routers | 123

```
Requirements | 123

Overview | 124

Configure the Devices | 126

Verification | 144
```

Example: Configuring Multichassis Link Aggregation on EX9200 Switches in the Core for Campus Networks | 148

```
Requirements | 149

Overview | 149

Configuration | 151

(Optional) Configuring RSTP | 177

(Optional) Configuring IGMP Snooping | 179
```

```
(Optional) Configuring VRRP | 181
       (Optional) Configuring MAC Address Synchronization | 184
       (Optional) Configuring OSPF | 186
       (Optional) Configuring PIM | 187
       (Optional) Configuring DHCP Relay | 190
       Verification | 193
    Example: Configure Optional Features For Multichassis Link Aggregation | 204
       (Optional) Configuring RSTP | 205
       (Optional) Configuring IGMP Snooping | 207
       (Optional) Configuring VRRP | 208
       (Optional) Configuring MAC Address Synchronization | 210
       (Optional) Configuring OSPF | 212
       (Optional) Configuring PIM | 213
       (Optional) Configuring DHCP Relay | 217
    Example: Simplifying Multichassis Link Aggregation on EX9200 Switches in the Core for
       Campus Networks | 220
       Requirements | 220
       Overview | 220
       Configuration | 223
       Verification | 254
    Example: Configuring CoS for FCoE Transit Switch Traffic Across an MC-LAG | 280
       Requirements | 281
       Overview | 282
       Configuration | 288
       Verification | 301
    Example: EVPN-MPLS Interworking With an MC-LAG Topology | 314
       Requirements | 314
       Overview and Topology | 315
       PE1 and PE2 Configuration | 317
       PE3 Configuration | 333
Other MC-LAG Configurations
```

Other MC-LAG Configurations

Other MC-LAG Configurations | 340

Configuring Active-Active Bridging and VRRP over IRB in Multichassis Link Aggregation on MX Series Routers | 340

Configuring MC-LAG | 341

Configuring the Interchassis Link-Protection Link | 342

Configuring Multiple Chassis | 342

Configuring the Service ID | 343

Configuring IGMP Snooping for Active-Active MC-LAG | 345

Configuring IGMP Snooping in MC-LAG Active-Active Mode | 346

Configuring Manual and Automatic Link Switchover for MC-LAG Interfaces on MX Series Routers | 347

Forcing MC-LAG Links or Interfaces with Limited LACP Capability to Be Up | 349

Increasing ARP and Network Discovery Protocol Entries for Enhanced MC-LAG and Layer 3 VXLAN Topologies | 350

Understanding the Need for an Increase in ARP and Network Discovery Protocol (NDP) Entries | 351

Increasing ARP and Network Discovery Protocol Entries for Enhanced MC-LAG Using IPv4
Transport | 351

Increasing ARP and Network Discovery Protocol Entries for Enhanced MC-LAG Using IPv6
Transport | 353

Increasing ARP for EVPN-VXLAN Gateway for Border-Leaf in Edge Routed Bridge (ERB) or Spine in Centrally Routed Bridge (CRB) for IPv4 Tenant Traffic | 355

Increasing ARP and Network Discovery Protocol Entries for EVPN-VXLAN gateway for Border-Leaf in Edge Routed Bridge (ERB) or Spine in Centrally Routed Bridge (CRB) for IPv6 Tenant Traffic | 357

Synchronizing and Committing Configurations | 360

Configure Devices for Configuration Synchronization | 360

Create a Global Configuration Group | 363

Create a Local Configuration Group | 366

Create a Remote Configuration Group | 369

Create Apply Groups for the Local, Remote, and Global Configurations | 371

Synchronizing and Committing Configurations | 372

Troubleshooting Remote Device Connections | 372

MC-LAG Upgrade Guidelines

MC-LAG Upgrade Guidelines | 376

Best Practices and Usage Notes | 379

Troubleshooting Multichassis Link Aggregation

Troubleshooting Multichassis Link Aggregation | 392

MAC Addresses Learned on Multichassis Aggregated Ethernet Interfaces Are Not Removed from the MAC Address Table | 393

MC-LAG Peer Does Not Go into Standby Mode | 394

Secondary MC-LAG Peer with Status Control Set to Standby Becomes Inactive | 395

Redirect Filters Take Priority over User-Defined Filters | 395

Operational Command Output Is Wrong | 396

ICCP Connection Might Take Up to 60 Seconds to Become Active | 397

MAC Address Age Learned on a Multichassis Aggregated Ethernet Interface Is Reset to Zero | 398

MAC Address Is Not Learned Remotely in a Default VLAN | 399

Snooping Entries Learned on Multichassis Aggregated Ethernet Interfaces Are Not Removed | 399

ICCP Does Not Come Up After You Add or Delete an Authentication Key | 400

Local Status Is Standby When It Should Be Active | 401

Packets Loop on the Server When ICCP Fails | 401

Both MC-LAG Peers Use the Default System ID After a Reboot or an ICCP Configuration Change | 402

No Commit Checks Are Done for ICL-PL Interfaces | 403

Double Failover Scenario | 403

Multicast Traffic Floods the VLAN When the ICL-PL Interface Goes Down and Up | 404

Layer 3 Traffic Sent to the Standby MC-LAG Peer Is Not Redirected to Active MC-LAG Peer | 405

Aggregated Ethernet Interfaces Go Down | 405

Flooding of Upstream Traffic | 406

ARP and MAC Table Entries Become Out of Sync in an MC-LAG Configuration | 407

Configuring Interface Diagnostics Tools to Test the Physical Layer Connections | 408

Configuring Loopback Testing | 408

Configuring BERT Testing | 411

Starting and Stopping a BERT Test | 415

Configuration Statements and Operational Commands

Junos CLI Reference Overview | 418

About This Guide

Use this guide to configure and monitor multichassis link aggregation groups (MC-LAGs).



Understanding MC-LAG

Understanding Multichassis Link Aggregation Groups | 2

MC-LAG Concepts | 4

Advanced MC-LAG Concepts | 8

Understanding Multichassis Link Aggregation Groups

IN THIS SECTION

Benefits of MC-LAGs | 3

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.

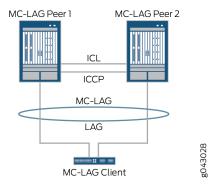
On one end of an MC-LAG, there is an MC-LAG client device, such as a server, that has one or more physical links in a link aggregation group (LAG). This client device uses the link as a LAG. On the other side of the MC-LAG, there can be a maximum of two MC-LAG peers. Each of the MC-LAG peers has one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

- Reduces operational expenses by providing active-active links within a Link Aggregation Group (LAG).
- Provides faster layer 2 convergence upon link and device failures.
- Adds node-level redundancy to the normal link-level redundancy that a LAG provides.
- Improves network resiliency, which reduces network down time as well as expenses.

MC-LAG Concepts

IN THIS SECTION

- ICCP and ICL | 4
- Multichassis Link Protection | 4
- Service ID | 5
- Failure Handling | 5
- Load Balancing | 6
- MC-LAG Packet Forwarding | 6
- Virtual Router Redundancy Protocol (VRRP) over IRB and MAC Address Synchronization | 6

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.

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.

Service ID

You must configure the same service ID on each MC-LAG peer when the MC-LAG logical interfaces are part of a bridge domain. The service ID, which is configured under the switch-options hierarchy, is used to synchronize applications such as IGMP, ARP, and MAC learning across MC-LAG members. If you are configuring virtual switch instances, configure a different service ID for each virtual switch instance.

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 the LACP state changes on both active and standby peers.

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.

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.

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.

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 have Junos OS Release 15.2R1 or earlier installed in your device, 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.**

If you have already manually configured a static ARP or ND entry and upgrade to a later release, the static entry is deleted when ICCP goes down. If you configured ICCP on the IRB static entry, then ICCP

might not come up. As a workaround, you can disable the automatic creation of static ARP and ND entries by issuing the following command: **set protocols I2-learning no-mclag-ifa-sync**.

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
```

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.

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.

Advanced MC-LAG Concepts

IN THIS SECTION

- Understanding Configuration Synchronization | 8
- Understanding Multichassis Link Aggregation Group Configuration Consistency Check | 14
- Unknown Unicast and IGMP Snooping | 27
- Layer 3 Unicast Feature Support | 28
- MAC Address Management | 28
- MAC Aging | 29
- Address Resolution Protocol Active-Active MC-LAG Support Methodology | 29
- DHCP Relay with Option 82 | 29
- Layer 2 Unicast Features Supported | 31
- Protocol Independent Multicast | 31
- IGMP Report Synchronization | 33
- IGMP Snooping in MC-LAG Active-Active Mode | 33
- Understanding MC-LAGs on an FCoE Transit Switch | 40
- Understanding EVPN-MPLS Interworking with Junos Fusion Enterprise and MC-LAG | 44
- Understanding the Incremented Values of Statistical Counters for Loop-Free MC-LAG Networks | 50
- Enhanced Convergence | 54
- IPv6 Neighbor Discovery Protocol | 55
- Anycast Gateways | 56

Understanding Configuration Synchronization

IN THIS SECTION

- Benefits of Configuration Synchronization | 9
 - How Configuration Synchronization Works | 9

- How to Enable Configuration Synchronization | 9
- How Configuration Synchronization is Supported | 10
- Configuration Groups for Local, Remote and Global Configurations | 10
- Creating Conditional Groups for Certain Devices | 10
- Applying Configuration Groups | 11
- Device Configuration Details for Configuration Synchronization | 11
- How Configurations and Commits Are Synchronized Between Devices | 12

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.
- 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.

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

Applying Configuration Groups

To apply configuration groups, enable the apply-groups statement at the <code>[edit]</code> 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 <code>set apply-groups [GroupA GroupB GroupC]</code> 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.

Understanding Multichassis Link Aggregation Group Configuration Consistency Check

IN THIS SECTION

- Benefits of Using MC-LAG Consistency Check | 14
- How MC-LAG Consistency Checks Work | 14
- Configuration Consistency Requirements | 15
- When Remote Peers are Not Reachable | 16
- Enabling MC-LAG Configuration Consistency Checking | 16
- Learning the Status of a Configuration Consistency Check | 26
- Support for MC-LAG Configuration Consistency Checking | 27

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 enabled by default. Table 1 on page 16 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 1: 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
interface-mac-limit Specify the maximum number of MAC addresses to be associated with an interface.	Global	Identical	Desired
MC-AE interface mac limit Specify the maximum number of MAC addresses to be associated with the MC-AE interface.	MCAE ifl	Identical	Desired

Table 1: MC-LAG Parameters Checked for Configuration Consistency (Continued)

Configuration Knob	Hierarchy	Consistency Requirement	Enforcement
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
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

Table 1: MC-LAG Parameters Checked for Configuration Consistency (Continued)

Configuration Knob	Hierarchy	Consistency Requirement	Enforcement
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
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

Table 1: MC-LAG Parameters Checked for Configuration Consistency (Continued)

Configuration Knob	Hierarchy	Consistency Requirement	Enforcement
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

Table 1: MC-LAG Parameters Checked for Configuration Consistency (Continued)

Configuration Knob	Hierarchy	Consistency Requirement	Enforcement
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
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

Table 1: MC-LAG Parameters Checked for Configuration Consistency (Continued)

Configuration Knob	Hierarchy	Consistency Requirement	Enforcement
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
force-icl-down Specify that if the 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

Table 1: MC-LAG Parameters Checked for Configuration Consistency (Continued)

Configuration Knob	Hierarchy	Consistency Requirement	Enforcement
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 1: MC-LAG Parameters Checked for Configuration Consistency (Continued)

Configuration Knob	Hierarchy	Consistency Requirement	Enforcement
13-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

Table 1: MC-LAG Parameters Checked for Configuration Consistency (Continued)

Configuration Knob	Hierarchy	Consistency Requirement	Enforcement
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 target address.	IRB Interface	Identical	Mandatory
vrrp-group priority Configure a Virtual Router Redundancy Protocol (VRRP) router priority for becoming the primary default router. The router with the highest priority within the group becomes the primary.	VRRP Group	Unique	Mandatory
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

Table 1: MC-LAG Parameters Checked for Configuration Consistency (Continued)

Configuration Knob	Hierarchy	Consistency Requirement	Enforcement
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 1: 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.

SEE ALSO

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

Unknown Unicast and IGMP Snooping

- During an MC-LAG peer reboot, known multicast traffic is flooded until the IGMP snooping state is synchronized with the peer.
- 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.

Layer 3 Unicast Feature Support

Layer 3 unicast feature support includes the following:

- VRRP active-standby support enables Layer 3 routing over MC-AE interfaces.
- Routed VLAN interface (RVI) or IRB MAC address synchronization enables MC-LAG peers to forward Layer 3 packets arriving on MC-AE interfaces with either its own RVI or IRB MAC address or its peers RVI or IRB MAC address.
- 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.

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.

DHCP Relay with 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:

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.

- 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.

Layer 2 Unicast Features Supported

NOTE: MAC learning is disabled on the ICL automatically. 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.

How layer 2 unicast learning and aging works:

- 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.

Protocol Independent Multicast

IN THIS SECTION

- PIM Operation with Normal Mode Designated Router Election | 32
- PIM Operation with Dual Designated Router Mode | 32
- Failure Handling | 33

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.

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.

IGMP Snooping in MC-LAG Active-Active Mode

IN THIS SECTION

- IGMP Snooping in MC-LAG Active-Active Mode Functionality | 34
- Typically Supported Network Topology for IGMP Snooping with MC-LAG Active-Active Bridging | 35
- Control Plane State Updates Triggered by Packets Received on Remote Chassis | 36
- Data Forwarding | 36
- Pure Layer 2 Topology Without Integrated Routing and Bridging | 37
- Qualified Learning | 37
- Data Forwarding with Qualified Learning | 38
- Static Groups on Single-Homed Interfaces | 39

Router-Facing Interfaces as Multichassis Links | 39

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 2 on page 35 depicts a typical network topology over which IGMP snooping with MC-LAG active-active bridging is supported.

Multicast source Network cloud (IP core) 11 12 ICCP Primary Secondary active active Interchassis link **IRB** IRB ae0.0 ae0.1 ae0.1 ae0.0 Multichassis link 13 Host

Figure 2: 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 37.

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 aeO 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 3 on page 37 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.

Multicast source Network cloud (IP core) U1 U2 PIM-DR 110 12 11 18 17 ICCP Primary Secondary Interchassis link 1 active active Interchassis link 2 ae0.0 ae0.1 ae0.1 ae0.0 Multichassis link Host

Figure 3: 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 Id2.
- Interfaces I9 and I10 belong to vlan3, learning domain Id3.

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 Id2, P-A and S-A linked should appear to be one switch.

Since there are no multichassis links in learning domain Id3, for the IGMP speakers (hosts and routers) in learning domain Id3, P-A and S-A will not appear to be one switch.

This discussion assumes interchassis link ICL1 corresponds to learning domain Id1 and interchassis link ICL2 corresponds to learning domain Id2.

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 Id1), 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 Id1.

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 Id3 with interests in s,g, and that interface I10 in learning domain Id3 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 Id3.

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

Understanding MC-LAGs on an FCoE Transit Switch

IN THIS SECTION

- Supported MC-LAG Topology | 40
- FIP Snooping and FCoE Trusted Ports | 43
- CoS and Data Center Bridging (DCB) | 43

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 4 on page 41 shows an inverted-U topology using QFX3500 switches.

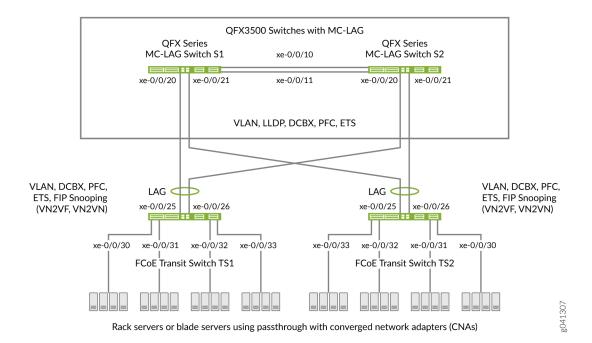


Figure 4: Supported Topology for an MC-LAG on an FCoE Transit Switch

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 4 on page 41, 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 4 on page 41, 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.

Understanding EVPN-MPLS Interworking with Junos Fusion Enterprise and MC-LAG

IN THIS SECTION

- Benefits of Using EVPN-MPLS with Junos Fusion Enterprise and MC-LAG | 47
- BUM Traffic Handling | 47
- Split Horizon | 48
- MAC Learning | 49
- Handling Down Link Between Cascade and Uplink Ports in Junos Fusion Enterprise | 49
- Layer 3 Gateway Support | 50

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 5 on page 45 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.

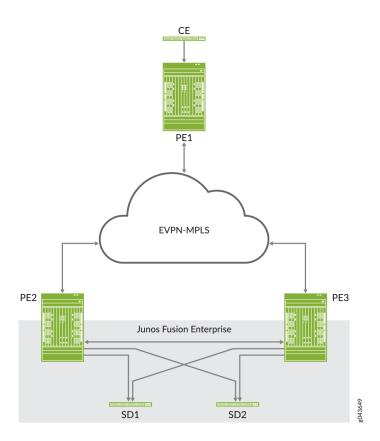
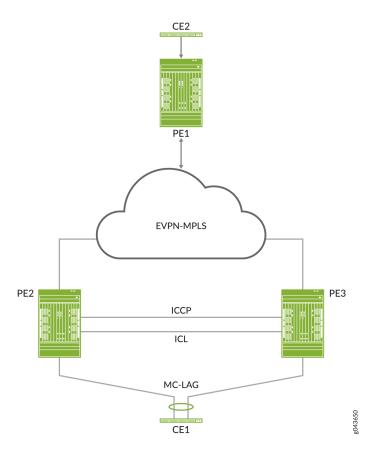


Figure 5: EVPN-MPLS Interworking with Junos Fusion Enterprise

Figure 6 on page 46 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 6: EVPN-MPLS Interworking with MC-LAG



Throughout this topic, Figure 5 on page 45 and Figure 6 on page 46 serve as references to illustrate various scenarios and points.

The use cases depicted in Figure 5 on page 45 and Figure 6 on page 46 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 5 on page 45 and Figure 6 on page 46, 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 2 on page 47 outlines the issues that arise, and how Juniper Networks resolves the issues in its implementation of the EVPN-MPLS interworking feature.

Table 2: 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 5 on page 45 and Figure 6 on page 46, 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 3 on page 48 explains the issue and how Juniper Networks resolves it in its implementation of the EVPN-MPLS interworking feature.

Table 3: 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 4 on page 49 describes the issue and how the EVPN-MPLS interworking implementation prevents the issue. The use cases shown in Figure 5 on page 45 and Figure 6 on page 46 illustrate the issue. In both use cases, PE1, PE2, and PE3 are EVPN peers, and PE2 and PE3 are MC-LAG peers.

Table 4: 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.
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 5 on page 45, 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.

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	7 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: It 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(92)	0	0	
GOT:	(0.)	v	Ů	
GOT:				
GOT: Firewall				
GOT:				
GOT: mac firewall	DISC(78)			
GOT: firewall discard	DISC(67)	0	0	
GOT: tram miss	DISC(07)	0	0	
GOT: firewall reject	PUNT(36)	155559	59137563	
GOT: firewall reject	PUNT(54)	0	0	
GOT. TITEWALL SCHA TO HOST	10111(34)	U	U	

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.

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.

Anycast Gateways

IN THIS SECTION

- Benefits of Anycast Gateways | 57
- Anycast Gateway Configuration Guidelines | 57
- Anycast Gateway Configuration Limitations | 59

In an EVPN-MPLS or MC-LAG environment with two Juniper Networks devices multihomed in all-active mode, you can configure IRB interfaces on the devices. With the IRB interfaces in place, the multihomed devices function as gateways that handle inter-subnet routing. To set up an IRB interface on a Juniper Networks device, you can configure the following:

- An IRB interface with:
 - An IPv4 or an IPv6 address

```
set interface irb unit logical-unit-number family inet ipv4-address set interface irb unit logical-unit-number family inet6 ipv6-address
```

• A media access control (MAC) address

```
set interface irb unit logical-unit-number mac mac-address
```

NOTE: In addition to explicitly configuring a MAC address using the above command syntax, you can use the MAC address that the Juniper Networks device automatically generates (chassis MAC).

- A virtual gateway address (VGA) with:
 - An IPv4 or an IPv6 address

set interfaces irb unit *logical-unit-number* family inet address *primary-ipv4-address*/8 virtual-gateway-address *gateway-ipv4-address*

set interfaces irb unit *logical-unit-number* family inet6 address *primary-ipv6-address*/104 virtual-gateway-address *gateway-ipv6-address*

A MAC address

set interface irb unit *logical-unit-number* virtual-gateway-v4-mac *mac-address* set interface irb unit *logical-unit-number* virtual-gateway-v6-mac *mac-address*

NOTE: In addition to explicitly configuring a MAC address using the above command syntax, you can use the MAC address that the Juniper Networks device automatically generates (chassis MAC).

When specifying an IP or MAC address for an IRB interface or VGA on the multihomed devices, you can now use an anycast address. This support of anycast addresses enables you to configure the same addresses for the IRB interface or VGA on each of the multihomed devices, thereby establishing the devices as anycast gateways.

Your IP address subnet scheme will determine whether you use the IRB interface command syntax or the VGA command syntax to set up your anycast gateway.

In an Ethernet VPN-Multiprotocol Label Switching (EVPN-MPLS) or multichassis link aggregation (MC-LAG) environment, you can configure two Juniper Networks devices multihomed in all-active mode as anycast gateways.

The following sections provide more information about anycast gateways.

Benefits of Anycast Gateways

With the two multihomed Juniper Networks devices acting as anycast gateways in an EVPN-MPLS
or MC-LAG network, a host in the same network that generates Layer 3 packets with destinations in
other networks can now send the packets to the local anycast gateway. Upon receipt of these Layer
3 packets, the anycast gateway routes the packets in the core network based on destination IP
lookup.

Anycast Gateway Configuration Guidelines

- In general, when configuring addresses for an anycast gateway:
 - For IPv4 or IPv6 addresses, you can specify any subnet.

- For MAC addresses, you can use the MAC address that the Juniper Networks device automatically generates (chassis MAC), or you can explicitly configure a MAC address using the CLI.
- Your IP address subnet scheme will determine whether you use the IRB interface command syntax or the VGA command syntax to set up your anycast gateway.

To set up your multihomed devices as anycast gateways, we provide the following configuration guidelines:

- Guideline 1—If the IP address for the anycast gateways is in the /30 or /31 (for IPv4) or /126 or /127 (for IPv6) subnet:
 - You must configure the same IP address for the IRB interface on each of the multihomed devices using one of the following commands.

```
set interface irb unit logical-unit-number family inet ipv4-address set interface irb unit logical-unit-number family inet6 ipv6-address
```

• You must explicitly configure the MAC address using the following command:

```
set interface irb unit logical-unit-number mac mac-address
```

- You must not configure a VGA (IP and MAC addresses).
- Guideline 2—If the IP address for the anycast gateways is a subnet other than /30, /31, /126, or /127 then a VGA can be configured:
 - You must configure the same IP address for the VGA on each of the multihomed devices using one of the following commands.

```
set interfaces irb unit logical-unit-number family inet address primary-ipv4-address/8 virtual-gateway-address gateway-ipv4-address
set interfaces irb unit logical-unit-number family inet6 address primary-ipv6-address/104 virtual-gateway-address gateway-ipv6-address
```

You must explicitly configure the MAC address using one of the following commands:

```
set interface irb unit logical-unit-number virtual-gateway-v4-mac set interface irb unit logical-unit-number virtual-gateway-v6-mac mac-address
```

 When specifying a MAC address for the VGA, we do not recommend using the same MAC address used for VRRP.

NOTE: You can also see No Link Title (No Link Title section) for similar guidelines to configure a leaf device as an anycast gateway in an EVPN-VXLAN edge-routed bridging (ERB) overlay fabric.

Anycast Gateway Configuration Limitations

When configuring the anycast gateway using guidelines described earlier in this topic, keep the following in mind:

- In general, we do not recommend reusing a VRRP MAC address as a MAC address for an IRB interface. However, if you must do so, as is the general practice when configuring VRRP on Juniper Networks devices, you must use a VRRP IPv4 MAC address for the IPv4 family and a VRRP IPv6 MAC address for the IPv6 family.
 - Given these parameters, the only configuration guideline with which this limitation will work is configuration guideline 2.
- When configuring anycast gateway addresses using guideline 1 in an EVPN-MPLS environment, you
 must also specify the default-gateway do-not-advertise configuration statements within a routing
 instance. For example:

set routing-instance routing-instance-name protocols evpn default-gateway do-not-advertise

- In an EVPN-MPLS environment, if your anycast gateway IP addresses are in different subnets and you specify the addresses within multiple routing instances:
 - If you configured an anycast gateway IP address using configuration guideline 1 in one routing instance, and another anycast gateway IP address using configuration guideline 2 in a different routing instance, you must also specify the default-gateway no-gateway-community configuration statements within the routing instance:

set routing-instance *routing-instance-name* protocols evpn default-gateway no-gateway-community

This additional configuration applies only to the routing instance that includes anycast gateway IP addresses configuring using guideline 1.

• For each routing instance in which you specified the anycast gateway IP address using configuration guideline 1, we recommend specifying a single non-VRRP MAC address.

- Automatic ESI generation is enabled by default on devices with EVPN multihoming for virtual gateway redundancy. We recommend that you disable the automatic ESI generation for EVPN-VXLAN networks with edge-routed bridging (ERB) overlays. In that case, you can include the no-auto-virtual-gateway-esi statement at the [edit interfaces irb unit logical-unit-number] hierarchy level.
 - Starting in Junos OS Release 22.1R1, MX960, MX2020, and MX10008 routers also enable automatic ESI generation by default for EVPN Layer 3 gateway IRB interface ESIs. However, the no-auto-virtual-gateway-esi statement is not supported with EVPN-MPLS networks. As a result, you will always see auto-generated ESIs for IRB interfaces in this case.
- In an EVPN-VXLAN environment with multihoming, you might use multiple EVPN routing instances on peer provider edge (PE) devices that share an Ethernet segment (ES). When you configure anycast gateways with the default-gateway statement, we don't support mixing the default behavior (advertise option) with the no-gateway-community option on the links that participate in the same ES.

As a result, if you configure the default-gateway statement with the no-gateway-community option in any EVPN routing instances on any peer PE device that share an ES, you must configure this statement:

- In all the routing instances that share the ES on a PE device,
- On all the peer PE devices that share the ES
- Only with either the no-gateway-community option or the do-not-advertise.

You can't omit setting the default-gateway statement or include the statement with the advertise option in any routing instance on any peer PE device.

We support setting an anycast gateway IP address on IRB interfaces on ACX5448 devices. However, for IRB interfaces with /30 or /31 IP addresses on connections between PE and customer edge (CE) device interfaces, the CE device doesn't have enough pool space for the BGP session IP address allocation. As a result, we don't support BGP with IRB interface /30 and /31 anycast IP addresses.

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.
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.

15.1X53-D60	On QFX Series switches, support for configuration synchronization started with Junos OS Release 15.1X53-D60.
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.
14.2R6	On MX Series routers and Junos Fusion, support for configuration synchronization started with Junos OS Release 14.2R6.
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.



Implementing MC-LAG

Getting Started with MC-LAG | 63

MC-LAG Examples | 78

Getting Started with MC-LAG

IN THIS SECTION

- Configuring Multichassis Link Aggregation on MX Series Routers | 63
- Configuring Multichassis Link Aggregation on EX Series Switches | 70
- Configuring ICCP for MC-LAG | 76

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 \$\mathcal{X}\$ aggregated-ether-options] hierarchy level along with one of the following statements at the [edit interfaces ae \$\mathcal{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 \$\mathcal{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 ae X 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 mc-ae 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)
```

```
[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
```

```
[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.

```
[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 minimuminterval 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 interfacename

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

```
[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
```

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)
```

[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.

[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

```
[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 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.

```
[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 milliseconds.

[edit protocols]

user@switch# set iccp peer peer-ip-address liveness-detection minimum-receive-interval
milliseconds

```
[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
```

```
[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 ICCP for MC-LAG

For multichassis link aggregation (MC-LAG), you must configure Inter-Control Center Communications Protocol (ICCP) to exchange information between two MC-LAG peers.

To enable ICCP, include the iccp statement at the [edit protocols] hierarchy level:

```
[edit protocols]
iccp {
    authentication-key string;
   local-ip-addr ipv4-address;
    peer ip-address{
        authentication-key string;
       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;
        redundancy-group-id-list [ redundancy-groups ];
```

```
session-establishment-hold-time value;
}
session-establishment-hold-time value;
traceoptions;
}
```

The local-ip-address statement sets the source address. This could be a specified address or interface address. The session-establishment-hold-time statement determines whether a chassis takes over as the primary at the ICCP session.

The authentication-key statement is provided by TCP Message Digest 5 (md5) option for an ICCP TCP session. The redundancy-group-id-list statement specifies the redundancy groups between ICCP peers and the liveness-detection hierarchy configures Bidirectional Forwarding Detection (BFD) protocol options.

NOTE: ICCP is based on TCP and it uses IP routes to reach the MC-LAG peer. To ensure that the ICCP session is as resilient as possible, we recommend that you configure alternative routes between the ICCP end-point IP addresses. Alternatively, configure a LAG interface that has two or more interfaces between the MC-LAG pairs to prevent session failure when there are no alternative routes.

For Inter-Control Center Communications Protocol (ICCP) in a multichassis link aggregation group (MC-LAG) configured in an active-active bridge domain, you must ensure that you configure the same peer IP address hosting the MC-LAG by including the peer *ip-address* statement at the [edit protocols iccp] hierarchy level and the multi-chassis-protection peer *ip-address* statement at the [edit interfaces interface-name] hierarchy level. Multichassis protection reduces the configuration at the logical interface level for MX Series routers with multichassis aggregated Ethernet (MC-AE) interfaces. If the ICCP is UP and the interchassis data link (ICL) comes UP, the router configured as standby will bring up the MC-AE interfaces shared with the peer active-active node specified by the peer statement.

For example, the following statements illustrate how the same peer IP address can be configured for both the ICCP peer and multichassis protection link:

```
set interfaces ae1 unit 0 multi-chassis-protection 10.255.34.112 interface ae0 set protocols iccp peer 10.255.34.112 redundancy-group-id-list 1
```

Although you can commit an MC-LAG configuration with various parameters defined for it, you can configure multichassis protection between two peers without configuring the ICCP peer address. You can also configure multiple ICCP peers and commit such a configuration.

MC-LAG Examples

IN THIS SECTION

- Example: Configuring Multichassis Link Aggregation on the QFX Series | 78
- Example: Configuring Multichassis Link Aggregation on the MX Series | 103
- Example: Configuring Multichassis Link Aggregation Between QFX Series Switches and MX Series
 Routers | 123
- Example: Configuring Multichassis Link Aggregation on EX9200 Switches in the Core for Campus Networks | 148
- Example: Configure Optional Features For Multichassis Link Aggregation | 204
- Example: Simplifying Multichassis Link Aggregation on EX9200 Switches in the Core for Campus
 Networks | 220
- Example: Configuring CoS for FCoE Transit Switch Traffic Across an MC-LAG | 280
- Example: EVPN-MPLS Interworking With an MC-LAG Topology | 314

Example: Configuring Multichassis Link Aggregation on the QFX Series

IN THIS SECTION

- Requirements | 79
- Overview | 79
- Configuration | 80
- Verification | 99
- Troubleshooting | 103

NOTE: Our content testing team has validated and updated this example.

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 13.2X51-D10 or later for the QFX5100 standalone switches, Release 15.1X53-D10 or later for QFX10002 standalone switches.
 - Revalidated on Junos OS Release 17.3R1 for QFX5100 and QFX10000 switches.
 - Revalidated on Junos OS Release 19.4R1 for 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.

Overview

IN THIS SECTION

Topology | 79

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 7 on page 80 shows the topology used in this example.

VLAN 10 = 10.10.1.0/24 VLAN 50 = 10.50.1.0/30 Service ID = 10 MGMT - 10.1.1.0/24 QFX10002 Backup Peer QFX10002 QFX1 QFX2 Chassis ID = 0 Chassis ID = 1 em0 = 10.1.1.2 irb.10 = 10.10.1.2 irb.50 = 10.50.1.2 em0 = 10.1.1.1 irb.10 = 10.10.1.1 ae0 irb.50 = 10.50.1.1 xe-0/0/1 ICL = VLAN 10 xe-0/0/1 xe-0/0/3 xe-0/0/3 xe-0/0/2 xe-0/0/2 ICCP = VLAN 50 xe-0/0/0 xe-0/0/0 Host1 Host2 10.10.1.101 10.10.1.102 ae1 VLAN 10 xe-0/0/0 xe-0/0/1 xe-0/0/2 QFX10002 QFX3 irb.10 = 10.10.1.3 Host3

10.10.1.103

Figure 7: Configuring a Multichassis LAG Between QFX1 and QFX2

Configuration

IN THIS SECTION

- CLI Quick Configuration | 80
- Configuring MC-LAG on Two Switches | 83
- Results | 91

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.

