



Junos[®] OS

Layer 2 Circuits Configuration Guide

Release
13.1



Published: 2013-02-28

Juniper Networks, Inc.
1194 North Mathilda Avenue
Sunnyvale, California 94089
USA
408-745-2000
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 Layer 2 Circuits Configuration Guide

13.1

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	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	Overview	
Chapter 1	Introduction to Layer 2 Circuits	3
	Layer 2 Circuit Overview	3
	Layer 2 Circuit Bandwidth Accounting and Call Admission Control	4
	Bandwidth Accounting and Call Admission Control Overview	4
	Selecting an LSP Based on the Bandwidth Constraint	5
	LSP Path Protection and CAC	5
	Secondary Paths and CAC	6
	Fast Reroute and CAC	6
	Link and Node Protection and CAC	6
	Layer 2 Circuits Trunk Mode	7
	Egress Protection LSPs for Layer 2 Circuits	7
Chapter 2	Introduction to Configuring Layer 2 Circuits	9
	Configuring LDP for Layer 2 Circuits	9
Part 2	Configuration	
Chapter 3	Configuring Layer 2 Circuits	13
	Configuring an IGP on the PE and P Routers	13
	Configuring IBGP Sessions Between PE Routers in VPNs	14
	Configuring a Signaling Protocol and LSPs for VPNs	15
	Using LDP for VPN Signaling	15
	Using RSVP for VPN Signaling	16
	Configuring Routing Instances on PE Routers in VPNs	18
	Configuring the Routing Instance Name for a VPN	19
	Configuring the Description	19
	Configuring the Instance Type	20

Configuring Interfaces for VPN Routing	20
General Configuration for VPN Routing	20
Configuring Interfaces for Layer 3 VPNs	21
Configuring Interfaces for Carrier-of-Carriers VPNs	21
Configuring Unicast RPF on VPN Interfaces	22
Configuring the Route Distinguisher	22
Configuring Automatic Route Distinguishers	23
Configuring Policies for the VRF Table on PE Routers in VPNs	23
Configuring the Route Target	24
Configuring the Route Origin	24
Configuring an Import Policy for the PE Router's VRF Table	25
Configuring an Export Policy for the PE Router's VRF Table	27
Applying Both the VRF Export and the BGP Export Policies	28
Configuring a VRF Target	29
Configuring Local Interface Switching in Layer 2 Circuits	30
Configuring the Interfaces for the Local Interface Switch	30
Enabling Local Interface Switching When the MTU Does Not Match	31
Configuring Interfaces for Layer 2 Circuits	32
Configuring the Address for the Neighbor of the Layer 2 Circuit	32
Configuring the Neighbor Interface for the Layer 2 Circuit	32
Configuring a Community for the Layer 2 Circuit	33
Configuring the Control Word for Layer 2 Circuits	33
Configuring the Encapsulation Type for the Layer 2 Circuit Neighbor Interface	35
Enabling the Layer 2 Circuit When the Encapsulation Does Not Match	35
Configuring the MTU for the Layer 2 Circuit Neighbor Interface	36
Configuring the Protect Interface	37
Configuring the Protect Interface From Switching Over to the Primary Interface	37
Configuring the Pseudowire Status TLV	38
Configuring Layer 2 Circuits over Both RSVP and LDP LSPs	38
Configuring the Virtual Circuit ID	39
Configuring the Interface Encapsulation Type for Layer 2 Circuits	39
Configuring ATM2 IQ Interfaces for Layer 2 Circuits	40
Configuring Static Layer 2 Circuits	40
Configuring Policies for Layer 2 Circuits	41
Configuring the Layer 2 Circuit Community	41
Configuring the Policy Statement for the Layer 2 Circuit Community	42
Example: Configuring a Policy for a Layer 2 Circuit Community	43
Verifying the Layer 2 Circuit Policy Configuration	44
Enabling BGP Path Selection for Layer 2 VPNs and VPLS	45
Configuring ATM Trunking on Layer 2 Circuits	47
Configuring Bandwidth Allocation and Call Admission Control in Layer 2 Circuits	48
Reducing APS Switchover Time in Layer 2 Circuits	49
Configuring Per-Packet Load Balancing	50
Configuring Fast APS Switchover	50

Chapter 4	Layer 2 Circuits Examples	53
	Example: Configuring Layer 2 Circuit Protect Interfaces	53
	Configuring Router PE1	54
	Configuring Router PE2	56
	Configuring Router CE1	57
	Configuring Router CE2	58
	Example: Configuring Layer 2 Circuit Switching Protection	58
	Example: Configuring an Egress Protection LSP for a Layer 2 Circuit	70
	Using the Layer 2 Interworking Interface to Interconnect a Layer 2 Circuit to a Layer 2 VPN	80
	Example: Interconnecting a Layer 2 Circuit with a Layer 2 VPN	81
	Applications for Interconnecting a Layer 2 Circuit with a Layer 2 Circuit	90
	Example: Interconnecting a Layer 2 Circuit with a Layer 2 Circuit	90
	Applications for Interconnecting a Layer 2 Circuit with a Layer 3 VPN	106
	Example: Interconnecting a Layer 2 Circuit with a Layer 3 VPN	106
Part 3	Administration	
Chapter 5	Layer 2 Circuit Reference	129
	Supported Layer 2 Circuit Standards	129
Chapter 6	Summary of Layer 2 Circuit Configuration Statements	131
	bandwidth (Protocols Layer 2 Circuit)	131
	backup-interface (Layer 2 Circuits)	132
	backup-neighbor	133
	community (Protocols Layer 2 Circuit)	134
	connection-protection	135
	control-word (Protocols Layer 2 Circuit Neighbor)	135
	description (Protocols Layer 2 Circuit Neighbor)	136
	egress-protection (Layer 2 circuit)	136
	egress-protection (MPLS)	137
	encapsulation (Physical Interface)	138
	encapsulation-type (Layer 2 Circuits)	143
	end-interface	144
	fast-aps-switch	145
	ignore-encapsulation-mismatch	146
	ignore-mtu-mismatch	146
	install-nexthop	147
	interface (Protocols Layer 2 Circuit)	148
	l2circuit	150
	l2vpn-use-bgp-rules	151
	local-switching (Layer 2 Circuits)	152
	mtu	153
	neighbor (Protocols Layer 2 Circuit)	155
	no-revert (Local Switching)	156
	no-revert (Neighbor Interface)	157
	ping-interval	158
	protect-interface	159
	protected-l2circuit	160

	protector-interface	161
	protector-pe	161
	pseudowire-status-tlv	162
	psn-tunnel-endpoint	163
	standby (Protocols Layer 2 Circuit)	164
	static (Protocols Layer 2 Circuit)	165
	traceoptions (Protocols Layer 2 Circuit)	166
	virtual-circuit-id	168
Part 4	Troubleshooting	
Chapter 7	Troubleshooting Layer 2 Circuits	171
	Tracing Layer 2 Circuit Operations	171
Part 5	Index	
	Index	175

List of Figures

Part 1	Overview	
Chapter 1	Introduction to Layer 2 Circuits	3
	Figure 1: Components of a Layer 2 Circuit	3
	Figure 2: Egress protection LSP	7
Part 2	Configuration	
Chapter 3	Configuring Layer 2 Circuits	13
	Figure 3: ATM Trunking on Layer 2 Circuits	47
Chapter 4	Layer 2 Circuits Examples	53
	Figure 4: Layer 2 Circuits Using Protect Interfaces	54
	Figure 5: Connection protection enabled between router PE1 and router PE2	59
	Figure 6: Connection Protection Using a Pseudowire Configured through Router PE3 as the Protection Path	60
	Figure 7: Connection Protection Using a Pseudowire Configured through Router PE3 as the Working Path	60
	Figure 8: Egress Protection LSP Configured from Router PE2 to Router PE3	71
	Figure 9: Physical Topology of a Layer 2 Circuit to a Layer 2 VPN Connection	82
	Figure 10: Logical Topology of a Layer 2 Circuit to a Layer 2 VPN Connection	83
	Figure 11: Physical Topology of a Layer 2 Circuit Terminating into a Layer 2 Circuit	91
	Figure 12: Logical Topology of a Layer 2 Circuit Terminating into a Layer 2 Circuit	92
	Figure 13: Physical Topology of a Layer 2 Circuit to Layer 3 VPN Interconnection	107
	Figure 14: Logical Topology of a Layer 2 Circuit to Layer 3 VPN Interconnection	108

List of Tables

About the Documentation	xi
Table 1: Notice Icons	xiii
Table 2: Text and Syntax Conventions	xiii

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:

- MX Series
- T Series
- M Series
- ACX Series
- PTX 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.xsl;
    }
  }
}
interfaces {
  fxp0 {
    disable;
    unit 0 {
      family inet {
        address 10.0.0.1/24;
      }
    }
  }
}
```

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

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

Merging a Snippet

To merge a snippet, follow these steps:

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

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

```
commit {
  file ex-script-snippet.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 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.

Table 2 on page xiii 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

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

Convention	Description	Examples
<i>Italic text like this</i>	<ul style="list-style-type: none"> Introduces or emphasizes important new terms. Identifies book 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 System Basics Configuration 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)	Enclose optional keywords or variables.	stub <default-metric metric>;
(pipe symbol)	Indicates a choice between the mutually exclusive keywords or variables on either side of the symbol. The set of choices is often enclosed in parentheses for clarity.	broadcast multicast <i>(string1 string2 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)	Enclose a variable for which you can substitute one or more values.	community name members [community-ids]
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.	
J-Web GUI Conventions		
Bold text like this	Represents J-Web 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 J-Web 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 send your comments to 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 Layer 2 Circuits on page 3](#)
- [Introduction to Configuring Layer 2 Circuits on page 9](#)

CHAPTER 1

Introduction to Layer 2 Circuits

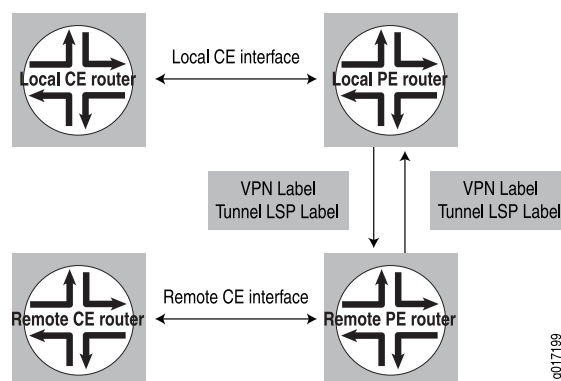
- [Layer 2 Circuit Overview on page 3](#)
- [Layer 2 Circuit Bandwidth Accounting and Call Admission Control on page 4](#)
- [Egress Protection LSPs for Layer 2 Circuits on page 7](#)

Layer 2 Circuit Overview

A Layer 2 circuit is a point-to-point Layer 2 connection transported using Multiprotocol Label Switching (MPLS) or other tunneling technology on the service provider's network. A Layer 2 circuit is similar to a circuit cross-connect (CCC), except that multiple virtual circuits (VCs) are transported over a single shared label-switched path (LSP) tunnel between two provider edge (PE) routers. In contrast, each CCC requires a separate dedicated LSP.

The Junos OS implementation of Layer 2 circuits supports only the remote form of a Layer 2 circuit; that is, a connection from a local customer edge (CE) router to a remote CE router. [Figure 1 on page 3](#) illustrates the components of a Layer 2 circuit.

Figure 1: Components of a Layer 2 Circuit



To establish a Layer 2 circuit, the Link Integrity Protocol (LIP) is used as the signaling protocol to advertise the ingress label to the remote PE routers. For this purpose, a targeted remote LDP neighbor session is established using the extended discovery mechanism described in LDP, and the session is brought up to the remote PE loopback IP address. Because LDP looks at the Layer 2 circuit configuration and initiates extended neighbor discovery for all the Layer 2 circuit neighbors (the remote PEs), no new

configuration is necessary in LDP. Each Layer 2 circuit is represented by the logical interface connecting the local PE router to the local customer edge (CE) router. Note that LDP must be enabled on the lo0.0 interface for extended neighbor discovery to function correctly.

Packets are sent to remote CE routers over an egress VPN label advertised by the remote PE router, using a targeted LDP session. The VPN label is sent over an LDP LSP to the remote PE router connected to the remote CE router. Return traffic from the remote CE router destined to the local CE router is sent using an ingress VPN label advertised by the local PE router, which is also sent over the LDP LSP to the local PE router from the remote PE router.

