



---

Junos<sup>®</sup> OS

# EVPNs Feature Guide for Routing Devices

Release  
13.2



---

Published: 2013-07-31

Juniper Networks, Inc.  
1194 North Mathilda Avenue  
Sunnyvale, California 94089  
USA  
408-745-2000  
[www.juniper.net](http://www.juniper.net)

This product includes the Envoy SNMP Engine, developed by Epilogue Technology, an Integrated Systems Company. Copyright © 1986-1997, Epilogue Technology Corporation. All rights reserved. This program and its documentation were developed at private expense, and no part of them is in the public domain.

This product includes memory allocation software developed by Mark Moraes, copyright © 1988, 1989, 1993, University of Toronto.

This product includes FreeBSD software developed by the University of California, Berkeley, and its contributors. All of the documentation and software included in the 4.4BSD and 4.4BSD-Lite Releases is copyrighted by the Regents of the University of California. Copyright © 1979, 1980, 1983, 1986, 1988, 1989, 1991, 1992, 1993, 1994. The Regents of the University of California. All rights reserved.

GateD software copyright © 1995, the Regents of the University. All rights reserved. Gate Daemon was originated and developed through release 3.0 by Cornell University and its collaborators. Gated is based on Kirton's EGP, UC Berkeley's routing daemon (routed), and DCN's HELLO routing protocol. Development of Gated has been supported in part by the National Science Foundation. Portions of the GateD software copyright © 1988, Regents of the University of California. All rights reserved. Portions of the GateD software copyright © 1991, D. L. S. Associates.

This product includes software developed by Maker Communications, Inc., copyright © 1996, 1997, Maker Communications, Inc.

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

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

Products made or sold by Juniper Networks or components thereof might be covered by one or more of the following patents that are owned by or licensed to Juniper Networks: U.S. Patent Nos. 5,473,599, 5,905,725, 5,909,440, 6,192,051, 6,333,650, 6,359,479, 6,406,312, 6,429,706, 6,459,579, 6,493,347, 6,538,518, 6,538,899, 6,552,918, 6,567,902, 6,578,186, and 6,590,785.

*Junos® OS EVPNs Feature Guide for Routing Devices*

13.2

Copyright © 2013, Juniper Networks, Inc.  
All rights reserved.

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

## YEAR 2000 NOTICE

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

## END USER LICENSE AGREEMENT

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

# Table of Contents

	About the Documentation .....	ix
	Documentation and Release Notes .....	ix
	Supported Platforms .....	ix
	Using the Examples in This Manual .....	ix
	Merging a Full Example .....	x
	Merging a Snippet .....	x
	Documentation Conventions .....	xi
	Documentation Feedback .....	xiii
	Requesting Technical Support .....	xiii
	Self-Help Online Tools and Resources .....	xiii
	Opening a Case with JTAC .....	xiv
<b>Part 1</b>	<b>Overview</b>	
<b>Chapter 1</b>	<b>Introduction to EVPNs .....</b>	<b>3</b>
	EVPN Overview .....	3
<b>Chapter 2</b>	<b>EVPN Standards .....</b>	<b>5</b>
	Supported EVPN Standards .....	5
<b>Part 2</b>	<b>Configuration</b>	
<b>Chapter 3</b>	<b>Configuring EVPNs .....</b>	<b>9</b>
	Configuring an IGP on the PE and P Routers .....	9
	Configuring IBGP Sessions Between PE Routers in VPNs .....	9
	Configuring a Signaling Protocol and LSPs for VPNs .....	11
	Using LDP for VPN Signaling .....	11
	Using RSVP for VPN Signaling .....	13
	Configuring EVPN Routing Instances .....	14
	Tracing EVPN Traffic and Operations .....	17
<b>Chapter 4</b>	<b>EVPN Configuration Statements .....</b>	<b>19</b>
	evpn .....	20
	family (Protocols BGP) .....	21
	ignore-encapsulation-mismatch .....	24
	instance-type .....	25
	interface (EVPN Routing Instances) .....	27
	interface (Routing Instances) .....	28
	interface-mac-limit .....	29
	label-allocation .....	30
	mac-statistics .....	31
	mac-table-size .....	32

no-mac-learning .....	33
packet-action .....	35
routing-instances .....	36
static-mac .....	37
traceoptions (Protocols EVPN) .....	38
vlan-id (routing instance) .....	40

## Part 3

## Index

Index .....	43
-------------	----

# List of Figures

Part 1	Overview	
Chapter 1	Introduction to EVPNs .....	3
	Figure 1: EVPN Connecting Data Center 1 and Data Center 2 .....	3



# List of Tables

<b>About the Documentation</b> .....	<b>ix</b>
Table 1: Notice Icons .....	xi
Table 2: Text and Syntax Conventions .....	xi





# About the Documentation

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

## 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:

- MX 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 the *CLI User Guide*.

## Documentation Conventions

Table 1 on page xi 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.