QFX1

```
set chassis aggregated-devices ethernet device-count 2
set interfaces xe-0/0/0 ether-options 802.3ad ae1
set interfaces xe-0/0/1 ether-options 802.3ad ae0
set interfaces xe-0/0/2 ether-options 802.3ad ae0
set interfaces xe-0/0/3 unit 0 family ethernet-switching interface-mode access
set interfaces xe-0/0/3 unit 0 family ethernet-switching vlan members v10
set interfaces ae0 aggregated-ether-options lacp active
set interfaces ae0 unit 0 family ethernet-switching interface-mode trunk
set interfaces ae0 unit 0 family ethernet-switching vlan members v50
set interfaces ae0 unit 0 family ethernet-switching vlan members v10
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 redundancy-group 1
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 v10
set interfaces em0 unit 0 family inet address 10.1.1.1/24
set interfaces irb unit 10 family inet address 10.10.1.1/24
set interfaces irb unit 50 family inet address 10.50.1.1/30
set multi-chassis multi-chassis-protection 10.50.1.2 interface ae0
set protocols iccp local-ip-addr 10.50.1.1
set protocols iccp peer 10.50.1.2 session-establishment-hold-time 340
set protocols iccp peer 10.50.1.2 redundancy-group-id-list 1
set protocols iccp peer 10.50.1.2 backup-liveness-detection backup-peer-ip 10.1.1.2
set protocols iccp peer 10.50.1.2 liveness-detection minimum-receive-interval 1000
set protocols iccp peer 10.50.1.2 liveness-detection transmit-interval minimum-interval 1000
set switch-options service-id 10
set vlans v10 vlan-id 10
set vlans v10 l3-interface irb.10
set vlans v50 vlan-id 50
set vlans v50 13-interface irb.50
```

QFX2

```
set chassis aggregated-devices ethernet device-count 2
set interfaces xe-0/0/0 ether-options 802.3ad ae1
set interfaces xe-0/0/1 ether-options 802.3ad ae0
set interfaces xe-0/0/2 ether-options 802.3ad ae0
set interfaces xe-0/0/3 unit 0 family ethernet-switching interface-mode access
set interfaces xe-0/0/3 unit 0 family ethernet-switching vlan members v10
set interfaces ae0 aggregated-ether-options lacp active
set interfaces ae0 unit 0 family ethernet-switching interface-mode trunk
set interfaces ae0 unit 0 family ethernet-switching vlan members v50
set interfaces ae0 unit 0 family ethernet-switching vlan members v10
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 redundancy-group 1
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 v10
set interfaces em0 unit 0 family inet address 10.1.1.2/24
set interfaces irb unit 10 family inet address 10.10.1.2/24
set interfaces irb unit 50 family inet address 10.50.1.2/30
set multi-chassis multi-chassis-protection 10.50.1.1 interface ae0
set protocols iccp local-ip-addr 10.50.1.2
set protocols iccp peer 10.50.1.1 session-establishment-hold-time 340
set protocols iccp peer 10.50.1.1 redundancy-group-id-list 1
set protocols iccp peer 10.50.1.1 backup-liveness-detection backup-peer-ip 10.1.1.1
set protocols iccp peer 10.50.1.1 liveness-detection minimum-receive-interval 1000
set protocols iccp peer 10.50.1.1 liveness-detection transmit-interval minimum-interval 1000
set switch-options service-id 10
set vlans v10 vlan-id 10
set vlans v10 l3-interface irb.10
set vlans v50 vlan-id 50
set vlans v50 13-interface irb.50
```

QFX3

```
set chassis aggregated-devices ethernet device-count 2
set interfaces xe-0/0/0 ether-options 802.3ad ae1
set interfaces xe-0/0/1 ether-options 802.3ad ae1
set interfaces xe-0/0/2 unit 0 family ethernet-switching interface-mode access
set interfaces xe-0/0/2 unit 0 family ethernet-switching vlan members v10
set interfaces ae1 aggregated-ether-options lacp active
set interfaces ae1 unit 0 family ethernet-switching interface-mode trunk
set interfaces ae1 unit 0 family ethernet-switching vlan members v10
set interfaces em0 unit 0 family inet address 10.1.1.3/24
set interfaces irb unit 10 family inet address 10.10.1.3/24
set vlans v10 vlan-id 10
set vlans v10 l3-interface irb.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 interfaces and multichassis protection link between MC-LAG peers:

1. Configure the number of LAGs on both QFX1 and QFX2.

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

2. Add member interfaces to the aggregated Ethernet interfaces on both QFX1 and QFX2.

```
QFX1 and QFX2:
[edit interfaces]
user@switch# set xe-0/0/0 ether-options 802.3ad ae1
[edit interfaces]
user@switch# set xe-0/0/1 ether-options 802.3ad ae0
[edit interfaces]
user@switch# set xe-0/0/2 ether-options 802.3ad ae0
```

3. Configure an access interface to the connected end host.

```
[edit interfaces]
user@switch# set xe-0/0/3 unit 0 family ethernet-switching interface-mode access
```

4. Add member interfaces to VLAN v10.

```
[edit interfaces]
user@switch# set interfaces xe-0/0/3 unit 0 family ethernet-switching vlan members v10
```

5. Configure a trunk interface between QFX1 and QFX2.

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

6. Enable VLANs on the MC-LAG between QFX1 and QFX2.

```
[edit]
user@switch# set vlans v10 vlan-id 10
```

```
[edit]
user@switch# set vlans v50 vlan-id 50
```

```
[edit interfaces]
user@switch# set ae0 unit 0 family ethernet-switching vlan members v10
```

```
[edit interfaces]
user@switch# set ae0 unit 0 family ethernet-switching vlan members v50
```

7. Configure an IRB 50.

```
[edit irb]
user@switch# set irb.50
```

8. Assign VLAN 50 to irb.50.

```
[edit]
user@switch# set vlans v50 l3-interface irb.50
```

9. Configure an IRB 10.

```
[edit irb]
user@switch# set irb.10
```

10. Assign VLAN 10 irb.10.

```
[edit]
user@switch# set vlans v10 l3-interface irb.10
```

11. Enable LACP on the MC-LAG interface on QFX1 and QFX2.

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

```
[edit interfaces]
user@switch# set ae0 aggregated-ether-options lacp active

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

12. Specify the same LACP system ID for the MC-LAG on QFX1 and QFX2.

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

13. Specify the same LACP administration key on both QFX1 and QFX2.

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

14. Specify the same multichassis aggregated Ethernet identification number on both MC-LAG peers on QFX1 and QFX2.

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

15. Specify a unique chassis ID for the MC-LAG on the MC-LAG peers on QFX1 and QFX2.

```
QFX1:
  [edit interfaces]
  user@switch# set ae1 aggregated-ether-options mc-ae chassis-id 0
```

```
QFX2:
[edit interfaces]
user@switch# set ae1 aggregated-ether-options mc-ae chassis-id 1
```

16. Specify the operating mode of the MC-LAG on both QFX1 and QFX2.

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

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

17. Specify the status control for MC-LAG on QFX1 and QFX2.

NOTE: You must configure status control on both QFX1 and QFX2 hosting the MC-LAG. If one peer is in active mode, the other must be in standby mode.

QFX1:

[edit interfaces]

user@switch# set ae1 aggregated-ether-options mc-ae status-control active

QFX2:

[edit interfaces]

user@switch# set ae1 aggregated-ether-options mc-ae status-control standby

18. Specify the number of seconds by which the bring-up of the multichassis aggregated Ethernet interface should be deferred after you reboot QFX1 and QFX2.

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

19. Configure Layer 3 connectivity between the MC-LAG peers on both QFX1 and QFX2.

[edit vlans]
user@switch# set v50 vlan-id 50

[edit vlans]
user@switch# set v50 l3-interface irb.50

[edit interfaces]

user@switch# set ae0 unit 0 family ethernet-switching interface-mode trunk vlan members v50

20. Configure a multichassis protection link between QFX1 and QFX2.

QFX1: [edit] user@switch# set multi-chassis multi-chassis-protection 10.50.1.2 interface ae0

```
QFX2:
[edit]
user@switch# set multi-chassis multi-chassis-protection 10.50.1.1 interface ae0
```

21. Configure the local IP address to be in the ICCP connection on QFX1 and QFX2.

```
QFX1:
[edit protocols]
user@switch# set iccp local-ip-addr 10.50.1.1
```

```
QFX2:
[edit protocols]
user@switch# set iccp local-ip-addr 10.50.1.2
```

22. (Optional) Configure the time during which an ICCP connection must succeed between MC-LAG peers on QFX1 and QFX2.

NOTE: On QFX 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.

```
QFX1:
[edit protocols]
user@switch# set iccp peer 10.50.1.2 session-establishment-hold-time 340
```

```
QFX2:
[edit protocols]
user@switch# set iccp peer 10.50.1.1 session-establishment-hold-time 340
```

23. Configure the redundancy groups for ICCP on QFX1 and QFX2.

```
QFX1:
[edit protocols]
user@switch# set iccp peer 10.50.1.2 redundancy-group-id-list 1
```

```
QFX2:
[edit protocols]
user@switch# set iccp peer 10.50.1.1 redundancy-group-id-list 1
```

24. (Optional) Configure the backup IP address to be used for backup liveness detection on both QFX1 and QFX2.

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.

QFX1:

[edit protocols]

user@switch# set iccp peer 10.50.1.2 backup-liveness-detection backup-peer-ip 10.1.1.2

QFX2:

[edit protocols]

user@switch# set iccp peer 10.50.1.1 backup-liveness-detection backup-peer-ip 10.1.1.1

25. Configure the peer IP address and minimum receive interval for a BFD session for ICCP on QFX1 and QFX2.

QFX1:

[edit protocols]

user@switch# set iccp peer 10.50.1.2 liveness-detection minimum-receive-interval 1000

QFX2:

[edit protocols]

user@switch# set iccp peer 10.50.1.1 liveness-detection minimum-receive-interval 1000

26. Configure the peer IP address and minimum transmit interval for BFD session for ICCP on QFX1 and QFX2.

QFX1:

[edit protocols]

user@switch# set iccp peer 10.50.1.2 liveness-detection transmit-interval minimum-interval 1000

QFX2:

[edit protocols]

user@switch# set iccp peer 10.50.1.1 liveness-detection transmit-interval minimum-interval 1000

27. To enable the service ID on QFX1 and QFX2:

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 10
```

Results

Here are the results of your configuration on QFX1.

```
chassis {
    aggregated-devices {
        ethernet {
            device-count 2;
        }
    }
}
interfaces {
    xe-0/0/0 {
        ether-options {
            802.3ad ae1;
        }
    }
    xe-0/0/1 {
        ether-options {
            802.3ad ae0;
        }
    }
    xe-0/0/2 {
        ether-options {
            802.3ad ae0;
        }
    }
    xe-0/0/3 {
        unit 0 {
            family ethernet-switching {
```

```
interface-mode access;
            vlan {
                members v10;
            }
        }
    }
}
ae0 {
    aggregated-ether-options {
        lacp {
            active;
        }
    }
    unit 0 {
        family ethernet-switching {
            interface-mode trunk;
            vlan {
                members [ v50 v10 ];
            }
        }
    }
}
ae1 {
    aggregated-ether-options {
        lacp {
            active;
            system-id 00:01:02:03:04:05;
            admin-key 3;
        }
        mc-ae {
            mc-ae-id 3;
            redundancy-group 1;
            chassis-id 0;
            mode active-active;
            status-control active;
            init-delay-time 240;
        }
    }
    unit 0 {
        family ethernet-switching {
            interface-mode trunk;
            vlan {
                members v10;
```

```
}
            }
        }
   }
    em0 {
        unit 0 {
            family inet {
                address 10.1.1.1/24;
            }
        }
   }
    irb {
        unit 10 {
            family inet {
                address 10.10.1.1/24;
            }
        }
        unit 50 {
            family inet {
                address 10.50.1.1/30;
            }
        }
   }
}
multi-chassis {
    multi-chassis-protection 10.50.1.2 {
        interface ae0;
   }
}
protocols {
    iccp {
        local-ip-addr 10.50.1.1;
        peer 10.50.1.2 {
            session-establishment-hold-time 340;
            redundancy-group-id-list 1;
            {\tt backup-liveness-detection}\ \{
                backup-peer-ip 10.1.1.2;
            }
            liveness-detection {
                minimum-receive-interval 1000;
                transmit-interval {
                    minimum-interval 1000;
```

```
}
    }
}
\verb|switch-options|| \{
    service-id 10;
}
vlans {
    v10 {
        vlan-id 10;
        13-interface irb.10;
        mcae-mac-synchronize;
    }
    v50 {
        vlan-id 50;
        13-interface irb.50;
    }
}
```

Display the results of the configuration on QFX2.

```
chassis {
    aggregated-devices {
        ethernet {
            device-count 2;
        }
    }
}
interfaces {
    xe-0/0/0 {
        ether-options {
            802.3ad ae1;
        }
    }
    xe-0/0/1 {
        ether-options {
            802.3ad ae0;
        }
    }
    xe-0/0/2 {
        ether-options {
            802.3ad ae0;
```

```
}
}
xe-0/0/3 {
    unit 0 {
        family ethernet-switching {
            interface-mode access;
            vlan {
                members v10;
            }
        }
    }
}
ae0 {
    aggregated-ether-options {
        lacp {
            active;
        }
    }
    unit 0 {
        family ethernet-switching {
            interface-mode trunk;
            vlan {
                members [ v50 v10 ];
            }
        }
    }
}
ae1 {
    aggregated-ether-options {
        lacp {
            active;
            system-id 00:01:02:03:04:05;
            admin-key 3;
        }
        mc-ae {
            mc-ae-id 3;
            redundancy-group 1;
            chassis-id 1;
            mode active-active;
            status-control standby;
            init-delay-time 240;
        }
```

```
unit 0 {
            family ethernet-switching {
                interface-mode trunk;
                vlan {
                    members v10;
                }
            }
        }
   }
    em0 {
        unit 0 {
            family inet {
                address 10.1.1.2/24;
        }
   }
    irb {
        unit 10 {
            family inet {
                address 10.10.1.2/24;
            }
        }
        unit 50 {
            family inet {
                address 10.50.1.2/30;
            }
        }
   }
}
multi-chassis {
    multi-chassis-protection 10.50.1.1 {
        interface ae0;
    }
}
protocols {
    iccp {
        local-ip-addr 10.50.1.2;
        peer 10.50.1.1 {
            session-establishment-hold-time 340;
            redundancy-group-id-list 1;
            backup-liveness-detection {
                backup-peer-ip 10.1.1.1;
```

```
liveness-detection {
                minimum-receive-interval 1000;
                transmit-interval {
                    minimum-interval 1000;
                }
            }
        }
    }
}
switch-options {
    service-id 10;
}
vlans {
    v10 {
        vlan-id 10;
        13-interface irb.10;
        mcae-mac-synchronize;
    }
    v50 {
        vlan-id 50;
        13-interface irb.50;
    }
}
```

Display the results of the configuration on QFX3.

```
chassis {
    aggregated-devices {
        ethernet {
            device-count 2;
        }
    }
}
interfaces {
    xe-0/0/0 {
        ether-options {
            802.3ad ae1;
        }
    }
    xe-0/0/1 {
        ether-options {
            802.3ad ae1;
```

```
}
   }
   xe-0/0/2 {
       unit 0 {
           family ethernet-switching \{
                interface-mode access;
                vlan {
                    members v10;
               }
           }
       }
   }
   ae1 {
       aggregated-ether-options {
           lacp {
               active;
           }
       }
       unit 0 {
           family ethernet-switching {
                interface-mode trunk;
                vlan {
                    members v10;
               }
           }
       }
   }
   em0 {
       unit 0 {
           family inet {
                address 10.1.1.3/24;
           }
       }
   }
   irb {
       unit 10 {
           family inet {
               address 10.10.1.3/24;
           }
       }
   }
}
vlans {
```

```
v10 {
    vlan-id 10;
    l3-interface irb.10;
}
```

Verification

IN THIS SECTION

- Verifying That ICCP Is Working on QFX1 | 99
- Verifying That LACP Is Active on QFX1 | 100
- Verifying That the MC-AE and ICL-PL Interfaces Are Up on QFX1 | 101
- Verifying That MAC Learning Is Occurring on QFX1 | 101
- Verifying That Host1 Can Connect to Host2 | 102

Verify that the configuration is working properly.

Verifying That ICCP Is Working on QFX1

Purpose

Verify that ICCP is running on QFX1.

Action

```
user@switch> show iccp

Redundancy Group Information for peer 10.50.1.2

TCP Connection : Established

Liveliness Detection : Up

Backup liveness peer status: Up

Redundancy Group ID Status

1 Up

Client Application: lacpd

Redundancy Group IDs Joined: 1
```

```
Client Application: l2ald_iccpd_client
Redundancy Group IDs Joined: 1
```

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 QFX1

Purpose

Verify that LACP is active on QFX1.

Action

```
user@switch> show lacp interfaces
Aggregated interface: ae0
   LACP state:
                    Role Exp
                                 Def Dist Col Syn Aggr Timeout Activity
     xe-0/0/1
                   Actor
                                 No
                                      Yes Yes Yes
                                                     Yes
                                                             Fast
                                                                     Active
                            No
     xe-0/0/1
                 Partner
                                      Yes Yes Yes
                                                    Yes
                                                                     Active
                            No
                                 No
                                                             Fast
     xe-0/0/2
                   Actor
                                 No
                                      Yes Yes Yes
                                                    Yes
                                                             Fast
                                                                     Active
                            No
     xe-0/0/2
                 Partner
                                      Yes Yes Yes Yes
                                                             Fast
                                                                     Active
                            No
                                 No
   LACP protocol:
                        Receive State Transmit State
                                                             Mux State
     xe-0/0/1
                              Current
                                       Fast periodic Collecting distributing
     xe-0/0/2
                              Current
                                       Fast periodic Collecting distributing
Aggregated interface: ae1
   LACP state:
                    Role
                           Exp
                                 Def
                                     Dist Col Syn Aggr Timeout Activity
     xe-0/0/0
                   Actor
                            No
                                 No
                                      Yes Yes Yes
                                                     Yes
                                                             Fast
                                                                     Active
     xe-0/0/0
                 Partner
                                                             Fast
                            No
                                 No
                                      Yes Yes Yes Yes
                                                                     Active
                        Receive State Transmit State
                                                             Mux State
   LACP protocol:
     xe-0/0/0
                              Current Fast periodic Collecting distributing
```

Meaning

This output shows that QFX1 is participating in LACP negotiation.

Verifying That the MC-AE and ICL-PL Interfaces Are Up on QFX1

Purpose

Verify that the MC-AE and ICL-PL interfaces are up on QFX1.

Action

```
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 : 10.50.1.2 ae0.0 up
```

Meaning

This output shows that the MC-AE interface on QFX1 is up and active.

Verifying That MAC Learning Is Occurring on QFX1

Purpose

Verify that MAC learning is working on QFX1.

Action

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

MAC flags (S - static MAC, D - dynamic MAC, L - locally learned, P - Persistent static, C -

Control MAC

SE - statistics enabled, NM - non configured MAC, R - remote PE MAC, O - ovsdb MAC)

Ethernet switching table : 3 entries, 3 learned
```

Routing instance : default-switch				
MAC	MAC	Age	Logical	NH
address	flags		interface	Index
00:50:56:93:73:cd	DR	-	ae0.0	0
00:50:56:93:87:58	DL	-	xe-0/0/3.0	0
00:50:56:93:89:a0	DLR	-	ae1.0	0
2	ddress	MAC Address flags 00:50:56:93:73:cd DR 10:50:56:93:87:58 DL	MAC Age address flags 00:50:56:93:73:cd DR - 10:50:56:93:87:58 DL -	MAC Age Logical dddress flags interface 00:50:56:93:73:cd DR - ae0.0 00:50:56:93:87:58 DL - xe-0/0/3.0

Meaning

The output shows three learned MAC addresses entries.

Verifying That Host1 Can Connect to Host2

Purpose

Verify that Host1 can ping Host2.

Action

```
[edit]
user@HOST1> ping 10.10.1.102
PING 10.10.1.102 (10.10.1.102): 56 data bytes
64 bytes from 10.10.1.102: icmp_seq=0 ttl=64 time=157.788 ms
64 bytes from 10.10.1.102: icmp_seq=1 ttl=64 time=153.965 ms
64 bytes from 10.10.1.102: icmp_seq=2 ttl=64 time=102.126 ms
...
```

Meaning

The output shows that HOST1 can successfully ping HOST2.

Troubleshooting

IN THIS SECTION

Troubleshooting a LAG That Is Down | 103

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 on the MX Series

IN THIS SECTION

- Requirements | 104
- Overview | 104
- Configuring the PE Routers | 105
- Configuring the CE Device | 115
- Configuring the Provider Router | 119
- Verification | 122

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

IN THIS SECTION

Topology Diagram | 105

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 8 on page 105, 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.

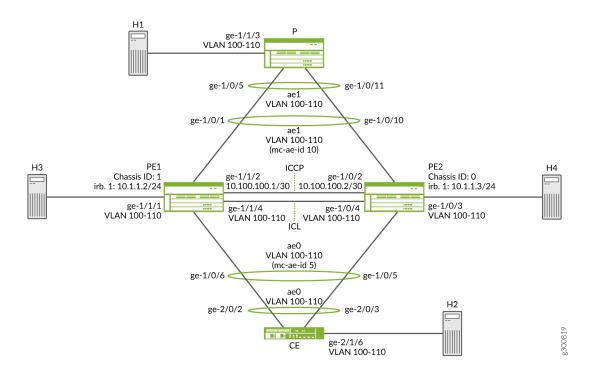
On one end of an MC-LAG is an MC-LAG client device, such as a server, that has one or more physical links in a LAG. This client device does not need to detect the MC-LAG. On the other side of an MC-LAG are two MC-LAG routers. Each of the routers has one or more physical links connected to a single client 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 8 on page 105 shows the topology used in this example.

Figure 8: MC-LAG Active-Active Mode on MX Series Routers



Configuring the PE Routers

IN THIS SECTION

- CLI Quick Configuration | 106
- Configuring the PE1 Router | 108
- Results | **112**

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 00:00:00:00:00:00:05
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 00:00:00:00:05
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 aeO 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-1/1/1.0
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-1/1/1.0;
    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

IN THIS SECTION

- CLI Quick Configuration | 115
- Configuring the CE Device | 116
- Results | 117

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-2/1/6 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 100-110
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 100-110
```

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 100-110;
}
```

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

Configuring the Provider Router

IN THIS SECTION

- CLI Quick Configuration | 119
- Configuring the PE Router | 120
- Results | 121

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 100-110
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 al1
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 PE 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 pim interfaces
- show vrrp
- show igmp
- show ospf
- show dhcp relay
- show 12-learning instance

Example: Configuring Multichassis Link Aggregation Between QFX Series Switches and MX Series Routers

IN THIS SECTION

- Requirements | 123
- Overview | 124
- Configure the Devices | 126
- Verification | 144

This example shows how to configure multichassis link aggregation groups (MC-LAGs) between a QFX Series switch and a MX Series router using active-active mode to support Layer 2 bridging. In active-active mode all member links carry traffic allowing traffic to be load balanced to both MC-LAG peers.

Requirements

This example uses the following hardware and software components:

- One Juniper Networks MX Series router (MX240, MX480, MX960)
- One Juniper Networks QFX Series switch (QFX10000, QFX5110, QFX5120)
- Two servers with LAG support; MX Series routers fill the server role in this example
- Junos OS Release 19.4R1 or later on the MC-LAG peers

Overview

IN THIS SECTION

Topology Diagram | 125

In the example topology two servers are connected to two provider edge (PE) devices, SO and R1. SO is a QFX Series switch while R1 is a MX Series router. Both PE devices have link aggregation groups (LAGs) connected to both servers. This example configures active-active mode for the MC-LAGs, meaning that both PE devices' LAG ports are active and carrying traffic at the same time.

The servers are not aware that their aggregated Ethernet links are connected to multiple PE devices. MC-LAG operation is opaque to the servers and both have a conventional Ethernet LAG interface configured.

On one end of an MC-LAG is an MC-LAG client device, for example, a server or switching/routing device, that has one or more physical links in a LAG. The client devices do not need to support MC-LAG as these devices only need to support a standard LAG interface. On the other side of the MC-LAG are two MC-LAG devices (PEs). Each of the PEs has one or more physical links connected to the client device. The PE devices coordinate with each other to ensure that data traffic is forwarded properly even when all client links are actively forwarding traffic.

In Figure 9 on page 125, the servers operate as if both LAG members were connected to a single provider device. Because the configured mode is active-active, all LAG members are in a forwarding state and the CE device load-balances the traffic to the peering PE devices.

The Interchassis Control Protocol (ICCP) sends messages between the PE devices to control the forwarding state of the MC-LAG. In addition, an interchassis link-protection link (ICL-PL) is used to forward traffic between the PE devices as needed when operating in active-active mode.

In this example you configure two MC-LAG on the PEs to support Layer 2 connectivity between the aggregated Ethernet interfaces on the servers. As part of the MC-LAG configuration you provision an aggregated Ethernet interface between the MC-LAG peers to support the ICL-PL and ICCP functionality.

Topology Diagram

Figure 9: QFX to MX MC-LAG Interoperability

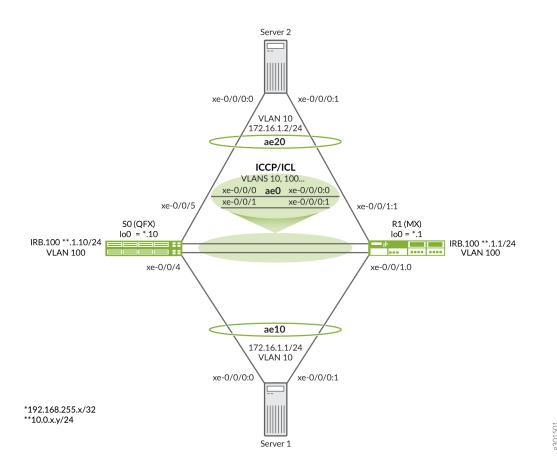


Figure 9 on page 125 shows the topology used in this example.

Key points about the topology include:

- 1. The SO node is a QFX10000 Series switch while the R1 node is a MX960 Series router.
- **2.** MX Series routers are used to fill the role of the 2 servers. Any switch, router, or server device that supports a conventional LACP based LAG interface can be used in this example.
- **3.** The servers are assigned VLAN 10 and have a shared subnet. You expect Layer 2 connectivity between the servers.
- **4.** The ICCP session between the PEs is anchored to an IRB interface. This is akin to BGP peering between loopback interfaces to survive link failures. However, here the IRBs are placed in a shared

- VLAN (VLAN 100) that provides Layer 2 connectivity between the PEs. This means that an IGP or static route is not needed for connectivity between the IRBs. As a result the IRBs share an IP subnet.
- 5. This example deploys a single LAG interface between the PEs (ae0) to support both the ICCP and ICL functionality. If desired you can run ICCP over a separate AE bundle. The use of multiple members in the AE bundle used for the ICCP/ICL links is highly recommended to ensure that they remain operational in the event of individual interface or link failures.
- **6.** While largely similar, the MC-LAG configuration differs slightly between the PE devices given they are different platforms. Demonstrating these configuration differences, and MC-LAG interoperability between the platforms, is the reason for this example. Be sure to keep track of which PE you are interacting with as you proceed through the example.

Configure the Devices

IN THIS SECTION

- CLI Quick Configuration | 126
- Configure the SO Switch | 129
- SO Results | **133**
- Configure the R1 Router | 136
- R1 Results | 140

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 <code>[edit]</code> hierarchy level. When done enter <code>commit</code> from configuration mode to activate the changes.

Switch SO

NOTE: In this example the SO device is a QFX10000 Series switch.

```
set system host-name mc-lag_r0
set chassis aggregated-devices ethernet device-count 10
set interfaces xe-0/0/0 gigether-options 802.3ad ae0
```

```
set interfaces xe-0/0/1 gigether-options 802.3ad ae0
set interfaces xe-0/0/4 gigether-options 802.3ad ae10
set interfaces xe-0/0/5 gigether-options 802.3ad ae20
set interfaces ae0 aggregated-ether-options lacp active
set interfaces ae0 unit 0 family ethernet-switching interface-mode trunk
set interfaces ae0 unit 0 family ethernet-switching vlan members all
set interfaces ae10 aggregated-ether-options lacp active
set interfaces ae10 aggregated-ether-options lacp system-id 01:01:01:01:01:01:01
set interfaces ae10 aggregated-ether-options lacp admin-key 10
set interfaces ae10 aggregated-ether-options mc-ae mc-ae-id 10
set interfaces ae10 aggregated-ether-options mc-ae redundancy-group 1
set interfaces ae10 aggregated-ether-options mc-ae chassis-id 0
set interfaces ae10 aggregated-ether-options mc-ae mode active-active
set interfaces ae10 aggregated-ether-options mc-ae status-control active
set interfaces ae10 unit 0 family ethernet-switching vlan members vlan10
set interfaces ae20 aggregated-ether-options lacp active
set interfaces ae20 aggregated-ether-options lacp system-id 02:02:02:02:02:02
set interfaces ae20 aggregated-ether-options lacp admin-key 20
set interfaces ae20 aggregated-ether-options mc-ae mc-ae-id 20
set interfaces ae20 aggregated-ether-options mc-ae redundancy-group 1
set interfaces ae20 aggregated-ether-options mc-ae chassis-id 1
set interfaces ae20 aggregated-ether-options mc-ae mode active-active
set interfaces ae20 aggregated-ether-options mc-ae status-control standby
set interfaces ae20 unit 0 family ethernet-switching vlan members vlan10
set interfaces irb unit 100 family inet address 10.0.1.10/24
set interfaces lo0 unit 0 family inet address 192.168.255.10/32
set multi-chassis multi-chassis-protection 10.0.1.1 interface ae0
set protocols iccp local-ip-addr 10.0.1.10
set protocols iccp peer 10.0.1.1 session-establishment-hold-time 50
set protocols iccp peer 10.0.1.1 redundancy-group-id-list 1
set protocols iccp peer 10.0.1.1 liveness-detection minimum-interval 1000
set switch-options service-id 100
set vlans vlan10 vlan-id 10
set vlans vlan100 vlan-id 100
set vlans vlan100 l3-interface irb.100
```

Router R1

NOTE: In this example the R1 device is a MX Series router.

```
set system host-name mc-lag_r1
set chassis aggregated-devices ethernet device-count 10
set interfaces xe-0/0/0:0 gigether-options 802.3ad ae0
set interfaces xe-0/0/0:1 gigether-options 802.3ad ae0
set interfaces xe-0/0/1:0 gigether-options 802.3ad ae10
set interfaces xe-0/0/1:1 gigether-options 802.3ad ae20
set interfaces ae0 aggregated-ether-options lacp active
set interfaces ae0 unit 0 family bridge interface-mode trunk
set interfaces ae0 unit 0 family bridge vlan-id-list 1-1000
set interfaces ae10 aggregated-ether-options lacp active
set interfaces ae10 aggregated-ether-options lacp system-id 01:01:01:01:01:01:01
set interfaces ae10 aggregated-ether-options lacp admin-key 10
set interfaces ae10 aggregated-ether-options mc-ae mc-ae-id 10
set interfaces ae10 aggregated-ether-options mc-ae redundancy-group 1
set interfaces ae10 aggregated-ether-options mc-ae chassis-id 1
set interfaces ae10 aggregated-ether-options mc-ae mode active-active
set interfaces ae10 aggregated-ether-options mc-ae status-control standby
set interfaces ae10 unit 0 multi-chassis-protection 10.0.1.10 interface ae0.0
set interfaces ae10 unit 0 family bridge interface-mode access
set interfaces ae10 unit 0 family bridge vlan-id 10
set interfaces ae20 aggregated-ether-options lacp active
set interfaces ae20 aggregated-ether-options lacp system-id 02:02:02:02:02:02
set interfaces ae20 aggregated-ether-options lacp admin-key 20
set interfaces ae20 aggregated-ether-options mc-ae mc-ae-id 20
set interfaces ae20 aggregated-ether-options mc-ae redundancy-group 1
set interfaces ae20 aggregated-ether-options mc-ae chassis-id 0
set interfaces ae20 aggregated-ether-options mc-ae mode active-active
set interfaces ae20 aggregated-ether-options mc-ae status-control active
set interfaces ae20 unit 0 multi-chassis-protection 10.0.1.10 interface ae0.0
set interfaces ae20 unit 0 family bridge interface-mode access
set interfaces ae20 unit 0 family bridge vlan-id 10
set interfaces irb unit 100 family inet address 10.0.1.1/24
set interfaces lo0 unit 0 family inet address 192.168.255.1/32
set protocols iccp local-ip-addr 10.0.1.1
set protocols iccp peer 10.0.1.10 session-establishment-hold-time 50
set protocols iccp peer 10.0.1.10 redundancy-group-id-list 1
set protocols iccp peer 10.0.1.10 liveness-detection minimum-interval 1000
```

```
set bridge-domains vlan10 vlan-id 10
set bridge-domains vlan100 vlan-id 100
set bridge-domains vlan100 routing-interface irb.100
set switch-options service-id 10
```

Server 1

NOTE: The servers in this example are MX routers. While this example focuses on configuring MC-LAG on the PE devices, the server configuration is provided for completeness. In this example server 2 has the same configuration, with the exception that it is assigned IPv4 address 172.16.1.2/24 and IPv6 address 2001:db8:172:16:1::2.

```
set system host-name server1
set chassis aggregated-devices ethernet device-count 10
set interfaces xe-0/0/0:0 gigether-options 802.3ad ae10
set interfaces xe-0/0/0:1 gigether-options 802.3ad ae10
set interfaces ae10 aggregated-ether-options lacp active
set interfaces ae10 unit 0 family inet address 172.16.1.1/24
set interfaces ae10 unit 0 family inet6 address 2001:db8:172:16:1::1/64
```

Configure the SO Switch

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 Switch SO:

1. Specify the number of aggregated Ethernet devices supported on the chassis. Only 3 LAGs are needed for the example, but having unused AE bundle capacity causes no issues.

```
[edit chassis]
user@S0# set aggregated-devices ethernet device-count 10
```

2. Configure the loopback (if desired, it's not used in this example), and IRB interfaces, along with the IRB interface's VLAN. In this example the IRB interface is used to anchor the ICCP session and is assigned to VLAN 100.

```
[edit]
user@S0# set interfaces lo0 unit 0 family inet address 192.168.255.10/32
user@S0# set interfaces irb unit 100 family inet address 10.0.1.10/24
user@S0# set vlans vlan100 vlan-id 100
user@S0# set vlans vlan100 l3-interface irb.100
```

3. Configure the ae0 interface to support ICCP and ICL. Be sure to include all MC-LAG VLANs, as well as the IRB VLAN used to support ICCP. You can specify a list of VLANs, but in this example the all keyword is used to quickly ensure all VLANs are supported over the ae0 interface. In this example only two VLANS are required on the ISL. The MC-LAG VLAN (10) and VLAN 100 that supports ICCP.

For proper operation unit 0 must be used for the ICL link on the QFX Series switch because, unlike an MX Series router, they don't support unit-level specification of the ICL link.

NOTE: The QFX Series switch supports only interface level specification of the ICL link and assumes the use of unit 0. Its is therefore important that you list all MC-LAG VLANs under unit 0 as is shown. The MX Series router can support both global or unit level specification of the ICL. The latter method is shown later in this example.

```
[edit interfaces]
user@S0# set xe-0/0/0 gigether-options 802.3ad ae0
user@S0# set xe-0/0/1 gigether-options 802.3ad ae0
user@S0# set ae0 aggregated-ether-options lacp active
user@S0# set ae0 unit 0 family ethernet-switching interface-mode trunk
user@S0# set ae0 unit 0 family ethernet-switching vlan members all
```

4. Specify the member interfaces used for the server facing aggregated Ethernet bundles.

```
[edit interfaces]
user@S0# set xe-0/0/4 gigether-options 802.3ad ae10
user@S0# set xe-0/0/5 gigether-options 802.3ad ae20
```

5. Configure the LACP and MC-LAG parameters for the MC-LAG that connects to server 1 (ae10). The MC-LAG is set for active-active mode and, in this example, S0 is set to be the active MC-LAG node using the status-control active statement. If S0 fails R1 will take over as the active node. The chassis-id statement is used by LACP for calculating the port number of the MC-LAG's physical member links. By convention the active node is assigned a chassis ID of 0 while the standby node is assigned 1. In a later step you configure R1 to be the active node for the MC-LAG connected to server 2.