**Related
Documentation**

- [Layer 3 VPN Overview](#)
- [Layer 2 VPN Overview](#)
- [Layer 2 VPN Applications](#)
- [Applications for Interconnecting a Layer 2 Circuit with a Layer 2 Circuit on page 90](#)
- [Applications for Interconnecting a Layer 2 Circuit with a Layer 3 VPN on page 106](#)
- [Example: Interconnecting a Layer 2 Circuit with a Layer 2 Circuit on page 90](#)
- [Example: Interconnecting a Layer 2 Circuit with a Layer 3 VPN on page 106](#)
- [Example: Interconnecting a Layer 2 Circuit with a Layer 2 VPN on page 81](#)

Layer 2 Circuit Bandwidth Accounting and Call Admission Control

The sections that follow discuss Layer 2 circuit bandwidth accounting and call admission control (CAC):

- [Bandwidth Accounting and Call Admission Control Overview on page 4](#)
- [Selecting an LSP Based on the Bandwidth Constraint on page 5](#)
- [LSP Path Protection and CAC on page 5](#)
- [Layer 2 Circuits Trunk Mode on page 7](#)

Bandwidth Accounting and Call Admission Control Overview

Some network environments require that a certain level of service be guaranteed across the entire length of a path transiting a service provider's network. For Layer 2 circuits transiting an MPLS core network, a customer requirement might be to assure that guarantees for bandwidth and class of service (CoS) be maintained across the core network. For example, an Asynchronous Transfer Mode (ATM) circuit can provide service guarantees for each traffic class. A Layer 2 circuit configured to transport that ATM circuit across the network could be expected to provide the same service guarantees.

Providing this type of service guarantee requires the following:

- The LSPs in the MPLS core network must be able to provide service guarantees for bandwidth, rerouting, and route failures. You accomplish these guarantees by

configuring multiclass LSPs. For more information about multiclass LSPs, see the Junos OS MPLS Applications Configuration Guide.

- The service guarantee must be maintained across the entire length of the link as it transits the service provider's network. Different Layer 2 circuits could have different bandwidth requirements. However, many Layer 2 circuits could be transported over the same E-LSP in the MPLS core network.
- CAC ensures that the LSP has sufficient bandwidth to accommodate the Layer 2 circuit. If there is not enough bandwidth over a particular LSP, the Layer 2 circuit is prevented from using that LSP.

Selecting an LSP Based on the Bandwidth Constraint

CAC of Layer 2 circuits is based on the bandwidth constraint. You must configure this constraint for each Layer 2 circuit interface. If there is a bandwidth constraint configured for a Layer 2 circuit, CAC bases the final selection of which LSP-forwarding next hop to use on the following:

- If multiple LSPs meet the bandwidth requirements, the first LSP found that can satisfy the bandwidth requirements for the Layer 2 circuit is selected.
- If there is more than one next hop mapped to the same LSP, then all the next hops that map to that LSP and pass CAC constraints are installed. This allows the Layer 2 circuit routes to restore themselves quickly in case of failure.
- The available bandwidth on the selected LSP is decremented by the bandwidth required for each Layer 2 circuit. Similarly, when the Layer 2 circuit route is changed or deleted (for example, when the route is disassociated from that particular LSP), the bandwidth on the corresponding LSP is incremented.
- There are no priorities among different Layer 2 circuits competing for the same LSP next hop in the core network.
- When an LSP's bandwidth changes, the Layer 2 circuits using that LSP repeat the CAC process again.

If the LSP bandwidth increases, some Layer 2 circuits that were not established might now successfully resolve over the LSP. Similarly, if the bandwidth of the LSP decreases, some Layer 2 circuits that were previously up might now be declared down because of insufficient bandwidth on the LSP.

- When no LSP is found to meet the bandwidth requirements of the Layer 2 circuit, it is considered to be a CAC failure, and an error is reported.

LSP Path Protection and CAC

CAC can take into account LSPs that have been configured with an MPLS path protection feature, such as secondary paths, fast reroute, or node and link protection. CAC can consider the bandwidth available on these auxiliary links and can accept the backup connection as valid if the main connection fails. However, there are limitations on how the path protection feature must be configured to prevent CAC from taking down the Layer 2 circuit when the LSP it is using is switched to a backup route.

For more information about MPLS path protection features, see the Junos OS MPLS Applications Configuration Guide.

The sections that follow discuss the path protection features that can be used in conjunction with CAC and how they must be configured:

- [Secondary Paths and CAC on page 6](#)
- [Fast Reroute and CAC on page 6](#)
- [Link and Node Protection and CAC on page 6](#)

Secondary Paths and CAC

The following describes the ways in which secondary paths would interact with Layer 2 circuit CAC:

- If an LSP is configured with both primary and secondary paths, if the paths have the same bandwidth, and if this bandwidth is enough to accommodate the Layer 2 circuit, the Layer 2 circuit route installs both next hops in the forwarding table.

CAC allows the Layer 2 circuit to be switched to the secondary path if the primary path fails.

- If the LSP has primary and secondary paths configured with different bandwidths, each path must run through CAC independently. If the active path for that LSP passes CAC constraints successfully, then that next hop is installed and the corresponding LSP is selected to transport the Layer 2 circuit traffic. The LSP's secondary paths are then checked for CAC, and installed if there is sufficient bandwidth.

However, if the active path for the LSP fails to meet the CAC constraints, then that LSP is not selected and the system looks for a different LSP to transport the Layer 2 circuit.

For example, an LSP has an active primary path with 30 megabits of bandwidth and a secondary path with 10 megabits of bandwidth. The Layer 2 circuit requires 15 megabits of bandwidth. The secondary path fails CAC, and only the next hop corresponding to the primary path is installed for the Layer 2 circuit route. The path protection originally provided by the secondary path is no longer available.

Fast Reroute and CAC

No CAC is done for fast reroute detours. However, as long as the protected path satisfies the CAC bandwidth constraints, the detour next hop is also selected and installed.

Link and Node Protection and CAC

You can configure CAC on Layer 2 circuit-based LSPs with bandwidth constraints and also enable link and node protection. However, if the primary LSP fails, CAC might not be applied to the bypass LSP, meaning the bypass LSP might not meet the bandwidth constraint for the Layer 2 circuit. To minimize the risk of losing traffic, the Layer 2 circuit continues to use the non-CAC bypass LSP while an attempt is made to establish a new Layer 2 circuit route over an LSP that does support CAC.

Layer 2 Circuits Trunk Mode

Using Layer 2 circuit trunk mode, you can configure Layer 2 circuits to carry ATM trunks, providing a way to link ATM switches over an MPLS core network.

Layer 2 circuit trunk mode allows you to configure the following CoS features:


- CoS queues in Layer 2 circuit trunk mode—For ATM2 IQ interfaces, you can configure ATM CoS queues for Layer 2 circuit trunk mode.
- Layer 2 circuit trunk mode scheduling—For ATM2 IQ interfaces configured to use Layer 2 circuit trunk mode, you can share a scheduler among 32 trunks on an ATM port.
- Two early packet discard (EPD) thresholds per queue—For ATM2 IQ interfaces configured to use Layer 2 circuit trunk mode, you can set two EPD thresholds that depend on the packet-loss priorities (PLPs) of the packets.

For a detailed overview and configuration documentation, see the Junos® OS Network Interfaces and Junos OS Class of Service Configuration Guide.

Egress Protection LSPs for Layer 2 Circuits

An egress protection LSP provides link protection for link between PE routers and CE devices as illustrated in [Figure 2 on page 7](#).

Figure 2: Egress protection LSP



9040633

Device CE1 is multihomed to router PE1 and router PE2. Device CE2 is multihomed to router PE3 and router PE4. There are two paths connecting devices CE1 and CE2. The working path is CE2-PE3-P1-PE1-CE1, using pseudowire PW1. The protecting path is CE2-PE3-P2-PE2-CE1, using pseudowire PW2. Normally, traffic flows through the working path. When the end-to-end OAM between devices CE1 and CE2 detects a failure on the working path, traffic will be switched from the working path to the protecting path.

In the topology shown in [Figure 2 on page 7](#), if there was a link or node failure in the core network (for example, a link failure from router P1 to PE1, from router PE3 to P1, or a node failure of router P1), MPLS fast reroute can be triggered on the transport LSPs between router PE3 and router PE1 to repair the connection within tens of milliseconds. Egress protection LSPs address the problem of when a link failure occurs at the edge of the network (for example, a link failure on router PE1 to device CE1).

An egress protection LSP has been configured from router PE1 to router PE2. In the event of a link failure between router PE1 and device CE1, traffic can be switched to the egress protection LSP. Traffic from device CE2 can now be routed through path PE3-P1-PE1-PE2 to reach device CE1.

CHAPTER 2

Introduction to Configuring Layer 2 Circuits

- [Configuring LDP for Layer 2 Circuits on page 9](#)

Configuring LDP for Layer 2 Circuits

Use LDP as the signaling protocol to advertise ingress labels to the remote PE routers. When configured, LDP examines the Layer 2 circuit configuration and initiates extended neighbor discovery for all the Layer 2 circuit neighbors (for example, remote PEs). This process is similar to how LDP works when tunneled over RSVP. You must run LDP on the **lo0.0** interface for extended neighbor discovery to function correctly.

For detailed information about how to configure LDP, see the Junos OS MPLS Applications Configuration Guide.

PART 2

Configuration

- [Configuring Layer 2 Circuits on page 13](#)
- [Layer 2 Circuits Examples on page 53](#)

CHAPTER 3

Configuring Layer 2 Circuits

- [Configuring an IGP on the PE and P Routers on page 13](#)
- [Configuring IBGP Sessions Between PE Routers in VPNs on page 14](#)
- [Configuring a Signaling Protocol and LSPs for VPNs on page 15](#)
- [Configuring Routing Instances on PE Routers in VPNs on page 18](#)
- [Configuring Policies for the VRF Table on PE Routers in VPNs on page 23](#)
- [Configuring Local Interface Switching in Layer 2 Circuits on page 30](#)
- [Configuring Interfaces for Layer 2 Circuits on page 32](#)
- [Configuring Static Layer 2 Circuits on page 40](#)
- [Configuring Policies for Layer 2 Circuits on page 41](#)
- [Enabling BGP Path Selection for Layer 2 VPNs and VPLS on page 45](#)
- [Configuring ATM Trunking on Layer 2 Circuits on page 47](#)
- [Configuring Bandwidth Allocation and Call Admission Control in Layer 2 Circuits on page 48](#)
- [Reducing APS Switchover Time in Layer 2 Circuits on page 49](#)

Configuring an IGP on the PE and P Routers

For Layer 2 VPNs, Layer 3 VPNs, virtual-router routing instances, VPLS, and Layer 2 circuits to function properly, the service provider's PE and P routers must be able to exchange routing information. To allow them to do this, 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 Configuration Guide](#)

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 (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 (**lo0**) on the local PE router. The IBGP session for the VPN runs through the loopback address. (You must also configure the **lo0** 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 either Layer 2 VPNs and VPLS 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 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.

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 15](#)
- [Using RSVP for VPN Signaling on page 16](#)

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 Junos OS Routing Protocols Configuration Guide.

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 **label-switched-path** and **interface** statements at the **[edit protocols mpls]** hierarchy level:

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

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 mpls** and **family inet** 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 Junos OS MPLS Applications Configuration Guide.

Configuring Routing Instances on PE Routers in VPNs

You need to configure a routing instance for each VPN on each of the PE routers participating in the VPN. The configuration procedures outlined in this section are applicable to Layer 2 VPNs, Layer 3 VPNs, and VPLS. The configuration procedures specific to each type of VPN are described in the corresponding sections in the other configuration chapters.

To configure routing instances for VPNs, include the following statements:

```
description text;
instance-type type;
interface interface-name;
route-distinguisher (as-number:number | ip-address:number);
vrf-import [ policy-names ];
vrf-export [ policy-names ];
vrf-target {
  export community-name;
  import community-name;
}
```

You can include these statements at the following hierarchy levels:

- **[edit routing-instances *routing-instance-name*]**
- **[edit logical-systems *logical-system-name* routing-instances *routing-instance-name*]**

To configure VPN routing instances, you perform the steps in the following sections:

- [Configuring the Routing Instance Name for a VPN on page 19](#)
- [Configuring the Description on page 19](#)
- [Configuring the Instance Type on page 20](#)
- [Configuring Interfaces for VPN Routing on page 20](#)
- [Configuring the Route Distinguisher on page 22](#)
- [Configuring Automatic Route Distinguishers on page 23](#)

Configuring the Routing Instance Name for a VPN

The name of the routing instance for a VPN can be a maximum of 128 characters and can contain letters, numbers, and hyphens. In Junos OS Release 9.0 and later, you can no longer specify **default** as the actual routing-instance name. You also cannot use any special characters (! @ # \$ % ^ & * , + < > ;) within the name of a routing instance.