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

Table 2: Text and Syntax Conventions

Convention	Description	Examples
<b>Bold text like this</b>	Represents text that you type.	To enter configuration mode, type the <b>configure</b> command:  user@host> <b>configure</b>
Fixed-width text like this	Represents output that appears on the terminal screen.	user@host> <b>show chassis alarms</b>  No alarms currently active

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

Convention	Description	Examples
<i>Italic text like this</i>	<ul style="list-style-type: none"> <li>Introduces or emphasizes important new terms.</li> <li>Identifies book names.</li> <li>Identifies RFC and Internet draft titles.</li> </ul>	<ul style="list-style-type: none"> <li>A policy <i>term</i> is a named structure that defines match conditions and actions.</li> <li><i>Junos OS System Basics Configuration Guide</i></li> <li>RFC 1997, <i>BGP Communities Attribute</i></li> </ul>
<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@# <b>set system domain-name</b> <i>domain-name</i>
<b>Text like this</b>	Represents names of configuration statements, commands, files, and directories; configuration hierarchy levels; or labels on routing platform components.	<ul style="list-style-type: none"> <li>To configure a stub area, include the <b>stub</b> statement at the [edit protocols ospf area area-id] hierarchy level.</li> <li>The console port is labeled <b>CONSOLE</b>.</li> </ul>
< > (angle brackets)	Enclose optional keywords or variables.	<b>stub</b> <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.	<b>broadcast   multicast</b>  ( <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.	<b>rsvp { # Required for dynamic MPLS only</b>
[ ] (square brackets)	Enclose a variable for which you can substitute one or more values.	<b>community name members [</b> <i>community-ids</i> <b>]</b>
Indentation and braces ( { } )	Identify 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.	
<b>GUI Conventions</b>		
<b>Bold text like this</b>	Represents graphical user interface (GUI) items you click or select.	<ul style="list-style-type: none"> <li>In the Logical Interfaces box, select <b>All Interfaces</b>.</li> <li>To cancel the configuration, click <b>Cancel</b>.</li> </ul>
> (bold right angle bracket)	Separates levels in a hierarchy of menu selections.	In the configuration editor hierarchy, select <b>Protocols&gt;Ospf</b> .

## Documentation Feedback

---

We encourage you to provide feedback, comments, and suggestions so that we can improve the documentation. You can send your comments to [techpubs-comments@juniper.net](mailto:techpubs-comments@juniper.net), or fill out the documentation feedback form at <https://www.juniper.net/cgi-bin/docbugreport/>. If you are using e-mail, be sure to include the following information with your comments:

- Document or topic name
- URL or page number
- Software release 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 JNASC 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: <https://www.juniper.net/alerts/>

- 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

# Overview

- [Introduction to EVPNs on page 3](#)
- [EVPN Standards on page 5](#)





## CHAPTER 1


# Introduction to EVPNs

- [EVPN Overview on page 3](#)

## EVPN Overview

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 is comprised of 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, an MX Series 3D Universal Edge Router can be configured to act as an MES. You can deploy multiple EVPNs within the service provider's network, each providing network connectivity to customers while ensuring that the traffic sharing that network remains private. [Figure 1 on page 3](#) illustrates a typical EVPN deployment. Traffic from Data Center 1 is transported over the Service Provider's network through MES1 to MES2 and then onto Data Center 2. DCS1, DCS2, DCS3, and DCS4 are the data center switches.

Figure 1: EVPN Connecting Data Center 1 and Data Center 2



The MESs are interconnected within the service provider's network 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 VLAN for each MAC VPN
- Automatic route distinguishers
- Active Standby multihoming

The following features are not supported for EVPNs:

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

**Related  
Documentation**

- [Supported EVPN Standards on page 5](#)

## CHAPTER 2

# EVPN Standards

- [Supported EVPN Standards on page 5](#)

### Supported EVPN Standards

---

Junos OS substantially 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*
- Internet draft draft-raggarwa-sajassi-l2vpn-evpn-01.txt, *BGP MPLS Based EVPN*

#### Related Documentation

- [EVPN Overview on page 3](#)



## PART 2

# Configuration

- [Configuring EVPNs on page 9](#)
- [EVPN Configuration Statements on page 19](#)



## CHAPTER 3

# Configuring EVPNs

- [Configuring an IGP on the PE and P Routers on page 9](#)
- [Configuring IBGP Sessions Between PE Routers in VPNs on page 9](#)
- [Configuring a Signaling Protocol and LSPs for VPNs on page 11](#)
- [Configuring EVPN Routing Instances on page 14](#)
- [Tracing EVPN Traffic and Operations on page 17](#)

## Configuring an IGP on the PE and P Routers

---

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

- *Example: Configuring IS-IS*
- *Examples: Configuring Static Routes*
- *OSPF Feature Guide for Routing Devices*

## Configuring IBGP Sessions Between PE Routers in VPNs

---

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. If you are using RSVP signaling, this IP address is the same address you specify in the **to** statement at the **[edit mpls label-switched-path *lsp-path-name*]** hierarchy level when you configure the MPLS LSP.

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

## Configuring a Signaling Protocol and LSPs for VPNs

For VPNs to function, you must enable a signaling protocol, either the LDP or RSVP on the provider edge (PE) routers and on the provider (P) routers. You also need to configure label-switched paths (LSPs) between the ingress and egress routers. In a typical VPN configuration, you need to configure LSPs from each PE router to all of the other PE routers participating in the VPN in a full mesh.



**NOTE:** As with any configuration involving MPLS, you cannot configure any of the core-facing interfaces on the PE routers over dense Fast Ethernet PICs.

