

EVPN Feature Guide for EX9200 Ethernet Switches



Modified: 2017-08-21

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

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

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

EVPN Feature Guide for EX9200 Ethernet Switches
Copyright © 2017 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/>. By downloading, installing or using such software, you agree to the terms and conditions of that EULA.

Table of Contents

	About the Documentation	xi
	Documentation and Release Notes	xi
	Supported Platforms	xi
	Using the Examples in This Manual	xi
	Merging a Full Example	xii
	Merging a Snippet	xii
	Documentation Conventions	xiii
	Documentation Feedback	xv
	Requesting Technical Support	xv
	Self-Help Online Tools and Resources	xv
	Opening a Case with JTAC	xvi
Part 1	Understanding Ethernet VPNs	
Chapter 1	Overview	3
	EVPN Overview for Switches	3
	EVPN Proxy ARP and ARP Suppression and Network Discovery Protocol and Network Discovery Protocol Suppression	5
Chapter 2	EVPN Standards	7
	Supported EVPN Standards	7
Chapter 3	Distributing Routes in VPNs	9
	Configuring an IGP on the PE and P Routers on EX9200 Switches	9
	Configuring IBGP Sessions Between PE Routers in VPNs on EX9200 Switches	9
	Configuring a Signaling Protocol and LSPs for VPNs on EX9200 Switches	11
	Using LDP for VPN Signaling	11
Part 2	Configuring EVPNs	
Chapter 4	Configuring Routing Instances for EVPN	15
	Configuring EVPN Routing Instances on EX9200 Switches	15
	Tracing EVPN Traffic and Operations	17
Chapter 5	Configuring Integrated Bridging and Routing Using EVPNs	19
	An EVPN with IRB Solution on EX9200 Switches Overview	19
	Need for an EVPN IRB Solution	20
	Implementing the EVPN IRB Solution	20
	Benefits of Implementing the EVPN IRB Solution	22
	Gateway MAC and IP Synchronization	22
	Layer 3 VPN Interworking	23

	IPv6 Support for IRB Interfaces with EVPN Using Neighborhood Discovery Protocol (NDP)	24
	Configuring an EVPN with IRB Solution on EX9200 Switches	24
	Example: Configuring an EVPN with IRB Solution on EX9200 Switches	26
Chapter 6	Configuring EVPN Multihoming	39
	EVPN Multihoming Overview	39
	Introduction to EVPN Multihoming	39
	Understanding EVPN Multihoming Concepts	41
	EVPN Multihoming Mode of Operation	43
	EVPN Multihoming Implementation	44
	New BGP NLRI	44
	New Extended Communities	47
	New EVPN Route Types	49
	Update to the MAC Forwarding Table	50
	Traffic Flow	51
	Aliasing	52
	EVPN Active-Active Multihoming and Multichassis Link Aggregation	54
	EVPN Active-Active Multihoming and IRB	54
	Sample Configuration	54
	Designated Forwarder Election	56
	DF Election Roles	56
	DF Election as Per RFC 7432	56
	Preference-Based DF Election	58
	DF Election for Virtual Switch	60
	Handling Failover	60
	ESIs on Physical, Aggregated Ethernet, and Logical Interfaces	61
	Convergence in an EVPN Network	61
	Configuring EVPN Active-Standby Multihoming	64
	Example: Configuring EVPN Active-Active Multihoming	67
	Example: Configuring an ESI on a Logical Interface With EVPN Multihoming	115
Part 3	Configuration Statements and Operational Commands	
Chapter 7	Configuration Statements	129
	encapsulation (Physical Interface)	130
	esi	136
	evpn	138
	extended-vlan-list	140
	instance-type	141
	interface (EVPN Routing Instances)	144
	interface (Routing Instances)	145
	interface-mac-limit (VPLS)	146
	ip-prefix-routes	148
	mac-statistics	152
	mac-table-size	154
	no-arp-suppression	155
	no-mac-learning	156
	packet-action	158
	routing-instances	161

Chapter 8

traceoptions (Protocols EVPN)	162
vlan-id (routing instance)	164
Operational Commands	165
clear evpn nd-table	166
clear ethernet-switching evpn nd-statistics	167
clear ethernet-switching evpn nd-table	168
show ethernet-switching evpn nd-statistics	169
show evpn arp-table	171
show evpn database	174
show evpn flood	177
show evpn instance	178
show evpn ip-prefix-database	186
show evpn l3-context	193
show evpn mac-table	196
show evpn nd-table	197
show evpn peer-gateway-macs	200
show vlans evpn nd-table	201

List of Figures

Part 2	Configuring EVPNs	
Chapter 5	Configuring Integrated Bridging and Routing Using EVPNs	19
	Figure 1: Inter-Subnet Traffic Forwarding	21
Chapter 6	Configuring EVPN Multihoming	39
	Figure 2: CE Device Multihomed to Two PE Routers	41
	Figure 3: Simple EVPN Topology	41
	Figure 4: Active-Active EVPN Multihoming	44
	Figure 5: EVPN Active-Active Multihoming Topology	69
	Figure 6: EVPN-MPLS Topology with Multihoming Active-Standby	116
	Figure 7: EVPN-MPLS Topology with Multihoming Active-Active	117

List of Tables

	About the Documentation	xi
	Table 1: Notice Icons	xiii
	Table 2: Text and Syntax Conventions	xiv
Part 2	Configuring EVPNs	
Chapter 6	Configuring EVPN Multihoming	39
	Table 3: Priority for NLRI Route Type	61
	Table 4: EVPN Multihoming Active-Standby: Configuring the Connection with CE1 on PE1 and PE2	116
	Table 5: Multihoming Active-Standby: Configuring Logical Interface on PE1 that Provides Services to a Different VLAN	117
	Table 6: Multihoming Active-Active: Configuring the Connection with CE1 on PE1 and PE2	117
	Table 7: EVPN Multihoming Active-Active: Configuring Interface on PE1 that Provides Services to a Different VLAN	118
Part 3	Configuration Statements and Operational Commands	
Chapter 8	Operational Commands	165
	Table 8: show ethernet-switching evpn nd-statistics Output Fields	169
	Table 9: show evpn arp-table Output Fields	171
	Table 10: show evpn database Output Fields	175
	Table 11: show evpn instance Output Fields	179
	Table 12: show evpn ip-prefix-database Output Fields	187
	Table 13: show evpn l3-context Output Fields	193
	Table 14: show evpn nd-table Output Fields	197
	Table 15: show vlans evpn nd-table Output Fields	201

About the Documentation

- Documentation and Release Notes on page xi
- Supported Platforms on page xi
- Using the Examples in This Manual on page xi
- Documentation Conventions on page xiii
- Documentation Feedback on page xv
- Requesting Technical Support on page xv

Documentation and Release Notes

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

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

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

Supported Platforms

For the features described in this document, the following platforms are supported:

- EX Series

Using the Examples in This Manual

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

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

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

Merging a Full Example

To merge a full example, follow these steps:

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

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

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

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

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

Merging a Snippet

To merge a snippet, follow these steps:

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

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

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

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

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

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

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

For more information about the **load** command, see [CLI Explorer](#).

Documentation Conventions

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

Table 1: Notice Icons







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

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

Table 2: Text and Syntax Conventions

Convention	Description	Examples
Bold text like this	Represents text that you type.	To enter configuration mode, type the configure command: user@host> configure
Fixed-width text like this	Represents output that appears on the terminal screen.	user@host> show chassis alarms No alarms currently active
<i>Italic text like this</i>	<ul style="list-style-type: none"> Introduces or emphasizes important new terms. Identifies guide names. Identifies RFC and Internet draft titles. 	<ul style="list-style-type: none"> A policy <i>term</i> is a named structure that defines match conditions and actions. <i>Junos OS CLI User Guide</i> RFC 1997, <i>BGP Communities Attribute</i>
<i>Italic text like this</i>	Represents variables (options for which you substitute a value) in commands or configuration statements.	Configure the machine's domain name: [edit] root@# set system domain-name <i>domain-name</i>
Text like this	Represents names of configuration statements, commands, files, and directories; configuration hierarchy levels; or labels on routing platform components.	<ul style="list-style-type: none"> To configure a stub area, include the stub statement at the [edit protocols ospf area area-id] hierarchy level. The console port is labeled CONSOLE.
< > (angle brackets)	Encloses optional keywords or variables.	stub <default-metric <i>metric</i> >;
(pipe symbol)	Indicates a choice between the mutually exclusive keywords or variables on either side of the symbol. The set of choices is often enclosed in parentheses for clarity.	broadcast multicast (<i>string1</i> <i>string2</i> <i>string3</i>)
# (pound sign)	Indicates a comment specified on the same line as the configuration statement to which it applies.	rsvp { # Required for dynamic MPLS only
[] (square brackets)	Encloses a variable for which you can substitute one or more values.	community name members [community-ids]
Indentation and braces ({ })	Identifies a level in the configuration hierarchy.	[edit] routing-options { static { route default { nexthop <i>address</i> ; retain; } } }
;(semicolon)	Identifies a leaf statement at a configuration hierarchy level.	

GUI Conventions

Table 2: Text and Syntax Conventions (*continued*)

Convention	Description	Examples
Bold text like this	Represents graphical user interface (GUI) items you click or select.	<ul style="list-style-type: none"> In the Logical Interfaces box, select All Interfaces. To cancel the configuration, click Cancel.
> (bold right angle bracket)	Separates levels in a hierarchy of menu selections.	In the configuration editor hierarchy, select Protocols>Ospf .

Documentation Feedback

We encourage you to provide feedback, comments, and suggestions so that we can improve the documentation. You can provide feedback by using either of the following methods:

- Online feedback rating system—On any page of the Juniper Networks TechLibrary site at <http://www.juniper.net/techpubs/index.html>, simply click the stars to rate the content, and use the pop-up form to provide us with information about your experience. Alternately, you can use the online feedback form at <http://www.juniper.net/techpubs/feedback/>.
- E-mail—Send your comments to techpubs-comments@juniper.net. Include the document or topic name, URL or page number, and software version (if applicable).

Requesting Technical Support

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

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

Self-Help Online Tools and Resources

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

- Find CSC offerings: <http://www.juniper.net/customers/support/>
- Search for known bugs: <http://www2.juniper.net/kb/>
- Find product documentation: <http://www.juniper.net/techpubs/>
- Find solutions and answer questions using our Knowledge Base: <http://kb.juniper.net/>
- Download the latest versions of software and review release notes: <http://www.juniper.net/customers/csc/software/>
- Search technical bulletins for relevant hardware and software notifications: <http://kb.juniper.net/InfoCenter/>
- Join and participate in the Juniper Networks Community Forum: <http://www.juniper.net/company/communities/>
- Open a case online in the CSC Case Management tool: <http://www.juniper.net/cm/>

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

Opening a Case with JTAC

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

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

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

PART 1

Understanding Ethernet VPNs

- [Overview on page 3](#)
- [EVPN Standards on page 7](#)
- [Distributing Routes in VPNs on page 9](#)

CHAPTER 1

Overview

- [EVPN Overview for Switches on page 3](#)
- [EVPN Proxy ARP and ARP Suppression and Network Discovery Protocol and Network Discovery Protocol Suppression on page 5](#)

EVPN Overview for Switches

An Ethernet VPN (EVPN) enables you to connect a group of dispersed customer sites using a Layer 2 virtual bridge. As with other types of VPNs, an EVPN comprises customer edge (CE) devices (host, router, or switch) connected to provider edge (PE) devices. The PE devices can include an MPLS edge switch (MES) that acts at the edge of the MPLS infrastructure. For the initial deployment of EVPNs using Juniper Networks equipment, you can configure an EX9200 switch to act as an MES. You can deploy multiple EVPNs within the network, each providing network connectivity to customers while ensuring that traffic sharing that network remains private.

The MESs are interconnected within the network by using label-switched paths (LSPs). The MPLS infrastructure allows you to take advantage of the MPLS functionality provided by the Junos operating system (Junos OS), including fast reroute, node and link protection, and standby secondary paths. For EVPNs, learning between MESs takes place in the control plane rather than in the data plane (as is the case with traditional network bridging). The control plane provides greater control over the learning process, allowing you to restrict which devices discover information about the network. You can also apply policies on the MESs, allowing you to carefully control how network information is distributed and processed. EVPNs utilize the BGP control plane infrastructure, providing greater scale and the ability to isolate groups of devices (hosts, servers, virtual machines, and so on) from each other.

The MESs attach an MPLS label to each MAC address learned from the CE devices. This label and MAC address combination is advertised to the other MESs in the control plane. Control plane learning enables load balancing and improves convergence times in the event of certain types of network failures. The learning process between the MESs and the CE devices is completed using the method best suited to each CE device (data plane learning, IEEE 802.1, LLDP, 802.1aq, and so on).

The policy attributes of an EVPN are similar to an IP VPN (for example, Layer 3 VPNs). Each EVPN routing instance requires that you configure a route distinguisher and one or

more route targets. A CE device attaches to an EVPN routing instance on an MES through an Ethernet interface that might be configured for one or more VLANs.

The following features are available for EVPNs:

- Ethernet connectivity between data centers spanning metropolitan area networks (MANs) and WANs
- One or more VLANs for each MAC VPN
- Automatic route distinguishers
- Dual-homed EVPN connection with active standby multihoming
- Starting with Junos OS Releases 16.1R4 and 16.2R2, the active-active mode for EVPN multihoming is supported.
- Starting with Junos OS Release 17.3R1, both pure type-5 routes and standard type-5 routes are supported on EX9200 switches. Use this feature, which advertises IP prefixes through EVPN, when the Layer 2 domain does not exist at the remote data centers or metro network peering points. For more information about how to configure, see [ip-prefix-routes](#).

The following features are not supported for EVPNs:

- Graceful restart, graceful Routing Engine switchover (GRES), and nonstop active routing (NSR)

Release History Table

Release	Description
17.3R1	Starting with Junos OS Release 17.3R1, both pure type-5 routes and standard type-5 routes are supported on EX9200 switches.
16.1R4	Starting with Junos OS Releases 16.1R4 and 16.2R2, the active-active mode for EVPN multihoming is supported.

Related Documentation

- [Supported EVPN Standards on page 7](#)
- [Example: Configuring EVPN Active-Active Multihoming on page 67](#)

EVPN Proxy ARP and ARP Suppression and Network Discovery Protocol and Network Discovery Protocol Suppression

Starting with Junos OS Release 17.2R1, MX Series routers and EX9200 that function as provider edge (PE) devices in an Ethernet VPN-MPLS (EVPN-MPLS) or Ethernet VPN-Virtual Extensible LAN (EVPN-VXLAN) environment support proxy Address Resolution Protocol (ARP) and ARP suppression and Network Discovery Protocol (NDP) and NDP suppression. Starting with Junos OS Release 17.3R1, QFX10000 Series switches that function as provider edge (PE) devices in an Ethernet VPN-Virtual Extensible LAN (EVPN-VXLAN) environment support proxy Address Resolution Protocol (ARP) and ARP suppression and Network Discovery Protocol (NDP) and NDP suppression.

By default, proxy ARP and ARP suppression and NDP and NDP suppression are enabled. For these features to work properly, the configuration of an integrated and routing (IRB) interface on the PE device is required.

Proxy ARP and Network Discovery Protocol snooping is enabled by default for all EVPN-MPLS or EVPN-VXLAN bridge domains and VLANs. ARP or Network Discovery Protocol (NDP) packets generated from a local CE are snooped. ARP and NDP packets generated from a remote PE through core-facing interfaces, however, are not snooped.

IRB interfaces configured on a PE device deliver ARP requests and NDP requests from both local and remote customer edge (CE) devices. When a PE device receives an ARP request or NDP request from a CE device, the PE device searches its media access control (MAC)-IP address bindings database for the requested IP address. If the PE device finds the MAC-IP address binding in its database, it responds to the request. If the device does not find the MAC-IP address binding, it swaps the source MAC address in the request with the MAC address of the IRB interface on which the request was received and sends the request to all interfaces.

Even when a PE device responds to an ARP request or NDP request, ARP packets and NDP might still be flooded across the WAN. ARP suppression and NDP suppression prevent this flooding from occurring.

You can disable the suppression of ARP packets and NDP packets by specifying the **no-arp-suppression** configuration statement. However, if you do so, be aware of the following implications:

- If the PE device does not find the MAC-IP address binding in its database, it cannot forward the ARP or NDP request. Instead, the device discards the ARP or NDP request, and the packets flood through the Layer 2 domain because the ARP or NDP request is considered to be Layer 2 broadcast traffic. Therefore, we recommend that ARP and NDP suppression remains enabled.
- Disabling the suppression of ARP packets and NDP packets on a PE device essentially disables the proxy ARP and NDP functionality.

Therefore, we recommend that ARP suppression and NDP suppression remains enabled.

Proxy ARP and ARP suppression and NDP and NDP suppression are supported with single-homed CE devices in active mode and multihomed CEs in active-active mode. In

a multihoming active-active situation, the database of MAC-IP address bindings are synchronized between the PE devices that act as designated forwarder (DF) and non-designated forwarder (non-DF).

Release History Table

Release	Description
17.3R1	Starting with Junos OS Release 17.3R1, QFX10000 Series switches that function as provider edge (PE) devices in an Ethernet VPN-Virtual Extensible LAN (EVPN-VXLAN) environment support proxy Address Resolution Protocol (ARP) and ARP suppression and Network Discovery Protocol (NDP) and NDP suppression.
17.2R1	Starting with Junos OS Release 17.2R1, MX Series routers and EX9200 that function as provider edge (PE) devices in an Ethernet VPN-MPLS (EVPN-MPLS) or Ethernet VPN-Virtual Extensible LAN (EVPN-VXLAN) environment support proxy Address Resolution Protocol (ARP) and ARP suppression and Network Discovery Protocol (NDP) and NDP suppression.

**Related
Documentation**

- [no-arp-suppression on page 155](#)

CHAPTER 2

EVPN Standards

- [Supported EVPN Standards on page 7](#)

Supported EVPN Standards

Junos OS supports the following RFCs and Internet drafts that define standards for EVPNs:

- RFC 4364, *BGP/MPLS IP Virtual Private Networks (VPNs)*
- RFC 4761, *Virtual Private LAN Service (VPLS) Using BGP for Auto-Discovery and Signaling*
- RFC 7432, *BGP MPLS-Based Ethernet VPN*

The following features are not supported:

- Automatic derivation of Ethernet segment (ES) values. Only static ES configurations are supported.
- Host proxy ARP.
- MAC mobility extended community.
- VLAN bundle service interface.

Related Documentation

- [EVPN Overview](#)
- [Accessing Standards Documents on the Internet](#)

CHAPTER 3

Distributing Routes in VPNs

- [Configuring an IGP on the PE and P Routers on EX9200 Switches on page 9](#)
- [Configuring IBGP Sessions Between PE Routers in VPNs on EX9200 Switches on page 9](#)
- [Configuring a Signaling Protocol and LSPs for VPNs on EX9200 Switches on page 11](#)

Configuring an IGP on the PE and P Routers on EX9200 Switches

For Layer 2 VPNs, Layer 3 VPNs, virtual-router routing instances, VPLS, EVPNs, and Layer 2 circuits to function properly, the service provider's PE and P routers must be able to exchange routing information. For this to happen, you must configure either an IGP (such as OSPF or IS-IS) or static routes on these routers. You configure the IGP on the master instance of the routing protocol process at the **[edit protocols]** hierarchy level, not within the routing instance used for the VPN—that is, not at the **[edit routing-instances]** hierarchy level.

When you configure the PE router, do not configure any summarization of the PE router's loopback addresses at the area boundary. Each PE router's loopback address should appear as a separate route.

Related Documentation

- [Understanding IS-IS Configuration](#)
- [Example: Configuring IS-IS](#)
- [OSPF Feature Guide](#)

Configuring IBGP Sessions Between PE Routers in VPNs on EX9200 Switches

You must configure an IBGP session between the PE routers to allow the PE routers to exchange information about routes originating and terminating in the VPN. The PE routers rely on this information to determine which labels to use for traffic destined for remote sites.

Configure an IBGP session for the VPN as follows:

```
[edit protocols]
bgp {
  group group-name {
```

```
type internal;  
local-address ip-address;  
family evpn {  
    signaling;  
}  
family (inet-vpn | inet6-vpn) {  
    unicast;  
}  
family l2vpn {  
    signaling;  
}  
neighbor ip-address;  
}
```

The IP address in the **local-address** statement is the address of the loopback interface on the local PE router. The IBGP session for the VPN runs through the loopback address. (You must also configure the loopback interface at the **[edit interfaces]** hierarchy level.)

The IP address in the **neighbor** statement is the loopback address of the neighboring PE router.

The **family** statement allows you to configure the IBGP session for Layer 2 VPNs, VPLS, EVPNs or for Layer 3 VPNs.

- To configure an IBGP session for Layer 2 VPNs and VPLS, include the **signaling** statement at the **[edit protocols bgp group *group-name* family l2vpn]** hierarchy level:

```
[edit protocols bgp group group-name family l2vpn]  
signaling;
```

- To configure an IBGP session for EVPNs, include the **signaling** statement at the **[edit protocols bgp group *group-name* family evpn]** hierarchy level:

```
[edit protocols bgp group group-name family evpn]  
signaling;
```

- To configure an IPv4 IBGP session for Layer 3 VPNs, configure the **unicast** statement at the **[edit protocols bgp group *group-name* family inet-vpn]** hierarchy level:

```
[edit protocols bgp group group-name family inet-vpn]  
unicast;
```

- To configure an IPv6 IBGP session for Layer 3 VPNs, configure the **unicast** statement at the **[edit protocols bgp group *group-name* family inet6-vpn]** hierarchy level:

```
[edit protocols bgp group group-name family inet6-vpn]  
unicast;
```



NOTE: You can configure both **family inet** and **family inet-vpn** or both **family inet6** and **family inet6-vpn** within the same peer group. This allows you to enable support for both IPv4 and IPv4 VPN routes or both IPv6 and IPv6 VPN routes within the same peer group.

- Related Documentation**
- [Configuring a Signaling Protocol and LSPs for VPNs on EX9200 Switches on page 11](#)

Configuring a Signaling Protocol and LSPs for VPNs on EX9200 Switches

For VPNs to function, you must enable the LDP signaling protocol on the provider edge (PE) routers and on the provider (P) routers.

To enable the LDP signaling protocol, perform the steps in the following section:

- [Using LDP for VPN Signaling on page 11](#)

Using LDP for VPN Signaling

To use LDP for VPN signaling, perform the following steps on the PE and provider (P) routers:

1. Configure LDP on the interfaces in the core of the network by including the **ldp** statement at the **[edit protocols]** hierarchy level.

You need to configure LDP only on the interfaces between PE routers or between PE and P routers. You can think of these as the “core-facing” interfaces. You do not need to configure LDP on the interface between the PE and customer edge (CE) routers.

```
[edit]
protocols {
  ldp {
    interface type-fpc/pic/port;
  }
}
```

2. Configure the MPLS address family on the interfaces on which you enabled LDP (the interfaces you configured in Step 1) by including the **family mpls** statement at the **[edit interfaces *type-fpc/pic/port* unit *logical-unit-number*]** hierarchy level.

```
[edit]
interfaces {
  type-fpc/pic/port {
    unit logical-unit-number {
      family mpls;
    }
  }
}
```

3. Configure OSPF or IS-IS on each PE and P router.

You configure these protocols at the master instance of the routing protocol, not within the routing instance used for the VPN.

- To configure OSPF, include the **ospf** statement at the **[edit protocols]** hierarchy level. At a minimum, you must configure a backbone area on at least one of the router's interfaces.

```
[edit]
```

```
protocols {
  ospf {
    area 0.0.0.0 {
      interface type-fpc/pic/port;
    }
  }
}
```

- To configure IS-IS, include the **isis** statement at the **[edit protocols]** hierarchy level and configure the loopback interface and International Organization for Standardization (ISO) family at the **[edit interfaces]** hierarchy level. At a minimum, you must enable IS-IS on the router, configure a network entity title (NET) on one of the router's interfaces (preferably the loopback interface, lo0), and configure the ISO family on all interfaces on which you want IS-IS to run. When you enable IS-IS, Level 1 and Level 2 are enabled by default. The following is the minimum IS-IS configuration. In the **address** statement, **address** is the NET.

```
[edit]
interfaces {
  lo0 {
    unit logical-unit-number {
      family iso {
        address address;
      }
    }
  }
  type-fpc/pic/port {
    unit logical-unit-number {
      family iso;
    }
  }
}
protocols {
  isis {
    interface all;
  }
}
```

For more information about configuring OSPF and IS-IS, see the *OSPF Feature Guide* and *IS-IS Feature Guide*.

Related Documentation

- [EVPN Overview for Switches on page 3](#)

PART 2

Configuring EVPNs

- [Configuring Routing Instances for EVPN on page 15](#)
- [Configuring Integrated Bridging and Routing Using EVPNs on page 19](#)
- [Configuring EVPN Multihoming on page 39](#)

Configuring Routing Instances for EVPN

- [Configuring EVPN Routing Instances on EX9200 Switches on page 15](#)
- [Tracing EVPN Traffic and Operations on page 17](#)

Configuring EVPN Routing Instances on EX9200 Switches

To configure an EVPN routing instance, complete the following configuration on the PE router (or on the MPLS edge switch) within the EVPN service provider's network:

1. Configure the EVPN routing instance name using the **routing-instances** statement at the **[edit]** hierarchy level:

```
routing-instances routing-instance-name {...}
```

2. Configure the **evpn** option for the **instance-type** statement at the **[edit routing-instances routing-instance-name]** hierarchy level:

```
instance-type evpn;
```

3. Configure the interfaces for handling EVPN traffic between the MES and the CE device using the **interface** statement at the **[edit routing-instances routing-instance-name]** hierarchy level:

```
interface interface-name;
```

4. Configure a VLAN identifier for the EVPN routing instance using the **vlan-id** statement at the **[edit routing-instances routing-instance-name]** hierarchy level:

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

5. Configure a route distinguisher on a PE router by including the **route-distinguisher** statement:

```
route-distinguisher (as-number:number | ip-address:number);
```

Each routing instance that you configure on a PE router must have a unique route distinguisher associated with it. VPN routing instances need a route distinguisher to help BGP to distinguish between potentially identical network layer reachability information (NLRI) messages received from different VPNs. If you configure different VPN routing instances with the same route distinguisher, the commit fails.

For a list of the hierarchy levels at which you can include this statement, see the statement summary for this statement.

The route distinguisher is a 6-byte value that you can specify in one of the following formats:

- **as-number:number**, where **as-number** is an autonomous system (AS) number (a 2-byte value) and **number** is any 4-byte value. The AS number can be in the range 1 through 65,535. We recommend that you use an Internet Assigned Numbers Authority (IANA)-assigned, nonprivate AS number, preferably the Internet service provider's (ISP's) own or the customer's own AS number.
 - **ip-address:number**, where **ip-address** is an IP address (a 4-byte value) and **number** is any 2-byte value. The IP address can be any globally unique unicast address. We recommend that you use the address that you configure in the **router-id** statement, which is a nonprivate address in your assigned prefix range.
6. Configure either import and export policies for the EVPN routing table, or configure the default policies using the **vrf-target** statement configured at the **[edit routing-instances routing-instance-name]** hierarchy level.

See *Configuring Policies for the VRF Table on PE Routers in VPNs*.

7. Configure each EVPN interface for the EVPN routing instance:
- a. Configure interface encapsulation for the CE facing interfaces at the **[edit interfaces interface-name encapsulation]** hierarchy level. . Supported encapsulations for EX9200 switches are: (**extended-vlan-bridge** | **flexible-ethernet-services** | **vlan-bridge**).
 - b. Configure **vlan-bridge** encapsulation on the logical interface at the **[edit interfaces interface-name flexible-vlan-tagging encapsulation flexible-ethernet-services unit 0 encapsulation]** hierarchy level.
 - c. (Optional) Allow the EVPN to establish a connection to the CE device even if the CE device interface encapsulation and the EVPN interface encapsulations do not match by including the **ignore-encapsulation-mismatch** statement at the **[edit routing-instances routing-instance-name protocols evpn interface interface-name]** hierarchy level.
8. Specify the maximum number of media access control (MAC) addresses that can be learned by the EVPN routing instance by including the **interface-mac-limit** statement.

You can configure the same limit for all interfaces configured for a routing instance by including this statement at the **[edit routing-instances routing-instance-name protocols evpn]** hierarchy level. You can also configure a limit for a specific interface by including this statement at the **[edit routing-instances routing-instance-name protocols evpn interface interface-name]** hierarchy level.

By default, packets with new source MAC addresses are forwarded after the MAC address limit is reached. You can alter this behavior by including the **packet-action drop** statement at either the **[edit routing-instances routing-instance-name protocols evpn interface-mac-limit]** or the **[edit routing-instances routing-instance-name protocols evpn interface interface-name]** hierarchy level. If you configure this statement, packets

from new source MAC addresses are dropped once the configured MAC address limit is reached.

9. Enable MAC accounting for the EVPN by including the `mac-statistics` statement at the `[edit routing-instances routing-instance-name protocols evpn]` hierarchy level.
10. Specify the number of addresses that can be stored in the MAC routing table using the `mac-table-size` statement at the `[edit routing-instances routing-instance-name protocols evpn]` hierarchy level.

You can optionally configure the `packet-action drop` option to specify that packets for new source MAC addresses be dropped once the MAC address limit is reached. If you do not configure this option, packets for new source MAC addresses are forwarded.

11. Disable MAC learning by including the `no-mac-learning` statement at either the `[edit routing-instances routing-instance-name protocols evpn]` hierarchy level to apply this behavior to all of the devices configured for an EVPN routing instance or at the `[edit routing-instances routing-instance-name protocols evpn interface interface-name]` hierarchy level to apply this behavior to just one of the CE devices.

Related Documentation

- [Configuring Policies for the VRF Table on PE Routers in VPNs](#)
- [Configuring Routing Instances on PE Routers in VPNs](#)
- [Tracing EVPN Traffic and Operations on page 17](#)

Tracing EVPN Traffic and Operations

To configure the EVPN routing instance to trace a variety of different parameters related to EVPN operation:

1. Specify the name of one or more EVPN trace files using the `file` option for the `traceoptions` statement at the `[edit routing-instances routing-instance-name protocols evpn]` hierarchy level:

```
traceoptions {
  file filename <files number> <size size> <world-readable | no-world-readable>;
}
```

The `file` option includes the following sub-options:

- **filename**—Specify the name of the file to receive the output of the tracing operation. Enclose the name within quotation marks. All files are placed in the directory `/var/log`.
- **files number**—(Optional) Maximum number of trace files. When a trace file named `trace-file` reaches its maximum **size**, it is renamed `trace-file.0`, then `trace-file.1`, and so on, until the specified maximum **number** of trace files specified is reached. Then the oldest trace file is overwritten.
- **size size**—(Optional) Maximum size of each trace file. When a trace file named `trace-file` reaches its maximum size, it is renamed `trace-file.0`, then `trace-file.1`, and

so on, until the maximum number of trace files is reached. Then the oldest trace file is overwritten.

- **world-readable | no-world-readable**—(Optional) Enable unrestricted file access or restrict file access to the user who created the file.

2. Specify the **flag** option for the **traceoptions** statement:

```
traceoptions {  
    flag flag <flag-modifier> <disable>;  
}
```

The **flag** option allows you to specify the scope of the trace by including one of the following sub-options:

- **all**—All EVPN tracing options
- **error**—Error conditions
- **general**—General events
- **mac-database**—MAC route database in the EVPN routing instance
- **nlri**—EVPN advertisements received or sent by means of the BGP
- **normal**—Normal events
- **oam**—OAM messages
- **policy**—Policy processing
- **route**—Routing information
- **state**—State transitions
- **task**—Routing protocol task processing
- **timer**—Routing protocol timer processing
- **topology**—EVPN topology changes caused by reconfiguration or advertisements received from other PE routers using BGP

You can also specify one of the following modifiers for any of the traceoptions flags:

- **detail**—Provide detailed trace information.
- **disable**—Disable this trace flag.
- **receive**—Trace received packets.
- **send**—Trace sent packets.

- Related Documentation**
- *Configuring EVPN Routing Instances*
 - [traceoptions on page 162](#)

CHAPTER 5

Configuring Integrated Bridging and Routing Using EVPNs

- [An EVPN with IRB Solution on EX9200 Switches Overview on page 19](#)
- [Configuring an EVPN with IRB Solution on EX9200 Switches on page 24](#)
- [Example: Configuring an EVPN with IRB Solution on EX9200 Switches on page 26](#)

An EVPN with IRB Solution on EX9200 Switches Overview

A data center service provider (DCSP) hosts the data center for its multiple customers on a common physical network. To each customer (also called a tenant), the service looks like a full-fledged data center that can expand to 4094 VLANs and all private subnets. For disaster recovery, high availability, and optimization of resource utilization, it is common for the DCSP to span the data center over multiple sites. To deploy data center services, a DCSP faces the following main challenges:

- Extending Layer 2 domains across more than one data center site. This requires optimal intra-subnet traffic forwarding.
- Supporting optimal inter-subnet traffic forwarding and optimal routing in the event of virtual machine (VM) motion.
- Supporting multiple tenants with independent VLAN and subnet space.

Ethernet VPN (EVPN) is targeted to handle all the preceding challenges, wherein:

- The basic EVPN functionality enables optimal intra-subnet traffic forwarding
- Implementing the integrated routing and bridging (IRB) solution in an EVPN deployment enables optimal inter-subnet traffic forwarding
- Configuring EVPN with virtual switch support enables multiple tenants with independent VLAN and subnet space

The following sections describe the integrated routing and bridging (IRB) solution for EVPNs:

- [Need for an EVPN IRB Solution on page 20](#)
- [Implementing the EVPN IRB Solution on page 20](#)

- [Benefits of Implementing the EVPN IRB Solution on page 22](#)
- [IPv6 Support for IRB Interfaces with EVPN Using Neighborhood Discovery Protocol \(NDP\) on page 24](#)

Need for an EVPN IRB Solution

EVPN is a technology used to provide Layer 2 extension and interconnection across an IP/MPLS core network to different physical sites belonging to a single Layer 2 domain. In a data center environment with EVPN, there is a need for both Layer 2 (intra-subnet traffic) and Layer 3 (inter-subnet traffic) forwarding and potentially interoperability with tenant Layer 3 VPNs.

With only a Layer 2 solution, there is no optimum forwarding of inter-subnet traffic, even when the traffic is local, for instance, when both the subnets are on the same server.

With only a Layer 3 solution, the following issues for intra-subnet traffic can arise:

- MAC address aliasing issue where duplicate MAC addresses are not detected.
- TTL issue for applications that use TTL 1 to confine traffic within a subnet.
- IPv6 link-local addressing and duplicate address detection that relies on Layer 2 connectivity.
- Layer 3 forwarding does not support the forwarding semantics of a subnet broadcast.
- Support of non-IP applications that require Layer 2 forwarding.

Because of the above mentioned shortcomings of a pure Layer 2 and Layer 3 solution, there is a need for a solution incorporating optimal forwarding of both Layer 2 and Layer 3 traffic in the data center environment when faced with operational considerations such as Layer 3 VPN interoperability and virtual machine (VM) mobility.

An EVPN-based integrated routing and bridging (IRB) solution provides optimum unicast and multicast forwarding for both intra-subnets and inter-subnets within and across data centers.

The EVPN IRB feature is useful for service providers operating in an IP/MPLS network that provides both Layer 2 VPN or VPLS services and Layer 3 VPN services who want to extend their service to provide cloud computation and storage services to their existing customers.

Implementing the EVPN IRB Solution

An EVPN IRB solution provides the following:

- Optimal forwarding for intra-subnet (Layer 2) traffic.
- Optimal forwarding for inter-subnet (Layer 3) traffic.
- Support for ingress replication for multicast traffic.
- Support for network-based as well as host-based overlay models.
- Support for consistent policy-based forwarding for both Layer 2 and Layer 3 traffic.

Junos OS supports several models of EVPN configuration to satisfy the individual needs of EVPN and data center cloud services customers. To provide flexibility and scalability, multiple VLANs can be defined within a particular EVPN instance. Likewise, one or more EVPN instances can be associated with a single Layer 3 VPN virtual routing and forwarding (VRF). In general, each data center tenant is assigned a unique Layer 3 VPN VRF, while a tenant could comprise one or more EVPN instances and one or more VLANs per EVPN instance. To support this model, each configured VLAN (including the default VLAN for an EVPN instance) requires an IRB interface to perform the Layer 2 and Layer 3 functions. Each VLAN or IRB interface maps to a unique IP subnet in the VRF.

There are two major functions that are supported for IRB in EVPN.

- Host MAC-IP synchronization

This includes:

- Advertising the IP address along with the MAC advertisement route in EVPN. This is done by using the IP field in the EVPN MAC advertisement route.
- The receiving PE router installs MAC into the EVPN instance (EVI) table and installs IP into the associated VRF.

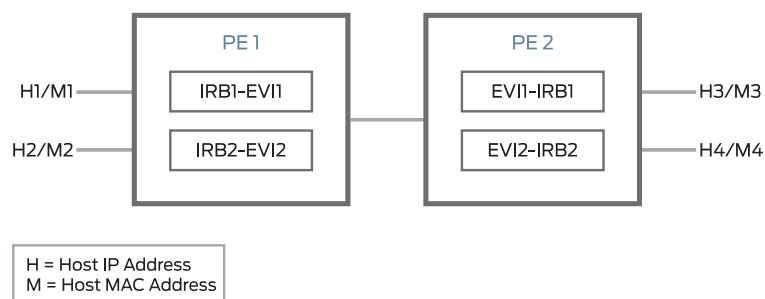
- Gateway MAC-IP synchronization

This includes:

- Advertising all local IRB MAC and IP addresses in an EVPN. This is achieved by including the default gateway extended community in the EVPN MAC advertisement route.
- The receiving PE creates a forwarding state to route packets destined for the gateway MAC, and a proxy ARP is done for the gateway IP with the MAC advertised in the route.

Figure 1 on page 21 illustrates the inter-subnet traffic forwarding between two provider edge (PE) devices—PE1 and PE2. The IRB1 and IRB2 interfaces on each PE device belong to a different subnet, but they share a common VRF.

Figure 1: Inter-Subnet Traffic Forwarding



The inter-subnet traffic forwarding is performed as follows:

1. PE2 advertises the H3-M3 and H4-M4 binding to PE1. Similarly, PE1 advertises the H1-M1 and H2-M2 binding to PE2.
2. PE1 and PE2 install the MAC address in the corresponding EVI MAC table, whereas the IP routes are installed in the shared VRF.
3. The advertising PE device is set as the next hop for the IP routes.
4. If H1 sends packets to H4, the packets are sent to IRB1 on PE1.
5. IP lookup for H4 happens in the shared VRF on PE1. Because the next hop for the H4 IP is PE2 (the advertising PE), an IP unicast packet is sent to PE2.
6. PE1 rewrites the MAC header based on the information in the VRF route, and PE2 performs a MAC lookup to forward the packet to H4.

Benefits of Implementing the EVPN IRB Solution

The main goal of the EVPN IRB solution is to provide optimal Layer 2 and Layer 3 forwarding. The solution is required to efficiently handle inter-subnet forwarding as well as virtual machine (VM) mobility. VM mobility refers to the ability of a VM to migrate from one server to another within the same or a different data center while retaining its existing MAC and IP address. Providing optimal forwarding for inter-subnet traffic and effective VM mobility involves solving two problems – the default gateway problem and the triangular routing problem.

- [Gateway MAC and IP Synchronization on page 22](#)
- [Layer 3 VPN Interworking on page 23](#)

Gateway MAC and IP Synchronization

In an EVPN IRB deployment, the IP default gateway for a VM is the IP address configured on the IRB interface of the provider edge (PE) router corresponding to the VLAN of which the VM is a member. The default gateway problem arises because a VM does not flush its ARP table when relocating from one server to another and continues sending packets with the destination MAC address set to that of the original gateway. If the old and new servers are not part of the same Layer 2 domain (the new Layer 2 domain could be within the current data center or a new data center), the gateway previously identified is no longer the optimal or local gateway. The new gateway needs to identify packets containing the MAC addresses of other gateways on remote PE routers and forward the traffic as if the packets were destined to the local gateway itself. At minimum, this functionality requires each PE router to advertise its gateway or IRB MAC and IP addresses to all other PE routers in the network. The gateway address exchange can be accomplished using the standard MAC route advertisement message (including the IP address parameter) and tagging that route with the default gateway extended community so that the remote

PE routers can distinguish the gateway MAC advertisement routes from normal MAC advertisement routes.

Layer 3 VPN Interworking

The inter-data center aspect of the EVPN IRB solution involves routing between VMs that are present in different data centers or routing between a host site completely outside of the data center environment and a VM within a data center. This solution relies on the ability of EVPN MAC route advertisements to carry both MAC address and IP address information. The local MAC learning functionality of the PE router is extended to also capture IP address information associated with MAC addresses learned locally. That IP-MAC address mapping information is then distributed to each PE router through normal EVPN procedures. When a PE router receives such MAC and IP information, it installs the MAC route in the EVPN instance as well as a host route for the associated IP address in the Layer 3 VPN VRF corresponding to that EVPN instance. When a VM moves from one data center to another, normal EVPN procedures result in the MAC and IP address being advertised from the new PE router which the VM resides behind. The host route installed in the VRF associated with an EVPN solicits Layer 3 traffic destined to that VM to the new PE router and avoids triangular routing between the source, the former PE router the VM resided behind, and the new PE router.

BGP scalability is a potential concern with the inter-data center triangular routing avoidance solution because of the potential for injection of many host routes into Layer 3 VPN. With the method previously described, in the worst case there is an IP host route for each MAC address learned through the local EVPN MAC learning procedures or through a MAC advertisement message received from a remote PE router. BGP route target filtering can be used to limit distribution of such routes.

The following functional elements are required to implement the inter-data center triangular routing avoidance using Layer 3 inter-subnet forwarding procedures:

1. The source host sends an IP packet using its own source MAC and IP address with the destination MAC of the IRB interface of the local PE router and the IP address of the destination host.
2. When the IRB interface receives the frame with its MAC as the destination, it performs a Layer 3 lookup in the VRF associated with the EVPN instance to determine where to route the packet.
3. In the VRF, the PE router finds the Layer 3 route derived from a MAC plus an IP EVPN route received from the remote PE router earlier. The destination MAC address is then changed to the destination MAC address corresponding to the destination IP.
4. The packet is then forwarded to the remote PE router serving the destination host using MPLS, using the label corresponding to the EVPN instance of which the destination host is a member.

5. The egress PE router receiving the packet performs a Layer 2 lookup for the destination host's MAC and sends the packet to the destination host on the attached subnet via the egress PE router's IRB interface.
6. Because the ingress PE router is performing Layer 3 routing, the IP TTL is decremented.

IPv6 Support for IRB Interfaces with EVPN Using Neighborhood Discovery Protocol (NDP)

Starting in Junos OS Release 17.3R1, IPv6 addresses are supported on IRB interfaces with EVPN using NDP. The following capabilities are introduced for IPv6 support with EVPN:

- IPv6 addresses on IRB interfaces in master routing instances
- Learning IPv6 neighbors from solicited neighbor advertisement (NA) messages
- Neighbor solicitation (NS) and neighbor advertisement (NA) packets on the IRB interfaces are disabled from network core
- Virtual gateway addresses are used as Layer 3 addresses
- Host MAC-IP synchronization for IPv6

You can configure the IPv6 addresses in the IRB interface at the **[edit interfaces irb]** hierarchy level.

Release History Table

Release	Description
17.3R1	Starting in Junos OS Release 17.3R1, IPv6 addresses are supported on IRB interfaces with EVPN

Related Documentation

- [Example: Configuring an EVPN with IRB Solution on EX9200 Switches on page 26](#)
- [EVPN Overview for Switches on page 3](#)

Configuring an EVPN with IRB Solution on EX9200 Switches

You can configure an Ethernet VPN (EVPN) with IRB solution to enable Layer 2 switching and Layer 3 routing operations within a single node, thus avoiding extra hops for inter-subnet traffic. The EVPN IRB solution eliminates the default gateway problem using the gateway MAC and IP synchronization, and avoids the triangular routing problem with Layer 3 interworking by creating IP host routes for virtual machines (VMs) in the tenant virtual routing and forwarding (VRF) routing instances.

Before you begin:

1. Configure the switch interfaces.
2. Configure the router ID and autonomous system number for the device.
3. Enable the chained composite next hop for EVPN.
4. Configure OSPF or any other IGP protocol.

5. Configure a BGP internal group.
6. Include the EVPN signaling network layer reachability information (NLRI) to the internal BGP group.
7. Configure LDP.
8. Configure MPLS.

To configure the PE device:

1. Configure the EVPN routing instance.

```
[edit routing-instances]
user@PE1# set evpn-instance instance-type evpn
```

2. Set the VLAN identifier for the bridging domain in the EVPN routing instance.

```
[edit routing-instances]
user@PE1# set evpn-instance vlan-id VLAN-ID
```

3. Configure the interface name for the EVPN routing instance.

```
[edit routing-instances]
user@PE1# set evpn-instance interface CE-facing-interface
```

4. Configure the IRB interface as the routing interface for the EVPN routing instance.

```
[edit routing-instances]
user@PE1# set evpn-instance l3-interface irb.0
```

5. Configure the route distinguisher for the EVPN routing instance.

```
[edit routing-instances]
user@PE1# set evpn-instance route-distinguisher route-distinguisher-value
```

6. Configure the VPN routing and forwarding (VRF) target community for the EVPN routing instance.

```
[edit routing-instances]
user@PE1# set evpn-instance vrf-target vrf-target-value
```

7. Assign the interface name that connects the PE device site to the VPN.

```
[edit routing-instances]
user@PE1# set evpn-instance protocols evpn interface CE-facing-interface
```

8. Configure the VRF routing instance.

```
[edit routing-instances]
user@PE1# set vrf-instance instance-type vrf
```

9. Configure the IRB interface as the routing interface for the VRF routing instance.

```
[edit routing-instances]
```

```
user@PE1# set vrf-instance interface irb.0
```

10. Configure the route distinguisher for the VRF routing instance.

```
[edit routing-instances]
user@PE1# set vrf-instance route-distinguisher route-distinguisher-value
```

11. Configure the VRF label for the VRF routing instance.

```
[edit routing-instances]
user@PE1# set vrf-instance vrf-table-label
```

12. Verify and commit the configuration.

For example:

```
[edit routing-instances]
user@PE1# set evpna instance-type evpn
user@PE1# set evpna vlan-id 10
user@PE1# set evpna interface ge-1/1/8.0
user@PE1# set evpna l3-interface irb.0
user@PE1# set evpna route-distinguisher 100.255.0.1:100
user@PE1# set evpna vrf-target target:100:100
user@PE1# set evpna protocols evpn interface ge-1/1/8.0
user@PE1# set vrf instance-type vrf
user@PE1# set vrf interface irb.0
user@PE1# set vrf route-distinguisher 100.255.0.1:300
user@PE1# set vrf vrf-target target:100:300
user@PE1# set vrf vrf-table-label

[edit]
user@PE1# commit
commit complete
```

Related Documentation • [Example: Configuring an EVPN with IRB Solution on EX9200 Switches on page 26](#)

Example: Configuring an EVPN with IRB Solution on EX9200 Switches

This example shows how to configure an integrated routing and bridging (IRB) solution in an Ethernet VPN (EVPN) deployment.