NOTE: In Junos OS Release 9.6 and later, you can include a slash (/) in a routing instance name only if a logical system is not configured. That is, you cannot include the slash character in a routing instance name if a logical system other than the default is explicitly configured.

Specify the routing-instance name with the **routing-instance** statement:

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

You can include this statement at the following hierarchy levels:

- **[edit]**
- **[edit logical-systems *logical-system-name*]**

Configuring the Description

To provide a text description for the routing instance, include the **description** statement. If the text includes one or more spaces, enclose them in quotation marks (" "). Any descriptive text you include is displayed in the output of the **show route instance detail** command and has no effect on the operation of the routing instance.

To configure a text description, include the **description** statement:

```
description text;
```

You can include this statement at the following hierarchy levels:

- **[edit routing-instances *routing-instance-name*]**
- **[edit logical-systems *logical-system-name* routing-instances *routing-instance-name*]**

Configuring the Instance Type

The instance type you configure varies depending on whether you are configuring Layer 2 VPNs, Layer 3 VPNs, VPLS, or virtual routers. Specify the instance type by including the **instance-type** statement:

- To enable Layer 2 VPN routing on a PE router, include the **instance-type** statement and specify the value **l2vpn**:

```
instance-type l2vpn;
```

- To enable VPLS routing on a PE router, include the **instance-type** statement and specify the value **vpls**:

```
instance-type vpls;
```

- Layer 3 VPNs require that each PE router have a VPN routing and forwarding (VRF) table for distributing routes within the VPN. To create the VRF table on the PE router, include the **instance-type** statement and specify the value **vrf**:

```
instance-type vrf;
```



NOTE: Routing Engine based sampling is not supported on VRF routing instances.

- To enable the virtual-router routing instance, include the **instance-type** statement and specify the value **virtual-router**:

```
instance-type virtual-router;
```

You can include this statement at the following hierarchy levels:

- [edit routing-instances *routing-instance-name*]
- [edit logical-systems *logical-system-name* routing-instances *routing-instance-name*]

Configuring Interfaces for VPN Routing

On each PE router, you must configure an interface over which the VPN traffic travels between the PE and CE routers.

The sections that follow describe how to configure interfaces for VPNs:

- [General Configuration for VPN Routing on page 20](#)
- [Configuring Interfaces for Layer 3 VPNs on page 21](#)
- [Configuring Interfaces for Carrier-of-Carriers VPNs on page 21](#)
- [Configuring Unicast RPF on VPN Interfaces on page 22](#)

General Configuration for VPN Routing

The configuration described in this section applies to all types of VPNs. For Layer 3 VPNs and carrier-of-carriers VPNs, complete the configuration described in this section before proceeding to the interface configuration sections specific to those topics.

To configure interfaces for VPN routing, include the **interface** statement:

```
interface interface-name;
```

You can include this statement at the following hierarchy levels:

- **[edit routing-instances *routing-instance-name*]**
- **[edit logical-systems *logical-system-name* routing-instances *routing-instance-name*]**

Specify both the physical and logical portions of the interface name, in the following format:

```
physical.logical
```

For example, in **at-1/2/1.2**, **at-1/2/1** is the physical portion of the interface name and **2** is the logical portion. If you do not specify the logical portion of the interface name, the value **0** is set by default.

A logical interface can be associated with only one routing instance. If you enable a routing protocol on all instances by specifying **interfaces all** when configuring the master instance of the protocol at the **[edit protocols]** hierarchy level, and if you configure a specific interface for VPN routing at the **[edit routing-instances *routing-instance-name*]** hierarchy level or at the **[edit logical-systems *logical-system-name* routing-instances *routing-instance-name*]** hierarchy level, the latter interface statement takes precedence and the interface is used exclusively for the VPN.

If you explicitly configure the same interface name at the **[edit protocols]** hierarchy level and at either the **[edit routing-instances *routing-instance-name*]** or **[edit logical-systems *logical-system-name* routing-instances *routing-instance-name*]** hierarchy levels, an attempt to commit the configuration fails.

Configuring Interfaces for Layer 3 VPNs

When you configure the Layer 3 VPN interfaces at the **[edit interfaces]** hierarchy level, you must also configure **family inet** when configuring the logical interface:

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

Configuring Interfaces for Carrier-of-Carriers VPNs

When you configure carrier-of-carriers VPNs, you need to configure the **family mpls** statement in addition to the **family inet** statement for the interfaces between the PE and CE routers. For carrier-of-carriers VPNs, configure the logical interface as follows:

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

```
}
```

If you configure **family mpls** on the logical interface and then configure this interface for a non-carrier-of-carriers routing instance, the **family mpls** statement is automatically removed from the configuration for the logical interface, since it is not needed.

Configuring Unicast RPF on VPN Interfaces

For VPN interfaces that carry IP version 4 or version 6 (IPv4 or IPv6) traffic, you can reduce the impact of denial-of-service (DoS) attacks by configuring unicast reverse path forwarding (RPF). Unicast RPF helps determine the source of attacks and rejects packets from unexpected source addresses on interfaces where unicast RPF is enabled.

You can configure unicast RPF on a VPN interface by enabling unicast RPF on the interface and including the **interface** statement at the **[edit routing-instances routing-instance-name]** hierarchy level.

You cannot configure unicast RPF on the core-facing interfaces. You can only configure unicast RPF on the CE router-to-PE router interfaces on the PE router. However, for virtual-router routing instances, unicast RPF is supported on all interfaces you specify in the routing instance.

For information about how to configure unicast RPF on VPN interfaces, see the Junos® OS Network Interfaces.

Configuring the Route Distinguisher

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 Layer 2 VPNs and VPLS, if you have configured the **l2vpn-use-bgp-rules** statement, you must configure a unique route distinguisher for each PE router participating in a specific routing instance.

For other types of VPNs, we recommend that you use a unique route distinguisher for each PE router participating in the routing instance. Although you can use the same route distinguisher on all PE routers for the same VPN routing instance (except for Layer 2 VPNs and VPLS), if you use a unique route distinguisher, you can determine the CE router from which a route originated within the VPN.

To configure a route distinguisher on a PE router, include the **route-distinguisher** statement:

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

You can include this statement at the following hierarchy levels:

- **[edit routing-instances routing-instance-name]**
- **[edit logical-systems logical-system-name routing-instances routing-instance-name]**

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.

Configuring Automatic Route Distinguishers

If you configure the **route-distinguisher-id** statement at the **[edit routing-options]** hierarchy level, a route distinguisher is automatically assigned to the routing instance. If you also configure the **route-distinguisher** statement in addition to the **route-distinguisher-id** statement, the value configured for **route-distinguisher** supersedes the value generated from **route-distinguisher-id**.

To assign a route distinguisher automatically, include the **route-distinguisher-id** statement:

```
route-distinguisher-id ip-address;
```

You can include this statement at the following hierarchy levels:

- **[edit routing-options]**
- **[edit logical-systems *logical-system-name* routing-options]**

A type 1 route distinguisher is automatically assigned to the routing instance using the format ***ip-address:number***. The IP address is specified by the **route-distinguisher-id** statement and the number is unique for the routing instance.

Related Documentation

- [Configuring Policies for the VRF Table on PE Routers in VPNs on page 23](#)
- [Configuring BGP Route Target Filtering for VPNs](#)

Configuring Policies for the VRF Table on PE Routers in VPNs

On each PE router, you must define policies that define how routes are imported into and exported from the router's VRF table. In these policies, you must define the route target, and you can optionally define the route origin.

To configure policy for the VRF tables, you perform the steps in the following sections:

- [Configuring the Route Target on page 24](#)
- [Configuring the Route Origin on page 24](#)
- [Configuring an Import Policy for the PE Router's VRF Table on page 25](#)

- [Configuring an Export Policy for the PE Router's VRF Table on page 27](#)
- [Applying Both the VRF Export and the BGP Export Policies on page 28](#)
- [Configuring a VRF Target on page 29](#)

Configuring the Route Target

As part of the policy configuration for the VPN routing table, you must define a route target, which defines which VPN the route is a part of. When you configure different types of VPN services (Layer 2 VPNs, Layer 3 VPNs, or VPLS) on the same PE router, be sure to assign unique route target values to avoid the possibility of adding route and signaling information to the wrong VPN routing table.

To configure the route target, include the **target** option in the **community** statement:

```
community name members target:community-id;
```

You can include this statement at the following hierarchy levels:

- **[edit policy-options]**
- **[edit logical-systems *logical-system-name* policy-options]**

name is the name of the community.

community-id is the identifier of the community. Specify it in one of the following formats:

- ***as-number:number***, where *as-number* is an AS number (a 2-byte value) and *number* is a 4-byte community value. The AS number can be in the range 1 through 65,535. We recommend that you use an IANA-assigned, nonprivate AS number, preferably the ISP's own or the customer's own AS number. The community value can be a number in the range 0 through 4,294,967,295 ($2^{32} - 1$).
- ***ip-address:number***, where *ip-address* is an IPv4 address (a 4-byte value) and *number* is a 2-byte community 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. The community value can be a number in the range 1 through 65,535.

Configuring the Route Origin

In the import and export policies for the PE router's VRF table, you can optionally assign the route origin (also known as the site of origin) for a PE router's VRF routes using a VRF export policy applied to multiprotocol external BGP (MP-EBGP) VPN IPv4 route updates sent to other PE routers.

Matching on the assigned route origin attribute in a receiving PE's VRF import policy helps ensure that VPN-IPv4 routes learned through MP-EBGP updates from one PE are not reimported to the same VPN site from a different PE connected to the same site.

To configure a route origin, complete the following steps:

1. Include the **community** statement with the **origin** option:

```
community name members origin:community-id;
```


You can include this statement at the following hierarchy levels:

- **[edit policy-options]**
- **[edit logical-systems *logical-system-name* policy-options]**

name is the name of the community.

community-id is the identifier of the community. Specify it in one of the following formats:

- **as-number:number**, where **as-number** is an AS number (a 2-byte value) and **number** is a 4-byte community value. The AS number can be in the range 1 through 65,535. We recommend that you use an IANA-assigned, nonprivate AS number, preferably the ISP's own or the customer's own AS number. The community value can be a number in the range 0 through 4,294,967,295 ($2^{32} - 1$).
 - **ip-address:number**, where **ip-address** is an IPv4 address (a 4-byte value) and **number** is a 2-byte community 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. The community value can be a number in the range 1 through 65,535.
2. Include the community in the import policy for the PE router's VRF table by configuring the **community** statement with the **community-id** identifier defined in Step 1 at the **[edit policy-options policy-statement import-policy-name term import-term-name from]** hierarchy level. See [“Configuring an Import Policy for the PE Router's VRF Table” on page 25](#).

If the policy's **from** clause does not specify a community condition, the **vrf-import** statement in which the policy is applied cannot be committed. The Junos OS commit operation does not pass the validation check.
 3. Include the community in the export policy for the PE router's VRF table by configuring the **community** statement with the **community-id** identifier defined in Step 1 at the **[edit policy-options policy-statement export-policy-name term export-term-name then]** hierarchy level. See [“Configuring an Export Policy for the PE Router's VRF Table” on page 27](#).

See Route Origin for VPNs for a configuration example.

Configuring an Import Policy for the PE Router's VRF Table

Each VPN can have a policy that defines how routes are imported into the PE router's VRF table. An import policy is applied to routes received from other PE routers in the VPN. A policy must evaluate all routes received over the IBGP session with the peer PE router. If the routes match the conditions, the route is installed in the PE router's **routing-instance-name.inet.0** VRF table. An import policy must contain a second term that rejects all other routes.

Unless an import policy contains only a **then reject** statement, it must include a reference to a community. Otherwise, when you try to commit the configuration, the commit fails. You can configure multiple import policies.

An import policy determines what to import to a specified VRF table based on the VPN routes learned from the remote PE routers through IBGP. The IBGP session is configured at the **[edit protocols bgp]** hierarchy level. If you also configure an import policy at the **[edit protocols bgp]** hierarchy level, the import policies at the **[edit policy-options]** hierarchy level and the **[edit protocols bgp]** hierarchy level are combined through a logical AND operation. This allows you to filter traffic as a group.