To enable a signaling protocol, perform the steps in one of the following sections:

- [Using LDP for VPN Signaling on page 11](#)
- [Using RSVP for VPN Signaling on page 13](#)

### 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 service provider's 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 for Routing Devices* and *IS-IS Feature Guide for Routing Devices*.

## Using RSVP for VPN Signaling

To use RSVP for VPN signaling, perform the following steps:

1. On each PE router, configure traffic engineering.

To do this, you must configure an interior gateway protocol (IGP) that supports traffic engineering (either IS-IS or OSPF) and enable traffic engineering support for that protocol.

To enable OSPF traffic engineering support, include the **traffic-engineering** statement at the **[edit protocols ospf]** hierarchy level:

```
[edit protocols ospf]
traffic-engineering {
  shortcuts;
}
```

For IS-IS, traffic engineering support is enabled by default.

2. On each PE and P router, enable RSVP on the interfaces that participate in the label-switched path (LSP).

On the PE router, these interfaces are the ingress and egress points to the LSP. On the P router, these interfaces connect the LSP between the PE routers. Do not enable RSVP on the interface between the PE and the CE routers, because this interface is not part of the LSP.

To configure RSVP on the PE and P routers, include the **interface** statement at the **[edit protocols rsvp]** hierarchy level. Include one **interface** statement for each interface on which you are enabling RSVP.

```
[edit protocols]
rsvp {
  interface interface-name;
  interface interface-name;
}
```

3. On each PE router, configure an MPLS LSP to the PE router that is the LSP's egress point.

To do this, include the **interface** and **label-switched-path** statements at the **[edit protocols mpls]** hierarchy level:

```
[edit protocols]
mpls {
  interface interface-name;
  label-switched-path path-name {
    to ip-address;
  }
}
```

In the **to** statement, specify the address of the LSP's egress point, which is an address on the remote PE router.

In the **interface** statement, specify the name of the interface (both the physical and logical portions). Include one **interface** statement for the interface associated with the LSP.

When you configure the logical portion of the same interface at the **[edit interfaces]** hierarchy level, you must also configure the **family inet** and **family mpls** statements:

```
[edit interfaces]
interface-name {
  unit logical-unit-number {
    family inet;
    family mpls;
  }
}
```

4. On all P routers that participate in the LSP, enable MPLS by including the **interface** statement at the **[edit mpls]** hierarchy level.

Include one **interface** statement for each connection to the LSP.

```
[edit]
mpls {
  interface interface-name;
  interface interface-name;
}
```

5. Enable MPLS on the interface between the PE and CE routers by including the **interface** statement at the **[edit mpls]** hierarchy level.

Doing this allows the PE router to assign an MPLS label to traffic entering the LSP or to remove the label from traffic exiting the LSP.

```
[edit]
mpls {
  interface interface-name;
}
```

For information about configuring MPLS, see the *Minimum MPLS Configuration*.

**Related Documentation**

- [Minimum MPLS Configuration](#)

---

## Configuring EVPN Routing Instances

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 **routing-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 each interface using the **interface** statement at the **[edit routing-instances routing-instance-name protocols evpn]** hierarchy level.
    - b. (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.

- c. (Optional) Specify a static MAC address for a logical interface in a bridge domain using the [static-mac](#) 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. Specify the MPLS label allocation setting for the EVPN by including the [label-allocation](#) statement with the [per-instance](#) option at the **[edit routing-instances routing-instance-name protocols evpn]** hierarchy level.

If you configure this statement, one MPLS label is allocated for the specified EVPN routing instance.

10. Enable MAC accounting for the EVPN by including the [mac-statistics](#) statement at the **[edit routing-instances routing-instance-name protocols evpn]** hierarchy level.
11. 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.

12. 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>;
  flag flag <flag-modifier>;
}
```

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.

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 on page 14](#)
- [traceoptions on page 38](#)



## CHAPTER 4

# EVPN Configuration Statements

## evpn

```
Syntax  evpn {
        interface interface-name {
            ignore-encapsulation-mismatch;
            interface-mac-limit limit {
                packet-action drop;
            }
            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 logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols],  
[edit routing-instances *routing-instance-name* protocols]

**Release Information** Statement introduced in Junos OS Release 13.2 for EVPNs on MX 3D Series routers.

**Description** Enables an Ethernet VPN (EVPN) on the routing instance.

**Options** 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**

- [Configuring EVPN Routing Instances on page 14](#)
- [Tracing EVPN Traffic and Operations on page 17](#)

