

Network Configuration Example

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



Modified: 2016-12-21

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

Copyright © 2016, Juniper Networks, Inc. All rights reserved.

Juniper Networks, Junos, Steel-Belted Radius, NetScreen, and ScreenOS are registered trademarks of Juniper Networks, Inc. in the United States and other countries. The Juniper Networks Logo, the Junos logo, and JunosE are trademarks of Juniper Networks, Inc. All other trademarks, service marks, registered trademarks, 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.

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

Copyright © 2016, 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 <http://www.juniper.net/support/eula.html>. By downloading, installing or using such software, you agree to the terms and conditions of that EULA.

Table of Contents

Chapter 1	Configuring MC-LAG on EX9200 Switches in the Core for Campus Networks	5
	About This Network Configuration Example	5
	Use Case for Configuring MC-LAG on the Core for Campus Networks	6
	Use Case for Simplifying MC-LAG Configuration	6
	MC-LAG Technical Overview	7
	ICCP and ICL	8
	Active/Standby and Active/Active Modes	9
	MC-LAG Interface	10
	Additional MC-LAG Specific Configuration	12
	Data Traffic Forwarding Rules in Active/Active MC-LAG Topologies	12
	Failure Handling During a Split-Brain State	13
	Layer 2 Feature Support	14
	MAC Address Management	14
	MAC Aging	15
	Spanning Tree Protocol	15
	Layer 2 Multicast Feature Support	16
	IGMP Snooping on an Active/Active MC-LAG	16
	Layer 3 Feature Support	17
	VRRP over IRB	17
	MAC Address Synchronization	18
	Address Resolution Protocol Synchronization for Active/Active MC-LAG Support	19
	DHCP Relay with Option 82	20
	Layer 3 Multicast Feature Support	21
	PIM Operation	21
	Layer 3 Multicast Configuration Guidelines	21
	MC-LAG Upgrade Guidelines	22
	Summary of MC-LAG Configuration Guidelines	23
	Understanding Multichassis Link Aggregation Group (MC-LAG) Configuration Synchronization	23
	Understanding Configuration Groups	24
	Understanding Conditional Groups	24
	Understanding Apply Groups	25
	Understanding Peer Configuration Details for MC-LAG Configuration Synchronization	25
	Understanding How Configurations Are Synchronized Between MC-LAG Peers	25

Understanding Multichassis Link Aggregation Group (MC-LAG) Configuration Consistency Check	27
Example: Configuring Multichassis Link Aggregation on EX9200 Switches in the Core for Campus Networks	28
Example: Simplifying Multichassis Link Aggregation on EX9200 Switches in the Core for Campus Networks	53

CHAPTER 1

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

- [About This Network Configuration Example on page 5](#)
- [Use Case for Configuring MC-LAG on the Core for Campus Networks on page 6](#)
- [Use Case for Simplifying MC-LAG Configuration on page 6](#)
- [MC-LAG Technical Overview on page 7](#)
- [Example: Configuring Multichassis Link Aggregation on EX9200 Switches in the Core for Campus Networks on page 28](#)
- [Example: Simplifying Multichassis Link Aggregation on EX9200 Switches in the Core for Campus Networks on page 53](#)

About This Network Configuration Example

This network configuration example describes the configuration of multichassis LAG (MC-LAG) on EX9200 switches in the core for campus networks, discusses considerations and recommendations for MC-LAG best practices, and provides two configuration examples. The [“Example: Configuring Multichassis Link Aggregation on EX9200 Switches in the Core for Campus Networks” on page 28](#) explains how to configure a high performance and highly available connection to end users and applications. The [“Example: Simplifying Multichassis Link Aggregation on EX9200 Switches in the Core for Campus Networks” on page 53](#) explains how to simplify MC-LAG using the configuration synchronization and configuration consistency check features, which were both introduced in Junos OS Release 16.1R1 for the EX9200 switch.

Juniper Networks validated network configuration examples are extensively tested using both simulation and live network elements to ensure comprehensive validation of all published solutions. Customer use cases, common domain examples, and field experience are combined to generate prescriptive configurations to guide customer and partner implementations of Juniper Networks solutions.

Use Case for Configuring MC-LAG on the Core for Campus Networks

The core is the heart of the campus network, and in today's mission critical enterprise environments, the flow of business requires that the network is always available. Increasing traffic loads and link resiliency are key considerations for campus network builders. The multichassis LAG (MC-LAG) feature set on the Juniper Networks EX9200 family of switches is an ideal solution for providing options for optimizing link utilization and ensuring high availability in the campus core.

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.

A common campus deployment model for MC-LAG with the EX9200 positions the EX9200 at the campus core using a collapsed core and aggregation model where access layer switches are logically grouped into a Virtual Chassis and uplink directly to the EX9200. In this collapsed model, the EX9200 is providing Layer 2 and Layer 3 services to the downstream network. With this scenario, MC-LAG is used between the core switches to provide a resilient, high bandwidth path to the downstream access layer. With the EX9200 providing routing at the campus core, MC-LAG is configured to support multiple VLANs with associated IRB interfaces, presented to the access network as a standard LAG group.

This configuration gives operators the benefits of increased bandwidth and link efficiency between the campus core and access layers, link resiliency between layers, along with the survivability provided by independent control and management planes.

Related Documentation

- [MC-LAG Technical Overview on page 7](#)
- [Example: Configuring Multichassis Link Aggregation on EX9200 Switches in the Core for Campus Networks on page 28](#)

Use Case for Simplifying MC-LAG Configuration

On the EX9200 switch, multichassis link aggregation (MC-LAG) enables a device to form a logical LAG interface across two physical chassis. Multichassis link aggregation groups provide node-level redundancy, multihoming support, and loop-free Layer 2 network without running the Spanning Tree Protocol (STP).

On the EX9200 switch, MC-LAG provides design flexibility and reliability with independent control and management planes. For MC-LAG to operate correctly, several configuration items should be configured in an identical manner on the MC-LAG peers. Because of the amount of configuration required, it is possible to make configuration mistakes or forget to configure required MC-LAG parameters on the peers. To simplify MC-LAG configuration,

the configuration consistency check and configuration synchronization features were introduced in Junos OS Release 16.1R1 for the EX Series.

Configuration consistency check verifies the MC-LAG configuration on each peer, flags any misconfigurations during the commit process, and prevents the MC-LAG interface from getting into an undesirable state because of inconsistent configuration between the MC-LAG peers. If there is an inconsistency, the corresponding MC-LAG interface is brought down, along with the reason why the consistency check failed. When you correct the configuration and issue another commit to fix the problem, the MC-LAG interface is brought back up.

Configuration synchronization reduces the chances of configuration inconsistencies by providing a single point of configuration for the MC-LAG peers. This feature uses configuration groups, so any configuration that is changed inside a configuration group is synchronized across MC-LAG peers that are defined as part of the group.

Related Documentation


- [MC-LAG Technical Overview on page 7](#)
- [Example: Simplifying Multichassis Link Aggregation on EX9200 Switches in the Core for Campus Networks on page 53](#)

MC-LAG Technical Overview

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 the Spanning Tree Protocol (STP).

[Figure 1 on page 7](#) illustrates the basic MC-LAG topology. On one end of the MC-LAG, there are two MC-LAG peers. Each of the MC-LAG peers has one or more physical links connected to the client device, such as a server or access switch. The client device, which is at the other end of the MC-LAG link, does not need to have an MC-LAG configured and does not need to be aware of MC-LAG. From its perspective, it is connecting to a single device through a LAG. 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.

Figure 1: Basic MC-LAG Topology



This topic provides an overview of MC-LAG and discusses the following:

- [ICCP and ICL on page 8](#)
- [Active/Standby and Active/Active Modes on page 9](#)
- [MC-LAG Interface on page 10](#)
- [Additional MC-LAG Specific Configuration on page 12](#)
- [Data Traffic Forwarding Rules in Active/Active MC-LAG Topologies on page 12](#)
- [Failure Handling During a Split-Brain State on page 13](#)
- [Layer 2 Feature Support on page 14](#)
- [Layer 2 Multicast Feature Support on page 16](#)
- [Layer 3 Feature Support on page 17](#)
- [Layer 3 Multicast Feature Support on page 21](#)
- [MC-LAG Upgrade Guidelines on page 22](#)
- [Summary of MC-LAG Configuration Guidelines on page 23](#)
- [Understanding Multichassis Link Aggregation Group \(MC-LAG\) Configuration Synchronization on page 23](#)
- [Understanding Multichassis Link Aggregation Group \(MC-LAG\) Configuration Consistency Check on page 27](#)

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. It uses TCP as a transport protocol and requires Bidirectional Forwarding Detection (BFD) for fast convergence. 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.

When configuring ICCP and the ICL, we recommend that you:

- 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.
- Use separate ports and choose different FPCs for the ICL and ICCP interfaces. Although you can use a single link for the ICCP interface, an aggregated Ethernet interface is preferred.
- Configure the ICCP liveness-detection interval (the 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 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 have 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 so results in faster ICCP connection establishment. The recommended value is 50 seconds.
- Configure a hold-down timer on the ICL member links that is greater than the configured BFD timer for the ICCP interface. This prevents the ICL from being advertised as being down before the ICCP link is down. If the ICL goes down before the ICCP link, this causes a flap of the MC-LAG interface on the status-control standby node, which leads to a delay in convergence.