To configure an import policy for the PE router's VRF table, follow these steps:

1. To define an import policy, include the **policy-statement** statement. For all PE routers, an import policy must always include the **policy-statement** statement, at a minimum:

```
policy-statement import-policy-name {  
  term import-term-name {  
    from {  
      protocol bgp;  
      community community-id;  
    }  
    then accept;  
  }  
  term term-name {  
    then reject;  
  }  
}
```

You can include the **policy-statement** statement at the following hierarchy levels:

- **[edit policy-options]**
- **[edit logical-systems *logical-system-name* policy-options]**

The ***import-policy-name*** policy evaluates all routes received over the IBGP session with the other PE router. If the routes match the conditions in the **from** statement, the route is installed in the PE router's *routing-instance-name*.inet.0 VRF table. The second term in the policy rejects all other routes.

For more information about creating policies, see the Routing Policy Configuration Guide.

2. You can optionally use a regular expression to define a set of communities to be used for the VRF import policy.

For example you could configure the following using the **community** statement at the **[edit policy-options policy-statement *policy-statement-name*]** hierarchy level:

```
[edit policy-options vrf-import-policy-sample]  
community high-priority members *:50
```

Note that you cannot configure a regular expression as a part of a route target extended community. For more information about how to configure regular expressions for communities, see the Routing Policy Configuration Guide.

3. To configure an import policy, include the **vrf-import** statement:

```
vrf-import import-policy-name;
```

You can include this statement at the following hierarchy levels:

- [edit routing-instances *routing-instance-name*]
- [edit logical-systems *logical-system-name* routing-instances *routing-instance-name*]

Configuring an Export Policy for the PE Router's VRF Table

Each VPN can have a policy that defines how routes are exported from the PE router's VRF table. An export policy is applied to routes sent to other PE routers in the VPN. An export policy must evaluate all routes received over the routing protocol session with the CE router. (This session can use the BGP, OSPF, or Routing Information Protocol [RIP] routing protocols, or static routes.) If the routes match the conditions, the specified community target (which is the route target) is added to them and they are exported to the remote PE routers. An export policy must contain a second term that rejects all other routes.

Export policies defined within the VPN routing instance are the only export policies that apply to the VRF table. Any export policy that you define on the IBGP session between the PE routers has no effect on the VRF table. You can configure multiple export policies.

To configure an export policy for the PE router's VRF table, follow these steps:

1. For all PE routers, an export policy must distribute VPN routes to and from the connected CE routers in accordance with the type of routing protocol that you configure between the CE and PE routers within the routing instance.

To define an export policy, include the **policy-statement** statement. An export policy must always include the **policy-statement** statement, at a minimum:

```
policy-statement export-policy-name {
  term export-term-name {
    from protocol (bgp | ospf | rip | static);
    then {
      community add community-id;
      accept;
    }
  }
  term term-name {
    then reject;
  }
}
```



NOTE: Configuring the **community add** statement is a requirement for Layer 2 VPN VRF export policies. If you change the **community add** statement to the **community set** statement, the router at the egress of the Layer 2 VPN link might drop the connection.



NOTE: When configuring draft-rosen multicast VPNs operating in source-specific mode and using the `vrf-export` statement to specify the export policy, the policy must have a term that accepts routes from the `vrf-name.mdt.0` routing table. This term ensures proper PE autodiscovery using the `inet-mdt` address family.

When configuring draft-rosen multicast VPNs operating in source-specific mode and using the `vrf-target` statement, the VRF export policy is automatically generated and automatically accepts routes from the `vrf-name.mdt.0` routing table.

You can include the **policy-statement** statement at the following hierarchy levels:

- **[edit policy-options]**
- **[edit logical-systems *logical-system-name* policy-options]**

The ***export-policy-name*** policy evaluates all routes received over the routing protocol session with the CE router. (This session can use the BGP, OSPF, or RIP routing protocols, or static routes.) If the routes match the conditions in the **from** statement, the community target specified in the **then community add** statement is added to them and they are exported to the remote PE routers. The second term in the policy rejects all other routes.

For more information about creating policies, see the Routing Policy Configuration Guide.

2. To apply the policy, include the **vrf-export** statement:

```
vrf-export export-policy-name;
```

You can include this statement at the following hierarchy levels:

- **[edit routing-instances *routing-instance-name*]**
- **[edit logical-systems *logical-system-name* routing-instances *routing-instance-name*]**

Applying Both the VRF Export and the BGP Export Policies

When you apply a VRF export policy as described in “[Configuring an Export Policy for the PE Router's VRF Table](#)” on page 27, routes from VPN routing instances are advertised to other PE routers based on this policy, whereas the BGP export policy is ignored.

If you include the **vpn-apply-export** statement in the BGP configuration, both the VRF export and BGP group or neighbor export policies are applied (VRF first, then BGP) before routes are advertised in the VPN routing tables to other PE routers.

When you include the **vpn-apply-export** statement, be aware of the following:

- Routes imported into the `l3vpn.bgp.0` routing table retain the attributes of the original routes (for example, an OSPF route remains an OSPF route even when it is stored in the `l3vpn.bgp.0` routing table). You should be aware of this when you configure an

export policy for connections between an IBGP PE router and a PE router, a route reflector and a PE router, or AS boundary router (ASBR) peer routers.

- By default, all routes in the l3vpn.bgp.0 routing table are exported to the IBGP peers. If the last statement of the export policy is deny all and if the export policy does not specifically match on routes in the l3vpn.bgp.0 routing table, no routes are exported.

To apply both the VRF export and BGP export policies to VPN routes, include the **vpn-apply-export** statement:

```
vpn-apply-export;
```

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

Configuring a VRF Target

Including the **vrf-target** statement in the configuration for a VRF target community causes default VRF import and export policies to be generated that accept and tag routes with the specified target community. You can still create more complex policies by explicitly configuring VRF import and export policies. These policies override the default policies generated when you configure the **vrf-target** statement.

If you do not configure the **import** and **export** options of the **vrf-target** statement, the specified community string is applied in both directions. The **import** and **export** keywords give you more flexibility, allowing you to specify a different community for each direction.

The syntax for the VRF target community is not a name. You must specify it in the format **target:x:y**. A community name cannot be specified because this would also require you to configure the community members for that community using the **policy-options** statement. If you define the **policy-options** statements, then you can just configure VRF import and export policies as usual. The purpose of the **vrf-target** statement is to simplify the configuration by allowing you to configure most statements at the **[edit routing-instances]** hierarchy level.

To configure a VRF target, include the **vrf-target** statement:

```
vrf-target community;
```

You can include this statement at the following hierarchy levels:

- **[edit routing-instances *routing-instance-name*]**
- **[edit logical-systems *logical-system-name* routing-instances *routing-instance-name*]**

An example of how you might configure the **vrf-target** statement follows:

```
[edit routing-instances sample]
vrf-target target:69:102;
```

To configure the **vrf-target** statement with the **export** and **import** options, include the following statements:

```
vrf-target {
  export community-name;
```

```
import community-name;  
}
```

You can include this statement at the following hierarchy levels:

- [edit routing-instances *routing-instance-name*]
- [edit logical-systems *logical-system-name* routing-instances *routing-instance-name*]

Configuring Local Interface Switching in Layer 2 Circuits

You can configure a virtual circuit entirely on the local router, terminating the circuit on a local interface. Possible uses for this feature include being able to enable switching between Frame Relay DLCIs.

To configure a virtual circuit to terminate locally, include the **local-switching** statement:

```
local-switching {  
  interface interface-name {  
    description text;  
    end-interface {  
      interface interface-name;  
      no-revert;  
      protect-interface interface-name;  
    }  
    ignore-mtu-mismatch;  
    no-revert;  
    protect-interface interface-name;  
  }  
}
```

You can include this statement at the following hierarchy levels:

- [edit protocols l2circuit]
- [edit logical-systems *logical-system-name* protocols l2circuit]

The following sections describe how to configure local interface switching:

- [Configuring the Interfaces for the Local Interface Switch on page 30](#)
- [Enabling Local Interface Switching When the MTU Does Not Match on page 31](#)

Configuring the Interfaces for the Local Interface Switch

Local interface switching requires you to configure at least two interfaces:

- Starting interface—Include the **interface** statement at the [edit protocols l2circuit **local-switching**] hierarchy level.
- Ending interface—Include the **end-interface** statement at the [edit protocols l2circuit **local-switching interface *interface-name***] hierarchy level.

You can also configure virtual circuit interface protection for each local interface:

- Protect interface for the starting interface—Include the **protect-interface** statement at the **[edit protocols l2circuit local-switching interface *interface-name*]** hierarchy level.
- Protect interface for the ending interface—Include the **protect-interface** statement at the **[edit protocols l2circuit local-switching interface *interface-name* end-interface]** hierarchy level.

For more information about how to configure protect interfaces, see [“Configuring the Protect Interface” on page 37](#).

Typically, when the primary interface goes down, the pseudowire starts using the protect interface. By default, when the primary interface comes back online, the interface is switched-over back from the protect interface to the primary interface. To prevent the switchover back to the primary interface, unless the primary interface goes down, include the **no-revert** statement. This prevents loss of traffic during the switchover.



NOTE: If the protect interface fails, the interface is switched-over back to the primary interface, irrespective of whether or not the **no-revert** statement is included in the configuration.

You can configure the **no-revert** statement both for the starting interface and the ending interface.

```
[edit protocols l2circuit local-switching interface interface-name]
no-revert;
end-interface {
  interface interface-name;
  no-revert;
}
```



NOTE: The protect interface must be configured prior to configuring the **no-revert** statement.

Enabling Local Interface Switching When the MTU Does Not Match

You can configure a local switching interface to ignore the MTU configuration set for the associated physical interface. This enables you to bring up a circuit between two logical interfaces that are defined on physical interfaces with different MTU values.

To configure the local switching interface to ignore the MTU configured for the physical interface, include the **ignore-mtu-mismatch** statement:

```
ignore-mtu-mismatch;
```

You can include this statement at the following hierarchy levels:

- **[edit protocols l2circuit local-switching interface *interface-name*]**

- [edit logical-systems *logical-system-name* protocols l2circuit local-switching interface *interface-name*]

Configuring Interfaces for Layer 2 Circuits

The following sections describe how to configure interfaces for Layer 2 circuits:

- [Configuring the Address for the Neighbor of the Layer 2 Circuit on page 32](#)
- [Configuring the Neighbor Interface for the Layer 2 Circuit on page 32](#)
- [Configuring the Interface Encapsulation Type for Layer 2 Circuits on page 39](#)
- [Configuring ATM2 IQ Interfaces for Layer 2 Circuits on page 40](#)

Configuring the Address for the Neighbor of the Layer 2 Circuit

All the Layer 2 circuits using a particular remote PE router designated for remote CE routers are listed under the **neighbor** statement (“neighbor” designates the PE router). Each neighbor is identified by its IP address and is usually the end-point destination for the label-switched path (LSP) tunnel transporting the Layer 2 circuit.

To configure a PE router as a neighbor for a Layer 2 circuit, specify the neighbor address using the **neighbor** statement:

```
neighbor address {  
  ...  
}
```

You can include this statement at the following hierarchy levels:

- [edit protocols l2circuit]
- [edit logical-systems *logical-system-name* protocols l2circuit]

Configuring the Neighbor Interface for the Layer 2 Circuit

Each Layer 2 circuit is represented by the logical interface connecting the local provider edge (PE) router to the local customer edge (CE) router. This interface is tied to the Layer 2 circuit neighbor configured in [“Configuring the Address for the Neighbor of the Layer 2 Circuit” on page 32](#).

To configure the interface for a Layer 2 circuit neighbor, include the **interface** statement:

```
interface interface-name {  
  bandwidth (bandwidth | ctnumber bandwidth);  
  community community-name;  
  (control-word | no-control-word);  
  description text;  
  encapsulation-type type;  
  ignore-encapsulation-mismatch;  
  ignore-mtu-mismatch;  
  mtu mtu-number;  
  no-revert;  
  protect-interface interface-name;  
  pseudowire-status-tlv;  
  psn-tunnel-endpoint address;
```



```
virtual-circuit-id identifier;  
}
```

You can include this statement at the following hierarchy levels:

- **[edit protocols l2circuit neighbor *address*]**
- **[edit logical-systems *logical-system-name* protocols l2circuit neighbor *address*]**

The following sections describe how to configure the interface for the Layer 2 circuit neighbor:

- [Configuring a Community for the Layer 2 Circuit on page 33](#)
- [Configuring the Control Word for Layer 2 Circuits on page 33](#)
- [Configuring the Encapsulation Type for the Layer 2 Circuit Neighbor Interface on page 35](#)
- [Enabling the Layer 2 Circuit When the Encapsulation Does Not Match on page 35](#)
- [Configuring the MTU for the Layer 2 Circuit Neighbor Interface on page 36](#)
- [Configuring the Protect Interface on page 37](#)
- [Configuring the Protect Interface From Switching Over to the Primary Interface on page 37](#)
- [Configuring the Pseudowire Status TLV on page 38](#)
- [Configuring Layer 2 Circuits over Both RSVP and LDP LSPs on page 38](#)
- [Configuring the Virtual Circuit ID on page 39](#)

Configuring a Community for the Layer 2 Circuit

To configure a community for a Layer 2 circuit, include the **community** statement:

```
community community-name;
```

You can include this statement at the following hierarchy levels:

- **[edit protocols l2circuit neighbor *address* interface *interface-name*]**
- **[edit logical-systems *logical-system-name* protocols l2circuit neighbor *address* interface *interface-name*]**

For information about how to configure a routing policy for a Layer 2 circuit, see [“Configuring Policies for Layer 2 Circuits” on page 41](#).