- [Requirements on page 26](#)
- [Overview on page 27](#)
- [Configuration on page 27](#)
- [Verification on page 33](#)

Requirements

This example uses the following hardware and software components:

- Two EX9200 switches configured as PE routers

- Junos OS Release 14.2 or later running on all the PE routers

Before you begin:

1. Configure the switch interfaces.
2. Configure OSPF or any other IGP protocol.
3. Configure BGP.
4. Configure LDP.
5. Configure MPLS.

Overview

In an EVPN solution, multiple VLANs can be defined within a particular EVPN instance, and one or more EVPN instances can be associated with a single Layer 3 VPN VRF. In general, each data center tenant is assigned a unique Layer 3 VPN virtual route forwarding (VRF), although the tenant can comprise one or more EVPN instances or VLANs per EVPN instance.

To support this flexibility and scalability factor, the EVPN solution provides support for the IRB interfaces on EX9200 switches to facilitate optimal Layer 2 and Layer 3 forwarding along with virtual machine mobility. The IRB interfaces are configured on each configured VLAN including the default VLAN for an EVPN instance.

IRB is the ability to do Layer 2 switching and Layer 3 routing within a single node, thus avoiding extra hops for inter-subnet traffic. The EVPN IRB solution eliminates the default gateway problem using the gateway MAC and IP synchronization, and avoids the triangular routing problem with Layer 3 interworking by creating IP host routes for virtual machines (VMs) in the tenant VRFs.

Configuration

CLI Quick Configuration

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

```
CE1    set interfaces ge-1/1/7 vlan-tagging
       set interfaces ge-1/1/7 unit 0 vlan-id 10
       set interfaces ge-1/1/7 unit 0 family inet address 10.0.0.1/24
       set routing-options static route 198.51.100.0/24 next-hop 10.0.0.251

PE1    set interfaces ge-1/0/8 unit 0 family inet address 192.0.2.1/24
       set interfaces ge-1/0/8 unit 0 family mpls
       set interfaces ge-1/1/8 flexible-vlan-tagging
       set interfaces ge-1/1/8 encapsulation flexible-ethernet-services
       set interfaces ge-1/1/8 unit 0 encapsulation vlan-bridge
       set interfaces ge-1/1/8 unit 0 vlan-id 10
       set interfaces irb unit 0 family inet address 10.0.0.251/24
       set interfaces lo0 unit 0 family inet address 203.0.113.1/32
       set routing-options router-id 203.0.113.1
```

```
set routing-options autonomous-system 100
set routing-options forwarding-table chained-composite-next-hop ingress evpn
set protocols ldp interface all
set protocols ldp interface fxp0.0 disable
set protocols mpls interface all
set protocols mpls interface fxp0.0 disable
set protocols bgp group ibgp type internal
set protocols bgp group ibgp local-address 203.0.113.1
set protocols bgp group ibgp family evpn signaling
set protocols bgp group ibgp neighbor 203.0.113.2
set protocols ospf area 0.0.0.0 interface all
set protocols ospf area 0.0.0.0 interface fxp0.0 disable
set routing-instances evpna instance-type evpn
set routing-instances evpna vlan-id 10
set routing-instances evpna interface ge-1/1/8.0
set routing-instances evpna l3-interface irb.0
set routing-instances evpna route-distinguisher 203.0.113.1:100
set routing-instances evpna vrf-target target:100:100
set routing-instances evpna protocols evpn interface ge-1/1/8.0
set routing-instances vrf instance-type vrf
set routing-instances vrf interface irb.0
set routing-instances vrf route-distinguisher 203.0.113.1:300
set routing-instances vrf vrf-target target:100:300
set routing-instances vrf vrf-table-label
```

```
PE2 set interfaces ge-2/0/8 unit 0 family inet address 192.0.2.2/24
set interfaces ge-2/0/8 unit 0 family mpls
set interfaces ge-2/1/8 flexible-vlan-tagging
set interfaces ge-2/1/8 encapsulation flexible-ethernet-services
set interfaces ge-2/1/8 unit 0 encapsulation vlan-bridge
set interfaces ge-2/1/8 unit 0 vlan-id 20
set interfaces irb unit 0 family inet address 198.51.100.251/24
set interfaces lo0 unit 0 family inet address 203.0.113.2/32
set routing-options router-id 203.0.113.2
set routing-options autonomous-system 100
set routing-options forwarding-table chained-composite-next-hop ingress evpn
set protocols ldp interface all
set protocols ldp interface fxp0.0 disable
set protocols mpls interface all
set protocols mpls interface fxp0.0 disable
set protocols bgp group ibgp type internal
set protocols bgp group ibgp local-address 203.0.113.2
set protocols bgp group ibgp family evpn signaling
set protocols bgp group ibgp neighbor 203.0.113.1
set protocols ospf area 0.0.0.0 interface all
set protocols ospf area 0.0.0.0 interface fxp0.0 disable
set routing-instances evpna instance-type evpn
set routing-instances evpna vlan-id 20
set routing-instances evpna interface ge-2/1/8.0
set routing-instances evpna l3-interface irb.0
set routing-instances evpna route-distinguisher 203.0.113.2:100
set routing-instances evpna vrf-target target:200:100
set routing-instances evpna protocols evpn interface ge-2/1/8.0
set routing-instances vrf instance-type vrf
set routing-instances vrf interface irb.0
```

```

set routing-instances vrf route-distinguisher 203.0.113.2:300
set routing-instances vrf vrf-target target:200:300
set routing-instances vrf vrf-table-label

```

```

CE2  set interfaces ge-2/1/7 unit 0 vlan-id 20
      set interfaces ge-2/1/7 unit 0 family inet address 198.51.100.2/24
      set routing-options static route 10.0.0.0/24 next-hop 198.51.100.251

```

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

To configure Router PE1:



NOTE: Repeat this procedure for Router PE2, after modifying the appropriate interface names, addresses, and other parameters.

1. Configure Router PE1 interfaces.

```

[edit interfaces]
user@PE1# set ge-1/0/8 unit 0 family inet address 192.0.2.1/24
user@PE1# set ge-1/0/8 unit 0 family mpls

user@PE1# set ge-1/1/8 flexible-vlan-tagging
user@PE1# set ge-1/1/8 encapsulation flexible-ethernet-services
user@PE1# set ge-1/1/8 unit 0 encapsulation vlan-bridge
user@PE1# set ge-1/1/8 unit 0 vlan-id 10

user@PE1# set irb unit 0 family inet address 10.0.0.251/24

user@PE1# set lo0 unit 0 family inet address 203.0.113.1/32

```

2. Set the router ID and autonomous system number for Router PE1.

```

[edit routing-options]
user@PE1# set router-id 203.0.113.1
user@PE1# set autonomous-system 100

```

3. Configure the chained composite next hop for EVPN.

```

[edit routing-options]
user@PE1# set forwarding-table chained-composite-next-hop ingress evpn

```

4. Enable LDP on all interfaces of Router PE1, excluding the management interface.

```

[edit protocols]
user@PE1# set ldp interface all
user@PE1# set ldp interface fxp0.0 disable

```

5. Enable MPLS on all the interfaces of Router PE1, excluding the management interface.

```
[edit protocols]
user@PE1# set mpls interface all
user@PE1# set mpls mpls interface fxp0.0 disable
```

6. Configure the BGP group for Router PE1.

```
[edit protocols]
user@PE1# set bgp group ibgp type internal
```

7. Assign local and neighbor addresses to the ibgp BGP group for Router PE1 to peer with Router PE2.

```
[edit protocols]
user@PE1# set bgp group ibgp local-address 203.0.113.1
user@PE1# set bgp group ibgp neighbor 203.0.113.2
```

8. Include the EVPN signaling Network Layer Reachability Information (NLRI) to the ibgp BGP group.

```
[edit protocols]
user@PE1# set bgp group ibgp family evpn signaling
```

9. Configure OSPF on all the interfaces of Router PE1, excluding the management interface.

```
[edit protocols]
user@PE1# set ospf area 0.0.0.0 interface all
user@PE1# set ospf area 0.0.0.0 interface fxp0.0 disable
```

10. Configure the EVPN routing instance.

```
[edit routing-instances]
user@PE1# set evpna instance-type evpn
```

11. Set the VLAN identifier for the bridging domain in the evpna routing instance.

```
[edit routing-instances]
user@PE1# set evpna vlan-id 10
```

12. Configure the interface name for the evpna routing instance.

```
[edit routing-instances]
user@PE1# set evpna interface ge-1/1/8.0
```

13. Configure the IRB interface as the routing interface for the evpna routing instance.

```
[edit routing-instances]
user@PE1# set evpna l3-interface irb.0
```

14. Configure the route distinguisher for the evpna routing instance.

```
[edit routing-instances]
user@PE1# set evpna route-distinguisher 203.0.113.1:100
```

15. Configure the VPN routing and forwarding (VRF) target community for the evpna routing instance.

```
[edit routing-instances]
user@PE1# set evpna vrf-target target:100:100
```

16. Assign the interface name that connects the PE1 site to the VPN.

```
[edit routing-instances]
user@PE1# set evpna protocols evpn interface ge-1/1/8.0
```

17. Configure the VRF routing instance.

```
[edit routing-instances]
user@PE1# set vrf instance-type vrf
```

18. Configure the IRB interface as the routing interface for the vrf routing instance.

```
[edit routing-instances]
user@PE1# set vrf interface irb.0
```

19. Configure the route distinguisher for the vrf routing instance.

```
[edit routing-instances]
user@PE1# set vrf route-distinguisher 203.0.113.1:300
```

20. Configure the VRF label for the vrf routing instance.

```
[edit routing-instances]
user@PE1# set vrf vrf-table-label
```

Results

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

```
user@PE1# show interfaces
ge-1/0/8 {
  unit 0 {
    family inet {
      address 192.0.2.1/24;
    }
    family mpls;
  }
}
ge-1/1/8 {
  flexible-vlan-tagging;
```

```
encapsulation flexible-ethernet-services;
unit 0 {
    encapsulation vlan-bridge;
    vlan-id 10;
}
}
irb {
    unit 0 {
        family inet {
            address 10.0.0.251/24;
        }
    }
}
lo0 {
    unit 0 {
        family inet {
            address 203.0.113.1/32 {
            }
        }
    }
}
```

```
user@PE1# show routing-options
router-id 203.0.113.1;
autonomous-system 100;
forwarding-table {
    chained-composite-next-hop {
        ingress {
            evpn;
        }
    }
}
```

```
user@PE1# show protocols
ldp {
    interface all;
    interface fxp0.0 {
        disable;
    }
}
mpls {
    interface all;
    interface fxp0.0 {
        disable;
    }
}
bgp {
    group ibgp {
        type internal;
        local-address 203.0.113.1;
        family evpn {
            signaling;
        }
        neighbor 203.0.113.2;
    }
}
ospf {
```



```

    area 0.0.0.0 {
        interface all;
        interface fxp0.0 {
            disable;
        }
    }
}

user@PE1# show routing-instances
evpna {
    instance-type evpn;
    vlan-id 10;
    interface ge-1/1/8.0;
    l3-interface irb.0;
    route-distinguisher 203.0.113.1:100;
    vrf-target target:100:100;
    protocols {
        evpn {
            interface ge-1/1/8.0;
        }
    }
}
vrf {
    instance-type vrf;
    interface irb.0;
    route-distinguisher 203.0.113.1:300;
    vrf-target target:100:300;
    vrf-table-label;
}

```

Verification

Confirm that the configuration is working properly.

- [Verifying Local IRB MACs on page 33](#)
- [Verifying Remote IRB MACs on page 34](#)
- [Verifying Local IRB IPs on page 35](#)
- [Verifying Remote IRB IPs on page 36](#)
- [Verifying CE-CE Inter-Subnet Forwarding on page 37](#)

Verifying Local IRB MACs

Purpose Verify that the local IRB MACs are learned from L2ALD.

Action On Router PE1, determine the MAC address of the local IRB interface.

From operational mode, run the **show interfaces irb extensive | match "Current address"** command.

```

user@PE1> show interfaces irb extensive | match "Current address"
Current address: a8:d0:e5:54:0d:10, Hardware address: a8:d0:e5:54:0d:10

```

From operational mode, run the **show route table evpn.evpn.0 extensive | find "a8:d0:e5:54:0d:10"** command.

```
user@PE1> show route table evpn.evpn.0 extensive | find "a8:d0:e5:54:0d:10"
2:10.255.0.1:100::0::100::a8:d0:e5:54:0d:10/384 (1 entry, 1 announced)
TSI:
Page 0 idx 0, (group PE type Internal) Type 1 val 0x2736568 (adv_entry)
  Advertised metrics:
    Flags: Nexthop Change
    Nexthop: Self
    Localpref: 100
    AS path: [100] I
    Communities: target:100:100 evpn-default-gateway
Path 2:10.255.0.1:100::0::100::a8:d0:e5:54:0d:10 Vector len 4. Val: 0
  *EVPN Preference: 170
    Next hop type: Indirect
    Address: 0x26f8354
    Next-hop reference count: 6
    Protocol next hop: 10.255.0.1
    Indirect next hop: 0x0 - INH Session ID: 0x0
    State: <Active Int Ext>
    Age: 23:29:08
    Validation State: unverified
    Task: evpna-evpn
    Announcement bits (1): 1-BGP_RT_Background
    AS path: I
    Communities: evpn-default-gateway
    Route Label: 299776
```

Meaning The MAC-only route for the local IRB interface appears in the EVPN instance route table on Router PE1 and is learned from EVPN and tagged with the default gateway extended community.

Verifying Remote IRB MACs

Purpose Verify that the remote IRB MACs are learned from BGP.

Action On Router PE1, determine the MAC address of the local IRB interface.

From operational mode, run the **show interfaces irb extensive | match "Current address"** command.

```
user@PE1> show interfaces irb extensive | match "Current address"
Current address: a8:d0:e5:54:0d:10, Hardware address: a8:d0:e5:54:0d:10
```

On Router PE2, verify that the remote IRB MACs are learned.

From operational mode, run the **show route table evpn.evpn.0 extensive | find "a8:d0:e5:54:0d:10"** command.

```
user@PE2> show route table evpn.evpn.0 extensive | find "a8:d0:e5:54:0d:10"
```

```

2:10.255.0.1:100::0::100::a8:d0:e5:54:0d:10/384 (1 entry, 1 announced)
  *BGP Preference: 170/-101
    Route Distinguisher: 2.91.223.24:100
    Next hop type: Indirect
    Address: 0x26f8d6c
    Next-hop reference count: 10
    Source: 10.255.0.1
    Protocol next hop: 10.255.0.1
    Indirect next hop: 0x2 no-forward INH Session ID: 0x0
    State: <Secondary Active Int Ext>
    Local AS: 100 Peer AS: 100
    Age: 23:22:17 Metric2: 1
    Validation State: unverified
    Task: BGP_100.10.255.0.1
    Announcement bits (1): 0-evpna-evpn
    AS path: I
    Communities: target:100:100 evpn-default-gateway
    Import Accepted
    Route Label: 299776
    Localpref: 100
    Router ID: 10.255.0.1
    Primary Routing Table bgp.evpn.0
    Indirect next hops: 1
      Protocol next hop: 10.255.0.1 Metric: 1
      Indirect next hop: 0x2 no-forward INH Session ID: 0x0
      Indirect path forwarding next hops: 1
        Next hop type: Router
        Next hop: 1.0.0.1 via ge-1/0/8.0
        Session Id: 0x1
      10.255.0.1/32 Originating RIB: inet.3
        Metric: 1 Node path count: 1
        Forwarding nexthops: 1
        Nexthop: 1.0.0.1 via ge-1/0/8.0

```

Meaning The MAC-only route for the remote IRB interface appears in the EVPN instance route table on Router PE2 and is learned from BGP and tagged with the default gateway extended community.

Verifying Local IRB IPs

Purpose Verify that the local IRB IPs are learned locally by RPD.

Action On Router PE1, determine the MAC and IP addresses of the local IRB interface.

From operational mode, run the **show interfaces irb extensive | match "Current address"** command.

```

user@PE1> show interfaces irb extensive | match "Current address"
Current address: a8:d0:e5:54:0d:10, Hardware address: a8:d0:e5:54:0d:10

```

From operational mode, run the **show interfaces irb.0 terse | match inet** command.

```

user@PE1> show interfaces irb.0 terse | match inet
irb.0                up    up    inet    10.0.0.251/24

```

From operational mode, run the **show route table evpn.evpn.0 extensive | find "a8:d0:e5:54:0d:10::10.0.0.251"** command.

```
user@PE2> show route table evpn.evpn.0 extensive | find "a8:d0:e5:54:0d:10::10.0.0.251"
2:10.255.0.1:100::0::100::a8:d0:e5:54:0d:10::10.0.0.251/384 (1 entry, 1 announced)
TSI:
Page 0 idx 0, (group PE type Internal) Type 1 val 0x27365a0 (adv_entry)
  Advertised metrics:
    Flags: Nexthop Change
    Nexthop: Self
    Localpref: 100
    AS path: [100] I
    Communities: target:100:100 evpn-default-gateway
Path 2:10.255.0.1:100::0::100::a8:d0:e5:54:0d:10::10.0.0.251 Vector len 4. Val:
0
    *EVPN   Preference: 170 <<<<<
           Next hop type: Indirect
           Address: 0x26f8354
           Next-hop reference count: 6
           Protocol next hop: 10.255.0.1
           Indirect next hop: 0x0 - INH Session ID: 0x0
           State: <Active Int Ext>
           Age: 23:48:46
           Validation State: unverified
           Task: evpn-evpn
           Announcement bits (1): 1-BGP_RT_Background
           AS path: I
           Communities: evpn-default-gateway
           Route Label: 299776
```

Meaning The MAC plus IP route for the local IRB interface appears in the EVPN instance route table on Router PE1 and is learned from EVPN and tagged with the default gateway extended community.

Verifying Remote IRB IPs

Purpose Verify that the remote IRB IPs are learned from BGP.

Action On Router PE1, determine the MAC and IP addresses of the local IRB interface.

From operational mode, run the **show interfaces irb extensive | match "Current address"** command.

```
user@PE1> show interfaces irb extensive | match "Current address"
Current address: a8:d0:e5:54:0d:10, Hardware address: a8:d0:e5:54:0d:10
```

From operational mode, run the **show interfaces irb.0 terse | match inet** command.

```
user@PE1> show interfaces irb.0 terse | match inet
irb.0                up    up    inet    10.0.0.251/24
```

On Router PE2, verify that the remote IRB IPs are learnt.

From operational mode, run the **show route table evpn.evpn.0 extensive | find "a8:d0:e5:54:0d:10::10.0.0.251"** command.

```
user@PE2> show route table evpn.evpn.0 extensive | find "a8:d0:e5:54:0d:10::10.0.0.251"
2:10.255.0.1:100::0::100::a8:d0:e5:54:0d:10::10.0.0.251/384 (1 entry, 1 announced)
  *BGP Preference: 170/-101
    Route Distinguisher: 2.91.223.216:100
    Next hop type: Indirect
    Address: 0x26f8d6c
    Next-hop reference count: 10
    Source: 10.255.0.1
    Protocol next hop: 10.255.0.1
    Indirect next hop: 0x2 no-forward INH Session ID: 0x0
    State: <Secondary Active Int Ext>
    Local AS: 100 Peer AS: 100
    Age: 23:56:36 Metric2: 1
    Validation State: unverified
    Task: BGP_100.10.255.0.1
    Announcement bits (1): 0-evpn-evpn
    AS path: I
    Communities: target:100:100 evpn-default-gateway
    Import Accepted
    Route Label: 299776
    Localpref: 100
    Router ID: 10.255.0.1
    Primary Routing Table bgp.evpn.0
    Indirect next hops: 1
      Protocol next hop: 10.255.0.1 Metric: 1
      Indirect next hop: 0x2 no-forward INH Session ID: 0x0
      Indirect path forwarding next hops: 1
        Next hop type: Router
        Next hop: 1.0.0.1 via ge-1/0/8.0
        Session Id: 0x1
      10.255.0.1/32 Originating RIB: inet.3
        Metric: 1 Node path count: 1
        Forwarding nexthops: 1
        Nexthop: 1.0.0.1 via ge-1/0/8.0
```

Meaning The MAC plus IP route for the remote IRB interface appears in the EVPN instance route table on Router PE2 and is tagged with the default gateway extended community.

Verifying CE-CE Inter-Subnet Forwarding

Purpose Verify inter-subnet forwarding between Routers CE1 and CE2.

Action From operational mode, run the **show route table inet.0** command.

```
user@CE1> show route table inet.0
inet.0: 3 destinations, 3 routes (3 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

0.0.0.0/0          *[Static/5] 00:15:09
                   > to 10.0.0.251 via ge-1/1/7.0
10.0.0.0/24        *[Direct/0] 1d 23:24:30
                   > via ge-1/1/7.0
10.0.0.1/32        *[Local/0] 1d 23:24:38
                   Local via ge-1/1/7.0
```

From operational mode, run the **ping** command.

```
user@CE1> ping 198.51.100.2 interval 0.1 count 10
PING 198.51.100.2 (20.0.0.2): 56 data bytes
64 bytes from 198.51.100.2: icmp_seq=0 ttl=63 time=0.919 ms
64 bytes from 198.51.100.2: icmp_seq=1 ttl=63 time=0.727 ms
64 bytes from 198.51.100.2: icmp_seq=2 ttl=63 time=0.671 ms
64 bytes from 198.51.100.2: icmp_seq=3 ttl=63 time=0.671 ms
64 bytes from 198.51.100.2: icmp_seq=4 ttl=63 time=0.666 ms
64 bytes from 198.51.100.2: icmp_seq=5 ttl=63 time=0.704 ms
64 bytes from 198.51.100.2: icmp_seq=6 ttl=63 time=0.763 ms
64 bytes from 198.51.100.2: icmp_seq=7 ttl=63 time=0.750 ms
64 bytes from 198.51.100.2: icmp_seq=8 ttl=63 time=12.967 ms
64 bytes from 198.51.100.2: icmp_seq=9 ttl=63 time=0.752 ms

--- 198.51.100.2 ping statistics ---
10 packets transmitted, 10 packets received, 0% packet loss
round-trip min/avg/max/stddev = 0.666/1.959/12.967/3.670 ms
```

Meaning Ping from Router CE1 to Router CE2 is successful.

Related Documentation

- [An EVPN with IRB Solution on EX9200 Switches Overview on page 19](#)

CHAPTER 6

Configuring EVPN Multihoming

- [EVPN Multihoming Overview on page 39](#)
- [Configuring EVPN Active-Standby Multihoming on page 64](#)
- [Example: Configuring EVPN Active-Active Multihoming on page 67](#)
- [Example: Configuring an ESI on a Logical Interface With EVPN Multihoming on page 115](#)

EVPN Multihoming Overview

- [Introduction to EVPN Multihoming on page 39](#)
- [Understanding EVPN Multihoming Concepts on page 41](#)
- [EVPN Multihoming Mode of Operation on page 43](#)
- [EVPN Multihoming Implementation on page 44](#)
- [Designated Forwarder Election on page 56](#)
- [ESIs on Physical, Aggregated Ethernet, and Logical Interfaces on page 61](#)
- [Convergence in an EVPN Network on page 61](#)

Introduction to EVPN Multihoming

An Ethernet VPN (EVPN) comprises of customer edge (CE) devices that are connected to provider edge (PE) devices, which form the edge of the MPLS infrastructure. A CE device can be a host, a router, or a switch. The PE devices provide Layer 2 virtual bridge connectivity between the CE devices. There can be multiple EVPNs in the provider network. Learning between the PE routers occurs in the control plane using BGP, unlike traditional bridging, where learning occurs in the data plane.



.....

NOTE: In releases earlier than Junos OS Release 15.1, EVPN functionality support on MX Series routers was limited to routers using MPC and MIC interfaces only. Starting with Junos OS Release 15.1, MX Series routers using DPCs can be leveraged to provide EVPN support on the CE device-facing interface.

DPC support for EVPN is provided with the following considerations:

- DPCs provide support for EVPN in the active-standby mode of operation including support for the following:
 - EVPN instance (EVI)
 - Virtual switch
 - Integrated routing and bridging (IRB) interfaces
- DPCs intended for providing the EVPN active-standby support must be the CE device-facing line card. The PE device in the EVPN domain must be MPC interfaces or MIC interfaces.

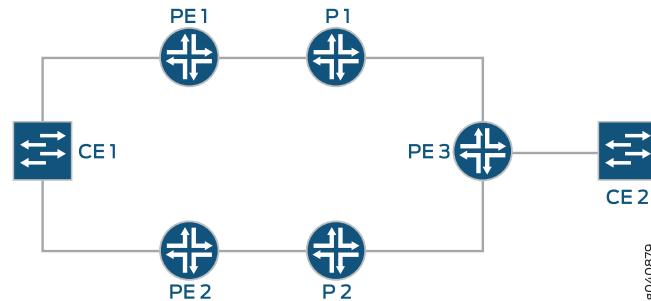
.....

The EVPN multihoming feature enables you to connect a customer site to two or more PE devices to provide redundant connectivity. A CE device can be multihomed to different PE devices or the same PE device. A redundant PE device can provide network service to the customer site as soon as a failure is detected. Thus, EVPN multihoming helps to maintain EVPN service and traffic forwarding to and from the multihomed site in the event of the following types of network failures:

- PE device to CE device link failure
- PE device failure
- MPLS-reachability failure between the local PE device and a remote PE device

Figure 2 on page 41 illustrates how a CE device can multihomed to two PE routers. Device CE 1 is multihomed to Routers PE 1 and PE 2. Device CE 2 has two potential paths to reach Device CE 1, and depending on the multihoming mode of redundancy, only one path or both the paths are active at any time. The multihoming mode of operation also determines which PE router or routers forward traffic to the CE device. The PE router forwarding traffic to the CE device (also called a designated forwarder) uses MPLS LSP or GRE tunnels to forward traffic. If a failure occurs over this path, a new designated forwarder is elected to forward the traffic to Device CE 1.

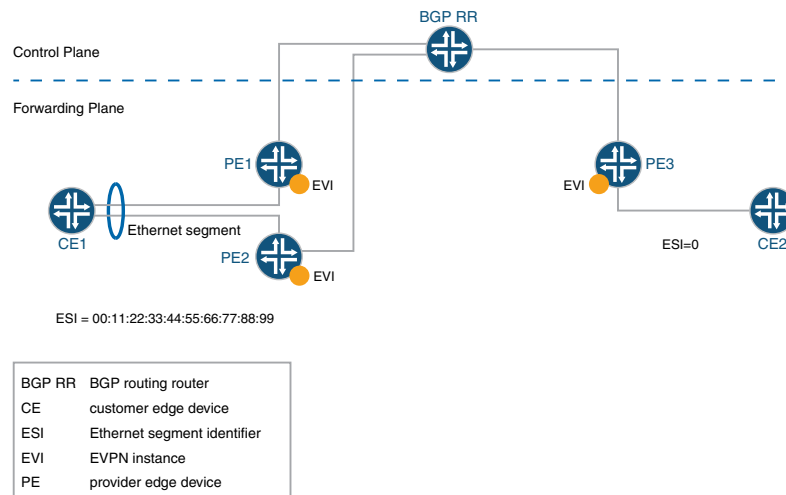
Figure 2: CE Device Multihomed to Two PE Routers



Understanding EVPN Multihoming Concepts

Figure 3 on page 41 shows a simple EVPN network topology to define EVPN multihoming concepts. .

Figure 3: Simple EVPN Topology



- **Ethernet segment**—When a CE device is multihomed to two or more PE routers, the set of Ethernet links constitutes an Ethernet segment. An Ethernet segment appears as a link aggregation group (LAG) to the CE device.

The links from Routers PE1 and PE2 to Device CE1 form an Ethernet segment.

In active-standby multihoming, the links that constitute an Ethernet segment form a bridge domain. In active-active multihoming, an Ethernet segment appears as a LAG to the CE device.

- **ESI**—An Ethernet segment must have a unique nonzero identifier, called the Ethernet segment identifier (ESI). The ESI is encoded as a 10-octet integer. When manually configuring an ESI value, the most significant octet, known as the type byte, must be 00. When a single-homed CE device is attached to an Ethernet segment, the entire ESI value is zero.

The Ethernet segment of the multihomed Device CE1 has an ESI value of 00:11:22:33:44:55:66:77:88:99 assigned. The single-homed Device CE2 has an ESI value of 0.

- **EVI**—An EVPN instance (EVI) is an EVPN routing and forwarding instance spanning all the PE routers participating in that VPN. An EVI is configured on the PE routers on a per-customer basis. Each EVI has a unique route distinguisher and one or more route targets.

An EVI is configured on Routers PE1, PE2, and PE3.

- **Ethernet tag**—An Ethernet tag identifies a particular broadcast domain, such as a VLAN. An EVPN instance consists of one or more broadcast domains. Ethernet tags are assigned to the broadcast domains of a given EVPN instance by the provider of that EVPN. Each PE router in that EVPN instance performs a mapping between broadcast domain identifiers understood by each of its attached CE devices and the corresponding Ethernet tag.
- **Ethernet segment route**—The PE routers that are connected to a multihomed CE device use BGP Ethernet segment route messages to discover that each of the PE routers is connected to the same Ethernet segment. The PE routers advertise the Ethernet segment route, which consists of an ESI and ES-import extended community.

Routers PE1 and PE2 advertise an ES route with an ES-import extended community (along with other extended communities like the route target). The PE routers also construct a filter that is based on an ES-import extended community, which results in only these PE routers importing the ES route and identifying that they are connected to the same Ethernet segment.

- **Extended community**— An extended community is similar in most ways to a regular community. EVPNs use extended communities because the 4-octet regular community value does not provide enough expansion and flexibility. An extended community is an 8-octet value divided into two main sections.
- **BUM traffic**—This type of traffic is sent to multiple destinations, including broadcast traffic, unknown unicast traffic that is broadcast in the Ethernet segment, and multicast traffic.
- **DF**—When a CE device is multihomed to two or more PE routers, either one or all of the multihomed PE routers are used to reach the customer site depending on the multihoming mode of operation. The PE router that assumes the primary role for forwarding BUM traffic to the CE device is called the designated forwarder (DF).
- **BDF**—Each router in the set of other PE routers advertising the autodiscovery route per Ethernet segment for the same ESI, and serving as the backup path in case the DF encounters a failure, is called a backup designated forwarder (BDF). A BDF is also called a non-DF router.
- **DF election**—On every Ethernet segment, the PE routers participate in a procedure called *designated forwarder* election to select the DF and the BDF PE routers.

EVPN Multihoming Mode of Operation

The different modes of operation for EVPN multihoming include:

- **Single**—When a PE router is connected to a single-homed customer site, this mode is in operation. The *single* mode is the default mode of operation, and does not require Ethernet segment values to be configured.
- **Active-standby**—When only a single PE router, among a group of PE routers attached to an Ethernet segment, is allowed to forward traffic to and from that Ethernet segment, the Ethernet segment is defined to be operating in the *active-standby* redundancy mode.

To configure the active-standby mode, include the ESI value and the **single-active** statement under the **[edit interfaces]** hierarchy level.

- **Active-active**—When all PE routers attached to an Ethernet segment are allowed to forward traffic to and from the Ethernet segment, the Ethernet segment is defined to be operating in the *active-active* redundancy mode.



NOTE: In Junos OS Release 14.2 and earlier, the EX9200 Series switch supports only the active-standby mode of operation for EVPN multihoming.



NOTE: Starting with Junos OS Release 14.1x53-D30, QFX5100 switches support the active-active mode of operation for EVPN multihoming. In this scenario, QFX5100 switches function as top-of-rack (ToR) switches in the data center for virtual networks. EVPN multihoming active-active functionality is used to provide access to the bare-metal servers connected to the top-of-rack switches.




NOTE: Starting with Junos OS Release 14.1R4, 14.2, 15.1F6, and 16.1R1, Junos OS supports the active-active mode for EVPN multihoming on MX Series routers.

Starting with Junos OS Releases 16.1R4 and 16.2R2, all EX9200 switches support the active-active mode for EVPN multihoming.

To configure the active-active mode, include the ESI value and the **all-active** statement at the **[edit interfaces]** hierarchy level.

Figure 4 on page 44 shows a reference topology for EVPN active-active multihoming. The ESI1 Ethernet segment for Device CE2 is multihomed to Routers PE1, PE2, and PE3. The Ethernet segment on the CE device can either be configured as a link aggregation group (LAG) or as an ECMP path. Devices CE1 and CE3 are the single-homed customer edge devices and have an ESI value of 0.

Figure 4: Active-Active EVPN Multihoming



EVPN Multihoming Implementation

The EVPN active-standby multihoming mode of operation provides redundancy for access link failures and PE node failure for the multihomed CE device, and is based on the EVPN *draft-ietf-l2vpn-evpn-03*.

The Junos OS implementation of the EVPN multihoming active-standby and active-active modes of operation includes the following:

- [New BGP NLRI](#)s on page 44
- [New Extended Communities](#) on page 47
- [New EVPN Route Types](#) on page 49
- [Update to the MAC Forwarding Table](#) on page 50
- [Traffic Flow](#) on page 51
- [Aliasing](#) on page 52
- [EVPN Active-Active Multihoming and Multichassis Link Aggregation](#) on page 54
- [EVPN Active-Active Multihoming and IRB](#) on page 54
- [Sample Configuration](#) on page 54

New BGP NLRI

To support EVPN multihoming, the following new BGP network layer reachability information (NLRI) routes have been introduced:

- [Autodiscovery Route per Ethernet Segment](#) on page 44
- [Ethernet Segment Route](#) on page 46
- [Autodiscovery Route per EVPN Instance](#) on page 47

Autodiscovery Route per Ethernet Segment

- [Autodiscovery Route Features](#) on page 45
- [Autodiscovery Route Advertisement](#) on page 45
- [Autodiscovery Route Withdrawal](#) on page 45

Autodiscovery Route Features

The autodiscovery route NLRI features include:

- This is a Type 1 mandatory route, used for fast convergence and for advertising the split horizon label. It is also known as the mass withdraw route.
- Type 1 route distinguishers are used with the IP address (loopback) of the originating PE router as the route distinguisher value.
- This route carries the ESI in the NLRI (nonzero when it is a multihomed PE, zero otherwise).
- The split horizon label is per ESI only, and carries an explicit NULL (0).
- The bit in the active-standby flag field in the ESI label extended community is used for signaling the active-standby mode (bit set).
- The 3-byte label values in the NLRI and the Ethernet tag is zero.
- This route is advertised and imported by all multihomed and remote PE routers that share the same EVI on the advertising ESI.

Autodiscovery Route Advertisement

- Active-standby mode

In active-standby mode, the designated forwarder (DF) advertises the autodiscovery route per Ethernet segment with an ESI MPLS label extended community that has the standby bit set to 1. The autodiscovery route is advertised per ESI, and the ESI label is set to 0 when active-standby mode is in operation.

The autodiscovery route is imported by all the multihomed and remote PE routers that are part of the EVI. On receiving the autodiscovery route, the PE routers in the network topology learn that active-standby multihoming mode is in operation for the ESI advertised.

- Active-active mode

In active-active mode, each of the multihomed PE device advertises a mandatory autodiscovery route per Ethernet segment as in the active-standby state. However, in the active-active state, the autodiscovery route per Ethernet segment is modified such that the active-standby bit carried in the MPLS extended community is cleared to indicate that the active-active mode is in operation. The autodiscovery route per Ethernet segment in the active-active mode also includes the split horizon label.

In [Figure 4 on page 44](#), for the ESI1 Ethernet segment, Routers PE1, PE2, and PE3 advertise the autodiscovery route. Router PE4 receives this autodiscovery route.

Autodiscovery Route Withdrawal

The autodiscovery route per Ethernet segment withdrawal is also referred to as the mass withdrawal. The mass withdrawal feature is used when there is a link failure on the ESI, or when the ESI configuration changes.

When the link between a multihomed CE device and a multihomed PE device fails, the PE device withdraws the autodiscovery route per Ethernet segment. In such a case, the mass withdrawal feature is handled in the following ways by the other PE devices:

- Remote PE device

When a remote PE device receives the BGP update for mass withdrawal, the following is performed at the remote PE device:

1. The current next hop to reach the remote ESI or CE device is deleted.
2. A new next hop through the remaining multihomed PE devices is created to reach the remote ESI or CE device.
3. All the MAC routes behind the CE device are updated with the newly created next hop.

- Other multihomed PE device

As a result of the mass withdrawal, load balancing on the multihomed CE device happens because of the following:

- When the other multihomed PE devices receive the same set of MAC addresses on the link to the concerned ESI.

In this case, the local routes are preferred. If the remote routes learned from the DF PE device gets withdrawn, it does not affect routes pointing to the local ESI.

- When the other multihomed PE devices have not received the same set of MAC addresses on the link to the concerned ESI.

In this case, the PE devices install the MAC routes pointing to the concerned ESI, although the MACs are remotely learned from the DF PE device. When the DF PE device withdraws these routes, the withdrawn routes are flushed. Packets that are destined to the flushed MAC addresses are flooded on all the local segments.

Ethernet Segment Route

- [Ethernet Segment Route Features on page 46](#)
- [Ethernet Segment Route Advertisement on page 47](#)

Ethernet Segment Route Features

The Ethernet segment route NRLI features include:

- This is a Type 4 route. The purpose of this route is to enable the PE routers connected to the same Ethernet segment to automatically discover each other with minimal configuration on exchanging this route.
- This route is associated with an ES-import extended community with an ESI value condensed to 6 bytes, similar to a route target.
- This route is advertised and imported only by PE routers that are multihomed on the advertising Ethernet segment.

Ethernet Segment Route Advertisement

The Ethernet segment route is exchanged among all the PE routers within a data center with the ES-import extended community. The ES-import extended community is constructed based on the ESI PE routers that are multihomed, and the Ethernet segment route carries the ESI value related to the Ethernet segment on which the PE routers are multihomed.

The Ethernet segment routes are filtered based on the ES-import extended community, such that only the PE routers that are multihomed on the same Ethernet segment import this route. Each PE router that is connected to a particular Ethernet segment constructs an import filtering rule to import a route that carries the ES-import extended community.

Autodiscovery Route per EVPN Instance

In active-active mode, each of the multihomed PE devices advertise an autodiscovery route per EVPN instance (EVI) with a valid MPLS label. This route is advertised per ESI and is imported by the remote PE devices. The MPLS label included in the autodiscovery route per EVI is used later for aliasing.

New Extended Communities

An extended community is similar in most ways to a regular community. Some networking implementations, such as virtual private networks (VPNs), use extended communities because the 4-octet regular community value does not provide enough expansion and flexibility. An extended community is an 8-octet value divided into two main sections.

To support active-standby multihoming, the following extended communities have been introduced:

- [ESI-Import on page 47](#)
- [Split Horizon on page 48](#)

ESI-Import

This extended community is attached to the ES route, and is populated from the ESI-import value extracted from the configured ESI value under the interface. To solve the problem of a conflict with another regular route target, the type is set to 0x06, which has been allocated by IANA.

The ESI-import extended community route target populates the list of import route targets configured for the special instance from where the ES route using this community is advertised.

Therefore, incoming ESI routes with the same ESI-import value in the extended community are imported by the PE routers, if the PE router is configured with an Ethernet segment that has the same ESI value. Once the PE router receives a set of these ESI routes that have the same ESI-import extended community value, the DF and BDF election can be done locally.



NOTE: When the ESI-import extended community is not created implicitly, a policy must be configured to attach all the route targets to the autodiscovery route per Ethernet segment.

Split Horizon

With reference to [Figure 4 on page 44](#) for example, when a CE device that is multihomed to two or more PE devices on an Ethernet segment (ESI1) and operating in the active-active redundancy mode sends a BUM packet to one of the non-DF PE devices (say PE1), then Device PE1 forwards that packet to all or a subset of the other PE devices in that EVPN instance, including the DF PE device for that Ethernet segment. In this case the DF PE device that the CE device is multihomed to drops the packet without forwarding it back to the CE device. This filtering is referred to as split horizon.

- **Split horizon signaling**

The split horizon extended community is attached to the autodiscovery route per Ethernet segment. The value of the extended community is the split horizon or the Poisson label itself, which is 3 bytes, and is advertised as an opaque attribute.

- **Split horizon advertisement**

- In active-standby mode, the standby bit in the split horizon extended community is set to 1, and the ESI split horizon label is set to 0.
- In the active-active mode, the split horizon extended community is modified to clear the standby bit to 0 and includes a valid ESI label used for split horizon purposes.

- **Split horizon MPLS routes**

The DF PE device advertises an autodiscovery route per Ethernet segment with a split horizon label A, and an inclusive multicast route with label B for BUM traffic forwarding. On the DF, the BUM packet from the core can come with following labels:

- When the non-DF PE devices receive a BUM packet on their single-homed ESIs, the BUM packet is sent to the DF PE device with multicast label B.
- When the non-DF PE devices receive a BUM packet on ESI1, the BUM packet is sent to the DF PE device with two MPLS labels — the multicast label B as the outer label, and the split horizon label A as the inner label.

In the EVPN multihoming scenario, the multicast label B has the S-bit set to 1 when it is the only label in the label stack. In this case, the BUM packet needs to be flooded on all the local ESIs on the DF PE device. But the label B has the S-bit set to 0 when split horizon label A is the innermost label in the label stack. In this case, the BUM packets need to be flooded on all local ESIs on the DF PE device, except the ESI that maps to the split horizon label A.

Assuming that packets originated from a multihomed CE device to a non-DF PE device on multihomed segment ESI1, when the non-DF PE device sends this packet to the DF PE device, the ESI label that the DF advertised to the non-DF PE device in its autodiscovery route per Ethernet segment is pushed first. The non-DF PE device also pushes the inclusive multicast label that the DF PE device advertised in its inclusive

multicast route and further pushes the LSP label. The MPLS header thus contains two labels within a 32-bit field.

The base EVPN functionality uses a table-next hop to stitch the MPLS table with its corresponding EVPN EVI table. In the EVPN EVI table, the mac-lookup is performed to switch the packet.

The following routes are programmed in the mpls.0 table for EVPN multicast:

- The (multicast-label, S=1) route points to the EVPN-EVI table-next hop.
- The (multicast-label, S=0) route points to the MPLS table-next hop. This route loops the packet back to the MPLS table after popping the multicast-label.
- The (split horizon-label) route points to the EVPN-EVI table-next hop. This is the same table-next hop that is used by the multicast-label, S=1 route.

New EVPN Route Types

EVPN multihoming mode supports the following EVPN route types:

- Autodiscovery route per Ethernet segment
- Autodiscovery route per EVPN instance (EVI)
- Ethernet segment route

These route types conform to the following naming convention:

<route-type>:<RD>::<esi>::<route-specific>/304

For example:

Autodiscovery route per Ethernet
segment—1:10.255.0.2:0::112233445566778899::0/304

Autodiscovery route per EVI—1:100.100.100.1:1::222222222222222222::0/304

Ethernet segment route—4:10.255.0.1:0::112233445566778899:10.255.0.1/304

where:

- **route-type**—Type of EVPN route.
 - 1—Autodiscovery route per Ethernet segment.
 - 1—Autodiscovery route per EVI.
 - 4—Ethernet segment route.
 - 5—Route with VXLAN/MPLS encapsulation

- **RD**—Route distinguisher value.

The route distinguisher value is set to the IP address of the PE router followed by 0.

- **esi**—Ethernet segment identifier. Displayed as 10 bytes of hexadecimal bytes, and leading 00 bytes are not displayed.

- **route-specific**—Differs per route type.
 - Autodiscovery route per Ethernet segment and autodiscovery route per EVI—This value is an MPLS label.



NOTE: The MPLS label is displayed in the extensive output, although it is not included in the prefix.

- Ethernet segment route—This value is the originating IP address.
- **304**—Maximum number of bits in an EVPN route. This is not very useful information and could be removed from the display. However, it might be useful in quickly identifying an EVPN route, either visually or with match operators.

Update to the MAC Forwarding Table

In active-standby EVPN multihoming, the MAC addresses are treated as routable addresses, and the MP-IBGP protocol is used to carry the customer MAC addresses. MAC learning at the PE routers does not occur in the data plane but in the control plane. This leads to more control applied in terms of the learning mechanism.

A PE router performs MAC learning in the data plane for packets coming from a customer network for a particular EVI. For CE MAC addresses that are behind other PE routers, the MAC addresses are advertised in BGP NLRI using a new MAC advertisement route type.

The MAC learning is of two types:

- Local MAC learning—PE routers must support the local MAC learning process through standard protocols.
- Remote MAC learning—Once the local learning process is completed, the PE routers can advertise the locally learned MAC address to remote PE router nodes through MP-IBGP. This process of receiving the remote MAC addresses of attached customers through MP-IBGP is known as the remote MAC learning process.

The MAC advertisement route type is used to advertise locally learned MAC addresses in BGP to remote PE routers. If an individual MAC address is advertised, the IP address field corresponds to that MAC address. If the PE router sees an ARP request for an IP address from a CE device, and if the PE router has the MAC address binding for that IP address, the PE router performs ARP proxy and responds to the ARP request.



NOTE: The ARP proxy is performed only for the gateway and not for the host.

The MPLS label field depends on the type of allocation. The PE router can advertise a single MPLS label for all MAC addresses per EVI, which requires the least number of MPLS labels and saves the PE router memory. However, when forwarding to the customer network, the PE router must perform a MAC lookup which can cause a delay and increase the number of CPU cycles.

Traffic Flow

In EVPN multihoming, traffic flow is performed in the forwarding-plane. Flood routes are created for flooding the packets, and are used in the following scenarios:

- When a packet is received on a local ESI
- When a packet is received from the core

The traffic flows in EVPN multihoming can be based on the two traffic types:

- Unicast traffic

Unicast traffic is a point-to-point communication with one sender and one receiver. In a multihomed EVPN, unicast traffic is forwarded as follows:

- In active-standby mode
 - CE to core—Traffic is learned and forwarded by the DF PE router.
 - Core to CE—The remote PE router learns the MAC addresses from the DF, and forwards all unicast traffic to the DF PE router.
- In active-active mode
 - CE to core—Traffic is load-balanced to all the connected multihomed PE devices.
 - Core to CE—Traffic from the remote PE devices is load-balanced to all the multihomed PE devices connected to the remote CE device.