Active/Standby and Active/Active Modes

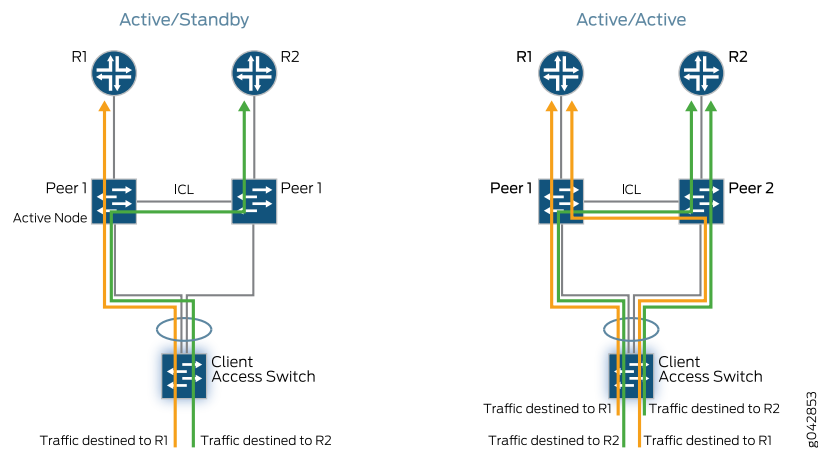
MC-LAG can be configured in active/standby mode, in which only one device actively forwards traffic, or in active/active mode, in which both devices actively forward traffic.

In active/standby mode, only one of the MC-LAG peers is active at any given time. The other MC-LAG peer is in backup (standby) mode. The active MC-LAG peer uses the Link Aggregation Control Protocol (LACP) to advertise to client devices that its child link is available for forwarding data traffic.

In active/active mode, all member links are active on the MC-LAG. In this mode, media access control (MAC) addresses learned on one MC-LAG peer are propagated to the other MC-LAG peer.

[Figure 2 on page 10](#) illustrates the difference between active/standby and active/active.

Figure 2: MC-LAG Active/Standby Versus Active/Active



This network configuration example uses active/active as the preferred mode for the following reasons:

- Traffic is load-balanced in active/active mode, resulting in a link-level efficiency of 100 percent.
- Convergence is faster in active/active mode than in active/standby mode. In active/active mode, information is exchanged between devices during operations. After a failure, the operational switch or router does not need to relearn any routes and continues to forward traffic.
- Active/active mode enables you to configure Layer 3 protocols on integrated routing and bridging (IRB) interfaces, providing a hybrid Layer 2 and Layer 3 environment on the core switch.

MC-LAG Interface

You configure an MC-LAG interface under the same configuration hierarchy as a LAG interface. You must configure the following:

- LACP—Configure LACP on the LAG. 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 all member links for an MC-LAG to work correctly.
- LACP system ID—Configure the same LACP system ID for the MC-LAG on each MC-LAG peer.
- MC-LAG specific options—MC-LAG specific options are configured under the **mc-ae** option. [Table 1 on page 11](#) describes the **mc-ae** options.

Table 1: mc-ae Statement Options

mc-ae Option	Description
mc-ae-id	Specifies which MC-LAG group the aggregated Ethernet interface belongs to.
redundancy-group	<p>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.</p> <p>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, this recommendation applies to each logical system—that is, configure one redundancy group between MC-LAG nodes in each logical system.</p>
init-delay-time	<p>Specifies the number of seconds by which to delay bringing the MC-LAG interface back to the up state when the MC-LAG peer is rebooted. By delaying the bring-up of the interface until after protocol convergence, you can prevent packet loss during the recovery of failed links and devices.</p> <p>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.</p>
chassis-id	Used by LACP for calculating the port number of the MC-LAG physical member links. Each MC-LAG peer should have a unique chassis ID.
mode	Indicates whether an MC-LAG is in active/standby mode or active/active mode. Chassis that are in the same group must be in the same mode. In this configuration example, the mode is active/active.
status-control	Specifies whether this node becomes active or goes into standby mode when an ICL failure occurs. Must be active on one node and standby on the other node.
events iccp-peer-down force-icl-down	Forces the ICL down if the peer of this node goes down.
events iccp-peer-down prefer-status-control-active	Allows the LACP system ID to be retained during a reboot, which provides better convergence after a failover. Note that if you configure both nodes as prefer-status-control-active , as this configuration example shows, you must also configure ICCP peering using the peer's loopback address to make sure that the ICCP session does not go down due to physical link failure.

Additional MC-LAG Specific Configuration

In addition to configuring ICCP, the ICL, and the MC-LAG interfaces, you must configure the following:

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

You can configure multichassis link protection under the **multi-chassis** hierarchy or under the logical interface configuration for each 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, as they are in this example. 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.


Data Traffic Forwarding Rules in Active/Active MC-LAG Topologies

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

- Single-homed link terminating on an MC-LAG peer device
- MC-LAG links
- ICL

These links are shown in [Figure 3 on page 12](#), which is used to illustrate the traffic forwarding rules that apply to MC-LAG active/active.

Figure 3: MC-LAG Traffic Forwarding Rules



The traffic forwarding rules are:

- Traffic received on MC-LAG peer N1 from the MC-LAG interface could be flooded to the ICL link to reach N2. When it reaches N2, it is not flooded back to the MC-LAG interface.
- Traffic received on SH1 could be flooded to the MC-LAG interface and the ICL by way of N1. When N2 receives SH1 traffic across the ICL link, it is not flooded to the MC-LAG interface.
- When receiving a packet from the ICL link, the MC-LAG peers forward the traffic to all local SH links. If the corresponding MC-LAG link on the peer is down, the receiving peer also forwards the traffic to its MC-LAG links.



NOTE: ICCP is used to signal MC-LAG link state between the peers.

- When N2 receives traffic from the ICL link, the traffic is not forwarded to the N2 upstream link if the upstream link is an MC-LAG link and the corresponding MC-LAG link on N1 is up.

Failure Handling During a Split-Brain State

Configuring ICCP adjacency over aggregated links 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 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 the 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 node 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 system 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 nodes.

To avoid this problem of the split-brain state and resultant convergence delays, configure *one* of the following two options:

- Enable backup liveness detection on the management (fxp0) interface. This is the preferred option.

For example:

```
[edit ]
user@switch# set protocols iccp peer 3.3.3.1 backup-liveness-detection backup-peer-ip
10.207.64.233
```

When you configure **backup-liveness-detection**, an out-of-band channel is established between the nodes, through the management network, to test the liveness of the Routing Engine. When both ICCP and backup liveness detection fail, the remote node is considered down, so the LACP system ID is not changed on the local node.

You must also configure the **master-only** statement on the IP address of the fxp0 interface for backup liveness detection, on both the master and backup Routing Engines, to ensure that the connection is not reset during GRES in the remote peer.

- Configure **prefer-status-control-active** under the **mc-ae** options for the MC-LAG on both nodes.

For example:

```
[edit ]
user@switch# set interfaces ae1 aggregated-ether-options mc-ae chassis-id 1 events
iccp-peer-down prefer-status-control-active
```

When you configure **prefer-status-control-active**, if ICCP goes down and backup liveness detection is up, the LACP system ID is not changed. Thus, if ICCP alone fails, the LACP system ID is not changed on the active node but it is changed on the standby node.

Layer 2 Feature Support

Support for the following Layer 2 features are discussed in this section:

- [MAC Address Management on page 14](#)
- [MAC Aging on page 15](#)
- [Spanning Tree Protocol on page 15](#)

MAC Address Management

Without proper MAC address management, an MC-LAG configuration could result in unnecessary flooding. For example:

- When an MC-LAG is configured to be active/active, upstream and downstream traffic could go through different MC-LAG peer devices. This means that the MAC address learned on one peer would have to be relearned on the other peer, causing unnecessary flooding.
- 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. MAC address replication is performed as follows:

- MAC addresses learned on an MC-LAG of one MC-LAG peer are replicated as learned on the same MC-LAG of the other MC-LAG peer.
- MAC addresses learned on single-homed clients of one MC-LAG peer are replicated as learned on the ICL interface of the other MC-LAG peer.
- MAC address learning from the data path is disabled on the ICL. MAC address learning on the ICL depends on software installing MAC addresses replicated through ICCP.

MAC Aging

MAC aging support in the Juniper Networks Junos[®] operating system (Junos OS) extends aggregated Ethernet logic for a specified MC-LAG. Aging of MAC addresses occurs when the MAC address is not seen on both of the MC-LAG peers. A MAC address in software is not deleted until all Packet Forwarding Engines have deleted the MAC address.

Spanning Tree Protocol

STP can be used to prevent loops in MC-LAG topologies. A potential loop, such as one that can happen due to miscabling at the core or access switching layer or due to a bug in server software, is broken by STP blocking one of the interfaces in the downstream network.

If your network topology requires RSTP or VSTP to prevent loops, configure the two MC-LAG nodes with same Spanning Tree Protocol (STP) virtual root ID using the Reverse Layer 2 Gateway Protocol (RL2GP). This root ID should be superior to all bridges in the downstream network while downstream bridges have to be capable of running STP. Because both the MC-LAG nodes are root bridges (virtual), the MC-LAG interface remains in the forwarding state. A downstream bridge receives bridge protocol data units (BPDUs) from both the nodes and thus receives twice the number of BPDUs on its aggregated Ethernet interface. If both MC-LAG nodes use the same aggregated Ethernet interface name, the STP port number will be identical, which reduces the STP load on the downstream bridge.

This network configuration example provides an example of configuring RSTP with RL2GP.