The multichassis aggregated Ethernet identification number (mc-ae-id) specifies which link aggregation group the aggregated Ethernet interface belongs to. The ae10 interfaces on SO and R1 are configured with mc-ae-id 10. In like fashion the ae20 interface is configured with mc-ae-id 20.

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 ae10 and ae20 interfaces on S0 and R1 are configured with the same redundancy group, **redundancy-group 1**.

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.

```
[edit interfaces ae10]

user@SO# set aggregated-ether-options lacp active

user@SO# set aggregated-ether-options lacp system-id 01:01:01:01:01:01

user@SO# set aggregated-ether-options lacp admin-key 10

user@SO# set aggregated-ether-options mc-ae mc-ae-id 10

user@SO# set aggregated-ether-options mc-ae redundancy-group 1

user@SO# set aggregated-ether-options mc-ae chassis-id 0

user@SO# set aggregated-ether-options mc-ae mode active-active

user@SO# set aggregated-ether-options mc-ae status-control active

user@SO# set unit 0 family ethernet-switching vlan members vlan10
```

6. Configure the LACP and MC-LAG parameters for the MC-LAG that connects to server 2 (ae20). The MC-LAG is set for active-active mode and, in this example, S0 is set to be the standby MC-LAG node. In the event of R1 failure S0 takes over as the active node.

```
[edit interfaces ae20]
user@S0# set aggregated-ether-options lacp active
user@S0# set interfaces ae20 aggregated-ether-options lacp system-id 02:02:02:02:02:02
user@S0# set aggregated-ether-options lacp admin-key 20
user@S0# set aggregated-ether-options mc-ae mc-ae-id 20
user@S0# set aggregated-ether-options mc-ae redundancy-group 1
user@S0# set aggregated-ether-options mc-ae chassis-id 1
```

```
user@SO# set aggregated-ether-options mc-ae mode active-active
user@SO# set aggregated-ether-options mc-ae status-control standby
user@SO# set unit 0 family ethernet-switching vlan members v20
```

7. Configure the VLAN for the AE 10 and AE 20 bundles.

```
[edit]
user@S0# set vlans vlan10 vlan-id 10
```

8. Configure the switch-options service ID.

The ports within a bridge domain share the same flooding or broadcast characteristics in order to perform Layer 2 bridging.

The global service-id statement is required to link related bridge domains across peers (in this case SO and R1), and must be configured with the same value.

```
[edit switch-options]
user@S0# set service-id 100
```

9. Configure the ICCP parameters. The local and peer parameters are set to reflect the values configured previously for the local and remote IRB interfaces, respectively. Configuring ICCP peering to an IRB (or loopback) interface ensures that the ICCP session can remain up in the face of individual link failures.

```
[edit protocols iccp]
user@S0# set local-ip-addr 10.0.1.10
user@S0# set peer 10.0.1.1 session-establishment-hold-time 50
user@S0# set peer 10.0.1.1 redundancy-group-id-list 1
user@S0# set peer 10.0.1.10 liveness-detection minimum-interval 1000
```

10. Configure the service ID at the global level. You must configure the same unique network-wide service ID in the set of PE routers providing the service. This service ID is required when the multichassis aggregated Ethernet interfaces are part of a bridge domain.

```
[edit switch-options]
user@SO# set service-id 100
```

11. Configure the ae0 interface to function as the ICL for the MC-LAG bundles supported by S0.

```
[edit multi-chassis]
user@S0# set multi-chassis-protection 10.0.1.1 interface ae0
```

NOTE: On the QFX platform you must specify a physical interface device as the ICL protection link. Logical unit level mapping of an ICL to a MC-LAG bundle is not supported. For proper operation you must ensure that unit 0 is used to support the bridging of the MC-LAG VLANs on the ICL.

SO Results

From configuration mode, confirm your configuration by entering the show command. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
[edit]
user@S0# show
. . .chassis {
    aggregated-devices {
        ethernet {
            device-count 10;
        }
    }
}
interfaces {
    xe-0/0/0 {
        gigether-options {
            802.3ad ae0;
        }
    }
    xe-0/0/1 {
        gigether-options {
            802.3ad ae0;
        }
    }
    xe-0/0/4 {
        gigether-options {
            802.3ad ae10;
```

```
}
}
xe-0/0/5 {
    gigether-options {
        802.3ad ae20;
    }
}
ae0 {
    aggregated-ether-options {
        lacp {
            active;
        }
    }
    unit 0 {
        family ethernet-switching {
            interface-mode trunk;
            vlan {
                members all;
            }
        }
    }
}
ae10 {
    aggregated-ether-options {
        lacp {
            active;
            system-id 01:01:01:01:01:01;
            admin-key 10;
        }
        mc-ae {
            mc-ae-id 10;
            redundancy-group 1;
            chassis-id 0;
            mode active-active;
            status-control active;
        }
    }
    unit 0 {
        family ethernet-switching {
            vlan {
                members vlan10;
            }
```

```
}
    }
    ae20 {
        aggregated-ether-options {
            lacp {
                active;
                system-id 02:02:02:02:02:02;
                admin-key 20;
            }
            mc-ae {
                mc-ae-id 20;
                redundancy-group 1;
                chassis-id 1;
                mode active-active;
                status-control standby;
            }
        }
        unit 0 {
            family ethernet-switching \{
                vlan {
                    members vlan10;
                }
            }
        }
    }
    irb {
        unit 100 {
            family inet {
                address 10.0.1.10/24;
            }
        }
    }
    lo0 {
        unit 0 {
            family inet {
                address 192.168.255.10/32;
            }
        }
    }
}
multi-chassis {
    multi-chassis-protection 10.0.1.1 \{
        interface ae0;
```

```
}
}
protocols {
    iccp {
        local-ip-addr 10.0.1.10;
        peer 10.0.1.1 {
            session-establishment-hold-time 50;
            redundancy-group-id-list 1;
            liveness-detection {
                minimum-interval 1000;
            }
        }
    }
}
switch-options {
    service-id 100;
}
vlans {
    vlan10 {
        vlan-id 10;
    }
    vlan100 {
        vlan-id 100;
        13-interface irb.100;
    }
}
```

Configure the R1 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 R1:

1. Specify the number of aggregated Ethernet interfaces to be created on the chassis. Only 3 LAGs are needed, but having additional LAG capacity causes no issue.

```
[edit chassis]
user@R1# set aggregated-devices ethernet device-count 10
```

2. Configure the loopback (if desired, it's not needed in this example) and IRB interfaces, along with IRB interface's VLAN. In this example the IRB interface is used to anchor the ICCP session.

```
[edit]
user@R1# set interfaces lo0 unit 0 family inet address 192.168.255.1/32
user@R1# set interfaces irb unit 100 family inet address 10.0.1.1/24
user@R1# set bridge-domains vlan100 vlan-id 100
user@R1# set bridge-domains vlan100 routing-interface irb.100
```

3. Configure the aeO interface to support both ICL and ICCP functionality. A vlan-id-list is used to support a range of VLANs that include VLAN 100 for ICCP and VLAN 10 for the MC-LAGs. Unlike the QFX Series switch, the all used as a shortcut to support all VLANs is not supported on MX Series routers.

NOTE: The ICL link must support all MC-LAG VLANs as well as the VLAN used for ICCP. In this example this means at a minimum you must list VLAN 10 and VLAN 100 given the ae0 link supports both ISL and ICCP in this example.

```
[edit interfaces]
user@R1# set xe-0/0/0:0 gigether-options 802.3ad ae0
user@R1# set xe-0/0/0:1 gigether-options 802.3ad ae0
user@R1# set ae0 aggregated-ether-options lacp active
user@R1# set ae0 unit 0 family bridge interface-mode trunk
user@R1# set ae0 unit 0 family bridge vlan-id-list 2-1000
```

4. Specify the members to be included within the server facing aggregated Ethernet bundles at RO.

```
[edit interfaces]
user@R1# set xe-0/0/1:0 gigether-options 802.3ad ae10
user@R1# set xe-0/0/1:1 gigether-options 802.3ad ae20
```

5. Configure the LACP and MC-LAG parameters for the MC-LAG that connects to server 1 (ae10). The MC-LAG is set for active-active mode and, in this example, R1 is set to be the standby MC-LAG node using the status-control standby statement. This makes S0 the active MC-LAG node for ae10 when it's operational. If S0 fails R1 takes over as the active node. The chassis-id statement is used by LACP for calculating the port number of the MC-LAG's physical member links. By convention the active node is assigned chassis ID of 0 while the standby node is assigned 1.

The multichassis aggregated Ethernet identification number (mc-ae-id) specifies which link aggregation group the aggregated Ethernet interface belongs to. The ae10 interfaces on S0 and R1 are configured with **mc-ae-id 10**. In like fashion the ae20 interface is configured with **mc-ae-id 20**.

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 ae10 and ae20 interfaces on S0 and R1 are configured with the same redundancy group, **redundancy-group 1**.

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.

This example demonstrates MX Series router support for the specification of the ICL interface at the unit level (under the MC-LAG unit as shown below). If desired the ICL protection link can be specified globally at the physical device level (with unit 0 assumed) at the [edit multi-chassis multi-chassis-protection] hierarchy, as was shown for the QFX Series switch SO.

```
[edit interfaces ae10]

user@R1# set aggregated-ether-options lacp active

user@R1# set aggregated-ether-options lacp system-id 01:01:01:01:01:01

user@R1# set aggregated-ether-options lacp admin-key 10

user@R1# set aggregated-ether-options mc-ae mc-ae-id 10

user@R1# set aggregated-ether-options mc-ae redundancy-group 1

user@R1# set aggregated-ether-options mc-ae chassis-id 1

user@R1# set aggregated-ether-options mc-ae mode active-active

user@R1# set aggregated-ether-options mc-ae status-control standby

user@R1# set ae10 unit 0 family bridge interface-mode access

user@R1# set ae10 unit 0 family bridge vlan-id 10

user@R1# set ae10 unit 0 multi-chassis-protection 10.0.1.10 interface ae0.0
```

NOTE: On the MX platform you can specify the ICL interface using either a global level physical device declaration at the edit multi-chassis multi-chassis-protection hierarchy, or as shown here, at the logical unit level within the MC-LAG bundle. QFX Series switches support only global level specification of the physical device.

6. Configure the LACP and MC-LAG parameters for the MC-LAG that connects to server 2 (ae20). The MC-LAG is set for active-active mode and, In this example, R1 is set to be the active MC-LAG node. In the event of R1 failure S0 takes over as the active node for the ae20 MC-LAG.

```
[edit interfaces ae20]

user@R1# set aggregated-ether-options lacp active

user@R1# set aggregated-ether-options lacp system-id 02:02:02:02:02:02

user@R1# set aggregated-ether-options lacp admin-key 20

user@R1# set aggregated-ether-options mc-ae mc-ae-id 20

user@R1# set aggregated-ether-options mc-ae redundancy-group 1

user@R1# set aggregated-ether-options mc-ae chassis-id 0

user@R1# set aggregated-ether-options mc-ae mode active-active

user@R1# set aggregated-ether-options mc-ae status-control active

user@R1# set unit 0 family bridge interface-mode access

user@R1# set unit 0 family bridge vlan-id 10

user@R1# set unit 0 multi-chassis-protection 10.0.1.10 interface ae0.0
```

7. Configure the VLAN for the ae10 and ae20 bundles.

NOTE: On the MX Series router you define VLANs under the [edit bridge-domains] hierarchy. On the WFX Series switch this is done at the [edit vlans] hierarchy. This is one of the differences between the QFX Series switch and the MX Series router.

```
[edit bridge-domains]
user@R1# set vlan10 vlan-id 10
```

8. Configure the switch-options service ID.

The ports within a bridge domain share the same flooding or broadcast characteristics in order to perform Layer 2 bridging.

The global service-id statement is required to link related bridge domains across peers (in this case SO and R1), and must be configured with the same value.

```
[edit switch-options]
user@R1# set service-id 100
```

9. Configure the ICCP parameters. The local and peer parameters are set to reflect the values configured previously on the local and remote IRB interfaces, respectively. Configuring ICCP

peering to an IRB (or loopback) interface ensures that the ICCP session can remain up in the face of individual link failures.

```
[edit protocols iccp]
user@R1# set local-ip-addr 10.0.1.1
user@R1# set peer 10.0.1.10 session-establishment-hold-time 50
user@R1# set peer 10.0.1.10 redundancy-group-id-list 1
user@R1# set peer 10.0.1.10 liveness-detection minimum-interval 1000
```

10. Configure the service ID at the global level. You must configure the same unique network-wide configuration for a service in the set of PE devices providing the service. This service ID is required if the multichassis aggregated Ethernet interfaces are part of a bridge domain.

```
[edit switch-options]
user@R1# set service-id 100
```

R1 Results

From configuration mode, confirm your configuration by entering the show command. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
[edit]
user@R1# show
. . .
chassis {
    aggregated-devices {
        ethernet {
            device-count 10;
        }
    }
}
interfaces {
    xe-0/0/0:0 {
        gigether-options {
            802.3ad ae0;
        }
    }
    xe-0/0/0:1 {
        gigether-options {
```

```
802.3ad ae0;
    }
}
xe-0/0/0:2 {
    gigether-options {
        802.3ad ae1;
    }
}
xe-0/0/0:3 {
    gigether-options {
        802.3ad ae1;
    }
}
xe-0/0/1:0 {
    gigether-options {
        802.3ad ae10;
    }
}
xe-0/0/1:1 {
    {\tt gigether\text{-}options}\ \{
        802.3ad ae20;
    }
}
ae0 {
    {\tt aggregated-ether-options}\ \{
        lacp {
            active;
        }
    }
    unit 0 {
        family bridge {
             interface-mode trunk;
             vlan-id-list 2-1000;
        }
    }
}
ae10 {
    aggregated-ether-options {
        lacp {
             active;
             system-id 01:01:01:01:01:01;
             admin-key 10;
```

```
mc-ae {
            mc-ae-id 10;
            redundancy-group 1;
            chassis-id 1;
            mode active-active;
            status-control standby;
        }
    }
    unit 0 {
        multi-chassis-protection 10.0.1.10 {
            interface ae0.0;
        }
        family bridge {
            interface-mode access;
            vlan-id 10;
        }
    }
}
ae20 {
    aggregated-ether-options {
        lacp {
            active;
            system-id 02:02:02:02:02:02;
            admin-key 20;
        }
        mc-ae {
            mc-ae-id 20;
            redundancy-group 1;
            chassis-id 0;
            mode active-active;
            status-control active;
        }
    }
    unit 0 {
        multi-chassis-protection 10.0.1.10 {
            interface ae0.0;
        }
        family bridge {
            interface-mode access;
            vlan-id 10;
        }
    }
```

```
irb {
        unit 100 {
            family inet \{
                address 10.0.1.1/24;
            }
        }
   }
    lo0 {
        unit 0 {
            family inet {
                address 192.168.255.1/32;
            }
        }
   }
}
protocols {
    iccp {
        local-ip-addr 10.0.1.1;
        peer 10.0.1.10 {
            session-establishment-hold-time 50;
            redundancy-group-id-list 1;
            liveness-detection {
                minimum-interval 1000;
            }
        }
   }
}
bridge-domains {
    vlan10 {
        vlan-id 10;
   }
    vlan100 {
        vlan-id 100;
        routing-interface irb.100;
   }
}
switch-options {
    service-id 100;
}
```

Verification

Confirm that the configuration is working properly by running the following operational mode commands:

- show iccp
- show interfaces mc-ae
- show interfaces aeX (0, 10, and 20)
- On the QFX Series switch use the show vlans and the show ethernet-switching table commands
- On the MX Series router use the show bridge mac-table command
- Verify Layer 2 connectivity between the servers

Select verification commands are run to show the expected output. We start with the show iccp command on SO. If the ICCP session is not established issue the ping command between the IRB interfaces to ensure the expected Layer 2 connectivity over the aeO ICCP/ICL link:

```
user@SO# show iccp

Redundancy Group Information for peer 10.0.1.1

TCP Connection : Established

Liveliness Detection : Up

Redundancy Group ID Status

1 Up

Client Application: 12ald_iccpd_client

Redundancy Group IDs Joined: 1

Client Application: lacpd

Redundancy Group IDs Joined: 1
```

Next, we run the show interfaces mc-ae extensive command on SO. The output confirms the expected active-active status and status control active/standby state for both MC-LAGs. Recall that SO is the status control active node for ae10 and the standby node for ae20 in this example:

```
user@S0# show interfaces mc-lag extensive

Member Link : ae10

Current State Machine's State: mcae active state

Local Status : active

Local State : up

Peer Status : active
```

Peer State : up

Logical Interface : ae10.0
Topology Type : bridge
Local State : up
Peer State : up

Peer Ip/MCP/State : 10.0.1.1 ae0.0 up

MCAE Configuration

Redundancy Group : 1 MCAE ID : 10

MCAE Mode : active_active

Status Control : active Chassis ID : \emptyset

LACP Configuration

System ID : 01:01:01:01:01

Admin Key : 10

LACP Information

Local Partner System ID : 2c:6b:f5:20:55:c0
Peer Partner System ID : 2c:6b:f5:20:55:c0

Member Link : ae20

Current State Machine's State: mcae active state

Local Status : active
Local State : up
Peer Status : active
Peer State : up
Logical Interface : ae20.0
Topology Type : bridge

Local State : up
Peer State : up

Peer Ip/MCP/State : 10.0.1.1 ae0.0 up

MCAE Configuration

Redundancy Group : 1 MCAE ID : 20

MCAE Mode : active_active
Status Control : standby

Chassis ID : 1

CH03313 1D .

LACP Configuration

System ID : 02:02:02:02:02

Admin Key : 20

LACP Information

```
Local Partner System ID : 2c:6b:f5:13:24:c0

Peer Partner System ID : 2c:6b:f5:13:24:c0
```

The show interfaces command is used to confirm the ICCP/ICL, and MC-LAG bundles are up. For brevity only the output for the ae10 bundle is shown. All AE interfaces (ae0, ae10, and ae20) should be up:

```
user@S0# show interfaces ae10
 Physical interface: ae10 (MC-AE-10, active), Enabled, Physical link is Up
 Interface index: 670, SNMP ifIndex: 561
  Link-level type: Ethernet, MTU: 1514, Speed: 10Gbps, BPDU Error: None, Ethernet-Switching
Error: None, MAC-REWRITE Error: None,
  Loopback: Disabled, Source filtering: Disabled, Flow control: Disabled, Minimum links needed:
1, Minimum bandwidth needed: 1bps
  Device flags : Present Running
  Interface flags: SNMP-Traps Internal: 0x4000
  Current address: 02:05:86:72:a9:f7, Hardware address: 02:05:86:72:a9:f7
  Last flapped : 2021-04-08 11:56:43 PDT (02:37:24 ago)
  Input rate
                : 0 bps (0 pps)
  Output rate
               : 992 bps (0 pps)
  Logical interface ae10.0 (Index 578) (SNMP ifIndex 562)
    Flags: Up SNMP-Traps 0x24024000 Encapsulation: Ethernet-Bridge
    Statistics
                      Packets
                                     pps
                                                 Bytes
                                                                bps
    Bundle:
                                                                776
        Input:
                         9454
                                                963310
        Output:
                         9429
                                       0
                                              1204640
                                                                968
    Adaptive Statistics:
        Adaptive Adjusts:
        Adaptive Scans :
        Adaptive Updates:
    Protocol eth-switch, MTU: 1514
```

The show vlans detail and show ethernet-switching table commands are used to confirm VLAN definition and mapping for the ICCP/ICL, and MC-LAG interfaces on the SO device:

```
user@SO# show vlans detail

Routing instance: default-switch

VLAN Name: vlan10 State: Active

Tag: 10

Internal index: 3, Generation Index: 3, Origin: Static
```

```
MAC aging time: 300 seconds
VXLAN Enabled : No
Interfaces:
    ae0.0*,tagged,trunk
    ae10.0*,untagged,access
    ae20.0*,untagged,access
Number of interfaces: Tagged 1
                                  , Untagged 2
Total MAC count: 2
Routing instance: default-switch
VLAN Name: vlan100
                                         State: Active
Tag: 100
Internal index: 2, Generation Index: 2, Origin: Static
MAC aging time: 300 seconds
Layer 3 interface: irb.100
VXLAN Enabled : No
Interfaces:
    ae0.0*,tagged,trunk
Number of interfaces: Tagged 1
                                  , Untagged 0
Total MAC count:
show ethernet-switching table
MAC flags (S - static MAC, D - dynamic MAC, L - locally learned, P - Persistent static, C -
Control MAC
           SE - statistics enabled, NM - non configured MAC, R - remote PE MAC, O - ovsdb MAC)
Ethernet switching table : 2 entries, 2 learned
Routing instance : default-switch
    Vlan
                                            MAC
                                                               Logical
                                                                                     NH
                                                       Age
RTR
                       address
                                                              interface
                                                                                     Index
                                            flags
    name
ID
    vlan10
                           2c:6b:f5:13:24:c2 DLR
                                                                 ae20.0
0
    vlan10
                           2c:6b:f5:20:55:c3 DL
                                                                 ae10.0
```

Lastly, you ping between server 1 and 2 to confirm Layer 2 connectivity:

```
user@Server1# ping 172.16.1.2 count 2
PING 172.16.1.2 (172.16.1.2): 56 data bytes
```

```
64 bytes from 172.16.1.2: icmp_seq=0 ttl=64 time=56.529 ms
64 bytes from 172.16.1.2: icmp_seq=1 ttl=64 time=117.029 ms

--- 172.16.1.2 ping statistics ---
2 packets transmitted, 2 packets received, 0% packet loss
round-trip min/avg/max/stddev = 56.529/86.779/117.029/30.250 ms
```

```
user@Server1# ping 2001:db8:172:16:1::2 count 2
PING6(56=40+8+8 bytes) 2001:db8:172:16:1::1 --> 2001:db8:172:16:1::2
16 bytes from 2001:db8:172:16:1::2, icmp_seq=0 hlim=64 time=154.164 ms
16 bytes from 2001:db8:172:16:1::2, icmp_seq=1 hlim=64 time=167.032 ms
--- 2001:db8:172:16:1::2 ping6 statistics ---
2 packets transmitted, 2 packets received, 0% packet loss
round-trip min/avg/max/std-dev = 154.164/160.598/167.032/6.434 ms
```

Example: Configuring Multichassis Link Aggregation on EX9200 Switches in the Core for Campus Networks

IN THIS SECTION

- Requirements | 149
- Overview | 149
- Configuration | 151
- (Optional) Configuring RSTP | 177
- (Optional) Configuring IGMP Snooping | 179
- (Optional) Configuring VRRP | 181
- (Optional) Configuring MAC Address Synchronization | 184
- (Optional) Configuring OSPF | 186
- (Optional) Configuring PIM | 187
- (Optional) Configuring DHCP Relay | 190
- Verification | 193

MC-LAG in a campus configuration allows you to bond two or more physical links into a logical link between core-aggregation or aggregation-access switches. MC-LAG improves availability by providing active/active links between multiple switches over a standard Link Aggregation Group (LAG), eliminates the need for the Spanning Tree Protocol (STP), and provides faster Layer 2 convergence upon link and device failures. With multiple active network paths, MC-LAG enables you to load balance traffic across the multiple physical links. If a link fails, the traffic can be forwarded through the other available links and the aggregated link remains available.

Requirements

This example uses the following hardware and software components:

- Junos OS Release 13.2R5.10 for EX Series
- Two EX9200 switches

NOTE: This configuration example has been tested using the software release listed and is assumed to work on all later releases.

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

- Configure aggregated Ethernet interfaces on a switch. See *Configuring an Aggregated Ethernet Interface*.
- Configure the Link Aggregation Control Protocol (LACP) on aggregated Ethernet interfaces on a switch. See Configuring Aggregated Ethernet LACP (CLI Procedure).

Overview

IN THIS SECTION

Topology | 150

In this example, you configure an MC-LAG across two switches, consisting of two aggregated Ethernet interfaces, 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.

Topology

The topology used in this example consists of two switches hosting an MC-LAG. The two switches are connected to an EX4600 switch and an MX80 router. Figure 10 on page 150 shows the topology of this example.

Figure 10: Topology Diagram

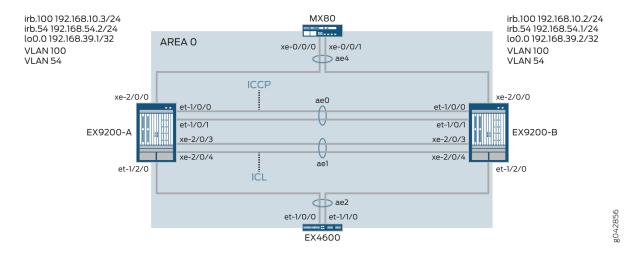


Table 5 on page 151 details the topology used in this configuration example.

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

Hostname	Base Hardware	Multichassis Link Aggregation Group
EX9200-A EX9200-B	EX9200 EX9200	ae0 is configured as an aggregated Ethernet interface, and is used as an ICCP link. The following interfaces are part of ae0: et-1/0/0 and et-1/0/1 on EX9200-A and et-1/0/0 and et-1/0/0 and et-1/0/1 on EX9200-B. ae1 is configured as an aggregated Ethernet interface and is used as an ICL link, and the following two interfaces are part of ae1: xe-2/0/3 and xe-2/0/4 on EX9200-A and xe-2/0/3 and xe-2/0/4 on EX9200-B. ae2 is configured as an MC-LAG, and the following interfaces are part of ae2: et-1/2/0 on EX9200-B. ae4 is configured as an MC-LAG, and the following interfaces are part of ae4: xe-2/0/0 on EX9200-A and xe-2/0/0 on EX9200-B.

Configuration

IN THIS SECTION

- CLI Quick Configuration | 152
- Configuring MC-LAG on Switch A | 155
- Configuring MC-LAG on Switch B | 161
- Results | **167**

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.

EX9200-A

```
set chassis aggregated-devices ethernet device-count 20
set interfaces et-1/0/0 ether-options 802.3ad ae0
set interfaces et-1/0/1 ether-options 802.3ad ae0
set interfaces et-1/2/0 ether-options 802.3ad ae2
set interfaces xe-2/0/3 hold-time up 100
set interfaces xe-2/0/3 hold-time down 9000
set interfaces xe-2/0/3 ether-options 802.3ad ae1
set interfaces xe-2/0/4 hold-time up 100
set interfaces xe-2/0/4 hold-time down 9000
set interfaces xe-2/0/4 ether-options 802.3ad ae1
set interfaces xe-2/0/0 ether-options 802.3ad ae4
set interfaces ae0 aggregated-ether-options lacp active
set interfaces ae0 aggregated-ether-options lacp periodic fast
set interfaces ae0 unit 0 family inet address 192.168.90.1/24
set interfaces ae1 description ICL-LINK
set interfaces ae1 aggregated-ether-options lacp active
set interfaces ae1 aggregated-ether-options lacp periodic fast
set interfaces ae1 unit 0 family ethernet-switching interface-mode trunk
set interfaces ae1 unit 0 family ethernet-switching vlan members all
set interfaces ae2 aggregated-ether-options lacp active
set interfaces ae2 aggregated-ether-options lacp periodic fast
set interfaces ae2 aggregated-ether-options lacp system-id 00:01:02:03:04:05
set interfaces ae2 aggregated-ether-options lacp admin-key 3
set interfaces ae2 aggregated-ether-options mc-ae mc-ae-id 3
set interfaces ae2 aggregated-ether-options mc-ae redundancy-group 1
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 520
set interfaces ae2 aggregated-ether-options mc-ae events iccp-peer-down prefer-status-control-
active
set interfaces ae2 unit 0 family ethernet-switching interface-mode trunk
set interfaces ae2 unit 0 family ethernet-switching vlan members all
set interfaces ae4 aggregated-ether-options lacp active
```

```
set interfaces ae4 aggregated-ether-options lacp periodic fast
set interfaces ae4 aggregated-ether-options lacp system-id 00:01:02:03:04:06
set interfaces ae4 aggregated-ether-options lacp admin-key 7
set interfaces ae4 aggregated-ether-options mc-ae mc-ae-id 7
set interfaces ae4 aggregated-ether-options mc-ae redundancy-group 1
set interfaces ae4 aggregated-ether-options mc-ae chassis-id 0
set interfaces ae4 aggregated-ether-options mc-ae mode active-active
set interfaces ae4 aggregated-ether-options mc-ae status-control active
set interfaces ae4 aggregated-ether-options mc-ae init-delay-time 520
set interfaces ae4 aggregated-ether-options mc-ae events iccp-peer-down prefer-status-control-
active
set interfaces ae4 unit 0 family ethernet-switching interface-mode trunk
set interfaces ae4 unit 0 family ethernet-switching vlan members v54
set vlans rack_1 vlan-id 100
set vlans rack_1 vlan-id 54
set vlans rack_1 l3-interface irb.100
set vlans v54 13-interface irb.54
set interfaces irb unit 54 family inet address 192.168.54.2/24 arp 192.168.54.1 12-interface
ae1.0
set interfaces irb unit 54 family inet address 192.168.54.2/24 arp 192.168.54.1 mac
3c:8a:b0:85:78:70
set interfaces irb unit 100 family inet address 192.168.10.3/24 arp 192.168.10.2 12-interface
ae1.0
set interfaces irb unit 100 family inet address 192.168.10.3/24 arp 192.168.10.2 mac
3c:8a:b0:85:78:70
set interfaces lo0 unit 0 family inet address 192.168.39.1/32
set protocols iccp local-ip-addr 192.168.39.1
set protocols iccp peer 192.168.39.2 session-establishment-hold-time 50
set protocols iccp peer 192.168.39.2 redundancy-group-id-list 1
set protocols iccp peer 192.168.39.2 backup-liveness-detection backup-peer-ip 10.105.5.6
set protocols iccp peer 192.168.39.2 liveness-detection minimum-interval 2000
set protocols iccp peer 192.168.39.2 liveness-detection multiplier 4
set multi-chassis multi-chassis-protection 192.168.39.2 interface ae1
set switch-options service-id 1
```

EX9200-B

```
set chassis aggregated-devices ethernet device-count 20
set interfaces et-1/0/0 ether-options 802.3ad ae0
set interfaces et-1/0/1 ether-options 802.3ad ae0
set interfaces et-1/2/0 ether-options 802.3ad ae2
set interfaces xe-2/0/0 ether-options 802.3ad ae4
```

```
set interfaces xe-2/0/3 hold-time up 100
set interfaces xe-2/0/3 hold-time down 9000
set interfaces xe-2/0/3 ether-options 802.3ad ae1
set interfaces xe-2/0/4 hold-time up 100
set interfaces xe-2/0/4 hold-time down 9000
set interfaces xe-2/0/4 ether-options 802.3ad ae1
set interfaces ae0 aggregated-ether-options lacp active
set interfaces ae0 aggregated-ether-options lacp periodic fast
set interfaces ae0 unit 0 family inet address 192.168.90.2/24
set interfaces ae1 description ICL-LINK
set interfaces ae1 aggregated-ether-options lacp active
set interfaces ae1 aggregated-ether-options lacp periodic fast
set interfaces ae1 unit 0 family ethernet-switching interface-mode trunk
set interfaces ae1 unit 0 family ethernet-switching vlan members all
set interfaces ae2 aggregated-ether-options lacp active
set interfaces ae2 aggregated-ether-options lacp periodic fast
set interfaces ae2 aggregated-ether-options lacp system-id 00:01:02:03:04:05
set interfaces ae2 aggregated-ether-options lacp admin-key 3
set interfaces ae2 aggregated-ether-options mc-ae mc-ae-id 3
set interfaces ae2 aggregated-ether-options mc-ae redundancy-group 1
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 init-delay-time 520
set interfaces ae2 aggregated-ether-options mc-ae events
set interfaces ae2 aggregated-ether-options mc-ae status-control standby
set interfaces ae2 unit 0 family ethernet-switching interface-mode trunk
set interfaces ae2 unit 0 family ethernet-switching vlan members all
set interfaces ae4 aggregated-ether-options lacp active
set interfaces ae4 aggregated-ether-options lacp periodic fast
set interfaces ae4 aggregated-ether-options lacp system-id 00:01:02:03:04:06
set interfaces ae4 aggregated-ether-options lacp admin-key 7
set interfaces ae4 aggregated-ether-options mc-ae mc-ae-id 7
set interfaces ae4 aggregated-ether-options mc-ae redundancy-group 1
set interfaces ae4 aggregated-ether-options mc-ae chassis-id 1
set interfaces ae4 aggregated-ether-options mc-ae mode active-active
set interfaces ae4 aggregated-ether-options mc-ae status-control standby
set interfaces ae4 aggregated-ether-options mc-ae init-delay-time 520
set interfaces ae4 unit 0 family ethernet-switching interface-mode trunk
set interfaces ae4 unit 0 family ethernet-switching vlan members v54
set vlans rack_1 vlan-id 100
set vlans rack_1 13-interface irb.100
set vlans v54 vlan-id 54
set vlans v54 13-interface irb.54
```

```
set interfaces irb unit 54 family inet address 192.168.54.1/24 arp 192.168.54.2 12-interface
ae1.0
set interfaces irb unit 54 family inet address 192.168.54.1/24 arp 192.168.54.2 mac
00:1f:12:b6:6f:f0
set interfaces irb unit 100 family inet address 192.168.10.2/24 arp 192.168.10.3 12-interface
ae1.0
set interfaces irb unit 100 family inet address 192.168.10.2/24 arp 192.168.10.3 mac
00:1f:12:b6:6f:f0
set interfaces lo0 unit 0 family inet address 192.168.39.2/32
set protocols iccp local-ip-addr 192.168.39.2
set protocols iccp peer 192.168.39.1 session-establishment-hold-time 50
set protocols iccp peer 192.168.39.1 redundancy-group-id-list 1
set protocols iccp peer 192.168.39.1 backup-liveness-detection backup-peer-ip 10.105.5.5
set protocols iccp peer 192.168.39.1 liveness-detection minimum-interval 2000
set protocols iccp peer 192.168.39.1 liveness-detection multiplier 4
set multi-chassis multi-chassis-protection 192.168.39.1 interface ae1
set switch-options service-id 1
```

Configuring MC-LAG on Switch A

Step-by-Step Procedure

The following example requires you to navigate various levels in the configuration hierarchy.