- BUM traffic

Traffic that is sent to multiple destinations, including broadcast traffic, unknown unicast traffic that is broadcast in the Ethernet segment, and multicast traffic is known as BUM traffic. In a multihomed EVPN, BUM traffic is forwarded as follows:

- In active-standby mode
 - CE to core—The CE device floods any BUM traffic to all the links in the Ethernet segment. The DF PE router with the active path forwards the BUM packets to the core. The BDF PE router in the standby mode drops all the traffic from the CE device, because the EVPN multihomed status of the interface is in blocking state.
 - Core to CE—The remote PE routers flood all BUM traffic to both the DF and BDF PE routers. Only the DF forwards the BUM traffic to the CE device. The BDF PE router drops all the traffic, because the EVPN multihomed status of the interface is in blocking state.
- In active-active mode

Based on the requirements, flooding and switching among local ESIs can be enabled or disabled in the active-active mode. This is referred to as the no-local-switching behavior.

The core of EVPN service provides a full-mesh connectivity among the multihomed PE devices. Because of this, EVPN uses split horizon in the core, so a packet received

from the core is never switched or flooded back to the core. Instead, ingress replication is used to replicate the packets to the remote PE devices.

To flood packets to remote PE devices, the multicast and the split horizon next hops are used. The multicast next hop tunnels the packet with the inclusive multicast label, and the split horizon next hop tunnels the packet with a multicast-label and a split horizon label. One such next hop is required per multihomed ESI per remote PE device.

The following flood routes are used in the active-active mode:

- All-CE flood route

This flood route is used by the local ESIs for the following:

- Flooding the packet on the local ESIs (when local-switching is allowed).
- Flooding the packet to the remote PE devices. The remote PE devices flood the packet on their local ESIs.

Because BUM traffic is forwarded only by the Designated Forwarder (DF), and not by the non-DF multihomed PE devices, the non-DFs use the split horizon next hop to flood this packet to other PE devices. However, the multihomed local ESIs for which the PE device is a non-DF does not participate in the flooding.

The all-CE flood route is not used by the non-DF ESIs, and the next hop for these flood routes is created accordingly. In such cases, the non-DF ESI flood route is used.

- All-VE flood route

This flood route is used when the packet is received from the core. It is used for flooding the packet received from the core to the local ESIs. Because the packet received from the core can come with multicast-label only or with both multicast-label and split horizon label, appropriate forwarding rules must be followed to drop the packet on the multihomed ESI that maps to the split horizon label.

- Non-DF flood route

This flood route is used for the following:

- Flooding the packet on the local ESIs.
- Flooding the packet to the remote PE devices using ingress replication with SH-label for the DF for the ESI.

Aliasing

Starting in Junos OS Release 15.1, Junos OS supports aliasing in an EVPN. Aliasing is the ability of a remote PE device to load balance Layer 2 unicast traffic on all the other PE devices that have same Ethernet segment towards a CE device.

- [Aliasing in the Active-Active Mode on page 53](#)
- [Aliasing and Autodiscovery Routes on page 53](#)

- [Aliasing and Label Route on page 53](#)
- [Aliasing and Unicast Packet Forwarding on page 54](#)

Aliasing in the Active-Active Mode

In [Figure 4 on page 44](#), aliasing in the active-active mode works as follows:

1. ESI1 is configured on Routers PE1, PE2, and PE3. Routers PE1, PE2, and PE3 advertise the autodiscovery route per Ethernet segment for ESI1.
2. Device CE1 sends Layer 2 traffic with source MAC address (MAC1) to Router PE1.
3. Router PE1 learns the MAC1 address on (ESI1, vlan X) and advertises it to all PE routers using BGP.
4. Router PE4 receives the MAC1 route through BGP.
5. Because Router PE4 also received the autodiscovery route per EVI from Routers PE2 and PE3, it knows that MAC1 must be reachable through Routers PE2 and PE3. Router PE4 builds its forwarding state to load-balance the Layer 2 traffic for MAC1 among Routers PE1, PE2, and PE3.

Aliasing and Autodiscovery Routes

Autodiscovery routes from Routers PE2 and PE3 can come in any order. As a result, these routes are installed by the Layer 2 process as follows:

1. After receiving MAC1 from Router PE1, and if any autodiscovery routes have not been received by Router PE4, MAC1 is programmed by PE4 with a next hop pointing toward Router PE1. When PE4 receives the autodiscovery route from Router PE2 for the same ESI, the next hop is installed so the traffic for MAC1 is load-balanced to Routers PE1 and PE2. When PE4 receives the autodiscovery route from Router PE3 for the same ESI, the next hop is updated to load-balance the traffic for MAC1 among Routers PE1, PE2, and PE3.
2. If Router PE4 has already received the autodiscovery routes from more than one PE device (PE1, PE2, and PE3), PE4 installs the MAC routes with the multi-destination next hop.

Aliasing and Label Route

Any PE device that advertises the autodiscovery route per EVI with a valid MPLS label programs the advertised label in the mpls.0 routing table. For instance, if Router PE2 advertised the autodiscovery route per EVI with label A, the mpls.0 entry is as follows:

Label A route points to the EVPN-EVI table-next hop.

When the remote Router PE4 sends a unicast data packet toward Router PE2 with this label A, lookup is done in Router PE2's forwarding table, and as a result of this lookup, the packet is forwarded on ESI1.

Aliasing and Unicast Packet Forwarding

When the unicast packets for MAC1 come from the remote Router PE4 to Router PE2, there could be two cases:

- Router PE2 also received the same set of MACs on its link to ES11—In this case, local routes are preferred and as a result of the MAC lookup, packets are forwarded to ES11.
- Router PE2 has not received the same set of MACs on its link to ES11—In this case, Router PE2 still installs MAC routes pointing to ES11, although MACs are remotely learned from Router PE1. As a result, the packets are forwarded to ES11.

EVPN Active-Active Multihoming and Multichassis Link Aggregation

When a CE device is configured with a LAG toward the PE devices, the following two options are available to run LACP on the PE devices:

- Configure the same LACP system ID on all the PE devices.
- Configure multichassis link aggregation on the PE devices.

When multichassis link aggregation is configured with EVPN, a reduced set of procedures for active-active multichassis link aggregation are required. These procedures provide link and node level redundancy. The multichassis link aggregation is completely transparent to the CE device, and is realized as pure LAG. Multichassis link aggregation operates at the port level as well. This essentially means that if multichassis link aggregation is configured as active-active, all VLANs on the multichassis link aggregation ports work in the active-active multihoming mode.

When multichassis link aggregation is configured along with EVPN, the following is considered:

- Both multichassis link aggregation and EVPN ESI must be enabled to work in the active-active mode only.
- The following functions are not required for multichassis link aggregation with EVPN:
 - Mac synchronization—This is performed in the BGP control plane of EVPN.
 - ICL linking—This is handled by the aliasing feature of EVPN.
 - ARP synchronization—This is handled by the BGP control plane with IRB functionality.

EVPN Active-Active Multihoming and IRB

When IRB is configured, the EVPN routes contain both MAC and IP information. The active-active multihoming requires ARP synchronization among the multihomed PE devices because the ARP responses can get hashed to a particular PE device.

Sample Configuration

The following is a sample configuration for EVPN active-standby multihoming on the following types of interfaces:

- Ethernet interface configuration

```

ge-0/1/2 {
  encapsulation ethernet-bridge;
  esi XX:XX:XX:XX:XX:XX:XX:XX:XX:XX;
  unit 0 {
    family bridge;
  }
}

```

- Single VLAN interface configuration

```

ge-0/1/3 {
  encapsulation extended-vlan-bridge;
  esi XX:XX:XX:XX:XX:XX:XX:XX:XX:XX;
  vlan-tagging
  unit 0 {
    family bridge;
    vlan-id 1;
  }
}

```



NOTE:

- An ESI value of 0 and all FFs are reserved and are not used for configuring a multihomed Ethernet segment.
- Two interfaces in the same EVI cannot be configured with the same ESI value.

The following is a sample routing instance configuration for EVPN active-standby multihoming:

- Routing instance configuration

```

routing-instances {
  evpn-0 {
    instance-type evpn;
    route-distinguisher value;
    vrf-target value;
    vlan-id vlan-ID;
    interface ge-0/1/2.0;
    interface ge-1/1/1.0;
    interface ge-2/2/2.0;
    protocols {
      evpn {
        designated-forwarder-election hold-time time;
      }
    }
  }
}

```



NOTE: With the active-standby mode configuration, the autodiscovery route per Ethernet segment is advertised with the active-standby bit set to 1 for each Ethernet segment.

Designated Forwarder Election

The following sections discuss DF election:

- [DF Election Roles on page 56](#)
- [DF Election as Per RFC 7432 on page 56](#)
- [Preference-Based DF Election on page 58](#)
- [DF Election for Virtual Switch on page 60](#)
- [Handling Failover on page 60](#)

DF Election Roles

The designated forwarder (DF) election process involves selecting the designated forwarder (DF) PE router and the backup designated forwarder (BDF) or a non-DF (non-designated forwarder PE router roles.

- **DF**—The MAC address from the customer site is reachable only through the PE router announcing the associated MAC advertisement route. This PE router is the primary PE router that is selected to forward BUM traffic to the multihomed CE device, and is called the designated forwarder (DF) PE router.
- **BDF**—Each PE router in the set of other PE routers advertising the autodiscovery route per Ethernet segment for the same ESI, and serving as the backup path in case the DF encounters a failure, is called a backup designated forwarder (BDF) or a non-DF (non-designated forwarder) PE router.

As a result of the DF election process, if a local PE router is elected as the BDF, the multihomed interface connecting to the customer site is put into a blocking state for the active-standby mode. The interface remains in the blocking state until the PE router is elected as the DF for the Ethernet segment that the interface belongs to.

DF Election as Per RFC 7432

- [DF Election Procedure on page 56](#)
- [DF Election Trigger on page 57](#)

DF Election Procedure

The default procedure for DF election at the granularity of the ESI and EVI is referred to as service carving. With *service carving*, it is possible to elect multiple DFs per Ethernet segment (one per EVI) in order to perform load-balancing of multidestination traffic destined for a given Ethernet segment. The load-balancing procedures carve up the EVI space among the PE nodes evenly, in such a way that every PE is the DF for a disjoint set of EVIs.

The procedure for service carving is as follows:

1. When a PE router discovers the ESI of the attached Ethernet segment, it advertises an autodiscovery route per Ethernet segment with the associated ES-import extended community attribute.
2. The PE router then starts a timer (default value of 3 seconds) to allow the reception of the autodiscovery routes from other PE nodes connected to the same Ethernet segment. This timer value must be the same across all the PE routers connected to the same Ethernet segment.

The default wait timer can be overwritten using the **designated-forwarder-election hold-time** configuration statement.

3. When the timer expires, each PE router builds an ordered list of the IP addresses of all the PE nodes connected to the Ethernet segment (including itself), in increasing numeric order. Every PE router is then given an ordinal indicating its position in the ordered list, starting with 0 as the ordinal for the PE with the numerically lowest IP address. The ordinals are used to determine which PE node is the DF for a given EVI on the Ethernet segment.
4. The PE router that is elected as the DF for a given EVI unblocks traffic for the Ethernet tags associated with that EVI. The DF PE unblocks multideestination traffic in the egress direction toward the Ethernet segment. All the non-DF PE routers continue to drop multideestination traffic (for the associated EVIs) in the egress direction toward the Ethernet segment.

In [Figure 4 on page 44](#), the election of the DF for active-active multihoming is performed among Routers PE1, PE2, and PE3. As a result of this DF election, each one of these routers can become the DF for a particular VLAN from a range of VLANs configured on ESI1. The DF is responsible for forwarding BUM traffic on that ESI and VLAN for which it is elected as the DF. The non-DF PE routers block the BUM traffic on that particular Ethernet segment.

DF Election Trigger

In general, a DF election process is triggered in the following conditions:

- When an interface is newly configured with a nonzero ESI, or when the PE router transitions from an isolated-from-the-core (no BGP session) state to a connected-to-the-core (has established BGP session) state, a wait timer is imposed. By default, the interface is put into a blocking state until the PE router is elected as the DF.
- After completing a DF election process, a PE router receives a new Ethernet segment route or detects the withdrawal of an existing Ethernet segment route, without an imposed wait timer.

- When an interface of a non-DF PE router recovers from a link failure, the PE router has no knowledge of the wait time imposed by other PE routers. As a result, no wait timer is imposed for the recovered PE router to avoid traffic loss.

Preference-Based DF Election

The DF election based on RFC 7432 does not meet some of the operational requirements needed by some service providers. As a solution to this, starting with Junos OS Release 17.3, the DF election in a multihoming EVPN network can be controlled by using an administrative preference value for every EVPN instance (EVI) or instance service identifier (I-SID).

In the default DF election procedure (as specified in RFC 7432), the DF is elected randomly from one of the multihoming devices with modulo operation. With the preference-based DF election, the DF is elected manually using interface configuration options, such as the preference value, the Don't Preempt (DP) bit, and IP address.

- [Preference-Based DF Election Procedure on page 58](#)
- [DF Election Algorithm Mismatch on page 59](#)
- [DF Election Algorithm Migration on page 59](#)
- [Changing Preference For Maintenance on page 59](#)

Preference-Based DF Election Procedure

The preference-based DF election allows for manually electing a DF. This is useful when there is a need to choose the DF based on interface attributes, such as the bandwidth associated with an interface.

The preference-based DF election is executed as follows:

1. The DF election type and preference value are configured under an ESI. By default, the preference-based DF election type is MOD based.
2. The configured preference value and DP bit are advertised to the multihoming PE devices using the DF election extended community in the type 4 routes.
3. After receiving the type 4 route, the PE device builds the list of candidate DF devices, in the order of the preference value, DP bit, and IP address.
4. When the DF timer expires, the PE device selects the DF based on the highest preference value.

By default, the DF is elected based on highest preference per EVI. However, the preference-based DF election allows for electing the DF based on the lowest preference value when the **designated-forwarder-preference-least** statement is included under the `[edit routing-instances routing-instance-name protocols evpn]`.



NOTE: The designated-forwarder-preference-least configuration should be same on both the multihoming EVIs, otherwise there can be two DFs causing traffic loss or loop.

5. When the same preference value is configured, then the PE device selects the DF based on the DP bit. When the DP bit is also the same, the DF is elected based on the lowest IP address.

DF Election Algorithm Mismatch

When there is a mismatch between a locally configured DF election algorithm and a remote PE device's DF election algorithm, then all the PE devices should fall back to the default DF election as specified in the RFC 7432.

DF Election Algorithm Migration

During the migration of the old DF election to the new DF election, it is expected to change the configuration during the maintenance window by bringing down the ESI, and changing the DF election algorithm.

To do the migration, do the following:

1. After a software upgrade, on the non-DF device bring down all the interfaces which have the same ESI.
2. Configure the new DF election algorithm on the DF PE.
3. Configure the DF election algorithm on other multihoming PE devices.
4. Bring up all the interfaces on the non-DF PE devices.

Changing Preference For Maintenance

After migrating the DF election algorithm, and all the multihoming PE device are running the preference-based DF election algorithm, maintenance tasks required on the existing DF can be executed by simply changing the configured preference value. This changes the DF for a given ESI.

To change the DF for a given ESI:

- Change the preference value to a higher value on the current non-DF device.
- Change the preference value to a lower value on the current DF device.



NOTE: Changing the preference value for an ESI can lead to some traffic loss for short duration to consider the delay in the updated BGP route propagation with the new preference value.

DF Election for Virtual Switch

The virtual switch allows multiple bridge domains in a single EVPN instance (EVI). The virtual switch also supports trunk and access ports. Junos OS allows flexible Ethernet services on the port, therefore different VLANs on a single port can be part of different EVIs.

The DF election for virtual switch depends on the following:

- Port mode—Sub-interface, trunk interface, and access port
- EVI mode—Virtual switch with EVPN and EVPN-EVI

In the virtual switch, multiple Ethernet tags can be associated with a single EVI, wherein the numerically lowest Ethernet tag value in the EVI is used for the DF election.

Handling Failover

A failover can occur when:

- The DF PE router loses its DF role.
- There is a link or port failure on the DF PE router.

On losing the DF role, the customer-facing interface on the DF PE router is put in the blocking state.

In the case of link or port failure, a DF election process is triggered, resulting in the BDF PE router to be selected as the DF. At that time, unicast traffic and BUM flow of traffic are affected as follows:

- [Unicast Traffic on page 60](#)
- [BUM Traffic on page 60](#)

Unicast Traffic

- CE to Core—The CE device continues to flood traffic on all the links. The previous BDF PE router changes the EVPN multihomed status of the interface from the blocking state to the forwarding state, and traffic is learned and forwarded through this PE router.
- Core to CE—The failed DF PE router withdraws the autodiscovery route per Ethernet segment and the locally-learned MAC routes, causing the remote PE routers to redirect traffic to the BDF.



NOTE: The transition of the BDF PE router to the DF role can take some time, causing the EVPN multihomed status of the interface to continue to be in the blocking state, resulting in traffic loss.

BUM Traffic

- CE to Core—All the traffic is routed toward the BDF.

- Core to CE—The remote PE routers flood the BUM traffic in the core.

ESIs on Physical, Aggregated Ethernet, and Logical Interfaces

In releases before Junos OS Release 15.1F6 and 17.1R1 for MX Series routers and Junos OS Release 17.3R1 for EX9200 switches, you can specify an ESI only on a physical or aggregated Ethernet interface, for example, **set interfaces ae0 esi**

00:11:22:33:44:55:66:77:88:99. If you specify an ESI on a physical or aggregated Ethernet interface, keep in mind that an ESI is a factor in the designated forwarder (DF) election process. For example, assume that you configure EVPN multihoming active-standby on aggregated Ethernet interface ae0, and given the ESI configured on ae0 and other determining factors, the DF election results in ae0 being in the down state. Further, all logical interfaces configured on aggregated Ethernet interface ae0, for example, **set interfaces ae0 unit 1** and **set interfaces ae0 unit 2** are also in the down state, which renders logical interfaces ae0.1 and ae0.2 unable to provide services to their respective customer sites (VLANs).

To better utilize logical interfaces in EVPN multihoming active-standby or active-active mode, starting with Junos OS Releases 15.1F6 and 17.1R1 for MX Series routers and Junos OS Release 17.3R1 for EX9200 switches, you can specify an ESI on a logical interface.. As a result, even if a logical interface is a non-DF, other logical interfaces on the same physical or aggregated Ethernet interface are still able to provide services to their respective customer sites (VLANs).

For more information, see [“Example: Configuring an ESI on a Logical Interface With EVPN Multihoming” on page 115](#).

Convergence in an EVPN Network

When there are changes in the network topology in a large-scale EVPN system, the convergence time might be significant. You can prioritize NLRI updates that are critical to route selection in routing policies to improve convergence. [Table 3 on page 61](#) lists the NLRI route types and the priority that must be configured in the routing policy.

Table 3: Priority for NLRI Route Type

NLRI Route Type	Description	Priority
NLRI Route Type 1	Ethernet auto-discovery route—Type 1 supports fast convergence and aliasing and is used to signal MAC mass withdrawal.	High
NLRI Route Type 2	MAC/IP advertisement route—Type 2 is used to advertise MAC addresses and IP addresses in EVPN networks.	Low
NLRI Route Type 3	Inclusive multicast Ethernet tag—Type 3 is used to set up a path for BUM traffic.	Low
NLRI Route Type 4	Ethernet segment route—Type 4 is used in the selection of a designated forwarder.	High

To prioritize the NLRI route type, set the **bgp-output-queue-priority priority** for **nlri-route-type** at the **[edit policy-options policy-statement]** hierarchy level on all provider

edge routers and route reflectors in the EVPN network. In this example, a high priority was configured for NRLI route type 1 and NRLI route type 4.

```
user@PE1#show policy-options
policy-statement evpn-rt-priority-policy {
  term 1 {
    from {
      family evpn;
      nlri-route-type 1;
    }
    then {
      bgp-output-queue-priority priority 16;
    }
  }
  term 2 {
    from {
      family evpn;
      nlri-route-type 2;
    }
    then {
      bgp-output-queue-priority priority 1;
    }
  }
  term 3 {
    from {
      family evpn;
      nlri-route-type 3;
    }
    then {
      bgp-output-queue-priority priority 2;
    }
  }
  term 4 {
    from {
      family evpn;
      nlri-route-type 4;
    }
    then {
      bgp-output-queue-priority priority 16;
    }
  }
}
```



NOTE: There are 17 prioritized output queues: an expedited queue that has the highest priority, and 16 numbered queues for which 1 is the lowest priority and 16 is the highest.

For more information about how to configure routing policies, see *Routing Policies, Firewall Filters, and Traffic Policers Feature Guide*.

Release History Table

Release	Description
16.1R4	Starting with Junos OS Releases 16.1R4 and 16.2R2, all EX9200 switches support the active-active mode for EVPN multihoming.
15.1F6	To better utilize logical interfaces in EVPN multihoming active-standby or active-active mode, starting with Junos OS Releases 15.1F6 and 17.1R1 for MX Series routers and Junos OS Release 17.3R1 for EX9200 switches, you can specify an ESI on a logical interface.. As a result, even if a logical interface is a non-DF, other logical interfaces on the same physical or aggregated Ethernet interface are still able to provide services to their respective customer sites (VLANs).
15.1	Starting in Junos OS Release 15.1, Junos OS supports aliasing in an EVPN.
14.1x53-D30	Starting with Junos OS Release 14.1x53-D30, QFX5100 switches support the active-active mode of operation for EVPN multihoming.
14.1R4	Starting with Junos OS Release 14.1R4, 14.2, 15.1F6, and 16.1R1, Junos OS supports the active-active mode for EVPN multihoming on MX Series routers.

**Related
Documentation**

- [Example: Configuring EVPN Active-Standby Multihoming](#)
- [Example: Configuring EVPN Active-Active Multihoming on page 67](#)

Configuring EVPN Active-Standby Multihoming

You can configure an Ethernet VPN (EVPN) with multihoming support to provide multihoming functionality with active-standby redundancy mode of operation in the EVPN and virtual switch routing instance. This mode enables autodiscovery of Ethernet segments, Ethernet segment route construction, and Ethernet segment identifier (ESI) label assignment.

When configuring active-standby EVPN multihoming, be aware of the following limitations:

- An interface or ESI can be attached to more than one EVPN instance (EVI), with a maximum limit of 200 EVIs per ESI.
- For an EVPN routing instance, only one logical interface per physical interface or ESI can be attached to an EVI.
- For a virtual switch routing instance, only one logical interface per physical interface or ESI can be configured under a bridge domain.
- All the PE routers in the network topology should be running Junos OS Release 14.1 or later releases, which are based on the EVPN draft-ietf-l2vpn-evpn-03. Junos OS releases prior to 14.1 support the older version of the EVPN draft, causing interoperability issues when Junos OS Release 14.1 and a previous release are running.

Before you begin:

1. Configure the router interfaces.
2. Configure the router ID and autonomous system number for the device.
3. Configure OSPF or any other IGP protocol.
4. Configure a BGP internal group.
5. Include the EVPN signaling network layer reachability information (NLRI) to the internal BGP group.
6. Configure LDP.
7. Configure MPLS.
8. Configure RSVP MPLS LSP or GRE tunnels.

To configure the PE device:

1. Enable EVPN active-standby multihoming on the multihomed interfaces.

```
[edit interfaces]
user@PE1# set interface-name vlan-tagging
user@PE1# set interface-name encapsulation flexible-ethernet-services
user@PE1# set interface-name esi esi-value
user@PE1# set interface-name esi single-active
user@PE1# set interface-name unit 0 encapsulation vlan-bridge
user@PE1# set interface-name unit 0 vlan-id VLAN-ID
```

For example:


```
[edit interfaces]
user@PE1# set ge-0/0/4 vlan-tagging
user@PE1# set ge-0/0/4 encapsulation flexible-ethernet-services
user@PE1# set ge-0/0/4 esi 00:22:44:66:88:00:22:44:66:88
user@PE1# set ge-0/0/4 esi single-active
user@PE1# set ge-0/0/4 unit 0 encapsulation vlan-bridge
user@PE1# set ge-0/0/4 unit 0 vlan-id 300
```

2. In configuration mode, go to the following hierarchy level:

```
[edit]
user@PE1# edit routing-instances
```

3. Configure the virtual switch routing instance.

```
[edit routing-instances]
user@PE1# set virtual-switch-instance instance-type virtual-switch
```

4. Configure the extended VLAN list for the virtual switch routing instance.

```
[edit routing-instances]
user@PE1# set virtual-switch-instance protocols evpn extended-vlan-list VLAN-ID
```

5. Set the type for the bridging domain in the virtual switch routing instance.

```
[edit routing-instances]
user@PE1# set virtual-switch-instance bridge-domains bridge-domain-name
domain-type bridge
```

6. Set the VLAN identifier for the bridging domain in the virtual switch routing instance.

```
[edit routing-instances]
user@PE1# set virtual-switch-instance bridge-domains bridge-domain-name vlan-id
VLAN-ID
```

7. Configure the interfaces for the virtual switch routing instance.

```
[edit routing-instances]
user@PE1# set virtual-switch-instance bridge-domains bridge-domain-name interface
interface-name
user@PE1# set virtual-switch-instance bridge-domains bridge-domain-name
routing-interface interface-name
```

8. Configure the route distinguisher for the virtual switch routing instance.

```
[edit routing-instances]
user@PE1# set virtual-switch-instance route-distinguisher route-distinguisher-value
```

9. Configure the VPN routing and forwarding (VRF) target community for the virtual switch routing instance.

```
[edit routing-instances]
user@PE1# set virtual-switch-instance vrf-target vrf-target
```

10. Configure the EVPN routing instance.

```
[edit routing-instances]
user@PE1# set evpn-instance instance-type evpn
```

11. Set the VLAN identifier for the bridging domain in the EVPN routing instance.

```
[edit routing-instances]
user@PE1# set evpn-instance vlan-id VLAN-ID
```

12. Configure the interface names for the EVPN routing instance.

```
[edit routing-instances]
user@PE1# set evpn-instance interface interface-name
user@PE1# set evpn-instance routing-interface interface-name
```

13. Configure the route distinguisher for the EVPN routing instance.

```
[edit routing-instances]
user@PE1# set evpn-instance route-distinguisher route-distinguisher-value
```

14. Configure the VPN routing and forwarding (VRF) target community for the EVPN routing instance.

```
[edit routing-instances]
user@PE1# set evpn-instance vrf-target vrf-target
```

15. Configure the VRF routing instance.

```
[edit routing-instances]
user@PE1# set vrf-instance instance-type vrf
```

16. Configure the interface names for the VRF routing instance.

```
[edit routing-instances]
user@PE1# set vrf-instance interface interface-name
```

17. Configure the route distinguisher for the VRF routing instance.

```
[edit routing-instances]
user@PE1# set vrf-instance route-distinguisher route-distinguisher-value
```

18. Configure the VPN routing and forwarding (VRF) target community for the VRF routing instance.

```
[edit routing-instances]
user@PE1# set vrf-instance vrf-target vrf-target
user@PE1# set vrf-instance vrf-table-label
```

19. Verify and commit the configuration.

For example:

```
[edit routing-instances]
```

```

user@PE1# set ALPHA instance-type virtual-switch
user@PE1# set ALPHA route-distinguisher 10.255.0.1:100
user@PE1# set ALPHA vrf-target target:100:100
user@PE1# set ALPHA protocols evpn extended-vlan-list 100
user@PE1# set ALPHA bridge-domains ONE domain-type bridge
user@PE1# set ALPHA bridge-domains ONE vlan-id 100
user@PE1# set ALPHA bridge-domains ONE interface ae0.0
user@PE1# set ALPHA bridge-domains ONE interface ge-0/0/2.0
user@PE1# set ALPHA bridge-domains ONE routing-interface irb.0
user@PE1# set BETA instance-type evpn
user@PE1# set BETA vlan-id 300
user@PE1# set BETA interface ge-0/0/4.0
user@PE1# set BETA interface ae1.0
user@PE1# set BETA routing-interface irb.1
user@PE1# set BETA route-distinguisher 10.255.0.1:300
user@PE1# set BETA vrf-target target:300:300
user@PE1# set DELTA instance-type vrf
user@PE1# set DELTA interface irb.0
user@PE1# set DELTA interface irb.1
user@PE1# set DELTA route-distinguisher 10.255.0.1:200
user@PE1# set DELTA vrf-target target:200:200
user@PE1# set DELTA vrf-table-label

[edit]
user@PE1# commit
commit complete

```

Related Documentation

- [Example: Configuring EVPN Active-Standby Multihoming](#)

Example: Configuring EVPN Active-Active Multihoming

This example shows how to configure Ethernet VPN (EVPN) for multihomed customer edge devices in the active-active redundancy mode, so the Layer 2 unicast traffic can be load-balanced across all the multihomed links on and toward the CE device.

- [Requirements on page 67](#)
- [Overview and Topology on page 68](#)
- [Configuration on page 69](#)
- [Verification on page 98](#)

Requirements

This example uses the following hardware and software components:

- Five MX Series 3D Universal Edge Routers with MPC interfaces only, where:
 - Three devices are configured as provider edge (PE) routers connected to a common multihomed customer site.
 - One device is configured as a remote PE router connected to a single-homed customer site.

- Six customer edge (CE) devices, with one multihomed CE device and the rest of the CE devices being single-homed to each of the PE routers.
- Junos OS Release 16.1 or later running on all the PE routers.



NOTE: The EVPN multihoming active-active mode of operation is supported in Junos OS Releases 16.1 and later releases.

Starting with Junos OS Release 16.1R4, EVPN multihoming active-active mode is supported on all EX9200 switches. For information about configuration specific to EX9200 switches, see [“Configuration on EX9200 Switches” on page 90](#).

Before you begin:

1. Configure the router interfaces.
2. Configure IS-IS or any other IGP protocol.
3. Configure BGP.
4. Configure LDP.
5. Configure MPLS.
6. Configure RSVP MPLS LSP or GRE tunnels.

Overview and Topology

Starting with Junos OS Release 15.1, the EVPN solution on MX Series routers with MPC interfaces is extended to provide multihoming functionality in the active-active redundancy mode of operation. This feature enables load balancing of Layer 2 unicast traffic across all the multihomed links on and toward a customer edge device.

The EVPN active-active multihoming feature provides link-level and node-level redundancy along with effective utilization of resources.

To enable EVPN active-active multihoming, include the **all-active** statement at the **[edit interfaces esi]** hierarchy level.

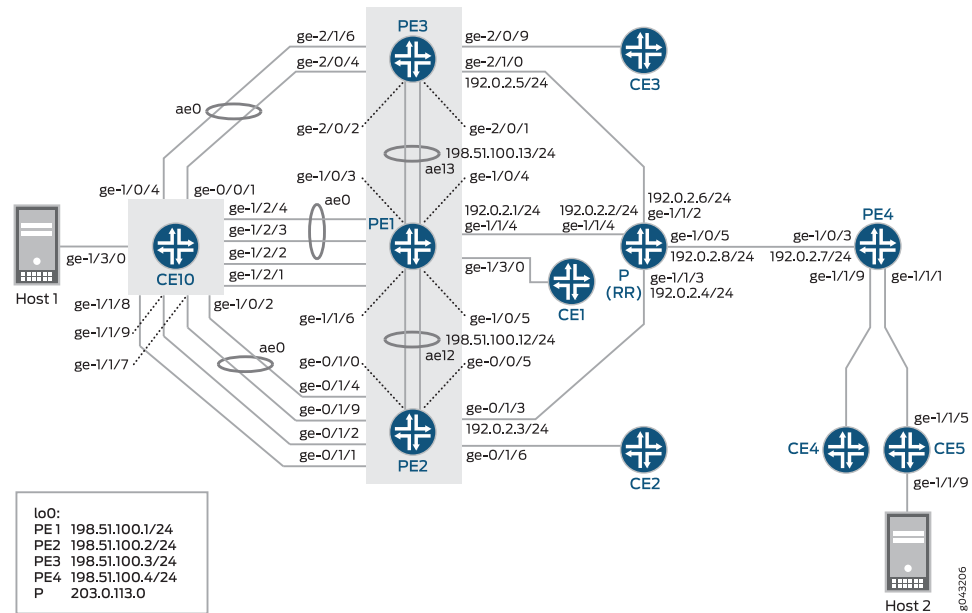
In [Figure 5 on page 69](#), the core consists of four provider edge (PE) routers and a provider router (P) that is configured as a route reflector (RR). Router CE10 is multihomed to Routers PE1, PE2, and PE3. Each PE router is also connected to a single-homed customer site.

There are three routing instances running in the topology – VS-1, VS-2, and mhevpn, along with the default routing instance. The routing instances share nine VLANs with three ESIs each. The VS-1 and VS-2 routing instances are configured as a virtual switch type of routing instance, and the mhevpn routing instance is an EVPN routing instance.

Three aggregated bundles – ae0, ae1, and ae2 – are used to connect the multihomed CE device, CE10, to Routers PE1, PE2, and PE3. These aggregated bundles are configured for

three ESIs each. The aggregated bundles, ae12 and ae13, are used to interconnect Routers PE1 and PE2, and PE1 and PE3, respectively.

Figure 5: EVPN Active-Active Multihoming Topology



Configuration

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

```
CE10 set chassis aggregated-devices ethernet device-count 20
set chassis network-services enhanced-ip
set interfaces ge-1/0/0 gigether-options 802.3ad ae0
set interfaces ge-1/0/2 gigether-options 802.3ad ae0
set interfaces ge-1/0/4 gigether-options 802.3ad ae0
set interfaces ge-1/1/7 gigether-options 802.3ad ae0
set interfaces ge-1/1/8 gigether-options 802.3ad ae2
set interfaces ge-1/1/9 gigether-options 802.3ad ae1
set interfaces ge-1/2/1 gigether-options 802.3ad ae2
set interfaces ge-1/2/2 gigether-options 802.3ad ae1
set interfaces ge-1/2/3 gigether-options 802.3ad ae0
set interfaces ge-1/2/4 gigether-options 802.3ad ae0
set interfaces ge-1/3/0 flexible-vlan-tagging
set interfaces ge-1/3/0 encapsulation flexible-ethernet-services
set interfaces ge-1/3/0 unit 0 encapsulation vlan-bridge
set interfaces ge-1/3/0 unit 0 vlan-id 10
set interfaces ge-1/3/0 unit 1 encapsulation vlan-bridge
set interfaces ge-1/3/0 unit 1 vlan-id 20
set interfaces ge-1/3/0 unit 2 encapsulation vlan-bridge
set interfaces ge-1/3/0 unit 2 vlan-id 30
```

```
set interfaces ge-1/3/0 unit 110 encapsulation vlan-bridge
set interfaces ge-1/3/0 unit 110 vlan-id 110
set interfaces ge-1/3/0 unit 120 encapsulation vlan-bridge
set interfaces ge-1/3/0 unit 120 vlan-id 120
set interfaces ge-1/3/0 unit 130 encapsulation vlan-bridge
set interfaces ge-1/3/0 unit 130 vlan-id 130
set interfaces ge-1/3/0 unit 210 encapsulation vlan-bridge
set interfaces ge-1/3/0 unit 210 vlan-id 210
set interfaces ge-1/3/0 unit 220 encapsulation vlan-bridge
set interfaces ge-1/3/0 unit 220 vlan-id 220
set interfaces ge-1/3/0 unit 230 encapsulation vlan-bridge
set interfaces ge-1/3/0 unit 230 vlan-id 230
set interfaces ae0 flexible-vlan-tagging
set interfaces ae0 encapsulation flexible-ethernet-services
set interfaces ae0 aggregated-ether-options minimum-links 1
set interfaces ae0 unit 0 encapsulation vlan-bridge
set interfaces ae0 unit 0 vlan-id 10
set interfaces ae0 unit 110 encapsulation vlan-bridge
set interfaces ae0 unit 110 vlan-id 110
set interfaces ae0 unit 210 encapsulation vlan-bridge
set interfaces ae0 unit 210 vlan-id 210
set interfaces ae1 flexible-vlan-tagging
set interfaces ae1 encapsulation flexible-ethernet-services
set interfaces ae1 aggregated-ether-options minimum-links 1
set interfaces ae1 unit 1 encapsulation vlan-bridge
set interfaces ae1 unit 1 vlan-id 20
set interfaces ae1 unit 120 encapsulation vlan-bridge
set interfaces ae1 unit 120 vlan-id 120
set interfaces ae1 unit 220 encapsulation vlan-bridge
set interfaces ae1 unit 220 vlan-id 220
set interfaces ae2 flexible-vlan-tagging
set interfaces ae2 encapsulation flexible-ethernet-services
set interfaces ae2 aggregated-ether-options minimum-links 1
set interfaces ae2 unit 2 encapsulation vlan-bridge
set interfaces ae2 unit 2 vlan-id 30
set interfaces ae2 unit 130 encapsulation vlan-bridge
set interfaces ae2 unit 130 vlan-id 130
set interfaces ae2 unit 230 encapsulation vlan-bridge
set interfaces ae2 unit 230 vlan-id 230
set routing-options forwarding-table export load-balancing-policy
set protocols l2-learning global-mac-table-aging-time 18000
set policy-options policy-statement load-balancing-policy then load-balance per-packet
set bridge-domains bd10 domain-type bridge
set bridge-domains bd10 vlan-id 10
set bridge-domains bd10 interface ae0.0
set bridge-domains bd10 interface ge-1/3/0.0
set bridge-domains bd110 domain-type bridge
set bridge-domains bd110 vlan-id 110
set bridge-domains bd110 interface ae0.110
set bridge-domains bd110 interface ge-1/3/0.110
set bridge-domains bd120 domain-type bridge
set bridge-domains bd120 vlan-id 120
set bridge-domains bd120 interface ge-1/3/0.120
set bridge-domains bd120 interface ae1.120
set bridge-domains bd130 domain-type bridge
set bridge-domains bd130 vlan-id 130
```

```

set bridge-domains bd130 interface ge-1/3/0.130
set bridge-domains bd130 interface ae2.130
set bridge-domains bd20 domain-type bridge
set bridge-domains bd20 vlan-id 20
set bridge-domains bd20 interface ge-1/3/0.1
set bridge-domains bd20 interface ae1.1
set bridge-domains bd210 domain-type bridge
set bridge-domains bd210 vlan-id 210
set bridge-domains bd210 interface ae0.210
set bridge-domains bd210 interface ge-1/3/0.210
set bridge-domains bd220 domain-type bridge
set bridge-domains bd220 vlan-id 220
set bridge-domains bd220 interface ge-1/3/0.220
set bridge-domains bd220 interface ae1.220
set bridge-domains bd230 domain-type bridge
set bridge-domains bd230 vlan-id 230
set bridge-domains bd230 interface ge-1/3/0.230
set bridge-domains bd230 interface ae2.230
set bridge-domains bd30 domain-type bridge
set bridge-domains bd30 vlan-id 30
set bridge-domains bd30 interface ge-1/3/0.2
set bridge-domains bd30 interface ae2.2

```

```

CE5  set chassis aggregated-devices ethernet device-count 20
      set chassis network-services enhanced-ip
      set interfaces ge-1/1/5 gigether-options 802.3ad ae0
      set interfaces ge-1/1/9 flexible-vlan-tagging
      set interfaces ge-1/1/9 encapsulation flexible-ethernet-services
      set interfaces ge-1/1/9 unit 0 encapsulation vlan-bridge
      set interfaces ge-1/1/9 unit 0 vlan-id 10
      set interfaces ge-1/1/9 unit 1 encapsulation vlan-bridge
      set interfaces ge-1/1/9 unit 1 vlan-id 20
      set interfaces ge-1/1/9 unit 2 encapsulation vlan-bridge
      set interfaces ge-1/1/9 unit 2 vlan-id 30
      set interfaces ge-1/1/9 unit 110 encapsulation vlan-bridge
      set interfaces ge-1/1/9 unit 110 vlan-id 110
      set interfaces ge-1/1/9 unit 120 encapsulation vlan-bridge
      set interfaces ge-1/1/9 unit 120 vlan-id 120
      set interfaces ge-1/1/9 unit 130 encapsulation vlan-bridge
      set interfaces ge-1/1/9 unit 130 vlan-id 130
      set interfaces ge-1/1/9 unit 210 encapsulation vlan-bridge
      set interfaces ge-1/1/9 unit 210 vlan-id 210
      set interfaces ge-1/1/9 unit 220 encapsulation vlan-bridge
      set interfaces ge-1/1/9 unit 220 vlan-id 220
      set interfaces ge-1/1/9 unit 230 encapsulation vlan-bridge
      set interfaces ge-1/1/9 unit 230 vlan-id 230
      set interfaces ae0 flexible-vlan-tagging
      set interfaces ae0 encapsulation flexible-ethernet-services
      set interfaces ae0 unit 0 encapsulation vlan-bridge
      set interfaces ae0 unit 0 vlan-id 10
      set interfaces ae0 unit 1 encapsulation vlan-bridge
      set interfaces ae0 unit 1 vlan-id 20
      set interfaces ae0 unit 2 encapsulation vlan-bridge
      set interfaces ae0 unit 2 vlan-id 30
      set interfaces ae0 unit 110 encapsulation vlan-bridge

```

```
set interfaces ae0 unit 110 vlan-id 110
set interfaces ae0 unit 120 encapsulation vlan-bridge
set interfaces ae0 unit 120 vlan-id 120
set interfaces ae0 unit 130 encapsulation vlan-bridge
set interfaces ae0 unit 130 vlan-id 130
set interfaces ae0 unit 210 encapsulation vlan-bridge
set interfaces ae0 unit 210 vlan-id 210
set interfaces ae0 unit 220 encapsulation vlan-bridge
set interfaces ae0 unit 220 vlan-id 220
set interfaces ae0 unit 230 encapsulation vlan-bridge
set interfaces ae0 unit 230 vlan-id 230
set interfaces lo0 unit 6 family inet address 203.0.113.0/24
set interfaces lo0 unit 6 family iso address
    47.0005.80ff.f800.0000.0108.0000.6006.0070.0600
set routing-options forwarding-table export load-balancing-policy
set protocols l2-learning global-mac-table-aging-time 18000
set policy-options policy-statement load-balancing-policy then load-balance per-packet
set bridge-domains bd10 domain-type bridge
set bridge-domains bd10 vlan-id 10
set bridge-domains bd10 interface ae0.0
set bridge-domains bd10 interface ge-1/1/9.0
set bridge-domains bd110 domain-type bridge
set bridge-domains bd110 vlan-id 110
set bridge-domains bd110 interface ae0.110
set bridge-domains bd110 interface ge-1/1/9.110
set bridge-domains bd120 domain-type bridge
set bridge-domains bd120 vlan-id 120
set bridge-domains bd120 interface ge-1/1/9.120
set bridge-domains bd120 interface ae0.120
set bridge-domains bd130 domain-type bridge
set bridge-domains bd130 vlan-id 130
set bridge-domains bd130 interface ge-1/1/9.130
set bridge-domains bd130 interface ae0.130
set bridge-domains bd20 domain-type bridge
set bridge-domains bd20 vlan-id 20
set bridge-domains bd20 interface ae0.1
set bridge-domains bd20 interface ge-1/1/9.1
set bridge-domains bd210 domain-type bridge
set bridge-domains bd210 vlan-id 210
set bridge-domains bd210 interface ae0.210
set bridge-domains bd210 interface ge-1/1/9.210
set bridge-domains bd220 domain-type bridge
set bridge-domains bd220 vlan-id 220
set bridge-domains bd220 interface ge-1/1/9.220
set bridge-domains bd220 interface ae0.220
set bridge-domains bd230 domain-type bridge
set bridge-domains bd230 vlan-id 230
set bridge-domains bd230 interface ge-1/1/9.230
set bridge-domains bd230 interface ae0.230
set bridge-domains bd30 domain-type bridge
set bridge-domains bd30 vlan-id 30
set bridge-domains bd30 interface ge-1/1/9.2
set bridge-domains bd30 interface ae0.2
```

PE1 set chassis aggregated-devices ethernet device-count 20


```

set chassis network-services enhanced-ip
set interfaces ge-1/0/3 gigether-options 802.3ad ae13
set interfaces ge-1/0/4 gigether-options 802.3ad ae13
set interfaces ge-1/0/5 gigether-options 802.3ad ae12
set interfaces ge-1/1/4 unit 0 family inet address 192.0.2.1/24
set interfaces ge-1/1/4 unit 0 family iso
set interfaces ge-1/1/4 unit 0 family mpls
set interfaces ge-1/1/6 gigether-options 802.3ad ae12
set interfaces ge-1/2/1 flexible-vlan-tagging
set interfaces ge-1/2/1 encapsulation flexible-ethernet-services
set interfaces ge-1/2/1 esi 00:33:33:33:33:33:33:33
set interfaces ge-1/2/1 esi all-active
set interfaces ge-1/2/1 unit 0 encapsulation vlan-bridge
set interfaces ge-1/2/1 unit 0 vlan-id 30
set interfaces ge-1/2/1 unit 130 family bridge interface-mode trunk
set interfaces ge-1/2/1 unit 130 family bridge vlan-id-list 130
set interfaces ge-1/2/1 unit 230 family bridge interface-mode trunk
set interfaces ge-1/2/1 unit 230 family bridge vlan-id-list 230
set interfaces ge-1/2/2 flexible-vlan-tagging
set interfaces ge-1/2/2 encapsulation flexible-ethernet-services
set interfaces ge-1/2/2 esi 00:22:22:22:22:22:22:22
set interfaces ge-1/2/2 esi all-active
set interfaces ge-1/2/2 unit 0 encapsulation vlan-bridge
set interfaces ge-1/2/2 unit 0 vlan-id 20
set interfaces ge-1/2/2 unit 120 family bridge interface-mode trunk
set interfaces ge-1/2/2 unit 120 family bridge vlan-id-list 120
set interfaces ge-1/2/2 unit 220 family bridge interface-mode trunk
set interfaces ge-1/2/2 unit 220 family bridge vlan-id-list 220
set interfaces ge-1/2/3 gigether-options 802.3ad ae0
set interfaces ge-1/2/4 gigether-options 802.3ad ae0
set interfaces ge-1/3/0 flexible-vlan-tagging
set interfaces ge-1/3/0 encapsulation flexible-ethernet-services
set interfaces ge-1/3/0 unit 0 encapsulation vlan-bridge
set interfaces ge-1/3/0 unit 0 vlan-id 10
set interfaces ge-1/3/0 unit 100 family bridge interface-mode trunk
set interfaces ge-1/3/0 unit 100 family bridge vlan-id-list 110
set interfaces ge-1/3/0 unit 100 family bridge vlan-id-list 120
set interfaces ge-1/3/0 unit 100 family bridge vlan-id-list 130
set interfaces ge-1/3/0 unit 200 family bridge interface-mode trunk
set interfaces ge-1/3/0 unit 200 family bridge vlan-id-list 210
set interfaces ge-1/3/0 unit 200 family bridge vlan-id-list 220
set interfaces ge-1/3/0 unit 200 family bridge vlan-id-list 230
set interfaces ae0 flexible-vlan-tagging
set interfaces ae0 encapsulation flexible-ethernet-services
set interfaces ae0 esi 00:11:11:11:11:11:11:11
set interfaces ae0 esi all-active
set interfaces ae0 unit 0 encapsulation vlan-bridge
set interfaces ae0 unit 0 vlan-id 10
set interfaces ae0 unit 110 family bridge interface-mode trunk
set interfaces ae0 unit 110 family bridge vlan-id-list 110
set interfaces ae0 unit 210 family bridge interface-mode trunk
set interfaces ae0 unit 210 family bridge vlan-id-list 210
set interfaces ae12 flexible-vlan-tagging
set interfaces ae12 encapsulation flexible-ethernet-services
set interfaces ae12 aggregated-ether-options minimum-links 1
set interfaces ae12 unit 0 vlan-id 1200