NOTE: STP is not supported on the ICL. If you enable STP globally, disable it on the ICL. This also means RSTP and VSTP cannot be configured on the ICL or ICL-PL.



NOTE: When configuring RSTP or VSTP in Junos, the MC-AE nodes must have the same system identifier configured as well as the highest bridge priority in the topology.

Layer 2 Multicast Feature Support

Layer 2 unknown multicast and IGMP snooping are supported. Key elements of this support are as follows:

- 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 as a multicast router port.
- IGMP membership learned on MC-LAG links is propagated across peers.
- During an MC-LAG peer reboot, known multicast traffic is flooded until the IGMP snooping state is synced with the peer.

IGMP Snooping on an Active/Active MC-LAG

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

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

When configuring IGMP snooping, keep these guidelines in mind:

- You must configure the ICL interface as a multicast router 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, you must enable PIM and IGMP on the IRB interface configured on the MC-LAG peers.

- You must configure the **multichassis-lag-replicate-state** statement under the **multicast-snooping-options** hierarchy for Internet Group Management Protocol (IGMP) snooping to work properly in an MC-LAG environment.

Layer 3 Feature Support

To provide Layer 3 routing functions to downstream clients, the MC-LAG network peers must be configured to provide the same gateway address to the downstream clients. To the upstream routers, the MC-LAG network peers could be viewed as either equal-cost multipath (ECMP) or two routes with different preference values. The following two methods can be used to enable Layer 3 functionality across an MC-LAG:

- VRRP over IRB—Configure different IP addresses on IRB interfaces on the MC-LAG peers and run the Virtual Router Redundancy Protocol (VRRP) over the IRB interfaces. The virtual IP address is the gateway IP address for the MC-LAG clients.
- MAC address synchronization—Configure the same IP address on the IRB interfaces on the MC-LAG peers, and configure the MAC address synchronization feature using the **mcae-mac-synchronize** statement. The IP address will be the gateway IP address for the MC-LAG clients.

We recommend that you use the VRRP over IRB method. Use MAC address synchronization only when you cannot configure VRRP over IRB. This network configuration example uses VRRP over IRB.

The following Layer 3 features are supported:


- [VRRP over IRB on page 17](#)
- [MAC Address Synchronization on page 18](#)
- [Address Resolution Protocol Synchronization for Active/Active MC-LAG Support on page 19](#)
- [DHCP Relay with Option 82 on page 20](#)

VRRP over IRB

Junos OS supports active/active MC-LAGs by using VRRP in active/standby mode. VRRP in active/standby mode enables Layer 3 routing over the multichassis aggregated Ethernet (MC-AE) interfaces on the MC-LAG peers. In this mode, the MC-LAG peers act as virtual routers. The peers share the virtual IP address that corresponds to the default route configured on the host or server connected to the MC-LAG. This virtual IP address (of the IRB interface) maps to either of the VRRP MAC addresses or to the logical interfaces of the MC-LAG peers. The host or server uses the VRRP MAC address to send any Layer 3 upstream packets.

At any time, one of the VRRP devices is the master (active), and the other is a backup (standby). Usually, a VRRP backup node does not forward incoming packets. However, when VRRP over IRB is configured in an MC-LAG active/active environment, both the VRRP master and the VRRP backup forward Layer 3 traffic arriving on the MC-AE interface, as shown in [Figure 4 on page 18](#). If the master fails, all the traffic shifts to the MC-AE link on the backup.

Figure 4: VRRP Forwarding in MC-LAG Configuration



NOTE: You must configure VRRP on both MC-LAG peers for both the active and standby members to accept and route packets.

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



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

For example, the following configures static ARP entries for IRB.21, where ae0.21 is the ICL interface:

```
set interfaces irb unit 21 family inet address 192.168.10.2/24 arp
192.168.10.3 12-interface ae0.21
```

MAC Address Synchronization

MAC address synchronization enables an MC-LAG peer to forward Layer 3 packets arriving on MC-AE interfaces with either its own IRB MAC address or its peer's IRB MAC address. Each MC-LAG peer installs its own IRB MAC address as well as the peer's IRB 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 MAC address is installed on the MC-LAG peer as if it was learned on the ICL.



NOTE: Use MAC address synchronization only if you are not planning to run routing protocols on the IRB interfaces. MAC address synchronization does not support routing protocols on the IRB interfaces. If you need routing capability, configure both VRRP and routing protocols on each MC-LAG peer.

Control packets destined for a particular MC-LAG peer that arrive on an MC-AE interface of its MC-LAG peer are not forwarded on the ICL interface. Additionally, using the gateway IP address as a source address when you issue either a ping, traceroute, telnet, or FTP request is not supported.



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

To enable the MAC address synchronization feature, issue the **set vlan *vlan-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.

Additional guidelines for implementing MAC address synchronization include:

- 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.
- Using Bidirectional Forwarding Detection (BFD) and MAC address synchronization together is not supported because ARP fails.

Address Resolution Protocol Synchronization for Active/Active MC-LAG Support

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 MC-AE interface.



NOTE: In some cases, ARP messages received by one MC-LAG peer are replicated to the other MC-LAG peer through ICCP. This optimization feature is applicable only for ARP replies, not ARP requests, received by the MC-LAG peers.



NOTE: 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.



NOTE: During graceful Routing Engine switchover (GRES), ARP entries that were learned remotely will be purged and then learned again.



NOTE: 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-l2-validate` statement on IRB interfaces in an MC-LAG configuration, as follows:

```
user@host# set interfaces irb arp-l2-validate
```

This option turns on validation of ARP and MAC table entries, automatically applying updates if they become out of sync.

DHCP Relay with Option 82



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

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

- Use the interface description instead of the interface name.
- Do not use the hostname as part of the circuit ID or remote ID strings.
- 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 3 Multicast Feature Support

The Protocol Independent Multicast (PIM) protocol and the Internet Group Management Protocol (IGMP) provide support for Layer 3 multicast.

PIM Operation

In standard mode of designated router election, 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 (RP) 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 MC-AE 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.

Layer 3 Multicast Configuration Guidelines

When you configure Layer 3 multicast, keep in mind the following guidelines:

- Enable PIM on the IRB interfaces on both MC-LAG nodes.
- 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.
- On the MC-LAG peer that has **status-control-active** configured, configure a high IP address or a high DR priority.

MC-LAG Upgrade Guidelines

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



NOTE: After a reboot, the MC-LAG 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.

1. Make sure that both of the MC-LAG peers (node1 and node2) are in the active/active state 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    : 20.1.1.2 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.

You can also use unified in-service software upgrade (ISSU) to upgrade the MC-LAG peers. On a device with dual Routing Engines, such as an EX9200, unified ISSU enables you to upgrade between two different Junos OS releases with no disruption on the control plane and with minimal disruption of traffic.

The guidelines for upgrading an MC-LAG using unified ISSU are similar to those for a regular upgrade:

- You must upgrade each MC-LAG peer independently. A unified ISSU performed on one peer does not trigger an upgrade of the other peer.

- We recommend that you upgrade each peer sequentially. Wait until one peer is fully upgraded before initiating a unified ISSU on the other peer.

In addition, graceful Routing Engine switchover (GRES) and nonstop active routing (NSR) must be enabled on each peer.

Summary of MC-LAG Configuration Guidelines

Table 2 on page 23 summarizes key configuration guidelines for an active/active MC-LAG configuration.

Table 2: Summary of Configuration Guidelines

Element	Configuration Guidelines
ICCP and ICL	<ul style="list-style-type: none"> • Use the peer loopback IP address for peering to avoid any direct link failure between MC-LAG peers. As long as the logical connection between the peers remains up, ICCP stays up. • Use separate ports and choose different FPCs for the ICL and ICCP interfaces. Although you can use a single link for the ICCP interface, an aggregated Ethernet interface is preferred. • Configure the ICCP liveness-detection interval (the BFD timer) to be at least 8 seconds, if you have configured ICCP connectivity through an IRB interface. • Configure a session establishment hold time for ICCP. Doing so results in faster ICCP connection establishment. The recommended value is 50 seconds. • Configure a hold-down timer on the ICL member links that is greater than the configured BFD timer for the ICCP interface. Doing so can minimize convergence delay.
MC-LAG Interface	<ul style="list-style-type: none"> • Configure LACP on all member links. • Configure only one redundancy group between MC-LAG nodes. • Configure either backup liveness detection or prefer-status-control-active on both MC-LAG peers to avoid LACP system ID flap during a reboot. Backup liveness detection is the preferred method of avoiding the LACP system ID flap. Use the prefer-status-control-active method only when you can ensure that ICCP goes down only when the node goes down.
Layer 2 Multicast	<ul style="list-style-type: none"> • Configure the ICL as a multicast router interface. • Configure the multichassis-lag-replicate-state statement under the multicast-snooping-options hierarchy.
Layer 3	<ul style="list-style-type: none"> • Use VRRP over IRB or MAC address synchronization to enable Layer 3 routing. We recommend using VRRP over IRB. If you use MAC address synchronization, routing protocols on IRBs are not supported. • Configure static ARP on the IRB peers to enable IRB-to-IRB connectivity across the ICL. • We recommend enabling the arp-l2-validate statement on IRBs as follows: <pre>user@host# set interfaces irb arp-l2-validate</pre>
Layer 3 Multicast	<ul style="list-style-type: none"> • Enable PIM on the IRB interfaces on both MC-LAG nodes. • Configure the ICL as a multicast router interface. • On the MC-LAG peer that has status-control-active configured, configure the IP address with a high IP address or a high DR priority.