1. Configure the number of aggregated Ethernet interfaces to be created on Switch A.

```
[edit chassis]
user@switch# set aggregated-devices ethernet device-count 20
```

2. Add member interfaces to the aggregated Ethernet interfaces that will be used for the Inter-Chassis Control Protocol (ICCP) interface.

```
[edit interfaces]
user@switch# set et-1/0/0 ether-options 802.3ad ae0
user@switch# set et-1/0/1 ether-options 802.3ad ae0
```

3. Specify the member interfaces that belong to interface ae2.

```
[edit interfaces]
user@switch# set et-1/2/0 ether-options 802.3ad ae2
```

4. Configure the member interfaces for the interchassis link (ICL) with a hold-time value that is higher than the configured BFD timer to prevent the ICL from being advertised as being down before the ICCP link is down.

If the ICL goes down before the ICCP link goes down, the MC-LAG interface configured as the standby status-control peer goes up and down. The interface going up and down causes a delay in convergence.

```
[edit interfaces]
user@switch# set xe-2/0/3 hold-time up 100
user@switch# set xe-2/0/3 hold-time down 9000
user@switch# set xe-2/0/3 ether-options 802.3ad ae1
user@switch# set xe-2/0/4 hold-time up 100
user@switch# set xe-2/0/4 hold-time down 9000
user@switch# set xe-2/0/4 ether-options 802.3ad ae1
```

5. Specify the members that belong to ae4.

Specify the members that belong to ae4.

```
[edit interfaces]
user@switch# set xe-2/0/0 ether-options 802.3ad ae4
```

6. Configure ae0 as a Layer 3 interface.

```
[edit interfaces]
user@switch# set ae0 aggregated-ether-options lacp active
user@switch# set ae0 aggregated-ether-options lacp periodic fast
user@switch# set ae0 unit 0 family inet address 192.168.90.1/24
```

7. Configure ae1 as a Layer 2 interface.

```
[edit interfaces]
user@switch# set ae1 description ICL-LINK
```

```
user@switch# set ae1 aggregated-ether-options lacp active
user@switch# set ae1 aggregated-ether-options lacp periodic fast
```

8. Configure a trunk interface between EX9200-A and EX9200-B.

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

9. Configure the LACP parameters on ae2.

```
[edit interfaces]
user@switch# set ae2 aggregated-ether-options lacp active
user@switch# set ae2 aggregated-ether-options lacp periodic fast
```

10. Configure the LACP system ID.

```
[edit interfaces]
user@switch# set ae2 aggregated-ether-options lacp system-id 00:01:02:03:04:06
```

11. Configure the MC-AE interface properties.

```
[edit interfaces]
user@switch# set ae2 aggregated-ether-options lacp admin-key 3
user@switch# set ae2 aggregated-ether-options mc-ae mc-ae-id 3
user@switch# set ae2 aggregated-ether-options mc-ae redundancy-group 1
```

12. Specify a unique chassis ID for the MC-LAG that the aggregated Ethernet interface belongs to.

```
[edit interfaces]
user@switch# set ae2 aggregated-ether-options mc-ae chassis-id 0
```

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

```
[edit interfaces]
user@switch# set ae2 aggregated-ether-options mc-ae mode active-active
```

14. Configure the status control on the switch that hosts the MC-LAG.

If one switch is in active mode, then the other switch must be in standby mode.

```
[edit interfaces]
user@switch# set ae2 aggregated-ether-options mc-ae status-control active
```

15. Specify the time in seconds by when routing adjacencies must form.

```
[edit interfaces]
user@switch# set ae2 aggregated-ether-options mc-ae init-delay-time 520
```

16. Specify that if a peer of the MC-LAG group goes down, the peer that is configured as status-control active becomes the active peer.

```
[edit interfaces]
user@switch# set ae2 aggregated-ether-options mc-ae events iccp-peer-down prefer-status-
control-active
```

17. Configure ae2 as a trunk port with membership in all VLANs.

```
[edit interfaces]
user@switch# set ae2 unit 0 family ethernet-switching interface-mode trunk
user@switch# set ae2 unit 0 family ethernet-switching vlan members all
```

18. Configure the LACP parameters on ae4.

```
[edit interfaces]
user@switch# set ae4 aggregated-ether-options lacp active
user@switch# set ae4 aggregated-ether-options lacp periodic fast
```

19. Specify the LACP administration key.

```
[edit interfaces]
user@switch# set ae4 aggregated-ether-options lacp system-id 00:01:02:03:04:06
user@switch# set ae4 aggregated-ether-options lacp admin-key 7
```

```
user@switch# set ae4 aggregated-ether-options mc-ae mc-ae-id 7
user@switch# set ae4 aggregated-ether-options mc-ae redundancy-group 1
```

20. Specify a unique chassis ID for the MC-LAG that the aggregated Ethernet interface belongs to.

```
[edit interfaces]
user@switch# set ae4 aggregated-ether-options mc-ae chassis-id 0
user@switch# set ae4 aggregated-ether-options mc-ae mode active-active
```

21. Configure the status control on the switch that hosts the MC-LAG.

If one switch is in active mode, then the other switch must be in standby mode.

```
[edit interfaces]
user@switch# set ae4 aggregated-ether-options mc-ae status-control active
user@switch# set ae4 aggregated-ether-options mc-ae init-delay-time 520
user@switch# set ae4 aggregated-ether-options mc-ae events iccp-peer-down prefer-status-
control-active
```

22. Configure ae4 as a Layer 2 interface.

```
[edit interfaces]
user@switch# set ae4 unit 0 family ethernet-switching interface-mode trunk
user@switch# set ae4 unit 0 family ethernet-switching vlan members v54
```

23. Configure VLAN rack_1 and configure a Layer 3 IRB interface on VLAN rack_1.

```
[edit vlans]
user@switch# set rack_1 vlan-id 100
user@switch# set rack_1 l3-interface irb.100
```

24. Configure VLAN rack_1.

```
[edit vlans]
user@switch# set rack_1 vlan-id 54
```

25. Configure VLAN 54 and configure a Layer 3 IRB on VLAN 54.

```
[edit vlans]
user@switch# set v54 vlan-id 54
user@switch# set v54 l3-interface irb.54
```

26. Configure an IRB interface on VLAN 54.

You must configure static ARP on the MC-LAG peers to allow routing protocols to traverse over the IRB interface.

```
[edit interfaces]
user@switch# set irb unit 54 family inet address 192.168.54.2/24 arp 192.168.54.1 12-
interface ae1.0
user@switch# set irb unit 54 family inet address 192.168.54.2/24 arp 192.168.54.1 mac
3c:8a:b0:85:78:70
```

27. Configure static ARP on the MC-LAG peers to allow routing protocols to traverse over the IRB interface

```
[edit interfaces]
user@switch# set irb unit 100 family inet address 192.168.10.3/24 arp 192.168.10.2 l2-
interface ae1.0
user@switch# set irb unit 100 family inet address 192.168.10.3/24 arp 192.168.10.2 mac
3c:8a:b0:85:78:70
```

28. Configure a loopback interface.

```
[edit interfaces]
user@switch# set lo0 unit 0 family inet address 192.168.39.2/32
```

29. Configure ICCP using the loopback address.

```
[edit protocols]
user@switch# set iccp local-ip-addr 192.168.39.1
```

30. Configure the session establishment hold time for ICCP to connect faster.

```
[edit protocols]
user@switch# set iccp peer 192.168.39.2 session-establishment-hold-time 50
user@switch# set iccp peer 192.168.39.2 redundancy-group-id-list 1
user@switch# set iccp peer 192.168.39.2 backup-liveness-detection backup-peer-ip 10.105.5.6
```

31. To enable Bidirectional Forwarding Detection (BFD), configure the minimum receive interval.

We recommend a minimum receive interval value of 6 seconds.

```
[edit protocols]
user@switch# set iccp peer 192.168.39.2 liveness-detection minimum-interval 2000
user@switch# set iccp peer 192.168.39.2 liveness-detection multiplier 4
[edit multi-chassis]
user@switch# set multi-chassis-protection 192.168.39.2 interface ae1
```

32. 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 1
```

Configuring MC-LAG on Switch B

Step-by-Step Procedure

The following example requires you to navigate various levels in the configuration hierarchy.

1. Configure the number of aggregated Ethernet interfaces to be created on Switch A.

```
[edit chassis]
user@switch# set aggregated-devices ethernet device-count 20
```

2. Add member interfaces to the aggregated Ethernet interfaces that will be used for the Inter-Chassis Control Protocol (ICCP) interface.

```
[edit interfaces]
user@switch# set et-1/0/0 ether-options 802.3ad ae0
user@switch# set et-1/0/1 ether-options 802.3ad ae0
```

3. Specify the member interfaces that belong to interface ae2.

```
[edit interfaces]
user@switch# set et-1/2/0 ether-options 802.3ad ae2
```

4. Configure the member interfaces for the interchassis link (ICL) with a hold-time value that is higher than the configured BFD timer to prevent the ICL from being advertised as being down before the ICCP link is down.

If the ICL goes down before the ICCP link goes down, the MC-LAG interface configured as the standby status-control peer goes up and down. The interface going up and down causes a delay in convergence.

```
[edit interfaces]
user@switch# set xe-2/0/3 hold-time up 100
user@switch# set xe-2/0/3 hold-time down 9000
user@switch# set xe-2/0/3 ether-options 802.3ad ae1
user@switch# set xe-2/0/4 hold-time up 100
user@switch# set xe-2/0/4 hold-time down 9000
user@switch# set xe-2/0/4 ether-options 802.3ad ae1
```

5. Specify the members that belong to ae4.

```
[edit interfaces]
user@switch# set xe-2/0/0 ether-options 802.3ad ae4
```

6. Configure ae0 as a Layer 3 interface.

```
[edit interfaces]
user@switch# set ae0 aggregated-ether-options lacp active
```

```
user@switch# set ae0 aggregated-ether-options lacp periodic fast user@switch# set ae0 unit 0 family inet address 192.168.90.2/24
```

7. Configure ae1 as a Layer 2 interface.

```
[edit interfaces]
user@switch# set ae1 description ICL-LINK
user@switch# set ae1 aggregated-ether-options lacp active
user@switch# set ae1 aggregated-ether-options lacp periodic fast
```

8. Configure a trunk interface between EX9200-A and EX9200-B.

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

9. Configure the LACP parameters on ae2.

```
[edit interfaces]
user@switch# set ae2 aggregated-ether-options lacp active
user@switch# set ae2 aggregated-ether-options lacp periodic fast
```

10. Configure the LACP system ID.

```
[edit interfaces]
user@switch# set ae2 aggregated-ether-options lacp system-id 00:01:02:03:04:05
```

11. Configure the MC-AE interface properties.

```
[edit interfaces]
user@switch# set ae2 aggregated-ether-options lacp admin-key 3
user@switch# set ae2 aggregated-ether-options mc-ae mc-ae-id 3
user@switch# set ae2 aggregated-ether-options mc-ae redundancy-group 1
```

12. Specify a unique chassis ID for the MC-LAG that the aggregated Ethernet interface belongs to.

```
[edit interfaces]
user@switch# set ae2 aggregated-ether-options mc-ae chassis-id 1
```

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

```
[edit interfaces]
user@switch# set ae2 aggregated-ether-options mc-ae mode active-active
```

14. Specify the time in seconds by when routing adjacencies must form.

```
[edit interfaces]
user@switch# set ae2 aggregated-ether-options mc-ae init-delay-time 520
```

15. Configure the status control on the switch that hosts the MC-LAG.

If one switch is in active mode, then the other switch must be in standby mode.

```
[edit interfaces]
user@switch# set ae2 aggregated-ether-options mc-ae status-control standby
```

16. Configure ae2 as a trunk port with membership in all VLANs.

```
[edit interfaces]
user@switch# set ae2 unit 0 family ethernet-switching interface-mode trunk
user@switch# set ae2 unit 0 family ethernet-switching vlan members all
```

17. Configure the LACP parameters on ae4.

```
[edit interfaces]
user@switch# set ae4 aggregated-ether-options lacp active
user@switch# set ae4 aggregated-ether-options lacp periodic fast
```

18. Specify the LACP administration key.

```
[edit interfaces]
user@switch# set ae4 aggregated-ether-options lacp system-id 00:01:02:03:04:06
user@switch# set ae4 aggregated-ether-options lacp admin-key 7
user@switch# set ae4 aggregated-ether-options mc-ae mc-ae-id 7
user@switch# set ae4 aggregated-ether-options mc-ae redundancy-group 1
```

19. Specify a unique chassis ID for the MC-LAG that the aggregated Ethernet interface belongs to.

```
[edit interfaces]
user@switch# set ae4 aggregated-ether-options mc-ae chassis-id 1
user@switch# set ae4 aggregated-ether-options mc-ae mode active-active
```

20. Configure the status control on the switch that hosts the MC-LAG.

If one switch is in active mode, then the other switch must be in standby mode.

```
[edit interfaces]
user@switch# set ae4 aggregated-ether-options mc-ae status-control standby
user@switch# set ae4 aggregated-ether-options mc-ae init-delay-time 520
user@switch# set ae4 aggregated-ether-options mc-ae events iccp-peer-down prefer-status-
control-active
```

21. Configure ae4 as a Layer 2 interface.

```
[edit interfaces]
user@switch# set ae4 unit 0 family ethernet-switching interface-mode trunk
user@switch# set ae4 unit 0 family ethernet-switching vlan members v54
```

22. Configure VLAN rack_1 and configure a Layer 3 IRB interface on VLAN rack_1.

```
[edit vlans]
user@switch# set rack_1 vlan-id 100
user@switch# set rack_1 l3-interface irb.100
```

23. Configure VLAN 54 and configure an IRB on VLAN 54.

```
[edit vlans]
user@switch# set v54 vlan-id 54
user@switch# set v54 l3-interface irb.54
```

24. Configure static ARP on the MC-LAG peers to allow routing protocols to traverse over the IRB interface.

```
[edit interfaces]
user@switch# set irb unit 54 family inet address 192.168.54.1/24 arp 192.168.54.2 l2-
interface ae1.0
user@switch# set irb unit 54 family inet address 192.168.54.1/24 arp 192.168.54.2 mac mac
00:1f:12:b6:6f:f0
```

25. Configure static Address Resolution Protocol (ARP) on the MC-LAG IRB peers to allow routing protocols to traverse the IRB interface.

```
[edit interfaces]
user@switch# set irb unit 100 family inet address 192.168.10.2/24 arp 192.168.10.3 12-
interface ae1.0
user@switch# set irb unit 100 family inet address 192.168.10.2/24 arp 192.168.10.3 mac
00:1f:12:b6:6f:f0
```

26. Configure a loopback interface.

```
[edit interfaces]
user@switch# set lo0 unit 0 family inet address 192.168.39.2/32
```

27. Configure ICCP using the loopback address.

```
[edit protocols]
user@switch# set iccp local-ip-addr 192.168.39.2
```

28. Configure the session establishment hold time for ICCP to connect faster.

```
[edit protocols]
user@switch# set iccp peer 192.168.39.1 session-establishment-hold-time 50
user@switch# set iccp peer 192.168.39.1 redundancy-group-id-list 1
user@switch# set iccp peer 192.168.39.1 backup-liveness-detection backup-peer-ip 10.105.5.5
```

29. To enable Bidirectional Forwarding Detection (BFD), configure the minimum receive interval.

We recommend a minimum receive interval value of 6 seconds.

```
[edit protocols]
user@switch# set iccp peer 192.168.39.1 liveness-detection minimum-interval 2000
user@switch# set iccp peer 192.168.39.1 liveness-detection multiplier 4
[edit multi-chassis]
user@switch# set multi-chassis-protection 192.168.39.1 interface ae1
```

30. 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 1
```

Results

Display the results of the configuration on EX9200-A.

```
user@switch> show chassis
  chassis {
    redundancy {
        graceful-switchover;
    }
    aggregated-devices {
        ethernet {
            device-count 20;
        }
}
```

```
}
```

```
user@switch> show interfaces
   interfaces {
       et-1/0/0 {
           ether-options {
               802.3ad ae0;
           }
       }
       et-1/0/1 {
           ether-options {
               802.3ad ae0;
           }
       }
       et-1/2/0 {
           ether-options {
               802.3ad ae2;
           }
       }
       xe-2/0/3 {
           hold-time up 100 down 7000;
           ether-options {
               802.3ad ae1;
           }
       }
       xe-2/0/4 {
           hold-time up 100 down 7000;
           ether-options {
               802.3ad ae1;
           }
       }
       ae0 {
           aggregated-ether-options {
               lacp {
                   active;
                   periodic fast;
               }
           }
           unit 0 {
               family inet {
```

```
address 192.168.90.1/24;
       }
    }
}
ae1 {
    description ICL-LINK;
    aggregated-ether-options {
        lacp {
            active;
            periodic fast;
        }
    }
    unit 0 {
        family ethernet-switching {
            interface-mode trunk;
            vlan {
                members all;
            }
        }
    }
}
ae2 {
    aggregated-ether-options {
        lacp {
            active;
            periodic fast;
            system-id 00:01:02:03:04:05;
            admin-key 3;
        }
        mc-ae {
            mc-ae-id 3;
            redundancy-group 1;
            chassis-id 0;
            mode active-active;
            status-control active;
            init-delay-time 520;
            events {
                iccp-peer-down {
                    prefer-status-control-active;
                }
            }
        }
```

```
unit 0 {
        family ethernet-switching {
            interface-mode trunk;
            vlan {
                members all;
            }
        }
    }
}
ae4 {
    aggregated-ether-options {
        lacp {
            active;
            periodic fast;
            system-id 00:01:02:03:04:06;
            admin-key 7;
        }
        mc-ae {
            mc-ae-id 7;
            redundancy-group 1;
            chassis-id 0;
            mode active-active;
            status-control standby;
            init-delay-time 520;
            events {
                iccp-peer-down {
                    prefer-status-control-active;
                }
            }
        }
    }
        unit 0 {
            family ethernet-switching {
                interface-mode trunk;
                vlan {
                    members [ rack_1 v54 ];
                }
            }
    irb {
        arp-l2-validate;
            unit 54 {
                family inet {
                    address 192.168.54.2/24 {
```

```
arp 192.168.54.1 12-interface ae1.0 mac 3c:8a:b0:85:78:70;
           }
       }
   }
        unit 100 {
            family inet {
                address 192.168.10.3/24 {
                    arp 192.168.10.2 12-interface ae1.0 mac 3c:8a:b0:85:78:70;
               }
           }
       }
}
100 {
    unit 0 {
        family inet {
            address 192.168.39.1/32;
       }
   }
}
```

```
user@switch> show multi-chassis
multi-chassis {
    multi-chassis-protection 192.168.39.2 {
        interface ae1;
    }
}
```

```
user@switch> show protocols

protocols {
    iccp {
        local-ip-addr 192.168.39.1;
        peer 192.168.39.2 {
            session-establishment-hold-time 50;
            redundancy-group-id-list 1;
            backup-liveness-detection {
                backup-peer-ip 10.105.5.6;
        }
        liveness-detection {
                 minimum-interval 2000;
                 multiplier 3;
        }
}
```

```
}
}

lldp {
  interface all;
}
layer2-control {
  nonstop-bridging;
}
```

```
user@switch> show switch-options
switch-options {
    service-id 1;
}
```

```
user@switch> show vlans
  vlans {
    rack_1 {
        vlan-id 100;
        l3-interface irb.100;

    }
        v54 {
            vlan-id 54;
            l3-interface irb.54;
          }
}
```

Display the results of the configuration on EX9200-B.

```
user@switch> show chassis
  chassis {
    redundancy {
        graceful-switchover;
    }
    aggregated-devices {
        ethernet {
            device-count 20;
    }
}
```

```
}
}
```

```
user@switch> show interfaces
   interfaces {
       et-1/0/0 {
            ether-options {
                802.3ad ae0;
            }
       }
       et-1/0/1 {
            \hbox{ ether-options } \{
                802.3ad ae0;
            }
       }
       et-1/2/0 {
            ether-options {
               802.3ad ae2;
            }
       }
       xe-2/0/3 {
            hold-time up 100 down 7000;
            ether-options {
                802.3ad ae1;
            }
       }
       xe-2/0/4 {
            hold-time up 100 down 7000;
            ether-options {
                802.3ad ae1;
            }
       }
       ae0 {
            aggregated-ether-options {
                lacp {
                    active;
                    periodic fast;
            }
                unit 0 {
```

```
family inet {
            address 192.168.90.2/24;
        }
}
ae1 {
    description ICL-LINK;
    aggregated-ether-options {
        lacp {
            active;
            periodic fast;
        }
    }
    unit 0 {
        family ethernet-switching {
            interface-mode trunk;
            vlan {
                members all;
            }
        }
    }
}
ae2 {
    aggregated-ether-options {
        lacp {
            active;
            periodic fast;
            system-id 00:01:02:03:04:05;
            admin-key 3;
        }
        mc-ae {
            mc-ae-id 3;
            redundancy-group 1;
            chassis-id 1;
            mode active-active;
            status-control active;
            init-delay-time 520;
            events {
                iccp-peer-down {
                    prefer-status-control-active;
                }
            }
        }
```

```
unit 0 {
        family ethernet-switching {
            interface-mode trunk;
            vlan {
                members all;
            }
        }
}
ae4 {
    aggregated-ether-options {
        lacp {
            active;
            periodic fast;
            system-id 00:01:02:03:04:06;
            admin-key 7;
        }
        mc-ae {
            mc-ae-id 7;
            redundancy-group 1;
            chassis-id 1;
            mode active-active;
            status-control standby;
            init-delay-time 520;
        }
    }
        unit 0 {
            family ethernet-switching {
                interface-mode trunk;
                vlan {
                    members [rack_1 v54 ];
                }
            }
    irb {
    arp-l2-validate;
        unit 54 {
            family inet {
                address 192.168.54.1/24 {
                    arp 192.168.54.2 12-interface ae1.0 mac 00:1f:12:b6:6f:f0;
                }
            }
        }
            unit 100 {
```

```
family inet {
        address 192.168.10.2/24 {
            arp 192.168.10.3 12-interface ae1.0 mac 00:1f:12:b6:6f:f0;
        }
    }
}

lo0 {
    unit 0 {
        family inet {
            address 192.168.39.2/32;
        }
}
```

```
user@switch> show multi-chassis
  multi-chassis {
    multi-chassis-protection 192.168.39.1 {
        interface ae1;
    }
}
```

```
user@switch> show protocols
   protocols {
       iccp {
            local-ip-addr 192.168.39.2;
            peer 192.168.39.1 {
                session-establishment-hold-time 50;
                redundancy-group-id-list 1;
                backup-liveness-detection {
                    backup-peer-ip 10.105.5.5;
                }
                liveness-detection {
                    minimum-interval 2000;
                    multiplier 3;
                }
           }
       }
       11dp {
            interface all;
       layer2-control {
```

```
nonstop-bridging;
}
```

```
user@switch> show switch-options
switch-options {
    service-id 1;
}
```

(Optional) Configuring RSTP

IN THIS SECTION

- CLI Quick Configuration | 177
- Configuring Switch A and Switch B | 178

CLI Quick Configuration

Switch A and Switch B

```
set protocols rstp interface ae2
set protocols rstp interface ae4
```

```
set protocols rstp system-identifier 00:01:02:03:04:05
set protocols rstp bridge-priority 0
```

Configuring Switch A and Switch B

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 Switch A and Switch B:

1. Enable the Rapid Spanning Tree Protocol on the ae2 and ae4 interfaces for optional loop prevention.

```
[edit protocols]
user@switch# set rstp interface ae2
user@switch# set rstp interface ae4
```

2. Configure the system identifier.

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

3. Set Rapid Spanning Tree Protocol priority to 0. This will make the MC-AE node the highest priority.

```
[edit protocols]
user@switch# set rstp bridge-priority 0
```

Switch A and Switch B

From configuration mode, confirm your configuration by entering the show protocols rstp commands. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
user@switch> show protocols rstp
rstp {
    system-identifier 00:01:02:03:04:05;
    interface ae2;
```

```
interface ae4;
}
```

(Optional) Configuring IGMP Snooping

IN THIS SECTION

- CLI Quick Configuration | 179
- Configuring Switch A and Switch B | 179

CLI Quick Configuration

Switch A and Switch B

```
set protocols igmp-snooping vlan rack_1
set protocols igmp-snooping vlan v54
set multicast-snooping-options multichassis-lag-replicate-state
set protocols igmp-snooping vlan rack_1 interface ae1.0 multicast-router-interface
set protocols igmp-snooping vlan v54 interface ae1.0 multicast-router-interface
```

Configuring Switch A and Switch B

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 Switch A and Switch B:

1. Enable IGMP snooping for all VLANs.

```
[edit protocols]
user@switch# set igmp-snooping vlan rack_1
user@switch# set igmp-snooping vlan v54
```

2. Synchronize multicast states across MC-LAG peers when bridge domains are configured.

At the global level, IGMP join and leave messages are replicated from the MC-LAG interface active link to the standby link to enable faster recovery of membership information after a failover.

```
[edit multicast-snooping-options]
user@switch# set multichassis-lag-replicate-state
```

3. Configure the ICL-PL interface as a router-facing interface.

```
[edit protocols]
user@switch# set igmp-snooping vlan rack_1 interface ae1.0 multicast-router-interface
user@switch# set igmp-snooping vlan v54 interface ae1.0 multicast-router-interface
```

Switch A and Switch B

From configuration mode, confirm your configuration by entering the show protocols igmp and show multicast-snooping-options commands. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
user@switch> show protocols igmp
igmp-snooping {
    vlan rack_1 {
        interface ae1.0 {
            multicast-router-interface;
        }
    }
    vlan v54 {
        interface ae1.0 {
            multicast-router-interface;
        }
    }
}
```

```
user@switch> show multicast-snooping-options
multicast-snooping-options {
    multichassis-lag-replicate-state;
}
```

(Optional) Configuring VRRP

IN THIS SECTION

- CLI Quick Configuration | 181
 - Configuring Switch A | 182
- Configuring Switch B | **183**

NOTE: You cannot configure both VRRP and MAC address synchronization.

CLI Quick Configuration

Switch A

```
set interfaces irb unit 100 family inet address 192.168.10.3/24 vrrp-group 1 virtual-address 192.168.10.1
set interfaces irb unit 100 family inet address 192.168.10.3/24 vrrp-group 1 priority 150
set interfaces irb unit 100 family inet address 192.168.10.3/24 vrrp-group 1 accept-data
set interfaces irb unit 54 family inet address 192.168.54.2/24 vrrp-group 4 virtual-address
192.168.54.3
set interfaces irb unit 54 family inet address 192.168.54.2/24 vrrp-group 4 priority 200
```

Switch B

```
set interfaces irb unit 100 family inet address 192.168.10.2/24 vrrp-group 1 virtual-address 192.168.10.1
set interfaces irb unit 100 family inet address 192.168.10.2/24 vrrp-group 1 priority 200
set interfaces irb unit 100 family inet address 192.168.10.2/24 vrrp-group 1 accept-data
set interfaces irb unit 54 family inet address 192.168.54.1/24 vrrp-group 4 virtual-address
192.168.54.3
set interfaces irb unit 54 family inet address 192.168.54.1/24 vrrp-group 4 priority 150
```

Configuring Switch A

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 Switch A:

1. Enable VRRP on the MC-LAGs by creating an IRB interface for each MC-LAG, assign a virtual IP address that is shared between each switch in the VRRP group, and assign an individual IP address for each individual member in the VRRP group.

```
[edit interfaces]
user@switch# set irb unit 100 family inet address 192.168.10.3/24 vrrp-group 1 virtual-
address 192.168.10.1
user@switch# set irb unit 100 family inet address 192.168.10.3/24 vrrp-group 1 priority 150
user@switch# set irb unit 100 family inet address 192.168.10.3/24 vrrp-group 1 accept-data
```

2. Enable VRRP on the MC-LAGs by creating an IRB interface for each MC-LAG, assign a virtual IP address that is shared between each switch in the VRRP group, and assign an individual IP address for each individual member in the VRRP group.

```
[edit interfaces]
user@switch# set irb unit 54 family inet address 192.168.54.2/24 vrrp-group 4 virtual-address
192.168.54.3
user@switch# set irb unit 54 family inet address 192.168.54.2/24 vrrp-group 4 priority 200
```

Switch A

From configuration mode, confirm your configuration by entering the show interfaces irb unit 100 family inet address 192.168.10.3/24 vrrp-group and show interfaces irb unit 100 family inet address 192.168.54.2/24 vrrp-group commands. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
user@switch> show interfaces irb unit 100 family inet address 192.168.10.3/24 vrrp-group
vrrp-group 1 {
   virtual-address 192.168.10.1;
   priority 150;
```

```
accept-data;
}
```

```
user@switch> show interfaces irb unit 100 family inet address 192.168.54.2/24 vrrp-group
vrrp-group 4 {
    virtual-address 192.168.54.3;
    priority 150;
}
```

Configuring Switch B

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 Switch A:

1. Enable VRRP on the MC-LAGs by creating an IRB interface for each MC-LAG, assign a virtual IP address that is shared between each switch in the VRRP group, and assign an individual IP address for each individual member in the VRRP group.

```
[edit interfaces]
user@switch# set irb unit 100 family inet address 192.168.10.2/24 vrrp-group 1 virtual-
address 192.168.10.1
user@switch# set irb unit 100 family inet address 192.168.10.2/24 vrrp-group 1 priority 150
user@switch# set irb unit 100 family inet address 192.168.10.2/24 vrrp-group 1 accept-data
```

2. Enable VRRP on the MC-LAGs by creating an IRB interface for each MC-LAG, assign a virtual IP address that is shared between each switch in the VRRP group, and assign an individual IP address for each individual member in the VRRP group.

```
[edit interfaces]
user@switch# set irb unit 54 family inet address 192.168.54.1/24 vrrp-group 4 virtual-address
192.168.54.3
user@switch# set irb unit 54 family inet address 192.168.54.1/24 vrrp-group 4 priority 150
```

Switch B

From configuration mode, confirm your configuration by entering the show protocols rstp commands. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
user@switch> show interfaces irb unit 100 family inet address 192.168.10.2/24 vrrp-group
vrrp-group 1 {
    virtual-address 192.168.10.1;
    priority 200;
    accept-data;
}
```

```
user@switch> show interfaces irb unit 100 family inet address 192.168.54.1/24 vrrp-group
vrrp-group 4 {
    virtual-address 192.168.54.3;
    priority 150;
}
```

(Optional) Configuring MAC Address Synchronization

IN THIS SECTION

- CLI Quick Configuration | 185
- Configuring Switch A and Switch B | 185

NOTE: You cannot configure both MAC synchronization and VRRP.

You must configure the same IP address on the IRB interface in the VLAN on both MC-LAG peers.

CLI Quick Configuration

Switch A and Switch B

```
set vlans v100 mcae-mac-synchronize
set vlans v54 mcae-mac-synchronize
```

Configuring Switch A and Switch B

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 Switch A:

1. 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
[edit]
user@switch# set vlans v54 mcae-mac-synchronize
```

Switch A and Switch B

From configuration mode, confirm your configuration by entering the show vlans v100 and show vlans v54 commands. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
user@switch> show vlans v100
v100 {
    vlan-id 100;
    l3-interface irb.100;
    mcae-mac-synchronize;
}
```

```
user@switch> show vlans v54
v54 {
```

```
vlan-id 54;
13-interface irb.54;
mcae-mac-synchronize;
}
```

(Optional) Configuring OSPF

IN THIS SECTION

- CLI Quick Configuration | 186
- Configuring Switch A and Switch B | 186

CLI Quick Configuration

Switch A and Switch B

```
set protocols ospf area 0.0.0.0 interface lo0.0
set protocols ospf area 0.0.0.0 interface ae0.0
set protocols ospf area 0.0.0.0 interface irb.54
set protocols ospf area 0.0.0.0 interface irb.100
```

Configuring Switch A and Switch B

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 Switch A and Switch B:

1. Configure an OSPF area.

```
[edit protocols]
user@switch# set ospf area 0.0.0.0 interface lo0.0
user@switch# set ospf area 0.0.0.0 interface ae0.0
user@switch# set ospf area 0.0.0.0 interface irb.54
user@switch# set ospf area 0.0.0.0 interface irb.100
```

Switch A and Switch B

From configuration mode, confirm your configuration by entering the show protocols ospf commands. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
user@switch> show protocols ospf
ospf {
    area 0.0.0.0 {
        interface lo0.0;
        interface ae0.0;
        interface irb.54;
        interface irb.100;
    }
}
```

(Optional) Configuring PIM

IN THIS SECTION

- CLI Quick Configuration | 187
- Configuring Switch A | 188
- Configuring Switch B | 189

CLI Quick Configuration

Switch A

```
set protocols pim interface irb.54
set protocols pim interface irb.100
set protocols pim interface lo0.0
set protocols pim rp bootstrap-priority 150
set protocols pim rp local address 192.168.39.1
```

Switch B

```
set protocols pim interface irb.54
set protocols pim interface irb.100
set protocols pim interface lo0.0
set protocols pim rp bootstrap-priority 200
set protocols pim rp local address 192.168.39.2
```

Configuring Switch A

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 Switch A:

1. Configure Protocol Independent Multicast (PIM) as the multicast protocol.

```
[edit protocols]
user@switch# set pim interface irb.54
user@switch# set pim interface irb.100
```

2. Configure the loopback interface.

```
[edit protocols]
user@switch# set pim interface lo0.0
```

3. Configure the switch as a secondary rendezvous point (RP).

A lower priority setting indicates that the secondary RP is in a bootstrap configuration.

```
[edit protocols]
user@switch# set pim rp bootstrap-priority 150
user@switch# set pim rp local address 192.168.39.1
```

Switch A

From configuration mode, confirm your configuration by entering the show protocols pim commands. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
user@switch> show protocols pim
pim {
    rp {
        bootstrap-priority 150;
        local {
            address 192.168.39.1;
        }
    }
    interface irb.54;
    interface irb.100;
    interface lo0.0;
}
```

Configuring Switch B

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 Switch A:

1. Configure Protocol Independent Multicast (PIM) as the multicast protocol.

```
[edit protocols]
user@switch# set pim interface irb.54
user@switch# set pim interface irb.100
```

2. Configure the loopback interface.

```
[edit protocols]
user@switch# set pim interface lo0.0
```

3. Configure the switch as a secondary rendezvous point (RP).

A lower priority setting indicates that the secondary RP is in a bootstrap configuration.

```
[edit protocols]
user@switch# set pim rp bootstrap-priority 200
user@switch# set pim rp local address 192.168.39.2
```

Switch B

From configuration mode, confirm your configuration by entering the show protocols pim commands. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
user@switch> show protocols pim
pim {
    rp {
        bootstrap-priority 200;
        local {
            address 192.168.39.2;
        }
    }
    interface irb.54;
    interface irb.100;
    interface lo0.0;
}
```

(Optional) Configuring DHCP Relay

IN THIS SECTION

- CLI Quick Configuration | 191
- Configuring Switch A and Switch B | 191

CLI Quick Configuration

Switch A and Switch B

```
set forwarding-options dhcp-relay forward-snooped-clients all-interfaces
set forwarding-options dhcp-relay overrides allow-snooped-clients
set forwarding-options dhcp-relay server-group GVP-DHCP 10.105.5.202
set forwarding-options dhcp-relay active-server-group GVP-DHCP
set forwarding-options dhcp-relay route-suppression destination
set forwarding-options dhcp-relay group Floor1 interface irb.100
set forwarding-options dhcp-relay relay-option-82 circuit-id use-interface-description device
```

Configuring Switch A and Switch B

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 Switch A and Switch B:

1. Configure forward snooped unicast packets on all interfaces.