```

```
set interfaces ae12 unit 0 family inet address 198.51.100.12/24
set interfaces ae12 unit 0 family iso
set interfaces ae12 unit 0 family mpls
set interfaces ae13 flexible-vlan-tagging
set interfaces ae13 encapsulation flexible-ethernet-services
set interfaces ae13 aggregated-ether-options minimum-links 1
set interfaces ae13 unit 0 vlan-tags outer 1300
set interfaces ae13 unit 0 vlan-tags inner 13
set interfaces ae13 unit 0 family inet address 198.51.100.13/24
set interfaces ae13 unit 0 family iso
set interfaces ae13 unit 0 family mpls
set interfaces irb unit 0 family inet address 192.0.2.9/24
set interfaces irb unit 0 mac 00:99:99:99:01:99
set interfaces irb unit 1 family inet address 192.0.2.10/24
set interfaces irb unit 1 mac 00:99:99:99:02:99
set interfaces irb unit 2 family inet address 192.0.2.11/24
set interfaces irb unit 2 mac 00:99:99:99:03:99
set interfaces irb unit 10 family inet address 192.0.2.12/24
set interfaces irb unit 10 mac 00:99:99:99:01:90
set interfaces lo0 unit 0 family inet address 198.51.100.1/24 primary
set interfaces lo0 unit 0 family iso
set routing-options router-id 198.51.100.1
set routing-options autonomous-system 65221
set routing-options forwarding-table export load-balancing-policy
set protocols rsvp interface all
set protocols rsvp interface fxp0.0 disable
set protocols mpls label-switched-path pe1tope2 from 198.51.100.1
set protocols mpls label-switched-path pe1tope2 to 198.51.100.2
set protocols mpls label-switched-path pe1tope2 primary direct_to_pe2
set protocols mpls label-switched-path pe1tope3 from 198.51.100.1
set protocols mpls label-switched-path pe1tope3 to 198.51.100.3
set protocols mpls label-switched-path pe1tope3 primary direct_to_pe3
set protocols mpls label-switched-path pe1tope4 from 198.51.100.1
set protocols mpls label-switched-path pe1tope4 to 198.51.100.4
set protocols mpls label-switched-path pe1tope4 primary direct_to_pe4
set protocols mpls path pe4_to_pe3 198.51.100.4 strict
set protocols mpls path pe4_to_pe3 198.51.100.3 strict
set protocols mpls path direct_to_pe2 198.51.100.5 strict
set protocols mpls path direct_to_pe3 198.51.100.6 strict
set protocols mpls path direct_to_pe4 198.51.100.9 strict
set protocols mpls path pe2_to_pe3 198.51.100.2 strict
set protocols mpls path pe2_to_pe3 198.51.100.3 strict
set protocols mpls interface all
set protocols mpls interface fxp0.0 disable
set protocols bgp group RR type internal
set protocols bgp group RR local-address 198.51.100.1
set protocols bgp group RR family evpn signaling
set protocols bgp group RR neighbor 203.0.113.0
set protocols isis level 1 disable
set protocols isis interface all level 2 metric 10
set protocols isis interface fxp0.0 disable
set protocols isis interface lo0.0 level 2 metric 0
set protocols ldp deaggregate
set protocols ldp interface all
set protocols ldp interface fxp0.0 disable
set protocols evpn
```

```

set protocols l2-learning global-mac-table-aging-time 18000
set policy-options policy-statement load-balancing-policy then load-balance per-packet
set routing-instances VS-1 instance-type virtual-switch
set routing-instances VS-1 interface ge-1/2/1.130
set routing-instances VS-1 interface ge-1/2/2.120
set routing-instances VS-1 interface ge-1/3/0.100
set routing-instances VS-1 interface ae0.110
set routing-instances VS-1 route-distinguisher 198.51.100.1:101
set routing-instances VS-1 vrf-target target:100:101
set routing-instances VS-1 protocols evpn extended-vlan-list 110
set routing-instances VS-1 protocols evpn extended-vlan-list 120
set routing-instances VS-1 protocols evpn extended-vlan-list 130
set routing-instances VS-1 protocols evpn default-gateway do-not-advertise
set routing-instances VS-1 bridge-domains bd-110 vlan-id 110
set routing-instances VS-1 bridge-domains bd-110 routing-interface irb.0
set routing-instances VS-1 bridge-domains bd-120 vlan-id 120
set routing-instances VS-1 bridge-domains bd-120 routing-interface irb.1
set routing-instances VS-1 bridge-domains bd-130 vlan-id 130
set routing-instances VS-1 bridge-domains bd-130 routing-interface irb.2
set routing-instances VS-2 instance-type virtual-switch
set routing-instances VS-2 interface ge-1/2/1.230
set routing-instances VS-2 interface ge-1/2/2.220
set routing-instances VS-2 interface ge-1/3/0.200
set routing-instances VS-2 interface ae0.210
set routing-instances VS-2 route-distinguisher 198.51.100.1:201
set routing-instances VS-2 vrf-target target:100:201
set routing-instances VS-2 protocols evpn extended-vlan-list 210
set routing-instances VS-2 protocols evpn extended-vlan-list 220
set routing-instances VS-2 protocols evpn extended-vlan-list 230
set routing-instances VS-2 bridge-domains bd-a vlan-id-list 210
set routing-instances VS-2 bridge-domains bd-a vlan-id-list 220
set routing-instances VS-2 bridge-domains bd-a vlan-id-list 230
set routing-instances mhevnp instance-type evpn
set routing-instances mhevnp vlan-id 10
set routing-instances mhevnp interface ge-1/2/1.0
set routing-instances mhevnp interface ge-1/2/2.0
set routing-instances mhevnp interface ge-1/3/0.0
set routing-instances mhevnp interface ae0.0
set routing-instances mhevnp routing-interface irb.10
set routing-instances mhevnp route-distinguisher 198.51.100.1:1
set routing-instances mhevnp vrf-target target:100:1
set routing-instances mhevnp protocols evpn default-gateway do-not-advertise
set routing-instances vrf instance-type vrf
set routing-instances vrf interface irb.0
set routing-instances vrf interface irb.1
set routing-instances vrf interface irb.2
set routing-instances vrf interface irb.10
set routing-instances vrf route-distinguisher 198.51.100.1:11
set routing-instances vrf vrf-target target:100:11
set routing-instances vrf vrf-table-label

```

```

PE2 set chassis aggregated-devices ethernet device-count 20
set chassis network-services enhanced-ip
set interfaces ge-0/0/5 gigether-options 802.3ad ae12
set interfaces ge-0/1/0 gigether-options 802.3ad ae12

```

```
set interfaces ge-0/1/1 flexible-vlan-tagging
set interfaces ge-0/1/1 encapsulation flexible-ethernet-services
set interfaces ge-0/1/1 esi 00:33:33:33:33:33:33:33:33
set interfaces ge-0/1/1 esi all-active
set interfaces ge-0/1/1 unit 0 encapsulation vlan-bridge
set interfaces ge-0/1/1 unit 0 vlan-id 30
set interfaces ge-0/1/1 unit 130 family bridge interface-mode trunk
set interfaces ge-0/1/1 unit 130 family bridge vlan-id-list 130
set interfaces ge-0/1/1 unit 230 family bridge interface-mode trunk
set interfaces ge-0/1/1 unit 230 family bridge vlan-id-list 230
set interfaces ge-0/1/2 flexible-vlan-tagging
set interfaces ge-0/1/2 encapsulation flexible-ethernet-services
set interfaces ge-0/1/2 esi 00:22:22:22:22:22:22:22:22
set interfaces ge-0/1/2 esi all-active
set interfaces ge-0/1/2 unit 0 encapsulation vlan-bridge
set interfaces ge-0/1/2 unit 0 vlan-id 20
set interfaces ge-0/1/2 unit 120 family bridge interface-mode trunk
set interfaces ge-0/1/2 unit 120 family bridge vlan-id-list 120
set interfaces ge-0/1/2 unit 220 family bridge interface-mode trunk
set interfaces ge-0/1/2 unit 220 family bridge vlan-id-list 220
set interfaces ge-0/1/3 unit 0 family inet address 192.0.2.3/24
set interfaces ge-0/1/3 unit 0 family iso
set interfaces ge-0/1/3 unit 0 family mpls
set interfaces ge-0/1/4 gigether-options 802.3ad ae0
set interfaces ge-0/1/6 flexible-vlan-tagging
set interfaces ge-0/1/6 encapsulation flexible-ethernet-services
set interfaces ge-0/1/6 unit 0 encapsulation vlan-bridge
set interfaces ge-0/1/6 unit 0 vlan-id 10
set interfaces ge-0/1/6 unit 100 family bridge interface-mode trunk
set interfaces ge-0/1/6 unit 100 family bridge vlan-id-list 110
set interfaces ge-0/1/6 unit 100 family bridge vlan-id-list 120
set interfaces ge-0/1/6 unit 100 family bridge vlan-id-list 130
set interfaces ge-0/1/6 unit 200 family bridge interface-mode trunk
set interfaces ge-0/1/6 unit 200 family bridge vlan-id-list 210
set interfaces ge-0/1/6 unit 200 family bridge vlan-id-list 220
set interfaces ge-0/1/6 unit 200 family bridge vlan-id-list 230
set interfaces ge-0/1/9 gigether-options 802.3ad ae0
set interfaces ae0 flexible-vlan-tagging
set interfaces ae0 encapsulation flexible-ethernet-services
set interfaces ae0 esi 00:11:11:11:11:11:11:11:11
set interfaces ae0 esi all-active
set interfaces ae0 unit 0 encapsulation vlan-bridge
set interfaces ae0 unit 0 vlan-id 10
set interfaces ae0 unit 110 family bridge interface-mode trunk
set interfaces ae0 unit 110 family bridge vlan-id-list 110
set interfaces ae0 unit 210 family bridge interface-mode trunk
set interfaces ae0 unit 210 family bridge vlan-id-list 210
set interfaces ae12 flexible-vlan-tagging
set interfaces ae12 encapsulation flexible-ethernet-services
set interfaces ae12 aggregated-ether-options minimum-links 1
set interfaces ae12 unit 0 vlan-id 1200
set interfaces ae12 unit 0 family inet address 198.51.100.5/24
set interfaces ae12 unit 0 family iso
set interfaces ae12 unit 0 family mpls
set interfaces irb unit 0 family inet address 192.0.2.9/24
set interfaces irb unit 0 mac 00:99:99:99:01:99
```

```
set interfaces irb unit 1 family inet address 192.0.2.10/24
set interfaces irb unit 1 mac 00:99:99:99:02:99
set interfaces irb unit 2 family inet address 192.0.2.11/24
set interfaces irb unit 2 mac 00:99:99:99:03:99
set interfaces irb unit 10 family inet address 192.0.2.12/24
set interfaces irb unit 10 mac 00:99:99:99:01:90
set interfaces lo0 unit 0 family inet address 198.51.100.2/32 primary
set interfaces lo0 unit 0 family iso
set routing-options router-id 198.51.100.2
set routing-options autonomous-system 65221
set routing-options forwarding-table export load-balancing-policy
set protocols rsvp interface all
set protocols rsvp interface fxp0.0 disable
set protocols mpls label-switched-path pe2tope1 from 198.51.100.2
set protocols mpls label-switched-path pe2tope1 to 198.51.100.1
set protocols mpls label-switched-path pe2tope1 primary direct_to_pe1
set protocols mpls label-switched-path pe2tope3 from 198.51.100.2
set protocols mpls label-switched-path pe2tope3 to 198.51.100.3
set protocols mpls label-switched-path pe2tope3 primary direct_to_pe3
set protocols mpls label-switched-path pe2tope4 from 198.51.100.2
set protocols mpls label-switched-path pe2tope4 to 198.51.100.4
set protocols mpls label-switched-path pe2tope4 primary direct_to_pe4
set protocols mpls path direct_to_pe1 198.51.100.12 strict
set protocols mpls path direct_to_pe3 198.51.100.7 strict
set protocols mpls path direct_to_pe4 198.51.100.8 strict
set protocols mpls interface all
set protocols mpls interface fxp0.0 disable
set protocols bgp group RR type internal
set protocols bgp group RR local-address 198.51.100.2
set protocols bgp group RR family evpn signaling
set protocols bgp group RR neighbor 203.0.113.0
set protocols isis level 1 disable
set protocols isis interface all level 2 metric 10
set protocols isis interface fxp0.0 disable
set protocols isis interface lo0.0 level 2 metric 0
set protocols ldp deaggregate
set protocols ldp interface all
set protocols ldp interface fxp0.0 disable
set protocols l2-learning global-mac-table-aging-time 18000
set policy-options policy-statement load-balancing-policy then load-balance per-packet
set routing-instances VS-1 instance-type virtual-switch
set routing-instances VS-1 interface ge-0/1/1.130
set routing-instances VS-1 interface ge-0/1/2.120
set routing-instances VS-1 interface ge-0/1/6.100
set routing-instances VS-1 interface ae0.110
set routing-instances VS-1 route-distinguisher 198.51.100.2:101
set routing-instances VS-1 vrf-target target:100:101
set routing-instances VS-1 protocols evpn extended-vlan-list 110
set routing-instances VS-1 protocols evpn extended-vlan-list 120
set routing-instances VS-1 protocols evpn extended-vlan-list 130
set routing-instances VS-1 protocols evpn default-gateway do-not-advertise
set routing-instances VS-1 bridge-domains bd-110 vlan-id 110
set routing-instances VS-1 bridge-domains bd-110 routing-interface irb.0
set routing-instances VS-1 bridge-domains bd-120 vlan-id 120
set routing-instances VS-1 bridge-domains bd-120 routing-interface irb.1
set routing-instances VS-1 bridge-domains bd-130 vlan-id 130
```

```

set routing-instances VS-1 bridge-domains bd-130 routing-interface irb.2
set routing-instances VS-2 instance-type virtual-switch
set routing-instances VS-2 interface ge-0/1/1.230
set routing-instances VS-2 interface ge-0/1/2.220
set routing-instances VS-2 interface ge-0/1/6.200
set routing-instances VS-2 interface ae0.210
set routing-instances VS-2 route-distinguisher 198.51.100.2:201
set routing-instances VS-2 vrf-target target:100:201
set routing-instances VS-2 protocols evpn extended-vlan-list 210
set routing-instances VS-2 protocols evpn extended-vlan-list 220
set routing-instances VS-2 protocols evpn extended-vlan-list 230
set routing-instances VS-2 bridge-domains bd-a vlan-id-list 210
set routing-instances VS-2 bridge-domains bd-a vlan-id-list 220
set routing-instances VS-2 bridge-domains bd-a vlan-id-list 230
set routing-instances mhevnp instance-type evpn
set routing-instances mhevnp vlan-id 10
set routing-instances mhevnp interface ge-0/1/1.0
set routing-instances mhevnp interface ge-0/1/2.0
set routing-instances mhevnp interface ge-0/1/6.0
set routing-instances mhevnp interface ae0.0
set routing-instances mhevnp routing-interface irb.10
set routing-instances mhevnp route-distinguisher 198.51.100.2:1
set routing-instances mhevnp vrf-target target:100:1
set routing-instances mhevnp protocols evpn default-gateway do-not-advertise
set routing-instances vrf instance-type vrf
set routing-instances vrf interface irb.0
set routing-instances vrf interface irb.1
set routing-instances vrf interface irb.2
set routing-instances vrf interface irb.10
set routing-instances vrf route-distinguisher 198.51.100.2:11
set routing-instances vrf vrf-target target:100:11
set routing-instances vrf vrf-table-label

```

```

PE3 set chassis aggregated-devices ethernet device-count 20
set chassis network-services enhanced-ip
set interfaces ge-2/0/1 gigether-options 802.3ad ae13
set interfaces ge-2/0/2 gigether-options 802.3ad ae13
set interfaces ge-2/0/4 gigether-options 802.3ad ae0
set interfaces ge-2/0/9 flexible-vlan-tagging
set interfaces ge-2/0/9 encapsulation flexible-ethernet-services
set interfaces ge-2/0/9 unit 0 encapsulation vlan-bridge
set interfaces ge-2/0/9 unit 0 vlan-id 10
set interfaces ge-2/0/9 unit 100 family bridge interface-mode trunk
set interfaces ge-2/0/9 unit 100 family bridge vlan-id-list 110
set interfaces ge-2/0/9 unit 100 family bridge vlan-id-list 120
set interfaces ge-2/0/9 unit 100 family bridge vlan-id-list 130
set interfaces ge-2/0/9 unit 200 family bridge interface-mode trunk
set interfaces ge-2/0/9 unit 200 family bridge vlan-id-list 210
set interfaces ge-2/0/9 unit 200 family bridge vlan-id-list 220
set interfaces ge-2/0/9 unit 200 family bridge vlan-id-list 230
set interfaces ge-2/1/0 unit 0 family inet address 192.0.2.5/24
set interfaces ge-2/1/0 unit 0 family iso
set interfaces ge-2/1/0 unit 0 family mpls
set interfaces ge-2/1/6 gigether-options 802.3ad ae0
set interfaces ae0 flexible-vlan-tagging

```

```
set interfaces ae0 encapsulation flexible-ethernet-services
set interfaces ae0 esi 00:11:11:11:11:11:11:11
set interfaces ae0 esi all-active
set interfaces ae0 unit 0 encapsulation vlan-bridge
set interfaces ae0 unit 0 vlan-id 10
set interfaces ae0 unit 110 family bridge interface-mode trunk
set interfaces ae0 unit 110 family bridge vlan-id-list 110
set interfaces ae0 unit 210 family bridge interface-mode trunk
set interfaces ae0 unit 210 family bridge vlan-id-list 210
set interfaces ae13 flexible-vlan-tagging
set interfaces ae13 encapsulation flexible-ethernet-services
set interfaces ae13 aggregated-ether-options minimum-links 1
set interfaces ae13 unit 0 vlan-tags outer 1300
set interfaces ae13 unit 0 vlan-tags inner 13
set interfaces ae13 unit 0 family inet address 198.51.100.6/24
set interfaces ae13 unit 0 family iso
set interfaces ae13 unit 0 family mpls
set interfaces irb unit 0 family inet address 192.0.2.9/24
set interfaces irb unit 0 mac 00:99:99:99:01:99
set interfaces irb unit 1 family inet address 192.0.2.10/24
set interfaces irb unit 1 mac 00:99:99:99:02:99
set interfaces irb unit 2 family inet address 192.0.2.11/24
set interfaces irb unit 2 mac 00:99:99:99:03:99
set interfaces irb unit 10 family inet address 192.0.2.12/24
set interfaces irb unit 10 mac 00:99:99:99:01:90
set interfaces lo0 unit 0 family inet address 198.51.100.3/32 primary
set interfaces lo0 unit 0 family iso
set routing-options router-id 198.51.100.3
set routing-options autonomous-system 65221
set routing-options forwarding-table export load-balancing-policy
set protocols rsvp interface all
set protocols rsvp interface fxp0.0 disable
set protocols mpls label-switched-path pe3tope1 from 198.51.100.3
set protocols mpls label-switched-path pe3tope1 to 198.51.100.1
set protocols mpls label-switched-path pe3tope1 primary direct_to_pe1
set protocols mpls label-switched-path pe3tope2 from 198.51.100.3
set protocols mpls label-switched-path pe3tope2 to 198.51.100.2
set protocols mpls label-switched-path pe3tope2 primary direct_to_pe2
set protocols mpls label-switched-path pe3tope4 from 198.51.100.3
set protocols mpls label-switched-path pe3tope4 to 198.51.100.4
set protocols mpls label-switched-path pe3tope4 primary direct_to_pe4
set protocols mpls path direct_to_pe1 198.51.100.13 strict
set protocols mpls path direct_to_pe2 198.51.100.10 strict
set protocols mpls path direct_to_pe4 198.51.100.11 strict
set protocols mpls interface all
set protocols mpls interface fxp0.0 disable
set protocols bgp group RR type internal
set protocols bgp group RR local-address 198.51.100.3
set protocols bgp group RR family evpn signaling
set protocols bgp group RR neighbor 203.0.113.0
set protocols isis level 1 disable
set protocols isis interface all level 2 metric 10
set protocols isis interface fxp0.0 disable
set protocols isis interface lo0.0 level 2 metric 0
set protocols ldp deaggregate
set protocols ldp interface all
```

```
set protocols ldp interface fxp0.0 disable
set protocols l2-learning global-mac-table-aging-time 18000
set policy-options policy-statement load-balancing-policy then load-balance per-packet
set routing-instances VS-1 instance-type virtual-switch
set routing-instances VS-1 interface ge-2/0/0.100
set routing-instances VS-1 interface ge-2/0/9.100
set routing-instances VS-1 interface ae0.110
set routing-instances VS-1 route-distinguisher 198.51.100.3:101
set routing-instances VS-1 vrf-target target:100:101
set routing-instances VS-1 protocols evpn extended-vlan-list 110
set routing-instances VS-1 protocols evpn extended-vlan-list 120
set routing-instances VS-1 protocols evpn extended-vlan-list 130
set routing-instances VS-1 protocols evpn default-gateway do-not-advertise
set routing-instances VS-1 bridge-domains bd-110 vlan-id 110
set routing-instances VS-1 bridge-domains bd-110 routing-interface irb.0
set routing-instances VS-1 bridge-domains bd-120 vlan-id 120
set routing-instances VS-1 bridge-domains bd-120 routing-interface irb.1
set routing-instances VS-1 bridge-domains bd-130 vlan-id 130
set routing-instances VS-1 bridge-domains bd-130 routing-interface irb.2
set routing-instances VS-2 instance-type virtual-switch
set routing-instances VS-2 interface ge-2/0/0.200
set routing-instances VS-2 interface ge-2/0/9.200
set routing-instances VS-2 interface ae0.210
set routing-instances VS-2 route-distinguisher 198.51.100.3:201
set routing-instances VS-2 vrf-target target:100:201
set routing-instances VS-2 protocols evpn extended-vlan-list 210
set routing-instances VS-2 protocols evpn extended-vlan-list 220
set routing-instances VS-2 protocols evpn extended-vlan-list 230
set routing-instances VS-2 bridge-domains bd-a vlan-id-list 210
set routing-instances VS-2 bridge-domains bd-a vlan-id-list 220
set routing-instances VS-2 bridge-domains bd-a vlan-id-list 230
set routing-instances mhevpn instance-type evpn
set routing-instances mhevpn vlan-id 10
set routing-instances mhevpn interface ge-2/0/0.0
set routing-instances mhevpn interface ge-2/0/9.0
set routing-instances mhevpn interface ae0.0
set routing-instances mhevpn routing-interface irb.10
set routing-instances mhevpn route-distinguisher 198.51.100.3:1
set routing-instances mhevpn vrf-target target:100:1
set routing-instances mhevpn protocols evpn default-gateway do-not-advertise
set routing-instances vrf instance-type vrf
set routing-instances vrf interface irb.0
set routing-instances vrf interface irb.1
set routing-instances vrf interface irb.2
set routing-instances vrf interface irb.10
set routing-instances vrf route-distinguisher 198.51.100.3:11
set routing-instances vrf vrf-target target:100:11
set routing-instances vrf vrf-table-label
```

```
PE4 set chassis aggregated-devices ethernet device-count 20
set chassis network-services enhanced-ip
set interfaces ge-1/0/3 unit 0 family inet address 192.0.2.7/24
set interfaces ge-1/0/3 unit 0 family iso
set interfaces ge-1/0/3 unit 0 family mpls
set interfaces ge-1/1/1 flexible-vlan-tagging
```



```
set interfaces ge-1/1/1 encapsulation flexible-ethernet-services
set interfaces ge-1/1/1 esi 00:44:44:44:44:44:44:44:44
set interfaces ge-1/1/1 esi all-active
set interfaces ge-1/1/1 unit 0 encapsulation vlan-bridge
set interfaces ge-1/1/1 unit 0 vlan-id 10
set interfaces ge-1/1/1 unit 100 family bridge interface-mode trunk
set interfaces ge-1/1/1 unit 100 family bridge vlan-id-list 110
set interfaces ge-1/1/1 unit 100 family bridge vlan-id-list 120
set interfaces ge-1/1/1 unit 100 family bridge vlan-id-list 130
set interfaces ge-1/1/1 unit 200 family bridge interface-mode trunk
set interfaces ge-1/1/1 unit 200 family bridge vlan-id-list 210
set interfaces ge-1/1/1 unit 200 family bridge vlan-id-list 220
set interfaces ge-1/1/1 unit 200 family bridge vlan-id-list 230
set interfaces ge-1/1/9 flexible-vlan-tagging
set interfaces ge-1/1/9 encapsulation flexible-ethernet-services
set interfaces ge-1/1/9 unit 0 encapsulation vlan-bridge
set interfaces ge-1/1/9 unit 0 vlan-id 10
set interfaces ge-1/1/9 unit 100 family bridge interface-mode trunk
set interfaces ge-1/1/9 unit 100 family bridge vlan-id-list 110
set interfaces ge-1/1/9 unit 100 family bridge vlan-id-list 120
set interfaces ge-1/1/9 unit 100 family bridge vlan-id-list 130
set interfaces ge-1/1/9 unit 200 family bridge interface-mode trunk
set interfaces ge-1/1/9 unit 200 family bridge vlan-id-list 210
set interfaces ge-1/1/9 unit 200 family bridge vlan-id-list 220
set interfaces ge-1/1/9 unit 200 family bridge vlan-id-list 230
set interfaces irb unit 0 family inet address 192.0.2.9/24
set interfaces irb unit 0 mac 00:99:99:99:01:99
set interfaces irb unit 1 family inet address 192.0.2.10/24
set interfaces irb unit 1 mac 00:99:99:99:02:99
set interfaces irb unit 2 family inet address 192.0.2.11/24
set interfaces irb unit 2 mac 00:99:99:99:03:99
set interfaces irb unit 10 family inet address 192.0.2.12/24
set interfaces irb unit 10 mac 00:99:99:99:01:90
set interfaces lo0 unit 0 family inet address 198.51.100.4/32 primary
set routing-options router-id 198.51.100.4
set routing-options autonomous-system 65221
set routing-options forwarding-table export load-balancing-policy
set protocols rsvp interface all
set protocols rsvp interface fxp0.0 disable
set protocols mpls label-switched-path pe4tope1 from 198.51.100.4
set protocols mpls label-switched-path pe4tope1 to 198.51.100.1
set protocols mpls label-switched-path pe4tope1 primary direct_to_pe1
set protocols mpls label-switched-path pe4tope2 from 198.51.100.4
set protocols mpls label-switched-path pe4tope2 to 198.51.100.2
set protocols mpls label-switched-path pe4tope2 primary direct_to_pe2
set protocols mpls label-switched-path pe4tope3 from 198.51.100.4
set protocols mpls label-switched-path pe4tope3 to 198.51.100.3
set protocols mpls label-switched-path pe4tope3 primary direct_to_pe3
set protocols mpls path pe2_to_pe3 198.51.100.2 strict
set protocols mpls path pe2_to_pe3 198.51.100.3 strict
set protocols mpls path direct_to_pe1 198.51.100.14 strict
set protocols mpls path direct_to_pe2 198.51.100.15 strict
set protocols mpls path direct_to_pe3 198.51.100.16 strict
set protocols mpls interface all
set protocols mpls interface fxp0.0 disable
set protocols bgp group RR type internal
```

```
set protocols bgp group RR local-address 198.51.100.4
set protocols bgp group RR family evpn signaling
set protocols bgp group RR neighbor 203.0.113.0
set protocols isis level 1 disable
set protocols isis interface all level 2 metric 10
set protocols isis interface fxp0.0 disable
set protocols isis interface lo0.0 level 2 metric 0
set protocols ldp deaggregate
set protocols ldp interface all
set protocols ldp interface fxp0.0 disable
set protocols l2-learning global-mac-table-aging-time 18000
set policy-options policy-statement load-balancing-policy then load-balance per-packet
set routing-instances VS-1 instance-type virtual-switch
set routing-instances VS-1 interface ge-1/1/1.100
set routing-instances VS-1 interface ge-1/1/9.100
set routing-instances VS-1 route-distinguisher 198.51.100.4:101
set routing-instances VS-1 vrf-target target:100:101
set routing-instances VS-1 protocols evpn extended-vlan-list 110
set routing-instances VS-1 protocols evpn extended-vlan-list 120
set routing-instances VS-1 protocols evpn extended-vlan-list 130
set routing-instances VS-1 protocols evpn default-gateway do-not-advertise
set routing-instances VS-1 bridge-domains bd-110 vlan-id 110
set routing-instances VS-1 bridge-domains bd-110 routing-interface irb.0
set routing-instances VS-1 bridge-domains bd-120 vlan-id 120
set routing-instances VS-1 bridge-domains bd-120 routing-interface irb.1
set routing-instances VS-1 bridge-domains bd-130 vlan-id 130
set routing-instances VS-1 bridge-domains bd-130 routing-interface irb.2
set routing-instances VS-2 instance-type virtual-switch
set routing-instances VS-2 interface ge-1/1/1.200
set routing-instances VS-2 interface ge-1/1/9.200
set routing-instances VS-2 route-distinguisher 198.51.100.4:201
set routing-instances VS-2 vrf-target target:100:201
set routing-instances VS-2 protocols evpn extended-vlan-list 210
set routing-instances VS-2 protocols evpn extended-vlan-list 220
set routing-instances VS-2 protocols evpn extended-vlan-list 230
set routing-instances VS-2 bridge-domains bd-a vlan-id-list 210
set routing-instances VS-2 bridge-domains bd-a vlan-id-list 220
set routing-instances VS-2 bridge-domains bd-a vlan-id-list 230
set routing-instances mhevpn instance-type evpn
set routing-instances mhevpn vlan-id 10
set routing-instances mhevpn interface ge-1/1/1.0
set routing-instances mhevpn interface ge-1/1/9.0
set routing-instances mhevpn routing-interface irb.10
set routing-instances mhevpn route-distinguisher 198.51.100.4:1
set routing-instances mhevpn vrf-target target:100:1
set routing-instances mhevpn protocols evpn default-gateway do-not-advertise
set routing-instances vrf instance-type vrf
set routing-instances vrf interface irb.0
set routing-instances vrf interface irb.1
set routing-instances vrf interface irb.2
set routing-instances vrf interface irb.10
set routing-instances vrf route-distinguisher 198.51.100.4:11
set routing-instances vrf vrf-target target:100:11
set routing-instances vrf vrf-table-label
```

```

P (RR)  set interfaces ge-1/0/5 unit 0 family inet address 192.0.2.8/24
        set interfaces ge-1/1/2 unit 0 family inet address 192.0.2.6/24
        set interfaces ge-1/1/3 unit 0 family inet address 192.0.2.4/24
        set interfaces ge-1/1/4 unit 0 family inet address 192.0.2.2/24
        set interfaces lo0 unit 0 family inet address 203.0.113.0
        set protocols bgp group RR type internal
        set protocols bgp group RR local-address 203.0.113.0
        set protocols bgp group RR family evpn signaling
        set protocols bgp group RR cluster 1.2.3.4
        set protocols bgp group RR neighbor 198.51.100.1
        set protocols bgp group RR neighbor 198.51.100.2
        set protocols bgp group RR neighbor 198.51.100.3
        set protocols bgp group RR neighbor 198.51.100.4
        set protocols isis interface all level 1 disable
        set protocols ldp interface all
        set routing-options router-id 203.0.113.0
        set routing-options autonomous-system 65221

```

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

To configure Router PE1:



NOTE: Repeat this procedure for all other multihomed PE routers after modifying the appropriate interface names, addresses, and other parameters.

1. Configure the MX Series router to operate in the enhanced-ip mode because the EVPN active-active functionality is supported on routers with MPCs and MIC interfaces only.

A system reboot is required on committing this configuration.

```

[edit chassis]
user@PE1# set network-services enhanced-ip

```

2. Specify the number of aggregated Ethernet interfaces to be created.

```

[edit chassis]
user@PE1# set aggregated-devices ethernet device-count 20

```

3. Configure Router PE1 interfaces within the ae0 aggregated bundle toward the multihomed customer site, Router CE10.

- a. Assign interfaces ge-1/2/3 and ge-1/2/4 within the ae0 aggregated bundle.

```

[edit interfaces]
user@PE1# set ge-1/2/3 gigether-options 802.3ad ae0
user@PE1# set ge-1/2/4 gigether-options 802.3ad ae0

```

- b. Configure the ae0 aggregated bundle parameters for VLAN tagging and encapsulation.

```
[edit interfaces]
user@PE1# set ae0 flexible-vlan-tagging
user@PE1# set ae0 encapsulation flexible-ethernet-services
```

- c. Assign an ESI value for the first Ethernet segment and enable EVPN active-active multihoming for the ae0 aggregated bundle.

```
[edit interfaces]
user@PE1# set ae0 esi 00:11:11:11:11:11:11:11
user@PE1# set ae0 esi all-active
```

- d. Configure a trunk interface on the bridge network for the ae0 aggregated bundle.

```
[edit interfaces]
user@PE1# set ae0 unit 0 encapsulation vlan-bridge
user@PE1# set ae0 unit 0 vlan-id 10
user@PE1# set ae0 unit 110 family bridge interface-mode trunk
user@PE1# set ae0 unit 110 family bridge vlan-id-list 110
user@PE1# set ae0 unit 210 family bridge interface-mode trunk
user@PE1# set ae0 unit 210 family bridge vlan-id-list 210
```

4. Configure the other Router PE1 interfaces toward the multihomed customer site, Router CE10.

- a. Configure the VLAN tagging and encapsulation parameters for the ge-1/2/2 PE1 interface.

```
[edit interfaces]
user@PE1# set ge-1/2/2 flexible-vlan-tagging
user@PE1# set ge-1/2/2 encapsulation flexible-ethernet-services
```

- b. Assign an ESI value for the second Ethernet segment and enable EVPN active-active multihoming for the ge-1/2/2 PE1 interface.

```
[edit interfaces]
user@PE1# set ge-1/2/2 esi 00:22:22:22:22:22:22:22
user@PE1# set ge-1/2/2 esi all-active
```

- c. Configure a trunk interface on the bridge network for the ge-1/2/2 PE1 interface.

```
[edit interfaces]
user@PE1# set ge-1/2/2 unit 0 encapsulation vlan-bridge
user@PE1# set ge-1/2/2 unit 0 vlan-id 20
user@PE1# set ge-1/2/2 unit 120 family bridge interface-mode trunk
user@PE1# set ge-1/2/2 unit 120 family bridge vlan-id-list 120
user@PE1# set ge-1/2/2 unit 220 family bridge interface-mode trunk
user@PE1# set ge-1/2/2 unit 220 family bridge vlan-id-list 220
```

- d. Configure the VLAN tagging and encapsulation parameters for the ge-1/2/1 PE1 interface.

```
[edit interfaces]
user@PE1# set ge-1/2/1 flexible-vlan-tagging
user@PE1# set ge-1/2/1 encapsulation flexible-ethernet-services
```

- e. Assign an ESI value for the third Ethernet segment and enable EVPN active-active multihoming for the ge-1/2/1 PE1 interface.

```
[edit interfaces]
user@PE1# set ge-1/2/1 esi 00:33:33:33:33:33:33:33
user@PE1# set ge-1/2/1 esi all-active
```

- f. Configure a trunk interface on the bridge network for the ge-1/2/1 PE1 interface.

```
[edit interfaces]
user@PE1# set ge-1/2/1 unit 0 encapsulation vlan-bridge
user@PE1# set ge-1/2/1 unit 0 vlan-id 30
user@PE1# set ge-1/2/1 unit 130 family bridge interface-mode trunk
user@PE1# set ge-1/2/1 unit 130 family bridge vlan-id-list 130
user@PE1# set ge-1/2/1 unit 230 family bridge interface-mode trunk
user@PE1# set ge-1/2/1 unit 230 family bridge vlan-id-list 230
```

5. Configure Router PE1 interfaces toward Router PE2.

- a. Assign the interfaces ge-1/0/5 and ge-1/1/6 within the ae12 aggregated bundle.

```
[edit interfaces]
user@PE1# set ge-1/0/5 gigether-options 802.3ad ae12
user@PE1# set ge-1/1/6 gigether-options 802.3ad ae12
```

- b. Specify the minimum number of links for the ae12 aggregated bundle to be labeled “up”.

```
[edit interfaces]
user@PE1# set ae12 aggregated-ether-options minimum-links 1
```

- c. Assign an IP address for the ae12 aggregated bundle and enable MPLS and IS-IS protocol families on the bundle.

```
[edit interfaces]
user@PE1# set ae12 unit 0 family inet address 198.51.100.12/24
user@PE1# set ae12 unit 0 family iso
user@PE1# set ae12 unit 0 family mpls
```

- d. Configure the VLAN tagging and encapsulation parameters for the ae12 aggregated bundle and assign VLAN ID 1200 for the bundle.

```
[edit interfaces]
user@PE1# set ae12 flexible-vlan-tagging
user@PE1# set ae12 encapsulation flexible-ethernet-services
user@PE1# set ae12 unit 0 vlan-id 1200
```

6. Configure Router PE1 interfaces toward Router PE3.

- a. Assign the interfaces ge-1/0/3 and ge-1/0/4 within the ae13 aggregated bundle.

```
[edit interfaces]
user@PE1# set ge-1/0/3 gigether-options 802.3ad ae13
user@PE1# set ge-1/0/4 gigether-options 802.3ad ae13
```

- b. Specify the minimum number of links for the ae13 aggregated bundle to be labeled “up”.

```
[edit interfaces]
user@PE1# set ae13 aggregated-ether-options minimum-links 1
```

- c. Assign an IP address for the ae13 aggregated bundle and enable MPLS and IS-IS protocol families on the bundle.

```
[edit interfaces]
user@PE1# set ae13 unit 0 family inet address 198.51.100.13/24
user@PE1# set ae13 unit 0 family iso
user@PE1# set ae13 unit 0 family mpls
```

- d. Configure the VLAN tagging and encapsulation parameters for the ae12 aggregated bundle and assign the inner and outer VLAN tags for the bundle.

```
[edit interfaces]
user@PE1# set ae13 flexible-vlan-tagging
user@PE1# set ae13 encapsulation flexible-ethernet-services
user@PE1# set ae13 unit 0 vlan-tags outer 1300
user@PE1# set ae13 unit 0 vlan-tags inner 13
```

7. Configure the Router PE1 interface toward the single-homed customer site, Router CE1.

- a. Configure the VLAN tagging and encapsulation parameters for the ge-1/3/0 PE1 interface.

```
[edit interfaces]
user@PE1# set ge-1/3/0 flexible-vlan-tagging
user@PE1# set ge-1/3/0 encapsulation flexible-ethernet-services
```

- b. Configure a trunk interface on the bridge network for the ge-1/3/0 PE1 interface.

```
[edit interfaces]
user@PE1# set ge-1/3/0 unit 0 encapsulation vlan-bridge
user@PE1# set ge-1/3/0 unit 0 vlan-id 10
user@PE1# set ge-1/3/0 unit 100 family bridge interface-mode trunk
user@PE1# set ge-1/3/0 unit 100 family bridge vlan-id-list 110
user@PE1# set ge-1/3/0 unit 100 family bridge vlan-id-list 120
user@PE1# set ge-1/3/0 unit 100 family bridge vlan-id-list 130
user@PE1# set ge-1/3/0 unit 200 family bridge interface-mode trunk
user@PE1# set ge-1/3/0 unit 200 family bridge vlan-id-list 210
user@PE1# set ge-1/3/0 unit 200 family bridge vlan-id-list 220
user@PE1# set ge-1/3/0 unit 200 family bridge vlan-id-list 230
```

8. Configure the Router PE1 interface toward Router P (RR) and enable the MPLS and IS-IS protocol families for the interface.

```
[edit interfaces]
user@PE1# set ge-1/1/4 unit 0 family inet address 192.0.2.1/24
user@PE1# set ge-1/1/4 unit 0 family iso
user@PE1# set ge-1/1/4 unit 0 family mpls
```

9. Configure an IRB interface for Router PE1.

```
[edit interfaces]
user@PE1# set irb unit 0 family inet address 192.0.2.9/24
user@PE1# set irb unit 0 mac 00:99:99:99:01:99
user@PE1# set irb unit 1 family inet address 192.0.2.10/24
user@PE1# set irb unit 1 mac 00:99:99:99:02:99
user@PE1# set irb unit 2 family inet address 192.0.2.11/24
user@PE1# set irb unit 2 mac 00:99:99:99:03:99
user@PE1# set irb unit 10 family inet address 192.0.2.12/24
user@PE1# set irb unit 10 mac 00:99:99:99:01:90
```

10. Configure the loopback interface for Router PE1.

```
[edit interfaces]
user@PE1# set lo0 unit 0 family inet address 198.51.100.1/24 primary
user@PE1# set lo0 unit 0 family iso
```

11. Assign a router ID and the autonomous system number for Router PE1.

```
[edit routing-options]
user@PE1# set router-id 198.51.100.1
user@PE1# set autonomous-system 65221
```
12. Assign a load-balancing policy to the forwarding table of Router PE1.

```
[edit routing-options]
user@PE1# set forwarding-table export load-balancing-policy
```
13. Configure IS-IS on Router PE1.

```
[edit protocols]
user@PE1# set isis level 1 disable
user@PE1# set isis interface all level 2 metric 10
user@PE1# set isis interface fxp0.0 disable
user@PE1# set isis interface lo0.0 level 2 metric 0
```
14. Configure an internal BGP group for Router PE1 to peer with route reflector, Router P.

```
[edit protocols]
user@PE1# set bgp group RR type internal
user@PE1# set bgp group RR local-address 198.51.100.1
user@PE1# set bgp group RR neighbor 203.0.113.0
```
15. Enable EVPN signaling for the RR BGP group on Router PE1.

```
[edit protocols]
user@PE1# set bgp group RR family evpn signaling
```
16. Configure RSVP, LDP, MPLS, EVPN, and L2 learning on Router PE1.

```
[edit protocols]
user@PE1# set rsvp interface all
user@PE1# set rsvp interface fxp0.0 disable

user@PE1# set ldp deaggregate
user@PE1# set ldp interface all
user@PE1# set ldp interface fxp0.0 disable

user@PE1# set mpls interface all
user@PE1# set mpls interface fxp0.0 disable

user@PE1# set evpn

user@PE1# set l2-learning global-mac-table-aging-time 18000
```
17. Configure label-switched paths between the PE routers.

```
[edit protocols]
user@PE1# set mpls label-switched-path pe1tope2 from 198.51.100.1
```

```
user@PE1# set mpls label-switched-path pe1tope2 to 198.51.100.2
user@PE1# set mpls label-switched-path pe1tope2 primary direct_to_pe2
```

```
user@PE1# set mpls label-switched-path pe1tope3 from 198.51.100.1
user@PE1# set mpls label-switched-path pe1tope3 to 198.51.100.3
user@PE1# set mpls label-switched-path pe1tope3 primary direct_to_pe3
```

```
user@PE1# set mpls label-switched-path pe1tope4 from 198.51.100.1
user@PE1# set mpls label-switched-path pe1tope4 to 198.51.100.4
user@PE1# set mpls label-switched-path pe1tope4 primary direct_to_pe4
```

18. Configure MPLS paths from Router PE1 to other PE routers.

```
user@PE1# set mpls path pe4_to_pe3 198.51.100.4 strict
user@PE1# set mpls path pe4_to_pe3 198.51.100.3 strict
```

```
user@PE1# set mpls path direct_to_pe2 198.51.100.5 strict
user@PE1# set mpls path direct_to_pe3 198.51.100.6 strict
```

```
user@PE1# set mpls path direct_to_pe4 198.51.100.9 strict
```

```
user@PE1# set mpls path pe2_to_pe3 198.51.100.2 strict
user@PE1# set mpls path pe2_to_pe3 198.51.100.3 strict
```

19. Configure the load-balancing policy to enable load balancing per packet.

```
[edit policy-options]
user@PE1# set policy-statement load-balancing-policy then load-balance per-packet
```

20. Configure the first virtual switch routing instance.

- a. Configure the routing-instance type and assign Router PE1 interfaces to the routing instance.

```
[edit routing-instances]
user@PE1# set VS-1 instance-type virtual-switch
user@PE1# set VS-1 interface ge-1/2/1.130
user@PE1# set VS-1 interface ge-1/2/2.120
user@PE1# set VS-1 interface ge-1/3/0.100
user@PE1# set VS-1 interface ae0.110
```

- b. Configure the route distinguisher and the VPN routing and forwarding (VRF) target for the VS-1 routing instance.

```
[edit routing-instances]
user@PE1# set VS-1 route-distinguisher 198.51.100.1:101
user@PE1# set VS-1 vrf-target target:100:101
```

- c. Configure EVPN and assign VLANs to the VS-1 routing instance.

```
[edit routing-instances]
user@PE1# set VS-1 protocols evpn extended-vlan-list 110
user@PE1# set VS-1 protocols evpn extended-vlan-list 120
user@PE1# set VS-1 protocols evpn extended-vlan-list 130
user@PE1# set VS-1 protocols evpn default-gateway do-not-advertise
```


- d. Configure the bridge domains and their associated VLANs and IRB interfaces for the VS-1 routing instance.

```
[edit routing-instances]
user@PE1# set VS-1 bridge-domains bd-110 vlan-id 110
user@PE1# set VS-1 bridge-domains bd-110 routing-interface irb.0
user@PE1# set VS-1 bridge-domains bd-120 vlan-id 120
user@PE1# set VS-1 bridge-domains bd-120 routing-interface irb.1
user@PE1# set VS-1 bridge-domains bd-130 vlan-id 130
user@PE1# set VS-1 bridge-domains bd-130 routing-interface irb.2
```

21. Configure the second virtual switch routing instance.

- a. Configure the routing-instance type and assign Router PE1 interfaces to the routing instance.

```
[edit routing-instances]
user@PE1# set VS-2 instance-type virtual-switch
user@PE1# set VS-2 interface ge-1/2/1.230
user@PE1# set VS-2 interface ge-1/2/2.220
user@PE1# set VS-2 interface ge-1/3/0.200
user@PE1# set VS-2 interface ae0.210
```

- b. Configure the route distinguisher and the VPN routing and forwarding (VRF) target for the VS-2 routing instance.

```
[edit routing-instances]
user@PE1# set VS-2 route-distinguisher 198.51.100.1:201
user@PE1# set VS-2 vrf-target target:100:201
```

- c. Configure EVPN and assign VLANs to the VS-2 routing instance.

```
[edit routing-instances]
user@PE1# set VS-2 protocols evpn extended-vlan-list 210
user@PE1# set VS-2 protocols evpn extended-vlan-list 220
user@PE1# set VS-2 protocols evpn extended-vlan-list 230
```

- d. Configure the bridge domains and their associated VLANs for the VS-2 routing instance.

```
[edit routing-instances]
user@PE1# set routing-instances VS-2 bridge-domains bd-a vlan-id-list 210
user@PE1# set routing-instances VS-2 bridge-domains bd-a vlan-id-list 220
user@PE1# set routing-instances VS-2 bridge-domains bd-a vlan-id-list 230
```

22. Configure the multihomed EVPN routing instance.

- a. Configure the routing-instance type and assign VLANs and Router PE1 interfaces to the routing instance.

```
[edit routing-instances]
user@PE1# set mhevpn instance-type evpn
user@PE1# set mhevpn vlan-id 10
user@PE1# set mhevpn interface ge-1/2/1.0
user@PE1# set mhevpn interface ge-1/2/2.0
user@PE1# set mhevpn interface ge-1/3/0.0
user@PE1# set mhevpn interface ae0.0
user@PE1# set mhevpn routing-interface irb.10
```

- b. Configure the route distinguisher and the VPN routing and forwarding (VRF) target for the mhevpn routing instance.

```
[edit routing-instances]
user@PE1# set mhevpn route-distinguisher 198.51.100.1:1
user@PE1# set mhevpn vrf-target target:100:1
```

- c. Configure EVPN to the mhevpn routing instance.

```
[edit routing-instances]
user@PE1# set routing-instances mhevpn protocols evpn default-gateway
do-not-advertise
```

23. Configure the default routing instance.

- a. Configure the routing-instance type and assign IRB interfaces to the routing instance.

```
[edit routing-instances]
user@PE1# set vrf instance-type vrf
user@PE1# set vrf interface irb.0
user@PE1# set vrf interface irb.1
user@PE1# set vrf interface irb.2
user@PE1# set vrf interface irb.10
```

- b. Configure the route distinguisher and the VPN routing and forwarding (VRF) target for the vrf routing instance.

```
[edit routing-instances]
user@PE1# set vrf route-distinguisher 198.51.100.1:11
user@PE1# set vrf vrf-target target:100:11
user@PE1# set vrf vrf-table-label
```

Configuration on EX9200 Switches

Step-by-Step Procedure

Several configuration statements used to configure active-active mode differ on EX9200 switches from those used on MX Series routers. This procedure shows which configuration statements are specific to EX9200 switches. All other configuration in this example applies both to EX9200 switches and MX Series routers.