## family (Protocols BGP)

```
Syntax  family {
    (inet | inet6 | inet-vpn | inet6-vpn | iso-vpn) {
        (any | flow | labeled-unicast | multicast | unicast) {
            accepted-prefix-limit {
                maximum number;
                teardown <percentage-threshold> idle-timeout (forever | minutes);
            }
            add-path {
                send {
                    path-count number;
                    prefix-policy [ policy-names ];
                }
                receive;
            }
            algp [disable];
            loops number;
            prefix-limit {
                maximum number;
                teardown <percentage> <idle-timeout (forever | minutes)>;
            }
            protection;
            rib-group group-name;
            topology name {
                community {
                    target identifier;
                }
            }
        }
        flow {
            no-validate policy-name;
        }
        labeled-unicast {
            accepted-prefix-limit {
                maximum number;
                teardown <percentage> <idle-timeout (forever | minutes)>;
            }
            aggregate-label {
                community community-name;
            }
            explicit-null {
                connected-only;
            }
            prefix-limit {
                maximum number;
                teardown <percentage> <idle-timeout (forever | minutes)>;
            }
            resolve-vpn;
            rib (inet.3 | inet6.3);
            rib-group group-name;
            traffic-statistics {
                file filename <world-readable | no-world-readable>;
                interval seconds;
            }
        }
    }
}
```

```

    }
  }
  route-target {
    accepted-prefix-limit {
      maximum number;
      proxy-generate <route-target-policy route-target-policy-name>;
      teardown <percentage> <idle-timeout (forever | minutes)>;
    }
    advertise-default;
    external-paths number;
    prefix-limit {
      maximum number;
      teardown <percentage> <idle-timeout (forever | minutes)>;
    }
  }
}
(evpn | inet-mdt | inet-mvpn | inet6-mvpn | l2vpn) {
  signaling {
    accepted-prefix-limit {
      maximum number;
      teardown <percentage-threshold> idle-timeout (forever | minutes);
    }
    add-path {
      send {
        path-count number;
        prefix-policy [ policy-names ];
      }
      receive;
    }
    aigp [disable];
    damping;
    loops number;
    prefix-limit {
      maximum number;
      teardown <percentage> <idle-timeout (forever | minutes)>;
    }
    rib-group group-name;
  }
}
}

```

**Hierarchy Level** [edit logical-systems *logical-system-name* protocols bgp],  
 [edit logical-systems *logical-system-name* protocols bgp group *group-name*],  
 [edit logical-systems *logical-system-name* protocols bgp group *group-name* neighbor *address*],  
 [edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols bgp],  
 [edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols bgp group *group-name*],  
 [edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols bgp group *group-name* neighbor *address*],  
 [edit protocols bgp],  
 [edit protocols bgp group *group-name*],  
 [edit protocols bgp group *group-name* neighbor *address*],  
 [edit routing-instances *routing-instance-name* protocols bgp],  
 [edit routing-instances *routing-instance-name* protocols bgp group *group-name*],  
 [edit routing-instances *routing-instance-name* protocols bgp group *group-name* neighbor *address*]

<b>Release Information</b>	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p> <p><b>inet-mvpn</b> and <b>inet6-mpvn</b> statements introduced in Junos OS Release 8.4.</p> <p><b>inet-mdt</b> statement introduced in Junos OS Release 9.4.</p> <p>Support for the <b>loops</b> statement introduced in Junos OS Release 9.6.</p> <p><b>evpn</b> statement introduced in Junos OS Release 13.2.</p>
<b>Description</b>	<p>Enable multiprotocol BGP (MP-BGP) by configuring BGP to carry network layer reachability information (NLRI) for address families other than unicast IPv4, to specify MP-BGP to carry NLRI for the IPv6 address family, or to carry NLRI for VPNs.</p>
<b>Options</b>	<p><b>any</b>—Configure the family type to be both unicast and multicast.</p> <p><b>evpn</b>—Configure NLRI parameters for Ethernet VPNs (EVPNs).</p> <p><b>inet</b>—Configure NLRI parameters for IPv4.</p> <p><b>inet6</b>—Configure NLRI parameters for IPv6.</p> <p><b>inet-mdt</b>—Configure NLRI parameters for the multicast distribution tree (MDT) subaddress family identifier (SAFI) for IPv4 traffic in Layer 3 VPNs.</p> <p><b>inet-mvpn</b>—Configure NLRI parameters for IPv4 for multicast VPNs.</p> <p><b>inet6-mvpn</b>—Configure NLRI parameters for IPv6 for multicast VPNs.</p> <p><b>inet-vpn</b>—Configure NLRI parameters for IPv4 for Layer 3 VPNs.</p> <p><b>inet6-vpn</b>—Configure NLRI parameters for IPv6 for Layer 3 VPNs.</p> <p><b>iso-vpn</b>—Configure NLRI parameters for IS-IS for Layer 3 VPNs.</p> <p><b>l2vpn</b>—Configure NLRI parameters for IPv4 for MPLS-based Layer 2 VPNs and VPLS.</p> <p><b>labeled-unicast</b>—Configure the family type to be labeled-unicast. This means that the BGP peers are being used only to carry the unicast routes that are being used by labeled-unicast for resolving the labeled-unicast routes. This statement is supported only with <b>inet</b> and <b>inet6</b>.</p> <p><b>multicast</b>—Configure the family type to be multicast. This means that the BGP peers are being used only to carry the unicast routes that are being used by multicast for resolving the multicast routes.</p> <p><b>unicast</b>—Configure the family type to be unicast. This means that the BGP peers only carry the unicast routes that are being used for unicast forwarding purposes. The default family type is <b>unicast</b>.</p> <p>The remaining statements are explained separately.</p>
<b>Required Privilege Level</b>	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>

- Related Documentation**
- [Configuring IBGP Sessions Between PE Routers in VPNs on page 9](#)
  - *Understanding Multiprotocol BGP*
  - *autonomous-system*
  - *local-as*