Understanding Multichassis Link Aggregation Group (MC-LAG) Configuration Synchronization

Starting with Junos OS Release 16.1R1, 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. Using configuration groups to simplify the configuration process, create one configuration group for the local MC-LAG peer, one for the remote MC-LAG peer, and one for the global configuration, which is essentially a configuration that is common to both MC-LAG peers.

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

To enable MC-LAG configuration synchronization, perform the following steps:

1. Create configuration groups for local, remote, and global configurations.
2. Create conditional groups.
3. Create apply groups.
4. Enable NETCONF over SSH.
5. Configure MC-LAG peer details and user authentication details for MC-LAG configuration synchronization.
6. Enable the **peers-synchronize** statement or issue the **commit peers-synchronize** command to synchronize and commit the configurations between local and remote MC-LAG peers.

Understanding Configuration Groups

You can create configuration groups for local, remote, and global configurations. A local configuration group is used by the local MC-LAG peer, a remote configuration group is used by the remote MC-LAG peer, and a global configuration group is shared between the local and remote MC-LAG peers.

For example, you could create a local configuration group called Group A, which would include the configuration used by the local MC-LAG peer (Switch A), a remote configuration group called Group B, which would include the configuration used by the remote MC-LAG peer (Switch B), and a global configuration group called Group C, which would include the configuration that is common to both MC-LAG peers.

Create configuration groups at the **[edit groups]** hierarchy level.



NOTE: MC-LAG configuration synchronization does not support nested groups.

Understanding Conditional Groups

You can create conditional groups to specify when a particular configuration group is applied. To do this, issue the **set groups name-of-group when peers [static-hostname-of-local-peer static-hostname-of-remote-peer]** command. The **when** statement defines the conditions under which the configuration group is applied. The

peers statement enables you to specify the conditions, which in this case, are the static hostnames of the MC-LAG peers. For example, to specify that peers Switch A and Switch B will apply the configuration group called Group C, issue the **set groups GroupC when peers [SwitchA SwitchB]** command.

Understanding Apply Groups

To apply the configuration groups, enable the **apply-groups** statement at the **[edit]** hierarchy level. For example, to apply the local configuration group (Group A, for example), remote configuration group (Group B), and global configuration group (Group C), issue the **set apply-groups [GroupA GroupB GroupC]** command.

Understanding Peer Configuration Details for MC-LAG Configuration Synchronization

To synchronize configurations between two MC-LAG peers, you need to configure the hostname or IP address, username, and password for both the local and remote MC-LAG peers. 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 each MC-LAG peer. 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. To synchronize a configuration from Switch B to Switch A, issue the **set peers SwitchA user administrator authentication test123** command on Switch B. If you only want to synchronize configurations from Switch A to Switch B, you do not need to configure the **peers** statement on Switch B.

The configuration details from the peers statements are also used to establish a NETCONF over SSH connection between the MC-LAG peers. To enable NETCONF over SSH, issue the **set system services netconf ssh** command on both MC-LAG peers.

Understanding How Configurations Are Synchronized Between MC-LAG Peers

The local (or requesting) MC-LAG peer 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) MC-LAG peer. Each MC-LAG peer then performs a syntax check on the configuration file being committed. If no errors are found, the configuration is activated and becomes the current operational configuration on both MC-LAG peers. The commits are propagated using a remote procedural call (RPC).

The following events occur during configuration synchronization:

1. The local MC-LAG peer sends the sync-peers.conf file (the configuration that will be shared with the peer specified in the conditional group) to the remote MC-LAG peer.
2. The remote MC-LAG peer loads the configuration, sends the results of the load to the local MC-LAG peer, exports its configuration to the local MC-LAG peer, and replies that the commit is complete.
3. The local MC-LAG peer reads the reply from the remote MC-LAG peer.
4. If successful, the configuration is committed.

Configuration synchronization is not successful if either a) the remote MC-LAG peer is unavailable or b) the remote MC-LAG peer 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 MC-LAG peers 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 MC-LAG peer to another. For example, if you configured the **peers** statement on the local MC-LAG peer, and want to synchronize the configuration with the remote MC-LAG peer, you can simply issue the **commit** command on the local MC-LAG peer. However, if you issue the **commit** command on the local MC-LAG peer and the remote MC-LAG peer is not reachable, you will receive a warning message saying that the remote MC-LAG peer is not reachable and only the configuration on the local MC-LAG peer 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: 'mc-lag-peer' failed
warning: Cannot connect to remote peers, ignoring it
commit complete
```

If you do not have the **peers** statement configured with the remote MC-LAG peer information and you issue the **commit** command, only the configuration on the local MC-LAG peer is committed. If the remote MC-LAG peer is unreachable and there are other failures, the commit is unsuccessful on both the local and remote MC-LAG peers.



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 MC-LAG peers 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 MC-LAG peers configured in the **peers** statement. If you issue the **commit peers-synchronize** command on the local MC-LAG peer to synchronize the configuration with the remote MC-LAG peer and the remote MC-LAG peer is reachable but there are other failures, the commit fails on both the local and remote MC-LAG peers.

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

Starting with Junos OS Release 16.1R1, the configuration consistency check feature was introduced. Configuration consistency check uses the Inter-Chassis Control Protocol (ICCP) to exchange MC-LAG configuration parameters (chassis ID, service ID, and so on) and checks for any configuration inconsistencies across MC-LAG peers. An example of an inconsistency is configuring identical chassis IDs on both peers instead of configuring unique chassis IDs on both peers. When there is an inconsistency, you are notified and can take action to resolve it. Configuration consistency check is invoked after you issue a commit on an MC-LAG peer.

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 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 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.

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.



NOTE: For information on the MC-LAG parameters that are checked for consistency, as well as the commands you can issue to verify configuration consistency check, see [Understanding Multichassis Link Aggregation Group Configuration Consistency Check](#).

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 MC-LAG 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 explicitly 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 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 becomes reachable, 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.

Consistency check is not enabled by default. To enable consistency check, issue the `set multi-chassis mc-lag consistency-check` command.

Related Documentation

- [Example: Configuring Multichassis Link Aggregation on EX9200 Switches in the Core for Campus Networks on page 28](#)

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

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 on page 29](#)
- [Overview on page 29](#)
- [Configuration on page 30](#)
- [Verification on page 44](#)

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 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 5 on page 30](#) shows the topology of this example.

Figure 5: Topology Diagram

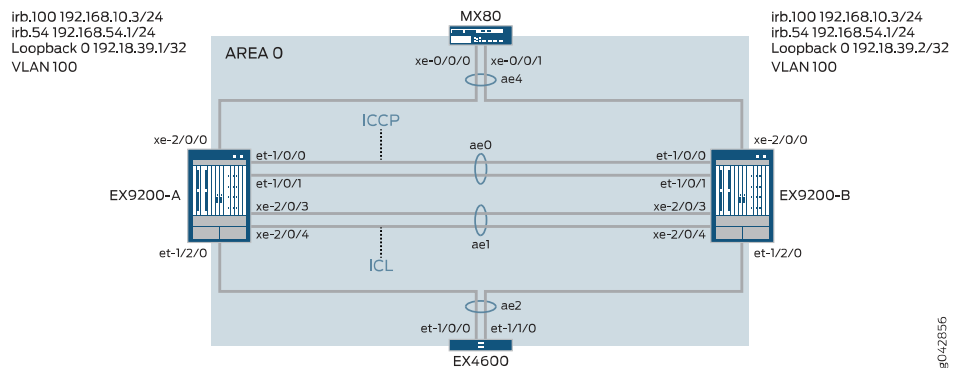


Table 3 on page 30 details the topology used in this configuration example.

Table 3: Components of the Topology for Configuring a Multichassis LAG Between Two Switches

Hostname	Base Hardware	Multichassis Link Aggregation Group
EX9200-A	EX9200	<p>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-1/0/1 on EX9200-B.</p> <p>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.</p> <p>ae2 is configured as an MC-LAG, and the following interfaces are part of ae2: et-1/2/0 on EX9200-A and et-1/2/0 on EX9200-B.</p> <p>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.</p>
EX9200-B	EX9200	

Configuration

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

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 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 et-1/2/0 ether-options 802.3ad ae2
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 irb unit 100 family inet address 192.168.10.3/24 arp 192.168.10.2 l2-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 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 lo0 unit 0 family inet address 192.18.39.1/32
set protocols iccp local-ip-addr 192.18.39.1
set protocols iccp peer 192.18.39.2 session-establishment-hold-time 50
set protocols iccp peer 192.18.39.2 redundancy-group-id-list 1
set protocols iccp peer 192.18.39.2 backup-liveness-detection backup-peer-ip 10.105.5.6
set protocols iccp peer 192.18.39.2 liveness-detection minimum-interval 2000
set protocols iccp peer 192.18.39.2 liveness-detection multiplier 4
set multi-chassis multi-chassis-protection 192.18.39.2 interface ae1
set switch-options service-id 1
set vlans rack_1 vlan-id 100
set vlans rack_1 l3-interface irb.100
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
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
set interfaces xe-2/0/0 ether-options 802.3ad ae4
```

```
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 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 v51
set interfaces ae4 unit 0 family ethernet-switching vlan members v54
set vlans v54 vlan-id 54
set vlans v54 l3-interface irb.54
set interfaces irb unit 54 family inet address 192.168.54.2/24 arp 192.168.54.1 l2-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 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
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
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
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 192.18.39.1
```