1. To configure a trunk interface, include the **family ethernet-switching** statements instead of the **family bridge** statements in all occurrences.

```
[edit interfaces]
user@PE#1# set interfaces ge-1/2/1 unit 130 family ethernet-switching
interface-mode trunk
```

2. To configure the Layer 2 Ethernet switching domain, include the **vlan members (vlan-id | name)** statement instead of the **vlan-id-list vlan-id** statement in all occurrences.

```
[edit interfaces]
user@PE#1# set interfaces ge-1/2/1 unit 130 family ethernet-switching vlan members
130
```

3. To configure the VLAN domain and associated VLANs for each routing instance, include the **vlan name** statement, instead of the **bridge-domains** statement in all occurrences.

[edit]

```
user@PE#1# set routing-instances VS-1 vlans bd-110 vlan-id 110
```

4. To configure the IRB interface in each routing instance, include the **l3-interface irb-interface-name** statement instead of the **routing-interface** statement in all occurrences.

[edit]

```
user@PE#1# set routing-instances VS-1 vlans bd-110 l3-interface irb.0
```

Results

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

```
user@PE1 show chassis
aggregated-devices {
  ethernet {
    device-count 20;
  }
}
network-services enhanced-ip;

user@PE1 show interfaces
ge-1/0/3 {
  gigether-options {
    802.3ad ae13;
  }
}
ge-1/0/4 {
  gigether-options {
    802.3ad ae13;
  }
}
ge-1/0/5 {
  gigether-options {
    802.3ad ae12;
  }
}
ge-1/1/4 {
  unit 0 {
    family inet {
      address 192.0.2.1/24;
    }
    family iso;
    family mpls;
  }
}
```

```
ge-1/1/6 {
  gigaether-options {
    802.3ad ae12;
  }
}
ge-1/2/1 {
  flexible-vlan-tagging;
  encapsulation flexible-ethernet-services;
  esi {
    00:33:33:33:33:33:33:33:33;
    all-active;
  }
  unit 0 {
    encapsulation vlan-bridge;
    vlan-id 30;
  }
  unit 130 {
    family bridge {
      interface-mode trunk;
      vlan-id-list 130;
    }
  }
  unit 230 {
    family bridge {
      interface-mode trunk;
      vlan-id-list 230;
    }
  }
}
ge-1/2/2 {
  flexible-vlan-tagging;
  encapsulation flexible-ethernet-services;
  esi {
    00:22:22:22:22:22:22:22:22;
    all-active;
  }
  unit 0 {
    encapsulation vlan-bridge;
    vlan-id 20;
  }
  unit 120 {
    family bridge {
      interface-mode trunk;
      vlan-id-list 120;
    }
  }
  unit 220 {
    family bridge {
      interface-mode trunk;
      vlan-id-list 220;
    }
  }
}
ge-1/2/3 {
  gigaether-options {
    802.3ad ae0;
```

```

    }
  }
  ge-1/2/4 {
    gigether-options {
      802.3ad ae0;
    }
  }
  ge-1/3/0 {
    flexible-vlan-tagging;
    encapsulation flexible-ethernet-services;
    unit 0 {
      encapsulation vlan-bridge;
      vlan-id 10;
    }
    unit 100 {
      family bridge {
        interface-mode trunk;
        vlan-id-list [ 110 120 130 ];
      }
    }
    unit 200 {
      family bridge {
        interface-mode trunk;
        vlan-id-list [ 210 220 230 ];
      }
    }
  }
}
ae0 {
  flexible-vlan-tagging;
  encapsulation flexible-ethernet-services;
  esi {
    00:11:11:11:11:11:11:11;
    all-active;
  }
  unit 0 {
    encapsulation vlan-bridge;
    vlan-id 10;
  }
  unit 110 {
    family bridge {
      interface-mode trunk;
      vlan-id-list 110;
    }
  }
  unit 210 {
    family bridge {
      interface-mode trunk;
      vlan-id-list 210;
    }
  }
}
ae12 {
  flexible-vlan-tagging;
  encapsulation flexible-ethernet-services;
  aggregated-ether-options {
    minimum-links 1;
  }
}

```

```
}
unit 0 {
  vlan-id 1200;
  family inet {
    address 198.51.100.12/24;
  }
  family iso;
  family mpls;
}
}
ae13 {
  flexible-vlan-tagging;
  encapsulation flexible-ethernet-services;
  aggregated-ether-options {
    minimum-links 1;
  }
  unit 0 {
    vlan-tags outer 1300 inner 13;
    family inet {
      address 198.51.100.13/24;
    }
    family iso;
    family mpls;
  }
}
}
irb {
  unit 0 {
    family inet {
      address 192.0.2.9/24;
    }
    mac 00:99:99:99:01:99;
  }
  unit 1 {
    family inet {
      address 192.0.2.10/24;
    }
    mac 00:99:99:99:02:99;
  }
  unit 2 {
    family inet {
      address 192.0.2.11/24;
    }
    mac 00:99:99:99:03:99;
  }
  unit 10 {
    family inet {
      address 192.0.2.12/24;
    }
    mac 00:99:99:99:01:90;
  }
}
}
lo0 {
  unit 0 {
    family inet {
      address 198.51.100.1/24 {
        primary;
      }
    }
  }
}
```

```
    }  
  }  
  family iso;  
}  
}  
  
user@PE1# show routing-options  
router-id 198.51.100.1;  
autonomous-system 65221;  
forwarding-table {  
  export load-balancing-policy;  
}  
  
user@PE1# show protocols  
rsrp {  
  interface all;  
  interface fxp0.0 {  
    disable;  
  }  
}  
mpls {  
  label-switched-path pe1tope2 {  
    from 198.51.100.1;  
    to 198.51.100.2;  
    primary direct_to_pe2;  
  }  
  label-switched-path pe1tope3 {  
    from 198.51.100.1;  
    to 198.51.100.3;  
    primary direct_to_pe3;  
  }  
  label-switched-path pe1tope4 {  
    from 198.51.100.1;  
    to 198.51.100.4;  
    primary direct_to_pe4;  
  }  
  path pe4_to_pe3 {  
    198.51.100.4 strict;  
    198.51.100.3 strict;  
  }  
  path direct_to_pe2 {  
    198.51.100.5 strict;  
  }  
  path direct_to_pe3 {  
    198.51.100.6 strict;  
  }  
  path direct_to_pe4 {  
    198.51.100.9 strict;  
  }  
  path pe2_to_pe3 {  
    198.51.100.2 strict;  
    198.51.100.3 strict;  
  }  
  interface all;  
  interface fxp0.0 {  
    disable;  
  }  
}
```

```
}
bgp {
  group RR {
    type internal;
    local-address 198.51.100.1;
    family evpn {
      signaling;
    }
    neighbor 203.0.113.0;
  }
}
isis {
  level 1 disable;
  interface all {
    level 2 metric 10;
  }
  interface fxp0.0 {
    disable;
  }
  interface lo0.0 {
    level 2 metric 0;
  }
}
ldp {
  deaggregate;
  interface all;
  interface fxp0.0 {
    disable;
  }
}
evpn {
}
l2-learning {
  global-mac-table-aging-time 18000;
}

user@PE1# show policy-options
policy-statement load-balancing-policy {
  then {
    load-balance per-packet;
  }
}

user@PE1# show routing-instances
VS-1 {
  instance-type virtual-switch;
  interface ge-1/2/1.130;
  interface ge-1/2/2.120;
  interface ge-1/3/0.100;
  interface ae0.110;
  route-distinguisher 198.51.100.1:101;
  vrf-target target:100:101;
  protocols {
    evpn {
      extended-vlan-list [ 110 120 130 ];
      default-gateway do-not-advertise;
    }
  }
}
```



```

}
bridge-domains {
  bd-110 {
    vlan-id 110;
    routing-interface irb.0;
  }
  bd-120 {
    vlan-id 120;
    routing-interface irb.1;
  }
  bd-130 {
    vlan-id 130;
    routing-interface irb.2;
  }
}
}
VS-2 {
  instance-type virtual-switch;
  interface ge-1/2/1.230;
  interface ge-1/2/2.220;
  interface ge-1/3/0.200;
  interface ae0.210;
  route-distinguisher 198.51.100.1:201;
  vrf-target target:100:201;
  protocols {
    evpn {
      extended-vlan-list [ 210 220 230 ];
    }
  }
  bridge-domains {
    bd-a {
      vlan-id-list [ 210 220 230 ];
    }
  }
}
mhevpn {
  instance-type evpn;
  vlan-id 10;
  interface ge-1/2/1.0;
  interface ge-1/2/2.0;
  interface ge-1/3/0.0;
  interface ae0.0;
  routing-interface irb.10;
  route-distinguisher 198.51.100.1:1;
  vrf-target target:100:1;
  protocols {
    evpn {
      default-gateway do-not-advertise;
    }
  }
}
vrf {
  instance-type vrf;
  interface irb.0;
  interface irb.1;
  interface irb.2;

```

```
interface irb.10;  
  route-distinguisher 198.51.100.1:11;  
  vrf-target target:100:11;  
  vrf-table-label;  
}
```

Verification

Confirm that the configuration is working properly.

- [Verifying VPN Services in the Core on page 98](#)
- [Verifying the EVPN Instance Status on page 100](#)
- [Verifying the Autodiscovery Routes per Ethernet Segment on page 105](#)
- [Verifying the Ethernet Segment Route on page 107](#)
- [Verifying the DF Status on page 109](#)
- [Verifying the BDF Status on page 110](#)
- [Verifying the Remote IRB and Host IP on page 111](#)

Verifying VPN Services in the Core

Purpose Ensure that the protocols in the VPN core are functioning properly.

Action From operational mode, enter the **show isis adjacency** command.

```
user@PE1> show isis adjacency
Interface          System      L State      Hold (secs) SNPA
ge-1/2/4.0         CE10        2 Up          24
5c:5e:ab:e:6f:4
```

From operational mode, enter the **show bgp summary** command.

```
user@PE1> show bgp summary
Groups: 1 Peers: 1 Down peers: 0
Table          Tot Paths  Act Paths Suppressed  History Damp State  Pending
bgp.evpn.0
              45         45         0         0         0         0
Peer          AS      InPkt    OutPkt    OutQ    Flaps  Last Up/Dwn
State|#Active/Received/Accepted/Damped...
203.0.113.0   65221      90       26       0       0       3:18
Establ
  bgp.evpn.0: 45/45/45/0
  VS-1.evpn.0: 19/19/19/0
  VS-2.evpn.0: 19/19/19/0
  mhevpn.evpn.0: 13/13/13/0
  __default_evpn__.evpn.0: 4/4/4/0
```

```
user@P> show bgp summary
Groups: 1 Peers: 4 Down peers: 0
Table          Tot Paths  Act Paths Suppressed  History Damp State  Pending
bgp.evpn.0
              68         68         0         0         0         0
Peer          AS      InPkt    OutPkt    OutQ    Flaps  Last Up/Dwn
State|#Active/Received/Accepted/Damped...
198.51.100.1  65221      25       90       0       0       3:04
Establ
  bgp.evpn.0: 22/22/22/0
198.51.100.2  65221      32       80       0       0       6:12
Establ
  bgp.evpn.0: 22/22/22/0
198.51.100.3  65221      31       62       0       0       6:58
Establ
  bgp.evpn.0: 12/12/12/0
198.51.100.4  65221      28       88       0       0       6:04
Establ
  bgp.evpn.0: 12/12/12/0
```

From operational mode, enter the **show mpls lsp** command.

```
user@PE1> show mpls lsp
Ingress LSP: 3 sessions
To          From          State Rt P    ActivePath  LSPName
198.51.100.2 198.51.100.1 Up    0 *    direct_to_pe2 pe1tope2
198.51.100.3 198.51.100.1 Up    0 *    direct_to_pe3 pe1tope3
198.51.100.4 198.51.100.1 Up    0 *    direct_to_pe4 pe1tope4
Total 3 displayed, Up 3, Down 0

Egress LSP: 3 sessions
To          From          State Rt Style Labelin Labelout LSPName
198.51.100.1 198.51.100.3 Up    0 1 FF      3      - pe3tope1
198.51.100.1 198.51.100.4 Up    0 1 FF      3      - pe4tope1
198.51.100.1 198.51.100.2 Up    0 1 FF      3      - pe2tope1
```

Total 3 displayed, Up 3, Down 0

Transit LSP: 0 sessions

Total 0 displayed, Up 0, Down 0

From operational mode, enter the **show interface ae* terse** command.

```
user@PE1> show interface ae* terse
Interface           Admin Link Proto   Local Remote
ae0                  up    up
ae0.0                up    up    bridge
ae0.110              up    up    bridge
ae0.210              up    up    bridge
ae0.32767            up    up    multiservice
ae12                  up    up
ae12.0                up    up    inet    198.51.100.12/24
                        iso
                        mpls
                        multiservice
ae12.32767            up    up    multiservice
ae13                  up    up
ae13.0                up    up    inet    198.51.100.13/24
                        iso
                        mpls
                        multiservice
ae13.32767            up    up    multiservice
```

Meaning The protocols IS-IS, BGP and MPLS are up and running. The aggregated bundles configured on Router PE1 are up.

Verifying the EVPN Instance Status

Purpose Verify the EVPN routing instances and their status.

Action From operational mode, run the **show evpn instance extensive** command.

```
user@PE1> show evpn instance extensive
Instance: VS-1
Route Distinguisher: 198.51.100.1:101
Per-instance MAC route label: 301664
MAC database status           Local Remote
Total MAC addresses:          0      0
Default gateway MAC addresses: 3      0
Number of local interfaces: 4 (3 up)
Interface name  ESI                               Mode      Status
ae0.110         00:11:11:11:11:11:11:11:11 all-active Up
ge-1/2/1.130    00:33:33:33:33:33:33:33:33 all-active Up
ge-1/2/2.120    00:22:22:22:22:22:22:22:22 all-active Up
ge-1/3/0.100    00:00:00:00:00:00:00:00:00 single-homed Up
Number of IRB interfaces: 3 (3 up)
Interface name  VLAN ID  Status  L3 context
irb.0           110      Up      vrf
irb.1           120      Up      vrf
irb.2           130      Up      vrf
Number of bridge domains: 3
```

VLAN ID	Intfs / up	Mode	MAC sync	IM route label
110	2 1	Extended	Enabled	301984
120	2 1	Extended	Enabled	302000
130	2 1	Extended	Enabled	302016

Number of neighbors: 3

198.51.100.2

Received routes

MAC address advertisement: 0

MAC+IP address advertisement: 0

Inclusive multicast: 3

Ethernet auto-discovery: 6

198.51.100.3

Received routes

MAC address advertisement: 0

MAC+IP address advertisement: 0

Inclusive multicast: 1

Ethernet auto-discovery: 2

198.51.100.4

Received routes

MAC address advertisement: 0

MAC+IP address advertisement: 0

Inclusive multicast: 3

Ethernet auto-discovery: 2

Number of ethernet segments: 4

ESI: 00:11:11:11:11:11:11:11:11:11

Status: Resolved by IFL ae0.110

Local interface: ae0.110, Status: Up/Forwarding

Number of remote PEs connected: 2

Remote PE	MAC label	Aliasing label	Mode
198.51.100.3	0	305584	all-active
198.51.100.2	0	306000	all-active

Designated forwarder: 198.51.100.3

Backup forwarder: 198.51.100.1

Backup forwarder: 198.51.100.2

Advertised MAC label: 301792

Advertised aliasing label: 301792

Advertised split horizon label: 301808

ESI: 00:22:22:22:22:22:22:22:22:22

Status: Resolved by IFL ge-1/2/2.120

Local interface: ge-1/2/2.120, Status: Up/Forwarding

Number of remote PEs connected: 1

Remote PE	MAC label	Aliasing label	Mode
198.51.100.2	0	306032	all-active

Designated forwarder: 198.51.100.1

Backup forwarder: 198.51.100.2

Advertised MAC label: 301824

Advertised aliasing label: 301824

Advertised split horizon label: 301840

ESI: 00:33:33:33:33:33:33:33:33:33

Status: Resolved by IFL ge-1/2/1.130

Local interface: ge-1/2/1.130, Status: Up/Forwarding

Number of remote PEs connected: 1

Remote PE	MAC label	Aliasing label	Mode
198.51.100.2	0	306064	all-active

Designated forwarder: 198.51.100.1

Backup forwarder: 198.51.100.2

Advertised MAC label: 301856

Advertised aliasing label: 301856

Advertised split horizon label: 301872

ESI: 00:44:44:44:44:44:44:44:44:44

Status: Resolved by NH 1048613

```

Number of remote PEs connected: 1
  Remote PE      MAC label  Aliasing label  Mode
  198.51.100.4   0          305152         all-active

Instance: VS-2
Route Distinguisher: 198.51.100.1:201
Per-instance MAC route label: 301696
MAC database status
  Total MAC addresses: 0      0
  Default gateway MAC addresses: 0      0
Number of local interfaces: 4 (3 up)
  Interface name  ESI                                     Mode                Status
  ae0.210         00:11:11:11:11:11:11:11:11:11         all-active          Up
  ge-1/2/1.230    00:33:33:33:33:33:33:33:33:33         all-active          Up
  ge-1/2/2.220    00:22:22:22:22:22:22:22:22:22         all-active          Up
  ge-1/3/0.200    00:00:00:00:00:00:00:00:00:00         single-homed        Down
Number of IRB interfaces: 0 (0 up)
Number of bridge domains: 3
  VLAN ID  Intfs / up  Mode        MAC sync  IM route label
  210      2 1         Extended    Enabled    302032
  220      2 1         Extended    Enabled    302048
  230      2 1         Extended    Enabled    302064
Number of neighbors: 3
  198.51.100.2
    Received routes
      MAC address advertisement: 0
      MAC+IP address advertisement: 0
      Inclusive multicast: 3
      Ethernet auto-discovery: 6
  198.51.100.3
    Received routes
      MAC address advertisement: 0
      MAC+IP address advertisement: 0
      Inclusive multicast: 1
      Ethernet auto-discovery: 2
  198.51.100.4
    Received routes
      MAC address advertisement: 0
      MAC+IP address advertisement: 0
      Inclusive multicast: 3
      Ethernet auto-discovery: 2
Number of ethernet segments: 4
ESI: 00:11:11:11:11:11:11:11:11:11
Status: Resolved by IFL ae0.210
Local interface: ae0.210, Status: Up/Forwarding
Number of remote PEs connected: 2
  Remote PE      MAC label  Aliasing label  Mode
  198.51.100.3   0          305648         all-active
  198.51.100.2   0          306096         all-active
Designated forwarder: 198.51.100.1
Backup forwarder: 198.51.100.2
Backup forwarder: 198.51.100.3
Advertised MAC label: 301888
Advertised aliasing label: 301888
Advertised split horizon label: 301808
ESI: 00:22:22:22:22:22:22:22:22:22
Status: Resolved by IFL ge-1/2/2.220
Local interface: ge-1/2/2.220, Status: Up/Forwarding
Number of remote PEs connected: 1
  Remote PE      MAC label  Aliasing label  Mode
  198.51.100.2   0          306112         all-active

```

```

    Designated forwarder: 198.51.100.1
    Backup forwarder: 198.51.100.2
    Advertised MAC label: 301904
    Advertised aliasing label: 301904
    Advertised split horizon label: 301840
    ESI: 00:33:33:33:33:33:33:33:33
    Status: Resolved by IFL ge-1/2/1.230
    Local interface: ge-1/2/1.230, Status: Up/Forwarding
    Number of remote PEs connected: 1
      Remote PE      MAC label  Aliasing label  Mode
      198.51.100.2    0          306128          all-active
    Designated forwarder: 198.51.100.1
    Backup forwarder: 198.51.100.2
    Advertised MAC label: 301920
    Advertised aliasing label: 301920
    Advertised split horizon label: 301872
    ESI: 00:44:44:44:44:44:44:44:44
    Status: Resolved by NH 1048616
    Number of remote PEs connected: 1
      Remote PE      MAC label  Aliasing label  Mode
      198.51.100.4    0          305184          all-active

Instance: __default_evpn__
Route Distinguisher: 198.51.100.1:0
VLAN ID: None
Per-instance MAC route label: 301760
MAC database status
  Total MAC addresses:          Local  Remote
  Default gateway MAC addresses: 0       0
Number of local interfaces: 0 (0 up)
Number of IRB interfaces: 0 (0 up)
Number of bridge domains: 0
Number of neighbors: 2
  198.51.100.2
    Received routes
      Ethernet auto-discovery:          0
      Ethernet Segment:                  3
  198.51.100.3
    Received routes
      Ethernet auto-discovery:          0
      Ethernet Segment:                  1
Number of ethernet segments: 0

Instance: mhevpn
Route Distinguisher: 198.51.100.1:1
VLAN ID: 10
Per-instance MAC route label: 301728
MAC database status
  Total MAC addresses:          Local  Remote
  Default gateway MAC addresses: 1       0
Number of local interfaces: 4 (3 up)
  Interface name  ESI                                     Mode          Status
  ae0.0           00:11:11:11:11:11:11:11:11             all-active     Up
  ge-1/2/1.0       00:33:33:33:33:33:33:33:33             all-active     Up
  ge-1/2/2.0       00:22:22:22:22:22:22:22:22             all-active     Up
  ge-1/3/0.0       00:00:00:00:00:00:00:00:00             single-homed    Down
Number of IRB interfaces: 1 (1 up)
  Interface name  VLAN ID  Status  L3 context
  irb.10          10       Up      vrf
Number of bridge domains: 1
  VLAN ID  Intfs / up  Mode          MAC sync  IM route label

```

```

10          4 3      Extended      Enabled  302080
Number of neighbors: 3
198.51.100.2
  Received routes
    MAC address advertisement:      0
    MAC+IP address advertisement:    0
    Inclusive multicast:             1
    Ethernet auto-discovery:         6
198.51.100.3
  Received routes
    MAC address advertisement:      0
    MAC+IP address advertisement:    0
    Inclusive multicast:             1
    Ethernet auto-discovery:         2
198.51.100.4
  Received routes
    MAC address advertisement:      0
    MAC+IP address advertisement:    0
    Inclusive multicast:             1
    Ethernet auto-discovery:         2
Number of ethernet segments: 4
ESI: 00:11:11:11:11:11:11:11:11
Status: Resolved by IFL ae0.0
Local interface: ae0.0, Status: Up/Forwarding
Number of remote PEs connected: 2
  Remote PE      MAC label  Aliasing label  Mode
  198.51.100.3    0          305680          all-active
  198.51.100.2    0          306144          all-active
Designated forwarder: 198.51.100.2
Backup forwarder: 198.51.100.1
Backup forwarder: 198.51.100.3
Advertised MAC label: 301936
Advertised aliasing label: 301936
Advertised split horizon label: 301808
ESI: 00:22:22:22:22:22:22:22:22
Status: Resolved by IFL ge-1/2/2.0
Local interface: ge-1/2/2.0, Status: Up/Forwarding
Number of remote PEs connected: 1
  Remote PE      MAC label  Aliasing label  Mode
  198.51.100.2    0          306160          all-active
Designated forwarder: 198.51.100.1
Backup forwarder: 198.51.100.2
Advertised MAC label: 301952
Advertised aliasing label: 301952
Advertised split horizon label: 301840
ESI: 00:33:33:33:33:33:33:33:33
Status: Resolved by IFL ge-1/2/1.0
Local interface: ge-1/2/1.0, Status: Up/Forwarding
Number of remote PEs connected: 1
  Remote PE      MAC label  Aliasing label  Mode
  198.51.100.2    0          306176          all-active
Designated forwarder: 198.51.100.1
Backup forwarder: 198.51.100.2
Advertised MAC label: 301968
Advertised aliasing label: 301968
Advertised split horizon label: 301872
ESI: 00:44:44:44:44:44:44:44:44
Status: Resolved by NH 1048612
Number of remote PEs connected: 1
  Remote PE      MAC label  Aliasing label  Mode
  198.51.100.4    0          305200          all-active

```


Meaning The output provides the following information:

- List of EVPN and virtual switch routing instances
- Mode of operation of each interface
- Neighbors of each routing instance
- Number of different routes received from each neighbor
- ESI attached to each routing instance
- Number of Ethernet segments on each routing instance
- DF election roles for each ESI in an EVI
- VLAN ID and MAC labels for each routing instance
- IRB interface details
- Number of default gateway MAC addresses received for the virtual switch routing instance (VS-1 and VS-2)

Verifying the Autodiscovery Routes per Ethernet Segment

Purpose Verify that the autodiscovery routes per Ethernet segment are received.

Action From operational mode, run the **show route table mhevpn.evpn.0** command.

Router PE1

```
user@PE1> show route table mhevpn.evpn.0
mhevpn.evpn.0: 17 destinations, 17 routes (17 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

1:198.51.100.1:1::111111111111111111::0/304
    *[EVPN/170] 00:11:37
    Indirect
1:198.51.100.1:1::222222222222222222::0/304
    *[EVPN/170] 00:11:37
    Indirect
1:198.51.100.1:1::333333333333333333::0/304
    *[EVPN/170] 00:11:37
    Indirect
1:198.51.100.2:0::111111111111111111::FFFF:FFFF/304
    *[BGP/170] 00:11:33, localpref 100, from 203.0.113.0
    AS path: I, validation-state: unverified
    > to 198.51.100.5 via ae12.0, label-switched-path pe1tope2
1:198.51.100.2:0::222222222222222222::FFFF:FFFF/304
    *[BGP/170] 00:11:33, localpref 100, from 203.0.113.0
    AS path: I, validation-state: unverified
    > to 198.51.100.5 via ae12.0, label-switched-path pe1tope2
1:198.51.100.2:0::333333333333333333::FFFF:FFFF/304
    *[BGP/170] 00:11:33, localpref 100, from 203.0.113.0
    AS path: I, validation-state: unverified
    > to 198.51.100.5 via ae12.0, label-switched-path pe1tope2
1:198.51.100.2:1::111111111111111111::0/304
```

```

* [BGP/170] 00:11:33, localpref 100, from 203.0.113.0
  AS path: I, validation-state: unverified
  > to 198.51.100.5 via ae12.0, label-switched-path pe1tope2
1:198.51.100.2:1::22222222222222222222::0/304
* [BGP/170] 00:11:33, localpref 100, from 203.0.113.0
  AS path: I, validation-state: unverified
  > to 198.51.100.5 via ae12.0, label-switched-path pe1tope2
1:198.51.100.2:1::33333333333333333333::0/304
* [BGP/170] 00:11:33, localpref 100, from 203.0.113.0
  AS path: I, validation-state: unverified
  > to 198.51.100.5 via ae12.0, label-switched-path pe1tope2
1:198.51.100.3:0::11111111111111111111::FFFF:FFFF/304
* [BGP/170] 00:11:37, localpref 100, from 203.0.113.0
  AS path: I, validation-state: unverified
  > to 198.51.100.6 via ae13.0, label-switched-path pe1tope3
1:198.51.100.3:1::11111111111111111111::0/304
* [BGP/170] 00:11:37, localpref 100, from 203.0.113.0
  AS path: I, validation-state: unverified
  > to 198.51.100.6 via ae13.0, label-switched-path pe1tope3
3:198.51.100.1:1:10::198.51.100.1/304
* [EVPN/170] 00:13:38
  Indirect
3:198.51.100.2:1:10::198.51.100.2/304
* [BGP/170] 00:11:33, localpref 100, from 203.0.113.0
  AS path: I, validation-state: unverified
  > to 198.51.100.5 via ae12.0, label-switched-path pe1tope2
3:198.51.100.3:1:10::198.51.100.3/304
* [BGP/170] 00:11:37, localpref 100, from 203.0.113.0
  AS path: I, validation-state: unverified
  > to 198.51.100.6 via ae13.0, label-switched-path pe1tope3

```

Router PE2

```

user@PE2> show route table mhevpn.evpn.0
mhevpn.evpn.0: 17 destinations, 17 routes (17 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

1:198.51.100.1:0::11111111111111111111::FFFF:FFFF/304
* [BGP/170] 01:10:26, localpref 100, from 203.0.113.0
  AS path: I, validation-state: unverified
  > to 198.51.100.12 via ae12.0, label-switched-path pe2tope1
1:198.51.100.1:0::22222222222222222222::FFFF:FFFF/304
* [BGP/170] 01:10:26, localpref 100, from 203.0.113.0
  AS path: I, validation-state: unverified
  > to 198.51.100.12 via ae12.0, label-switched-path pe2tope1
1:198.51.100.1:0::33333333333333333333::FFFF:FFFF/304
* [BGP/170] 01:10:26, localpref 100, from 203.0.113.0
  AS path: I, validation-state: unverified
  > to 198.51.100.12 via ae12.0, label-switched-path pe2tope1
1:198.51.100.1:1:11111111111111111111::0/304
* [BGP/170] 01:10:26, localpref 100, from 203.0.113.0
  AS path: I, validation-state: unverified
  > to 198.51.100.12 via ae12.0, label-switched-path pe2tope1
1:198.51.100.1:1:22222222222222222222::0/304
* [BGP/170] 01:10:26, localpref 100, from 203.0.113.0
  AS path: I, validation-state: unverified
  > to 198.51.100.12 via ae12.0, label-switched-path pe2tope1
1:198.51.100.1:1:33333333333333333333::0/304
* [BGP/170] 01:10:26, localpref 100, from 203.0.113.0
  AS path: I, validation-state: unverified
  > to 198.51.100.12 via ae12.0, label-switched-path pe2tope1

```

```

1:198.51.100.2:1::11111111111111111111::0/304
    *[EVPN/170] 01:10:26
    Indirect
1:198.51.100.2:1::22222222222222222222::0/304
    *[EVPN/170] 01:10:26
    Indirect
1:198.51.100.2:1::33333333333333333333::0/304
    *[EVPN/170] 01:10:26
    Indirect
1:198.51.100.3:0::11111111111111111111::FFFF:FFFF/304
    *[BGP/170] 01:10:26, localpref 100, from 203.0.113.0
    AS path: I, validation-state: unverified
    > to 198.51.100.7 via ge-0/0/4.0, label-switched-path pe2tope3
1:198.51.100.3:1::11111111111111111111::0/304
    *[BGP/170] 01:10:26, localpref 100, from 203.0.113.0
    AS path: I, validation-state: unverified
    > to 198.51.100.7 via ge-0/0/4.0, label-switched-path pe2tope3
1:198.51.100.4:0::44444444444444444444::FFFF:FFFF/304
    *[BGP/170] 01:10:17, localpref 100, from 203.0.113.0
    AS path: I, validation-state: unverified
    > to 198.51.100.8 via ge-0/1/8.0, label-switched-path pe2tope4
1:198.51.100.4:1::44444444444444444444::0/304
    *[BGP/170] 01:10:17, localpref 100, from 203.0.113.0
    AS path: I, validation-state: unverified
    > to 198.51.100.8 via ge-0/1/8.0, label-switched-path pe2tope4
3:198.51.100.1:1::10::198.51.100.1/304
    *[BGP/170] 01:10:26, localpref 100, from 203.0.113.0
    AS path: I, validation-state: unverified
    > to 198.51.100.12 via ae12.0, label-switched-path pe2tope1
3:198.51.100.2:1::10::198.51.100.2/304
    *[EVPN/170] 01:12:14
    Indirect
3:198.51.100.3:1::10::198.51.100.3/304
    *[BGP/170] 01:10:26, localpref 100, from 203.0.113.0
    AS path: I, validation-state: unverified
    > to 198.51.100.7 via ge-0/0/4.0, label-switched-path pe2tope3
3:198.51.100.4:1::10::198.51.100.4/304
    *[BGP/170] 01:10:17, localpref 100, from 203.0.113.0
    AS path: I, validation-state: unverified
    > to 198.51.100.8 via ge-0/1/8.0, label-switched-path pe2tope4

```

Meaning The remote type 1 autodiscovery route is received for the ESI attached to Router PE2, which is the other PE router connected to the multihomed CE device.

Verifying the Ethernet Segment Route

Purpose Verify that the local and advertised autodiscovery routes per Ethernet segment and the Ethernet segment routes are received.

Action From operational mode, run the **show route table __default_evpn__evpn.0** command.

Router PE1

```

user@PE1> show route table __default_evpn__evpn.0
__default_evpn__evpn.0: 10 destinations, 10 routes (10 active, 0 holddown, 0
hidden)

```

```

+ = Active Route, - = Last Active, * = Both

1:198.51.100.1:0::111111111111111111::FFFF:FFFF/304
    *[EVPN/170] 00:25:18
    Indirect
1:198.51.100.1:0::222222222222222222::FFFF:FFFF/304
    *[EVPN/170] 00:25:18
    Indirect
1:198.51.100.1:0::333333333333333333::FFFF:FFFF/304
    *[EVPN/170] 00:25:18
    Indirect
4:198.51.100.1:0::111111111111111111:198.51.100.1/304
    *[EVPN/170] 00:25:18
    Indirect
4:198.51.100.1:0::222222222222222222:198.51.100.1/304
    *[EVPN/170] 00:25:18
    Indirect
4:198.51.100.1:0::333333333333333333:198.51.100.1/304
    *[EVPN/170] 00:25:18
    Indirect
4:198.51.100.2:0::111111111111111111:198.51.100.2/304
    *[BGP/170] 00:25:14, localpref 100, from 203.0.113.0
    AS path: I, validation-state: unverified
    > to 198.51.100.5 via ae12.0, label-switched-path pe1tope2
4:198.51.100.2:0::222222222222222222:198.51.100.2/304
    *[BGP/170] 00:25:14, localpref 100, from 203.0.113.0
    AS path: I, validation-state: unverified
    > to 198.51.100.5 via ae12.0, label-switched-path pe1tope2
4:198.51.100.2:0::333333333333333333:198.51.100.2/304
    *[BGP/170] 00:25:14, localpref 100, from 203.0.113.0
    AS path: I, validation-state: unverified
    > to 198.51.100.5 via ae12.0, label-switched-path pe1tope2
4:198.51.100.3:0::111111111111111111:198.51.100.3/304
    *[BGP/170] 00:25:18, localpref 100, from 203.0.113.0
    AS path: I, validation-state: unverified
    > to 198.51.100.6 via ae13.0, label-switched-path pe1tope3

```

Router PE2

```

user@PE2> show route table __default_evpn__.evpn.0
__default_evpn__.evpn.0: 10 destinations, 10 routes (10 active, 0 holddown, 0
hidden)
+ = Active Route, - = Last Active, * = Both

1:198.51.100.2:0::111111111111111111::FFFF:FFFF/304
    *[EVPN/170] 01:17:59
    Indirect
1:198.51.100.2:0::222222222222222222::FFFF:FFFF/304
    *[EVPN/170] 01:17:59
    Indirect
1:198.51.100.2:0::333333333333333333::FFFF:FFFF/304
    *[EVPN/170] 01:17:59
    Indirect
4:198.51.100.1:0::111111111111111111:198.51.100.1/304
    *[BGP/170] 01:17:59, localpref 100, from 203.0.113.0
    AS path: I, validation-state: unverified
    > to 198.51.100.12 via ae12.0, label-switched-path pe2tope1
4:198.51.100.1:0::222222222222222222:198.51.100.1/304
    *[BGP/170] 01:17:59, localpref 100, from 203.0.113.0
    AS path: I, validation-state: unverified
    > to 198.51.100.12 via ae12.0, label-switched-path pe2tope1

```

```

4:198.51.100.1:0::33333333333333333333333333333333:198.51.100.1/304
    * [BGP/170] 01:17:59, localpref 100, from 203.0.113.0
      AS path: I, validation-state: unverified
      > to 198.51.100.12 via ae12.0, label-switched-path pe2tope1
4:198.51.100.2:0::11111111111111111111111111111111:198.51.100.2/304
    * [EVPN/170] 01:17:59
      Indirect
4:198.51.100.2:0::22222222222222222222222222222222:198.51.100.2/304
    * [EVPN/170] 01:17:59
      Indirect
4:198.51.100.2:0::33333333333333333333333333333333:198.51.100.2/304
    * [EVPN/170] 01:17:59
      Indirect
4:198.51.100.3:0::11111111111111111111111111111111:198.51.100.3/304
    * [BGP/170] 01:17:59, localpref 100, from 203.0.113.0
      AS path: I, validation-state: unverified
      > to 198.51.100.7 via ge-0/0/4.0, label-switched-path pe2tope3

```

Meaning The output displays the local and remote type 1 (autodiscovery) and type 4 (Ethernet segment) routes.

Verifying the DF Status

Purpose Confirm which PE router is the designated forwarder (DF) for each routing instance.

Action From operational mode, run the **show evpn instance designated-forwarder** command.

```

user@PE1> show evpn instance designated forwarder
Instance: VS-1
  Number of ethernet segments: 4
    ESI: 00:11:11:11:11:11:11:11:11:11
      Designated forwarder: 198.51.100.3
    ESI: 00:22:22:22:22:22:22:22:22:22
      Designated forwarder: 198.51.100.1
    ESI: 00:33:33:33:33:33:33:33:33:33
      Designated forwarder: 198.51.100.1
    ESI: 00:44:44:44:44:44:44:44:44:44
      Designated forwarder: No local attachment to ethernet segment

Instance: VS-2
  Number of ethernet segments: 4
    ESI: 00:11:11:11:11:11:11:11:11:11
      Designated forwarder: 198.51.100.1
    ESI: 00:22:22:22:22:22:22:22:22:22
      Designated forwarder: 198.51.100.1
    ESI: 00:33:33:33:33:33:33:33:33:33
      Designated forwarder: 198.51.100.1
    ESI: 00:44:44:44:44:44:44:44:44:44
      Designated forwarder: No local attachment to ethernet segment

Instance: mhevpn
  Number of ethernet segments: 4
    ESI: 00:11:11:11:11:11:11:11:11:11
      Designated forwarder: 198.51.100.2
    ESI: 00:22:22:22:22:22:22:22:22:22
      Designated forwarder: 198.51.100.1

```

```
ESI: 00:33:33:33:33:33:33:33:33:33
  Designated forwarder: 198.51.100.1
ESI: 00:44:44:44:44:44:44:44:44:44
  Designated forwarder: No local attachment to ethernet segment
```

Meaning The designated forwarder is displayed for each routing instance and ESI.

Verifying the BDF Status

Purpose Confirm which PE router is the backup designated forwarder (BDF) for each routing instance.

Action From operational mode, run the **show evpn instance backup-forwarder** command.

```
user@PE1> show evpn instance backup-forwarder
Instance: VS-1
  Number of ethernet segments: 4
    ESI: 00:11:11:11:11:11:11:11:11:11
      Backup forwarder: 198.51.100.1
      Backup forwarder: 198.51.100.2
    ESI: 00:22:22:22:22:22:22:22:22:22
      Backup forwarder: 198.51.100.2
    ESI: 00:33:33:33:33:33:33:33:33:33
      Backup forwarder: 198.51.100.2
    ESI: 00:44:44:44:44:44:44:44:44:44
      Backup forwarder: No local attachment to ethernet segment

Instance: VS-2
  Number of ethernet segments: 4
    ESI: 00:11:11:11:11:11:11:11:11:11
      Backup forwarder: 198.51.100.2
      Backup forwarder: 198.51.100.3
    ESI: 00:22:22:22:22:22:22:22:22:22
      Backup forwarder: 198.51.100.2
    ESI: 00:33:33:33:33:33:33:33:33:33
      Backup forwarder: 198.51.100.2
    ESI: 00:44:44:44:44:44:44:44:44:44
      Backup forwarder: No local attachment to ethernet segment

Instance: mhevpn
  Number of ethernet segments: 4
    ESI: 00:11:11:11:11:11:11:11:11:11
      Backup forwarder: 198.51.100.1
      Backup forwarder: 198.51.100.3
    ESI: 00:22:22:22:22:22:22:22:22:22
      Backup forwarder: 198.51.100.2
    ESI: 00:33:33:33:33:33:33:33:33:33
      Backup forwarder: 198.51.100.2
    ESI: 00:44:44:44:44:44:44:44:44:44
      Backup forwarder: No local attachment to ethernet segment
```

Meaning The backup designated forwarder is displayed for each routing instance and ESI.

Verifying the Remote IRB and Host IP

Purpose Verify that the remote IRB IP and the host IP are received.