---

## ignore-encapsulation-mismatch

---

<b>Syntax</b>	ignore-encapsulation-mismatch;
<b>Hierarchy Level</b>	<pre>[edit logical-systems <i>logical-system-name</i> protocols l2circuit local-switching interface <i>interface-name</i>], [edit logical-systems <i>logical-system-name</i> protocols l2circuit neighbor <i>address</i> interface <i>interface-name</i>], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols vpls mesh-group <i>mesh-group-name</i> neighbor <i>neighbor-id</i>], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols vpls neighbor <i>neighbor-id</i>], [edit protocols l2circuit local-switching interface <i>interface-name</i>], [edit protocols l2circuit neighbor <i>address</i> interface <i>interface-name</i>], [edit routing-instances <i>routing-instance-name</i> protocols evpn interface <i>interface-name</i>], [edit routing-instances <i>routing-instance-name</i> protocols vpls mesh-group <i>mesh-group-name</i> neighbor <i>neighbor-id</i>], [edit routing-instances <i>routing-instance-name</i> protocols vpls neighbor <i>neighbor-id</i>]</pre>
<b>Release Information</b>	Statement introduced in Junos OS Release 9.2. Statement extended to support local switching in Junos OS Release 10.4. Statement introduced for EVPNs in Junos OS Release 13.2 for MX 3D Series.
<b>Description</b>	Allow a Layer 2 circuit, VPLS, or EVPN to be established even though the encapsulation configured on the CE device interface does not match the encapsulation configured on the Layer 2 circuit, VPLS, or EVPN interface.
<b>Required Privilege Level</b>	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Configuring EVPN Routing Instances on page 14</a></li><li>• <i>Configuring Interfaces for Layer 2 Circuits</i></li></ul>

## instance-type

<b>Syntax</b>	<code>instance-type type;</code>
<b>Hierarchy Level</b>	[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> ], [edit routing-instances <i>routing-instance-name</i> ]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4. <b>virtual-switch</b> and <b>layer2-control</b> 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. <b>evpn</b> option introduced in Junos OS Release 13.2 for MX 3D Series routers.
<b>Description</b>	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 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.
- **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.

- **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 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.

<b>Required Privilege Level</b>	routing—To view this statement in the configuration.
	routing-control—To add this statement to the configuration.

- |                              |  |
|------------------------------|--|
| <b>Related Documentation</b> | • <i>Example: Using Virtual Routing Instances to Route Among VLANs on EX Series Switches</i> |
|                              | • <i>Configuring Routing Instances on PE Routers in VPNs</i>                                 |
|                              | • <a href="#">Configuring EVPN Routing Instances on page 14</a>                              |
|                              | • <i>Configuring Virtual Routing Instances (CLI Procedure)</i>                               |
|                              | • <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>Layer 2 Routing Instance Types</i>  |



## interface (EVPN Routing Instances)

<b>Syntax</b>	<pre>interface <i>interface-name</i> {   ignore-encapsulation-mismatch;   interface-mac-limit <i>limit</i> {     packet-action drop;   }   no-mac-learning;   static-mac <i>mac-address</i>; }</pre>
<b>Hierarchy Level</b>	[edit routing-instances <i>routing-instance-name</i> protocols evpn]
<b>Release Information</b>	Statement introduced in Junos OS Release 13.2 for EVPNs on MX 3D Series routers.
<b>Description</b>	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.
<b>Options</b>	<p><i>interface-name</i>—Name of the interface.</p> <p>The remaining statements are explained separately.</p>
<b>Required Privilege Level</b>	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Configuring EVPN Routing Instances on page 14</a></li> <li>• <a href="#">evpn on page 20</a></li> <li>• <a href="#">instance-type on page 25</a></li> </ul>

## interface (Routing Instances)

---

<b>Syntax</b>	<code>interface <i>interface-name</i> {     description <i>text</i>; }</code>
<b>Hierarchy Level</b>	[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> ], [edit routing-instances <i>routing-instance-name</i> ]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.2 for EX Series switches. Statement introduced in Junos OS Release 12.3 for ACX Series routers. Statement introduced in Junos OS Release 13.2 for MX 3D Series routers.
<b>Description</b>	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 <b>vrf</b> is specified for the <b>instance-type</b> statement included in the routing instance configuration, this statement is required.
<b>Options</b>	<i>interface-name</i> —Name of the interface.
<b>Required Privilege Level</b>	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Configuring Routing Instances on PE Routers in VPNs</a></li><li>• <a href="#">Configuring EVPN Routing Instances on page 14</a></li><li>• <a href="#">Example: Configuring MPLS-Based Layer 3 VPNs on EX Series Switches</a></li><li>• <a href="#">interface (VPLS Routing Instances)</a></li></ul>

## interface-mac-limit

<b>Syntax</b>	<pre>interface-mac-limit <i>limit</i> {     <b>packet-action</b> drop; }</pre>
<b>Hierarchy Level</b>	<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>
<b>Release Information</b>	<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>
<b>Description</b>	<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>
<b>Options</b>	<p><b>limit</b>—Specify the number of MAC addresses that can be learned from each interface.</p> <p><b>Range:</b> 16 through 65,536 MAC addresses</p> <p><b>Default:</b> 512 addresses</p> <p>The remaining statement is explained separately.</p>
<b>Required Privilege Level</b>	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Configuring EVPN Routing Instances on page 14</a></li> <li>• <a href="#">Configuring VPLS Routing Instances</a></li> <li>• <a href="#">interface on page 27</a></li> <li>• <a href="#">mac-table-size on page 32</a></li> </ul>