Configuring the Control Word for Layer 2 Circuits

To emulate the virtual circuit (VC) encapsulation for Layer 2 circuits, a 4-byte control word is added between the Layer 2 protocol data unit (PDU) being transported and the VC label that is used for demultiplexing. For most protocols, a null control word consisting of all zeroes is sent between Layer 2 circuit neighbors.

However, individual bits are available in a control word that can carry Layer 2 protocol control information. The control information is mapped into the control word, which allows the header of a Layer 2 protocol to be stripped from the frame. The remaining data and control word can be sent over the Layer 2 circuit, and the frame can be reassembled with the proper control information at the egress point of the circuit.

The following Layer 2 protocols map Layer 2 control information into special bit fields in the control word:

- Frame Relay—The control word supports the transport of discard eligible (DE), forward explicit congestion notification (FECN), and backward explicit congestion notification (BECN) information. For configuration information, see [“Configuring the Control Word for Frame Relay Interfaces” on page 34](#).



NOTE: Frame Relay is not supported on the ACX Series routers.

- ATM AAL5 mode—The control word supports the transport of sequence number processing, ATM cell loss priority (CLP), and explicit forward congestion indication (EFCI) information. When you configure an AAL5 mode Layer 2 circuit, the control information is carried by default and no additional configuration is needed.
- ATM cell-relay mode—The control word supports sequence number processing only. When you configure a cell-relay mode Layer 2 circuit, the sequence number information is carried by default and no additional configuration is needed.

The Junos OS implementation of sequence number processing for ATM cell-relay mode and AAL5 mode is not the same as that described in Sec. 3.1.2 of the IETF draft *Encapsulation Methods for Transport of Layer 2 Frames Over IP and MPLS Networks*. The differences are as follows:

- A packet with a sequence number of 0 is considered as out of sequence.
- A packet that does not have the next incremental sequence number is considered out of sequence.
- When out-of-sequence packets arrive, the sequence number in the Layer 2 circuit control word increments by one and becomes the expected sequence number for the neighbor.

The following sections discuss how to configure the control word for Layer 2 circuits:

- [Configuring the Control Word for Frame Relay Interfaces on page 34](#)
- [Disabling the Control Word for Layer 2 Circuits on page 35](#)

Configuring the Control Word for Frame Relay Interfaces

On interfaces with Frame Relay CCC encapsulation, you can configure Frame Relay control bit translation to support Frame Relay services over IP and MPLS backbones by using CCC, Layer 2 VPNs, and Layer 2 circuits. When you configure translation of Frame Relay control bits, the bits are mapped into the Layer 2 circuit control word and preserved across the IP or MPLS backbone.

For information about how to configure the control bits, see the Junos® OS Network Interfaces and the Junos OS Feature Guides.

Disabling the Control Word for Layer 2 Circuits

The Junos OS can typically determine whether a neighboring router supports the control word. However, if you want to explicitly disable its use on a specific interface, include the **no-control-word** statement:

```
no-control-word;
```

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

Configuring the Encapsulation Type for the Layer 2 Circuit Neighbor Interface

You can specify the Layer 2 circuit encapsulation type for the interface receiving traffic from a Layer 2 circuit neighbor. The encapsulation type is carried in the LDP-signaling messages exchanged between Layer 2 circuit neighbors when pseudowires are created. The encapsulation type you configure for each Layer 2 circuit neighbor varies depending on the type of networking equipment or the type of Layer 2 protocol you have deployed in your network. If you do not specify an encapsulation type for the Layer 2 circuit, the encapsulation of the CE device interface is used by default.

Specify the encapsulation type for the Layer 2 circuit neighbor interface by including the **encapsulation-type** statement:

```
encapsulation-type (atm-aal5 | atm-cell | atm-cell-port-mode | atm-cell-vc-mode |
  atm-cell-vp-mode | cesop | cisco-hdlc | ethernet | ethernet-vlan | frame-relay |
  frame-relay-port-mode | interworking | ppp | satop-e1 | satop-e3 | satop-t1 | satop-t3);
```

You can include this statement at the following hierarchy levels:

- **[edit protocols l2circuit neighbor address interface *interface-name*]**
- **[edit logical-systems *logical-system-name* protocols l2circuit neighbor address interface *interface-name*]**

Enabling the Layer 2 Circuit When the Encapsulation Does Not Match

You can configure the Junos OS to allow a Layer 2 circuit to be established even though the encapsulation configured on the CE device interface does not match the encapsulation configured on the Layer 2 circuit interface by including the **ignore-encapsulation-mismatch** statement. You can configure the **ignore-encapsulation-mismatch** statement for the connection to the remote connection by including the statement at the **[edit protocols l2circuit neighbor address interface *interface-name*]** hierarchy level or for the local connection by including this statement at the **[edit protocols l2circuit local-switching interface *interface-name*]** hierarchy level.

```
ignore-encapsulation-mismatch;
```

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

Configuring the MTU for the Layer 2 Circuit Neighbor Interface

The following sections describe how to configure the MTU for the Layer 2 circuit neighbor interface:

- [Enabling the Layer 2 Circuit When the MTU Does Not Match on page 36](#)
- [Configuring the MTU Advertised for a Layer 2 Circuit on page 36](#)

Enabling the Layer 2 Circuit When the MTU Does Not Match

You can configure the Junos OS to allow a Layer 2 circuit to be established even though the MTU configured on the PE router does not match the MTU configured on the remote PE router by including the `ignore-mtu-mismatch` statement:

`ignore-mtu-mismatch;`

You can include this statement at the following hierarchy levels:

- `[edit protocols l2circuit neighbor address interface interface-name]`
- `[edit logical-systems logical-system-name protocols l2circuit neighbor address interface interface-name]`

Configuring the MTU Advertised for a Layer 2 Circuit

By default, the MTU used to advertise a Layer 2 circuit is determined by taking the interface MTU for the associated physical interface and subtracting the encapsulation overhead for sending IP packets based on the encapsulation.

However, encapsulations that support multiple logical interfaces (and multiple Layer 2 circuits) rely on the same interface MTU (since they are all associated with the same physical interface). This can prove to be a limitation for VLAN Layer 2 circuits using the same Ethernet interface or for Layer 2 circuit DLCIs using the same Frame Relay interface.

This can also affect multivendor environments. For example, if you have three PE devices supplied by different vendors and one of the devices only supports an MTU of 1500, even if the other devices support larger MTUs you must configure the MTU as 1500 (the smallest MTU of the three PE devices).

You can explicitly configure which MTU is advertised for a Layer 2 circuit, even if the Layer 2 circuit is sharing a physical interface with other Layer 2 circuits. When you explicitly configure an MTU for a Layer 2 circuit, be aware of the following:

- An explicitly configured MTU is signaled to the remote PE device. The configured MTU is also compared to the MTU received from the remote PE device. If there is a conflict, the Layer 2 circuit is taken down.
- If you configure an MTU for an ATM cell relay interface on an ATM II PIC, the configured MTU is used to compute the cell bundle size advertised for that Layer 2 circuit, instead of the default interface MTU.
- A configured MTU is used only in the control plane. It is not enforced in the data plane. You need to ensure that the CE device for a given Layer 2 circuit uses the correct MTU for data transmission.

To configure the MTU for a Layer 2 circuit, include the **mtu** statement:

```
mtu mtu-number;
```

You can include this statement at the following hierarchy levels:

- [edit protocols l2circuit neighbor *address* interface *interface-name*]
- [edit logical-systems *logical-system-name* protocols l2circuit neighbor *address* interface *interface-name*]

Configuring the Protect Interface

You can configure a protect interface for the logical interface linking a virtual circuit to its destination, whether the destination is remote or local. A protect interface provides a backup for the protected interface in case of failure. Network traffic uses the primary interface only so long as the primary interface functions. If the primary interface fails, traffic is switched to the protect interface. The protect interface is optional.

To configure the protect interface, include the **protect-interface** statement:

```
protect-interface interface-name;
```



NOTE: The protect interface must be configured prior to configuring the **no-revert** statement.

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

For an example of how to configure a protect interface for a Layer 2 circuit, see “[Example: Configuring Layer 2 Circuit Protect Interfaces](#)” on page 53.

Configuring the Protect Interface From Switching Over to the Primary Interface

Typically, when the primary interface goes down, the pseudowire starts using the protect interface. By default, when the primary interface comes back online, the interface is switched-over back from the protect interface to the primary interface. To prevent the switchover back to the primary interface, unless the protect interface goes down, include the **no-revert** statement. This prevents loss of traffic during the switchover.



NOTE: If the protect interface fails, the interface is switched-over back to the primary interface, irrespective of whether or not the **no-revert** statement is included in the configuration.

You can configure the **no-revert** statement at the [edit protocols l2circuit neighbor *address* interface *interface-name*] hierarchy level:

```
[edit protocols l2circuit neighbor address interface interface-name]  
no-revert;
```

Configuring the Pseudowire Status TLV

The pseudowire status type length variable (TLV) is used to communicate the status of a pseudowire back and forth between two PE routers. For Layer 2 circuit configurations, you can configure the PE router to negotiate the pseudowire with its neighbor using the pseudowire status TLV. This same functionality is also available for LDP VPLS neighbor configurations. The pseudowire status TLV is configurable for each pseudowire connection and is disabled by default. The pseudowire status negotiation process assures that a PE router reverts back to the label withdraw method for pseudowire status if its remote PE router neighbor does not support the pseudowire status TLV.

Unlike the control word, a PE router's ability to support the pseudowire status TLV is communicated when the initial label mapping message is sent to its remote PE router. Once the PE router transmits its support for the pseudowire status TLV to its remote PE router, it includes the pseudowire status TLV in every label mapping message sent to the remote PE router. If you disable support for the pseudowire status TLV on the PE router, a label withdraw message is sent to the remote PE router and then a new label mapping message without the pseudowire status TLV follows.

To configure the pseudowire status TLV for the pseudowire to the neighbor PE router, include the **pseudowire-status-tlv** statement:

```
pseudowire-status-tlv;
```

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

Configuring Layer 2 Circuits over Both RSVP and LDP LSPs

You can configure two Layer 2 circuits between the same two routers, and have one Layer 2 circuit traverse an RSVP LSP and the other traverse an LDP LSP. To accomplish this, you need to configure two loopback addresses on the local router. You configure one of the loopback address for the Layer 2 circuit traversing the RSVP LSP. You configure the other loopback address to handle the Layer 2 circuit traversing the LDP LSP. For information about how to configure multiple loop back interfaces, see Configuring Logical Units on the Loopback Interface for Routing Instances in Layer 3 VPNs.

You also need to configure a packet switched network (PSN) tunnel endpoint for one of the Layer 2 circuits. It can be either the Layer 2 circuit traversing the RSVP LSP or the one traversing the LDP LSP. The PSN tunnel endpoint address is the destination address for the LSP on the remote router.

To configure the address for the PSN tunnel endpoint, include the **psn-tunnel-endpoint** statement:

```
psn-tunnel-endpoint address;
```

You can include this statement at the following hierarchy levels:

- **[edit logical-systems *logical-system-name* protocols l2circuit neighbor *address* interface *interface-name*]**
- **[edit protocols l2circuit neighbor *address* interface *interface-name*]**

By default, the PSN tunnel endpoint for a Layer 2 circuit is identical to the neighbor address, which is also the same as the LDP neighbor address.

The tunnel endpoints on the remote router do not need to be loopback addresses.

Example: PSN Tunnel Endpoint

The following example illustrates how you might configure a PSN tunnel endpoint:

```
[edit protocols l2circuit]
neighbor 10.255.0.6 {
  interface t1-0/2/2.0 {
    psn-tunnel-endpoint 20.20.20.20;
    virtual-circuit-id 1;
  }
  interface t1-0/2/1.0 {
    virtual-circuit-id 10;
  }
}
```

The Layer 2 circuit configured for the **t1-0/2/2.0** interface resolves in the inet3 routing table to **20.20.20.20**. This could be either an RSVP route or a static route with an LSP next hop.