Router PE1 Action

From operational mode, run the **show route table mhevpn** command.

```
user@PE1> show route table mhevpn
mhevpn.evpn.0: 17 destinations, 17 routes (17 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

1:198.51.100.1:1::1111111111111111::0/304
    *[EVPN/170] 01:34:26
    Indirect
1:198.51.100.1:1::2222222222222222::0/304
    *[EVPN/170] 01:34:26
    Indirect
1:198.51.100.1:1::3333333333333333::0/304
    *[EVPN/170] 01:34:26
    Indirect
1:198.51.100.2:0::1111111111111111::FFFF:FFFF/304
    *[BGP/170] 01:34:22, localpref 100, from 203.0.113.0
    AS path: I, validation-state: unverified
    > to 198.51.100.5 via ae12.0, label-switched-path pe1tope2
1:198.51.100.2:0::2222222222222222::FFFF:FFFF/304
    *[BGP/170] 01:34:22, localpref 100, from 203.0.113.0
    AS path: I, validation-state: unverified
    > to 198.51.100.5 via ae12.0, label-switched-path pe1tope2
1:198.51.100.2:0::3333333333333333::FFFF:FFFF/304
    *[BGP/170] 01:34:22, localpref 100, from 203.0.113.0
    AS path: I, validation-state: unverified
    > to 198.51.100.5 via ae12.0, label-switched-path pe1tope2
1:198.51.100.2:1::1111111111111111::0/304
    *[BGP/170] 01:34:22, localpref 100, from 203.0.113.0
    AS path: I, validation-state: unverified
    > to 198.51.100.5 via ae12.0, label-switched-path pe1tope2
1:198.51.100.2:1::2222222222222222::0/304
    *[BGP/170] 01:34:22, localpref 100, from 203.0.113.0
    AS path: I, validation-state: unverified
    > to 198.51.100.5 via ae12.0, label-switched-path pe1tope2
1:198.51.100.2:1::3333333333333333::0/304
    *[BGP/170] 01:34:22, localpref 100, from 203.0.113.0
    AS path: I, validation-state: unverified
    > to 198.51.100.5 via ae12.0, label-switched-path pe1tope2
1:198.51.100.3:0::1111111111111111::FFFF:FFFF/304
    *[BGP/170] 01:34:26, localpref 100, from 203.0.113.0
    AS path: I, validation-state: unverified
    > to 198.51.100.6 via ae13.0, label-switched-path pe1tope3
1:198.51.100.3:1::1111111111111111::0/304
    *[BGP/170] 01:34:26, localpref 100, from 203.0.113.0
    AS path: I, validation-state: unverified
    > to 198.51.100.6 via ae13.0, label-switched-path pe1tope3
3:198.51.100.1:1:10::198.51.100.1/304
    *[EVPN/170] 01:36:27
    Indirect
3:198.51.100.2:1:10::198.51.100.2/304
    *[BGP/170] 01:34:22, localpref 100, from 203.0.113.0
```

```

        AS path: I, validation-state: unverified
        > to 198.51.100.5 via ae12.0, label-switched-path pe1tope2
3:198.51.100.3:1::10::198.51.100.3/304
    *[BGP/170] 01:34:26, localpref 100, from 203.0.113.0
        AS path: I, validation-state: unverified
        > to 198.51.100.6 via ae13.0, label-switched-path pe1tope3

```

Router PE2

From operational mode, run the **show route table mhevpn** command.

```

user@PE2> show route table mhevpn
mhevpn.evpn.0: 17 destinations, 17 routes (17 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

1:198.51.100.1:0::111111111111111111::FFFF:FFFF/304
    *[BGP/170] 01:35:11, localpref 100, from 203.0.113.0
        AS path: I, validation-state: unverified
        > to 198.51.100.12 via ae12.0, label-switched-path pe2tope1
1:198.51.100.1:0::222222222222222222::FFFF:FFFF/304
    *[BGP/170] 01:35:11, localpref 100, from 203.0.113.0
        AS path: I, validation-state: unverified
        > to 198.51.100.12 via ae12.0, label-switched-path pe2tope1
1:198.51.100.1:0::333333333333333333::FFFF:FFFF/304
    *[BGP/170] 01:35:11, localpref 100, from 203.0.113.0
        AS path: I, validation-state: unverified
        > to 198.51.100.12 via ae12.0, label-switched-path pe2tope1
1:198.51.100.1:1::111111111111111111::0/304
    *[BGP/170] 01:35:11, localpref 100, from 203.0.113.0
        AS path: I, validation-state: unverified
        > to 198.51.100.12 via ae12.0, label-switched-path pe2tope1
1:198.51.100.1:1::222222222222222222::0/304
    *[BGP/170] 01:35:11, localpref 100, from 203.0.113.0
        AS path: I, validation-state: unverified
        > to 198.51.100.12 via ae12.0, label-switched-path pe2tope1
1:198.51.100.1:1::333333333333333333::0/304
    *[BGP/170] 01:35:11, localpref 100, from 203.0.113.0
        AS path: I, validation-state: unverified
        > to 198.51.100.12 via ae12.0, label-switched-path pe2tope1
1:198.51.100.2:1::111111111111111111::0/304
    *[EVPN/170] 01:35:11
        Indirect
1:198.51.100.2:1::222222222222222222::0/304
    *[EVPN/170] 01:35:11
        Indirect
1:198.51.100.2:1::333333333333333333::0/304
    *[EVPN/170] 01:35:11
        Indirect
1:198.51.100.3:0::111111111111111111::FFFF:FFFF/304
    *[BGP/170] 01:35:11, localpref 100, from 203.0.113.0
        AS path: I, validation-state: unverified
        > to 198.51.100.7 via ge-0/0/4.0, label-switched-path pe2tope3
1:198.51.100.3:1::111111111111111111::0/304
    *[BGP/170] 01:35:11, localpref 100, from 203.0.113.0
        AS path: I, validation-state: unverified
        > to 198.51.100.7 via ge-0/0/4.0, label-switched-path pe2tope3
1:198.51.100.4:0::444444444444444444::FFFF:FFFF/304
    *[BGP/170] 01:35:02, localpref 100, from 203.0.113.0
        AS path: I, validation-state: unverified
        > to 198.51.100.8 via ge-0/1/8.0, label-switched-path pe2tope4
1:198.51.100.4:1::444444444444444444::0/304

```



```

* [BGP/170] 01:35:02, localpref 100, from 203.0.113.0
  AS path: I, validation-state: unverified
> to 198.51.100.8 via ge-0/1/8.0, label-switched-path pe2tope4
3:198.51.100.1:1::10::198.51.100.1/304
* [BGP/170] 01:35:11, localpref 100, from 203.0.113.0
  AS path: I, validation-state: unverified
> to 198.51.100.12 via ae12.0, label-switched-path pe2tope1
3:198.51.100.2:1::10::198.51.100.2/304
* [EVPN/170] 01:36:59
  Indirect
3:198.51.100.3:1::10::198.51.100.3/304
* [BGP/170] 01:35:11, localpref 100, from 203.0.113.0
  AS path: I, validation-state: unverified
> to 198.51.100.7 via ge-0/0/4.0, label-switched-path pe2tope3
3:198.51.100.4:1::10::198.51.100.4/304
* [BGP/170] 01:35:02, localpref 100, from 203.0.113.0
  AS path: I, validation-state: unverified
> to 198.51.100.8 via ge-0/1/8.0, label-switched-path pe2tope4

```

Router PE3

From operational mode, run the **show route table mhevpn** command.

```

user@PE3> show route table mhevpn
mhevpn.evpn.0: 19 destinations, 19 routes (19 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

1:198.51.100.1:0::111111111111111111::FFFF:FFFF/304
* [BGP/170] 01:36:10, localpref 100, from 203.0.113.0
  AS path: I, validation-state: unverified
> to 198.51.100.13 via ae13.0, label-switched-path pe3tope1
1:198.51.100.1:0::222222222222222222::FFFF:FFFF/304
* [BGP/170] 01:36:10, localpref 100, from 203.0.113.0
  AS path: I, validation-state: unverified
> to 198.51.100.13 via ae13.0, label-switched-path pe3tope1
1:198.51.100.1:0::333333333333333333::FFFF:FFFF/304
* [BGP/170] 01:36:10, localpref 100, from 203.0.113.0
  AS path: I, validation-state: unverified
> to 198.51.100.13 via ae13.0, label-switched-path pe3tope1
1:198.51.100.1:1::111111111111111111::0/304
* [BGP/170] 01:36:10, localpref 100, from 203.0.113.0
  AS path: I, validation-state: unverified
> to 198.51.100.13 via ae13.0, label-switched-path pe3tope1
1:198.51.100.1:1::222222222222222222::0/304
* [BGP/170] 01:36:10, localpref 100, from 203.0.113.0
  AS path: I, validation-state: unverified
> to 198.51.100.13 via ae13.0, label-switched-path pe3tope1
1:198.51.100.1:1::333333333333333333::0/304
* [BGP/170] 01:36:10, localpref 100, from 203.0.113.0
  AS path: I, validation-state: unverified
> to 198.51.100.13 via ae13.0, label-switched-path pe3tope1
1:198.51.100.2:0::111111111111111111::FFFF:FFFF/304
* [BGP/170] 01:36:06, localpref 100, from 203.0.113.0
  AS path: I, validation-state: unverified
> to 198.51.100.10 via ge-2/0/3.0, label-switched-path pe3tope2
1:198.51.100.2:0::222222222222222222::FFFF:FFFF/304
* [BGP/170] 01:36:06, localpref 100, from 203.0.113.0
  AS path: I, validation-state: unverified
> to 198.51.100.10 via ge-2/0/3.0, label-switched-path pe3tope2
1:198.51.100.2:0::333333333333333333::FFFF:FFFF/304
* [BGP/170] 01:36:06, localpref 100, from 203.0.113.0

```

```

        AS path: I, validation-state: unverified
        > to 198.51.100.10 via ge-2/0/3.0, label-switched-path pe3tope2
1:198.51.100.2:1::11111111111111111111::0/304
        *[BGP/170] 01:36:06, localpref 100, from 203.0.113.0
        AS path: I, validation-state: unverified
        > to 198.51.100.10 via ge-2/0/3.0, label-switched-path pe3tope2
1:198.51.100.2:1::22222222222222222222::0/304
        *[BGP/170] 01:36:06, localpref 100, from 203.0.113.0
        AS path: I, validation-state: unverified
        > to 198.51.100.10 via ge-2/0/3.0, label-switched-path pe3tope2
1:198.51.100.2:1::33333333333333333333::0/304
        *[BGP/170] 01:36:06, localpref 100, from 203.0.113.0
        AS path: I, validation-state: unverified
        > to 198.51.100.10 via ge-2/0/3.0, label-switched-path pe3tope2
1:198.51.100.3:1::11111111111111111111::0/304
        *[EVPN/170] 01:36:25
        Indirect
1:198.51.100.4:0::44444444444444444444::FFFF:FFFF/304
        *[BGP/170] 01:35:58, localpref 100, from 203.0.113.0
        AS path: I, validation-state: unverified
        > to 198.51.100.11 via ge-2/1/3.0, label-switched-path pe3tope4
1:198.51.100.4:1::44444444444444444444::0/304
        *[BGP/170] 01:35:58, localpref 100, from 203.0.113.0
        AS path: I, validation-state: unverified
        > to 198.51.100.11 via ge-2/1/3.0, label-switched-path pe3tope4
3:198.51.100.1:1::10::198.51.100.1/304
        *[BGP/170] 01:36:10, localpref 100, from 203.0.113.0
        AS path: I, validation-state: unverified
        > to 198.51.100.13 via ae13.0, label-switched-path pe3tope1
3:198.51.100.2:1::10::198.51.100.2/304
        *[BGP/170] 01:36:06, localpref 100, from 203.0.113.0
        AS path: I, validation-state: unverified
        > to 198.51.100.10 via ge-2/0/3.0, label-switched-path pe3tope2
3:198.51.100.3:1::10::198.51.100.3/304
        *[EVPN/170] 01:37:33
        Indirect
3:198.51.100.4:1::10::198.51.100.4/304
        *[BGP/170] 01:35:58, localpref 100, from 203.0.113.0
        AS path: I, validation-state: unverified
        > to 198.51.100.11 via ge-2/1/3.0, label-switched-path pe3tope4

```

Meaning The output displays the local and remote IRB interfaces. It also displays the local and remote hosts that are installed in the VRF table.

Release History Table

Release	Description
16.1R4	Starting with Junos OS Release 16.1R4, EVPN multihoming active-active mode is supported on all EX9200 switches.

Related Documentation

Example: Configuring an ESI on a Logical Interface With EVPN Multihoming

When a customer edge (CE) device in an Ethernet VPN-Multiprotocol Label Switching (EVPN-MPLS) environment is multihomed to two or more provider edge (PE) devices, the set of Ethernet links that connect the devices comprise an Ethernet segment. An Ethernet segment identifier (ESI) is a 10-octet integer that identifies this segment. A sample ESI is 00:11:22:33:44:55:66:77:88:99.

In releases before Junos OS Release 15.1F6 and 17.1R1 for MX Series routers and Junos OS Release 17.3R1 for EX9200 switches, you can specify an ESI only on a physical or aggregated Ethernet interface, for example, **set interfaces ae0 esi 00:11:22:33:44:55:66:77:88:99**. If you specify an ESI on a physical or aggregated Ethernet interface, keep in mind that an ESI is a factor in the designated forwarder (DF) election process. For example, assume that you configure EVPN multihoming active-standby on aggregated Ethernet interface ae0, and given the ESI configured on ae0 and other determining factors, the DF election results in ae0 being in the down state. Further, all logical interfaces configured on ae0, for example, **set interfaces ae0 unit 1** and **set interfaces ae0 unit 2** are also in the down state, which renders logical interfaces ae0.1 and ae0.2 unable to provide services to their respective customer sites (VLANs).

Starting with Junos OS Releases 15.1F6 and 17.1R1 for MX Series routers and Junos OS Release 17.3R1 for EX9200 switches, you can specify an ESI on a logical interface. If you specify an ESI on a logical interface, the DF election process now occurs at the individual logical interface level, which enables you to better utilize logical interfaces. For example, assume that you configure logical interfaces ae0.1 and ae0.2 on aggregated Ethernet interface ae0. You configure EVPN multihoming active-standby on both logical interfaces and given the ESI configured on ae0.1 and other determining factors, the DF election results in ae0.1 being in the down state. Despite logical interface ae0.1 being down, logical interface ae0.2 and other logical interfaces configured on ae0 can be in the up state and provide services to their respective customer sites (VLANs).

This topic shows how to configure an ESI on logical interfaces in both EVPN multihoming active-standby and active-active modes.

- [Requirements on page 115](#)
- [Overview and Topology on page 116](#)
- [EVPN Multihoming Active-Standby Configuration on page 118](#)
- [Verification on page 125](#)

Requirements

Both EVPN multihoming active-standby and multihoming active-active examples use the following hardware and software components:

- An EX9200 switch running Junos OS Release 17.3R1 or later (PE1)
- An MX Series router running Junos OS Release 15.1F6 or later, or Junos OS Release 17.1R1 or later (PE2)

Overview and Topology

EVPN Multihoming Active-Standby

Figure 6 on page 116 shows an EVPN-MPLS topology in which CE1 is multihomed to PE1 and PE2 to provide redundant paths to CE2. On CE1, the connections to PE1 and PE2 are configured as separate aggregated Ethernet interfaces. Table 4 on page 116 shows how the connections with CE1 are configured on PE1 and PE2. Note that the EVPN multihoming mode, ESI, and VLAN ID are actually configured on logical interfaces ae0.1 on each PE device.

Figure 6: EVPN-MPLS Topology with Multihoming Active-Standby

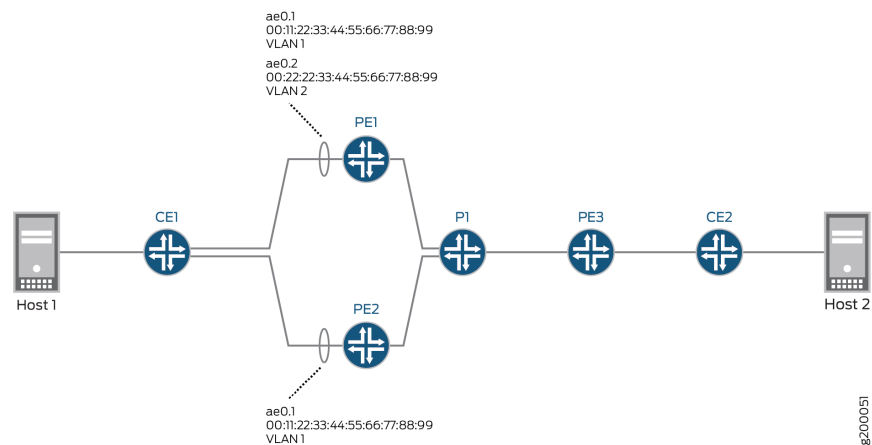


Table 4: EVPN Multihoming Active-Standby: Configuring the Connection with CE1 on PE1 and PE2

Device	Physical Interface	Aggregated Ethernet Interface	Logical Interface	EVPN Multihoming Mode	ESI	VLAN ID
PE1	xe-2/0/0	ae0	ae0.1	single-active	00:11:22:33:44:55:66:77:88:99	1
PE2	xe-3/0/2	ae0	ae0.1	single-active	00:11:22:33:44:55:66:77:88:99	1

Based on the DF election, logical interface ae0.1 on PE2 is up, and logical interface ae0.1 on PE1 is down.

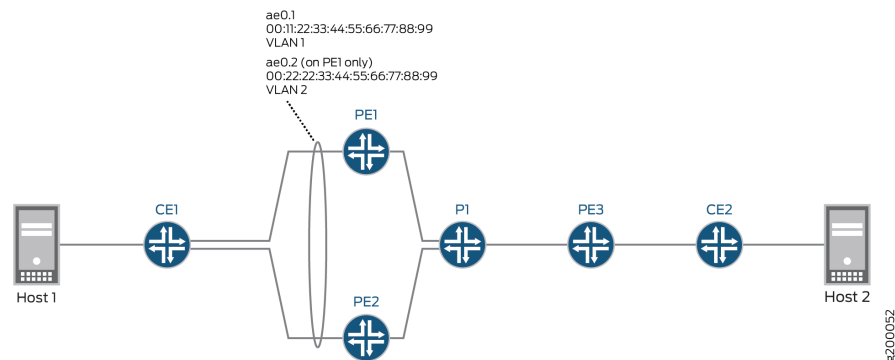
Table 5 on page 117 also shows the configuration for logical interface ae0.2 on PE1. Note that logical interface ae0.2 provides services for a different VLAN and is configured with a different ESI than logical interface ae0.1, which is configured on the same aggregated Ethernet interface. As a result, logical interface ae0.2 is up and providing services to VLAN 2 despite the fact that logical interface ae0.1 is in the down state.

Table 5: Multihoming Active-Standby: Configuring Logical Interface on PE1 that Provides Services to a Different VLAN

Device	Physical Interface	Aggregated Ethernet Interface	Logical Interface	EVPN Multihoming Mode	ESI	VLAN ID
PE1	xe-2/0/0	ae0	ae0.2	single-active	00:22:33:44:55:66:77:88:99	2

EVPN Multihoming Active-Active

Figure 7 on page 117 shows an EVPN-MPLS topology in which CE1 is multihomed to PE1 and PE2 to provide redundant paths to CE2. On CE1, the connections to PE1 and PE2 are configured as one aggregated Ethernet interface. Table 6 on page 117 shows how the connections with CE1 are configured on PE1 and PE2. Note that the EVPN multihoming mode, ESI, and VLAN ID are actually configured on logical interfaces ae0.1 on each device.

Figure 7: EVPN-MPLS Topology with Multihoming Active-Active**Table 6: Multihoming Active-Active: Configuring the Connection with CE1 on PE1 and PE2**

Device	Physical Interface	Aggregated Ethernet Interface	Logical Interface	EVPN Multihoming Mode	ESI	VLAN ID
PE1	xe-2/0/0	ae0	ae0.1	all-active	00:11:22:33:44:55:66:77:88:99	1
PE2	xe-3/0/2	ae0	ae0.1	all-active	00:11:22:33:44:55:66:77:88:99	1

Based on the DF election, logical interface ae0.1 on PE1 is in the up state, and logical interface ae0.1 on PE2 is in the down state.

Table 7 on page 118 also shows the configuration for logical interface ae0.2 on PE1. Note that logical interface ae0.2 provides services for a different VLAN and is configured with a different ESI than logical interface ae0.1, which is in the same aggregated Ethernet interface. As a result, logical interface ae0.2 is down and unable to provide services to VLAN 2 despite the fact that logical interface ae0.1 is in the up state.

Table 7: EVPN Multihoming Active-Active: Configuring Interface on PE1 that Provides Services to a Different VLAN

Device	Physical Interface	Aggregated Ethernet Interface	Logical Interface	EVPN Multihoming Mode	ESI	VLAN ID
PE1	xe-2/0/0	ae0	ae0.2	all-active	00:22:22:33:44:55:66:77:88:99	2

EVPN Multihoming Active-Standby Configuration



NOTE: The configurations for PE1 (EX9200) and PE2 (MX Series router) focus on configuring EVPN multihoming active-standby and ESIs on logical interfaces. The configurations do not include all EVPN-related configurations for physical interfaces, aggregated Ethernet interfaces, logical interfaces, and routing instances. For a more comprehensive configuration of EVPN multihoming active-standby in an EVPN-MPLS environment, see *Example: Configuring EVPN Active-Standby Multihoming*. Note that the referenced example shows how to configure ESIs on physical and aggregated Ethernet interfaces.

CLI Configuration

```
set interfaces xe-2/0/0 gigether-options 802.3ad ae0
set interfaces ae0 flexible-vlan-tagging
set interfaces ae0 encapsulation flexible-ethernet-services
set interfaces ae0 unit 1 encapsulation vlan-bridge
set interfaces ae0 unit 1 vlan-id 1
set interfaces ae0 unit 1 esi 00:11:22:33:44:55:66:77:88:99
set interfaces ae0 unit 1 esi single-active
set interfaces ae0 unit 2 encapsulation vlan-bridge
set interfaces ae0 unit 2 vlan-id 2
set interfaces ae0 unit 2 esi 00:22:22:33:44:55:66:77:88:99
set interfaces ae0 unit 2 esi single-active
set interfaces irb unit 1 family inet address 192.0.2.1/24
set interfaces irb unit 2 family inet address 192.0.2.2/24
set routing-instances blue instance-type evpn
set routing-instances blue vlan-id 1
set routing-instances blue interface ae0.1
set routing-instances blue l3-interface irb.1
...
set routing-instances blue protocols evpn interface ae0.1
set routing-instances green instance-type evpn
set routing-instances green vlan-id 2
set routing-instances green interface ae0.2
set routing-instances green l3-interface irb.2
...
set routing-instances green protocols evpn interface ae0.2
set routing-instances vrf instance-type vrf
set routing-instances vrf interface irb.1
set routing-instances vrf interface irb.2
...
```

```

PE2    set interfaces xe-3/0/2 gigether-options 802.3ad ae0
       set interfaces ae0 flexible-vlan-tagging
       set interfaces ae0 encapsulation flexible-ethernet-services
       set interfaces ae0 unit 1 encapsulation vlan-bridge
       set interfaces ae0 unit 1 vlan-id 1
       set interfaces ae0 unit 1 esi 00:11:22:33:44:55:66:77:88:99
       set interfaces ae0 unit 1 esi single-active
       set interfaces irb unit 1 family inet address 192.0.2.1/24
       set routing-instances blue instance-type evpn
       set routing-instances blue vlan-id 1
       set routing-instances blue interface ae0.1
       set routing-instances blue routing-interface irb.1
       ...
       set routing-instances blue protocols evpn interface ae0.1
       set routing-instances vrf instance-type vrf
       set routing-instances vrf interface irb.1
       ...

```

Step-by-Step Procedure

To configure EVPN multihoming active-standby on PE1:

1. Specify an Ethernet interface as a member of aggregated Ethernet interface ae0.

```

[edit interfaces]
user@switch# set xe-2/0/0 gigether-options 802.3ad ae0

```

2. Configure aggregated Ethernet interface ae0 to simultaneously transmit 802.1Q VLAN single-tag and dual-tag frames and to support different types of Ethernet encapsulation at the logical interface level.

```

[edit interfaces]
user@switch# set ae0 flexible-vlan-tagging
user@switch# set ae0 encapsulation flexible-ethernet-services

```

3. On aggregated Ethernet interface ae0, configure logical interfaces ae0.1 and ae0.2. Configure the logical interfaces to use VLAN bridge encapsulation, and map the logical interfaces to VLANs 1 and 2, respectively. Also, assign an ESI to the logical interfaces, and enable EVPN multihoming active-standby.

```

[edit interfaces]
user@switch# set ae0 unit 1 encapsulation vlan-bridge
user@switch# set ae0 unit 1 vlan-id 1
user@switch# set ae0 unit 1 esi 00:11:22:33:44:55:66:77:88:99
user@switch# set ae0 unit 1 esi single-active
user@switch# set ae0 unit 2 encapsulation vlan-bridge
user@switch# set ae0 unit 2 vlan-id 2
user@switch# set ae0 unit 2 esi 00:22:22:33:44:55:66:77:88:99
user@switch# set ae0 unit 2 esi single-active

```

4. Configure IRB interfaces irb.1 and irb.2, and assign an IP address to each interface.

```

[edit interfaces]
set interfaces irb unit 1 family inet address 192.0.2.1/24
set interfaces irb unit 2 family inet address 192.0.2.2/24

```

5. Configure an EVPN routing instance named blue. Map the routing instance to VLAN 1, logical interface ae0.1, and IRB interface irb.1. Configure EVPN logical interface ae0.1 for the EVPN routing instance.

```
[edit routing-instances]
user@switch# set blue instance-type evpn
user@switch# set blue vlan-id 1
user@switch# set blue interface ae0.1
user@switch# set blue l3-interface irb.1
user@switch# set blue protocols evpn interface ae0.1
```

6. Configure an EVPN routing instance named green. Map the routing instance to VLAN 2, logical interface ae0.2, and IRB interface irb.2. Configure logical interface ae0.2 for the EVPN routing instance.

```
[edit routing-instances]
user@switch# set green instance-type evpn
user@switch# set green vlan-id 2
user@switch# set green interface ae0.2
user@switch# set green l3-interface irb.2
user@switch# set green protocols evpn interface ae0.2
```

7. Configure a VRF routing instance, and add IRB interfaces irb.1 and irb.2 to the routing instance.

```
[edit routing-instances]
set vrf instance-type vrf
set vrf interface irb.1
set vrf interface irb.2
```

Step-by-Step Procedure

To configure EVPN multihoming active-standby on PE2:

1. Specify an Ethernet interface as a member of aggregated Ethernet interface ae0.

```
[edit interfaces]
user@router# set xe-3/0/2 gigether-options 802.3ad ae0
```

2. Configure aggregated Ethernet interface ae0 to simultaneously transmit 802.1Q VLAN single-tag and dual-tag frames and to support different types of Ethernet encapsulation at the logical interface level.

```
[edit interfaces]
user@router# set ae0 flexible-vlan-tagging
user@router# set ae0 encapsulation flexible-ethernet-services
```

3. On aggregated Ethernet interface ae0, configure logical interface ae0.1. Configure the logical interface to use VLAN bridge encapsulation, and map the logical interface to VLAN 1 and 2. Also, assign an ESI to the logical interface, and enable EVPN multihoming active-standby.

```
[edit interfaces]
user@router# set ae0 unit 1 encapsulation vlan-bridge
user@router# set ae0 unit 1 vlan-id 1
```



```
user@router# set ae0 unit 1 esi 00:11:22:33:44:55:66:77:88:99
user@router# set ae0 unit 1 esi single-active
```

4. Configure IRB interface irb.1, and assign an IP address to the interface.

```
[edit interfaces]
user@router# set irb unit 1 family inet address 192.0.2.1/24
```

5. Configure an EVPN routing instance named blue. Map the routing instance to VLAN 1, logical interface ae0.1, and IRB interface irb.1. Configure EVPN logical interface ae0.1 for the EVPN routing instance.

```
[edit routing-instances]
user@router# set blue instance-type evpn
user@router# set blue vlan-id 1
user@router# set blue interface ae0.1
user@router# set blue routing-interface irb.1
user@router# set blue protocols evpn interface ae0.1
```

6. Configure a VRF routing instance, and add IRB interface irb.1 to the routing instance.

```
[edit routing-instances]
user@router# set vrf instance-type vrf
user@router# set vrf interface irb.1
```

EVPN Multihoming Active-Active Configuration

CLI Quick Configuration



NOTE: The configurations for PE1 (EX9200) and PE2 (MX Series router) focus on configuring EVPN multihoming active-active and ESIs on logical interfaces. The configurations do not include all EVPN-related configurations for physical interfaces, aggregated Ethernet interfaces, logical interfaces, and routing instances. For a more comprehensive configuration of EVPN multihoming active-active, see [“Example: Configuring EVPN Active-Active Multihoming” on page 67](#). Note that the referenced example shows how to configure ESIs on physical and aggregated Ethernet interfaces.

```
PE1 set interfaces xe-2/0/0 gigether-options 802.3ad ae0
set interfaces ae0 flexible-vlan-tagging
set interfaces ae0 encapsulation flexible-ethernet-services
set interfaces ae0 unit 1 encapsulation vlan-bridge
set interfaces ae0 unit 1 vlan-id 1
set interfaces ae0 unit 1 esi 00:11:22:33:44:55:66:77:88:99
set interfaces ae0 unit 1 esi all-active
set interfaces ae0 unit 2 encapsulation vlan-bridge
set interfaces ae0 unit 2 vlan-id 2
set interfaces ae0 unit 2 esi 00:22:22:33:44:55:66:77:88:99
set interfaces ae0 unit 2 esi all-active
set interfaces irb unit 1 family inet address 192.0.2.1/24
```

```
set interfaces irb unit 2 family inet address 192.0.2.2/24
set routing-instances blue instance-type evpn
set routing-instances blue vlan-id 1
set routing-instances blue interface ae0.1
set routing-instances blue l3-interface irb.1
...
set routing-instances blue protocols evpn interface ae0.1
set routing-instances green instance-type evpn
set routing-instances green vlan-id 2
set routing-instances green interface ae0.2
set routing-instances green l3-interface irb.2
...
set routing-instances green protocols evpn interface ae0.2
set routing-instances vrf instance-type vrf
set routing-instances vrf interface irb.1
set routing-instances vrf interface irb.2
...
```

```
PE2 set interfaces xe-3/0/2 gigether-options 802.3ad ae0
set interfaces ae0 flexible-vlan-tagging
set interfaces ae0 encapsulation flexible-ethernet-services
set interfaces ae0 unit 1 encapsulation vlan-bridge
set interfaces ae0 unit 1 vlan-id 1
set interfaces ae0 unit 1 esi 00:11:22:33:44:55:66:77:88:99
set interfaces ae0 unit 1 esi all-active
set interfaces irb unit 1 family inet address 192.0.2.1/24
set routing-instances blue instance-type evpn
set routing-instances blue vlan-id 1
set routing-instances blue interface ae0.1
set routing-instances blue routing-interface irb.1
...
set routing-instances blue protocols evpn interface ae0.1
set routing-instances vrf instance-type vrf
set routing-instances vrf interface irb.1
...
```

**Step-by-Step
Procedure**

To configure EVPN multihoming active-active on PE1:

1. Specify an Ethernet interface as a member of aggregated Ethernet interface ae0.

```
[edit interfaces]
```

```
user@switch# set xe-2/0/0 gigether-options 802.3ad ae0
```

2. Configure aggregated Ethernet interface ae0 to simultaneously transmit 802.1Q VLAN single-tag and dual-tag frames and to support different types of Ethernet encapsulation at the logical interface level.

```
[edit interfaces]
```

```
user@switch# set ae0 flexible-vlan-tagging
```

```
user@switch# set ae0 encapsulation flexible-ethernet-services
```

3. On aggregated Ethernet interface ae0, configure logical interfaces ae0.1 and ae0.2. Configure the logical interfaces to use VLAN bridge encapsulation, and map the logical interfaces to VLANs 1 and 2, respectively. Also, assign an ESI to the logical interfaces, and enable EVPN multihoming active-active.

```
[edit interfaces]
user@switch# set ae0 unit 1 encapsulation vlan-bridge
user@switch# set ae0 unit 1 vlan-id 1
user@switch# set ae0 unit 1 esi 00:11:22:33:44:55:66:77:88:99
user@switch# set ae0 unit 1 esi all-active
user@switch# set ae0 unit 2 encapsulation vlan-bridge
user@switch# set ae0 unit 2 vlan-id 2
user@switch# set ae0 unit 2 esi 00:22:22:33:44:55:66:77:88:99
user@switch# set ae0 unit 2 esi all-active
```

4. Configure IRB interfaces irb.1 and irb.2, and assign an IP address to each interface.

```
[edit interfaces]
set interfaces irb unit 1 family inet address 192.0.2.1/24
set interfaces irb unit 2 family inet address 192.0.2.2/24
```

5. Configure an EVPN routing instance named blue. Map the routing instance to VLAN 1, logical interface ae0.1, and IRB interface irb.1. Configure logical interface ae0.1 for the EVPN routing instance.

```
[edit routing-instances]
user@switch# set blue instance-type evpn
user@switch# set blue vlan-id 1
user@switch# set blue interface ae0.1
user@switch# set blue l3-interface irb.1
user@switch# set blue protocols evpn interface ae0.1
```

6. Configure an EVPN routing instance named green. Map the routing instance to VLAN 2, logical interface ae0.2, and IRB interface irb.2. Configure logical interface ae0.2 for the EVPN routing instance.

```
[edit routing-instances]
user@switch# set green instance-type evpn
user@switch# set green vlan-id 2
user@switch# set green interface ae0.2
user@switch# set green l3-interface irb.2
user@switch# set green protocols evpn interface ae0.2
```

7. Configure a VRF routing instance, and add IRB interfaces irb.1 and irb.2 to the routing instance.

```
[edit routing-instances]
set vrf instance-type vrf
set vrf interface irb.1
set vrf interface irb.2
```

**Step-by-Step
Procedure**

To configure EVPN multihoming active-active on PE2:

1. Specify Ethernet interface xe-3/0/2 as a member of aggregated Ethernet interface ae0.

```
[edit interfaces]
user@router# set xe-3/0/2 aggregated-options 802.3ad ae0
```

2. Configure aggregated Ethernet interface ae0 to simultaneously transmit 802.1Q VLAN single-tag and dual-tag frames and to support different types of Ethernet encapsulation at the logical interface level.

```
[edit interfaces]
user@router# set ae0 flexible-vlan-tagging
user@router# set ae0 encapsulation flexible-ethernet-services
```

3. On aggregated Ethernet interface ae0, configure logical interface ae0.1. Configure the logical interface to use VLAN bridge encapsulation, and map the logical interface to VLANs 1. Also, assign an ESI to the logical interface, and enable EVPN multihoming active-active.

```
[edit interfaces]
user@router# set ae0 unit 1 encapsulation vlan-bridge
user@router# set ae0 unit 1 vlan-id 1
user@router# set ae0 unit 1 esi 00:11:22:33:44:55:66:77:88:99
user@router# set ae0 unit 1 esi all-active
```

4. Configure IRB interface irb.1, and assign an IP address to the interface.

```
[edit interfaces]
user@router# set irb unit 1 family inet address 192.0.2.1/24
```

5. Configure an EVPN routing instance named blue. Map the routing instance to VLAN 1, logical interface ae0.1, and IRB interface irb.1. Configure logical interface ae0.1 for the EVPN routing instance.

```
[edit routing-instances]
user@router# set blue instance-type evpn
user@router# set blue vlan-id 1
user@router# set blue interface ae0.1
user@router# set blue routing-interface irb.1
user@router# set blue protocols evpn interface ae0.1
```

6. Configure a VRF routing instance, and add IRB interface irb.1 to the routing instance.

```
[edit routing-instances]
user@router# set vrf instance-type vrf
user@router# set vrf interface irb.1
```

Verification

Verifying that a Logical Interface Has Correct ESI and EVPN Multihoming Mode

Purpose Verify that logical interface ae0.1 is configured with the correct ESI and EVPN multihoming mode.

Action From operational mode, enter the **show interfaces ae0.1** command.

```
user@switch> show interfaces ae0.1
  Logical interface ae0.1 (Index 332) (SNMP ifIndex 706)
  Flags: Up SNMP-Traps 0x20004000 VLAN-Tag [ 0x8100.100 ] Encapsulation:
VLAN-Bridge
  Input packets : 0
  Output packets: 0
  Protocol bridge, MTU: 1522
  Flags: Is-Primary
  Ethernet segment value: 00:11:22:33:44:55:66:77:88:99, Mode: Single-active
```

Meaning The output shows that logical interface ae0.1 is configured with ESI 00:11:22:33:44:55:66:77:88:99 and the EVPN multihoming mode is single-active for multihoming active-standby mode. For a scenario with multihoming active-active mode, the output would display all-active.

Verifying the EVPN Routing Instance Status

Purpose Verify the status of the various elements configured in an EVPN routing instance.

Action From operational mode, enter the **show evpn instance extensive** command.

```

user@switch> show evpn instance extensive
Instance: blue
...
Number of local interfaces: (1 up)
  Interface name  ESI                               Mode          Status
AC-Role
  ae0.1           00:11:22:33:44:55:066:77:88:99  single-active  Up
Root
Number of IRB interfaces: 1 (1 up)
  Interface name  VLAN ID  Status  L3 context
  irb.1           1        Up      vrf
Number of protect interfaces: 0
Number of bridge domains: 1
  VLAN  Domain ID  Intfs / up  IRB intf  Mode          MAC sync  IM route
  label SG sync  IM core nexthop
  1      1          0      0          Local switching
...
Number of ethernet segments: 1
ESI: 00:11:22:33:44:55:066:77:88:99
Status: Resolved by IFL ae0.1
Local interface: ae0.1, Status: Up/Forwarding
Number of remote PEs connected: 1
  Remote PE      MAC label  Aliasing label  Mode
  10.255.0.1      0          300928          all-active
DF Election Algorithm: MOD based
Designated forwarder: 10.255.0.1
Backup forwarder: 10.255.0.2
Last designated forwarder update: Jul 28 13:04:29
Advertised split horizon label: 300128

```

Meaning The output shows that for the EVPN routing instance named blue, the logical interface ae0.1 and IRB interface irb.1 are up and running. It also shows that VLAN 1 and Ethernet segment 00:11:22:33:44:55:066:77:88:99 are properly mapped to the routing instance. It also shows the status of the DF election.

Release History Table

Release	Description
15.1F6	Starting with Junos OS Releases 15.1F6 and 17.1R1 for MX Series routers and Junos OS Release 17.3R1 for EX9200 switches, you can specify an ESI on a logical interface.

Related Documentation

- [EVPN Multihoming Overview on page 39](#)

PART 3

Configuration Statements and Operational Commands



- Configuration Statements on page 129
- Operational Commands on page 165

CHAPTER 7

Configuration Statements

- [encapsulation \(Physical Interface\) on page 130](#)
- [esi on page 136](#)
- [evpn on page 138](#)
- [extended-vlan-list on page 140](#)
- [instance-type on page 141](#)
- [interface \(EVPN Routing Instances\) on page 144](#)
- [interface \(Routing Instances\) on page 145](#)
- [interface-mac-limit \(VPLS\) on page 146](#)
- [ip-prefix-routes on page 148](#)
- [mac-statistics on page 152](#)
- [mac-table-size on page 154](#)
- [no-arp-suppression on page 155](#)
- [no-mac-learning on page 156](#)
- [packet-action on page 158](#)
- [routing-instances on page 161](#)
- [traceoptions \(Protocols EVPN\) on page 162](#)
- [vlan-id \(routing instance\) on page 164](#)

encapsulation (Physical Interface)

Syntax	encapsulation (atm-ccc-cell-relay atm-pvc cisco-hdlc cisco-hdlc-ccc cisco-hdlc-tcc ethernet-bridge ethernet-ccc ethernet-over-atm ethernet-tcc ethernet-vpls ethernet-vpls-fr ether-vpls-over-atm-llc ethernet-vpls-ppp extended-frame-relay-ccc extended-frame-relay-ether-type-tcc extended-frame-relay-tcc extended-vlan-bridge extended-vlan-ccc extended-vlan-tcc extended-vlan-vpls flexible-ethernet-services flexible-frame-relay frame-relay frame-relay-ccc frame-relay-ether-type frame-relay-ether-type-tcc frame-relay-port-ccc frame-relay-tcc generic-services multilink-frame-relay-uni-nni ppp ppp-ccc ppp-tcc vlan-ccc vlan-vci-ccc vlan-vpls);
Hierarchy Level	[edit interfaces <i>interface-name</i>], [edit interfaces rlsq <i>number:number</i>]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 11.1 for EX Series switches. Statement introduced in Junos OS Release 12.1X48 for PTX Series Packet Transport Routers (flexible-ethernet-services , ethernet-ccc , and ethernet-tcc options only).
Description	Specify the physical link-layer encapsulation type.
	<div>  <p>NOTE: Not all encapsulation types are supported on the switches. See the switch CLI.</p> </div>
Default	ppp—Use serial PPP encapsulation.
Options	<div>  <p>NOTE: Frame Relay, ATM, PPP, SONET, and SATSOP options are not supported on the EX Series switches.</p> </div> <p>atm-ccc-cell-relay—Use ATM cell-relay encapsulation.</p> <p>atm-pvc—Defined in RFC 2684, <i>Multiprotocol Encapsulation over ATM Adaptation Layer 5</i>. When you configure physical ATM interfaces with ATM PVC encapsulation, an RFC 2684-compliant ATM Adaptation Layer 5 (AAL5) tunnel is set up to route the ATM cells over a Multiprotocol Label Switching (MPLS) path that is typically established between two MPLS-capable routers using the Label Distribution Protocol (LDP).</p> <p>cisco-hdlc—Use Cisco-compatible High-Level Data Link Control (HDLC) framing. E1, E3, SONET/SDH, T1, and T3 interfaces can use Cisco HDLC encapsulation. Two related versions are supported:</p>

- CCC version (**cisco-hdlc-ccc**)—The logical interface does not require an encapsulation statement. When you use this encapsulation type, you can configure the **ccc** family only.
- TCC version (**cisco-hdlc-tcc**)—Similar to CCC and has the same configuration restrictions, but used for circuits with different media on either side of the connection.

cisco-hdlc-ccc—Use Cisco-compatible HDLC framing on CCC circuits.

cisco-hdlc-tcc—Use Cisco-compatible HDLC framing on TCC circuits for connecting different media.

ethernet-bridge—Use Ethernet bridge encapsulation on Ethernet interfaces that have bridging enabled and that must accept all packets.

ethernet-ccc—Use Ethernet CCC encapsulation on Ethernet interfaces that must accept packets carrying standard Tag Protocol ID (TPID) values. For 8-port, 12-port, and 48-port Fast Ethernet PICs, CCC is not supported.

ethernet-over-atm—For interfaces that carry IPv4 traffic, use Ethernet over ATM encapsulation. When you use this encapsulation type, you cannot configure multipoint interfaces. As defined in RFC 2684, *Multiprotocol Encapsulation over ATM Adaptation Layer 5*, this encapsulation type allows ATM interfaces to connect to devices that support only bridge protocol data units (BPDUs). Junos OS does not completely support bridging, but accepts BPDU packets as a default gateway. If you use the router as an edge device, then the router acts as a default gateway. It accepts Ethernet LLC/SNAP frames with IP or ARP in the payload, and drops the rest. For packets destined to the Ethernet LAN, a route lookup is done using the destination IP address. If the route lookup yields a full address match, the packet is encapsulated with an LLC/SNAP and MAC header, and the packet is forwarded to the ATM interface.

ethernet-tcc—For interfaces that carry IPv4 traffic, use Ethernet TCC encapsulation on interfaces that must accept packets carrying standard TPID values. For 8-port, 12-port, and 48-port Fast Ethernet PICs, TCC is not supported.

ethernet-vpls—Use Ethernet VPLS encapsulation on Ethernet interfaces that have VPLS enabled and that must accept packets carrying standard TPID values. On M Series routers, except the M320 router, the 4-port Fast Ethernet TX PIC and the 1-port, 2-port, and 4-port, 4-slot Gigabit Ethernet PICs can use the Ethernet VPLS encapsulation type.

ethernet-vpls-fr—Use in a VPLS setup when a CE device is connected to a PE device over a time division multiplexing (TDM) link. This encapsulation type enables the PE device to terminate the outer Layer 2 Frame Relay connection, use the 802.1p bits inside the inner Ethernet header to classify the packets, look at the MAC address from the Ethernet header, and use the MAC address to forward the packet into a given VPLS instance.

ethernet-vpls-ppp—Use in a VPLS setup when a CE device is connected to a PE device over a time division multiplexing (TDM) link. This encapsulation type enables the PE device to terminate the outer Layer 2 PPP connection, use the 802.1p bits inside the inner Ethernet header to classify the packets, look at the MAC address from the Ethernet header, and use it to forward the packet into a given VPLS instance.

ether-vpls-over-atm-llc—For ATM intelligent queuing (IQ) interfaces only, use the Ethernet virtual private LAN service (VPLS) over ATM LLC encapsulation to bridge Ethernet interfaces and ATM interfaces over a VPLS routing instance (as described in RFC 2684, *Multiprotocol Encapsulation over ATM Adaptation Layer 5*). Packets from the ATM interfaces are converted to standard ENET2/802.3 encapsulated Ethernet frames with the frame check sequence (FCS) field removed.

extended-frame-relay-ccc—Use Frame Relay encapsulation on CCC circuits. This encapsulation type allows you to dedicate DLCIs 1 through 1022 to CCC. When you use this encapsulation type, you can configure the **ccc** family only.

extended-frame-relay-ether-type-tcc—Use extended Frame Relay ether type TCC for Cisco-compatible Frame Relay for DLCIs 1 through 1022. This encapsulation type is used for circuits with different media on either side of the connection.

extended-frame-relay-tcc—Use Frame Relay encapsulation on TCC circuits to connect different media. This encapsulation type allows you to dedicate DLCIs 1 through 1022 to TCC.

extended-vlan-bridge—Use extended VLAN bridge encapsulation on Ethernet interfaces that have IEEE 802.1Q VLAN tagging and bridging enabled and that must accept packets carrying TPID 0x8100 or a user-defined TPID.

extended-vlan-ccc—Use extended VLAN encapsulation on CCC circuits with Gigabit Ethernet and 4-port Fast Ethernet interfaces that must accept packets carrying 802.1Q values. Extended VLAN CCC encapsulation supports TPIDs 0x8100, 0x9100, and 0x9901. When you use this encapsulation type, you can configure the **ccc** family only. For 8-port, 12-port, and 48-port Fast Ethernet PICs, extended VLAN CCC is not supported. For 4-port Gigabit Ethernet PICs, extended VLAN CCC is not supported.

extended-vlan-tcc—For interfaces that carry IPv4 traffic, use extended VLAN encapsulation on TCC circuits with Gigabit Ethernet interfaces on which you want to use 802.1Q tagging. For 4-port Gigabit Ethernet PICs, extended VLAN TCC is not supported.

extended-vlan-vpls—Use extended VLAN VPLS encapsulation on Ethernet interfaces that have VLAN 802.1Q tagging and VPLS enabled and that must accept packets carrying TPIDs 0x8100, 0x9100, and 0x9901. On M Series routers, except the M320 router, the 4-port Fast Ethernet TX PIC and the 1-port, 2-port, and 4-port, 4-slot Gigabit Ethernet PICs can use the Ethernet VPLS encapsulation type.



NOTE: The built-in Gigabit Ethernet PIC on an M7i router does not support extended VLAN VPLS encapsulation.

flexible-ethernet-services—For Gigabit Ethernet IQ interfaces and Gigabit Ethernet PICs with small form-factor pluggable transceivers (SFPs) (except the 10-port Gigabit Ethernet PIC and the built-in Gigabit Ethernet port on the M7i router), and for Gigabit Ethernet interfaces, use flexible Ethernet services encapsulation when you want to configure multiple per-unit Ethernet encapsulations. Aggregated Ethernet bundles can use this encapsulation type. This encapsulation type allows you to configure any combination of route, TCC, CCC, Layer 2 virtual private networks (VPNs), and VPLS encapsulations on a single physical port. If you configure flexible Ethernet services encapsulation on the physical interface, VLAN IDs from 1 through 511 are no longer reserved for normal VLANs.

flexible-frame-relay—For IQ interfaces only, use flexible Frame Relay encapsulation when you want to configure multiple per-unit Frame Relay encapsulations. This encapsulation type allows you to configure any combination of TCC, CCC, and standard Frame Relay encapsulations on a single physical port. Also, each logical interface can have any DLCI value from 1 through 1022.

frame-relay—Use Frame Relay encapsulation is defined in RFC 1490, *Multiprotocol Interconnect over Frame Relay*. E1, E3, link services, SONET/SDH, T1, T3, and voice services interfaces can use Frame Relay encapsulation.

frame-relay-ccc—Use Frame Relay encapsulation on CCC circuits. This encapsulation is same as standard Frame Relay for DLCIs 0 through 511. DLCIs 512 through 1022 are dedicated to CCC. The logical interface must also have **frame-relay-ccc** encapsulation. When you use this encapsulation type, you can configure the **ccc** family only.

frame-relay-ether-type—Use Frame Relay ether type encapsulation for compatibility with the Cisco Frame Relay. IETF frame relay encapsulation identifies the payload format using NLPID and SNAP formats. Cisco-compatible Frame Relay encapsulation uses the Ethernet type to identify the type of payload.