## label-allocation


---

<b>Syntax</b>	label-allocation per-instance;
<b>Hierarchy Level</b>	[edit routing-instances <i>routing-instance-name</i> protocols evpn]
<b>Release Information</b>	Statement introduced in Junos OS Release 13.2 on MX 3D Series routers.
<b>Description</b>	Specifies the MPLS label allocation setting for the EVPN routing instance.
<b>Options</b>	<b>per-instance</b> —Allocates a single MPLS label for the EVPN routing instance.
<b>Required Privilege Level</b>	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Configuring EVPN Routing Instances on page 14</a></li></ul>

## mac-statistics

<b>Syntax</b>	mac-statistics;
<b>Hierarchy Level</b>	<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 <a href="#">evpn</a>],</p> <p>[edit switch-options],</p> <p>[edit switch-options],</p> <p>[edit vlans <i>vlan-name</i> switch-options]</p>
<b>Release Information</b>	<p>Statement introduced in Junos OS Release 8.4.</p> <p>Support for the <b>switch-options</b> statement added in Junos OS Release 9.2.</p> <p>Support for top-level configuration for the <b>virtual-switch</b> 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>
<b>Description</b>	(MX Series routers and EX Series switches 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.
<b>Default</b>	disabled
<b>Required Privilege Level</b>	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <i>Layer 2 Learning and Forwarding for Bridge Domains Overview</i></li> <li>• <i>Layer 2 Learning and Forwarding for VLANs Overview</i></li> <li>• <i>Layer 2 Learning and Forwarding for Bridge Domains Functioning as Switches with Layer 2 Trunk Ports</i></li> <li>• <i>Layer 2 Learning and Forwarding for VLANs Acting as a Switch for a Layer 2 Trunk Port</i></li> <li>• <a href="#">Configuring EVPN Routing Instances on page 14</a></li> </ul>

## mac-table-size

<b>Syntax</b>	<code>mac-table-size size {     <code>packet-action</code> drop; }</code>
<b>Hierarchy Level</b>	[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]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 13.2 for EVPNs on MX 3D Series routers.
<b>Description</b>	Specify the size of the MAC address table.
<b>Options</b>	<p><b>size</b>—Specify the size of the MAC address table.</p> <p><b>Range:</b></p> <ul style="list-style-type: none"> <li>• (M Series and T Series routers only) 16 through 65,536 MAC addresses</li> <li>• (MX Series routers only) 16 through 1,048,575 MAC addresses</li> <li>• (T4000 routers with Type 5 FPCs only) 16 through 262,143 MAC addresses</li> </ul>
	<div>  <p><b>NOTE:</b> 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> </div>
	<p><b>Default:</b> 512 MAC addresses</p> <p>The remaining statement is explained separately.</p>
<b>Required Privilege Level</b>	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Configuring EVPN Routing Instances on page 14</a></li> <li>• <a href="#">Configuring VPLS Routing Instances</a></li> <li>• <a href="#">Configuring Improved VPLS MAC Address Learning on T4000 Routers with Type 5 FPCs</a></li> <li>• <a href="#">enhanced-mode</a></li> <li>• <a href="#">evpn on page 20</a></li> </ul>

## no-mac-learning

<b>Syntax</b>	no-mac-learning;
<b>Hierarchy Level</b>	<p>[edit bridge-domains <i>bridge-domain-name</i> bridge-options],</p> <p>[edit bridge-domains <i>bridge-domain-name</i> bridge-options interface <i>interface-name</i>],</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> bridge-domains <i>bridge-domain-name</i> bridge-options interface <i>interface-name</i>],</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> bridge-domains <i>bridge-domain-name</i> bridge-options interface <i>interface-name</i>],</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 bridge-domains <i>bridge-domain-name</i> bridge-options interface <i>interface-name</i>],</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> bridge-domains <i>bridge-domain-name</i> bridge-options interface <i>interface-name</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols <a href="#">evpn</a>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols evpn <a href="#">interface</a> <i>interface-name</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> switch-options],</p> <p>[edit switch-options],</p> <p>[edit switch-options interface <i>interface-name</i>],</p> <p>[edit vlans <i>vlan-name</i> switch-options],</p> <p>[edit vlans <i>vlan-name</i> switch-options interface <i>interface-name</i>]</p>
<b>Release Information</b>	<p>Statement introduced in Junos OS Release 8.4.</p> <p>Support for the <b>switch-options</b> statement added in Junos OS Release 9.2.</p> <p>Support for top-level configuration for the <b>virtual-switch</b> 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.</p> <p>Support for logical systems added in Junos OS Release 9.6.</p> <p>[edit switch-options], [edit switch-options interface <i>interface-name</i>], [edit vlans <i>vlan-name</i> switch-options], and [edit vlans <i>vlan-name</i> switch-options interface <i>interface-name</i>] hierarchy levels introduced in Junos OS Release 12.3 R2 for EX Series switches.</p> <p>Support for EVPNs added in Junos OS Release 13.2 for MX 3D Series routers.</p>
<b>Description</b>	(MX Series routers and EX Series switches only) Disable MAC learning for a virtual switch, for a bridge domain or VLAN, for a specific logical interface in a bridge domain or VLAN, for a set of bridge domains or VLANs associated with a Layer 2 trunk port, or for 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.