Configuring the Virtual Circuit ID

You configure a virtual circuit ID on each interface. Each virtual circuit ID uniquely identifies the Layer 2 circuit among all the Layer 2 circuits to a specific neighbor. The key to identifying a particular Layer 2 circuit on a PE router is the neighbor address and the virtual circuit ID. An LDP-FEC-to-label binding is associated with a Layer 2 circuit based on the virtual circuit ID in the FEC and the neighbor that sent this binding. The LDP-FEC-to-label binding enables the dissemination of the VPN label used for sending traffic on that Layer 2 circuit to the remote CE device.

You also configure a virtual circuit ID for each redundant pseudowire. A redundant pseudowire is identified by the backup neighbor address and the virtual circuit ID. For more information, see *Configuring Redundant Pseudowires for Layer 2 Circuits and VPLS*.

To configure the virtual circuit ID, include the **virtual-circuit-id** statement:

```
virtual-circuit-id identifier;
```

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

Configuring the Interface Encapsulation Type for Layer 2 Circuits

The Layer 2 encapsulation type is carried in the LDP forwarding equivalence class (FEC). You can configure either circuit cross-connect (CCC) or translational cross-connect (TCC) encapsulation types for Layer 2 circuits. For more information, see the *Junos OS MPLS Applications Configuration Guide and Router Interfaces*.

To configure the interface encapsulation for a Layer 2 circuit, include the **encapsulation** statement:

```
encapsulation encapsulation;
```

You can include this statement at the following hierarchy levels:

- **[edit interfaces *interface-name*]**
- **[edit logical-systems *logical-system-name* interfaces *interface-name*]**

Configuring ATM2 IQ Interfaces for Layer 2 Circuits

You can configure Asynchronous Transfer Mode 2 (ATM2) intelligent queuing (IQ) interfaces for Layer 2 circuits by using Layer 2 circuit ATM Adaptation Layer 5 (AAL5) transport mode, Layer 2 circuit ATM cell relay mode, and the Layer 2 circuit ATM trunk mode.

The configuration statements are as follows:

- **atm-l2circuit-mode aal5**
- **atm-l2circuit-mode cell**
- **atm-l2circuit-mode trunk**

For more information about these statements, see the Junos OS System Basics Configuration Guide. For more information about how to configure ATM2 IQ interfaces, see the Junos® OS Network Interfaces.

The Junos OS implementation of sequence number processing for Layer 2 circuit ATM cell relay mode and Layer 2 circuit AAL5 mode differs from that described in the Internet draft draft-martini-l2circuit-encap-mpls-11.txt, *Encapsulation Methods for Transport of Layer 2 Frames over MPLS Networks* (expires August 2006).

The Junos OS implementation has the following differences:

1. A packet with a sequence number of 0 is treated as out of sequence.
2. A packet that does not have the next incremental sequence number is considered out of sequence.

When out-of-sequence packets arrive, the expected sequence number for the neighbor is set to the sequence number in the Layer 2 circuit control word.

Configuring Static Layer 2 Circuits

You can configure static Layer 2 circuit pseudowires. Static pseudowires are designed for networks that do not support LDP or do not have LDP enabled. You configure a static pseudowire by configuring static values for the in and out labels needed to enable a pseudowire connection. The **ignore-mtu-mismatch**, **ignore-vlan-id**, and **ignore-encapsulation-mismatch** statements are not relevant for static pseudowire configurations since the peer router cannot forward this information.

When you configure static pseudowires, you need to manually compare the encapsulation, TDM bit rate, and control word of the router with the remote peer router and ensure that they match, otherwise the static pseudowire might not work.

To configure static Layer 2 circuit pseudowires, include the **static** statement:

```
static {  
    incoming-label label;  
    outgoing-label label;  
}
```

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

You can configure a static pseudowire as a standalone Layer 2 circuit or in conjunction with a redundant pseudowire. You configure the static pseudowire statement at the **[edit protocols l2circuit neighbor address interface *interface-name*]** hierarchy level. You configure the redundant pseudowire at the **[edit protocols l2circuit neighbor address interface *interface-name* backup-neighbor *neighbor*]** hierarchy level. If you configure a static pseudowire to a neighbor and also configure a redundant pseudowire, the redundant pseudowire must also be static.

For information about how to configure redundant pseudowires, see *Configuring Redundant Pseudowires for Layer 2 Circuits and VPLS*.

Configuring Policies for Layer 2 Circuits

You can configure Junos routing policies to control the flow of packets over Layer 2 circuits. This capability allows you to provide different level of service over a set of equal-cost Layer 2 circuits. For example, you can configure a circuit for high-priority traffic, a circuit for average-priority traffic, and a circuit for low-priority traffic. By configuring Layer 2 circuit policies, you can ensure that higher-value traffic has a greater likelihood of reaching its destination.

The following sections explain how to configure Layer 2 circuit policies:

- [Configuring the Layer 2 Circuit Community on page 41](#)
- [Configuring the Policy Statement for the Layer 2 Circuit Community on page 42](#)
- [Verifying the Layer 2 Circuit Policy Configuration on page 44](#)

Configuring the Layer 2 Circuit Community

To configure a community for Layer 2 circuits, include the **community** statement.

```
community community-name {  
    members [ community-ids ];  
}
```

You can include this statement at the following hierarchy levels:

- **[edit policy-options]**
- **[edit logical-systems *logical-system-name* policy-options]**

name identifies the community or communities.

community-ids identifies the type of community or extended community:

- A normal community uses the following community ID format:

as-number:community-value

as-number is the autonomous system (AS) number of the community member.

community-value is the identifier of the community member. It can be a number from 0 through 65,535.

- An extended community uses the following community ID format:

type:administrator:assigned-number

type is the type of target community. The target community identifies the route's destination.

administrator is either an AS number or an IP version 4 (IPv4) address prefix, depending on the type of community.

assigned-number identifies the local provider.

You also need to configure the community for the Layer 2 circuit interface; see [“Configuring a Community for the Layer 2 Circuit” on page 33](#).

Configuring the Policy Statement for the Layer 2 Circuit Community

To configure a policy to send community traffic over a specific LSP, include the **policy-statement** statement:

```
policy-statement policy-name {  
  term term-name {  
    from community community-name;  
    then {  
      install-nexthop (except | lsp lsp-name | lsp-regex lsp-regular-expression);  
      accept;  
    }  
  }  
}
```

You can include this statement at the following hierarchy levels:

- [edit policy-options]
- [edit logical-systems *logical-system-name* policy-options]

To prevent the installation of any matching next hops, include the **install-nexthop** statement with the **except** option:

```
install-nexthop except;
```

You can include this statement at the following hierarchy levels:

- [edit policy-options policy-statement *policy-name* term *term-name* then]
- [edit logical-systems *logical-system-name* policy-options policy-statement *policy-name* term *term-name* then]

To assign traffic from a community to a specific LSP, include the **install-nexthop** statement with the **lsp *lsp-name*** option and the **accept** statement:

```
install-nexthop lsp lsp-name;
accept;
```

You can include these statements at the following hierarchy levels:

- [edit policy-options policy-statement *policy-name* term *term-name* then]
- [edit logical-systems *logical-system-name* policy-options policy-statement *policy-name* term *term-name* then]

You can also use a regular expression to select an LSP from a set of similarly named LSPs for the **install-nexthop** statement. To configure a regular expression, include the **install-nexthop** statement with the **lsp-regex** option and the **accept** statement:

```
install-nexthop lsp-regex lsp-regular-expression;
accept;
```

You can include these statements at the following hierarchy levels:

- [edit policy-options policy-statement *policy-name* term *term-name* then]
- [edit logical-systems *logical-system-name* policy-options policy-statement *policy-name* term *term-name* then]

Example: Configuring a Policy for a Layer 2 Circuit Community

The following example illustrates how you might configure a regular expression in a Layer 2 circuit policy. You create three LSPs to handle gold-tier traffic from a Layer 2 circuit. The LSPs are named **alpha-gold**, **beta-gold**, and **delta-gold**. You then include the **install-nexthop** statement with the **lsp-regex** option with the LSP regular expression **.*-gold** at the [edit policy-options policy-statement *policy-name* term *term-name* then] hierarchy level:

```
[edit policy-options]
policy-statement gold-traffic {
  term to-gold-LSPs {
    from community gold;
    then {
      install-nexthop lsp-regex .*-gold;
      accept;
    }
  }
}
```

The community **gold** Layer 2 circuits can now use any of the **-gold** LSPs. Given equal utilization across the three **-gold** LSPs, LSP selection is made at random.

You need to apply the policy to the forwarding table. To apply a policy to the forwarding table, configure the **export** statement at the [edit routing-options forwarding-table] hierarchy level:

```
[edit routing-options forwarding-table]
export policy-name;
```

Verifying the Layer 2 Circuit Policy Configuration

To verify that you have configured a policy for the Layer 2 circuit, issue the **show route table mpls detail** command. It should display the community for ingress routes that corresponds to the Layer 2 circuits, as shown by the following example:

```
user@host> show route table mpls detail
so-1/0/1.0 (1 entry, 1 announced)
*L2VPN Preference: 7
Next hop: via so-1/0/0.0 weight 1, selected
Label-switched-path to-community-gold
Label operation: Push 100000 Offset: -4
Next hop: via so-1/0/0.0 weight 1
Label-switched-path to-community-silver
Label operation: Push 100000 Offset: -4
Protocol next hop: 10.255.245.45
Push 100000 Offset: -4
Indirect next hop: 85333f0 314
State: <Active Int>
Local AS: 100
Age: 22
Task: Common L2 VC
Announcement bits (2): 0-KRT 1-Common L2 VC
AS path: I
Communities: 100:1
```

For more information about how to configure routing policies, see the Routing Policy Configuration Guide.

Enabling BGP Path Selection for Layer 2 VPNs and VPLS

Layer 2 VPNs and VPLS share the same path selection process for determining the optimal path to reach all of the destinations shared within a single routing instance. For Layer 2 VPN and VPLS topologies, the path selection process is straightforward if there is just a single path from each PE router to each CE device. However, the path selection process becomes more complex if the PE routers receive two or more valid paths to reach a specific CE device.

The following network scenarios provide examples of what might cause a PE router to receive more than one valid path to reach a specific CE device:

- **Multihoming**—One or more CE devices within a routing instance are multihomed to two or more PE routers. Each multihomed CE device has at least two valid paths.
- **Route reflectors**—There are multiple route reflectors deployed within the same network and they are supporting PE routers within the same routing instance. Due to time delays in large complex networks, the route reflectors can separately receive a different valid path to reach a CE device at different times. When they readvertise these valid paths, a PE router could receive two or more separate but apparently valid paths to the same CE device.

By default, Juniper Networks routers use just the designated forwarder path selection algorithm to select the best path to reach each Layer 2 VPN or VPLS routing instance destination (for more information, see *VPLS Path Selection Process for PE Routers*). However, you can also configure the routers in your network to use both the BGP path selection algorithm and the designated forwarder path selection algorithm as follows:

- On the Provider routers within the service providers network, the standard BGP path selection algorithm is used (for more information, see *Understanding BGP Path Selection*). Using the standard BGP path selection for Layer 2 VPN and VPLS routes allows a service provider to leverage the existing Layer 3 VPN network infrastructure to also support Layer 2 VPNs and VPLS. The BGP path selection algorithm also helps to ensure that the service provider's network behaves predictably with regard to Layer 2 VPN and VPLS path selection. This is particularly important in networks employing route reflectors and multihoming.

When a Provider router receives multiple paths for the same destination prefix (for example, a multihomed CE device), one path is selected based on the BGP path selection algorithm and placed in the `bgp.l2vpn.0` routing table and the appropriate *instance.l2vpn.0* routing table.

- When a PE router receives all of the available paths to each CE device, it runs the designated forwarder path selection algorithm to select the preferred path to reach each CE device, independently of the results of the earlier BGP path selection algorithm run on the Provider router. The VPLS designated forwarder algorithm uses the D-bit, preference, and PE router identifier to determine which of the valid paths to each CE device to use. The PE router might select a path to reach a CE device which is different from the path selected by the BGP-based Provider routers. In this scenario, the following is the expected behavior for traffic sent to the multihomed CE device:
 - If the path selected by the remote PE router is available, traffic will traverse the network to the multihomed CE device using the remote PE router's preferred path (again, ignoring the path selected by the BGP-based Provider routers).
 - If the path selected by the remote PE router fails:
 1. The Provider routers switch the traffic destined for the multihomed CE device to the alternate path as soon as failure is detected.
 2. The Provider routers notify the remote PE routers of the path failure.
 3. The remote PE routers update their routing tables accordingly.