NOTE: When the encapsulation type is set to Cisco-compatible Frame Relay encapsulation, ensure that the LMI type is set to ANSI or Q933-A.

frame-relay-ether-type-tcc—Use Frame Relay ether type TCC for Cisco-compatible Frame Relay on TCC circuits to connect different media. This encapsulation is Cisco-compatible Frame Relay for DLCIs 0 through 511. DLCIs 512 through 1022 are dedicated to TCC.

frame-relay-port-ccc—Use Frame Relay port CCC encapsulation to transparently carry all the DLCIs between two customer edge (CE) routers without explicitly configuring each DLCI on the two provider edge (PE) routers with Frame Relay transport. The connection between the two CE routers can be either user-to-network interface (UNI) or network-to-network interface (NNI); this is completely transparent to the PE routers. When you use this encapsulation type, you can configure the **ccc** family only.

frame-relay-tcc—This encapsulation is similar to Frame Relay CCC and has the same configuration restrictions, but used for circuits with different media on either side of the connection.

generic-services—Use generic services encapsulation for services with a hierarchical scheduler.

multilink-frame-relay-uni-nni—Use MLFR UNI NNI encapsulation. This encapsulation is used on link services, voice services interfaces functioning as FRF.16 bundles, and their constituent T1 or E1 interfaces, and is supported on LSQ and redundant LSQ interfaces.

ppp—Use serial PPP encapsulation. This encapsulation is defined in RFC 1661, *The Point-to-Point Protocol (PPP) for the Transmission of Multiprotocol Datagrams over Point-to-Point Links*. PPP is the default encapsulation type for physical interfaces. E1, E3, SONET/SDH, T1, and T3 interfaces can use PPP encapsulation.

ppp-ccc—Use serial PPP encapsulation on CCC circuits. When you use this encapsulation type, you can configure the **ccc** family only.

ppp-tcc—Use serial PPP encapsulation on TCC circuits for connecting different media. When you use this encapsulation type, you can configure the **tcc** family only.

vlan-ccc—Use Ethernet VLAN encapsulation on CCC circuits. VLAN CCC encapsulation supports TPID 0x8100 only. When you use this encapsulation type, you can configure the **ccc** family only.

vlan-vci-ccc—Use ATM-to-Ethernet interworking encapsulation on CCC circuits. When you use this encapsulation type, you can configure the **ccc** family only. All logical interfaces configured on the Ethernet interface must also have the encapsulation type set to **vlan-vci-ccc**.

vlan-vpls—Use VLAN VPLS encapsulation on Ethernet interfaces with VLAN tagging and VPLS enabled. Interfaces with VLAN VPLS encapsulation accept packets carrying standard TPID values only. On M Series routers, except the M320 router, the 4-port Fast Ethernet TX PIC and the 1-port, 2-port, and 4-port, 4-slot Gigabit Ethernet PICs can use the Ethernet VPLS encapsulation type.



NOTE:

- Label-switched interfaces (LSIs) do not support VLAN VPLS encapsulation. Therefore, you can only use VLAN VPLS encapsulation on a PE-router-to-CE-router interface and not a core-facing interface.
 - Starting with Junos OS release 13.3, a commit error occurs when you configure **vlan-vpls** encapsulation on a physical interface and configure family **inet** on one of the logical units. Previously, it was possible to commit this invalid configuration.
-

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

Related Documentation

- *Configuring Interface Encapsulation on Physical Interfaces*
- *Configuring CCC Encapsulation for Layer 2 VPNs*
- *Configuring Layer 2 Switching Cross-Connects Using CCC*
- *Configuring TCC Encapsulation for Layer 2 VPNs and Layer 2 Circuits*
- *Configuring ATM Interface Encapsulation*
- *Configuring ATM-to-Ethernet Interworking*
- *Configuring VLAN and Extended VLAN Encapsulation*
- *Configuring VLAN and Extended VLAN Encapsulation*
- *Configuring Encapsulation for Layer 2 Wholesale VLAN Interfaces*
- *Configuring Interfaces for Layer 2 Circuits*
- *Configuring Interface Encapsulation on PTX Series Packet Transport Routers*
- *Configuring MPLS LSP Tunnel Cross-Connects Using CCC*
- *Configuring TCC*
- *Configuring VPLS Interface Encapsulation*
- *Configuring Interfaces for VPLS Routing*
- *Defining the Encapsulation for Switching Cross-Connects*
- *Configuring an MPLS-Based Layer 2 VPN (CLI Procedure)*

esi

Syntax	<code>esi identifier { all-active single-active; }</code>
Hierarchy Level	<code>[edit interfaces <i>interface-name</i>] [edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i>]</code>
Release Information	Statement introduced in Junos OS Release 14.1. Statement introduced in Junos OS Releases 15.1F6 and 17.1 on MX Series routers with MPCs.
Description	<p>Configure an Ethernet Segment Identifier (ESI) on a physical, aggregated Ethernet, or logical interface in either EVPN multihoming active-standby or active-active mode.</p> <p>In releases before Junos OS Release 15.1F6 and 17.1R1 for MX Series routers and Junos OS Release 17.3R1 for EX9200 switches, you can specify an ESI only on a physical or aggregated Ethernet interface, for example, set interfaces ae0 esi 00:11:22:33:44:55:66:77:88:99. If you specify an ESI on a physical or aggregated Ethernet interface, keep in mind that an ESI is a factor in the designated forwarder (DF) election process. For example, assume that you configure EVPN multihoming active-standby on aggregated Ethernet interface ae0, and given the ESI configured on ae0 and other determining factors, the DF election results in ae0 being in the down state. Further, all logical interfaces configured on aggregated Ethernet interface ae0, for example, set interfaces ae0 unit 1 and set interfaces ae0 unit 2 are also in the down state, which renders logical interfaces ae0.1 and ae0.2 unable to provide services to their respective customer sites (VLANs).</p> <p>Starting with Junos OS Releases 15.1F6 and 17.1R1 for MX Series routers and Junos OS Release 17.3R1 for EX9200 switches, you can specify an ESI on a logical interface. If you specify an ESI on a logical interface, the DF election process now occurs at the individual logical interface level, which enables you to better utilize logical interfaces. For example, assume that you configure logical interfaces ae0.1 and ae0.2 on aggregated Ethernet interface ae0. You configure EVPN multihoming active-standby on both logical interfaces and given the ESI configured on ae0.1 and other determining factors, the DF election results in ae0.1 being in the down state. Despite logical interface ae0.1 being down, logical interface ae0.2 and other logical interfaces configured on ae0 can be in the up state and provide services to their respective customer sites (VLANs).</p>
Options	esi —Ten octet value. ESI value 0 and all fixed filters are reserved, and not used for configuring a multihomed Ethernet segment.



NOTE:

- Two interfaces (physical, logical, or aggregated Ethernet) cannot be configured with the same ESI value.

- The left most octet must be configured as 00. The other 9 octets are fully configurable.

.....

all-active—Configure the EVPN active-active multihoming mode of operation.

single-active—Configure the EVPN active-standby multihoming mode of operation.

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

Related • [evpn on page 138](#)
Documentation • [Example: Configuring an ESI on a Logical Interface With EVPN Multihoming on page 115](#)

evpn

```
Syntax evpn {
    designated-forwarder-election-hold-time seconds;
    designated-forwarder-preference least;
    extended-vni-list {
        vni-options
    }
    no-default-gateway-ext-comm
    encapsulation encapsulation-type;
    extended-vni-list;
    multicast-mode client | ingress-replication;
    vni-options {
        vni xxx vrf-target export target:xxx:xx
        vni xxx vrf-export name
    }
    extended-vni-all;
    no-default-gateway-ext-comm;
}
extended-vlan-list vlan-id | [vlan-id set];
extended-isid-list (single-isid | isid-list | isid-range | all);
pbb-evpn-core;
interface interface-name {
    ignore-encapsulation-mismatch;
    interface-mac-limit limit {
        packet-action drop;
    }
    mac-pinning (EVPN Routing Instances)
    no-mac-learning;
    static-mac mac-address;
}
interface-mac-limit limit {
    packet-action drop;
}
label-allocation per-instance;
mac-statistics;
mac-table-size limit {
    packet-action drop;
}
no-mac-learning;
traceoptions {
    file filename <files number> <size size> <world-readable | no-world-readable>;
    flag flag <flag-modifier>;
}
}
```

Hierarchy Level [edit routing-instances *routing-instance-name* protocols]
[edit protocols]

Release Information Statement introduced in Junos OS Release 13.2 for EVPNs on MX 3D Series routers.
designated-forwarder-election-hold-time *seconds* statement introduced in Junos OS Release 14.1.
extended-vlan-list *vlan-id* | [*vlan-id set*] statement introduced in Junos OS Release 14.1.

Statement introduced in Junos OS Release 14.1-X53-D30 for QFX Series switches.




NOTE: The `extended-vlan-list` statement is not supported on QFX10000 switches.

`designated-forwarder-preference-least` statement introduced in Junos OS Release 17.3 for MX Series routers with MPC and MIC interfaces.

Description	Enable an Ethernet VPN (EVPN) on the routing instance.
Options	<p><code>designated-forwarder-election-hold-time <i>seconds</i></code>—Time in seconds to wait before electing a designated forwarder (DF). Range: 1 through 1800 seconds</p> <p><code>designated-forwarder-preference-least</code>—Use least preference value for DF election. By default, the preference-based DF election is based on the highest preference value configured.</p> <p><code>evpn-etree</code>—Configure an Ethernet VPN E-TREE service.</p> <p>The remaining statements are explained separately.</p>
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • <i>Implementing EVPN-VXLAN for Data Centers</i> • <i>Configuring EVPN Routing Instances</i> • Tracing EVPN Traffic and Operations on page 17 • <i>Implementing EVPN-VXLAN for Data Centers</i> • <i>Provider Backbone Bridging (PBB) and EVPN Integration Overview</i>

extended-vlan-list

Syntax	<code>extended-vlan-list <i>vlan-id</i> [<i>vlan-id set</i>];</code>
Hierarchy Level	<code>[edit routing-instances <i>routing-instance-name</i> instance-type virtual-switch protocols evpn]</code>
Release Information	Statement introduced in Junos OS Release 14.1. Statement introduced in Junos OS Release 14.2 on EX Series switches.
Description	Specify the VLAN or range of VLANs that are extended over the WAN, wherein all the single VLAN bridge domains corresponding to these VLANs are stretched.
<div> NOTE: The <code>extended-vni-list</code> statement is an exclusive command. You cannot configure the <code>extended-vni-list</code> statement with either the <code>extended-vlan-list</code> or <code>extended-vni-all</code> statements.</div>	
Options	<p><i>vlan-id</i>—VLAN ID to be EVPN extended.</p> <p><i>vlan-id set</i>—List of VLAN IDs to be EVPN extended. Range: 1 through 4094 VLANs</p>
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none">• evpn on page 138

instance-type

Syntax	<code>instance-type type;</code>
Hierarchy Level	<code>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i>],</code> <code>[edit routing-instances <i>routing-instance-name</i>]</code>
Release Information	Statement introduced before Junos OS Release 7.4. virtual-switch and layer2-control options introduced in Junos OS Release 8.4. Statement introduced in Junos OS Release 9.2 for EX Series switches. Statement introduced in Junos OS Release 11.3 for the QFX Series. Statement introduced in Junos OS Release 12.3 for ACX Series routers. mpls-internet-multicast option introduced in Junos OS Release 11.1 for the EX Series, M Series, MX Series, and T Series. evpn option introduced in Junos OS Release 13.2 for MX 3D Series routers. forwarding option introduced in Junos OS Release 14.2 for the PTX Series. mpls-forwarding option introduced in Junos OS Release 16.1 for the MX Series. evpn-vpws option introduced in Junos OS Release 17.1 for MX Series routers.
Description	Define the type of routing instance.

Options



NOTE: On ACX Series routers, you can configure only the forwarding, virtual router, and VRF routing instances.

type—Can be one of the following:

- **evpn**—(MX 3D Series routers and QFX switches only) Enable an Ethernet VPN (EVPN) on the routing instance.
You cannot configure the **evpn** option under the `[edit logical-systems logical-system-name routing-instances routing-instance-name instance-type]` hierarchy level.
- **evpn-vpws**—Enable an Ethernet VPN (EVPN) Virtual Private Wire Service (VPWS) on the routing instance.
- **forwarding**—Provide support for filter-based forwarding, where interfaces are not associated with instances. All interfaces belong to the default instance. Other instances are used for populating RPD learned routes. For this instance type, there is no one-to-one mapping between an interface and a routing instance. All interfaces belong to the default instance inet.0.
- **l2backhaul-vpn**—Provide support for Layer 2 wholesale VLAN packets with no existing corresponding logical interface. When using this instance, the router learns both the outer tag and inner tag of the incoming packets, when the **instance-role** statement is defined as **access**, or the outer VLAN tag only, when the **instance-role** statement is defined as **nni**.

- **l2vpn**—Enable a Layer 2 VPN on the routing instance. You must configure the **interface**, **route-distinguisher**, **vrf-import**, and **vrf-export** statements for this type of routing instance.
- **layer2-control**—(MX Series routers only) Provide support for RSTP or MSTP in customer edge interfaces of a VPLS routing instance. This instance type cannot be used if the customer edge interface is multihomed to two provider edge interfaces. If the customer edge interface is multihomed to two provider edge interfaces, use the default BPDU tunneling.
- **mpls-forwarding**—(MX Series routers only) Allow filtering and translation of route distinguisher (RD) values in IPv4 and IPv6 VPN address families on both routes received and routes sent for selected BGP sessions. In particular, for Inter-AS VPN Option-B networks, this option can prevent the malicious injection of VPN labels from one peer AS boundary router to another.
- **mpls-internet-multicast**—(EX Series, M Series, MX Series, and T Series routers only) Provide support for ingress replication provider tunnels to carry IP multicast data between routers through an MPLS cloud, using MBGP or next-generation MVPN.
- **no-forwarding**—This is the default routing instance. Do not create a corresponding forwarding instance. Use this routing instance type when a separation of routing table information is required. There is no corresponding forwarding table. All routes are installed into the default forwarding table. IS-IS instances are strictly nonforwarding instance types.
- **virtual-router**—Enable a virtual router routing instance. This instance type is similar to a VPN routing and forwarding instance type, but used for non-VPN-related applications. You must configure the **interface** statement for this type of routing instance. You do not need to configure the **route-distinguisher**, **vrf-import**, and **vrf-export** statements.
- **virtual-switch**—(MX Series routers, EX9200 switches, and QFX switches only) Provide support for Layer 2 bridging. Use this routing instance type to isolate a LAN segment with its Spanning Tree Protocol (STP) instance and to separate its VLAN identifier space.
- **vpls**—Enable VPLS on the routing instance. Use this routing instance type for point-to-multipoint LAN implementations between a set of sites in a VPN. You must configure the **interface**, **route-distinguisher**, **vrf-import**, and **vrf-export** statements for this type of routing instance.
- **vrf**—VPN routing and forwarding (VRF) instance. Provides support for Layer 3 VPNs, where interface routes for each instance go into the corresponding forwarding table only. Required to create a Layer 3 VPN. Create a VRF table (**instance-name.inet.0**) that contains the routes originating from and destined for a particular Layer 3 VPN. For this instance type, there is a one-to-one mapping between an interface and a routing instance. Each VRF instance corresponds with a forwarding table. Routes on an interface go into the corresponding forwarding table. You must configure the **interface**, **route-distinguisher**, **vrf-import**, and **vrf-export** statements for this type of routing instance.

Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• <i>Configuring Routing Instances on PE Routers in VPNs</i>• <i>Configuring EVPN Routing Instances</i>• Configuring EVPN Routing Instances on EX9200 Switches on page 15• <i>Configuring Virtual Router Routing Instances</i>• <i>Example: Configuring Filter-Based Forwarding on the Source Address</i>• <i>Example: Configuring Filter-Based Forwarding on Logical Systems</i>• <i>vpws-service-id</i>

interface (EVPN Routing Instances)

Syntax	<pre>interface <i>interface-name</i> { ignore-encapsulation-mismatch; interface-mac-limit <i>limit</i> { packet-action drop; } mac-pinning (EVPN Routing Instances) no-mac-learning; static-mac <i>mac-address</i>; vpws-service-id }</pre>
Hierarchy Level	[edit routing-instances <i>routing-instance-name</i> protocols evpn]
Release Information	Statement introduced in Junos OS Release 13.2 for EVPNs on MX 3D Series routers. Statement introduced in Junos OS Release 14.2 on EX Series switches. Statement (mac-pinning) introduced in Junos OS Release 16.2 on MX Series routers. vpws-service-id statement introduced in Junos OS Release 17.1 on MX Series routers.
Description	Specify each interface over which the Ethernet VPN (EVPN) traffic travels between the PE device and customer edge (CE) device. The interfaces are bound to the EVPN routing instance.
Options	<p><i>interface-name</i>—Name of the interface.</p> <p>The remaining statements are explained separately. See CLI Explorer.</p>
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• <i>Configuring EVPN Routing Instances</i>• Configuring EVPN Routing Instances on EX9200 Switches on page 15• evpn on page 138• instance-type on page 141• <i>vpws-service-id</i>

interface (Routing Instances)

Syntax	<pre>interface <i>interface-name</i> { description <i>text</i>; }</pre>
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i>], [edit routing-instances <i>routing-instance-name</i>]
Release Information	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.2 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 12.3 for ACX Series routers.</p> <p>Statement introduced in Junos OS Release 13.2 for MX 3D Series routers.</p>
Description	Specify the interface over which the VPN traffic travels between the PE device and CE device. You configure the interface on the PE device. If the value vrf is specified for the instance-type statement included in the routing instance configuration, this statement is required.
Options	<p><i>interface-name</i>—Name of the interface.</p> <p>The remaining statement is explained separately. See CLI Explorer.</p>
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • Configuring Routing Instances on PE Routers in VPNs • Configuring EVPN Routing Instances • Configuring EVPN Routing Instances on EX9200 Switches on page 15 • interface (VPLS Routing Instances)

interface-mac-limit (VPLS)

Syntax	<pre>interface-mac-limit <i>limit</i> { packet-action drop; }</pre>
Hierarchy Level	<p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols vpls site <i>site-name</i> interfaces <i>interface-name</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols evpn],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols evpn interface <i>interface-name</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols vpls site <i>site-name</i> interfaces <i>interface-name</i>]</p>
Release Information	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Support for EVPNs introduced in Junos OS Release 13.2 on MX 3D Series routers.</p> <p>Support for EVPNs introduced in Junos OS Release 14.2 on EX Series switches.</p>
Description	<p>Specify the maximum number of media access control (MAC) addresses that can be learned by the EVPN or VPLS routing instance. You can configure the same limit for all interfaces configured for a routing instance. You can also configure a limit for a specific interface.</p> <p>Starting with Junos OS Release 12.3R4, if you do not configure the parameter to limit the number of MAC addresses to be learned by a VPLS instance, the default value is not effective. Instead, if you do not include the interface-mac-limit option at the [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols vpls site <i>site-name</i> interfaces <i>interface-name</i>], hierarchy level, this setting is not present in the configuration with the default value of 1024 addresses. If you upgrade a router running a Junos OS release earlier than Release 12.3R4 to Release 12.3R4 or later, you must configure the interface-mac-limit option with a valid value for it to be saved in the configuration.</p>
Options	<p>limit—Number of MAC addresses that can be learned from each interface.</p> <p>Range: 16 through 65,536 MAC addresses</p> <p>Default: 1024 addresses</p> <p>The remaining statement is explained separately. See CLI Explorer.</p>
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • Configuring EVPN Routing Instances • Configuring EVPN Routing Instances on EX9200 Switches on page 15 • Configuring VPLS Routing Instances • interface on page 144

- [mac-table-size on page 154](#)

ip-prefix-routes

Syntax `ip-prefix-routes {
 advertise (direct-nexthop | gateway-address);
 encapsulation (vxlan | mpls);
 <export routing-policy-name>;
 gateway-interface interface-name;
 vni number;
 }`

Hierarchy Level [edit routing-instances *routing-instance-name* protocols evpn]

Release Information Statement introduced in Junos OS Release 15.1X53-D60 and Junos OS Release 17.2R1 for QFX10000 switches.
 Statement introduced in Junos OS Release 17.3R1 for EX9200 switches.

Description In an Ethernet VPN (EVPN) environment, enable the switch to advertise the IP prefix associated with a specified customer domain as a type-5 route to remote data centers or in a metro transport network. You use this feature when the Layer 2 domain does not exist at the remote data centers or metro network peering points.

Two models for implementing type-5 routes are supported:

- Pure type-5 route without overlay next hop and type-2 route
- Type-5 route with gateway IRB interface as overlay next hop and type-2 route

A pure type-5 route advertises the summary IP prefix and includes a BGP extended community called a router MAC, which is used to carry the MAC address of the sending switch and to provide next-hop reachability information for the prefix. In contrast, a standard type-5 route requires a gateway IP address as a next-hop overlay and a supporting type-2 route to provide recursive route resolution.

On QFX10000 switches, only fully resolved next-hop—that is, EVPN pure type-5—routes are currently supported. QFX10000 switches support only EVPN Virtual Extensible LAN (VXLAN). MPLS encapsulation is not supported.

On EX9200 switches, both type-5 routes are supported. Only MPLS encapsulation is supported with pure type-5 routes. Both MPLS and VXLAN encapsulation are supported with the standard type-5 route with a gateway IP address as next-hop overlay.



CAUTION: Pure type-5 routing for EVPN-VXLAN was introduced in Junos OS Release 15.1X53-D30 for QFX10002 switches only. In that release, this statement is `ip-prefix-support forwarding-mode symmetric`. Starting with Junos OS Release 15.1X53-D60, the statement is `ip-prefix-routes advertise direct-nexthop`. Any configuration with the original `ip-prefix-support` statement is automatically upgraded to the new `ip-prefix-routes` statement when you upgrade to Junos OS Release 15.1X53-D60 or later.



NOTE: Pure type-5 routing is supported on all QFX10000 switches starting in Junos OS Release 15.1X53-D60.



NOTE: QFX10000 switches do not support advertising an IP prefix with a mask length of /32 as a pure type-5 route.

Options The **advertise**, **encapsulation** and **vni *number*** options are required. The **export *routing-policy-name*** option is optional.

advertise *direct-nexthop*—Enable the switch to send IP prefix information using an EVPN pure type-5 route, which includes a router MAC extended community used to send the MAC address of the switch. This router MAC extended community provides next-hop reachability without requiring an overlay next-hop or supporting type-2 route.



NOTE: For pure route type-5, QFX10000 switches support only VXLAN encapsulation, and EX9200 switches support only MPLS encapsulation.

advertise *gateway-address* (EX9200 switches only)—Enable the switch to advertise a gateway address in exported IP prefix routes. This gateway address provides overlay next-hop reachability.



NOTE: You must also specify a gateway address by including the **gateway-interface *interface-name*** statement.

encapsulation (*vxlan* | *mpls*)—Specify to encapsulate forwarded traffic in VXLAN or MPLS for transmission to the remote data center.



NOTE: The same type of encapsulation must be used end to end. Only VXLAN encapsulation is supported on QFX10000 switches.

export *routing-policy-name*—(Optional) Specify the name of the routing policy configured at the **[edit policy-options policy-statement *policy-statement-name*]** hierarchy level to apply to the routes for the specified customer domain. Applying an export policy allows you to further control the IP prefixes to advertise or to suppress through EVPN type-5 routes for each customer. You can apply a separate export routing policy to one or more customer domains. This allows each customer to each have its own policy.

gateway-interface *interface-name*—Specify the gateway interface to use as a next-hop overlay for a standard type-5 route. You must use this option in conjunction with the **advertise *gateway-address*** option.

vni *number*—Specify the identifier associated with a customer domain. Each customer domain must have a unique identifier.


Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• <i>Understanding EVPN Pure Route Type-5 on QFX Series Switches</i>• EVPN Overview for Switches on page 3• <i>policy-statement</i>

mac-statistics

Syntax	mac-statistics;
Hierarchy Level	<p>[edit bridge-domains <i>bridge-domain-name</i> bridge-options],</p> <p>[edit logical-systems <i>logical-system-name</i> bridge-domains <i>bridge-domain-name</i> bridge-options],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> bridge-domains <i>bridge-domain-name</i> bridge-options],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> switch-options],</p> <p>[edit logical-systems <i>logical-system-name</i> switch-options],</p> <p>[edit routing-instances <i>routing-instance-name</i> bridge-domains <i>bridge-domain-name</i> bridge-options],</p> <p>[edit routing-instances <i>routing-instance-name</i> switch-options],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols evpn],</p> <p>[edit switch-options],</p> <p>[edit switch-options],</p> <p>[edit vlans <i>vlan-name</i> switch-options]</p>
Release Information	<p>Statement introduced in Junos OS Release 8.4.</p> <p>Support for the switch-options statement added in Junos OS Release 9.2.</p> <p>Support for top-level configuration for the virtual-switch type of routing instance added in Junos OS Release 9.2. In Junos OS Release 9.1 and earlier, the routing instances hierarchy supported this statement only for a VPLS instance or a bridge domain configured within a virtual switch.</p> <p>Support for logical systems added in Junos OS Release 9.6.</p> <p>[edit switch-options] and [edit vlans <i>vlan-name</i> switch-options] hierarchy levels introduced in Junos OS Release 12.3R2 for EX Series switches.</p> <p>Support for EVPNs added in Junos OS Release 13.2 for MX 3D Series routers.</p> <p>[edit switch-options] and [edit vlans <i>vlan-name</i> switch-options] hierarchy levels introduced in Junos OS Release 13.2 for the QFX Series.</p>
Description	(MX Series routers, EX Series switches, and QFX Series only) For bridge domains or VLANs, enable MAC accounting either for a specific bridge domain or VLAN, or for a set of bridge domains or VLANs associated with a Layer 2 trunk port.
Default	disabled
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • <i>Understanding Layer 2 Learning and Forwarding for Bridge Domains</i> • <i>Layer 2 Learning and Forwarding for VLANs Overview</i> • <i>Understanding Layer 2 Learning and Forwarding for Bridge Domains Functioning as Switches with Layer 2 Trunk Ports</i>

- *Layer 2 Learning and Forwarding for VLANs Acting as a Switch for a Layer 2 Trunk Port*
- *Configuring EVPN Routing Instances*
- [Configuring EVPN Routing Instances on EX9200 Switches on page 15](#)

mac-table-size

Syntax	<code>mac-table-size size { packet-action drop; }</code>
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols vpls], [edit routing-instances <i>routing-instance-name</i> protocols evpn], [edit routing-instances <i>routing-instance-name</i> protocols vpls]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 13.2 for EVPNs on MX 3D Series routers. Statement introduced in Junos OS Release 14.2 for EX Series switches.
Description	Specify the size of the MAC address table.
Options	<p><i>size</i>—Size of the MAC address table.</p> <p>Range:</p> <ul style="list-style-type: none"> • (M Series and T Series routers only) 16 through 65,536 MAC addresses • (MX Series routers only) 16 through 1,048,575 MAC addresses • (T4000 routers with Type 5 FPCs only) 16 through 262,143 MAC addresses
	<p> NOTE: Before modifying the size of the MAC address table (to 262,143 addresses), you must enable network services mode by including the <code>enhanced-mode</code> statement at the [edit chassis network-services] hierarchy level and then reboot the router.</p>
	<p>Default: 512 MAC addresses</p> <p>The remaining statement is explained separately. See CLI Explorer.</p>
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • <i>Configuring EVPN Routing Instances</i> • Configuring EVPN Routing Instances on EX9200 Switches on page 15 • <i>Configuring VPLS Routing Instances</i> • <i>Configuring Improved VPLS MAC Address Learning on T4000 Routers with Type 5 FPCs</i> • <i>enhanced-mode</i> • evpn on page 138

no-arp-suppression

Syntax	no-arp-suppression;
Hierarchy Level	<pre>[edit bridge-domains <i>bridge-domain-name</i>] [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> vlans <i>vlan-name</i>] [edit routing-instances <i>routing-instance-name</i> vlans <i>vlan-name</i>] [edit vlans <i>vlan-name</i>] [edit routing-instances <i>instance-name</i> protocols evpn] [edit routing-instances <i>instance-name</i> bridge-domains <i>domain-name</i>]</pre>
Release Information	Statement introduced in Junos OS Release 17.2.
Description	<p>Disable the suppression of Address Resolution Protocol (ARP) requests from a customer edge (CE) device to a provider edge (PE) device that acts as ARP proxy in an Ethernet VPN-MPLS (EVPN-MPLS) or Ethernet VPN-Virtual Extensible LAN (EVPN-VXLAN) environment.</p> <p>When disabling ARP suppression, be aware of the following implications;</p> <ul style="list-style-type: none"> • If the PE device does not find the MAC-IP address binding in its database, it cannot forward the ARP request. Instead, the device discards the ARP request, and the packets flood through the Layer 2 domain because the ARP request is considered to be Layer 2 broadcast traffic. Therefore, we recommend that ARP suppression remains enabled. • Disabling the suppression of ARP packets on a PE device essentially disables the proxy ARP functionality. <p>Therefore, we recommend that ARP suppression remains enabled.</p>
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • EVPN Proxy ARP and ARP Suppression and Network Discovery Protocol and Network Discovery Protocol Suppression on page 5

no-mac-learning

Syntax no-mac-learning;

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

Release Information Statement introduced in Junos OS Release 8.4.
 Support for the **switch-options** statement added in Junos OS Release 9.2.
 Support for top-level configuration for the **virtual-switch** type of routing instance added in Junos OS Release 9.2. In Junos OS Release 9.1 and earlier, the routing instances hierarchy supported this statement only for a VPLS instance or bridge domain configured within a virtual switch.
 Support for logical systems added in Junos OS Release 9.6.
[edit switch-options], **[edit switch-options interface *interface-name*]**, **[edit vlans *vlan-name* switch-options]**, and **[edit vlans *vlan-name* switch-options interface *interface-name*]** hierarchy levels introduced in Junos OS Release 12.3 R2 for EX Series switches.
 Support for EVPNs added in Junos OS Release 13.2 for MX 3D Series routers.
 Hierarchy levels **[edit switch-options interface *interface-name*]** and **[edit vlans *vlan-name* switch-options]** introduced in Junos OS Release 13.2X50-D10 for EX Series switches.

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



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

Default	MAC learning is enabled.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • <i>Configuring EVPN Routing Instances</i> • Configuring EVPN Routing Instances on EX9200 Switches on page 15 • <i>Understanding Layer 2 Learning and Forwarding for Bridge Domains</i> • <i>Layer 2 Learning and Forwarding for VLANs Overview</i> • <i>Understanding Layer 2 Learning and Forwarding for Bridge Domains Functioning as Switches with Layer 2 Trunk Ports</i> • <i>Understanding Bridging and VLANs on EX Series Switches</i> • <i>Understanding Q-in-Q Tunneling on EX Series Switches</i>

packet-action

Syntax packet-action *action*;

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

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

supported this statement only for a VPLS instance or a bridge domain configured within a virtual switch.

Support for logical systems added in Junos OS Release 9.6.

[edit switch-options interface *interface-name* interface-mac-limit *limit*], [edit switch-options interface-mac-limit *limit*], [edit switch-options mac-table-size *limit*], [edit vlans *vlan-name* switch-options interface *interface-name* interface-mac-limit *limit*], [edit vlans *vlan-name* switch-options interface-mac-limit *limit*], and [edit vlans *vlan-name* switch-options mac-table-size *limit*] hierarchy levels introduced in Junos OS Release 12.3R2 for EX Series switches.

Support for EVPNs introduced in Junos OS Release 13.2 on MX Series 3D Universal Edge Routers.

Support at the [edit switch-options interface *interface-name* interface-mac-limit *limit*] hierarchy level and hierarchy levels under [edit vlans *vlan-name*] introduced in Junos OS Release 13.2X50-D10 for EX Series switches and Junos OS Release 13.2 for the QFX Series.

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

Default



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

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

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



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

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

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

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

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

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

- Related Documentation**
- *Configuring EVPN Routing Instances*
 - [Configuring EVPN Routing Instances on EX9200 Switches on page 15](#)
 - *Configuring MAC Limiting (CLI Procedure)*
 - *Configuring Persistent MAC Learning (CLI Procedure)*
 - *Understanding Layer 2 Learning and Forwarding for Bridge Domains*
 - *Layer 2 Learning and Forwarding for VLANs Overview*
 - *Understanding Layer 2 Learning and Forwarding for Bridge Domains Functioning as Switches with Layer 2 Trunk Ports*
 - *Layer 2 Learning and Forwarding for VLANs Overview*
 - *Layer 2 Learning and Forwarding for VLANs Acting as a Switch for a Layer 2 Trunk Port*

routing-instances

Syntax	<code>routing-instances <i>routing-instance-name</i> { ... }</code>
Hierarchy Level	[edit], [edit logical-systems <i>logical-system-name</i>]
Release Information	Statement introduced before Junos OS Release 7.4.
Description	Configure an additional routing entity for a router or switch. You can create multiple instances of BGP, IS-IS, OSPF, OSPF version 3 (OSPFv3), and RIP for a router or switch.
Default	Routing instances are disabled for the router or switch.
Options	<i>routing-instance-name</i> —Name of the routing instance, a maximum of 31 characters. The remaining statements are explained separately.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• <i>Configuring EVPN Routing Instances</i>• <i>Configuring Routing Instances on PE Routers in VPNs</i>

traceoptions (Protocols EVPN)

Syntax	<pre>traceoptions { file <i>filename</i> <files <i>number</i>> <size <i>size</i>> <world-readable no-world-readable>; flag <i>flag</i> <flag-modifier>; }</pre>
Hierarchy Level	[edit routing-instances <i>routing-instance-name</i> protocols evpn]
Release Information	Statement introduced in Junos OS Release 13.2 for MX 3D Series routers. Statement introduced in Junos OS Release 14.2 for EX Series switches.
Description	Trace traffic flowing through an EVPN routing instance.
Options	<p>file <i>filename</i>—Name of the file to receive the output of the tracing operation. Enclose the name in quotation marks (" ").</p> <p>files <i>number</i>—(Optional) Maximum number of trace files. When a trace file named <i>trace-file</i> reaches the maximum size as specified by the size option, it is renamed <i>trace-file.0</i>. When <i>trace-file</i> again reaches the maximum size, <i>trace-file.0</i> is renamed <i>trace-file.1</i> and <i>trace-file</i> is renamed <i>trace-file.0</i>. This renaming scheme continues until the maximum number of trace files is reached. Then the oldest trace file is overwritten.</p> <p>If you specify a maximum number of files, you also must specify a maximum file size with the size option.</p> <p>Range: 2 through 1000 files</p> <p>Default: 2 files</p> <p>flag <i>flag</i>—Tracing operation to perform. To specify more than one tracing operation, include multiple flag statements. You can specify the following tracing flags:</p> <ul style="list-style-type: none">• all—All EVPN tracing options• error—Error conditions• general—General events• mac-database—MAC route database in the EVPN routing instance• nlri—EVPN advertisements received or sent by means of BGP• normal—Normal events• oam—OAM messages• policy—Policy processing• route—Routing information• state—State transitions• task—Routing protocol task processing

- **timer**—Routing protocol timer processing
- **topology**—EVPN topology changes caused by reconfiguration or advertisements received from other provider edge (PE) routers using BGP

flag-modifier—(Optional) Modifier for the tracing flag. You can specify the following modifiers:

- **detail**—Provide detailed trace information.
- **disable**—Disable this trace flag.
- **receive**—Trace received packets.
- **send**—Trace sent packets.

no-world-readable—Do not allow any user to read the log file.

size size—(Optional) Maximum size of each trace file, in kilobytes (KB), megabytes (MB), or gigabytes (GB). When a trace file named **trace-file** reaches this size, it is renamed **trace-file.0**. When **trace-file** again reaches the maximum size, **trace-file.0** is renamed **trace-file.1** and **trace-file** is renamed **trace-file.0**. This renaming scheme continues until the maximum number of trace files (as specified by the **files** option) is reached. Then the oldest trace file is overwritten.

If you specify a maximum file size, you also must specify a maximum number of trace files with the **files** option.

Syntax: **xk** to specify kilobytes, **xm** to specify megabytes, or **xg** to specify gigabytes

Range: 10 KB through the maximum file size supported on your system

Default: 1 MB

world-readable—Allow any user to read the log file.

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

Related Documentation	<ul style="list-style-type: none"> • Tracing EVPN Traffic and Operations on page 17
------------------------------	--

vlan-id (routing instance)

Syntax	<code>vlan-id (vlan-id all none);</code>
Hierarchy Level	<code>[edit logical-systems logical-system-name routing-instances routing-instance-name],</code> <code>[edit routing-instances routing-instance-name]</code>
Release Information	Statement introduced in Junos OS Release 13.2. Statement introduced in Junos OS Release 14.2 for EX Series switches.
Description	Specify 802.1Q VLAN tag IDs to a routing instance.
Options	<p>vlan-id—A valid VLAN identifier.</p> <p>Range: For 4-port Fast Ethernet PICs, 512 through 1023. For 1-port and 10-port Gigabit Ethernet PICs configured to handle VPLS traffic, 512 through 4094.</p> <p>all—Include all VLAN identifiers specified on the logical interfaces included in the routing instance.</p> <p>none—Include no VLAN identifiers for the routing instance.</p>
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • Configuring EVPN Routing Instances • Configuring EVPN Routing Instances on EX9200 Switches on page 15

CHAPTER 8

Operational Commands

- `clear evpn nd-table`
- `clear ethernet-switching evpn nd-statistics`
- `clear ethernet-switching evpn nd-table`
- `show ethernet-switching evpn nd-statistics`
- `show evpn arp-table`
- `show evpn database`
- `show evpn flood`
- `show evpn instance`
- `show evpn ip-prefix-database`
- `show evpn l3-context`
- `show evpn mac-table`
- `show evpn nd-table`
- `show evpn peer-gateway-macs`
- `show vlans evpn nd-table`

clear evpn nd-table

Syntax	clear evpn nd-table
Release Information	Command introduced in Junos OS Release 16.2. Command introduced in Junos OS Release 17.1 for MX series Routers. Command introduced in Junos OS Release 17.3R1 for EX series switches.
Description	Clear the Neighbor Discovery (ND) proxy table from the Ethernet VPN (EVPN). ND proxy is a kernel module that implements IPv6 Neighbor Discovery proxying over Ethernet-like access networks.
Options	none —Clear learned ND proxy tables from the EVPN.
Required Privilege Level	clear
List of Sample Output	clear evpn nd-table on page 166
Output Fields	When you enter this command, you are provided feedback on the status of your request.

Sample Output

clear evpn nd-table

```
user@host> clear evpn nd-table
```

clear ethernet-switching evpn nd-statistics

Syntax	clear ethernet-switching evpn nd-statistics
Release Information	Command introduced in Junos OS Release 17.3R1 for EX series switches.
Description	Clear the Neighbor Discovery (ND) proxy table statistics for integrated routing and bridging (IRB) interfaces participating in the Ethernet VPN (EVPN). ND proxy is a kernel module that implements IPv6 Neighbor Discovery proxying over Ethernet-like access networks.
Options	none —Clear the ND proxy table statistics for IRB interfaces participating in the EVPN.
Required Privilege Level	clear
List of Sample Output	clear ethernet-switching evpn nd-statistics on page 167
Output Fields	When you enter this command, the ND proxy table statistics for IRB interfaces participating in the EVPN are cleared.

Sample Output

clear ethernet-switching evpn nd-statistics

```
user@host> clear ethernet-switching evpn nd-statistics
```

clear ethernet-switching evpn nd-table

Syntax	clear ethernet-switching evpn nd-table
Release Information	Command introduced in Junos OS Release 17.3R1 for the EX Series.
Description	Clear the Neighbor Discovery (ND) proxy table from the Ethernet VPN (EVPN) for integrated routing and bridging (IRB) interfaces. This command applies to EVPN instances of type virtual-switch .
Options	none —Clear information about the ND proxy tables for an EVPN.
Required Privilege Level	clear
List of Sample Output	clear ethernet-switching evpn nd-table on page 168
Output Fields	When you enter this command, you are provided feedback on the status of your request.

Sample Output

clear ethernet-switching evpn nd-table

```
user@host> clear ethernet-switching evpn nd-table
```


show ethernet-switching evpn nd-statistics

Syntax `show ethernet-switching evpn nd-statistics`
`<brief | count | detail | extensive | summary | display xml>`

Release Information Command introduced in Junos OS Release 17.3R1 for EX series switches.

Description Show Neighbor Discovery (ND) proxy table statistics for integrated routing and bridging (IRB) interfaces participating in the Ethernet VPN (EVPN). ND proxy is a kernel module that implements IPv6 Neighbor Discovery proxying over Ethernet-like access networks.

Options **none**—Display information about ND proxy table statistics for all IRB interfaces participating in the EVPN.

brief | count | detail | extensive | summary | display xml—(Optional) Display the specified level of output.

Required Privilege Level view

List of Sample Output [show ethernet-switching evpn nd-statistics on page 170](#)

Output Fields [Table 8 on page 169](#) lists the output fields for the **show ethernet-switching evpn nd-statistics** command. Output fields are listed in the approximate order in which they appear.

Table 8: show ethernet-switching evpn nd-statistics Output Fields

Field Name	Field Description	Level of Output
Interface	IRB interface on which the statistics has been applied	All levels
EVPN	Name of the EVPN	All levels
Bridge domain	Name of the bridge domain for the EVPN	All levels
ND routes add received	Total number of additional ND routes received	All levels
ND routes del received	Total number of deleted ND routes received	All levels
ND routes dropped	Total number of ND routes dropped	All levels
MAC+IPv6s sent to peer	Total number of MAC and IPv6 addresses sent to peer device	All levels

Sample Output

show ethernet-switching evpn nd-statistics

```
user@host> show ethernet-switching evpn nd-statistics interface irb.0

Interface : irb.0      EVPN : evpn1      Bridge domain: vlan10
ND routes add received      : 2
ND routes del received      : 0
ND routes dropped           : 0
MAC+IPv6s sent to peer      : 3
```

show evpn arp-table

Syntax	show evpn arp-table <address> <brief count detail extensive> <instance <i>instance-name</i> >
Release Information	Command introduced in Junos OS Release 15.1 for EX Series switches.
Description	Show Ethernet VPN (EVPN) Address Resolution Protocol (ARP) entries associated with learned MAC addresses.
Options	<p>none—Display brief information about the EVPN ARP table.</p> <p>address—(Optional) Display ARP information for the specified MAC address.</p> <p>brief count detail extensive—(Optional) Display the specified level of output.</p> <p>instance <instance-name>—(Optional) Display ARP information for the specified routing instance .</p>
Required Privilege Level	view
List of Sample Output	show evpn arp-table on page 172 show evpn arp-table 00:05:86:a0:dc:f0 (MAC address) on page 172 show evpn arp-table brief on page 172 show evpn arp-table detail on page 172 show evpn arp-table count on page 172 show evpn arp-table extensive on page 172 show evpn arp-table instance evpn1 on page 173
Output Fields	Table 9 on page 171 lists the output fields for the show evpn arp-table command. Output fields are listed in the approximate order in which they appear.

Table 9: show evpn arp-table Output Fields

Field Name	Field Description	Level of Output
INET address	The INET address related to the INET entries that are added to the ARP table.	All levels
MAC address	MAC addresses learned through ARP.	brief, detail, extensive, instance, mac-address,,
Logical Interface	Logical interface associated with the routing instance in which the ARP INET address is learned.	brief, instance, mac-address,,

Table 9: show evpn arp-table Output Fields (*continued*)

Field Name	Field Description	Level of Output
Routing instance	Routing instance in which the ARP INET address is learned.	all levels
Bridging domain	Bridging domain in which the ARP INET address is learned.	all levels
Learning interface	Interface on which the ARP INET address is learned.	detail, extensive
Count	Indicates the number of ARP INET addresses learned in a routing instance in a bridge domain.	count

Sample Output

show evpn arp-table

```

user@host> show evpn arp-table
INET          MAC          Logical      Routing      Bridging
address        address      interface    instance     domain
203.0.113.2    00:05:86:a0:dc:f0  irb.0        evpn1        __evpn1__

```

show evpn arp-table 00:05:86:a0:dc:f0 (MAC address)

```

user@host> show evpn arp-table 00:05:86:a0:dc:f0
INET          MAC          Logical      Routing      Bridging
address        address      interface    instance     domain
203.0.113.2    00:05:86:a0:dc:f0  irb.0        evpn1        __evpn1__

```

show evpn arp-table brief

```

user@host> show evpn arp-table brief
INET          MAC          Logical      Routing      Bridging
address        address      interface    instance     domain
203.0.113.2    00:05:86:a0:dc:f0  irb.0        evpn1        __evpn1__

```

show evpn arp-table detail

```

user@host> show evpn arp-table detail

INET address: 203.0.113.2
MAC address: 00:05:86:a0:dc:f0
Routing instance: evpn1
Bridging domain: __evpn1__
Learning interface: irb.0

```

show evpn arp-table count

```

user@switch> show evpn arp-table count
1 ARP INET addresses learned in routing instance evpn1 bridge domain __evpn1__

```

show evpn arp-table extensive

```

user@host> show evpn arp-table extensive

INET address: 203.0.113.2

```

```
MAC address: 00:05:86:a0:dc:f0
Routing instance: evpn1
Bridging domain: __evpn1__
Learning interface: irb.0
```

show evpn arp-table instance evpn1

```
user@host> show evpn arp-table instance evpn1
```

INET address	MAC address	Logical interface	Routing instance	Bridging domain
203.0.113.2	00:05:86:a0:dc:f0	irb.0	evpn1	__evpn1__

show evpn database

Syntax	<code>show evpn database</code> <code><extensive></code> <code><instance <i>instance-name</i>></code> <code><interface <i>interface-name</i>></code> <code><mac-address <i>address</i>></code> <code><neighbor <i>neighbor-name</i>></code> <code><origin <i>origin-name</i>></code> <code><state <i>state-name</i>></code> <code><vlan-id <i>vlan-id</i>></code>
Release Information	Command introduced in Junos OS Release 14.2 for EX Series switches. Command introduced in Junos OS Release 16.1 for MX Series routers. Command introduced in Junos OS Release 17.1 on MX Series routers with MPCs.
Description	Show Ethernet VPN (EVPN) database information.
Options	none —Display brief information about the EVPN database. extensive —(Optional) Display detailed information about the EVPN database. instance <i>instance-name</i> —(Optional) Display MAC addresses from the specified routing instance. interface <i>interface-name</i> —(Optional) Display the MAC address learned from the specified interface. mac-address <i>address</i> —(Optional) Display the specified MAC address. neighbor <i>neighbor-name</i> —(Optional) Display the MAC address learned from the specified neighbor. origin <i>origin-name</i> —(Optional) Display the MAC address with the specified origin. state <i>state-name</i> —(Optional) Display the MAC address with the specified state. vlan-id <i>vlan-id</i> —(Optional) Display the MAC address with the specified VLAN.
Required Privilege Level	view
Related Documentation	<ul style="list-style-type: none">• Example: Configuring an EVPN with IRB Solution on EX9200 Switches on page 26
List of Sample Output	show evpn database on page 175 show evpn database extensive on page 175
Output Fields	Table 10 on page 175 lists the output fields for the show evpn database command. Output fields are listed in the approximate order in which they appear.