<b>Default</b>	MAC learning is enabled. Use <b>no-mac-learning</b> to disable MAC learning.
<b>Required Privilege Level</b>	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Configuring EVPN Routing Instances on page 14</a></li><li>• <i>Layer 2 Learning and Forwarding for Bridge Domains Overview</i></li><li>• <i>Layer 2 Learning and Forwarding for VLANs Overview</i></li><li>• <i>Layer 2 Learning and Forwarding for Bridge Domains Functioning as Switches with Layer 2 Trunk Ports</i></li><li>• <i>Layer 2 Learning and Forwarding for VLANs Acting as a Switch for a Layer 2 Trunk Port</i></li></ul>



## packet-action

**Syntax** packet-action drop;

**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](#)],  
 [edit routing-instances *routing-instance-name* protocols evpn interface *interface-name* [interface-mac-limit](#)],  
 [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 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*]

**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*],

	<p>[edit vlans <i>vlan-name</i> switch-options interface-mac-limit <i>limit</i>], and [edit vlans <i>vlan-name</i> switch-options mac-table-size <i>limit</i>] hierarchy levels introduced in Junos OS Release 12.3R2 for EX Series switches.</p> <p>Support for EVPNs introduced in Junos OS Release 13.2 on MX 3D Series routers.</p>
<b>Description</b>	Specify that packets for new source MAC addresses be dropped after the MAC address limit is reached. If this statement is not configured, packets for new source MAC addresses are forwarded by default.
<b>Default</b>	Disabled. The default is for packets for new source MAC addresses to be forwarded after the MAC address limit is reached.
<b>Required Privilege Level</b>	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Configuring EVPN Routing Instances on page 14</a></li><li>• <i>Layer 2 Learning and Forwarding for Bridge Domains Overview</i></li><li>• <i>Layer 2 Learning and Forwarding for VLANs Overview</i></li><li>• <i>Layer 2 Learning and Forwarding for Bridge Domains Functioning as Switches with Layer 2 Trunk Ports</i></li><li>• <i>Layer 2 Learning and Forwarding for VLANs Overview</i></li><li>• <i>Layer 2 Learning and Forwarding for VLANs Acting as a Switch for a Layer 2 Trunk Port</i></li></ul>

---

## routing-instances

---

<b>Syntax</b>	routing-instances <i>routing-instance-name</i> { ... }
<b>Hierarchy Level</b>	[edit], [edit logical-systems <i>logical-system-name</i> ]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4.
<b>Description</b>	Configure an additional routing entity for a router. You can create multiple instances of BGP, IS-IS, OSPF, OSPF version 3 (OSPFv3), and RIP for a router.
<b>Default</b>	Routing instances are disabled for the router.
<b>Options</b>	<i>routing-instance-name</i> —Name of the routing instance, a maximum of 31 characters. The remaining statements are explained separately.
<b>Required Privilege Level</b>	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Configuring EVPN Routing Instances on page 14</a></li><li>• <i>Configuring Routing Instances on PE Routers in VPNs</i></li></ul>

## static-mac

<b>Syntax</b>	static-mac <i>mac-address</i> { vlan-id <i>number</i> ; }
<b>Hierarchy Level</b>	[edit bridge-domains <i>bridge-domain-name</i> bridge-options interface <i>interface-name</i> ], [edit logical-systems <i>logical-system-name</i> bridge-domains <i>bridge-domain-name</i> bridge-options interface <i>interface-name</i> ], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> bridge-domains <i>bridge-domain-name</i> bridge-options interface <i>interface-name</i> ], [edit routing-instances <i>routing-instance-name</i> bridge-domains <i>bridge-domain-name</i> bridge-options interface <i>interface-name</i> ], [edit routing-instances <i>routing-instance-name</i> protocols evpn interface <i>interface-name</i> ] [edit vlans <i>vlan-name</i> switch-options interface <i>interface-name</i> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 8.4. Support for logical systems added in Junos OS Release 9.6. [edit vlans <i>vlan-name</i> switch-options interface <i>interface name</i> ] hierarchy level introduced in Junos OS Release 12.3R2 for EX Series switches. Support for EVPNs added in Junos OS Release 13.2 for MX 3D Series routers. The <b>vlan-id</b> option is not available for EVPNs.
<b>Description</b>	Configure a static MAC address for a logical interface in a bridge domain or VLAN.  The <b>vlan-id</b> option can be specified for <b>static-macs</b> only if <b>vlan-id all</b> is configured for the bridging domain or VLAN.
<b>Options</b>	<b>mac-address</b> —MAC address  <b>vlan-id number</b> —(Optional) VLAN identifier to associate with static MAC address.
<b>Required Privilege Level</b>	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Configuring EVPN Routing Instances on page 14</a></li> <li>• <a href="#">Layer 2 Learning and Forwarding for Bridge Domains Overview</a></li> <li>• <a href="#">Layer 2 Learning and Forwarding for VLANs Overview</a></li> </ul>