For more information about the VPLS designated forwarder path selection algorithm, see VPLS Path Selection Process for PE Routers. This algorithm is also described in the Internet draft draft-kompella-l2vpn-vpls-multihoming-03.txt, *Multi-homing in BGP-based Virtual Private LAN Service*.

To enable the BGP path selection algorithm for Layer 2 VPN and VPLS routing instances, complete the following steps:

1. Run Junos OS Release 12.3 or later on all of the PE and Provider routers participating in Layer 2 VPN or VPLS routing instances.

Attempting to enable this functionality on a network with a mix of routers that both do and do not support this feature can result in anomalous behavior.

2. Specify a unique route distinguisher on each PE router participating in a Layer 2 VPN or VPLS routing instance.
3. Configure the `l2vpn-use-bgp-rules` statement on all of the PE and Provider routers participating in Layer 2 VPN or VPLS routing instances.

You can configure this statement at the `[edit protocols bgp path-selection]` hierarchy level to apply this behavior to all of the routing instances on the router or at the `[edit routing-instances routing-instance-name protocols bgp path-selection]` hierarchy level to apply this behavior to a specific routing instance.

Related Documentation

- Understanding BGP Path Selection
- VPLS Path Selection Process for PE Routers
- [l2vpn-use-bgp-rules on page 151](#)
- route-distinguisher

Configuring ATM Trunking on Layer 2 Circuits

You can configure Layer 2 circuits to transport ATM traffic from directly connected ATM switches across an MPLS core network. Traffic from an ATM switch is received on the local PE router. The ATM cells are given an MPLS label and then sent across the MPLS network to the remote PE router. The receiving router removes the MPLS label from the ATM cell and then forwards the cell the receiving ATM switch.



NOTE: ATM trunking on Layer 2 circuits is supported only on T Series and M320 routers and ATM2 IQ PICs.

Figure 3: ATM Trunking on Layer 2 Circuits




Figure 3 on page 47 illustrates how ATM switches could be linked together by a Layer 2 circuit. The PE1 Router is configured to receive ATM trunk traffic from ATM Switch 1. As each ATM cell is received on the PE1 Router, it is classified by means of the class-of-service (CoS) information in the cell header and then encapsulated as a labeled packet. The CoS information and cell loss priority (CLP) of the ATM cell are copied into the experimental (EXP) bits of the MPLS label. The labeled packet is then transported across the service provider network to the PE2 Router by means of a Layer 2 circuit.

On the PE2 Router, the label is removed and the plain ATM cell is forwarded to ATM Switch 2. The CoS and CLP are extracted from the EXP bits and are then used to select the correct output queue and determine whether the ATM cell should be dropped.

The ATM physical port on the router can support 32 logical trunks when network-to-network interface (NNI) is used and 8 logical trunks when user-to-network interface (UNI) is used. A trunk can carry traffic on 32 virtual path identifiers (VPs), numbered 0 through 31. Each ATM trunk is associated with an MPLS label and a logical interface. On the ingress router, one or more of these trunks are mapped to a Layer 2 circuit.

The configuration for the Layer 2 circuit between PE routers is conventional. Follow the procedures outlined in this chapter for configuring the circuit. However, there is some specific configuration you need to complete for the Layer 2 circuit to carry traffic from an ATM trunk.

First, enable ATM trunking for Layer 2 circuits. To enable ATM trunking for Layer 2 circuits, specify the **trunk** option for the **atm-l2circuit-mode** statement at the **[edit chassis fpc number pic number]** hierarchy level:

```
[edit chassis fpc number pic number]
```

```
atm-l2circuit-mode trunk (uni | nni);
```

Specify the **uni** option for UNI trunks and the **nni** option for NNI trunks. The default option is **uni**.

You also need to configure each ATM trunk for a specific logical interface. Each ATM trunk has a trunk identifier in the range from 0 to 31. This configuration step is in addition to the typical configuration steps you follow related to configuring interfaces for Layer 2 circuits, as described in [“Configuring Interfaces for Layer 2 Circuits” on page 32](#).

To associate a specific trunk identifier with a logical interface, include the **trunk-id** statement:

```
trunk-id number;
```

You can include this statement at the following hierarchy levels:

- **[edit interfaces *interface-name* unit *number*]**
- **[edit logical-systems *logical-system-name* interfaces *interface-name* unit *number*]**

Since ATM trunking is supported on ATM2 IQ PICs only, the only value you can configure for the **pic-type** statement is **atm2**. If you do not configure the **pic-type** statement but you do configure the **trunk** option for the **atm-l2circuit-mode** statement (at the **[chassis *fpc number* pic *number*]** hierarchy level), the **pic-type** statement defaults to **atm2**.

Configuring Bandwidth Allocation and Call Admission Control in Layer 2 Circuits

You can configure bandwidth allocation and call admission control (CAC) on Layer 2 circuits. This feature is available for RSVP-signaled LSPs traversing an MPLS network.

When you enable bandwidth allocation on a Layer 2 circuit, attempts to establish an RSVP-signaled LSP are preceded by a check of the available bandwidth on the network. This check is the CAC. The available bandwidth is compared to the bandwidth requested by the LSP. If there is insufficient bandwidth, the Layer 2 circuit is not established and an error message is generated. To apply CAC to a Layer 2 circuit, a bandwidth constraint must be configured.

You can specify the bandwidth for a Layer 2 circuit without configuring a bandwidth for each class type (queue). To specify the bandwidth allocation for a Layer 2 circuit, include the **bandwidth** statement:

```
bandwidth bandwidth;
```

Specify the bandwidth in bits per second.

You can include this statement at the following hierarchy levels:

- **[edit protocols l2circuit neighbor *address* interface *interface-name*]**
- **[edit logical-systems *logical-system-name* protocols l2circuit neighbor *address* interface *interface-name*]**

Alternatively, you can configure the bandwidth for each class type on a Layer 2 circuit. If you use this type of configuration, you cannot simultaneously configure the nonclass type of bandwidth configuration for the Layer 2 circuit (the commit operation fails).

To configure the bandwidth for each class type on an Layer 2 circuit, include the **bandwidth** statement:

```
bandwidth {  
  ct0 bandwidth;  
  ct1 bandwidth;  
  ct2 bandwidth;  
  ct3 bandwidth;  
}
```

You can include this statement at the following hierarchy levels:

- [edit protocols l2circuit neighbor *address* interface *interface-name*]
- [edit logical-systems *logical-system-name* protocols l2circuit neighbor *address* interface *interface-name*]

Specify the bandwidth for each class type in bits per second. It is not necessary to specify a bandwidth for all four class types.

Reducing APS Switchover Time in Layer 2 Circuits

On M320 routers with Channelized OC3/STM1 Circuit Emulation PIC with SFP, you can configure the **fast-aps-switch** statement at the [edit interfaces *interface-name* **sonet-options aps**] hierarchy level to reduce the Automatic Protection Switching (APS) switchover time in Layer 2 circuits. Configuring this statement reduces the APS switchover time only when the Layer 2 circuit encapsulation type for the interface receiving traffic from a Layer 2 circuit neighbor is Structure Agnostic time-division multiplexing (TDM) over Packet (SAToP).

The **fast-aps-switch** statement must be configured on both working and protect circuits. Additionally, to achieve reduction in APS switchover time:

- Per-packet load balancing must be configured.
- Bidirectional switching mode must be configured.
- If the **fast-aps-switch** statement is configured in revertive APS mode, configure an appropriate value for revert time. We recommend that you configure a revert time of 600 seconds for 672 through 1344 Layer 2 circuits.
- To prevent the logical interfaces in the data path from being shut down, configure appropriate hold-time values on all the interfaces in the data path that support TDM.

**NOTE:**

- The **fast-aps-switch** statement cannot be configured when the APS annex-b option is configured.
- The interfaces that have the **fast-aps-switch** statement configured cannot be used in virtual private LAN service (VPLS) environments.

The following tasks illustrate how to configure Junos OS to reduce APS switchover time.



NOTE: Per-packet load balancing can be configured for a limited set of routes or for all routes. To simplify the steps involved in configuring per-packet load balancing, steps for configuring per-packet load balancing for all routes is covered in this procedure.

- [Configuring Per-Packet Load Balancing on page 50](#)
- [Configuring Fast APS Switchover on page 50](#)

Configuring Per-Packet Load Balancing

To configure per-packet load balancing for all routes:

1. Configure the **per-packet** option for the **load-balance** statement at the **[edit policy-options policy-statement *policy-name* then]** hierarchy level.

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

For example:

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

2. Configure the policy name in the **export** statement at the **[edit routing-options forwarding-table]** hierarchy level.

```
[edit routing-options forwarding-table]
user@host# set export policy-name
```

For example:

```
[edit routing-options forwarding-table]
user@host# set export load-balancing-policy
```

Configuring Fast APS Switchover

To configure fast APS switchover:

1. On both the working and protect circuits, configure the **fast-aps-switch** statement at the **[edit interfaces *interface-name* sonet-options aps]** hierarchy level.

```
[edit interfaces interface-name sonet-options aps]
user@host# set fast-aps-switch
```

For example:

```
[edit interfaces cstm1-0/0/0 sonet-options aps]
user@host# set fast-aps-switch

[edit interfaces cstm1-0/1/0 sonet-options aps]
user@host# set fast-aps-switch
```

2. Configure bidirectional switching mode on both the working and protect circuits. To do this, configure the **switching-mode bidirectional** statement at the **[edit interfaces *interface-name* sonet-options aps]** hierarchy level on both the working and protect circuits.

```
[edit interfaces interface-name sonet-options aps]
user@host# set switching-mode bidirectional
```

For example:

```
[edit interfaces cstm1-0/1/0 sonet-options aps]
user@host# set switching-mode bidirectional

[edit interfaces cstm1-0/1/0 sonet-options aps]
user@host# set switching-mode bidirectional
```

3. If APS is configured in revertive mode, configure an appropriate value for revert time on both the working and protect circuits. To do this, configure the **revert-time** statement at the **[edit interfaces *interface-name* sonet-options aps]** hierarchy level on both the working and protect circuits.

```
[edit interfaces interface-name sonet-options aps]
user@host# set revert-time seconds
```

For example:

```
[edit interfaces cstm1-0/0/0 sonet-options aps]
user@host# set revert-time 600

[edit interfaces cstm1-0/1/0 sonet-options aps]
user@host# set revert-time 600
```

4. To prevent the logical interfaces in the data path from being shut down, configure appropriate hold-time values on *all* interfaces in the data path that support TDM.

```
[edit interfaces interface-name hold-time]
user@host# set up seconds down seconds
```

For example:

```
[edit interfaces cstm1-0/0/0 hold-time]
user@host# set up 1 down 400
```


CHAPTER 4

Layer 2 Circuits Examples


- [Example: Configuring Layer 2 Circuit Protect Interfaces on page 53](#)
- [Example: Configuring Layer 2 Circuit Switching Protection on page 58](#)
- [Example: Configuring an Egress Protection LSP for a Layer 2 Circuit on page 70](#)
- [Using the Layer 2 Interworking Interface to Interconnect a Layer 2 Circuit to a Layer 2 VPN on page 80](#)
- [Example: Interconnecting a Layer 2 Circuit with a Layer 2 VPN on page 81](#)
- [Applications for Interconnecting a Layer 2 Circuit with a Layer 2 Circuit on page 90](#)
- [Example: Interconnecting a Layer 2 Circuit with a Layer 2 Circuit on page 90](#)
- [Applications for Interconnecting a Layer 2 Circuit with a Layer 3 VPN on page 106](#)
- [Example: Interconnecting a Layer 2 Circuit with a Layer 3 VPN on page 106](#)

Example: Configuring Layer 2 Circuit Protect Interfaces

This example illustrates how you might configure a Layer 2 circuit with protect interfaces. Protect interfaces act as backups for their associated interfaces. The primary interface has priority over the protect interface and carries network traffic as long as it is functional. If the primary interface fails, the protect interface is activated. These interfaces can also share the same virtual path identifier (VPI) or virtual circuit identifier (VCI).

[Figure 4 on page 54](#) shows the network topology used in this example.