```
[edit fowarding-options]
user@switch# set dhcp-relay forward-snooped-clients all-interfaces
```

2. Create a binding entry to snoop unicast clients.

```
[edit forwarding-options]
user@switch# set dhcp-relay overrides allow-snooped-clients
```

3. Create a DHCP server group.

```
[edit forwarding-options]
user@switch# set dhcp-relay server-group GVP-DHCP 10.105.5.202
```

4. Apply a DHCP relay agent configuration to the named group of DHCP server addresses.

```
[edit forwarding-options]
user@switch# set dhcp-relay active-server-group GVP-DHCP
```

5. Configure the relay agent to suppress the installation of ARP and route entries for corresponding client binding.

```
[edit forwarding-options]
user@switch# set dhcp-relay route-suppression destination
```

6. Create a DHCP relay group that includes at least one interface.

DHCP runs on the interfaces defined in the DHCP groups.

```
[edit forwarding-options]
user@switch# set dhcp-relay group Floor1 interface irb.100
```

7. Configure DHCP relay with option 82.

```
[edit forwarding-options]
user@switch# set dhcp-relay relay-option-82 circuit-id use-interface-description device
```

Switch A and Switch B

From configuration mode, confirm your configuration by entering the show protocols pim commands. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
user@switch> show forwarding-options dhcp-relay
dhcp-relay {
    forward-snooped-clients all-interfaces;
    overrides {
        allow-snooped-clients;
    }
    relay-option-82 {
        circuit-id {
            use-interface-description device;
    }
}
```

```
}
}
server-group {
    GVP-DHCP {
        10.105.5.202;
    }
}

active-server-group GVP-DHCP;
route-suppression {
    destination;
}
group Floor1 {
    interface irb.100;
}
```

Verification

IN THIS SECTION

- Verifying ICCP on MC-LAG | 193
- Verifying LACP on MC-LAG | 195
- Verifying Aggregated Ethernet Interfaces in MC-LAG | 197
- Verifying MAC Learning on MC-LAG | 199
- Verifying VRRP in MC-LAG | 201
- Verifying OSPF on MC-LAG | 202

Confirm that the configuration is working properly.

Verifying ICCP on MC-LAG

Purpose

Verify that ICCP is running on each device in the MC-LAG.

Action

1. Verify that ICCP is running on Switch A.

```
root@EX92000-A> show iccp
Redundancy Group Information for peer 192.168.39.2
 TCP Connection
                   : Established
 Liveliness Detection : Up
Backup liveness peer status: Up
                      Status
 Redundancy Group ID
   1
                              Up
Client Application: 12ald_iccpd_client
  Redundancy Group IDs Joined: 1
Client Application: lacpd
 Redundancy Group IDs Joined: 1
Client Application: MCSNOOPD
  Redundancy Group IDs Joined: 1
```

2. Verify that ICCP is running on Switch B.

```
root@EX9200-B> show iccp

Redundancy Group Information for peer 192.168.39.1

TCP Connection : Established
Liveliness Detection : Up

Backup liveness peer status: Up
Redundancy Group ID Status

1 Up

Client Application: lacpd
Redundancy Group IDs Joined: 1

Client Application: l2ald_iccpd_client
Redundancy Group IDs Joined: 1
```

```
Client Application: MCSNOOPD

Redundancy Group IDs Joined: 1
```

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 LACP on MC-LAG

Purpose

Verify that LACP is working properly on each device in the MC-LAG.

Action

1. Verify that the LACP interfaces are up and running on Switch A.

```
root@EX9200-A> show lacp interfaces
Aggregated interface: ae0
   LACP state:
                    Role
                                 Def Dist Col Syn Aggr Timeout Activity
                           Exp
     et-1/0/0
                   Actor
                                                      Yes
                                                                     Active
                            No
                                  No
                                      Yes Yes Yes
                                                             Fast
     et-1/0/0
                 Partner
                                                                     Active
                                      Yes Yes Yes
                                                      Yes
                                                             Fast
                            No
                                  No
     et-1/0/1
                   Actor
                                                                     Active
                            No
                                  No
                                      Yes Yes Yes
                                                      Yes
                                                             Fast
     et-1/0/1
                                      Yes Yes Yes
                 Partner
                            No
                                  No
                                                      Yes
                                                             Fast
                                                                     Active
   LACP protocol:
                        Receive State Transmit State
                                                             Mux State
     et-1/0/0
                              Current
                                       Fast periodic Collecting distributing
     et-1/0/1
                                       Fast periodic Collecting distributing
                              Current
Aggregated interface: ae1
   LACP state:
                    Role
                           Exp
                                 Def Dist Col Syn Aggr Timeout Activity
     xe-2/0/3
                   Actor
                            No
                                  No
                                      Yes Yes Yes
                                                      Yes
                                                             Fast
                                                                     Active
     xe-2/0/3
                 Partner
                                  No Yes Yes Yes
                                                      Yes
                                                             Fast
                                                                     Active
                            No
     xe-2/0/4
                   Actor
                            No
                                      Yes Yes
                                               Yes
                                                      Yes
                                                             Fast
                                                                     Active
     xe-2/0/4
                 Partner
                            Nο
                                  No
                                      Yes Yes Yes
                                                      Yes
                                                             Fast
                                                                     Active
   LACP protocol:
                        Receive State Transmit State
                                                             Mux State
     xe-2/0/3
                              Current
                                       Fast periodic Collecting distributing
     xe-2/0/4
                                       Fast periodic Collecting distributing
                              Current
Aggregated interface: ae3
   LACP state:
                           Exp Def Dist Col Syn Aggr Timeout Activity
                    Role
```

```
xe-2/0/1
                     Actor
                              No
                                    No
                                         Yes Yes
                                                  Yes
                                                         Yes
                                                                 Fast
                                                                          Active
      xe-2/0/1
                   Partner
                              No
                                    No
                                         Yes
                                              Yes
                                                   Yes
                                                         Yes
                                                                 Fast
                                                                         Passive
      xe-2/0/2
                     Actor
                                         Yes Yes
                                                  Yes
                                                         Yes
                                                                          Active
                              No
                                    No
                                                                 Fast
      xe-2/0/2
                   Partner
                              No
                                    No
                                         Yes Yes Yes
                                                         Yes
                                                                 Fast
                                                                         Passive
    LACP protocol:
                          Receive State
                                        Transmit State
                                                                 Mux State
      xe-2/0/1
                                Current
                                          Fast periodic Collecting distributing
      xe-2/0/2
                                Current
                                          Fast periodic Collecting distributing
Aggregated interface: ae4
    LACP state:
                      Role
                             Exp
                                   Def Dist Col Syn Aggr Timeout Activity
     xe-2/0/0
                     Actor
                                    No
                                         Yes Yes
                                                         Yes
                                                                 Fast
                              No
                                                   Yes
                                                                          Active
     xe-2/0/0
                   Partner
                              No
                                    No
                                         Yes Yes Yes
                                                         Yes
                                                                 Fast
                                                                          Active
   LACP protocol:
                          Receive State Transmit State
                                                                 Mux State
      xe-2/0/0
                                          Fast periodic Collecting distributing
                                Current
```

2. Verify that the LACP interfaces are up and running on Switch B.

```
root@EX9200-B> show lacp interfaces
Aggregated interface: ae0
   LACP state:
                      Role
                            Exp
                                   Def
                                       Dist Col Syn
                                                        Aggr Timeout Activity
     et-1/0/0
                     Actor
                             No
                                    No
                                         Yes Yes
                                                   Yes
                                                         Yes
                                                                 Fast
                                                                         Active
      et-1/0/0
                  Partner
                             No
                                    No
                                         Yes
                                             Yes
                                                   Yes
                                                         Yes
                                                                 Fast
                                                                         Active
     et-1/0/1
                     Actor
                                         Yes Yes
                                                  Yes
                                                         Yes
                                                                         Active
                              No
                                    No
                                                                 Fast
      et-1/0/1
                  Partner
                              No
                                    No
                                         Yes Yes Yes
                                                         Yes
                                                                 Fast
                                                                         Active
   LACP protocol:
                          Receive State Transmit State
                                                                 Mux State
      et-1/0/0
                                Current
                                          Fast periodic Collecting distributing
      et-1/0/1
                                          Fast periodic Collecting distributing
                                Current
Aggregated interface: ae1
   LACP state:
                      Role
                            Exp
                                   Def Dist Col Syn Aggr Timeout Activity
      xe-2/0/3
                     Actor
                                             Yes
                                                                         Active
                             No
                                    No
                                         Yes
                                                   Yes
                                                         Yes
                                                                 Fast
      xe-2/0/3
                                                                         Active
                   Partner
                              No
                                    No
                                         Yes Yes
                                                   Yes
                                                         Yes
                                                                 Fast
     xe-2/0/4
                     Actor
                                         Yes Yes
                                                  Yes
                                                         Yes
                                                                         Active
                              No
                                    No
                                                                 Fast
     xe-2/0/4
                                         Yes Yes Yes
                                                                 Fast
                  Partner
                             No
                                    No
                                                         Yes
                                                                         Active
   LACP protocol:
                          Receive State Transmit State
                                                                 Mux State
      xe-2/0/3
                                          Fast periodic Collecting distributing
                                Current
      xe-2/0/4
                                          Fast periodic Collecting distributing
                                Current
Aggregated interface: ae2
   LACP state:
                      Role
                            Exp
                                   Def
                                       Dist Col Syn Aggr Timeout Activity
      et-1/2/0
                     Actor
                             No
                                         Yes Yes Yes
                                                         Yes
                                                                 Fast
                                                                         Active
```

```
Fast Passive
     et-1/2/0
                 Partner
                            No
                                 No
                                      Yes Yes Yes Yes
   LACP protocol:
                        Receive State Transmit State
                                                             Mux State
     et-1/2/0
                                       Fast periodic Collecting distributing
                              Current
Aggregated interface: ae4
   LACP state:
                    Role Exp
                                 Def Dist Col Syn Aggr Timeout Activity
     xe-2/0/0
                   Actor
                            No
                                 No
                                      Yes Yes Yes
                                                     Yes
                                                             Fast
                                                                     Active
     xe-2/0/0
                 Partner
                            No
                                 No
                                      Yes Yes Yes
                                                     Yes
                                                             Fast
                                                                     Active
   LACP protocol:
                        Receive State Transmit State
                                                             Mux State
     xe-2/0/0
                              Current Fast periodic Collecting distributing
```

Meaning

This output means that both devices and all related interfaces are properly participating in LACP negotiations.

Verifying Aggregated Ethernet Interfaces in MC-LAG

Purpose

Verify that all of the ae interfaces are configured properly in the MC&LAG.

Action

1. Verify the ae interfaces on Switch A.

```
user@EX9200-A> show interfaces mc-ae
Member Link
Current State Machine's State: mcae active state
Local Status
                             : active
Local State
                             : up
Peer Status
                             : active
 Peer State
                             : up
     Logical Interface
                             : ae2.0
    Topology Type
                             : bridge
    Local State
                             : up
    Peer State
                             : up
    Peer Ip/MCP/State
                             : 192.168.39.2 ae1.0 up
```

Member Link : ae4

Current State Machine's State: mcae active state

Local Status : active
Local State : up
Peer Status : active
Peer State : up
Logical Interface : ae4.0
Topology Type : bridge
Local State : up

Peer State : up

Peer Ip/MCP/State : 192.168.39.2 ae1.0 up

2. Verify the ae interfaces on Switch B.

root@EX9200-B> show interface mc-ae

Member Link : ae2

Current State Machine's State: mcae active state

Local Status : active
Local State : up
Peer Status : active
Peer State : up
Logical Interface : ae2.0
Topology Type : bridge
Local State : up

Peer Ip/MCP/State : 192.168.39.1 ae1.0 up

Member Link : ae4

Peer State

Current State Machine's State: mcae active state

Local Status : active
Local State : up
Peer Status : active
Peer State : up
Logical Interface : ae4.0
Topology Type : bridge

Local State : up
Peer State : up

```
Peer Ip/MCP/State : 192.168.39.1 ae1.0 up
```

Meaning

This output means that the mc-ae interfaces on each device are up and active.

Verifying MAC Learning on MC-LAG

Purpose

Verify that MAC learning between devices is happening in the MC-LAG.

Action

1. Show Ethernet switching table in Switch A.

```
root@EX9200-A> show ethernet-switching table
MAC flags (S - static MAC, D - dynamic MAC, L - locally learned, P - Persistent static
           SE - statistics enabled, NM - non configured MAC, R - remote PE MAC)
Ethernet switching table : 68 entries, 68 learned
Routing instance : default-switch
    Vlan
                       MAC
                                           MAC
                                                       Age
                                                              Logical
                                                              interface
    name
                       address
                                           flags
                                                             ae4.0
    dmzuplink
                       00:00:5e:00:01:ba DL
                                                             ae4.0
    dmzuplink
                       00:10:db:bc:f5:9d DR
    dmzuplink
                       00:10:db:ff:10:01
                                                             ae3.0
                                           DL
    dmzuplink
                       00:19:e2:57:33:81
                                           DR
                                                             ae4.0
    dmzuplink
                       00:26:88:92:ef:1d
                                                             ae4.0
                                          DR
    dmzuplink
                       28:8a:1c:74:fb:07
                                                             ae4.0
                                           DR
    dmzuplink
                        28:8a:1c:75:05:1f
                                                             ae4.0
    dmzuplink
                       28:c0:da:6a:1d:2a DR
                                                             ae4.0
                       2c:21:72:7d:40:01 DL
                                                             ae4.0
    dmzuplink
                        3c:8a:b0:77:a9:d6 DR
                                                             ae4.0
    dmzuplink
    dmzuplink
                        5c:5e:ab:0e:cd:e0
                                           DL
                                                             ae4.0
    dmzuplink
                        84:18:88:8d:9d:2a
                                           DL
                                                             ae4.0
MAC flags (S - static MAC, D - dynamic MAC, L - locally learned, P - Persistent static
```

SE - statistics enabled, NM - non configured MAC, R - remote PE MAC) Ethernet switching table : 68 entries, 68 learned Routing instance : default-switch Vlan MAC MAC Age Logical name address flags interface ae2.0 rack_1 00:50:56:9b:01:57 DR rack_1 00:50:56:9b:09:95 ae2.0 rack_1 00:50:56:9b:15:2e ae2.0 DL rack_1 00:50:56:9b:20:44 DL ae2.0 rack_1 00:50:56:9b:20:a7 DL ae2.0 rack_1 00:50:56:9b:22:a8 DR ae2.0 rack_1 00:50:56:9b:38:01 DL ae2.0 rack_1 00:50:56:9b:66:dc DL ae2.0 00:50:56:9b:75:60 DR rack_1 ae2.0 MAC flags (S - static MAC, D - dynamic MAC, L - locally learned, P - Persistent static SE - statistics enabled, NM - non configured MAC, R - remote PE MAC) Ethernet switching table : 68 entries, 68 learned Routing instance : default-switch Vlan MAC MAC Age Logical address name flags interface 80:71:1f:c1:85:f0 v54 DL ae4.0

2. Show Ethernet switching table in Switch B.

root@EX9200-B> show ethernet-switching table MAC flags (S - static MAC, D - dynamic MAC, L - locally learned, P - Persistent static SE - statistics enabled, NM - non configured MAC, R - remote PE MAC) Ethernet switching table : 66 entries, 66 learned Routing instance : default-switch Vlan MAC MAC Age Logical address interface name flags rack_1 00:50:56:9b:01:57 DL ae2.0 rack_1 00:50:56:9b:09:95 DR ae2.0

```
rack_1
                        00:50:56:9b:15:2e
                                                              ae2.0
                                           DR
                                                              ae2.0
    rack_1
                        00:50:56:9b:20:44
                                           DR
    rack_1
                        00:50:56:9b:20:a7
                                                              ae2.0
                                           DR
    rack_1
                        00:50:56:9b:22:a8
                                           DL
                                                              ae2.0
    rack_1
                        00:50:56:9b:38:01
                                                              ae2.0
    rack_1
                        00:50:56:9b:66:dc DR
                                                              ae2.0
    rack_1
                        00:50:56:9b:75:60
                                           DL
                                                              ae2.0
MAC flags (S - static MAC, D - dynamic MAC, L - locally learned, P - Persistent static
           SE - statistics enabled, NM - non configured MAC, R - remote PE MAC)
Ethernet switching table : 66 entries, 66 learned
Routing instance : default-switch
    Vlan
                        MAC
                                            MAC
                                                               Logical
                                                        Age
    name
                        address
                                            flags
                                                               interface
MAC flags (S - static MAC, D - dynamic MAC, L - locally learned, P - Persistent static
           SE - statistics enabled, NM - non configured MAC, R - remote PE MAC)
Ethernet switching table : 66 entries, 66 learned
Routing instance : default-switch
    Vlan
                        MAC
                                            MAC
                                                        Age
                                                               Logical
                                                               interface
    name
                        address
                                            flags
                        80:71:1f:c1:85:f0
                                           DR
                                                              ae4.0
    v54
```

Meaning

This output means that the MAC addresses are properly learned within the shared VLANs defined in the MC-LAG. This includes IRB interfaces to define the MC-LAG as well as the ICL interfaces used to configure VRRP.

Verifying VRRP in MC-LAG

Purpose

Verify that VRRP is up and active between the devices in the MC-LAG.

Action

1. Confirm that VRRP is up and active on Switch A.

```
root@EX9200-A> show vrrp
Interface
             State
                                 VR state VR Mode
                                                                   Address
                         Group
                                                   Timer
                                                            Type
irb.54
             up
                                 backup Active
                                                     D 3.090 lcl
                                                                     192.168.54.1
                                                                     192.168.54.3
                                                              vip
                                                                     192.168.54.2
                                                              mas
irb.100
                                backup Active
                                                     D 2.655 lcl
                                                                     192.168.10.3
             up
                                                                     192.168.10.1
                                                              vip
                                                                     192.168.10.2
                                                              mas
```

In this example, Switch A is the backup VRRP member.

2. Confirm that VRRP is up and active on Switch B.

root@EX9200-	B> show vr	rp					
Interface	State	Group	VR state	VR Mode	Time	r Type	Address
irb.54	up	4	master	Active	Α	0.900 lcl	192.168.54.2
						vip	192.168.54.3
irb.100	up	1	master	Active	Α	0.175 lcl	192.168.10.2
						vip	192.168.10.1

In this example, Switch B is the primary VRRP member.

Meaning

This output means that VRRP is up and running properly.

Verifying OSPF on MC-LAG

Purpose

Verify that OSPF is properly up and running with MC-LAG.

Action

1. Show OSPF neighbors on Switch A.

root@EX9200-A	> show ospf neighbo	r		
Address	Interface	State	ID	Pri Dead
192.168.90.2	ae0.0	Full	192.168.39.2	128 35
192.168.10.2	irb.100	Full	192.168.39.2	128 33
192.168.54.2	irb.54	Full	192.168.39.2	128 38

2. Show OSPF routing table on Switch A.

root@EX9200-A> sh Topology default	·			
Prefix	Path Route	NH	Metric NextHop	Nexthop
	Type Type	Туре	Interface	Address/LSP
192.168.39.2	Intra Router	IP	1 ae0.0	192.168.90.2
			irb.100	192.168.10.2
			irb.54	192.168.54.2
192.168.39.1/32	Intra Network	IP	0 100.0	
192.168.39.2/32	Intra Network	IP	1 ae0.0	192.168.90.2
			irb.100	192.168.10.2
			irb.54	192.168.54.2
192.168.10.0/24	Intra Network	IP	1 irb.100	
192.168.54.0/24	Intra Network	IP	1 irb.54	
192.168.90.0/24	Intra Network	ΙP	1 ae0.0	

3. Show OSPF neighbors on Switch B.

root@EX9200-E	3> show ospf neighbor			
Address	Interface	State	ID	Pri Dead
192.168.90.1	ae0.0	Full	192.168.39.1	128 32
192.168.10.3	irb.100	Full	192.168.39.1	128 34
192.168.54.1	irb.54	Full	192.168.39.1	128 37

4. Show OSPF routing table on Switch B.

root@EX9200-B> sh Topology default	·			
Prefix	Path Route	NH	Metric NextHop	Nexthop
	Type Type	Type	Interface	Address/LSP
192.168.39.1	Intra Router	IP	1 ae0.0	192.168.90.1
			irb.100	192.168.10.3
			irb.54	192.168.54.1
192.168.39.1/32	Intra Network	IP	1 ae0.0	192.168.90.1
			irb.100	192.168.10.3
			irb.54	192.168.54.1
192.168.39.2/32	Intra Network	IP	0 100.0	
192.168.10.0/24	Intra Network	ΙP	1 irb.100	
192.168.54.0/24	Intra Network	ΙP	1 irb.54	
192.168.90.0/24	Intra Network	IP	1 ae0.0	

Example: Configure Optional Features For Multichassis Link Aggregation

IN THIS SECTION

- (Optional) Configuring RSTP | 205
- (Optional) Configuring IGMP Snooping | 207
- (Optional) Configuring VRRP | 208
- (Optional) Configuring MAC Address Synchronization | 210
- (Optional) Configuring OSPF | 212
- (Optional) Configuring PIM | 213
- (Optional) Configuring DHCP Relay | 217

This example shows how to configure optional features that can be combined with MC LAG.

(Optional) Configuring RSTP

IN THIS SECTION

- CLI Quick Configuration | 205
- Configuring QFX1 and QFX2 | 205

CLI Quick Configuration

QFX1 and QFX2

```
set protocols rstp interface xe-0/0/3 edge
set protocols rstp interface ae0 disable
set protocols rstp interface all mode point-to-point
set protocols rstp bpdu-block-on-edge
```

Configuring QFX1 and QFX2

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. Configure the MC-LAG interfaces as edge ports on QFX1 and QFX2.

```
[edit]
user@switch# set protocols rstp interface xe-0/0/3 edge
```

2. Disable RSTP on the ICL-PL interfaces on QFX1 and QFX2:

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

3. Enable RSTP globally on all interfaces on QFX1 and QFX2.

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

4. Enable BPDU blocking on all interfaces except for the ICL-PL interfaces on QFX1 and QFX2.

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

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

QFX1 and QFX2

From configuration mode, confirm your configuration by entering the show protocols rstp commands. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
user@switch> show protocols rstp
rstp {
    interface xe-0/0/3 {
        edge;
    }
    interface ae0 {
        disable;
    }
    interface all {
        mode point-to-point;
        }
    bpdu-block-on-edge;
```

(Optional) Configuring IGMP Snooping

IN THIS SECTION

- CLI Quick Configuration | 207
- Configuring QFX1 and QFX2 | 207

CLI Quick Configuration

QFX1 and QFX2

```
set protocols igmp-snooping vlan all
```

Configuring QFX1 and QFX2

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 QFX1 and QFX2:

1. Enable IGMP snooping for all VLANs.

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

QFX1 and QFX2

From configuration mode, confirm your configuration by entering the show protocols igmp and show multicast-snooping-options commands. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
user@switch> show protocols igmp
igmp-snooping {
```

```
vlan all;
}
```

(Optional) Configuring VRRP

IN THIS SECTION

- CLI Quick Configuration | 208
- Configuring QFX1 | 208
- Configuring QFX2 | 209

CLI Quick Configuration

QFX1

```
set interfaces irb unit 50 family inet address 10.50.1.1/30 vrrp-group 1 virtual-address 10.50.1.10 set interfaces irb unit 50 family inet address 10.50.1.1/30 vrrp-group 1 priority 200 set interfaces irb unit 50 family inet address 10.50.1.1/30 vrrp-group 1 accept-data
```

QFX2

```
set interfaces irb unit 50 family inet address 10.50.1.2/30 vrrp-group 1 virtual-address 10.50.1.10 set interfaces irb unit 50 family inet address 10.50.1.2/30 vrrp-group 1 priority 150 set interfaces irb unit 50 family inet address 10.50.1.2/30 vrrp-group 1 accept-data
```

Configuring QFX1

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 QFX1:

1. Enable VRRP on the MC-LAGs by creating an IRB interface for each MC-LAG, assign a virtual IP address that is shared between each switch in the VRRP group, and assign an individual IP address for each individual member in the VRRP group.

```
[edit interfaces]
user@switch# set interfaces irb unit 50 family inet address 10.50.1.1/30 vrrp-group 1 virtual-
address 10.50.1.10
user@switch# set interfaces irb unit 50 family inet address 10.50.1.1/30 vrrp-group 1
priority 200
user@switch# set interfaces irb unit 50 family inet address 10.50.1.2/30 vrrp-group 1 accept-
data
```

QFX1

From configuration mode, confirm your configuration by entering the show interfaces irb unit 50 family inet address 10.50.1.1/30 vrrp-group command. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
user@switch> show interfaces irb unit 50 family inet address 10.50.1.1/30 vrrp-group
vrrp-group 1 {
    virtual-address 10.50.1.10;
    priority 200;
    accept-data;
}
```

Configuring QFX2

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 QFX1:

1. Enable VRRP on the MC-LAGs by creating an IRB interface for each MC-LAG, assign a virtual IP address that is shared between each switch in the VRRP group, and assign an individual IP address for each individual member in the VRRP group.

```
[edit interfaces]
user@switch# set interfaces irb unit 50 family inet address 10.50.1.2/30 vrrp-group 1 virtual-
```

```
address 10.50.1.10

user@switch# set interfaces irb unit 50 family inet address 10.50.1.2/30 vrrp-group 1

priority 150

user@switch# set interfaces irb unit 50 family inet address 10.50.1.2/30 vrrp-group 1 accept-data
```

QFX2

From configuration mode, confirm your configuration by entering the show interfaces irb unit 500 family inet address 10.1.1.10/8 vrrp-group commands. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
user@switch> show interfaces irb unit 50 family inet address 10.50.1.2/30 vrrp-group
vrrp-group 1 {
    virtual-address 10.50.1.10;
    priority 150;
    accept-data;
}
```

(Optional) Configuring MAC Address Synchronization

IN THIS SECTION

- CLI Quick Configuration | 211
- Configuring QFX1 and QFX2 | 211

NOTE: You cannot configure both MAC synchronization and VRRP.

You must configure the same IP address on the IRB interface in the VLAN on both MC-LAG peers.

CLI Quick Configuration

QFX1 and QFX2

```
set vlans v10 mcae-mac-synchronize
```

Configuring QFX1 and QFX2

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 QFX1:

1. Configure MAC address synchronization in the MC-LAG VLAN on both QFX1 and QFX2.

```
[edit]
user@switch# set vlans v10 mcae-mac-synchronize
```

QFX1 and QFX2

From configuration mode, confirm your configuration by entering the show vlans v10 command. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
user@switch> show vlans v10
v10 {
    vlan-id 10;
    l3-interface irb.10;
    mcae-mac-synchronize;
}
```

(Optional) Configuring OSPF

IN THIS SECTION

- CLI Quick Configuration | 212
- Configuring QFX1 and QFX2 | 212

CLI Quick Configuration

QFX1, QFX2, and QFX3

```
set protocols ospf area 0.0.0.0 interface irb.10
```

Configuring QFX1 and QFX2

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 QFX1 and QFX2:

1. Configure an OSPF area on QFX1, QFX2.

```
[edit protocols]
user@switch# set ospf area 0.0.0.0 interface irb.10
```

QFX1 and QFX2

From configuration mode, confirm your configuration by entering the show protocols ospf commands. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
user@switch> show protocols ospf
ospf {
   area 0.0.0.0 {
```

```
interface irb.10 {
    }
}
```

(Optional) Configuring PIM

IN THIS SECTION

- CLI Quick Configuration | 213
- Configuring QFX1 and QFX2 | 214

CLI Quick Configuration

QFX1

```
set protocols pim rp static address 10.0.0.3 group-ranges 224.0.0.0/4
set protocols pim interface irb.50 priority 200
set protocols pim interface irb.50 dual-dr
set protocols pim interface irb.50 family inet bfd-liveness-detection minimum-receive-interval
700
set protocols pim interface irb.50 family inet bfd-liveness-detection transmit-interval minimum-
interval 350
set protocols pim interface irb.50 family inet bfd-liveness-detection transmit-interval
threshold 500
```

QFX2

```
set protocols pim rp static address 10.0.0.3 group-ranges 224.0.0.0/4
set protocols pim interface irb.50 priority 500
set protocols pim interface irb.50 dual-dr
set protocols pim interface irb.50 family inet bfd-liveness-detection minimum-receive-interval
700
set protocols pim interface irb.50 family inet bfd-liveness-detection transmit-interval minimum-
interval 350
```

set protocols pim interface irb.50 family inet bfd-liveness-detection transmit-interval threshold 500

Configuring QFX1 and QFX2

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 PIM as the multicast protocol on QFX1:

1. Configure a static rendezvous point (RP) address on QFX1 and QFX2.

```
[edit protocols pim]
user@switch# set rp static address 10.0.0.3
```

2. Configure the address ranges of the multicast groups for which QFX1 and QFX2 can be a rendezvous point (RP).

```
[edit protocols pim rp static address 10.0.0.3]
user@switch# set group-ranges 224.0.0.0/4
```

3. Enable PIM on the VLAN interfaces for the MC-LAGs on QFX1 and QFX2.

```
[edit protocols pim]
user@switch# set interface irb.500 dual-dr
```

4. Configure each PIM interface's priority for being selected as the designated router (DR) on QFX1 and QFX2.

An interface with a higher priority value has a higher probability of being selected as the DR.

```
QFX1:
[edit protocols pim]
user@switch# set interface irb.500 priority 200
```

```
QFX2:
[edit protocols pim]
user@switch# set interface irb.500 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 QFX1 and QFX2.

```
[edit protocols pim]
user@switch# set interface irb.500 family inet bfd-liveness-detection minimum-receive-
interval 700
user@switch# set interface irb.500 family inet bfd-liveness-detection transmit-interval
minimum-interval 350
user@switch# set interface irb.500 family inet bfd-liveness-detection transmit-interval
threshold 500
```

Results

QFX1

From configuration mode, confirm your configuration by entering the show protocols pim commands. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
}
}
interface irb.50 {
    family inet {
        bfd-liveness-detection {
            minimum-receive-interval 700;
            transmit-interval {
                minimum-interval 350;
                threshold 500;
        }
    }
    priority 200;
    dual-dr;
}
```

QFX2

From configuration mode, confirm your configuration by entering the show protocols pim command. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
user@switch> show protocols pim
rp {
    static {
        address 10.0.0.3 {
            group-ranges {
                 224.0.0.0/4;
            }
        }
    }
}
interface irb.50 {
    family inet {
        bfd\text{-liveness-detection }\{
            minimum-receive-interval 700;
            transmit-interval {
                 minimum-interval 350;
                 threshold 500;
            }
```

```
}
priority 500;
dual-dr;
}
```

(Optional) Configuring DHCP Relay

IN THIS SECTION

- CLI Quick Configuration | 217
- Configuring QFX1 and QFX2 | 217

CLI Quick Configuration

QFX1 and QFX2

```
set forwarding-options dhcp-relay forward-snooped-clients all-interfaces
set forwarding-options dhcp-relay overrides allow-snooped-clients
set forwarding-options dhcp-relay server-group GVP-DHCP 10.105.5.202
set forwarding-options dhcp-relay active-server-group GVP-DHCP
set forwarding-options dhcp-relay route-suppression destination
set forwarding-options dhcp-relay group Floor1 interface irb.500
set forwarding-options dhcp-relay relay-option-82 circuit-id use-interface-description device
```

Configuring QFX1 and QFX2

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 QFX1 and QFX2:

1. Configure forward snooped unicast packets on all interfaces.

```
[edit fowarding-options]
user@switch# set dhcp-relay forward-snooped-clients all-interfaces
```

2. Create a binding entry to snoop unicast clients.

```
[edit forwarding-options]
user@switch# set dhcp-relay overrides allow-snooped-clients
```

3. Create a DHCP server group.

```
[edit forwarding-options]
user@switch# set dhcp-relay server-group GVP-DHCP 10.105.5.202
```

4. Apply a DHCP relay agent configuration to the named group of DHCP server addresses.

```
[edit forwarding-options]
user@switch# set dhcp-relay active-server-group GVP-DHCP
```

5. Configure the relay agent to suppress the installation of ARP and route entries for corresponding client binding.

```
[edit forwarding-options]
user@switch# set dhcp-relay route-suppression destination
```

6. Create a DHCP relay group that includes at least one interface.

DHCP runs on the interfaces defined in the DHCP groups.

```
[edit forwarding-options]
user@switch# set dhcp-relay group Floor1 interface irb.500
```

7. Configure DHCP relay with option 82.

```
[edit forwarding-options]
user@switch# set dhcp-relay relay-option-82 circuit-id use-interface-description device
```

QFX1 and QFX2

From configuration mode, confirm your configuration by entering the show protocols pim commands. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
user@switch> show forwarding-options dhcp-relay
dhcp-relay {
    forward-snooped-clients all-interfaces;
    overrides {
        allow-snooped-clients;
    }
    relay-option-82 {
        circuit-id {
            use-interface-description device;
        }
    }
    server-group {
        GVP-DHCP {
            10.105.5.202;
        }
    }
    active-server-group GVP-DHCP;
    route-suppression {
        destination;
    group Floor1 {
        interface irb.500;
   }
}
```

Example: Simplifying Multichassis Link Aggregation on EX9200 Switches in the Core for Campus Networks

IN THIS SECTION

- Requirements | 220
- Overview | 220
- Configuration | 223
- Verification | 254

Requirements

This example uses the following hardware and software components:

- Junos OS Release 16.1R1 for EX Series
- Two EX9200 switches

NOTE: This configuration example has been tested using the software release listed and is assumed to work on all later releases.

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

- Configure aggregated Ethernet interfaces on a switch. See Configuring an Aggregated Ethernet Interface.
- Configure the Link Aggregation Control Protocol (LACP) on aggregated Ethernet interfaces on a switch. See Configuring Aggregated Ethernet LACP (CLI Procedure).