## traceoptions (Protocols EVPN)

---

Syntax	<pre>traceoptions {     file <i>filename</i> &lt;files <i>number</i>&gt; &lt;size <i>size</i>&gt; &lt;world-readable   no-world-readable&gt;;     flag <i>flag</i> &lt;<i>flag-modifier</i>&gt;; }</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.
Description	Trace traffic flowing through an EVPN routing instance.
Options	<p><b>file <i>filename</i></b>—Name of the file to receive the output of the tracing operation. Enclose the name in quotation marks (" ").</p> <p><b>files <i>number</i></b>—(Optional) Maximum number of trace files. When a trace file named <b><i>trace-file</i></b> reaches the maximum size as specified by the <b>size</b> option, it is renamed <b><i>trace-file.0</i></b>. When <b><i>trace-file</i></b> again reaches the maximum size, <b><i>trace-file.0</i></b> is renamed <b><i>trace-file.1</i></b> and <b><i>trace-file</i></b> is renamed <b><i>trace-file.0</i></b>. 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 <b>size</b> option.</p> <p><b>Range:</b> 2 through 1000 files</p> <p><b>Default:</b> 2 files</p> <p><b>flag <i>flag</i></b>—Tracing operation to perform. To specify more than one tracing operation, include multiple <b>flag</b> statements. You can specify the following tracing flags:</p> <ul style="list-style-type: none"><li>• <b>all</b>—All EVPN tracing options</li><li>• <b>error</b>—Error conditions</li><li>• <b>general</b>—General events</li><li>• <b>mac-database</b>—MAC route database in the EVPN routing instance</li><li>• <b>nlri</b>—EVPN advertisements received or sent by means of BGP</li><li>• <b>normal</b>—Normal events</li><li>• <b>oam</b>—OAM messages</li><li>• <b>policy</b>—Policy processing</li><li>• <b>route</b>—Routing information</li><li>• <b>state</b>—State transitions</li><li>• <b>task</b>—Routing protocol task processing</li><li>• <b>timer</b>—Routing protocol timer processing</li></ul>

- **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.

<b>Required Privilege Level</b>	routing—To view this statement in the configuration.
	routing-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Tracing EVPN Traffic and Operations on page 17</a></li></ul>

## vlan-id (routing instance)

---

<b>Syntax</b>	vlan-id ( <i>vlan-id</i>   all   none);
<b>Hierarchy Level</b>	[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> ], [edit routing-instances <i>routing-instance-name</i> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 13.2.
<b>Description</b>	Specify 802.1Q VLAN tag IDs to a routing instance.
<b>Options</b>	<p><b>vlan-id</b>—A valid VLAN identifier.</p> <p><b>Range:</b> 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><b>all</b>—Include all VLAN identifiers specified on the logical interfaces included in the routing instance.</p> <p><b>none</b>—Include no VLAN identifiers for the routing instance.</p>
<b>Required Privilege Level</b>	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Configuring EVPN Routing Instances on page 14</a></li></ul>

## PART 3

# Index

- [Index on page 43](#)





# Index

## Symbols

#, comments in configuration statements.....	xii
( ), in syntax descriptions.....	xii
< >, in syntax descriptions.....	xii
[ ], in configuration statements.....	xii
{ }, in configuration statements.....	xii
(pipe), in syntax descriptions.....	xii

## A

autonomous system number	
route distinguisher.....	15

## B

BGP	
MP-BGP.....	21
braces, in configuration statements.....	xii
brackets	
angle, in syntax descriptions.....	xii
square, in configuration statements.....	xii

## C

comments, in configuration statements.....	xii
conventions	
text and syntax.....	xi
curly braces, in configuration statements.....	xii
customer support.....	xiii
contacting JTAC.....	xiii

## D

documentation	
comments on.....	xiii

## E

evpn statement.....	20
usage guidelines.....	15
EVPNs	
routing instance, configuring.....	14

## F

family evpn.....	10
family inet-vpn.....	10

family inet6-vpn.....	10
family l2vpn.....	10
family statement	
BGP.....	21
flow statement.....	21
font conventions.....	xi

## I

ignore-encapsulation-mismatch statement.....	24
usage guidelines.....	15
instance-type statement.....	25
usage guidelines.....	14
interface statement	
EVPNs.....	27
usage guidelines.....	15
IS-IS.....	28
OSPF.....	28
routing instances	
usage guidelines.....	15
VPNs.....	28
interface-mac-limit statement.....	29
usage guidelines.....	16

## L

label-allocation statement.....	30
usage guidelines.....	16

## M

mac-statistics statement.....	31
usage guidelines.....	16
mac-table-size statement.....	32
usage guidelines.....	16
manuals	
comments on.....	xiii
MP-BGP.....	21
multiprotocol BGP (MP-BGP).....	21

## N

no-mac-learning statement.....	33
usage guidelines.....	16

## O

OSPF	
enabling.....	28

## P

packet-action statement.....	35
usage guidelines.....	16
parentheses, in syntax descriptions.....	xii

**R**

route distinguisher	
autonomous system number.....	15
routing-instances statement.....	36
usage guidelines.....	14

**S**

static-mac statement.....	37
usage guidelines.....	16
support, technical See technical support	
syntax conventions.....	xi

**T**

technical support	
contacting JTAC.....	xiii
traceoptions statement	
EVPNs.....	38
usage guidelines.....	17

**V**

vlan-id statement.....	40
usage guidelines.....	15