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 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 et-1/2/0 ether-options 802.3ad ae2
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 iccp-peer-down
prefer-status-control-active
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 irb unit 100 family inet address 192.168.10.2/24 arp 192.168.10.3 l2-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 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 lo0 unit 0 family inet address 192.18.39.2/32
set protocols iccp local-ip-addr 192.18.39.2
set protocols iccp peer 192.18.39.1 session-establishment-hold-time 50
set protocols iccp peer 192.18.39.1 redundancy-group-id-list 1
set protocols iccp peer 192.18.39.1 backup-liveness-detection backup-peer-ip 10.105.5.5
set protocols iccp peer 192.18.39.1 liveness-detection minimum-interval 2000
set protocols iccp peer 192.18.39.1 liveness-detection multiplier 4
set multi-chassis multi-chassis-protection 192.18.39.1 interface ae1
set switch-options service-id 1
set vlans rack_1 vlan-id 100
set vlans rack_1 l3-interface irb.100
set forwarding-options dhcp-relay forward-snooped-clients all-interfaces
set 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
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
set interfaces xe-2/0/0 ether-options 802.3ad ae4
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 v51
set interfaces ae4 unit 0 family ethernet-switching vlan members v54
set vlans v54 vlan-id 54
set vlans v54 l3-interface irb.54
set interfaces irb unit 54 family inet address 192.168.54.1/24 arp 192.168.54.2 l2-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 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
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
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
set protocols pim interface lo0.0
set protocols pim rp bootstrap-priority 200
set protocols pim rp local address 192.18.39.2

```

Configuring MC-LAG on 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* in the *CLI User Guide*.

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. 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

```

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

```

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

```

5. 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

```

6. 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
```
7. 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
```
8. 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
```
9. Configure the LACP administration key on ae2.

```
[edit interfaces]
user@switch# set ae2 aggregated-ether-options lACP system-id 00:01:02:03:04:06
```
10. 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
```
11. 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
```
12. 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
```
13. 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
```
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. 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
```
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 an integrated routing and bridging (IRB) interface on VLAN 100.

To configure an MC-LAG IRB, 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 192.168.10.3/24 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
```
18. 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
```
19. Configure a loopback interface.
- ```
[edit interfaces]
user@switch# set lo0 unit 0 family inet address 192.18.39.1/32
```
20. Configure ICCP using the loopback address.
- ```
[edit protocols]
user@switch# set iccp local-ip-addr 192.18.39.1
```
21. 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@switch# set iccp peer 192.18.39.2 session-establishment-hold-time 50
user@switch# set iccp peer 192.18.39.2 redundancy-group-id-list 1
user@switch# set iccp peer 192.18.39.2 backup-liveness-detection backup-peer-ip 10.105.5.6
```
22. 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 192.18.39.2 liveness-detection minimum-interval 2000
user@switch# set iccp peer 192.18.39.2 liveness-detection multiplier 4
[edit multi-chassis]
user@switch# set multi-chassis-protection 192.18.39.2 interface ae1
```
23. 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
```
24. Configure VLAN 100.
- ```
[edit vlans]
user@switch# set rack_1 vlan-id 100
user@switch# set rack_1 l3-interface irb.100
```
25. Configure forward snooped unicast packets on all interfaces.
- ```
[edit forwarding-options]
user@switch# set dhcp-relay forward-snooped-clients all-interfaces
```

26. Create a binding entry to snoop unicast clients.  

```
[edit forwarding-options]
user@switch# set dhcp-relay overrides allow-snooped-clients
```
27. Create a DHCP server group.  

```
[edit forwarding-options]
user@switch# set dhcp-relay server-group GVP-DHCP 10.105.5.202
```
28. 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
```
29. 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
```
30. 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
```
31. Configure DHCP relay with option 82.  

```
[edit forwarding-options]
user@switch# set dhcp-relay relay-option-82 circuit-id use-interface-description device
```
32. 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
```
33. Configure the system identifier.  

```
[edit protocols]
user@switch# set rstp system-identifier 00:01:02:03:04:05
```
34. 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
```
35. Configure multicast with multichassis link aggregation between EX9200-A and an MX Series router.  

Specify the members that belong to ae4.

```
[edit interfaces]
user@switch# set xe-2/0/0 ether-options 802.3ad ae4
```
36. 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
```
37. 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
```

38. 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
```

39. 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
```

40. 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 v51
user@switch# set ae4 unit 0 family ethernet-switching vlan members v54
```

41. Configure VLAN 54.

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

42. Configure an IRB interface on VLAN 100.

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 l2-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
```

43. 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
```

44. Enable IGMP snooping for all VLANs.

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

45. 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
```

46. 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

```
47. 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

```
 48. 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

```
 49. Configure the loopback interface.


```

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

```
 50. 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 192.18.39.1

```
 51. (Optional) To change the MC-AE mode, use this command.


```

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

```

Configuring MC-LAG on 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* in the *CLI User Guide*.

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


```

[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. 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
```
4. Specify the member interfaces that belong to interface ae2.  

```
[edit interfaces]
user@switch# set et-1/2/0 ether-options 802.3ad ae2
```
  5. 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
```
  6. 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
```
  7. 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
```
  8. Configure the LACP parameters on ae2.  

```
[edit interfaces]
user@switch# set interfaces ae2 aggregated-ether-options lacp active
user@switch# set interfaces ae2 aggregated-ether-options lacp periodic fast
```
  9. Configure the LACP administration key on ae2.  

```
[edit interfaces]
user@switch# set ae2 aggregated-ether-options lacp system-id 00:01:02:03:04:05
```
  10. 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
```
  11. 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
```
  12. 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
```
  13. Specify time in seconds by when routing adjacencies must form.  

```
[edit interfaces]
user@switch# set ae2 aggregated-ether-options mc-ae init-delay-time 520
```
  14. 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
```
  15. Configure the status control on the switch that hosts the MC-LAG.



If one switch is in standby mode, then the other switch must be in active 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 an integrated routing and bridging (IRB) interface on VLAN 100.

To configure an MC-LAG IRB, 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
l2-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
```

18. 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 200
user@switch# set irb unit 100 family inet address 192.168.10.2/24 vrrp-group 1 accept-data
```

19. Configure a loopback interface.

```
[edit interfaces]
user@switch# set lo0 unit 0 family inet address 192.18.39.2/32
```

20. Configure ICCP using the loopback address.

```
[edit protocols]
user@switch# set iccp local-ip-addr 192.18.39.2
```

21. 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@switch# set iccp peer 192.18.39.1 session-establishment-hold-time 50
user@switch# set iccp peer 192.18.39.1 redundancy-group-id-list 1
user@switch# set iccp peer 192.18.39.1 backup-liveness-detection backup-peer-ip 10.105.5.5
```

22. 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 192.18.39.1 liveness-detection minimum-interval 2000
user@switch# set iccp peer 192.18.39.1 liveness-detection multiplier 4
```

```
[edit multi-chassis]
user@switch# set multi-chassis-protection 192.18.39.1 interface ae1
```

23. 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
```

24. Configure VLAN 100.

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

25. Configure forward snooped unicast packets on all interfaces.

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

26. Create a binding entry to snoop unicast clients.

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

27. Create a DHCP server group.

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

28. 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
```

29. 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
```

30. 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
```

31. Configure DHCP relay with option 82.

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

32. 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
```

33. Configure the system identifier.

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

34. 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
```

35. Configure multicast with multichassis link aggregation between EX9200-B and an MX Series router.

Specify the members that belong to ae4.

- ```
[edit interfaces]
user@switch# set xe-2/0/0 ether-options 802.3ad ae4
```
36. 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
```
37. 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
```
38. Specify a unique chassis ID for the MC-LAG that the aggregated Ethernet interface belongs to on each switch.
- ```
[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
```
39. Configure the status control on the switch that hosts the MC-LAG.
- If one switch is in standby mode, then the other switch must be in active 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
```
40. 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 v51
user@switch# set ae4 unit 0 family ethernet-switching vlan members v54
```
41. Configure VLAN 54.
- ```
[edit vlans]
user@switch# set v54 vlan-id 54
user@switch# set v54 l3-interface irb.54
```
42. Configure an IRB interface on VLAN 100.
- 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 192.168.54.1/24 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 00:1f:12:b6:6f:f0
```
43. Enable VRRP on the MC-LAGs on each switch 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 switch 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
```
44. Enable IGMP snooping for all VLANs.
- ```
[edit protocols]
```

- ```
user@switch# set igmp-snooping vlan rack_1
user@switch# set igmp-snooping vlan v54
```
45. 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
```
46. 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
```
47. 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
```
48. 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
```
49. Configure the loopback interface.
- ```
[edit protocols]
user@switch# set pim interface lo0.0
```
50. 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 192.18.39.2
```
51. (Optional) To change the MC-AE mode, use this command.
- ```
[edit interfaces]
user@switch# set ae2 aggregated-ether-options mc-ae mode active-standby
```

## Verification

Confirm that the configuration is working properly.

- [Verifying ICCP on MC-LAG on page 45](#)
- [Verifying LACP on MC-LAG on page 45](#)
- [Verifying Aggregated Ethernet Interfaces in MC-LAG on page 47](#)
- [Verifying MAC Learning on MC-LAG on page 49](#)
- [Verifying VRRP in MC-LAG on page 51](#)
- [Verifying OSPF on MC-LAG on page 52](#)

### 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@EX9200-A> show iccp
Redundancy Group Information for peer 192.18.39.2
 TCP Connection : Established
 Liveliness Detection : Up
```

```
Backup liveness peer status: 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
```

```
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 192.18.39.1
 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: 12ald_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 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 No No Yes Yes Yes Yes Fast Active
et-1/0/1 Partner No No Yes Yes Yes Yes Fast Active

LACP protocol:
et-1/0/0 Receive State Transmit State Mux State
 Current Fast periodic Collecting distributing
et-1/0/1 Current Fast periodic Collecting distributing

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 No Yes Yes Yes Yes Fast Active
xe-2/0/4 Actor No No Yes Yes Yes Yes Fast Active
xe-2/0/4 Partner No No Yes Yes Yes Yes Fast Active

LACP protocol:
xe-2/0/3 Receive State Transmit State Mux State
 Current Fast periodic Collecting distributing
xe-2/0/4 Current Fast periodic Collecting distributing

Aggregated interface: ae3
LACP state: Role Exp Def Dist Col Syn Aggr Timeout Activity
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 No No Yes Yes Yes Yes Fast Active
xe-2/0/2 Partner No No Yes Yes Yes Yes Fast Passive

LACP protocol:
xe-2/0/1 Receive State Transmit State Mux State
 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 No Yes Yes Yes Yes Fast Active
xe-2/0/0 Partner No No Yes Yes Yes Yes Fast Active

LACP protocol:
xe-2/0/0 Receive State Transmit State Mux State
 Current Fast periodic Collecting distributing