Table 10: show evpn database Output Fields

Field Name	Field Description	Level of Output
Instance	Name of the routing instances.	All Levels
VNI	VXLAN network identifier.	All Levels
MAC address	MAC address of the routing instance.	All Levels
Origin	Specific origin of the MAC address.	All Levels
Timestamp	Month, day and time of generated command output.	All Levels
IP address	IP address.	All Levels
Prefix	Prefix address of the VXLAN network identifier.	extensive
Source	Source address of the VXLAN network identifier.	extensive
Rank	Level number of the VXLAN network identifier.	extensive
Status	Status of the of the VXLAN network identifier.	extensive
Remote origin	Specific remote origin of the VXLAN network identifier.	extensive

Sample Output

show evpn database

```

user@host> show evpn database
Instance: bd_red
VNI  MAC address      Origin      Timestamp      IP address
50    1a:1b:1b:1b:1b:1b  192.0.2.3   Jun 12 21:57:17
50    1a:1c:1c:1c:1c:1c  ge-2/3/0.0  Jun 12 22:05:37
50    b0:c6:9a:e9:cd:d8  192.0.2.3   Jun 12 21:57:17
50    b0:c6:9a:ea:87:42  ge-2/3/0.0  Jun 12 22:05:38

Instance: evpn-0
VNI  MAC address      Origin      Timestamp      IP address
4000  1a:1a:1a:1a:01:01  ge-2/3/0.10 Jun 12 21:53:01
4000  1a:1a:1a:1a:01:02  ge-2/3/0.10 Jun 12 21:53:01

```

show evpn database extensive

```

user@host> show evpn database extensive
Instance: ALPHA

VN Identifier: 9100, Prefix: 10.0.0.0/24
Source: 10.255.0.2, Rank: 1, Status: Active
Timestamp: May 22 17:16:12 (0x555fc6cc)

VN Identifier: 9100, Prefix: 198.51.100.0/24
Source: 05:00:00:00:64:00:00:23:8c:00, Rank: 1, Status: Active

```

Remote origin: 10.255.0.2
Remote origin: 10.255.0.3
Timestamp: May 22 17:14:20 (0x555fc65c)

VN Identifier: 9100, Prefix: 192.0.2.0/24
Source: irb.0, Rank: 1, Status: Active
Timestamp: May 22 17:16:12 (0x555fc6cc)

VN Identifier: 9100, Prefix: 203.0.113.0/24
Source: 10.255.0.3, Rank: 1, Status: Active
Timestamp: May 22 17:16:12 (0x555fc6cc)

show evpn flood

Syntax	<pre>show evpn flood event-queue <instance <i>instance-name</i>> <logical-system <<i>logical-system-name</i>>> <route <i>route-name</i>></pre>
Release Information	Command introduced in Junos OS Release 14.2 for EX Series switches.
Description	Show Ethernet VPN (EVPN) flooding information.
Options	<p>none—Display brief information about EVPN flooding.</p> <p>brief detail extensive summary—(Optional) Display the specified level of output.</p> <p>event-queue—(Optional) Display the queue of pending EVPN flood events.</p> <p>instance <i>instance-name</i>—(Optional) Display flooding information for the specified routing instance.</p> <p>logical-system <<i>logical-system-name</i>>—(Optional) Display flooding information for the specified logical system or all logical systems.</p> <p>route <i>route-name</i>—(Optional) Display flooding information for the specified route.</p>
Required Privilege Level	view
Related Documentation	<ul style="list-style-type: none"> • Example: Configuring an EVPN with IRB Solution on EX9200 Switches on page 26

show evpn instance

Syntax	<pre>show evpn instance <brief extensive> <backup-forwarder> <designated-forwarder> <esi esi> <instance-name> <neighbor neighbor-address></pre>
Release Information	<p>Command introduced in Junos OS Release 14.1.</p> <p>Command introduced in Junos OS Release 14.2 for EX Series switches.</p>
Description	Show Ethernet VPN (EVPN) routing instance information.
Options	<p>none—Display brief information about the EVPN routing instance.</p> <p>brief extensive—(Optional) Display the specified level of output.</p> <p>backup-forwarder—(Optional) Display IP addresses of all the backup designated forwarder for the Ethernet segment.</p> <p>designated-forwarder—(Optional) Display the IP address of the designated forwarder for the Ethernet segment.</p> <p>esi esi—(Optional) Display brief information about the routing instance associated with the specified ESI value.</p> <p>instance-name—(Optional) Display information about the specified routing instance.</p> <p>neighbor neighbor-address—(Optional) Display the IP address of the EVPN neighbor.</p>
Required Privilege Level	view
List of Sample Output	<p>show evpn instance brief on page 181</p> <p>show evpn instance on page 181</p> <p>show evpn instance extensive (In Junos OS Release 16.1 and earlier) on page 181</p> <p>show evpn instance extensive (In Junos OS Release 16.2 and later) on page 183</p> <p>show evpn instance extensive (Preference-based DF election) on page 184</p> <p>show evpn instance esi backup-forwarder (Instance Name with Ethernet Segment Identifier) on page 185</p> <p>show evpn instance esi designated-forwarder (Instance Name with Ethernet Segment Identifier) on page 185</p>
Output Fields	<p>Table 11 on page 179 lists the output fields for the show evpn instance command. Output fields are listed in the approximate order in which they appear.</p>

Table 11: show evpn instance Output Fields

Field Name	Field Description	Level of Output
Instance	Names of the routing instances.	All Levels
Intfs	Total number of interfaces participating in each routing instance, and number of interfaces that are up.	brief
IRB intfs	Statistics on number of IRB interfaces for each routing instance: <ul style="list-style-type: none"> • Total—Total number of IRB interfaces. • Up—Number of active IRB interfaces. • Nbrs—Number of neighbor IRB interfaces. 	brief
MH ESIs	Number of Ethernet segments per routing instance that connect to a multihomed customer site.	brief
MAC addresses	Number of local and remote MAC addresses for each routing instance.	brief
Route Distinguisher	Unique route distinguisher associated with this routing instance.	none
VLAN ID	VLAN identifier.	none
Label allocation mode	Label allocation policy for the routing instance.	none
Encapsulation type	Encapsulation type for VXLAN EVPN instances (EVIs).	extensive
Per-instance MAC route label	Label of MAC route for each routing instance.	none
DF Election preference	Preference value used for the DF election: <ul style="list-style-type: none"> • Highest preference—Default DF election preference. • Lowest preference—Based on the configuration of the designated-forwarder-preference-least statement at the [edit routing-instance routing-instance-name protocols evpn] hierarchy level. 	extensive
Per-instance multicast route label	Label of multicast route for each routing instance.	none
Total MAC addresses	Total number of local and remote MAC addresses received for each routing instance.	extensive
Default gateway MAC addresses	Number of local and remote MAC addresses serving as a default gateway in the routing instance.	extensive
Number of local interfaces	Number of local interfaces belonging to this routing instance.	extensive

Table 11: show evpn instance Output Fields (*continued*)

Field Name	Field Description	Level of Output
Number of local interfaces up	Number of active local interfaces belonging to this routing instance.	extensive
Interface name	Name of interfaces that belong to this routing instance.	extensive
ESI	Ethernet segment identifier (ESI) value of the interfaces belonging to this routing instance.	extensive
DF Election Algorithm	DF election type: <ul style="list-style-type: none"> • MOD based—Default DF election algorithm based on the modulo operation. • Preference based—DF election based on the configured preference values for an ESI. 	extensive
Preference	Preference value for EVPN Multihoming DF election.	extensive
Mode	Mode of operation for each routing instance: <ul style="list-style-type: none"> • single-homed—Default mode and does not require Ethernet segment values to be configured. • single-active—EVPN active-standby multihoming mode of operation. 	extensive
SH label	Split horizon label used for the active-standby multihoming mode of operation.	extensive
Number of IRB interfaces	Number of IRB interfaces that belong to this routing instance.	extensive
L3 context	Names of routing instances that have the Layer 3 routes installed for an EVPN IRB interface, typically a VRF routing instance.	extensive
Number of neighbors	Number of neighbors connected to this routing instance and their IP addresses.	extensive
MAC address advertisement	Number of MAC address advertisements received from the neighbor.	extensive
MAC+IP address advertisement	Number of MAC and IP address advertisements received from the neighbor.	extensive
Inclusive multicast	Number of inclusive multicast routes received from the neighbor.	extensive
Ethernet auto-discovery	Number of autodiscovery routes per Ethernet segment received from the neighbor.	extensive
Number of ethernet segments	Total number of Ethernet segments for the routing instance.	extensive

Table 11: show evpn instance Output Fields (*continued*)

Field Name	Field Description	Level of Output
Designated forwarder	IP address of the designated forwarder (DF) for the Ethernet segment.	extensive
Backup forwarder	IP address of all the backup designated forwarders or routers that are not designated forwarders for the Ethernet segment. <i>NOTE:</i> Immediately after an EVPN interface Ethernet segment identifier value is changed and the new configuration is committed, the Designated forwarder information changes to DF not elected yet and the Backup forwarder information is not displayed until after the election is complete.	extensive

Sample Output

show evpn instance brief

```

user@host> show evpn instance brief

```

Instance	Intfs		IRB intfs		Nbrs	MH ESIs	MAC addresses	
	Total	Up	Total	Up			Local	Remote
ALPHA	2	2	1	1	2	1	3	4
BETA	2	2	1	1	2	1	2	4
__default_evpn__	0	0	0	0	1	0	0	0

show evpn instance

```

user@host> show evpn instance
Instance: black
Route Distinguisher: 101:101
VLAN ID: 100
Label allocation mode: Per-instance
Per-instance MAC route label: 299776
Per-instance multicast route label: 299792
Number of local interfaces: 1
Number of local interfaces up: 1
Interface name    Static MACs    ESI
cbp-0.0           0              0
Number of neighbors: 1
192.0.2.1
Received routes
MAC address advertisement: 1
Ethernet auto-discovery: 0
Inclusive multicast: 1

```

show evpn instance extensive (In Junos OS Release 16.1 and earlier)

```

user@host> show evpn instance extensive
Instance: ALPHA
Route Distinguisher: 10.255.0.1:100
Encapsulation type: VXLAN
Per-instance MAC route label: 300144
Per-instance multicast route label: 300160
MAC database status
Total MAC addresses: 3 4
Default gateway MAC addresses: 1 2

```

```

Number of local interfaces: 2 (2 up)
  Interface name  ESI                                     Mode          SH label
  ae0.0           00:11:22:33:44:55:66:77:88:99             single-active
  ge-0/0/2.0      00:00:00:00:00:00:00:00:00:00             single-homed
Number of IRB interfaces: 1 (1 up)
  Interface name  L3 context
  irb.0           DELTA
Number of neighbors: 2
  10.255.0.2
    Received routes
      MAC address advertisement:          2
      MAC+IP address advertisement:       3
      Inclusive multicast:                1
      Ethernet auto-discovery:            1
  10.255.0.3
    Received routes
      MAC address advertisement:          2
      MAC+IP address advertisement:       2
      Inclusive multicast:                1
      Ethernet auto-discovery:            0
Number of ethernet segments: 1
  ESI: 00:11:22:33:44:55:66:77:88:99
  Designated forwarder: 10.255.0.1
  Backup forwarder: 10.255.0.2

Instance: BETA
  Route Distinguisher: 10.255.0.1:300
  Encapsulation type: VXLAN
  VLAN ID: 300
  Per-instance MAC route label: 300176
  Per-instance multicast route label: 300192
  MAC database status
    Local Remote
  Total MAC addresses:          3      4
  Default gateway MAC addresses: 1      2
Number of local interfaces: 2 (2 up)
  Interface name  ESI                                     Mode          SH label
  ae1.0           00:00:00:00:00:00:00:00:00:00             single-homed
  ge-0/0/4.0      00:22:44:66:88:00:22:44:66:88             single-active
Number of IRB interfaces: 1 (1 up)
  Interface name  L3 context
  irb.1           DELTA
Number of neighbors: 2
  10.255.0.2
    Received routes
      MAC address advertisement:          2
      MAC+IP address advertisement:       3
      Inclusive multicast:                1
      Ethernet auto-discovery:            1
  10.255.0.3
    Received routes
      MAC address advertisement:          2
      MAC+IP address advertisement:       2
      Inclusive multicast:                1
      Ethernet auto-discovery:            0
Number of ethernet segments: 1
  ESI: 00:22:44:66:88:00:22:44:66:88
  Designated forwarder: 10.255.0.1
  Backup forwarder: 10.255.0.2

Instance: __default__evpn__
  Route Distinguisher: 10.255.0.1:0

```

```

Encapsulation type: VXLAN
VLAN ID: 0
Per-instance MAC route label: 300208
Per-instance multicast route label: 300224
MAC database status
  Total MAC addresses:          Local Remote
  Default gateway MAC addresses: 0      0
Number of local interfaces: 0 (0 up)
Number of IRB interfaces: 0 (0 up)
Number of neighbors: 1
  10.255.0.2
    Received routes
      Ethernet auto-discovery:          0
      Ethernet Segment:                 2
Number of ethernet segments: 0

```

show evpn instance extensive (In Junos OS Release 16.2 and later)

```
user@host> show evpn instance extensive
```

```
user@host> show evpn instance extensive
```

```

Instance: ALPHA
Route Distinguisher: 10.255.0.1:100
Encapsulation type: VXLAN
Per-instance MAC route label: 300144
Per-instance multicast route label: 300160
MAC database status
  Total MAC addresses:          Local Remote
  Default gateway MAC addresses: 1      2
Number of local interfaces: 2 (2 up)
  Interface name  ESI
  ae0.0           00:11:22:33:44:55:66:77:88:99  Mode          SH label
  ge-0/0/2.0      00:00:00:00:00:00:00:00:00:00  single-active
Number of IRB interfaces: 1 (1 up)
  Interface name  L3 context
  irb.0           DELTA
Number of neighbors: 2
  Address      MAC      MAC+IP      AD      IM      ES
  10.255.0.2   2        3          1       1       0
  10.255.0.3   2        2          0       1       0
Number of ethernet segments: 1
  ESI: 00:11:22:33:44:55:66:77:88:99
  Designated forwarder: 10.255.0.1
  Backup forwarder: 10.255.0.2

```

```

Instance: BETA
Route Distinguisher: 10.255.0.1:300
Encapsulation type: VXLAN
VLAN ID: 300
Per-instance MAC route label: 300176
Per-instance multicast route label: 300192
MAC database status
  Total MAC addresses:          Local Remote
  Default gateway MAC addresses: 1      2
Number of local interfaces: 2 (2 up)
  Interface name  ESI
  ae1.0           00:00:00:00:00:00:00:00:00:00  Mode          SH label
  ge-0/0/4.0      00:22:44:66:88:00:22:44:66:88  single-active
Number of IRB interfaces: 1 (1 up)
  Interface name  L3 context

```

```

    irb.1          DELTA
Number of neighbors: 2
  Address          MAC      MAC+IP      AD      IM      ES
  10.255.0.2       2        3        1        1        0
  10.255.0.3       2        2        0        1        0
Number of ethernet segments: 1
  ESI: 00:22:44:66:88:00:22:44:66:88
    Designated forwarder: 10.255.0.1
    Backup forwarder: 10.255.0.2

Instance: __default_evpn__
Route Distinguisher: 10.255.0.1:0
Encapsulation type: VXLAN
VLAN ID: 0
Per-instance MAC route label: 300208
Per-instance multicast route label: 300224
MAC database status
  Total MAC addresses:          Local Remote
  Default gateway MAC addresses: 0      0
Number of local interfaces: 0 (0 up)
Number of IRB interfaces: 0 (0 up)
Number of neighbors: 1
  Address          MAC      MAC+IP      AD      IM      ES
  10.255.0.2       0        0        0        0        2
Number of ethernet segments: 0

```

show evpn instance extensive (Preference-based DF election)

```

user@host> show evpn instance EVPN_1 extensive
Instance: EVPN_1
Route Distinguisher: 1:101
Per-instance MAC route label: 299792
+ DF Election preference: Lowest preference
MAC database status
  MAC advertisements:          Local Remote
  MAC+IP advertisements:      7      15
  Default gateway MAC advertisements: 4      8
Number of local interfaces: 2 (2 up)
  Interface name  ESI                      Mode      Status
  ae101.0         00:11:11:11:11:11:11:11:01  all-active Up
  ae102.0         00:11:11:11:11:11:11:11:02  all-active Up
Number of IRB interfaces: 4 (4 up)
  Interface name  VLAN  VNI  Status  L3 context
  irb.101         101   Up    master
  irb.102         102   Up    master
  irb.103         103   Up    master
  irb.104         104   Up    master
Number of bridge domains: 4
  VLAN  Domain ID  Intfs / up  IRB intf  Mode      MAC sync  IM route
  label
  101          1    1    irb.101  Extended  Enabled   299776
  102          1    1    irb.102  Extended  Enabled   299776
  103          1    1    irb.103  Extended  Enabled   299776
  104          1    1    irb.104  Extended  Enabled   299776
Number of neighbors: 2
  100.0.0.2
  Received routes

```



```

MAC address advertisement: 7
MAC+IP address advertisement: 4
Inclusive multicast: 4
Ethernet auto-discovery: 4
100.0.0.4
Received routes
MAC address advertisement: 8
MAC+IP address advertisement: 8
Inclusive multicast: 4
Ethernet auto-discovery: 0
Number of ethernet segments: 2
ESI: 00:11:11:11:11:11:11:11:11
Status: Resolved by IFL ae0.110
Local interface: ae0.110, Status: Up/Forwarding
Number of remote PEs connected: 2
  Remote PE      MAC label  Aliasing label  Mode
  100.100.100.3    0          0                single-active
  100.100.100.2    0          0                single-active
DF Election Algorithm: Preference based
Designated forwarder: 100.100.100.3, Preference: 200
Backup forwarder: 100.100.100.1, Preference: 800
Backup forwarder: 100.100.100.2, Preference: 400
ESI: 00:11:11:11:11:11:11:11:02
Status: Resolved by IFL ae102.0
Local interface: ae102.0, Status: Up/Forwarding
Number of remote PEs connected: 1
  Remote PE      MAC label  Aliasing label  Mode
  100.0.0.2      299792    299792          all-active
DF Election Algorithm: MOD based
Designated forwarder: 100.0.0.2
Backup forwarder: 100.0.0.1
Last designated forwarder update: Feb 21 22:23:32
Advertised split horizon label: 300800

```

show evpn instance esi backup-forwarder (Instance Name with Ethernet Segment Identifier)

```

user@host> show evpn instance ALPHA esi 00:11:22:33:44:55:66:77:88:99 backup-forwarder
Instance: ALPHA
Number of ethernet segments: 1
ESI: 00:11:22:33:44:55:66:77:88:99
Backup forwarder: 10.255.0.2

```

show evpn instance esi designated-forwarder (Instance Name with Ethernet Segment Identifier)

```

user@host> show evpn instance ALPHA esi 00:11:22:33:44:55:66:77:88:99 designated-forwarder
Instance: ALPHA
Number of ethernet segments: 1
ESI: 00:11:22:33:44:55:66:77:88:99
Designated forwarder: 10.255.0.1

```

show evpn ip-prefix-database

Syntax	<code>show evpn ip-prefix-database</code> <code><extensive></code> <code><direction (imported exported)></code> <code><esi <i>number</i>></code> <code><ethernet-tag <i>number</i>></code> <code><family (inet inet6)></code> <code><gateway <i>ip-address</i>></code> <code><l3-context <i>routing-instance-name</i>></code> <code><nexthop <i>ip-address</i>></code> <code><prefix <i>ip-prefix</i>></code>
Release Information	Command introduced in Junos OS Release 15.1X53-D30 for QFX10002 switches. Statement introduced in Junos OS Release 15.1X53-D60 for QFX10008 and QFX10016 switches. Statement introduced in Junos OS Release 17.3R1 for EX9200 switches.
Description	Display Ethernet VPN (EVPN) database information for imported and exported IP prefixes.
Options	none —Display standard information for all EVPN imported and exported IP prefixes. extensive —Display the specified level of output. direction (imported exported) —(Optional) Display the specified subset of IP prefixes. esi <i>number</i> —(Optional) Display the IP prefix associated with the specified Ethernet segment identifier. ethernet-tag <i>number</i> —(Optional) Display the IP prefix associated with the specified Ethernet tag. family (inet inet6) —(Optional) Display IP prefixes for the specified protocol family. gateway <i>ip-address</i> —(Optional) Display the IP Prefix associated with the overlay gateway IP address. l3-context <i>routing-instance-name</i> —(Optional) Display EVPN IP prefix information for the specified Layer 3 virtual routing and forwarding (VRF) instance. nexthop <i>ip-address</i> —(Optional) Display EVPN IP prefix information for the specified underlay next-hop IP address. prefix <i>ip-prefix</i> —(Optional) Display EVPN database information for the specified IP prefix.
Required Privilege Level	view
Related Documentation	<ul style="list-style-type: none">• show evpn l3-context on page 193• show route table

List of Sample Output [show evpn ip-prefix-database on page 187](#)
[show evpn ip-prefix-database extensive on page 189](#)

Output Fields [Table 12 on page 187](#) lists the output fields for the **show evpn ip-prefix-database** command. Output fields are listed in the approximate order in which they appear.

Table 12: show evpn ip-prefix-database Output Fields

Field Name	Field Description	Level of Output
L3 context	Name of virtual routing and forwarding (VRF) instance.	All levels
EVPN Exported Prefixes Prefix	List of exported IP prefixes.	All levels
EVPN route status	For exported prefixes only, status of route: Created .	level-of-output none
EVPN imported Prefixes	List of imported EVPN IP prefixes.	All levels
Etag	Ethernet tag	level-of-output none
Route distinguisher	IP address identifier for the IP prefix route.	All levels
VNI	Virtual network identifier for the Layer 3 virtual and routing forwarding (VRF) for the customer or tenant domain.	All levels
Router MAC	MAC address associated with the IP prefix.	All levels
Nexthop/Overlay GW/ESI	For imported IP prefixes, next-hop IP address.	level-of-output none
Change flags	Trace flags.	level-of-output extensive
Advertisement mode	For exported IP prefixes, type of next-hop address.	level-of-output extensive
Encapsulation	For exported IP prefixes, type of encapsulation	level-of-output extensive
Remote Advertisements	For imported IP prefixes, route distinguisher identifier, MAC address, virtual network identifier, and BGP next-hop address to remote destination.	level-of-output extensive

Sample Output

show evpn ip-prefix-database

```

user@host > show evpn ip-prefix-database
L3 context: VRF-100

IPv4->EVPN Exported Prefixes
Prefix
EVPN route status

```

100.1.0.0/22	Created
100.1.4.0/22	Created
100.1.8.0/22	Created
100.1.12.0/22	Created
100.1.16.0/22	Created

IPv6->EVPN Exported Prefixes

Prefix	EVPN route status
1234:100:1::/64	Created
1234:100:1:4::/64	Created
1234:100:1:8::/64	Created
1234:100:1:12::/64	Created
1234:100:1:16::/64	Created

EVPN->IPv4 Imported Prefixes

Prefix	Etag	IP route status
100.2.0.0/22	0	Created
Route distinguisher	VNI/Label	Router MAC
10.255.2.1:100	9101	00:05:86:e1:03:f0
10.255.2.2:100	9101	00:05:86:d0:a6:f0
100.2.4.0/22	0	Created
Route distinguisher	VNI/Label	Router MAC
10.255.2.1:100	9101	00:05:86:e1:03:f0
10.255.2.2:100	9101	00:05:86:d0:a6:f0
100.2.8.0/22	0	Created
Route distinguisher	VNI/Label	Router MAC
10.255.2.1:100	9101	00:05:86:e1:03:f0
10.255.2.2:100	9101	00:05:86:d0:a6:f0
100.2.12.0/22	0	Created
Route distinguisher	VNI/Label	Router MAC
10.255.2.1:100	9101	00:05:86:e1:03:f0
10.255.2.2:100	9101	00:05:86:d0:a6:f0
100.2.16.0/22	0	Created
Route distinguisher	VNI/Label	Router MAC
10.255.2.1:100	9101	00:05:86:e1:03:f0
10.255.2.2:100	9101	00:05:86:d0:a6:f0

EVPN->IPv6 Imported Prefixes

Prefix	Etag	IP route status
1234:100:2::/64	0	Created
Route distinguisher	VNI/Label	Router MAC
10.255.2.1:100	9101	00:05:86:e1:03:f0
10.255.2.2:100	9101	00:05:86:d0:a6:f0
1234:100:2:4::/64	0	Created
Route distinguisher	VNI/Label	Router MAC
10.255.2.1:100	9101	00:05:86:e1:03:f0
10.255.2.2:100	9101	00:05:86:d0:a6:f0
1234:100:2:8::/64	0	Created
Route distinguisher	VNI/Label	Router MAC
10.255.2.1:100	9101	00:05:86:e1:03:f0
10.255.2.2:100	9101	00:05:86:d0:a6:f0
1234:100:2:12::/64	0	Created
Route distinguisher	VNI/Label	Router MAC
10.255.2.1:100	9101	00:05:86:e1:03:f0
10.255.2.2:100	9101	00:05:86:d0:a6:f0
1234:100:2:16::/64	0	Created
Route distinguisher	VNI/Label	Router MAC
10.255.2.1:100	9101	00:05:86:e1:03:f0
10.255.2.2:100	9101	00:05:86:d0:a6:f0

show evpn ip-prefix-database extensive

```
user@host> show evpn ip-prefix-database extensive
L3 context: VRF-100
```

IPv4->EVPN Exported Prefixes

```
Prefix: 100.1.0.0/22
  EVPN route status: Created
  Change flags: 0x0
  Advertisement mode: Direct nexthop
  Encapsulation: VXLAN
  VNI: 9100
  Router MAC: 00:05:86:28:90:f0
```

```
Prefix: 100.1.4.0/22
  EVPN route status: Created
  Change flags: 0x0
  Advertisement mode: Direct nexthop
  Encapsulation: VXLAN
  VNI: 9100
  Router MAC: 00:05:86:28:90:f0
```

```
Prefix: 100.1.8.0/22
  EVPN route status: Created
  Change flags: 0x0
  Advertisement mode: Direct nexthop
  Encapsulation: VXLAN
  VNI: 9100
  Router MAC: 00:05:86:28:90:f0
```

```
Prefix: 100.1.12.0/22
  EVPN route status: Created
  Change flags: 0x0
  Advertisement mode: Direct nexthop
  Encapsulation: VXLAN
  VNI: 9100
  Router MAC: 00:05:86:28:90:f0
```

```
Prefix: 100.1.16.0/22
  EVPN route status: Created
  Change flags: 0x0
  Advertisement mode: Direct nexthop
  Encapsulation: VXLAN
  VNI: 9100
  Router MAC: 00:05:86:28:90:f0
```

IPv6->EVPN Exported Prefixes

```
Prefix: 1234:100:1::/64
  EVPN route status: Created
  Change flags: 0x0
  Advertisement mode: Direct nexthop
  Encapsulation: VXLAN
  VNI: 9100
  Router MAC: 00:05:86:28:90:f0
```

```
Prefix: 1234:100:1:4::/64
  EVPN route status: Created
  Change flags: 0x0
  Advertisement mode: Direct nexthop
```

Encapsulation: VXLAN
VNI: 9100
Router MAC: 00:05:86:28:90:f0

Prefix: 1234:100:1:8::/64
EVPN route status: Created
Change flags: 0x0
Advertisement mode: Direct nexthop
Encapsulation: VXLAN
VNI: 9100
Router MAC: 00:05:86:28:90:f0

Prefix: 1234:100:1:12::/64
EVPN route status: Created
Change flags: 0x0
Advertisement mode: Direct nexthop
Encapsulation: VXLAN
VNI: 9100
Router MAC: 00:05:86:28:90:f0

Prefix: 1234:100:1:16::/64
EVPN route status: Created
Change flags: 0x0
Advertisement mode: Direct nexthop
Encapsulation: VXLAN
VNI: 9100
Router MAC: 00:05:86:28:90:f0

EVPN->IPv4 Imported Prefixes

Prefix: 100.2.0.0/22, Ethernet tag: 0
IP route status: Created
Change flags: 0x0
Remote advertisements:
Route Distinguisher: 10.255.2.1:100
VNI: 9101
Router MAC: 00:05:86:e1:03:f0
BGP nexthop address: 10.255.2.1
Route Distinguisher: 10.255.2.2:100
VNI: 9101
Router MAC: 00:05:86:d0:a6:f0
BGP nexthop address: 10.255.2.2

Prefix: 100.2.4.0/22, Ethernet tag: 0
IP route status: Created
Change flags: 0x0
Remote advertisements:
Route Distinguisher: 10.255.2.1:100
VNI: 9101
Router MAC: 00:05:86:e1:03:f0
BGP nexthop address: 10.255.2.1
Route Distinguisher: 10.255.2.2:100
VNI: 9101
Router MAC: 00:05:86:d0:a6:f0
BGP nexthop address: 10.255.2.2

Prefix: 100.2.8.0/22, Ethernet tag: 0
IP route status: Created
Change flags: 0x0
Remote advertisements:
Route Distinguisher: 10.255.2.1:100

```

VNI: 9101
Router MAC: 00:05:86:e1:03:f0
BGP nexthop address: 10.255.2.1
Route Distinguisher: 10.255.2.2:100
VNI: 9101
Router MAC: 00:05:86:d0:a6:f0
BGP nexthop address: 10.255.2.2

Prefix: 100.2.12.0/22, Ethernet tag: 0
IP route status: Created
Change flags: 0x0
Remote advertisements:
  Route Distinguisher: 10.255.2.1:100
  VNI: 9101
  Router MAC: 00:05:86:e1:03:f0
  BGP nexthop address: 10.255.2.1
  Route Distinguisher: 10.255.2.2:100
  VNI: 9101
  Router MAC: 00:05:86:d0:a6:f0
  BGP nexthop address: 10.255.2.2

Prefix: 100.2.16.0/22, Ethernet tag: 0
IP route status: Created
Change flags: 0x0
Remote advertisements:
  Route Distinguisher: 10.255.2.1:100
  VNI: 9101
  Router MAC: 00:05:86:e1:03:f0
  BGP nexthop address: 10.255.2.1
  Route Distinguisher: 10.255.2.2:100
  VNI: 9101
  Router MAC: 00:05:86:d0:a6:f0
  BGP nexthop address: 10.255.2.2

EVPN->IPv6 Imported Prefixes

Prefix: 1234:100:2::/64, Ethernet tag: 0
IP route status: Created
Change flags: 0x0
Remote advertisements:
  Route Distinguisher: 10.255.2.1:100
  VNI: 9101
  Router MAC: 00:05:86:e1:03:f0
  BGP nexthop address: 10.255.2.1
  Route Distinguisher: 10.255.2.2:100
  VNI: 9101
  Router MAC: 00:05:86:d0:a6:f0
  BGP nexthop address: 10.255.2.2

Prefix: 1234:100:2:4::/64, Ethernet tag: 0
IP route status: Created
Change flags: 0x0
Remote advertisements:
  Route Distinguisher: 10.255.2.1:100
  VNI: 9101
  Router MAC: 00:05:86:e1:03:f0
  BGP nexthop address: 10.255.2.1
  Route Distinguisher: 10.255.2.2:100
  VNI: 9101
  Router MAC: 00:05:86:d0:a6:f0
  BGP nexthop address: 10.255.2.2

```

```
Prefix: 1234:100:2:8::/64, Ethernet tag: 0
IP route status: Created
Change flags: 0x0
Remote advertisements:
  Route Distinguisher: 10.255.2.1:100
  VNI: 9101
  Router MAC: 00:05:86:e1:03:f0
  BGP nexthop address: 10.255.2.1
  Route Distinguisher: 10.255.2.2:100
  VNI: 9101
  Router MAC: 00:05:86:d0:a6:f0
  BGP nexthop address: 10.255.2.2

Prefix: 1234:100:2:12::/64, Ethernet tag: 0
IP route status: Created
Change flags: 0x0
Remote advertisements:
  Route Distinguisher: 10.255.2.1:100
  VNI: 9101
  Router MAC: 00:05:86:e1:03:f0
  BGP nexthop address: 10.255.2.1
  Route Distinguisher: 10.255.2.2:100
  VNI: 9101
  Router MAC: 00:05:86:d0:a6:f0
  BGP nexthop address: 10.255.2.2

Prefix: 1234:100:2:16::/64, Ethernet tag: 0
IP route status: Created
Change flags: 0x0
Remote advertisements:
  Route Distinguisher: 10.255.2.1:100
  VNI: 9101
  Router MAC: 00:05:86:e1:03:f0
  BGP nexthop address: 10.255.2.1
  Route Distinguisher: 10.255.2.2:100
  VNI: 9101
  Router MAC: 00:05:86:d0:a6:f0
  BGP nexthop address: 10.255.2.2
```


show evpn l3-context

Syntax `show evpn l3-context`
`<brief | extensive>`
`<routing-instance-name>`

Release Information Statement introduced in Junos OS Release 15.1X53-D30 for QFX10002 switches.
Statement introduced in Junos OS Release 15.1X53-D60 for QFX10008 and QFX10016 switches in Junos OS Release 15.1X53-D60.
Statement introduced in Junos OS Release 17.3R1 for EX9200 switches.

Description Display EVPN database information for Layer 3 virtual routing and forwarding (VRF) instances.

Options **none**—Display standard EVPN database information for all Layer 3 VRF instances.
brief | extensive—(Optional) Display the specified level of output.
routing-instance-name—(Optional) Display EVPN database information for the specified Layer 3 VRF instance.

Additional Information

Required Privilege Level view

Related Documentation • [show evpn ip-prefix-database on page 186](#)

List of Sample Output [show evpn l3-context on page 194](#)
[show evpn l3-context extensive on page 194](#)

Output Fields [Table 13 on page 193](#) lists the output fields for the **show evpn l3-context** command. Output fields are listed in the approximate order in which they appear.

Table 13: show evpn l3-context Output Fields

Field Name	Field Description	Level of Output
L3 context	Name of VRF instance.	All levels
Type	Status of VRF instance	All levels
Fwd	Type of forwarding route	level-of-output none
Encap	Type of encapsulation.	level-of-output none
VNI	Virtual network identifier for VRF instance.	All levels

Table 13: show evpn l3-context Output Fields (*continued*)

Field Name	Field Description	Level of Output
Router MAC/GW intf	For a pure type-5 route, the router MAC provides next-hop reachability information and the MAC address of the sending device For a standard type-5 route, the GW interface provides the next-hop route.	All levels
Advertisement mode	Type of forwarding route.	level-of-output Extensive
IPv4 source VTEP address	IPv4 source address for the Virtual Extensible LAN (VXLAN) tunnel.	level-of-output Extensive
IPv6 source VTEP address	IPv6 source address for the Virtual Extensible LAN (VXLAN) tunnel.	level-of-output Extensive
IP->EVPN export policy	Name of export routing policy applied to forwarded IP routes.	level-of-output extensive
Flags	Enabled trace flags.	level-of-output extensive
Change flags	Changed trace flags.	level-of-output extensive
Composite nexthop	Status of next-hop route: Enabled or Disabled .	level-of-output extensive
Route distinguisher	Route distinguisher identifier of the VRF instance	level-of-output extensive
Reference count		level-of-output extensive

Sample Output

show evpn l3-context

```

user@DC1_SPINE1_RE> show evpn l3-context
L3 context          Type Fwd  Encap VNI/Label Router MAC/GW intf
VRF-100             Cfg  Sym  VXLAN 9100      00:05:86:28:90:f0
VRF-200             Cfg  Sym  VXLAN 9200      00:05:86:28:90:f0

```

show evpn l3-context extensive

```

user@DC1_SPINE1_RE> show evpn l3-context extensive
L3 context: VRF-100
Type: Configured
Advertisement mode: Direct nexthop, Router MAC: 00:05:86:28:90:f0
Encapsulation: VXLAN, VNI: 9100
IPv4 source VTEP address: 10.255.1.1
IPv6 source VTEP address: abcd::128:102:242:145

```

IP->EVPN export policy: EVPN-TYPE5-EXPORT-VRF-100
Flags: 0x19 <Configured IRB-MAC New>
Change flags: 0xe <Export-Policy Fwd-Mode Encap>
Composite nexthop support: Enabled
Route Distinguisher: 10.255.1.1:100
Reference count: 21

L3 context: VRF-200

Type: Configured
Advertisement mode: Direct nexthop, Router MAC: 00:05:86:28:90:f0
Encapsulation: VXLAN, VNI: 9200
IPv4 source VTEP address: 10.255.1.1
IPv6 source VTEP address: abcd::128:102:242:145
IP->EVPN export policy: EVPN-TYPE5-EXPORT-VRF-200
Flags: 0x19 <Configured IRB-MAC New>
Change flags: 0xe <Export-Policy Fwd-Mode Encap>
Composite nexthop support: Enabled
Route Distinguisher: 10.255.1.1:200
Reference count: 21

show evpn mac-table

Syntax	<code>show evpn mac-table</code> <code><age></code> <code><address></code> <code><brief count detail extensive summary></code> <code><instance <i>instance-name</i>></code> <code><interface <i>interface-name</i>></code> <code><isid <i><isid></i>></code> <code><logical-system <i><logical-system-name></i>></code> <code><vlan-id <i>vlan-id</i>></code>
Release Information	Command introduced in Junos OS Release 14.2 for EX Series switches.
Description	Show Ethernet VPN (EVPN) MAC table information.
Options	<p>none—Display brief information about the EVPN MAC table.</p> <p>age— (Optional) Display age of a single mac-address.</p> <p>address—(Optional) Display MAC table information for the specified MAC address.</p> <p>brief count detail extensive summary—(Optional) Display the specified level of output.</p> <p>instance <i>instance-name</i>—(Optional) Display MAC table information for a specific routing instance.</p> <p>isid <i><isid></i>—(Optional) Display MAC table information for the specified ISID or all ISIDs.</p> <p>logical-system <i><logical-system-name></i>—(Optional) Display MAC table information for the specified logical system or all logical systems.</p> <p>vlan-id <i>vlan-id</i>—(Optional) Display MAC table information for the specified VLAN.</p>
Required Privilege Level	view
Related Documentation	<ul style="list-style-type: none">• Example: Configuring an EVPN with IRB Solution on EX9200 Switches on page 26

show evpn nd-table

Syntax	show evpn nd-table <address> <brief count detail extensive> <instance <i>instance-name</i> >
Release Information	Command introduced in Junos OS Release 16.2. Command introduced in Junos OS Release 17.1 for MX series Routers. Command introduced in Junos OS Release 17.3R1 for EX series switches.
Description	Show Ethernet VPN (EVPN) Network Discovery Protocol (NDP) entries associated with learned MAC addresses.
Options	none —Display brief information about the EVPN NDP table. address —(Optional) Display NDP information for the specified MAC address. brief count detail extensive —(Optional) Display the specified level of output. instance <instance-name> —(Optional) Display NDP information for the specified routing instance .
Required Privilege Level	view
List of Sample Output	show evpn nd-table on page 198 show evpn nd-table 00:05:86:a0:dc:f0 (MAC address) on page 198 show evpn nd-table brief on page 198 show evpn nd-table detail on page 198 show evpn nd-table count on page 198 show evpn and-table extensive on page 199 show evpn ndp-table instance evpn1 on page 199
Output Fields	Table 14 on page 197 lists the output fields for the show evpn nd-table command. Output fields are listed in the approximate order in which they appear.

Table 14: show evpn nd-table Output Fields

Field Name	Field Description	Level of Output
INET address	The INET address related to the INET entries that are added to the NDP table.	All levels
MAC address	MAC addresses learned through NDP.	brief, detail, extensive, instance, mac-address,,

Table 14: show evpn nd-table Output Fields (*continued*)

Field Name	Field Description	Level of Output
Logical Interface	Logical interface associated with the routing instance in which the NDP INET address is learned.	brief, instance, mac-address,,
Routing instance	Routing instance in which the NDP INET address is learned.	all levels
Bridging domain	Bridging domain in which the NDP INET address is learned.	all levels
Learning interface	Interface on which the NDP INET address is learned.	detail, extensive
Count	Indicates the number of NDP INET addresses learned in a routing instance in a bridge domain.	count

Sample Output

show evpn nd-table

```

user@host> show evpn nd-table
INET          MAC          Logical      Routing      Bridging
address       address      interface    instance     domain
8001::2       00:05:86:a0:dc:f0  irb.0        evpn_1       __evpn_1__

```

show evpn nd-table 00:05:86:a0:dc:f0 (MAC address)

```

user@host> show evpn nd-table 00:05:86:a0:dc:f0
INET          MAC          Logical      Routing      Bridging
address       address      interface    instance     domain
8001::2       00:05:86:a0:dc:f0  irb.0        evpn1        __evpn1__

```

show evpn nd-table brief

```

user@host> show evpn nd-table brief
INET          MAC          Logical      Routing      Bridging
address       address      interface    instance     domain
8001::2       00:05:86:a0:dc:f0  irb.0        evpn1        __evpn1__

```

show evpn nd-table detail

```

user@host> show evpn nd-table detail

INET address: 8001::2
MAC address: 00:05:86:a0:dc:f0
Routing instance: evpn1
Bridging domain: __evpn1__
Learning interface: irb.0

```

show evpn nd-table count

```

user@switch> show evpn nd-table count
1 NDP INET addresses learned in routing instance evpn1 bridge domain __evpn1__

```

show evpn and-table extensive

```
user@host> show evpn nd-table extensive
```

```
INET address: 8001::2
MAC address: 00:05:86:a0:dc:f0
Routing instance: evpn1
Bridging domain: __evpn1__
Learning interface: irb.0
```

show evpn ndp-table instance evpn1

```
user@host> show evpn arp-table instance evpn1
```

INET address	MAC address	Logical interface	Routing instance	Bridging domain
8001::2	00:05:86:a0:dc:f0	irb.0	evpn1	__evpn1__

show evpn peer-gateway-macs

Syntax	<code>show evpn peer-gateway-macs</code> <code><address></code> <code><instance <i>instance-name</i>></code>
Release Information	Command introduced in Junos OS Release 14.2 for EX Series switches.
Description	Show Ethernet VPN (EVPN) peer gateway MAC information.
Options	none —Display brief information about the EVPN peer gateway MAC. address —(Optional) Display peer gateway information for the specified MAC address. instance <i>instance-name</i> —(Optional) Display peer gateway MAC information for the specified routing instance.
Required Privilege Level	view
Related Documentation	<ul style="list-style-type: none">• Example: Configuring an EVPN with IRB Solution on EX9200 Switches on page 26

show vlans evpn nd-table

Syntax show vlans evpn nd-table
 <brief | detail | extensive>
 <count>
 <instance-name>
 <mac-address>
 <vlan-name>

Release Information Command introduced in Junos OS Release 17.3R1 for EX Series Switches.

Description Display information about INET entries associated with MAC addresses learned through Network Discovery Protocol (NDP).

Options **none**—Display information for all INET entries.

brief | detail | extensive—(Optional) Display the specified level of output.

count—(Optional) Display the number of INET addresses learned in a routing instance.

instance—(Optional) Display information for a specified instance.

mac-address—(Optional) Display information for a specified MAC address.

vlan-name (all)—(Optional) Display information for a specified VLAN or for all VLANs.

Required Privilege Level view

List of Sample Output [show vlans evpn nd-table brief on page 202](#)
[show vlans evpn nd-table detail on page 202](#)
[show vlans evpn nd-table extensive on page 202](#)
[show vlans evpn nd-table count on page 202](#)
[show vlans evpn nd-table instance evpn1 on page 202](#)
[show vlans evpn nd-table instance 00:05:86:90:bd:f0 \(MAC address\) on page 203](#)
[show vlans evpn nd-table vlan-name vlan10 on page 203](#)

Output Fields [Table 15 on page 201](#) lists the output fields for the **show vlans evpn nd-table** command. Output fields are listed in the approximate order in which they appear.

Table 15: show vlans evpn nd-table Output Fields

Field Name	Field Description	Level of Output
INET address	The INET address related to the INET entries that are added to the NDP table.	All levels
MAC address	MAC addresses learned through NDP.	brief, detail, extensive, instance, mac-address,vlan-name,,

Table 15: show vlans evpn nd-table Output Fields (*continued*)

Field Name	Field Description	Level of Output
Logical Interface	Logical interface associated with the routing instance in which the NDP INET address is learned.	brief, instance, mac-addressvlan-name,,
Routing instance	Routing instance in which the NDP INET address is learned.	all levels
Bridging domain	Bridging domain in which the NDP INET address is learned.	all levels
Learning interface	Interface on which the NDP INET address is learned.	detail, extensive
Count	Indicates the number of NDP INET addresses learned in a routing instance in a bridge domain.	count

Sample Output

show vlans evpn nd-table brief

```

user@switch> show vlans evpn nd-table brief
INET          MAC          Logical      Routing      Bridging
address       address      interface    instance     domain
8002::2       00:05:86:a0:d5:00  irb.0        evpn1        vlan10

```

show vlans evpn nd-table detail

```

user@switch> show vlans evpn nd-table detail
INET address: 8002::2
MAC address: 00:05:86:a0:d5:00
Routing instance: evpn1
Bridging domain: vlan10
Learning interface: irb.0

```

show vlans evpn nd-table extensive

```

user@switch> show vlans evpn nd-table extensive
INET address: 8002::2
MAC address: 00:05:86:a0:d5:00
Routing instance: evpn1
Bridging domain: vlan10
Learning interface: irb.0

```

show vlans evpn nd-table count

```

user@switch> show vlans evpn nd-table count
1 ND INET addresses learned in routing instance evpn1 bridge domain vlan10

```

show vlans evpn nd-table instance evpn1

```

user@switch> show vlans evpn nd-table instance evpn1
INET          MAC          Logical      Routing      Bridging
address       address      interface    instance     domain
8002::2       00:05:86:a0:d5:00
                evpn1        vlan10
irb.0

```

show vlans evpn nd-table instance 00:05:86:90:bd:f0 (MAC address)

```

user@switch> show vlans evpn nd-table
INET          MAC          Logical   Routing   Bridging
address       address      interface instance  domain
8002::2
              00:05:86:a0:d5:00
              evpn1      __evpn1__

```

show vlans evpn nd-table vlan-name vlan10

```

user@switch> show vlans evpn nd-table vlan-name vlan10
INET          MAC          Logical   Routing   Bridging
address       address      interface instance  domain
8002::2
              00:05:86:a0:d5:00
              evpn1      vlan10

```