Overview

IN THIS SECTION

Topology | 222

In this example, you configure an MC-LAG across two switches, consisting of two aggregated Ethernet interfaces, multichassis protection using the ICL, ICCP for the peers hosting the MC-LAG, and Layer 3 connectivity between MC-LAG peers. Layer 3 connectivity is required for ICCP.

To simplify the MC-LAG configuration process, you will enable configuration synchronization and configuration consistency check. 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. 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. When there is an inconsistency, you are notified and can take action to resolve it. Configuration consistency check is invoked after you issue a commit on an MC-LAG peers.

On the EX9200-A switch, you will configure the following configuration synchronization and configuration consistency check parameters:

- Local, remote, and global configuration groups that are synchronized to the EX9200-B switch.
- Conditional groups.
- Apply groups.
- NETCONF over SSH.
- MC-LAG peer details and user authentication details for MC-LAG configuration synchronization.
- peers-synchronize statement to synchronize the configurations between local and remote MC-LAG
 peers by default.
- set multi-chassis mc-lag consistency-check command for consistency check.

On the EX9200-B switch, the configuration process is much shorter and simpler. You will configure the following configuration synchronization and configuration consistency check parameters:

- Apply groups.
- NETCONF over SSH.
- MC-LAG peer details and user authentication details for MC-LAG configuration synchronization.
- peers-synchronize statement to synchronize and commit the configurations between local and remote MC-LAG peers.
- multi-chassis mc-lag consistency-check statement to enable consistency check.

Topology

The topology used in this example consists of two switches hosting an MC-LAG. Figure 11 on page 222 shows the topology of this example.

Figure 11: Topology Diagram

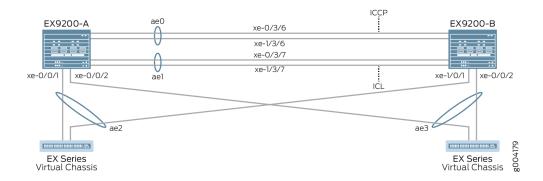


Table 6 on page 223 details the topology used in this configuration example.

Table 6: Components of the Topology for Configuring a Multichassis LAG Between Two Switches

Hostname	Base Hardware	Multichassis Link Aggregation Group
EX9200-B	EX9200 EX9200	ae0 is configured as an aggregated Ethernet interface, and is used as an ICCP link, and the following interfaces are part of ae0: xe-0/3/6 and xe-1/3/6. ae1 is configured as an aggregated Ethernet interface and is used as an ICL link, and the following interfaces are part of ae1: xe-0/3/7 and xe-1/3/7. ae2 is configured as an MC-LAG, and the following interfaces are part of ae2: xe-0/0/1 on Switch B and xe-1/0/1 on Switch A. ae3 is configured as an MC-LAG, and the following interface is part of ae3 on both Switch A and Switch B: xe-0/0/2.
Virtual Chassis Virtual Chassis	Not applicable. Virtual Chassis are shown only for illustration purposes.	The Virtual Chassis are connected to the two EX9200 switches through LAG interfaces. The Virtual Chassis configuration is not included in this example and is only shown to illustrate a sample topology.

Configuration

IN THIS SECTION

CLI Quick Configuration | 224

Configuring MC-LAG on EX9200-A | 227

- Configuring MC-LAG on EX9200-B | 241
- Results | 244

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.

EX9200-A

```
set system login user MCLAG_Admin uid 2000
set system login user MCLAG_Admin class super-user
set system login user MCLAG_Admin authentication encrypted-password "$ABC123"
set system static-host-mapping EX9200-A inet 10.92.76.2
set system static-host-mapping EX9200-B inet 10.92.76.4
set system services netconf ssh
set system commit peers-synchronize
set system commit peers EX9200-B user MCLAG_Admin
set system commit peers EX9200-B authentication "$ABC123"
set interfaces irb unit 100 family inet address 192.168.100.2/24 arp 192.168.100.3 12-interface
ae1
set interfaces irb unit 100 family inet address 192.168.100.2/24 arp 192.168.100.3 mac
28:8a:1c:e5:3b:f0
set interfaces irb unit 100 family inet address 192.168.100.2/24 vrrp-group 1 virtual-address
192.168.100.1
set interfaces irb unit 100 family inet address 192.168.100.2/24 vrrp-group 1 priority 150
set interfaces irb unit 100 family inet address 192.168.100.2/24 vrrp-group 1 accept-data
set interfaces lo0 unit 0 family inet address 172.16.32.5/32
set routing-options static route 0.0.0.0/0 next-hop 10.92.77.254
set protocols ospf area 0.0.0.0 interface lo0.0 passive
set protocols ospf area 0.0.0.0 interface ae0.0
set protocols lldp interface all
set chassis aggregated-devices ethernet device-count 20
set groups MC_Config_Global
set groups MC_Config_Global when peers EX9200-A
set groups MC_Config_Global when peers EX9200-B
set groups MC_Config_Global interfaces xe-0/3/6 ether-options 802.3ad ae0
set groups MC_Config_Global interfaces xe-1/3/6 ether-options 802.3ad ae0
```

```
set groups MC_Config_Global interfaces ae0 description "ICCP Layer 3 Link with 2
members, xe-0/3/6, xe-1/3/6"
set groups MC_Config_Global interfaces ae0 aggregated-ether-options lacp active
set groups MC_Config_Global interfaces ae0 aggregated-ether-options lacp periodic fast
set groups MC_Config_Global interfaces ae0 aggregated-ether-options lacp system-id
00:01:02:03:04:05
set groups MC_Config_Global interfaces ae0 aggregated-ether-options lacp admin-key 0
set groups MC_Config_Global interfaces xe-0/3/7 ether-options 802.3ad ae1
set groups MC_Config_Global interfaces xe-1/3/7 ether-options 802.3ad ae1
set groups MC_Config_Global interfaces ae1 description "ICL Layer 2 link with 2
members, xe-0/3/7, 1/3/7"
set groups MC_Config_Global interfaces ae1 unit 0 family ethernet-switching interface-mode trunk
set groups MC_Config_Global interfaces ae1 unit 0 family ethernet-switching vlan members all
set groups MC_Config_Global interfaces ae1 vlan-tagging
set groups MC_Config_Global interfaces ae1 aggregated-ether-options lacp active
set groups MC_Config_Global interfaces ae1 aggregated-ether-options lacp periodic fast
set groups MC_Config_Global interfaces ae1 aggregated-ether-options lacp system-id
00:01:02:03:04:06
set groups MC_Config_Global interfaces ae1 aggregated-ether-options lacp admin-key 1
set groups MC_Config_Global interfaces xe-0/0/1 ether-options 802.3ad ae2
set groups MC_Config_Global interfaces xe-1/0/1 ether-options 802.3ad ae2
set groups MC_Config_Global interfaces ae2 unit 0 description "MC-LAG interface with members
xe-0/0/1, xe-1/0/1"
set groups MC_Config_Global interfaces ae2 unit 0 family ethernet-switching interface-mode trunk
set groups MC_Config_Global interfaces ae2 unit 0 family ethernet-switching vlan members all
set groups MC_Config_Global interfaces ae2 aggregated-ether-options lacp active
set groups MC_Config_Global interfaces ae2 aggregated-ether-options lacp periodic fast
set groups MC_Config_Global interfaces ae2 aggregated-ether-options lacp system-id
00:01:02:03:04:07
set groups MC_Config_Global interfaces ae2 aggregated-ether-options lacp admin-key 2
set groups MC_Config_Global interfaces ae2 aggregated-ether-options mc-ae mc-ae-id 2
set groups MC_Config_Global interfaces ae2 aggregated-ether-options mc-ae redundancy-group 1
set groups MC_Config_Global interfaces ae2 aggregated-ether-options mc-ae mode active-active
set groups MC_Config_Global interfaces ae2 aggregated-ether-options mc-ae init-delay-time 520
set groups MC_Config_Global interfaces ae2 aggregated-ether-options mc-ae events iccp-peer-down
prefer-status-control-active
set groups MC_Config_Global interfaces xe-0/0/2 ether-options 802.3ad ae3
set groups MC_Config_Global interfaces ae3 unit 0 description "MC-LAG interface with members
xe-0/0/2 on both switches"
set groups MC_Config_Global interfaces ae3 unit 0 family ethernet-switching interface-mode trunk
set groups MC_Config_Global interfaces ae3 unit 0 family ethernet-switching vlan members all
set groups MC_Config_Global interfaces ae3 aggregated-ether-options lacp active
set groups MC_Config_Global interfaces ae3 aggregated-ether-options lacp periodic fast
```

```
set groups MC_Config_Global interfaces ae3 aggregated-ether-options lacp system-id
00:01:02:03:04:08
set groups MC_Config_Global interfaces ae3 aggregated-ether-options lacp admin-key 3
set groups MC_Config_Global interfaces ae3 aggregated-ether-options mc-ae mc-ae-id 3
set groups MC_Config_Global interfaces ae3 aggregated-ether-options mc-ae redundancy-group 1
set groups MC_Config_Global interfaces ae3 aggregated-ether-options mc-ae mode active-active
set groups MC_Config_Global interfaces ae3 aggregated-ether-options mc-ae init-delay-time 520
set groups MC_Config_Global interfaces ae3 aggregated-ether-options mc-ae events iccp-peer-down
prefer-status-control-active
set groups MC_Config_Global vlans v100 vlan-id 100
set groups MC_Config_Global vlans v100 l3-interface irb.100
set groups MC_Config_Global multi-chassis mc-lag consistency-check
set groups MC_Config_Global protocols rstp interface ae2
set groups MC_Config_Global protocols rstp interface ae3
set groups MC_Config_Global protocols rstp bridge-priority 0
set groups MC_Config_Global protocols rstp system-id 00:01:02:03:04:09
set groups MC_Config_Global switch-options service-id 1
set groups MC_Config_Local
set groups MC_Config_Local interfaces ae0 unit 0 family inet address 172.16.32.9/30
set groups MC_Config_Local interfaces ae2 aggregated-ether-options mc-ae chassis-id 0
set groups MC_Config_Local interfaces ae2 aggregated-ether-options mc-ae status-control active
set groups MC_Config_Local interfaces ae3 aggregated-ether-options mc-ae chassis-id 0
set groups MC_Config_Local interfaces ae3 aggregated-ether-options mc-ae status-control active
set groups MC_Config_Remote
set groups MC_Config_Remote interfaces ae0 unit 0 family inet address 172.16.32.10/30
set groups MC_Config_Remote interfaces ae2 aggregated-ether-options mc-ae chassis-id 1
set groups MC_Config_Remote interfaces ae2 aggregated-ether-options mc-ae status-control standby
set groups MC_Config_Remote interfaces ae3 aggregated-ether-options mc-ae chassis-id 1
set groups MC_Config_Remote interfaces ae3 aggregated-ether-options mc-ae status-control standby
set interfaces ae2 unit 0 multi-chassis-protection 172.16.32.6 interface ae1
set interfaces ae3 unit 0 multi-chassis-protection 172.16.32.6 interface ae1
set protocols iccp local-ip-addr 172.16.32.5
set protocols iccp peer 172.16.32.6 session-establishment-hold-time 50
set protocols iccp peer 172.16.32.6 redundancy-group-id-list 1
set protocols iccp peer 172.16.32.6 backup-liveness-detection backup-peer-ip 10.92.76.4
set protocols iccp peer 172.16.32.6 liveness-detection minimum-interval 2000
set protocols iccp peer 172.16.32.6 liveness-detection multiplier 4
set multi-chassis multi-chassis-protection 172.16.32.6 interface ae1
set apply-groups [ MC_Config_Global MC_Config_Local MC_Config_Remote ]
```

EX9200-B

```
set system login user MCLAG_Admin uid 2000
set system login user MCLAG_Admin class super-user
set system login user MCLAG_Admin authentication encrypted-password "$ABC123"
set system static-host-mapping EX9200-A inet 10.92.76.2
set system static-host-mapping EX9200-B inet 10.92.76.4
set system services netconf ssh
set system commit peers-synchronize
set system commit peers EX9200-A user MCLAG_Admin
set system commit peers EX9200-A authentication "$ABC123"
set interfaces irb unit 100 family inet address 192.168.100.3/24 arp 192.168.100.2 12-interface
set interfaces irb unit 100 family inet address 192.168.100.3/24 arp 192.168.100.2 mac
28:8a:1c:e3:f7:f0
set interfaces irb unit 100 family inet address 192.168.100.3/24 vrrp-group 1 virtual-address
set interfaces irb unit 100 family inet address 192.168.100.3/24 vrrp-group 1 priority 100
set interfaces irb unit 100 family inet address 192.168.100.3/24 vrrp-group 1 accept-data
set interfaces lo0 unit 0 family inet address 172.16.32.6/32
set routing-options static route 0.0.0.0/0 next-hop 10.92.77.254
set protocols ospf area 0.0.0.0 interface lo0 passive
set protocols ospf area 0.0.0.0 interface ae0
set protocols lldp interface all
set chassis aggregated-devices ethernet device-count 20
set interfaces ae2 unit 0 multi-chassis-protection 172.16.32.5 interface ae1
set interfaces ae3 unit 0 multi-chassis-protection 172.16.32.5 interface ae1
set protocols iccp local-ip-addr 172.16.32.6
set protocols iccp peer 172.16.32.5 session-establishment-hold-time 50
set protocols iccp peer 172.16.32.5 redundancy-group-id-list 1
set protocols iccp peer 172.16.32.5 backup-liveness-detection backup-peer-ip 10.92.76.2
set protocols iccp peer 172.16.32.5 liveness-detection minimum-interval 2000
set protocols iccp peer 172.16.32.5 liveness-detection multiplier 4
set apply-groups [ MC_Config_Global MC_Config_Local MC_Config_Remote ]
```

Configuring MC-LAG on EX9200-A

Step-by-Step Procedure

The following example requires you to navigate various levels in the configuration hierarchy.

1. Create a user account to access the switch, along with a user identifier (UID), a login class, and a password.

```
[edit system]
user@EX9200-A# set login user MCLAG_Admin uid 2000
user@EX9200-A# set login user MCLAG_Admin class super-user
user@EX9200-A# set login user MCLAG_Admin authentication encrypted-password "$ABC123"
```

2. Statically map EX9200-A to 10.92.76.2 and EX9200-B to 10.92.76.4.

```
[edit system]
user@EX9200-A# set static-host-mapping EX9200-A inet 10.92.76.2
user@EX9200-A# set static-host-mapping EX9200-B inet 10.92.76.4
```

3. Enable NETCONF service using SSH.

```
[edit system]
user@EX9200-A# set services netconf ssh
```

4. Enable the peers-synchronize statement to copy and load the MC-LAG configuration from EX9200-A to EX9200-B by default.

```
[edit system]
user@EX9200-A# set commit peers-synchronize
```

5. Configure the hostname, usernames, and authentication details for EX9200-B, the peer with which EX9200-A will be synchronizing the MC-LAG configuration.

```
[edit system]
user@EX9200-A# set commit peers EX9200-B user MCLAG_Admin
user@EX9200-A# set commit peers EX9200-B user authentication "$ABC123"
```

6. Configure an MC-LAG IRB and configure static Address Resolution Protocol (ARP) on the MC-LAG IRB peers to allow routing protocols to traverse the IRB interface.

```
[edit interfaces]
user@EX9200-A# set irb unit 100 family inet address 192.168.100.2/24 arp 192.168.100.3 l2-
```

interface ae1

user@EX9200-A# set irb unit 100 family inet address 192.168.100.2/24 arp 192.168.100.3 mac 28:8a:1c:e5:3b:f0

7. Enable VRRP on the MC-LAGs by assigning a virtual IP address that is shared between each switch in the VRRP group, and assigning an individual IP address for each individual member in the VRRP group.

```
[edit interfaces]
user@EX9200-A# set irb unit 100 family inet address 192.168.100.2/24 vrrp-group 1 virtual-
address 192.168.100.1
user@EX9200-A# set irb unit 100 family inet address 192.168.100.2/24 vrrp-group 1 priority
150
user@EX9200-A# set irb unit 100 family inet address 192.168.100.2/24 vrrp-group 1 accept-
data
```

8. Configure a loopback interface.

```
[edit interfaces]
user@EX9200-A# set lo0 unit 0 family inet address 172.16.32.5/32
```

9. Configure a default gateway.

```
[edit routing-options]
user@EX9200-A# set static route 0.0.0.0 next-hop 10.92.77.254
```

10. Configure an OSPF area that includes the loopback interface and the ICCP interface.

```
[edit protocols]
user@EX9200-A# set ospf area 0.0.0.0 interface lo0 passive
user@EX9200-A# set ospf area 0.0.0.0 interface ae0
```

11. Configure Link Layer Discovery Protocol for all interfaces.

```
[edit protocols]
user@EX9200-A# set lldp interface all
```

12. Configure the number of aggregated Ethernet interfaces to be created on EX9200-A.

```
[edit chassis]
user@EX9200-A# set aggregated-devices ethernet device-count 20
```

 Configure a configuration group for a global MC-LAG configuration that applies to both EX9200-A and EX9200-B.

The global configuration is synchronized between EX9200-A and EX9200-B.

```
[edit groups]
user@EX9200-A# set MC_Config_Global
```

14. Specify the peers that will apply the MC_Config_Global configuration group.

```
[edit groups]
user@EX9200-A# set MC_Config_Global when peers EX9200-A
user@EX9200-A# set MC_Config_Global when peers EX9200-B
```

15. Add member interfaces to the aggregated Ethernet interfaces that will be used for the Inter-Chassis Control Protocol (ICCP) interface.

```
[edit groups]
user@EX9200-A# set MC_Config_Global interfaces xe-0/3/6 ether-options 802.3ad ae0
user@EX9200-A# set MC_Config_Global interfaces xe-1/3/6 ether-options 802.3ad ae0
```

16. Configure the aggregated Ethernet interface (ae0) that will be used for the Inter-Chassis Control Protocol (ICCP) interface.

NOTE: You will be configuring the IP address for ae0 in a later step.

```
[edit groups]
user@EX9200-A# set MC_Config_Global interfaces ae0 description "ICCP Layer 3 Link with 2
members,xe-0/3/6,xe-1/3/6"
```

17. Configure the LACP parameters on ae0.

```
[edit groups]
user@EX9200-A# set MC_Config_Global interfaces ae0 aggregated-ether-options lacp active
user@EX9200-A# set MC_Config_Global interfaces ae0 aggregated-ether-options lacp periodic
fast
```

18. Configure the LACP system ID on ae0.

```
[edit groups]
user@EX9200-A# set MC_Config_Global interfaces ae0 aggregated-ether-options lacp system-id
00:01:02:03:04:05
```

19. Configure the LACP administrative key on ae0.

```
[edit groups]
user@EX9200-A# set MC_Config_Global interfaces ae0 aggregated-ether-options lacp admin-key 0
```

20. Add member interfaces to the aggregated Ethernet interface (ae1) that will be used for the ICL.

```
[edit groups]
user@EX9200-A# set MC_Config_Global interfaces xe-0/3/7 ether-options 802.3ad ae1
user@EX9200-A# set MC_Config_Global interfaces xe-1/3/7 ether-options 802.3ad ae1
```

21. Configure the aggregated Ethernet interface that will be used for the ICL.

```
[edit groups]
user@EX9200-A# set MC_Config_Global interfaces ae1 description "ICL Layer 2 link with 2
members,xe-0/3/7,1/3/7"
```

22. Configure ae1 as a Layer 2 interface.

```
[edit groups]
user@EX9200-A# set MC_Config_Global ae1 unit 0 family ethernet-switching interface-mode
trunk
user@EX9200-A# set MC_Config_Global ae1 unit 0 family ethernet-switching vlan members all
```

23. Enable the reception and transmission of 802.1Q VLAN-tagged frames on ae1.

```
[edit groups]
user@EX9200-A# set MC_Config_Global interfaces ae1 vlan-tagging
```

24. Configure the LACP parameters on ae1.

```
[edit groups]
user@EX9200-A# set MC_Config_Global interfaces ae1 aggregated-ether-options lacp active
user@EX9200-A# set MC_Config_Global interfaces ae1 aggregated-ether-options lacp periodic
fast
```

25. Configure the LACP system ID on ae1.

```
[edit groups]
user@EX9200-A# set MC_Config_Global interfaces ae1 aggregated-ether-options lacp system-id
00:01:02:03:04:06
```

26. Configure the LACP administrative key on ae1.

```
[edit groups]
user@EX9200-A# set MC_Config_Global interfaces ae1 aggregated-ether-options lacp admin-key 1
```

27. Add member interfaces to the aggregated Ethernet interface (ae2) that will be used as the MC-LAG interface.

```
[edit groups]
user@EX9200-A# set MC_Config_Global interfaces xe-0/0/1 ether-options 802.3ad ae2
user@EX9200-A# set MC_Config_Global interfaces xe-1/0/1 ether-options 802.3ad ae2
```

28. Configure the aggregated Ethernet interface (ae2) that will be used as an MC-LAG interface.

```
[edit groups]
user@EX9200-A# set MC_Config_Global interfaces ae2 description "MC-LAG interface with
members xe-0/0/1,xe-1/0/1"
```

29. Configure ae2 as a Layer 2 interface.

```
[edit groups]
user@EX9200-A# set MC_Config_Global interfaces ae2 unit 0 family ethernet-switching
interface-mode trunk
user@EX9200-A# set MC_Config_Global interfaces ae2 unit 0 family ethernet-switching vlan
members all
```

30. Configure the LACP parameters on ae2.

```
[edit groups]
user@EX9200-A# set MC_Config_Global interfaces ae2 aggregated-ether-options lacp active
user@EX9200-A# set MC_Config_Global interfaces ae2 aggregated-ether-options lacp periodic
fast
```

31. Configure the LACP system ID on ae2.

```
[edit groups]
user@EX9200-A# set MC_Config_Global interfaces ae2 aggregated-ether-options lacp system-id
00:01:02:03:04:07
```

32. Configure the LACP administrative key on ae2.

```
[edit groups]
user@EX9200-A# set MC_Config_Global interfaces ae2 aggregated-ether-options lacp admin-key 2
```

33. Configure the MC-AE interface properties on ae2.

```
[edit groups]
user@EX9200-A# set MC_Config_Global interfaces ae2 aggregated-ether-options mc-ae mc-ae-id 2
user@EX9200-A# set MC_Config_Global interfaces ae2 aggregated-ether-options mc-ae
redundancy-group 1
```

34. Specify the mode of ae2 to be active-active.

```
[edit groups]
user@EX9200-A# set MC_Config_Global interfaces ae2 aggregated-ether-options mc-ae mode
active-active
```

35. Specify the time in seconds to delay bringing the MC-AE interface to the up state after rebooting an MC-LAG peer.

By delaying the bring-up of the interface until after protocol convergence, you can prevent packet loss during the recovery of failed links and devices. This network configuration example uses a delay time of 520 seconds. This delay time might not be optimal for your network and should be adjusted to fit your network requirements.

```
[edit groups]
user@EX9200-A# set MC_Config_Global interfaces ae2 aggregated-ether-options mc-ae init-
delay-time 520
```

36. Specify that if a peer of the MC-LAG group goes down, the peer that is configured as status-control active becomes the active peer.

```
[edit groups]
user@EX9200-A# set MC_Config_Global interfaces ae2 aggregated-ether-options mc-ae events
iccp-peer-down prefer-status-control-active
```

37. Add member interfaces to the aggregated Ethernet interface (ae3) that will be used as the MC-LAG interface.

NOTE: EX9200-B uses the same interface name of xe-0/0/2.

```
[edit groups]
user@EX9200-A# set MC_Config_Global interfaces xe-0/0/2 ether-options 802.3ad ae3
```

38. Configure the aggregated Ethernet interface (ae3) that will be used as an MC-LAG interface.

```
[edit groups]
user@EX9200-A# set groups MC_Config_Global interfaces ae3 description "MC-LAG interface
with members xe-0/0/2 on both switches"
```

39. Configure ae3 as a Layer 2 interface.

```
[edit groups]
user@EX9200-A# set MC_Config_Global interfaces ae3 unit 0 family ethernet-switching
interface-mode trunk
user@EX9200-A# set MC_Config_Global interfaces ae3 unit 0 family ethernet-switching vlan
members all
```

40. Configure the LACP parameters on ae3.

```
[edit groups]
user@EX9200-A# set MC_Config_Global interfaces ae3 aggregated-ether-options lacp active
user@EX9200-A# set MC_Config_Global interfaces ae3 aggregated-ether-options lacp periodic
fast
```

41. Configure the LACP system ID on ae3.

```
[edit groups]
user@EX9200-A# set MC_Config_Global interfaces ae3 aggregated-ether-options lacp system-id
00:01:02:03:04:08
```

42. Configure the LACP administrative key on ae3.

```
[edit groups]
user@EX9200-A# set MC_Config_Global interfaces ae3 aggregated-ether-options lacp admin-key 3
```

43. Configure the MC-AE interface properties on ae3.

```
[edit groups]
user@EX9200-A# set MC_Config_Global interfaces ae3 aggregated-ether-options mc-ae mc-ae-id 3
```

user@EX9200-A# set MC_Config_Global interfaces ae3 aggregated-ether-options mc-ae
redundancy-group 1

44. Specify the mode of ae3 to be active-active.

```
[edit groups]
user@EX9200-A# set MC_Config_Global interfaces ae3 aggregated-ether-options mc-ae mode
active-active
```

45. Specify the time in seconds to delay bringing the MC-AE interface to the up state after rebooting an MC-LAG peer.

By delaying the bring-up of the interface until after protocol convergence, you can prevent packet loss during the recovery of failed links and devices. This network configuration example uses a delay time of 520 seconds. This delay time might not be optimal for your network and should be adjusted to fit your network requirements.

```
[edit groups]
user@EX9200-A# set MC_Config_Global interfaces ae3 aggregated-ether-options mc-ae init-
delay-time 520
```

46. Specify that if a peer of the MC-LAG group goes down, the peer that is configured as status-control active becomes the active peer.

```
[edit groups]
user@EX9200-A# set MC_Config_Global interfaces ae3 aggregated-ether-options mc-ae events
iccp-peer-down prefer-status-control-active
```

47. Configure VLAN 100 to connect end users.

```
[edit groups]
user@EX9200-A# set MC_Config_Global vlans v100 vlan-id 100
```

48. Configure the routed VLAN interface for VLAN 100.

```
[edit groups]
user@EX9200-A# set MC_Config_Global vlans v100 l3-interface irb.100
```

49. Enable consistency check.

```
[edit groups]
user@EX9200-A# set MC_Config_Global multi-chassis mc-lag consistency-check
```

50. Enable the Rapid Spanning Tree Protocol on the ae2 and ae3 interfaces (MC-LAG interfaces) for optional loop prevention.

```
[edit groups]
user@EX9200-A# set MC_Config_Global protocols rstp interfaces ae2
user@EX9200-A# set MC_Config_Global protocols rstp interfaces ae3
```

51. Configure the RSTP bridge priority.

Setting the bridge priority to 0 will make the MC-AE nodes of EX9200-A and EX9200-B the best priority.

```
[edit groups]
user@EX9200-A# set MC_Config_Global protocols rstp bridge-priority 0
```

52. Configure the RSTP system identifier value.

```
[edit groups]
user@EX9200-A# set MC_Config_Global protocols rstp system-id 00:01:02:03:04:09
```

53. Specify the switch service ID.

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

```
[edit groups]
user@EX9200-A# set MC_Config_Global switch-options service-id 1
```

54. Configure a configuration group for an MC-LAG configuration that applies to the local peer.

```
[edit groups]
user@EX9200-A# set MC_Config_Local
```

55. Configure the ICCP interface (ae0) as a Layer 3 interface.

```
[edit groups]
user@EX9200-A# set MC_Config_Local interfaces ae0 unit 0 family inet address 172.16.32.9/30
```

56. Specify a unique chassis ID for the MC-LAG (ae2) that the aggregated Ethernet interface belongs to.

```
[edit groups]
user@EX9200-A# set MC_Config_Local interfaces ae2 aggregated-ether-options mc-ae chassis-id
0
```

57. Specify the status-control setting of ae2 to be active.

```
[edit groups]
user@EX9200-A# set MC_Config_Local interfaces ae2 aggregated-ether-options mc-ae status-
control active
```

58. Specify a unique chassis ID for the MC-LAG (ae3) that the aggregated Ethernet interface belongs to.

```
[edit groups]
user@EX9200-A# set MC_Config_Local interfaces ae3 aggregated-ether-options mc-ae chassis-id
0
```

59. Specify the status-control setting of ae3 to be active..

```
[edit groups]
user@EX9200-A# set MC_Config_Local interfaces ae3 aggregated-ether-options mc-ae status-
control active
```

60. Configure a configuration group for an MC-LAG configuration that applies to the remote peer.

```
[edit groups]
user@EX9200-A# set MC_Config_Remote
```

61. Configure ae0 as a Layer 3 interface.

```
[edit groups]
user@EX9200-A# set MC_Config_Remote interfaces ae0 unit 0 family inet address
172.16.32.10/30
```

62. Specify a unique chassis ID for the MC-LAG (ae2) that the aggregated Ethernet interface belongs to.

```
[edit groups]
user@EX9200-A# set MC_Config_Remote interfaces ae2 aggregated-ether-options mc-ae chassis-
id 1
```

63. Specify the status-control setting of ae2 to be standby.

```
[edit groups]
user@EX9200-A# set MC_Config_Remote interfaces ae2 aggregated-ether-options mc-ae status-
control standby
```

64. Specify a unique chassis ID for the MC-LAG (ae3) that the aggregated Ethernet interface belongs to.

```
[edit groups]
user@EX9200-A# set MC_Config_Remote interfaces ae3 aggregated-ether-options mc-ae chassis-
id 1
```

65. Specify the status-control setting of ae3 to be standby.

```
[edit interfaces]
user@EX9200-A# set MC_Config_Remote interfaces ae3 aggregated-ether-options mc-ae status-
control standby
```

66. Specify that if a peer of the MC-LAG group goes down, the peer that is configured as status-control active becomes the active peer.

```
[edit interfaces]
user@EX9200-A# set MC_Config_Remote interfaces ae3 aggregated-ether-options mc-ae events
iccp-peer-down prefer-status-control-standby
```

67. Enable link protection between the two MC-LAG peers.

Assign interface ae1 to act as the ICL to protect the MC-AE interfaces, ae2 and ae3, in case of failure.

```
[edit interfaces]
user@EX9200-A# set ae2 unit 0 multi-chassis-protection 172.16.32.6 interface ae1
```

```
user@EX9200-A# set ae3 unit 0 multi-chassis-protection 172.16.32.6 interface ae1
```

68. Specify the local IP address of the ICCP interface.

```
[edit protocols]
user@EX9200-A# set iccp local-ip-addr 172.16.32.5
```

69. Configure the session establishment hold time for ICCP to connect faster.

NOTE: We recommend 50 seconds as the session establishment hold time value.

```
[edit protocols]
user@EX9200-A# set iccp peer 172.16.32.6 session-establishment-hold-time 50
user@EX9200-A# set iccp peer 172.16.32.6 redundancy-group-id-list 1
user@EX9200-A# set iccp peer 172.16.32.6 backup-liveness-detection backup-peer-ip 10.92.76.4
```

70. To enable BFD for ICCP, configure the minimum receive interval.

We recommend a minimum receive interval value of 6 seconds.

```
[edit protocols]
user@EX9200-A# set iccp peer 172.16.32.6 liveness-detection minimum-interval 2000
user@EX9200-A# set iccp peer 172.16.32.6 liveness-detection multiplier 4
```

71. Apply the groups configured earlier, so that the Junos configuration will inherit the statements from the MC_Config_Global, MC_Config_Local, and MC_Config_Remote configuration groups.

```
[edit]
user@EX9200-A# set apply-groups [ MC_Config_Global MC_Config_Local MC_Config_Remote ]
```

Configuring MC-LAG on EX9200-B

Step-by-Step Procedure

The following example requires you to navigate various levels in the configuration hierarchy.