```

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   | No                    | No  | Yes                                     | Yes | Yes | Yes  | Fast                   | Active   |
| et-1/0/1                  | Partner | No                    | No  | Yes                                     | Yes | Yes | Yes  | Fast                   | Active   |
| LACP protocol: et-1/0/0   |         | Receive State Current |     | Transmit State Fast periodic Collecting |     |     |      | Mux State distributing |          |
| et-1/0/1                  |         | Current               |     | Fast periodic Collecting                |     |     |      | distributing           |          |
| 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                    | No  | Yes                                     | Yes | Yes | Yes  | Fast                   | Active   |
| xe-2/0/4                  | Actor   | No                    | No  | Yes                                     | Yes | Yes | Yes  | Fast                   | Active   |
| xe-2/0/4                  | Partner | No                    | No  | Yes                                     | Yes | Yes | Yes  | Fast                   | Active   |
| LACP protocol: xe-2/0/3   |         | Receive State Current |     | Transmit State Fast periodic Collecting |     |     |      | Mux State distributing |          |
| xe-2/0/4                  |         | Current               |     | Fast periodic Collecting                |     |     |      | distributing           |          |
| Aggregated interface: ae2 |         |                       |     |                                         |     |     |      |                        |          |
| LACP state:               | Role    | Exp                   | Def | Dist                                    | Col | Syn | Aggr | Timeout                | Activity |
| et-1/2/0                  | Actor   | No                    | No  | Yes                                     | Yes | Yes | Yes  | Fast                   | Active   |
| et-1/2/0                  | Partner | No                    | No  | Yes                                     | Yes | Yes | Yes  | Fast                   | Passive  |
| LACP protocol: et-1/2/0   |         | Receive State Current |     | Transmit State Fast periodic Collecting |     |     |      | Mux State distributing |          |
| 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: xe-2/0/0   |         | Receive State Current |     | Transmit State Fast periodic Collecting |     |     |      | Mux State 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 : 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 State : up
 Peer Ip/MCP/State : 192.18.39.2 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 : 192.18.39.2 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 State : up
 Peer Ip/MCP/State : 192.18.39.1 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 : 192.18.39.1 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 name | MAC address       | MAC flags | Age | Logical interface |
|-----------|-------------------|-----------|-----|-------------------|
| dmzuplink | 00:00:5e:00:01:ba | DL        | -   | ae4.0             |
| dmzuplink | 00:10:db:bc:f5:9d | DR        | -   | ae4.0             |
| dmzuplink | 00:10:db:ff:10:01 | DL        | -   | ae3.0             |
| dmzuplink | 00:19:e2:57:33:81 | DR        | -   | ae4.0             |
| dmzuplink | 00:26:88:92:ef:1d | DR        | -   | ae4.0             |
| dmzuplink | 28:8a:1c:74:fb:07 | DR        | -   | ae4.0             |
| dmzuplink | 28:8a:1c:75:05:1f | DR        | -   | ae4.0             |
| dmzuplink | 28:c0:da:6a:1d:2a | DR        | -   | ae4.0             |
| dmzuplink | 2c:21:72:7d:40:01 | DL        | -   | ae4.0             |
| dmzuplink | 3c:8a:b0:77:a9:d6 | DR        | -   | ae4.0             |
| 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 name | MAC address       | MAC flags | Age | Logical interface |
|-----------|-------------------|-----------|-----|-------------------|
| rack_1    | 00:50:56:9b:01:57 | DR        | -   | ae2.0             |
| rack_1    | 00:50:56:9b:09:95 | DL        | -   | ae2.0             |
| rack_1    | 00:50:56:9b:15:2e | DL        | -   | ae2.0             |
| 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 |
| rack_1 | 00:50:56:9b:75:60 | DR | - | 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<br>name | MAC<br>address    | MAC<br>flags | Age | Logical<br>interface |
|--------------|-------------------|--------------|-----|----------------------|
| v54          | 80:71:1f:c1:85:f0 | 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<br>name | MAC<br>address    | MAC<br>flags | Age | Logical<br>interface |
|--------------|-------------------|--------------|-----|----------------------|
| 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 | DR           | -   | ae2.0                |
| rack_1       | 00:50:56:9b:20:44 | DR           | -   | ae2.0                |
| rack_1       | 00:50:56:9b:20:a7 | DR           | -   | ae2.0                |
| rack_1       | 00:50:56:9b:22:a8 | DL           | -   | ae2.0                |
| rack_1       | 00:50:56:9b:38:01 | DR           | -   | 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<br>name | MAC<br>address | MAC<br>flags | Age | Logical<br>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<br>name | MAC<br>address    | MAC<br>flags | Age | Logical<br>interface |
|--------------|-------------------|--------------|-----|----------------------|
| v54          | 80:71:1f:c1:85:f0 | DR           | -   | ae4.0                |

**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 | Group | VR state | VR Mode | Timer   | Type | Address |
|--------------|-------|-------|----------|---------|---------|------|---------|
| irb.54       | up    | 4     | backup   | Active  | D 3.090 | lcl  |         |
| 192.168.54.1 |       |       |          |         |         |      |         |
|              |       |       |          |         |         |      | vip     |
| 192.168.54.3 |       |       |          |         |         |      | mas     |
| 192.168.54.2 |       |       |          |         |         |      |         |
| irb.100      | up    | 1     | backup   | Active  | D 2.655 | lcl  |         |
| 192.168.10.3 |       |       |          |         |         |      |         |
|              |       |       |          |         |         |      | vip     |
| 192.168.10.1 |       |       |          |         |         |      | mas     |
| 192.168.10.2 |       |       |          |         |         |      |         |

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 vrrp

| Interface    | State | Group | VR state | VR Mode | Timer   | Type | Address |
|--------------|-------|-------|----------|---------|---------|------|---------|
| irb.54       | up    | 4     | master   | Active  | A 0.900 | lcl  |         |
| 192.168.54.2 |       |       |          |         |         |      |         |
|              |       |       |          |         |         |      | vip     |
| 192.168.54.3 |       |       |          |         |         |      |         |
| irb.100      | up    | 1     | master   | Active  | A 0.175 | lcl  |         |
| 192.168.10.2 |       |       |          |         |         |      |         |
|              |       |       |          |         |         |      | vip     |
| 192.168.10.1 |       |       |          |         |         |      |         |

In this example, Switch B is the master 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 neighbor
```

| Address      | Interface | State | ID          | Pri | Dead |
|--------------|-----------|-------|-------------|-----|------|
| 192.168.90.2 | ae0.0     | Full  | 192.18.39.2 | 128 | 35   |
| 192.168.10.2 | irb.100   | Full  | 192.18.39.2 | 128 | 33   |
| 192.168.54.2 | irb.54    | Full  | 192.18.39.2 | 128 | 38   |

2. Show OSPF routing table on Switch A.

```
root@EX9200-A> show ospf route
```

Topology default Route Table:

| Prefix          | Path  | Route   | NH   | Metric | NextHop   | Nexthop      |
|-----------------|-------|---------|------|--------|-----------|--------------|
|                 | Type  | Type    | Type |        | Interface | Address/LSP  |
| 192.18.39.2     | Intra | Router  | IP   | 1      | ae0.0     | 192.168.90.2 |
|                 |       |         |      |        | irb.100   | 192.168.10.2 |
|                 |       |         |      |        | irb.54    | 192.168.54.2 |
| 192.18.39.1/32  | Intra | Network | IP   | 0      | lo0.0     |              |
| 192.18.39.2/32  | 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 | IP   | 1      | ae0.0     |              |

3. Show OSPF neighbors on Switch B.

```
root@EX9200-B> show ospf neighbor
```

| Address      | Interface | State | ID          | Pri | Dead |
|--------------|-----------|-------|-------------|-----|------|
| 192.168.90.1 | ae0.0     | Full  | 192.18.39.1 | 128 | 32   |
| 192.168.10.3 | irb.100   | Full  | 192.18.39.1 | 128 | 34   |
| 192.168.54.1 | irb.54    | Full  | 192.18.39.1 | 128 | 37   |

4. Show OSPF routing table on Switch B.

```
root@EX9200-B> show ospf route
```

Topology default Route Table:

| Prefix      | Path  | Route  | NH   | Metric | NextHop   | Nexthop      |
|-------------|-------|--------|------|--------|-----------|--------------|
|             | Type  | Type   | Type |        | Interface | Address/LSP  |
| 192.18.39.1 | Intra | Router | IP   | 1      | ae0.0     | 192.168.90.1 |
|             |       |        |      |        | irb.100   | 192.168.10.3 |
|             |       |        |      |        | irb.54    | 192.168.54.1 |

|                 |               |    |           |              |
|-----------------|---------------|----|-----------|--------------|
| 192.18.39.1/32  | Intra Network | IP | 1 ae0.0   | 192.168.90.1 |
|                 |               |    | irb.100   | 192.168.10.3 |
|                 |               |    | irb.54    | 192.168.54.1 |
| 192.18.39.2/32  | Intra Network | IP | 0 lo0.0   |              |
| 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 | IP | 1 ae0.0   |              |

**Related Documentation**

- [Configuring Multichassis Link Aggregation](#)

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

- [Requirements on page 53](#)
- [Overview on page 53](#)
- [Configuration on page 55](#)
- [Verification on page 74](#)

### 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 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 peer.

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 6 on page 55](#) shows the topology of this example.

Figure 6: Topology Diagram




Table 4 on page 55 details the topology used in this configuration example.

Table 4: Components of the Topology for Configuring a Multichassis LAG Between Two Switches

| Hostname        | Base Hardware                                                             | Multichassis Link Aggregation Group                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
|-----------------|---------------------------------------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| EX9200-A        | EX9200                                                                    | <p>ae0 is configured as an aggregated Ethernet interface, and is used as an ICCP link, and the following interfaces are part of ae0:<br/>xe-0/3/6 and xe-1/3/6.</p> <p>ae1 is configured as an aggregated Ethernet interface and is used as an ICL link, and the following interfaces are part of ae1:<br/>xe-0/3/7 and xe-1/3/7.</p> <p>ae2 is configured as an MC-LAG, and the following interfaces are part of ae2:<br/>xe-0/0/1 on Switch B and xe-1/0/1 on Switch A.</p> <p>ae3 is configured as an MC-LAG, and the following interface is part of ae3 on both Switch A and Switch B:<br/>xe-0/0/2.</p> |
| EX9200-B        | EX9200                                                                    |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| Virtual Chassis | Not applicable. Virtual Chassis are shown only for illustration purposes. |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| Virtual Chassis |                                                                           |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
|                 |                                                                           | 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

**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 l2-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 l2-interface ae1
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
192.168.100.1
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. For information about navigating the CLI, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

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
```

13. 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. For information about navigating the CLI, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

- 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"
```
- 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 l2-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;
        }
    }
}
lldp {
    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;
    l3-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;  
      }  
    }  
  }  
}  
user@EX9200-A# show groups MC_Config_Remote  
interfaces {  
  ae0 {  
    unit 0 {  
      family inet {  
        address 172.16.32.10/30;  
      }  
    }  
  }  
  ae2 {  
    aggregated-ether-options {  
      mc-ae {  
        chassis-id 1;  
        status-control standby;  
      }  
    }  
  }  
  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 l2-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

- [Verifying ICCP on MC-LAG on page 75](#)
- [Verifying LACP on MC-LAG on page 76](#)
- [Verifying Aggregated Ethernet Interfaces in MC-LAG on page 78](#)
- [Verifying VRRP in MC-LAG on page 79](#)
- [Verifying OSPF on MC-LAG on page 79](#)
- [Verifying that Configuration Consistency Check Passed on page 80](#)
- [Verifying the Configuration Consistency Check Status for the Global Configuration on page 84](#)
- [Verifying the Configuration Consistency Check Status for the Interchassis Control Link on page 85](#)

- [Verifying the Configuration Consistency Check Status for the MC-LAG Interfaces on page 86](#)
- [Verifying the Configuration Consistency Check Status for the VLAN Configuration on page 90](#)
- [Verifying the Configuration Consistency Check Status for VRRP on page 91](#)

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.

```
user@EX9200-A> show iccp
```

```
Redundancy Group Information for peer 172.16.32.6
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: 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: l2ald_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  Exp  Def  Dist  Col  Syn  Aggr  Timeout  Activity
xe-0/3/6        Actor  No   No   Yes   Yes  Yes   Yes    Fast    Active
xe-0/3/6        Partner No   No   Yes   Yes  Yes   Yes    Fast    Active
xe-1/3/6        Actor  No   No   Yes   Yes  Yes   Yes    Fast    Active
xe-1/3/6        Partner No   No   Yes   Yes  Yes   Yes    Fast    Active

LACP protocol:   Receive State  Transmit State      Mux State
xe-0/3/6         Current    Fast periodic Collecting distributing
xe-1/3/6         Current    Fast periodic Collecting distributing

Aggregated interface: ae1
LACP state:      Role  Exp  Def  Dist  Col  Syn  Aggr  Timeout  Activity
xe-0/3/7        Actor  No   No   Yes   Yes  Yes   Yes    Fast    Active
xe-0/3/7        Partner No   No   Yes   Yes  Yes   Yes    Fast    Active
xe-1/3/7        Actor  No   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
xe-0/3/7         Current    Fast periodic Collecting distributing
xe-1/3/7         Current    Fast periodic Collecting distributing

Aggregated interface: ae2
LACP state:      Role  Exp  Def  Dist  Col  Syn  Aggr  Timeout  Activity
xe-0/0/1        Actor  No   Yes   No   No   No   Yes    Fast    Active
xe-0/0/1        Partner No   Yes   No   No   No   Yes    Fast    Passive

LACP protocol:   Receive State  Transmit State      Mux State
xe-0/0/1         Current    Fast periodic Collecting
distributing
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        Partner No   Yes   No   No   No   Yes    Fast    Passive

```

```

LACP protocol:      Receive State  Transmit State      Mux State
xe-0/0/2            Current    Fast periodic Collecting distributing

```

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  Exp  Def  Dist  Col  Syn  Aggr  Timeout  Activity
xe-0/3/6        Actor  No   No   Yes   Yes   Yes   Yes    Fast    Active
xe-0/3/6        Partner No   No   Yes   Yes   Yes   Yes    Fast    Active
xe-1/3/6        Actor  No   No   Yes   Yes   Yes   Yes    Fast    Active
xe-1/3/6        Partner No   No   Yes   Yes   Yes   Yes    Fast    Active

LACP protocol:      Receive State  Transmit State      Mux State
xe-0/3/6            Current    Fast periodic Collecting distributing
xe-1/3/6            Current    Fast periodic Collecting distributing

Aggregated interface: ae1
LACP state:      Role  Exp  Def  Dist  Col  Syn  Aggr  Timeout  Activity
xe-0/3/7        Actor  No   No   Yes   Yes   Yes   Yes    Fast    Active
xe-0/3/7        Partner No   No   Yes   Yes   Yes   Yes    Fast    Active
xe-1/3/7        Actor  No   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
xe-0/3/7            Current    Fast periodic Collecting distributing
xe-1/3/7            Current    Fast periodic Collecting distributing

Aggregated interface: ae2
LACP state:      Role  Exp  Def  Dist  Col  Syn  Aggr  Timeout  Activity
xe-1/0/1        Actor  No   Yes   No   No   No   Yes    Fast    Active
xe-1/0/1        Partner No   Yes   No   No   No   Yes    Fast    Passive

LACP protocol:      Receive State  Transmit State      Mux State
xe-0/0/1            Current    Fast periodic Collecting distributing
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            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 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 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 1c1
192.168.100.2
                                     vip
192.168.100.1

```

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

2. Confirm that VRRP is up and active on EX9200-B.

```

user@EX9200-B> show vrrp
Interface      State      Group  VR state VR Mode  Timer  Type  Address
irb.100        up          1    backup  Active   D 2.887 1c1
192.168.100.3
                                     vip
192.168.100.1
                                     mas
192.168.100.2

```

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      NH      Metric NextHop      Nexthop
           Type Type      Type      Interface Address/LSP
172.16.32.6 Intra Router IP        1 ae0.0      172.16.32.10
172.16.32.5/32 Intra Network IP        0 lo0.0

```

```

172.16.32.6/32      Intra Network   IP          1 ae0.0      172.16.32.10
172.16.32.8/30      Intra Network   IP          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> show ospf route
Topology default Route Table:

Prefix          Path  Route      NH      Metric NextHop      Nexthop
                Type  Type       Type
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 lo0.0
172.16.32.8/30   Intra Network IP      1 ae0.0

```

Meaning 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.