1. Create a user account to access the switch, along with a user identifier (UID), a login class, and a password.

```
[edit system]
user@EX9200-A# set login user MCLAG_Admin uid 2000
user@EX9200-B# set login user MCLAG_Admin class super-user
user@EX9200-B# set login user MCLAG_Admin authentication encrypted-password "$ABC123"
```

2. Statically map EX9200-A to 10.92.76.2 and EX9200-B to 10.92.76.4.

```
[edit system]
user@EX9200-B# set static-host-mapping EX9200-A inet 10.92.76.2
user@EX9200-B# set static-host-mapping EX9200-B inet 10.92.76.4
```

3. Enable NETCONF service using SSH.

```
[edit system]
user@EX9200-B# set services netconf ssh
```

4. Enable the peers-synchronize statement to copy and load the MC-LAG configuration from EX9200-B to EX9200-A by default.

```
[edit system]
user@EX9200-B# set commit peers-synchronize
```

5. Configure the hostname, usernames, and authentication details for EX9200-A, the peer with which EX9200-B will be synchronizing the MC-LAG configuration.

```
[edit system]
user@EX9200-B# set commit peers EX9200-A user MCLAG_Admin
user@EX9200-A# set commit peers EX9200-A authentication "$ABC123"
```

6. Configure an MC-LAG IRB and configure static Address Resolution Protocol (ARP) on the MC-LAG IRB peers to allow routing protocols to traverse the IRB interface.

```
[edit interfaces]
user@EX9200-B# set irb unit 100 family inet address 192.168.100.3/24 arp 192.168.100.2 l2-
interface ae1
user@EX9200-B# set irb unit 100 family inet address 192.168.100.3/24 arp 192.168.100.2 mac
28:8a:1c:e3:f7:f0
```

7. Enable VRRP on the MC-LAGs by assigning a virtual IP address that is shared between each switch in the VRRP group, and assigning an individual IP address for each individual member in the VRRP group.

```
[edit interfaces]
user@EX9200-B# set irb unit 100 family inet address 192.168.100.3/24 vrrp-group 1 virtual-
address 192.168.100.1
user@EX9200-B# set irb unit 100 family inet address 192.168.100.3/24 vrrp-group 1 priority
100
user@EX9200-B# set irb unit 100 family inet address 192.168.100.3/24 vrrp-group 1 accept-
data
```

8. Configure a loopback interface.

```
[edit interfaces]
user@EX9200-B# set lo0 unit 0 family inet address 172.16.32.6/32
```

9. Configure a default gateway.

```
[edit routing-options]
user@EX9200-B# set static route 0.0.0.0 next-hop 10.92.77.254
```

10. Configure an OSPF area that includes the loopback interface and the ICCP interface.

```
[edit protocols]
user@EX9200-B# set ospf area 0.0.0.0 interface lo0 passive
user@EX9200-B# set ospf area 0.0.0.0 interface ae0
```

11. Configure Link Layer Discovery Protocol for all interfaces.

```
[edit protocols]
user@EX9200-B# set lldp interface all
```

12. Configure the number of aggregated Ethernet interfaces to be created on EX9200-B.

```
[edit chassis]
user@EX9200-B# set aggregated-devices ethernet device-count 20
```

13. Enable link protection between the two MC-LAG peers.

Assign interface ae1 to act as the ICL to protect the MC-AE interfaces, ae2 and ae3, in case of failure.

```
[edit interfaces]
user@EX9200-B# set ae2 unit 0 multi-chassis-protection 172.16.32.5 interface ae1
user@EX9200-B# set ae3 unit 0 multi-chassis-protection 172.16.32.5 interface ae1
```

14. Specify the local IP address of the ICCP interface.

```
[edit protocols]
user@EX9200-B# set iccp local-ip-addr 172.16.32.6
```

15. Configure the session establishment hold time for ICCP to connect faster.

NOTE: We recommend 50 seconds as the session establishment hold time value.

```
[edit protocols]
user@EX9200-B# set iccp peer 172.16.32.5 session-establishment-hold-time 50
user@EX9200-B# set iccp peer 172.16.32.5 redundancy-group-id-list 1
user@EX9200-B# set iccp peer 172.16.32.5 backup-liveness-detection backup-peer-ip 10.92.76.2
```

16. To enable BFD for ICCP, configure the minimum receive interval.

We recommend a minimum receive interval value of 6 seconds.

```
[edit protocols]
user@EX9200-B# set iccp peer 172.16.32.5 liveness-detection minimum-interval 2000
user@EX9200-B# set iccp peer 172.16.32.5 liveness-detection multiplier 4
```

17. Apply the groups configured earlier, so that the Junos configuration will inherit the statements from the MC_Config_Global, MC_Config_Local, and MC_Config_Remote configuration groups.

```
[edit]
user@EX9200-B# set apply-groups [ MC_Config_Global MC_Config_Local MC_Config_Remote ]
```

Results

Display the results of the configuration on EX9200-A before you commit the configuration.

```
user@EX9200-A# show system services
netconf {
```

```
ssh;
}
```

```
user@EX9200-A# show system commit
peers-synchronize;
  peers {
     EX9200-B {
        user MCLAG_Admin;
        authentication "$ABC123";
     }
  }
}
```

```
user@EX9200-A# show interfaces
ae2 {
    unit 0 {
        multi-chassis-protection 172.16.32.6 {
            interface ae1;
        }
    }
}
ae3 {
    unit 0 {
        multi-chassis-protection 172.16.32.6 {
            interface ae1;
        }
    }
}
irb {
    unit 100 {
        family inet {
            address 192.168.100.2/24 {
                arp 192.168.100.3 12-interface ae1.0 mac 28:8a:1c:e5:3b:f0;
                vrrp-group 1 {
                    virtual-address 192.168.100.1;
                    priority 150;
                    accept-data;
                }
            }
```

```
}
}
lo0 {
  unit 0 {
    family inet {
       address 172.16.32.5/32;
    }
}
```

```
user@EX9200-A# show routing-options
static {
    route 0.0.0.0/0 next-hop 10.92.77.254;
}
```

```
user@EX9200-A# show protocols
ospf {
    area 0.0.0.0 {
        interface lo0.0 {
            passive;
        interface ae0.0;
    }
}
iccp {
    local-ip-addr 172.16.32.5;
    peer 172.16.32.6 {
        session-establishment-hold-time 50;
        redundancy-group-id-list 1;
        backup-liveness-detection {
            backup-peer-ip 10.92.76.4;
        liveness-detection {
            minimum-interval 2000;
            multiplier 4;
        }
    }
}
11dp {
```

```
interface all;
}
```

```
user@EX9200-A# show chassis
aggregated-devices {
    ethernet {
        device-count 20;
    }
}
```

```
user@EX9200-A# show groups MC_Config_Global
when {
    peers [ EX9200-A EX9200-B ];
}
interfaces {
    xe-0/3/6 {
        ether-options {
            802.3ad ae0;
        }
    }
    xe-1/3/6 {
        ether-options {
            802.3ad ae0;
        }
    }
    ae0 {
        description "ICCP Layer 3 Link with 2 members, xe-0/3/6, xe-1/3/6";
        aggregated-ether-options {
            lacp {
                active;
                periodic fast;
                system-id 00:01:02:03:04:05;
                admin-key 0;
            }
        }
    }
    xe-0/3/7 {
        ether-options {
            802.3ad ae1;
```

```
xe-1/3/7 {
    ether-options {
        802.3ad ae1;
    }
}
ae1 {
    description "ICL Layer 2 link with 2 members, xe-0/3/7,1/3/7";
    vlan-tagging;
    aggregated-ether-options {
        lacp {
            active;
            periodic fast;
            system-id 00:01:02:03:04:06;
            admin-key 1;
        }
    }
    unit 0 {
        family ethernet-switching {
            interface-mode trunk;
            vlan {
                members all;
            }
        }
    }
}
xe-0/0/1 {
    ether-options {
        802.3ad ae2;
    }
}
xe-1/0/1 {
    ether-options {
        802.3ad ae2;
    }
}
ae2 {
    description "MC-LAG interface with members xe-0/0/1,xe-1/0/1";
    aggregated-ether-options {
        lacp {
            active;
            periodic fast;
            system-id 00:01:02:03:04:07;
```

```
admin-key 2;
        }
        mc-ae {
            mc-ae-id 2;
            redundancy-group 1;
            mode active-active;
            init-delay-time 520;
            events {
                iccp-peer-down {
                    prefer-status-control-active;
                }
            }
        }
    }
    unit 0 {
        family ethernet-switching {
            interface-mode trunk;
            vlan {
                members all;
            }
        }
    }
}
xe-0/0/2 {
    ether-options {
        802.3ad ae3;
    }
}
ae3 {
    description "MC-LAG interface with members xe-0/0/2 on both switches"
    aggregated-ether-options {
        lacp {
            active;
            periodic fast;
            system-id 00:01:02:03:04:08;
            admin-key 3;
        }
        mc-ae {
            mc-ae-id 3;
            redundancy-group 1;
            mode active-active;
            init-delay-time 520;
            events {
```

```
iccp-peer-down {
                        prefer-status-control-active;
                    }
                }
            }
        }
        unit 0 {
            family ethernet-switching {
                interface-mode trunk;
                vlan {
                    members all;
                }
            }
        }
    }
}
multi-chassis {
    mc-lag {
        consistency-check;
    }
}
protocols {
    rstp {
        bridge-priority 0;
        system-id 00:01:02:03:04:09;
        interface ae2;
        interface ae3;
    }
}
switch-options {
    service-id 1;
}
vlans {
    v100 {
        vlan-id 100;
        13-interface irb.100;
   }
}
```

```
user@EX9200-A# show groups MC_Config_Local
interfaces {
```

```
ae0 {
        unit 0 {
            family inet \{
                address 172.16.32.9/30;
            }
        }
   }
    ae2 {
        aggregated-ether-options {
            mc-ae {
                chassis-id 0;
                status-control active;
            }
        }
   }
    ae3 {
        aggregated-ether-options {
            mc-ae {
                chassis-id 0;
                status-control active;
        }
   }
}
```

```
ae3 {
    aggregated-ether-options {
        mc-ae {
            chassis-id 1;
            status-control standby;
        }
    }
}
```

```
user@EX9200-A# show apply-groups
apply-groups [ MC_Config_Global MC_Config_Local MC_Config_Remote ];
```

Display the results of the configuration on EX9200-B before you commit the configuration.

```
user@EX9200-B# show system services
netconf {
    ssh;
}
```

```
user@EX9200-B# show system commit
peers-synchronize;
peers {
    EX9200-A {
      user MCLAG_Admin;
      authentication "$ABC123";
    }
}
```

```
user@EX9200-B# show interfaces
ae2 {
   unit 0 {
      multi-chassis-protection 172.16.32.5 {
         interface ae1;
      }
   }
}
```

```
ae3 {
    unit 0 {
        multi-chassis-protection 172.16.32.5 {
            interface ae1;
        }
    }
}
irb {
    unit 100 {
        family inet {
            address 192.168.100.3/24 {
                arp 192.168.100.2 12-interface ae1.0 mac 28:8a:1c:e3:f7:f0;
                vrrp-group 1 {
                    virtual-address 192.168.100.1;
                    priority 100;
                    accept-data;
                }
            }
        }
    }
}
lo0 {
    unit 0 {
        family inet {
            address 172.16.32.6/32;
        }
    }
}
```

```
user@EX9200-B# show routing-options
static {
    route 0.0.0.0/0 next-hop 10.92.77.254;
}
```

```
user@EX9200-B# show protocols
ospf {
    area 0.0.0.0 {
        interface lo0.0 {
           passive;
        }
}
```

```
interface ae0.0;
   }
}
iccp {
    local-ip-addr 172.16.32.6;
    peer 172.16.32.5 {
        session-establishment-hold-time 50;
        redundancy-group-id-list 1;
        backup-liveness-detection {
            backup-peer-ip 10.92.76.2;
        }
        liveness-detection {
            minimum-interval 2000;
            multiplier 4;
        }
   }
}
lldp {
    interface all;
}
```

```
user@EX9200-B# show chassis
aggregated-devices {
    ethernet {
        device-count 20;
    }
}
```

```
user@EX9200-B# show apply-groups
[ MC_Config_Global MC_Config_Local MC_Config_Remote ];
```

Verification

IN THIS SECTION

• Verifying ICCP on MC-LAG | 255

Verifying LACP on MC-LAG | 256

- Verifying Aggregated Ethernet Interfaces in MC-LAG | 259
- Verifying MAC Learning on MC-LAG | 260
- Verifying VRRP in MC-LAG | 262
- Verifying OSPF on MC-LAG | 262
- Verifying that Configuration Consistency Check Passed | 264
- Verifying the Configuration Consistency Check Status for the Global Configuration | 269
- Verifying the Configuration Consistency Check Status for the Interchassis Control Link | 272
- Verifying the Configuration Consistency Check Status for the MC-LAG Interfaces | 273
- Verifying the Configuration Consistency Check Status for the VLAN Configuration | 278
- Verifying the Configuration Consistency Check Status for VRRP | 279

Verifying ICCP on MC-LAG

Purpose

Verify that ICCP is running on each device in the MC-LAG.

Action

1. Verify that ICCP is running on EX9200-A.

```
Client Application: mclag_cfgchkd
Redundancy Group IDs Joined: 1
```

2. Verify that ICCP is running on EX9200-B.

```
user@EX9200-B> show iccp
Redundancy Group Information for peer 172.16.32.5
  TCP Connection
                       : Established
  Liveliness Detection : Up
Backup liveness peer status: Up
  Redundancy Group ID
                               Status
    1
                                Up
Client Application: lacpd
  Redundancy Group IDs Joined: 1
Client Application: 12ald_iccpd_client
  Redundancy Group IDs Joined: 1
Client Application: mclag_cfgchkd
  Redundancy Group IDs Joined: 1
```

Meaning

This output shows that the TCP connection between the peers hosting the MC-LAG is up, liveness detection is up, Backup liveness peer status is up, and LACPD, MCLAG_CFGCHKD, and L2ALD _ICCP_CLIENT client applications are running.

Verifying LACP on MC-LAG

Purpose

Verify that LACP is working properly on each device in the MC-LAG.

Action

1. Verify that the LACP interfaces are up and running on EX9200-A.

```
user@EX9200-A> show lacp interfaces
Aggregated interface: ae0
   LACP state:
                      Role
                                        Dist Col Syn
                                                        Aggr
                                                              Timeout Activity
                             Exp
                                   Def
      xe-0/3/6
                                                         Yes
                                                                          Active
                     Actor
                              No
                                    No
                                         Yes
                                              Yes
                                                   Yes
                                                                  Fast
     xe-0/3/6
                   Partner
                                         Yes Yes
                                                   Yes
                                                         Yes
                                                                  Fast
                                                                          Active
                              No
                                    No
     xe-1/3/6
                     Actor
                              No
                                    No
                                         Yes Yes
                                                   Yes
                                                         Yes
                                                                  Fast
                                                                          Active
     xe-1/3/6
                                                                          Active
                   Partner
                              No
                                    No
                                         Yes Yes Yes
                                                         Yes
                                                                  Fast
   LACP protocol:
                                                                 Mux State
                          Receive State Transmit State
      xe-0/3/6
                                Current
                                          Fast periodic Collecting distributing
      xe-1/3/6
                                          Fast periodic Collecting distributing
                                Current
Aggregated interface: ae1
   LACP state:
                      Role
                             Exp
                                   Def
                                        Dist Col
                                                   Syn
                                                        Aggr
                                                              Timeout Activity
      xe-0/3/7
                                                         Yes
                     Actor
                              No
                                    No
                                         Yes Yes
                                                   Yes
                                                                  Fast
                                                                          Active
     xe-0/3/7
                   Partner
                              No
                                    No
                                         Yes Yes
                                                   Yes
                                                         Yes
                                                                  Fast
                                                                          Active
      xe-1/3/7
                     Actor
                              No
                                         Yes
                                              Yes
                                                   Yes
                                                         Yes
                                                                  Fast
                                                                          Active
      xe-1/3/7
                   Partner
                              No
                                    No
                                         Yes Yes Yes
                                                         Yes
                                                                  Fast
                                                                          Active
   LACP protocol:
                          Receive State Transmit State
                                                                  Mux State
                                          Fast periodic Collecting distributing
      xe-0/3/7
                                Current
      xe-1/3/7
                                          Fast periodic Collecting distributing
                                Current
Aggregated interface: ae2
   LACP state:
                      Role
                                                              Timeout Activity
                             Exp
                                   Def Dist
                                              Col
                                                   Syn
                                                        Aggr
     xe-0/0/1
                                                         Yes
                     Actor
                              No
                                   Yes
                                          No
                                               No
                                                    No
                                                                  Fast
                                                                          Active
     xe-0/0/1
                   Partner
                              No
                                   Yes
                                               No
                                                    No
                                                         Yes
                                                                  Fast
                                                                         Passive
                                          No
   LACP protocol:
                          Receive State Transmit State
                                                                  Mux State
      xe-0/0/1
                                            Fast periodic Collecting distributing
                                Current
      xe-1/0/1
                          Port disabled
                                            Fast periodic Collecting distributing
Aggregated interface: ae3
   LACP state:
                      Role
                             Exp
                                   Def Dist Col
                                                   Syn
                                                        Aggr Timeout Activity
     xe-0/0/2
                     Actor
                              No
                                   Yes
                                          No
                                               No
                                                    No
                                                         Yes
                                                                  Fast
                                                                          Active
     xe-0/0/2
                                                                         Passive
                   Partner
                              No
                                   Yes
                                          No
                                               No
                                                    No
                                                         Yes
                                                                  Fast
   LACP protocol:
                                                                  Mux State
                          Receive State Transmit State
      xe-0/0/2
                                          Fast periodic Collecting distributing
                                Current
```

2. Verify that the LACP interfaces are up and running on EX9200-B.

```
user@EX9200-B> show lacp interfaces
Aggregated interface: ae0
   LACP state:
                     Role
                                   Def Dist Col Syn Aggr Timeout Activity
                            Exp
      xe-0/3/6
                    Actor
                             No
                                   No
                                         Yes Yes
                                                  Yes
                                                         Yes
                                                                 Fast
                                                                         Active
      xe-0/3/6
                  Partner
                                                                         Active
                             No
                                   No
                                         Yes Yes
                                                  Yes
                                                         Yes
                                                                 Fast
     xe-1/3/6
                    Actor
                                                                         Active
                              No
                                   No
                                         Yes Yes Yes
                                                         Yes
                                                                 Fast
     xe-1/3/6
                  Partner
                             No
                                   No
                                         Yes Yes Yes
                                                         Yes
                                                                 Fast
                                                                         Active
                         Receive State Transmit State
   LACP protocol:
                                                                 Mux State
                                          Fast periodic Collecting distributing
      xe-0/3/6
                                Current
      xe-1/3/6
                                Current
                                          Fast periodic Collecting distributing
Aggregated interface: ae1
   LACP state:
                     Role
                                                             Timeout Activity
                            Exp
                                   Def Dist Col Syn Aggr
     xe-0/3/7
                    Actor
                             No
                                         Yes
                                             Yes
                                                  Yes
                                                         Yes
                                                                 Fast
                                                                         Active
      xe-0/3/7
                  Partner
                                        Yes Yes
                                                  Yes
                                                         Yes
                                                                 Fast
                                                                         Active
                             No
                                   No
      xe-1/3/7
                    Actor
                                        Yes Yes Yes
                                                                 Fast
                                                                         Active
                             No
                                   No
                                                         Yes
      xe-1/3/7
                  Partner
                                         Yes Yes Yes
                                                                 Fast
                                                                         Active
   LACP protocol:
                         Receive State Transmit State
                                                                 Mux State
      xe-0/3/7
                                Current
                                          Fast periodic Collecting distributing
      xe-1/3/7
                                Current
                                          Fast periodic Collecting distributing
Aggregated interface: ae2
   LACP state:
                     Role
                                                        Aggr Timeout Activity
                            Exp
                                   Def Dist Col
                                                  Syn
     xe-1/0/1
                    Actor
                                                         Yes
                             No
                                   Yes
                                          No
                                              No
                                                   No
                                                                 Fast
                                                                         Active
      xe-1/0/1
                  Partner
                              No
                                   Yes
                                              No
                                                         Yes
                                                                 Fast
                                                                        Passive
                                          No
                                                   No
   LACP protocol:
                         Receive State Transmit State
                                                                 Mux State
      xe-0/0/1
                                          Fast periodic Collecting distributing
                                Current
      xe-1/0/1
                                Current
                                          Fast periodic Collecting distributing
Aggregated interface: ae3
   LACP state:
                     Role
                             Exp
                                   Def Dist
                                             Col
                                                  Syn
                                                        Aggr
                                                             Timeout Activity
      xe-0/0/2
                    Actor
                             No
                                   Yes
                                          No
                                              No
                                                   No
                                                         Yes
                                                                 Fast
                                                                         Active
     xe-0/0/2
                  Partner
                              No
                                   Yes
                                          No
                                              No
                                                   No
                                                         Yes
                                                                 Fast
                                                                       Passive
   LACP protocol:
                         Receive State Transmit State
                                                                 Mux State
      xe-0/0/2
                                          Fast periodic Collecting distributing
                                Current
```

This output means that both devices and all related interfaces are properly participating in LACP negotiations.

Verifying Aggregated Ethernet Interfaces in MC-LAG

Purpose

Verify that all of the ae interfaces are configured properly in the MC-LAG.

Action

1. Verify the ae interfaces on EX9200-A.

```
user@EX9200-A> show interfaces mc-ae
Member Link
                            : ae2
Current State Machine's State: mcae active state
Configuration Error Status : No Error
Local Status
                           : active
Local State
                          : up
Peer Status
                           : active
Peer State
                            : up
    Logical Interface
                          : ae2.0
    Topology Type
                           : bridge
    Local State
                          : up
    Peer State
                            : up
    Peer Ip/MCP/State : 172.16.32.6 ae1.0 up
Member Link
                            : ae3
Current State Machine's State: mcae active state
Configuration Error Status : No Error
Local Status
                           : active
Local State
                            : up
Peer Status
                            : active
Peer State
                            : up
    Logical Interface
                          : ae3.0
    Topology Type
                          : bridge
    Local State
                            : up
    Peer State
                            : up
    Peer Ip/MCP/State
                            : 172.16.32.6 ae1.0 up
```

2. Verify the ae interfaces on EX9200-B.

user@EX9200-B> show interface mc-ae
Member Link : ae2

Current State Machine's State: mcae active state

Configuration Error Status : No Error
Local Status : active
Local State : up
Peer Status : active
Peer State : up
Logical Interface : ae2.0

Topology Type : bridge
Local State : up
Peer State : up

Peer Ip/MCP/State : 172.16.32.5 ae1.0 up

Member Link : ae3

Current State Machine's State: mcae active state

Configuration Error Status : No Error
Local Status : active
Local State : down
Peer Status : active
Peer State : down
Logical Interface : ae3.0
Topology Type : bridge

Local State : up
Peer State : up

Peer Ip/MCP/State : 172.16.32.5 ae1.0 up

Meaning

This output means that the mc-ae interfaces on each device are up and active.

Verifying MAC Learning on MC-LAG

Purpose

Verify that MAC learning between devices is happening in the MC-LAG.

Action

1. Show the Ethernet switching table on EX9200-A.

```
user@EX9200-A> show ethernet-switching table
MAC flags (S - static MAC, D - dynamic MAC, L - locally learned, P - Persistent static, C -
Control MAC
           SE - statistics enabled, NM - non configured MAC, R - remote PE MAC, O - ovsdb MAC)
Ethernet switching table : 2 entries, 2 learned
Routing instance : EVPN-2
    Vlan
                       MAC
                                           MAC
                                                       Age
                                                              Logical
NH
         RTR
                                            flags
                                                              interface
    name
                       address
Index
         ID
   v100
                       10:0e:7e:b1:01:80
                                                             pip-7.040010000000
                                           DC
1048580 1048580
    v100
                       4c:96:14:e7:fd:81 DRC
                                                             ae10.200
0
          0
```

2. Show the Ethernet switching table on EX9200-B.

```
user@EX9200-B> show ethernet-switching table
MAC flags (S - static MAC, D - dynamic MAC, L - locally learned, P - Persistent static, C -
Control MAC
           SE - statistics enabled, NM - non configured MAC, R - remote PE MAC, O - ovsdb MAC)
Ethernet switching table : 2 entries, 2 learned
Routing instance : EVPN-2
Vlan name
                   MAC address
                                               MAC flags
                                                               Age
                                                                      Logical
interface
                       NH Index
                                      RTR
ΙD
         v100
                            10:0e:7e:b1:01:80
                                                         DC
                               1048581
                                              1048580
pip-7.060010000000
                    4c:96:14:e7:fd:81
                                                D
v100
ae10.200
                                               0
```

This output means that the MAC addresses are properly learned within the shared VLANs defined in the MC-LAG.

Verifying VRRP in MC-LAG

Purpose

Verify that VRRP is up and active between the devices in the MC-LAG.

Action

1. Confirm that VRRP is up and active on EX9200-A.

```
user@EX9200-A> show vrrp

Interface State Group VR state VR Mode Timer Type Address
irb.100 up 1 master Active A 0.789 lcl 192.168.100.2
vip 192.168.100.1
```

In this example, Switch A is the primary VRRP member.

2. Confirm that VRRP is up and active on EX9200-B.

In this example, Switch B is the backup VRRP member.

Meaning

This output means that VRRP is up and running properly.

Verifying OSPF on MC-LAG

Purpose

Verify that OSPF is properly up and running with MC-LAG.

Action

1. Show the OSPF neighbors on EX9200-A.

user@EX9200-	A> show ospf neighbor				
Address	Interface	State	ID	Pri	Dead
172.16.32.10	ae0.0	Full	172.16.32.6	128	33

2. Show the OSPF routing table on EX9200-A.

```
user@EX9200-A> show ospf route
Topology default Route Table:
Prefix
                  Path Route
                                                                 Nexthop
                                   NH
                                            Metric NextHop
                  Type Type
                                   Туре
                                                   Interface
                                                                 Address/LSP
172.16.32.6
                  Intra Router
                                   ΙP
                                                 1 ae0.0
                                                                 172.16.32.10
172.16.32.5/32
                 Intra Network
                                   ΙP
                                                 0 100.0
172.16.32.6/32
                  Intra Network
                                   ΙP
                                                 1 ae0.0
                                                                 172.16.32.10
172.16.32.8/30
                  Intra Network
                                   ΙP
                                                 1 ae0.0
```

3. Show the OSPF neighbors on EX9200-B.

user@EX9200-B> show ospf neighbor	
Address Interface State ID Pri	Dead
172.16.32.9 ae0.0 Full 172.16.32.5 128	31

4. Show the OSPF routing table on EX9200-B.

user@EX9200-B> sh Topology default	•			
Prefix	Path Route	NH	Metric NextHop	Nexthop
	Type Type	Type	Interface	Address/LSP
172.16.32.5	Intra Router	IP	1 ae0.0	172.16.32.9
172.16.32.5/32	Intra Network	IP	1 ae0.0	172.16.32.9
172.16.32.6/32	Intra Network	IP	0 100.0	
172.16.32.8/30	Intra Network	IP	1 ae0.0	

The output shows that the neighboring devices are fully adjacent.

Verifying that Configuration Consistency Check Passed

Purpose

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.

Action

1. Show the list of committed MC-LAG parameters that passed or failed configuration consistency check on EX9200-A.

onfiguration Item	Enforcement Level	Local Value	Peer
Value Result			
ICL interface	Mandatory	ae1	
ae1 PAS			
rstp-bridge-priority	Desirable	0	
PAS			
service-id	Mandatory	1	
1 PAS			
session-establishment-hol	time Mandatory	300	
300 PAS			
local-ip-addr	Mandatory	172.16.32.5	
172.16.32.6 PAS			
backup-liveness-detection	Mandatory	10.92.76.4	
10.92.76.2 PAS			
iccp/bfd multiplier	Mandatory	4	
4 PAS			
ofd minimum-interval	Mandatory	2000	
2000 PAS			
session-establishment-hol	time Mandatory	50	
50 PAS			

VALUA D.	1.	Enforcement Level	Local Value	Peer
	sult			
		Mandatany	2	
lacp admin-key 2	PASS	Mandatory	2	
z lacp system-id	LASS	Mandatory	00:01:02:03:04:07	
00:01:02:03:04:07	PASS	Tianua cor y	00.01.02.03.04.07	
lacp periodic	1 755	Mandatory	0	
0	PASS	Tiania con y	v	
lacp mode		Mandatory	0	
0	PASS	, and the same of		
prefer-status-contro		Desirable	TRUE	
	PASS			
mcae status-control		Mandatory	standby	
active	PASS			
mcae deployment mode	е	Mandatory	active-active	active-
active PASS				
mcae chassis-id		Mandatory	0	
1	PASS			
mcae redundancy-grou	up	Mandatory	1	
1	PASS			
Local Logical Interi	face:ae2.0			
Peer Logical Interfa				
Configuration Item		Enforcement Level	Local Value	Peer
	sult			
		Mandatory	100	
	PASS	Mandatory	100	
vlan membership		Mandatory Mandatory	100 trunk	
vlan membership 100				
vlan membership 100 interface-mode	PASS PASS			
vlan membership 100 interface-mode trunk	PASS PASS rface:ae3			
vlan membership 100 interface-mode trunk Local Physical Inter	PASS PASS rface:ae3	Mandatory		Peer
vlan membership 100 interface-mode trunk Local Physical Inter Peer Physical Inter	PASS PASS rface:ae3	Mandatory	trunk	Peer
vlan membership 100 interface-mode trunk Local Physical Inter Peer Physical Inter	PASS PASS rface:ae3 face :ae3	Mandatory	trunk	Peer
vlan membership 100 interface-mode trunk Local Physical Inter Peer Physical Inter Configuration Item	PASS PASS rface:ae3 face :ae3	Mandatory Enforcement Level	trunk Local Value	Peer
vlan membership 100 interface-mode trunk Local Physical Inter Peer Physical Interf Configuration Item Value Res	PASS PASS rface:ae3 face :ae3	Mandatory Enforcement Level	trunk Local Value	Peer

lacp system-id		Mandatory	00:01:02:03:04:08	
00:01:02:03:04:08	PASS		•	
lacp periodic	DACC	Mandatory	0	
0	PASS	Manufatan	٥	
lacp mode	DACC	Mandatory	0	
0	PASS	Destudi.	TDUE	
prefer-status-contro		Desirable	TRUE	
	PASS	Mandatan	- 4 - 11 db	
mcae status-control	DACC	Mandatory	standby	
active	PASS	Mandatory	active-active	activo-
mcae deployment mode active PASS		Mandatory	active-active	active-
mcae chassis-id		Mandatory	۵	
1	PASS	Mandatory	0	
		Mandatory	1	
<pre>mcae redundancy-grou 1</pre>	PASS	Mandatory	ı	
1	FA33			
Local Logical Interf	ace:ae3 0			
Peer Logical Interfa				
Configuration Item		Enforcement Level	Local Value	Peer
Value Res	ult	2 0. 0000 20.01	20002 70200	. 66.
				
vlan membership		Mandatory	100	
100	PASS			
interface-mode		Mandatory	trunk	
trunk	PASS			
Local VLAN:v100				
Peer VLAN :v100				
Local IRB:irb.100				
Peer IRB :irb.100				
Configuration Item		Enforcement Level	Local Value	Peer
Value Res	ult			, 551
vrrp-group id		Mandatory	1	
1	PASS	J J		
ipv4 address		Mandatory	192.168.100.2/24	
192.168.100.3/24	PASS			

2. Show the list of committed MC-LAG parameters that passed or failed configuration consistency check on EX9200-B.

Configuration Item		Enforcement Level	Local Value	Peer
Value Re	sult			
ICL interface		Mandatory	ae1	
ae1	PASS			
rstp-bridge-priorit	у	Desirable	0	
0	PASS			
service-id		Mandatory	1	
1	PASS			
session-establishme	nt-hold-time	Mandatory	300	
300	PASS			
local-ip-addr		Mandatory	172.16.32.6	
172.16.32.5	PASS			
backup-liveness-det	ection	Mandatory	10.92.76.2	
10.92.76.4	PASS			
iccp/bfd multiplier		Mandatory	4	
4	PASS			
bfd minimum-interva	1	Mandatory	2000	
2000	PASS			
session-establishme	nt-hold-time	Mandatory	50	
50	PASS			
Local Physical Inte	rface:ae2			
Peer Physical Inter	face :ae2			
Configuration Item		Enforcement Level	Local Value	Peer
Value Re	sult			
lacp admin-key		Mandatory	2	
2	PASS			
lacp system-id		Mandatory	00:01:02:03:04:07	
00:01:02:03:04:07	PASS			
lacp periodic		Mandatory	0	
0	PASS			
lacp mode		Mandatory	0	
0	PASS			

standby	PASS			
mcae deployment mode		Mandatory	active-active	active-
active PASS				
mcae chassis-id		Mandatory	1	
0	PASS			
mcae redundancy-group	р	Mandatory	1	
1	PASS			
prefer-status-contro	l-active	Desirable		
TRUE	PASS			
Local Logical Interfa	ace:ae2.0			
Peer Logical Interfac				
Configuration Item		Enforcement Level	Local Value	Peer
Value Res	ult			
vlan membership		Mandatory	100	
100	PASS	·		
interface-mode		Mandatory	trunk	
trunk	PASS	·		
Local Physical Inter	faceraes			
Peer Physical Interfa				
Configuration Item	acc .acs	Enforcement Level	Local Value	Peer
Value Resi	u1+	Lift of collectic Level	Local value	1 001
	uic			
		Mandatory	3	
lacp admin-key	PASS	nanuator y	3	
3	rass	Mandatanu	00.01.02.02.04.00	
lacp system-id 00:01:02:03:04:08	DACC	Mandatory	00:01:02:03:04:08	
	PASS	Mandatani	٥	
lacp periodic	DACC	Mandatory	0	
0	PASS	W. L.		
lacp mode	D. 0.0	Mandatory	0	
0	PASS			
mcae status-control		Mandatory	active	
standby	PASS			
mcae deployment mode		Mandatory	active-active	active-
active PASS				
mcae chassis-id		Mandatory	1	
0	PASS			
mcae redundancy-group	р	Mandatory	1	
1	PASS			

prefer-status-contro		Desirable		
TRUE	PASS			
Local Logical Interf	Face ae3 0			
Peer Logical Interfa				
Configuration Item		Enforcement Level	Local Value	Peer
_	sult			
vlan membership		Mandatory	100	
100	PASS			
interface-mode		Mandatory	trunk	
trunk	PASS			
Local VLAN:v100				
Peer VLAN :v100				
Local IRB:irb.100				
Peer IRB :irb.100		- 0		_
Configuration Item	1.	Enforcement Level	Local Value	Peer
	sult			
		Mandatory	1	
vrrp-group id 1	PASS	rialiuatoi y	I	
ipv4 address	1 733	Mandatory	192.168.100.3/24	
192.168.100.2/24	PASS	Handatol y	132.100.100.3/24	
132.100.100.2/27	1700			

The output shows that all configured and committed MC-LAG parameters have passed configuration consistency check.

Verifying the Configuration Consistency Check Status for the Global Configuration

Purpose

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

Parameters specific to the ICL, MC-LAG interfaces, and VLAN and VRRP configurations are shown later in this document.

Action

1. Show the list of committed global configuration parameters that passed or failed configuration consistency check on EX9200-A.

The output below shows all of the parameters that directly affect the MC-LAG configuration.

		mc-lag configuration-		_
Configuration Ite	m	Enforcement Level	Local Value	Peer
Value	Result			
	-			
ICL interface		Mandatory	ae1	
ae1	PASS			
rstp-bridge-prior	ity	Desirable	0	
0	PASS			
service-id		Mandatory	1	
1	PASS			
session-establish	ment-hold-time	Mandatory	300	
300	PASS			
local-ip-addr		Mandatory	172.16.32.5	
172.16.32.6	PASS			
backup-liveness-d	etection	Mandatory	10.92.76.4	
10.92.76.2	PASS			

iccp/bfd multiplier	Mandatory	4	
4 PASS			
bfd minimum-interval	Mandatory	2000	
2000 PASS			
session-establishment-hold-time	Mandatory	50	
50 PASS			

2. Show the list of committed global configuration parameters that passed or failed configuration consistency check on EX9200-B

Configuration Item		Enforcement Level	Local Value	Peer
/alue Res	ult			
ICL interface		Mandatory	ae1	
ae1	PASS			
rstp-bridge-priority		Desirable	0	
)	PASS			
service-id		Mandatory	1	
I	PASS			
session-establishmen	t-hold-time	Mandatory	300	
300	PASS			
local-ip-addr		Mandatory	172.16.32.6	
172.16.32.5	PASS			
oackup-liveness-dete	ction	Mandatory	10.92.76.2	
10.92.76.4	PASS			
iccp/bfd multiplier		Mandatory	4	
1	PASS			
ofd minimum-interval		Mandatory	2000	
2000	PASS			
session-establishmen	t-hold-time	Mandatory	50	
50	PASS			

Meaning

The output shows that the committed global configuration related to MC-LAG have passed configuration consistency check.

Verifying the Configuration Consistency Check Status for the Interchassis Control Link

Purpose

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

Action

1. Show the list of committed ICL configuration parameters that passed or failed configuration consistency check on EX9200-A

```
user@EX9200-A> show multi-chassis mc-lag configuration-consistency icl-config
Local Physical Interface:ae1
Peer Physical Interface :ae1
Local Logical Interface:ae1.0
Peer Logical Interface :ae1.0
Configuration Item
                                   Enforcement Level Local Value
                                                                            Peer
Value
                Result
_____
                     -----
vlan membership
                                   Mandatory
                                                       100
                                                                            100
                                                                                      PASS
interface-mode
                                   Mandatory
                                                       trunk
trunk
                     PASS
```

2. Show the list of committed ICL configuration parameters that passed or failed configuration consistency check on EX9200-B

```
user@EX9200-B> show multi-chassis mc-lag configuration-consistency icl-config
Local Physical Interface:ae1
Peer Physical Interface :ae1
```

Local Logical Interface:ae1.0					
Peer Logical Inter	face :ae1.0				
Configuration Item	Configuration Item		Local Value	Peer	
Value R	Result				
vlan membership		Mandatory	100	100	PASS
interface-mode		Mandatory	trunk		
trunk	PASS				

The output shows that the committed MC-LAG parameters related to the ICL have passed configuration consistency check.

Verifying the Configuration Consistency Check Status for the MC-LAG Interfaces

Purpose

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

Action

1. Show the list of committed MC-LAG interface configuration parameters that passed or failed configuration consistency check on EX9200-A.

Local Physical Inter Peer Physical Interf				
Configuration Item	acc .acz	Enforcement Level	Local Value	Peer
Value Result		Liff of Cement Level	Local value	1 661
	Juli			
lacp admin-key		Mandatory	2	
2	PASS	·		
lacp system-id		Mandatory	00:01:02:03:04:07	
00:01:02:03:04:07	PASS			
lacp periodic		Mandatory	0	
9	PASS			
lacp mode		Mandatory	0	
9	PASS			
prefer-status-control-active		Desirable	TRUE	
	PASS			
mcae status-control		Mandatory	standby	
active	PASS			
mcae deployment mode	9	Mandatory	active-active	active-
active PASS				
mcae chassis-id		Mandatory	0	
1	PASS			
mcae redundancy-group		Mandatory	1	
1	PASS			
Local Logical Interf	face:ae2 A			
Peer Logical Interface :ae2.0 Configuration Item		Enforcement Level	Local Value	Peer
Value Result		Em or coment Level	1001 (0100	1 001

100	PASS			
interface-mode		Mandatory	trunk	
trunk	PASS			
Local Physical Inte	rface:ae3			
Peer Physical Inter	face :ae3			
Configuration Item		Enforcement Level	Local Value	Peer
Value Re:	sult			
lacp admin-key		Mandatory	3	
3	PASS			
lacp system-id		Mandatory	00:01:02:03:04:05	
00:01:02:03:04:05	PASS			
lacp periodic		Mandatory	0	
0	PASS			
lacp mode		Mandatory	0	
0	PASS			
prefer-status-contro	ol-active	Desirable	TRUE	
	PASS			
mcae status-control		Mandatory	standby	
active	PASS			
mcae deployment mode	e	Mandatory	active-active	active-
active PASS				
mcae chassis-id		Mandatory	0	
1	PASS			
mcae redundancy-gro	qu	Mandatory	1	
1	PASS			
Local Logical Inter	face:ae3.0			
Peer Logical Interfa	ace :ae3.0			
Configuration Item		Enforcement Level	Local Value	Peer
Value Res	sult			
vlan membership		Mandatory	100	
100	PASS			
interface-mode		Mandatory	trunk	
trunk	PASS			

2. Show the list of committed MC-LAG interface configuration parameters that passed or failed configuration consistency check on EX9200-B.

Configuration Item /alue Res	sult	Enforcement Level	Local Value	_
	sult		Local value	Peer
acn admin-kay				
.acp admiti-key		Mandatory	2	
2	PASS			
acp system-id		Mandatory	00:01:02:03:04:05	
00:01:02:03:04:05	PASS			
acp periodic		Mandatory	0	
)	PASS			
acp mode		Mandatory	0	
)	PASS			
ncae status-control		Mandatory	active	
standby	PASS			
mcae deployment mode		Mandatory	active-active	active-
active PASS				
ncae chassis-id		Mandatory	1	
)	PASS			
ncae redundancy-grou	ıp	Mandatory	1	
	PASS			
orefer-status-contro	ol-active	Desirable		
RUE	PASS			
ocal Logical Interf	Face:ae2.0			
Peer Logical Interfa	ace :ae2.0			
Configuration Item		Enforcement Level	Local Value	Peer
/alue Res	sult			
lan membership		Mandatory	100	
00	PASS			
nterface-mode		Mandatory	trunk	
runk	PASS			

Configuration Item		Enforcement Level	Local Value	Peer
Value Re:	sult			
lacp admin-key		Mandatory	3	
3	PASS			
lacp system-id		Mandatory	00:01:02:03:04:08	
00:01:02:03:04:08	PASS			
lacp periodic		Mandatory	0	
0	PASS			
lacp mode		Mandatory	0	
0	PASS			
mcae status-control		Mandatory	active	
standby	PASS			
mcae deployment mod	е	Mandatory	active-active	active-
active PASS				
mcae chassis-id		Mandatory	1	
0	PASS			
mcae redundancy-gro	up	Mandatory	1	
1	PASS			
prefer-status-contr	ol-active	Desirable		
TRUE	PASS			
Local Logical Inter				
Peer Logical Interf	ace :ae3.0	5 C () 1		
Configuration Item	1.	Enforcement Level	Local Value	Peer
	sult			
		Mandatar	100	
vlan membership	DACC	Mandatory	100	
100	PASS	Wandatan	f al.	
interface-mode	DAGG	Mandatory	trunk	
trunk	PASS			

The output shows that the committed MC-LAG parameters related to the MC-AE interfaces have passed configuration consistency check.

Verifying the Configuration Consistency Check Status for the VLAN Configuration

Purpose

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

Action

1. Show the list of committed VLAN configuration parameters that passed or failed configuration consistency check on EX9200-A.

```
user@EX9200-A> show multi-chassis mc-lag configuration-consistency vlan-config
Local VLAN:v100
Peer VLAN :v100
Local IRB:irb.100
Peer IRB :irb.100
Configuration Item
                                    Enforcement Level Local Value
                                                                              Peer
                Result
_____
vrrp-group id
                                    Mandatory
                     PASS
ipv4 address
                                    Mandatory
                                                       192.168.100.2/24
192.168.100.3/24
                     PASS
```

2. Show the list of committed VLAN configuration parameters that passed or failed configuration consistency check on EX9200-B.

```
user@EX9200-B> show multi-chassis mc-lag configuration-consistency vlan-config
Peer VLAN :v100
```

```
Local IRB:irb.100
Peer IRB :irb.100
Configuration Item Enforcement Level Local Value Peer
Value Result
-----
vrrp-group id Mandatory 1
1 PASS
ipv4 address Mandatory 192.168.100.3/24
192.168.100.2/24 PASS
```

The output shows that the committed MC-LAG parameters related to the VLAN and IRB configurations have passed configuration consistency check.

Verifying the Configuration Consistency Check Status for VRRP

Purpose

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.

Action

1. Show the list of committed VRRP configuration parameters that passed or failed configuration consistency check on EX9200-A.

```
user@EX9200-A> show multi-chassis mc-lag configuration-consistency vrrp-config
Local VRRP Group:1

Peer VRRP Group :1
Configuration Item Enforcement Level Local Value Peer
Value Result
```

```
vrrp-group virtual-address Mandatory 192.168.100.001
192.168.100.001 PASS
vrrp-group priority Mandatory 150
100 PASS
```

2. Show the list of committed VRRP configuration parameters that passed or failed configuration consistency check on EX9200-B.

```
user@EX9200-B> show multi-chassis mc-lag configuration-consistency vrrp-config
Local VRRP Group:1
Peer VRRP Group :1
Configuration Item
                                  Enforcement Level Local Value
                                                                         Peer
Value
               Result
_____
                                  Mandatory
                                                    192.168.100.001
vrrp-group virtual-address
192.168.100.001
                    PASS
                                  Mandatory
                                                     100
vrrp-group priority
150
                    PASS
```

Meaning

The output shows that the committed MC-LAG parameters related to VRRP configuration have passed configuration consistency check.

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

IN THIS SECTION

- Requirements | 281
- Overview | 282
- Configuration | 288

Verification | 301

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. 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

IN THIS SECTION

Topology | 283

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 12 on page 283.

Figure 12: Supported Topology for an MC-LAG on an FCoE Transit Switch

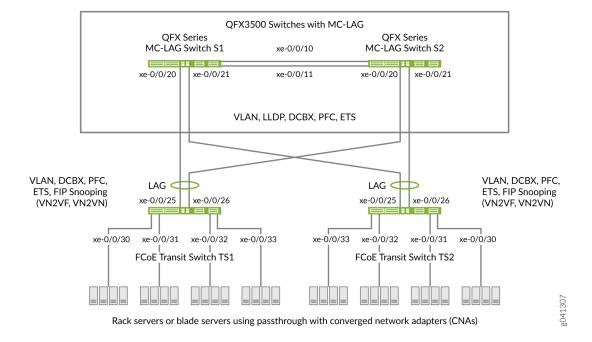


Table 7 on page 283 shows the configuration components for this example.

Table 7: 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 7: 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 fcoetrusted, 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 fcoetrusted, 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 fcoetrusted, 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 fcoetrusted, 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

Table 7: Components of the CoS for FCoE Traffic Across an MC-LAG Configuration Topology *(Continued)*

Component	Settings
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%
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

Table 7: Components of the CoS for FCoE Traffic Across an MC-LAG Configuration Topology (Continued)

Component	Settings
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

IN THIS SECTION

- CLI Quick Configuration | 288
- Configuring MC-LAG Switches S1 and S2 | 290
- Configuring FCoE Transit Switches TS1 and TS2 | 293
- Results | 297

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

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
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
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

IN THIS SECTION

- Verifying That the Output Queue Schedulers Have Been Created | 301
- Verifying That the Priority Group Output Scheduler (Traffic Control Profile) Has Been Created | 302
- Verifying That the Forwarding Class Set (Priority Group) Has Been Created | 303
- Verifying That Priority-Based Flow Control Has Been Enabled | 304
- Verifying That the Interface Class of Service Configuration Has Been Created | 305
- Verifying That the Interfaces Are Correctly Configured | 308
- Verifying That FIP Snooping Is Enabled on the FCoE VLAN on FCoE Transit Switches TS1 and TS2
 Access Interfaces | 311
- Verifying That the FIP Snooping Mode Is Correct on FCoE Transit Switches TS1 and TS2 | 312
- Verifying That IGMP Snooping Is Disabled on the FCoE VLAN | 313

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

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
                                  1
                                       <default-drop-profile>
  Low
                 any
  Medium high
                                  1
                                       <default-drop-profile>
                 any
                                       <default-drop-profile>
  High
                                  1
                  any
```

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
               Disabled
 000
 001
               Disabled
               Disabled
 010
               Enabled
                            2500
 011
  100
               Disabled
               Disabled
 101
               Disabled
 110
               Disabled
 111
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;
}
```

```
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;
                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

Example: Configuring CoS PFC for FCoE Traffic

Example: EVPN-MPLS Interworking With an MC-LAG Topology

IN THIS SECTION

- Requirements | 314
- Overview and Topology | 315
- PE1 and PE2 Configuration | 317
- PE3 Configuration | 333

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.

Requirements

This example uses the following hardware and software components:

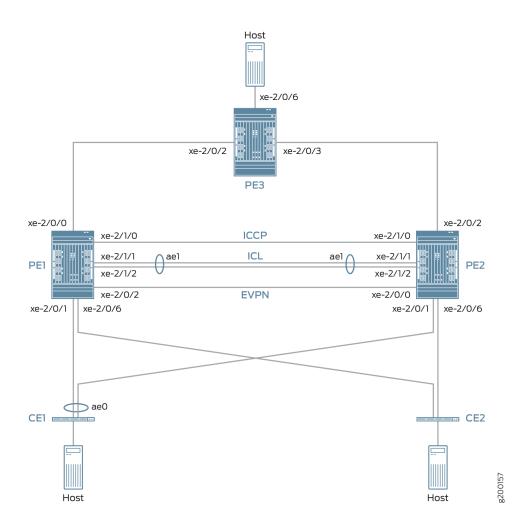
- Three EX9200 switches:
 - PE1 and PE2, which both function as MC-LAG peers in the MC-LAG topology and EVPN BGP peers in the EVPN-MPLS overlay network.
 - PE3, which functions as an EVPN BGP peer in the EVPN-MPLS overlay network.
- The EX9200 switches are running Junos OS Release 17.4R1 or later software.

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 13 on page 315 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 13: EVPN-MPLS Interworking With an MC-LAG Topology



The topology in Figure 13 on page 315 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 13 on page 315 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 8 on page 316 outlines key MC-LAG and EVPN (BGP and MPLS) attributes configured on PE1, PE2, and PE3.

Table 8: 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			

Table 8: Key MC-LAG and EVPN (BGP and MPLS) Attributes Configured on PE1, PE2, and PE3 (Continued)

Key Attributes	PE1	PE2	PE3
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

IN THIS SECTION

- CLI Quick Configuration | 318
- PE1: Configuring MC-LAG | 323
- PE1: Configuring EVPN-MPLS | 326
- PE2: Configuring MC-LAG | 328
- PE2: Configuring EVPN-MPLS | 331

To configure PE1 and PE2, perform these tasks:

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
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 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 1dp interface 100.0
```

```
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 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 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
```

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
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 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 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
```

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 00:11:11:11:11:11:11:11:11:11:11
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 00:11:22:22:22:22:22:22:22:22:22
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 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 13-interface irb.2
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 13-interface irb.3
```

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 00:11:11:11:11:11:11:11:11:11:11
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.

```
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 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 13-interface irb.2
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 13-interface irb.3
```

PE3 Configuration

IN THIS SECTION

- CLI Quick Configuration | 334
- PE3: Configuring EVPN-MPLS | 335

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 1dp interface 100.0
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 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 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
```

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 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 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 yrf-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 vlan-id 3
```



Other MC-LAG Configurations

Other MC-LAG Configurations | 340

Other MC-LAG Configurations

IN THIS SECTION

- Configuring Active-Active Bridging and VRRP over IRB in Multichassis Link Aggregation on MX Series
 Routers | 340
- Configuring IGMP Snooping in MC-LAG Active-Active Mode | 346
- Configuring Manual and Automatic Link Switchover for MC-LAG Interfaces on MX Series Routers | 347
- Forcing MC-LAG Links or Interfaces with Limited LACP Capability to Be Up | 349
- Increasing ARP and Network Discovery Protocol Entries for Enhanced MC-LAG and Layer 3 VXLAN Topologies | 350
- Synchronizing and Committing Configurations | 360

Configuring Active-Active Bridging and VRRP over IRB in Multichassis Link Aggregation on MX Series Routers

IN THIS SECTION

- Configuring MC-LAG | 341
- Configuring the Interchassis Link-Protection Link | 342
- Configuring Multiple Chassis | 342
- Configuring the Service ID | 343
- Configuring IGMP Snooping for Active-Active MC-LAG | 345

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:

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

```
service-id 10;
}
}
}
```

NOTE: Using a service name per bridge domain is not supported.

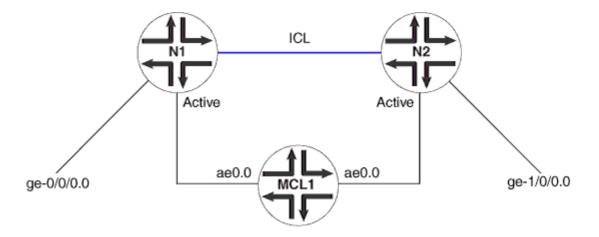
The bridge-level service ID is required to link related bridge domains across peers, and should be configured with the same value. The **service-id** values share the name space across all bridging and routing instances, and across peers. Thus, duplicate values for service IDs are not permitted across 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 14 on page 344, 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 14: N1 and N2 for the Same Service with Same Service ID

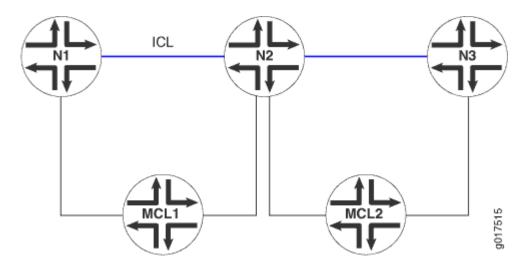


The following configuration restrictions apply:

- 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 15 on page 345, 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 15: 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 multicastsnooping-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.
- 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.

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-etheroptions 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 & 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.

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.

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

IN THIS SECTION

- Understanding the Need for an Increase in ARP and Network Discovery Protocol (NDP) Entries | 351
- Increasing ARP and Network Discovery Protocol Entries for Enhanced MC-LAG Using IPv4
 Transport | 351
- Increasing ARP and Network Discovery Protocol Entries for Enhanced MC-LAG Using IPv6
 Transport | 353
- Increasing ARP for EVPN-VXLAN Gateway for Border-Leaf in Edge Routed Bridge (ERB) or Spine in Centrally Routed Bridge (CRB) for IPv4 Tenant Traffic | 355
- Increasing ARP and Network Discovery Protocol Entries for EVPN-VXLAN gateway for Border-Leaf in Edge Routed Bridge (ERB) or Spine in Centrally Routed Bridge (CRB) for IPv6 Tenant Traffic | 357

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.

Increasing ARP and Network Discovery Protocol Entries for Enhanced MC-LAG 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 12-learning]
user@switch# set global-mac-table-aging-time seconds
```

For example:

```
[edit protocols 12-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 12-learning]
user@switch# set global-mac-ip-table-aging-timeseconds
```

For example:

```
[edit protocols 12-learning]
user@switch# set global-mac-ip-table-aging-time 1200
```

8. Reboot the device in order for these changes to take effect.

```
user@switch# request system reboot
```

Increasing ARP and Network Discovery Protocol Entries for Enhanced MC-LAG 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 12-learning]
user@switch# set global-mac-table-aging-time seconds
```

For example:

```
[edit protocols 12-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 12-learning]
user@switch# set global-mac-ip-table-aging-timeseconds
```

```
[edit protocols 12-learning]
user@switch# set global-mac-ip-table-aging-time 1200
```

8. Reboot the device in order for these changes to take effect.

```
user@switch# request system reboot
```

Increasing ARP for EVPN-VXLAN Gateway for Border-Leaf in Edge Routed Bridge (ERB) or Spine in Centrally Routed Bridge (CRB) for IPv4 Tenant Traffic

To increase the number of ARP entries using IPv4 tenant traffic, 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
```

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

4. On QFX10002-60C devices, 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 900
```

5. Specify the amount of time that elapses before the MAC table entries are timed out and entries are deleted from the table.

```
[edit protocols 12-learning]
user@switch# set global-mac-table-aging-time seconds
```

For example:

```
[edit protocols 12-learning]
user@switch# set global-mac-table-aging-time 3600
```

6. Specify the amount time that elapses before the entries in the MAC-IP bindings database are timed out and deleted.

```
[edit protocols 12-learning]
user@switch# set global-mac-ip-table-aging-timeseconds
```

For example:

```
[edit protocols 12-learning]
user@switch# set global-mac-ip-table-aging-time 1200
```

7. On QFX10002-60C devices, specify the amount time that elapses before the entries in the MAC-IP bindings database are timed out and deleted.

```
[edit protocols 12-learning]
user@switch# set global-mac-ip-table-aging-timeseconds
```

For example:

```
[edit protocols 12-learning]
user@switch# set global-mac-ip-table-aging-time 900
```

8. 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 12-learning]
user@switch# set global-mac-table-aging-time seconds
```

For example:

```
[edit protocols 12-learning]
user@switch# set global-mac-table-aging-time 3600
```

9. On QFX10002-60C devices, 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 12-learning]
user@switch# set global-mac-table-aging-time seconds
```

For example:

```
[edit protocols 12-learning]
user@switch# set global-mac-table-aging-time 1200
```

10. Reboot the device in order for these changes to take effect.

```
user@switch# request system reboot
```

Increasing ARP and Network Discovery Protocol Entries for EVPN-VXLAN gateway for Border-Leaf in Edge Routed Bridge (ERB) or Spine in Centrally Routed Bridge (CRB) for IPv6 Tenant Traffic

To increase the number of ARP and Network Discovery Protocol entries using IPv4 and IPv6 tenant traffic, 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. 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. Specify the amount of time that elapses before the MAC table entries are timed out and entries are deleted from the table.

```
[edit protocols 12-learning]
user@switch# set global-mac-table-aging-time seconds
```

```
[edit protocols 12-learning]
user@switch# set global-mac-table-aging-time 3600
```

5. Specify the amount time that elapses before the entries in the MAC-IP bindings database are timed out and deleted.

```
[edit protocols 12-learning]
user@switch# set global-mac-ip-table-aging-timeseconds
```

For example:

```
[edit protocols 12-learning]
user@switch# set global-mac-ip-table-aging-time 1200
```

6. 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 12-learning]
user@switch# set global-mac-table-aging-time seconds
```

For example:

```
[edit protocols 12-learning]
user@switch# set global-mac-table-aging-time 3600
```

7. 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 12-learning]
user@switch# set global-mac-table-aging-time seconds
```

For example:

```
[edit protocols 12-learning]
user@switch# set global-mac-table-aging-time 1200
```

8. Reboot the device in order for these changes to take effect.

```
user@switch# request system reboot
```

Synchronizing and Committing Configurations

IN THIS SECTION

- Configure Devices for Configuration Synchronization | 360
- Create a Global Configuration Group | 363
- Create a Local Configuration Group | 366
- Create a Remote Configuration Group | 369
- Create Apply Groups for the Local, Remote, and Global Configurations | 371
- Synchronizing and Committing Configurations | 372
- Troubleshooting Remote Device Connections | 372

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

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.

```
[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:

3. Enable a NETCONF connection using SSH between all devices (Switch A, Switch B, Switch C, and Switch D).

```
[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.

```
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;
        13-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 {
                 {\tt 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;
                 13-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.

```
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 12-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.

```
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 12-interface ae0.0 mac 00:00:5E:00:53:00;
```

```
}
}

property

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 12-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 terminated.

To apply the configuration groups:

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

IN THIS SECTION

- Problem | 373
- Resolution | 373

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

Resolution

1. Verify that the SSH connection to the remote device (Switch B) is working.

root@Switch A# ssh root@Switch B

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.



MC-LAG Upgrade Guidelines

MC-LAG Upgrade Guidelines | 376

MC-LAG Upgrade Guidelines

IN THIS SECTION

Upgrade Guidelines | 376

Upgrade Guidelines

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.



Best Practices and Usage Notes

Best Practices and Usage Notes | 379

Best Practices and Usage Notes

IN THIS SECTION

- Address Resolution Protocol Active-Active MC-LAG Support Methodology | 379
- DHCP Relay with Option 82 | 380
- Failure Handling | 380
- ICCP and ICL | 383
- IGMP Snooping on an Active-Active MC-LAG | 385
- Init Delay Time | 385
- Label Swap Function | 385
- Miswiring Detection Guidelines | 385
- Redundancy Groups | 387
- Segment Routing | 387
- Status Control | 387
- Virtual Router Redundancy Protocol (VRRP) over IRB and MAC Address Synchronization | 388
- VLANs | 389
- VRRP Active-Standby Support | 389

Address Resolution Protocol Active-Active MC-LAG Support Methodology

ARP and MAC address tables normally stay synchronized in MC-LAG configurations, but might get out of sync under certain network conditions (such as link flapping). To ensure these tables remain in sync while those conditions are being resolved, we recommend enabling the arp-12-validate statement on IRB interfaces in an MC-LAG configuration.

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.

- 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.

Failure Handling

Table 9 on page 381 describes the different ICCP failure scenarios for EX9200 switches. The dash means that the item is not applicable.

Table 9: 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.
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 10 on page 381 describes the different ICCP failure scenarios for QFX Series switches. The dash means that the item is not applicable.

Table 10: 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 for both active and standby MC-AE interfaces.
Down	Down or Up	Inactive	No change in LACP system ID.

Table 10: ICCP Failure Scenarios for QFX Series Switches (Continued)

ICCP Connection Status	ICL Status	Backup Liveness Peer Status	Action on Multichassis Aggregated Ethernet Interface with Status Set to Standby
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 primary and backup Routing Engines. This ensures that the connection is not reset during *GRES* in the remote peer.

For example, on the primary 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 primary 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;
    }
}
```

ICCP and ICL

BEST PRACTICE: We recommend that you use separate ports and choose different Flexible PIC Concentrators (*FPC*s) 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 ICCP and the ICL on the same interface, and configure prefer status control as active on the active MC-LAG peer.

For example, issue the prefer-status-control-active statement on the active MC-LAG peer.

• 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.

NOTE: On QFX 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.

- For better convergence during a MC-LAG peer (status-control of active) reboot, we
 recommend that you configure the icl-down-delay time instead of the hold-time. The icldown-delay time is the number of seconds that elapse between when the interchassis link
 (ICL) goes down and the multichassis aggregated Ethernet interfaces (MCAEs) change to
 standby mode.
- 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-detection statement on the management interface
 (fxp0) only.

NOTE: Liveness detection does not work if the management interface is within the management instance.

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.

IGMP Snooping on an Active-Active MC-LAG

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

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.

Init Delay Time

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.

Label Swap Function

NOTE: QFX Series switches configured as MC-LAG peers do not supportvlans label swap function.

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.

Configuration Guidelines and Caveats

- Configure the IP address on the active MC-LAG peer with a high IP address or a high DR priority. To ensure that the active MC-LAG peer retains the DR membership designation if PIM neighborship with the peer goes down.
- Using Bidirectional Forwarding Detection (BFD) and RVI or IRB MAC synchronization together is not supported because ARP fails.
- When using RVI or IRB MAC synchronization, make sure that you configure the primary IP address on both MC-LAG peers. Doing this ensures that both MC-LAG peers cannot become assert winners.
- The number of BFD sessions on RVIs or IRBs with PIM enabled is restricted to 100. Also,
 If you have more than 100 RVIs or IRBs configured, do not configure BFD, and make sure
 that the hello interval is 2 seconds.

Redundancy Groups

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.

Segment Routing

NOTE: QFX Series switches configured as MC-LAG peers do not support segment routing.

Status Control

BEST PRACTICE: You can configure **prefer-status-control-active** statement with the **status-control standby** configuration to prevent the LACP mc-ae system ID from reverting to the LACP default system ID during an ICCP failure. You should only configure this option if ICCP will never go down unless the remote peer goes down.

In the event that the peer configured with **status-control active** goes down abruptly, such as during a power off, we recommended that you configure the **hold-time** interval for the interchassis control link configured as **status-control standby** with a value greater than the ICCP BFD timeout. Doing this will prevent momentary traffic loss on the peer configured as **status-control standby**. Without the **hold-time** interval configuration on the ICL, the MC-AE interface on the peer configured as **status-control standby** will momentarily move to standby during a power-off of the remote peer.

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

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.

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

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.

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.

On the QFX Series, the all option at the [edit interfaces name unit number ethernet-switching vlan] hierarchy is not supported on multichassis aggregated Ethernet (MC-AE), aggregated Ethernet (AE), Gigabit Ethernet (GE), 10-Gigabit Ethernet (XE), 40-Gigabit Ethernet (ET), and 100-Gigabit Ethernet (ET) Layer 2 Ethernet interfaces. Instead, specify a range of VLAN IDs or VLAN names. VLAN names or VLAN IDs for ICCP purposes are not supported on MC-AE, single-homed AE, GE, XE, and ET interfaces connected to servers or other network devices. VLAN names or VLAN IDs used for ICCP purposes are not supported on MC-AE multi-homed access and trunk Layer 2 interfaces, single-homed AE, GE, XE, and ET Layer 2 Ethernet access or trunk interfaces. When you configure VLANs on an MC-AE or other revenue interfaces on the device, the VLAN range is supported but without the ICCP-dedicated VLAN.

QFX Series switches (except QFX10002, QFX10008, and QFX10016), EX4600 switches, and EX4650 switches do not support service provider style of configuration for MC-LAG.

VRRP Active-Standby Support

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.

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.

Starting in Junos OS Release 15.2R1, you do not need to configure a static ARP or ND entry for the remote IRB IP address. If you have already manually configured a static ARP or ND entry and upgrade to a later release, the static entry is deleted when ICCP goes down. If you configured ICCP on the IRB static entry, then ICCP might not come up. As a workaround, you can disable the automatic creation of static ARP and ND entries by issuing the following command: **set protocols I2-learning no-mclag-ifa-sync**.

Release History Table

Release	Description
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.



Troubleshooting Multichassis Link Aggregation

Troubleshooting Multichassis Link Aggregation | 392

Configuring Interface Diagnostics Tools to Test the Physical Layer Connections | 408

Troubleshooting Multichassis Link Aggregation

IN THIS SECTION

- MAC Addresses Learned on Multichassis Aggregated Ethernet Interfaces Are Not Removed from the MAC
 Address Table | 393
- MC-LAG Peer Does Not Go into Standby Mode | 394
- Secondary MC-LAG Peer with Status Control Set to Standby Becomes Inactive | 395
- Redirect Filters Take Priority over User-Defined Filters | 395
- Operational Command Output Is Wrong | 396
- ICCP Connection Might Take Up to 60 Seconds to Become Active | 397
- MAC Address Age Learned on a Multichassis Aggregated Ethernet Interface Is Reset to Zero | 398
- MAC Address Is Not Learned Remotely in a Default VLAN | 399
- Snooping Entries Learned on Multichassis Aggregated Ethernet Interfaces Are Not Removed | 399
- ICCP Does Not Come Up After You Add or Delete an Authentication Key | 400
- Local Status Is Standby When It Should Be Active | 401
- Packets Loop on the Server When ICCP Fails | 401
- Both MC-LAG Peers Use the Default System ID After a Reboot or an ICCP Configuration Change | 402
- No Commit Checks Are Done for ICL-PL Interfaces | 403
- Double Failover Scenario | 403
- Multicast Traffic Floods the VLAN When the ICL-PL Interface Goes Down and Up | 404
- Layer 3 Traffic Sent to the Standby MC-LAG Peer Is Not Redirected to Active MC-LAG Peer | 405
- Aggregated Ethernet Interfaces Go Down | 405
- Flooding of Upstream Traffic | 406
- ARP and MAC Table Entries Become Out of Sync in an MC-LAG Configuration | 407

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

IN THIS SECTION

- Problem | 393
- Solution | 394

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)
                    00:00:5E:00:53:01 Learn(R)
                                                      0 xe-0/0/9.0
 v10
                                      Flood
                                                     - All-members
 v20
                                                     - All-members
                                      Flood
 v30
 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

IN THIS SECTION

- Problem | 394
- Solution | 394

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

IN THIS SECTION

- Problem | 395
- Solution | 395

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

IN THIS SECTION

- Problem | 396
- Solution | 396

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

IN THIS SECTION

Problem | 396

Solution | 397

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

IN THIS SECTION

- Problem | 397
- Solution | 397

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

IN THIS SECTION

- Problem | 398
- Solution | 398

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.

Solution

This is expected behavior.

MAC Address Is Not Learned Remotely in a Default VLAN

IN THIS SECTION

- Problem | 399
- Solution | 399

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

IN THIS SECTION

- Problem | 400
- Solution | 400

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

IN THIS SECTION

- Problem | 400
- Solution | 400

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

IN THIS SECTION

- Problem | 401
- Solution | 401

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

IN THIS SECTION

- Problem | 402
- Solution | 402

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

IN THIS SECTION

- Problem | 402
- Solution | 402

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

IN THIS SECTION

- Problem | 403
- Solution | 403

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

IN THIS SECTION

- Problem | 403
- Solution | 404

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

IN THIS SECTION

- Problem | 404
- Solution | 404

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

IN THIS SECTION

- Problem | 405
- Solution | 405

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

IN THIS SECTION

- Problem | 406
- Solution | 406

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

IN THIS SECTION

- Problem | 406
- Solution | 407

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

IN THIS SECTION

- Problem | 407
- Solution | 407

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-I2-validate option on the switch's IRB interface, as follows:

user@switch> set interfaces irb arp-12-validate

The arp-12-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

IN THIS SECTION

- Configuring Loopback Testing | 408
- Configuring BERT Testing | 411
- Starting and Stopping a BERT Test | 415

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, NxDSO, 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 11 on page 409 shows the loopback modes supported on the various interface types.

Table 11: 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	
/ kDS0	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	

Table 11: Loopback Modes by Interface Type (Continued)

Interface	Loopback Modes	Usage Guidelines	
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 0.152 standard)
   pseudo-2e15-o151     Pattern is 2^15 - 1 (per 0.152 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)
   ...
```

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 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-2e23-o151 Pattern is 2^23 (per 0.151 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 0.152 standard)
pseudo-2e15-o151
                     Pattern is 2<sup>15</sup> - 1 (per 0.151 standard)
                     Pattern is 2^20 - 1 (per 0.151 standard)
pseudo-2e20-o151
pseudo-2e7
                     Pattern is 2<sup>7</sup> - 1
pseudo-2e9-o153
                     Pattern is 2<sup>9</sup> - 1 (per 0.153 standard)
repeating-1-in-4
                     1 bit in 4 is set
                     1 bit in 8 is set
repeating-1-in-8
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
                               1 bit in 8 is set
repeating-1-in-8
repeating-3-in-24
                               3 bits in 24 are set
pseudo-2e9-o153
                               Pattern is 2<sup>9</sup> - 1 (per 0.153 standard)
                               Pattern is 2<sup>11</sup> - 1 (per 0.152 standard)
pseudo-2e11-o152
                               Pattern is 2<sup>15</sup> - 1 (per 0.151 standard)
pseudo-2e15-o151
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 12 on page 413 shows the BERT capabilities for various interface types.

Table 12: BERT Capabilities by Interface Type

Interface	T1 BERT	T3 BERT	Comments
12-port T1/E1 Circuit Emulation	Yes (ports 0–11)	_	Limited algorithms

Table 12: 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 timeLimited 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 timeLimited algorithmsNo 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)



Configuration Statements and Operational Commands

Junos CLI Reference Overview | 418

Junos CLI Reference Overview

We've consolidated all Junos CLI commands and configuration statements in one place. Learn about the syntax and options that make up the statements and commands and understand the contexts in which you'll use these CLI elements in your network configurations and operations.

• Junos CLI Reference

Click the links to access Junos OS and Junos OS Evolved configuration statement and command summary topics.

- Configuration Statements
- CLI Commands