```

user@EX9200-A> show multi-chassis mc-lag configuration-consistency
Configuration Item      Enforcement Level  Local Value
Peer Value              Result
-----
ICL interface          Mandatory          ae1
ae1                     PASS
rstp-bridge-priority   Desirable          0
0                       PASS
service-id             Mandatory          1
1                       PASS
session-establishment-hold-time Mandatory          300
300                    PASS
local-ip-addr           Mandatory          172.16.32.5
172.16.32.6            PASS
backup-liveness-detection 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

Local Physical Interface:ae2

```


Peer Physical Interface :ae2			
Configuration Item		Enforcement Level	Local Value
Peer Value	Result		
-----	-----	-----	-----
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		
prefer-status-control-active		Desirable	TRUE
--	PASS		
mcae status-control		Mandatory	standby
active	PASS		
mcae deployment mode		Mandatory	active-active
active-active	PASS		
mcae chassis-id		Mandatory	0
1	PASS		
mcae redundancy-group		Mandatory	1
1	PASS		
Local Logical Interface:ae2.0			
Peer Logical Interface :ae2.0			
Configuration Item		Enforcement Level	Local Value
Peer Value	Result		
-----	-----	-----	-----
vlan membership		Mandatory	100
100	PASS		
interface-mode		Mandatory	trunk
trunk	PASS		
Local Physical Interface:ae3			
Peer Physical Interface :ae3			
Configuration Item		Enforcement Level	Local Value
Peer Value	Result		
-----	-----	-----	-----
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		
prefer-status-control-active		Desirable	TRUE
--	PASS		
mcae status-control		Mandatory	standby
active	PASS		
mcae deployment mode		Mandatory	active-active
active-active	PASS		
mcae chassis-id		Mandatory	0
1	PASS		
mcae redundancy-group		Mandatory	1
1	PASS		
Local Logical Interface:ae3.0			

```

Peer Logical Interface :ae3.0
Configuration Item      Enforcement Level  Local Value
Peer Value      Result
-----
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      Result
-----
vrrp-group id      Mandatory      1
1                PASS
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.

```

user@EX9200-B> show multi-chassis mc-lag configuration-consistency
Configuration Item      Enforcement Level  Local Value
Peer Value      Result
-----
ICL interface      Mandatory      ae1
ae1              PASS
rstp-bridge-priority  Desirable      0
0              PASS
service-id      Mandatory      1
1              PASS
session-establishment-hold-time  Mandatory      300
300            PASS
local-ip-addr      Mandatory      172.16.32.6
172.16.32.5      PASS
backup-liveness-detection  Mandatory      10.92.76.2
10.92.76.4      PASS
iccp/bfd multiplier  Mandatory      4
4              PASS
bfd minimum-interval  Mandatory      2000
2000            PASS
session-establishment-hold-time  Mandatory      50
50              PASS

Local Physical Interface:ae2
Peer Physical Interface :ae2
Configuration Item      Enforcement Level  Local Value
Peer Value      Result
-----
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		
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		Mandatory	1
1	PASS		
prefer-status-control-active		Desirable	--
TRUE	PASS		
Local Logical Interface:ae2.0			
Peer Logical Interface :ae2.0			
Configuration Item		Enforcement Level	Local Value
Peer Value	Result		
-----	-----	-----	-----
-----	-----		
vlan membership		Mandatory	100
100	PASS		
interface-mode		Mandatory	trunk
trunk	PASS		
Local Physical Interface:ae3			
Peer Physical Interface :ae3			
Configuration Item		Enforcement Level	Local Value
Peer Value	Result		
-----	-----	-----	-----
-----	-----		
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 mode		Mandatory	active-active
active-active	PASS		
mCAE chassis-id		Mandatory	1
0	PASS		
mCAE redundancy-group		Mandatory	1
1	PASS		
prefer-status-control-active		Desirable	--
TRUE	PASS		
Local Logical Interface:ae3.0			
Peer Logical Interface :ae3.0			
Configuration Item		Enforcement Level	Local Value
Peer Value	Result		
-----	-----	-----	-----
-----	-----		
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      Peer Value      Result      Enforcement Level  Local Value
-----
vrrp-group id          1              PASS        Mandatory          1
ipv4 address            192.168.100.2/24  PASS        Mandatory          192.168.100.3/24

```

Meaning 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.

```

user@EX9200-A> show multi-chassis mc-lag configuration-consistency global-config
Configuration Item      Peer Value      Result      Enforcement Level  Local Value

```

Configuration Item	Result	Enforcement Level	Local Value
ICL interface		Mandatory	ae1
ae1	PASS		
rstp-bridge-priority		Desirable	0
0	PASS		
service-id		Mandatory	1
1	PASS		
session-establishment-hold-time		Mandatory	300
300	PASS		
local-ip-addr		Mandatory	172.16.32.5
172.16.32.6	PASS		
backup-liveness-detection		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

```
user@EX9200-B> show multi-chassis mc-lag configuration-consistency global-config
```

Configuration Item	Result	Enforcement Level	Local Value
Peer Value			
ICL interface		Mandatory	ae1
ae1	PASS		
rstp-bridge-priority		Desirable	0
0	PASS		
service-id		Mandatory	1
1	PASS		
session-establishment-hold-time		Mandatory	300
300	PASS		
local-ip-addr		Mandatory	172.16.32.5
172.16.32.6	PASS		
backup-liveness-detection		Mandatory	10.92.76.2
10.92.76.4	PASS		
iccp/bfd multiplier		Mandatory	4
4	PASS		
bfd minimum-interval		Mandatory	2000
2000	PASS		
session-establishment-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	Peer Value	Result	Enforcement Level	Local Value
vlan membership	100	PASS	Mandatory	100
interface-mode	trunk	PASS	Mandatory	trunk

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 Interface :ae1.0
```

Configuration Item	Peer Value	Result	Enforcement Level	Local Value
vlan membership	100	PASS	Mandatory	100
interface-mode	trunk	PASS	Mandatory	trunk

- Meaning**
- 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.

```
user@EX9200-A> show multi-chassis mc-lag configuration-consistency mcae-config
```

```
Local Physical Interface:ae2
```

```
Peer Physical Interface :ae2
```

Configuration Item		Enforcement Level	Local Value
Peer Value	Result		
-----	-----	-----	-----
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		
prefer-status-control-active		Desirable	TRUE
--	PASS		
mcae status-control		Mandatory	standby
active	PASS		
mcae deployment mode		Mandatory	active-active
active-active	PASS		
mcae chassis-id		Mandatory	0
1	PASS		
mcae redundancy-group		Mandatory	1
1	PASS		

```
Local Logical Interface:ae2.0
```

```
Peer Logical Interface :ae2.0
```

Configuration Item		Enforcement Level	Local Value
Peer Value	Result		
-----	-----	-----	-----
vlan membership		Mandatory	100
100	PASS		

```

interface-mode
trunk
PASS
Mandatory
trunk

Local Physical Interface:ae3
Peer Physical Interface :ae3
Configuration Item
Peer Value      Result
-----
lacp admin-key
3               PASS
lacp system-id
00:01:02:03:04:05 PASS
lacp periodic
0              PASS
lacp mode
0             PASS
prefer-status-control-active
--           PASS
mcae status-control
active       PASS
mcae deployment mode
active-active PASS
mcae chassis-id
1           PASS
mcae redundancy-group
1           PASS

Local Logical Interface:ae3.0
Peer Logical Interface :ae3.0
Configuration Item
Peer Value      Result
-----
vlan membership
100            PASS
interface-mode
trunk          PASS

```

2. Show the list of committed MC-LAG interface configuration parameters that passed or failed configuration consistency check on EX9200-B.

```

user@EX9200-B> show multi-chassis mc-lag configuration-consistency mcae-config
Local Physical Interface:ae2
Peer Physical Interface :ae2
Configuration Item
Peer Value      Result
-----
lacp admin-key
2               PASS
lacp system-id
00:01:02:03:04:05 PASS
lacp periodic
0              PASS
lacp mode
0             PASS
mcae status-control
standby        PASS
mcae deployment mode
active-active  PASS

```


mcae chassis-id		Mandatory	1
0	PASS		
mcae redundancy-group		Mandatory	1
1	PASS		
prefer-status-control-active		Desirable	--
TRUE	PASS		
Local Logical Interface:ae2.0			
Peer Logical Interface :ae2.0			
Configuration Item		Enforcement Level	Local Value
Peer Value	Result		
-----	-----	-----	-----
vlan membership		Mandatory	100
100	PASS		
interface-mode		Mandatory	trunk
trunk	PASS		
Local Physical Interface:ae3			
Peer Physical Interface :ae3			
Configuration Item		Enforcement Level	Local Value
Peer Value	Result		
-----	-----	-----	-----
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 mode		Mandatory	active-active
active-active	PASS		
mcae chassis-id		Mandatory	1
0	PASS		
mcae redundancy-group		Mandatory	1
1	PASS		
prefer-status-control-active		Desirable	--
TRUE	PASS		
Local Logical Interface:ae3.0			
Peer Logical Interface :ae3.0			
Configuration Item		Enforcement Level	Local Value
Peer Value	Result		
-----	-----	-----	-----
vlan membership		Mandatory	100
100	PASS		
interface-mode		Mandatory	trunk
trunk	PASS		

Meaning 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	Peer Value	Result	Enforcement Level	Local Value
-----	-----	-----	-----	-----
vrrp-group id				
1		PASS	Mandatory	1
ipv4 address				
192.168.100.3/24		PASS	Mandatory	192.168.100.2/24

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	Peer Value	Result	Enforcement Level	Local Value
-----	-----	-----	-----	-----
vrrp-group id				
1		PASS	Mandatory	1
ipv4 address				
192.168.100.2/24		PASS	Mandatory	192.168.100.3/24

Meaning 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 Peer Value	Result	Enforcement Level	Local Value
-----	-----	-----	-----
vrrp-group virtual-address 192.168.100.001	PASS	Mandatory	192.168.100.001
vrrp-group priority 100	PASS	Mandatory	150

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 Peer Value	Result	Enforcement Level	Local Value
-----	-----	-----	-----
vrrp-group virtual-address 192.168.100.001	PASS	Mandatory	192.168.100.001
vrrp-group priority 150	PASS	Mandatory	100

Meaning The output shows that the committed MC-LAG parameters related to VRRP configuration have passed configuration consistency check.

Related Documentation

- [Configuring Multichassis Link Aggregation on EX Series Switches](#)