Figure 4: Layer 2 Circuits Using Protect Interfaces



The following sections describe how to configure a Layer 2 circuit to use a protect interface:

- [Configuring Router PE1 on page 54](#)
- [Configuring Router PE2 on page 56](#)
- [Configuring Router CE1 on page 57](#)
- [Configuring Router CE2 on page 58](#)

Configuring Router PE1

Configure an interface for traffic to Router CE1 from Router PE1 at the **[edit interfaces]** hierarchy level:

```
[edit interfaces]
so-0/2/2 {
  description "Router CE1 so-0/2/2";
  no-keepalives;
  encapsulation frame-relay-ccc;
  unit 1 {
    encapsulation frame-relay-ccc;
    point-to-point;
    dlci 600;
  }
  unit 2 {
    encapsulation frame-relay-ccc;
    point-to-point;
    dlci 602;
  }
}
```

Configure an interface for traffic to Router CE1 from Router PE1 at the **[edit interfaces]** hierarchy level. Logical interface **so-0/2/0.2** acts as the protect interface for **so-0/2/2.2**, and logical interface **so-0/2/0.1** acts as the protect interface for **so-0/2/2.1**:

```
[edit interfaces]
so-0/2/0 {
  description "to Router CE1 so-0/3/0";
```

```

no-keepalives;
encapsulation frame-relay-ccc;
unit 1 {
    encapsulation frame-relay-ccc;
    dlci 600;
}
unit 2 {
    encapsulation frame-relay-ccc;
    dlci 602;
}
}

```

Configure an interface for traffic to Router PE2 from Router PE1 at the **[edit interfaces]** hierarchy level:

```

[edit interfaces]
so-0/2/1 {
    description "to Router PE2 so-1/0/1";
    unit 0 {
        family inet {
            address 100.100.40.22/32 {
                destination 100.100.40.23;
            }
        }
        family iso;
        family mpls;
    }
}

```

Configure an interface for traffic to Router PE2 from Router PE1 at the **[edit interfaces]** hierarchy level:

```

[edit interfaces]
so-0/2/3 {
    description "Router PE2 so-1/0/3";
    unit 0 {
        family inet;
        family iso;
        family mpls;
    }
    lo0 {
        unit 0 {
            family inet {
                address 127.0.0.1/32;
                address 10.100.40.200/32;
            }
            family iso {
                address 47.0005.80ff.f800.0000.0108.0001.1921.6800.4213.00;
            }
        }
    }
}

```

Configure the Layer 2 circuit by including the **l2circuit** statement at the **[edit protocols]** hierarchy level. The logical interfaces for the Layer 2 circuits and their corresponding protect interfaces are included here:

```
[edit protocols]
l2circuit {
  neighbor 10.100.40.210 {
    interface so-0/2/2.2 {
      protect-interface so-0/2/0.2;
      virtual-circuit-id 2;
      no-control-word;
    }
    interface so-0/2/2.1 {
      protect-interface so-0/2/0.1;
      virtual-circuit-id 1;
      no-control-word;
    }
  }
}
```

Configuring Router PE2

Configure an interface for traffic to Router CE2 from Router PE2:

```
[edit interfaces]
so-1/0/0 {
  description "to Router CE2 so-0/2/0";
  no-keepalives;
  encapsulation frame-relay-ccc;
  unit 1 {
    encapsulation frame-relay-ccc;
    point-to-point;
    dlci 700;
  }
  unit 2 {
    encapsulation frame-relay-ccc;
    point-to-point;
    dlci 702;
  }
}
```

Configure an interface for traffic to Router PE1 from Router PE2:

```
[edit interfaces]
so-1/0/1 {
  description "to Router PE1 so-0/2/1";
  unit 0 {
    family inet {
      address 100.100.40.23/32 {
        destination 100.100.40.22;
      }
    }
    family iso;
    family mpls;
  }
}
```

Configure an interface for traffic to Router PE1 from Router PE2:

```
[edit interfaces]
so-1/0/3 {
```



```

description "to Router PE1 so-0/2/3";
unit 0 {
    family inet;
    family iso;
    family mpls;
}
}
lo0 {
    unit 0 {
        family inet {
            address 127.0.0.1/32;
            address 10.100.40.210/32;
        }
        family iso {
            address 47.0005.80ff.f800.0000.0108.0001.1921.6800.4216.00;
        }
    }
}
}

```

Configure the Layer 2 circuit at the **[edit protocols]** hierarchy level:

```

[edit protocols]
l2circuit {
    neighbor 10.100.40.200 {
        interface so-1/0/0.1 {
            virtual-circuit-id 1;
            no-control-word;
        }
        interface so-1/0/0.2 {
            virtual-circuit-id 2;
            no-control-word;
        }
    }
}
}

```

Configuring Router CE1

Configure an interface for traffic to Router PE1 from Router CE1:

```

[edit interfaces]
so-0/3/0 {
    description "to Router PE1 so-0/2/0";
    no-keepalives;
    encapsulation frame-relay;
    unit 1 {
        dlci 601;
        family inet {
            address 12.12.12.1/24;
        }
    }
}
}

```

Configure an interface for traffic to Router PE1 from Router CE1:

```

[edit interfaces]
so-0/3/1 {
    description "Router PE1 so-0/2/2";
}

```

```
no-keepalives;
encapsulation frame-relay;
unit 0 {
    dlc1 600;
    family inet {
        address 10.10.10.1/24;
        address 11.1.1.1/24;
    }
    family iso;
    family mpls;
}
unit 2 {
    dlc1 602;
    family inet {
        address 13.13.13.1/24;
    }
}
}
```

Configuring Router CE2

Configure an interface for traffic to Router PE2 from Router CE2:

```
[edit interfaces]
so-0/2/0 {
    description "to Router PE2 so-1/0/0";
    no-keepalives;
    encapsulation frame-relay;
    unit 1 {
        dlc1 700;
        family inet {
            address 10.10.10.2/24;
            address 11.1.1.2/24;
            address 12.12.12.2/24;
        }
    }
    unit 2 {
        dlc1 702;
        family inet {
            address 13.13.13.2/24;
        }
    }
}
```

Example: Configuring Layer 2 Circuit Switching Protection

Unlike Layer 2 circuit protect interfaces (see [“Example: Configuring Layer 2 Circuit Protect Interfaces” on page 53](#)), which provide traffic protection for paths configured between the PE routers and CE routers, Layer 2 circuit switching protection provides traffic protection for the paths configured between the PE routers. In the event the path used by a Layer 2 circuit fails, traffic can be switched to an alternate path (or protection path). Switching protection is supported for locally switched Layer 2 circuits and provides 1 to 1 protection for each Layer 2 circuit interface.

When you enable Layer 2 circuit switching protection, each Layer 2 circuit interface requires the following paths:

- Working path—Used by the Layer 2 circuit when working normally.
- Protection path—Used by the Layer 2 circuit when the working path fails.
- [Requirements on page 59](#)
- [Overview on page 59](#)
- [Configuration on page 60](#)

Requirements

This example uses the following hardware and software components:

- MX 3D routers
- Junos OS Release 12.3

Overview

Each working path can be configured to have either a protection path routed directly to the neighboring PE router (as shown in [Figure 5 on page 59](#)) or indirectly using a pseudowire configured through an intermediate PE router (as shown in [Figure 6 on page 60](#) and [Figure 7 on page 60](#)). The protection path provides failure protection for the traffic flowing between the PE routers. Ethernet OAM monitors the status of these paths. When OAM detects a failure, it reroutes the traffic from the failed working path to the protection path. You can configure OAM to revert the traffic automatically to the working path when it is restored. You can also manually switch traffic between the working path, the protection path, and back.




NOTE: Non-stop routing (NSR) and graceful routing engine switchover (GRES) do not support Layer 2 circuit switching protection.

Topology

[Figure 5 on page 59](#) illustrates Layer 2 circuit local switching. There are two OAM sessions running between router PE1 and router PE2. One OAM session is configured over the working path and the other is configured over the protection path.


Figure 5: Connection protection enabled between router PE1 and router PE2



In [Figure 6 on page 60](#) and [Figure 7 on page 60](#), there are two OAM sessions running between router PE1 and router PE2. For Figure 2, one OAM session is configured over the

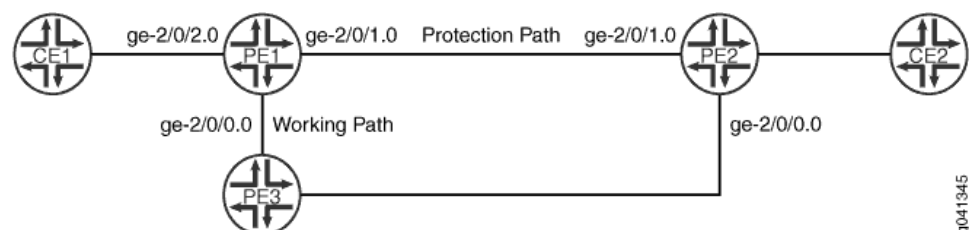
working path between router PE1 and router PE2. The other OAM session is configured over the protection path between router PE1 and router PE3 to router PE2.

Figure 6: Connection Protection Using a Pseudowire Configured through Router PE3 as the Protection Path



For [Figure 7 on page 60](#), one OAM session is configured over the working path, the pseudowire between router PE1 and router PE3, then to router PE2. The other OAM session is configured on the protect path between router PE1 and router PE2.

Figure 7: Connection Protection Using a Pseudowire Configured through Router PE3 as the Working Path



Configuration

The following sections describe how to configure each of the variations of Layer 2 circuit connection protection:

- [Configuring Connection Protection Between Two PE Routers on page 60](#)
- [Configuring Connection Protection Using Another PE Router for the Protection Path on page 64](#)
- [Configuring Connection Protection Using an Another PE Router for the Working Path on page 67](#)

Configuring Connection Protection Between Two PE Routers

Step-by-Step Procedure

To configure Layer 2 Circuit switching protection as shown in [Figure 5 on page 59](#) on router PE1:

1. Configure the Layer 2 circuit on router PE1.

```
[edit protocols l2circuit]
set local-switching interface ge-2/0/2.0 connection-protection
set local-switching interface ge-2/0/2.0 end-interface interface ge-2/0/1.0
set local-switching interface ge-2/0/2.0 end-interface backup-interface ge-2/0/1.1
```
2. Configure the routing policy on router PE1.

```
[edit policy-options]
```

set policy-statement protection-policy then load-balance per-packet

3. Enable the routing policy on router PE1.

```
[edit routing-options]
set routing-options forwarding-table export protection-policy
```

4. Configure OAM on Router PE1. OAM is used to monitor the working path between router PE1 and router PE2. In the event of a failure on the working path, traffic is switched automatically to the protection path. A connectivity fault management (CFM) session is configured on the working path and on the protection path. Begin by configuring the OAM maintenance domain.

```
[edit protocols oam ethernet]
set protocols oam ethernet connectivity-fault-management maintenance-domain
l2circuit-example-md level 5
```

5. Configure OAM on Router PE1 for the working path.

```
[edit protocols oam ethernet]
set protocols oam ethernet connectivity-fault-management maintenance-domain
l2circuit-example-md maintenance-association working continuity-check interval
100ms
set protocols oam ethernet connectivity-fault-management maintenance-domain
l2circuit-example-md maintenance-association working mep 1000 interface
ge-2/0/1.0
set protocols oam ethernet connectivity-fault-management maintenance-domain
l2circuit-example-md maintenance-association working mep 1000 interface
working
set protocols oam ethernet connectivity-fault-management maintenance-domain
l2circuit-example-md maintenance-association working mep 1000 direction down
set protocols oam ethernet connectivity-fault-management maintenance-domain
l2circuit-example-md maintenance-association working mep 1000 remote-mep
103
```

6. Configure OAM on Router PE1 for the protection path.

```
[edit protocols oam ethernet]
set protocols oam ethernet connectivity-fault-management maintenance-domain
l2circuit-example-md maintenance-association protection continuity-check
interval 100ms
set protocols oam ethernet connectivity-fault-management maintenance-domain
l2circuit-example-md maintenance-association protection mep 1001 interface
ge-2/0/1.1
set protocols oam ethernet connectivity-fault-management maintenance-domain
l2circuit-example-md maintenance-association protection mep 1001 interface
protect
set protocols oam ethernet connectivity-fault-management maintenance-domain
l2circuit-example-md maintenance-association protection mep 1001 direction
down
set protocols oam ethernet connectivity-fault-management maintenance-domain
l2circuit-example-md maintenance-association protection mep 1001 remote-mep
104
```

7. Configure the OAM maintenance domain on Router PE2.

```
set protocols oam ethernet connectivity-fault-management maintenance-domain
l2circuit-example-md level 5
```

8. Configure OAM on Router PE2 for the working path.

```
set protocols oam ethernet connectivity-fault-management maintenance-domain
l2circuit-example-md maintenance-association working continuity-check interval
100ms
set protocols oam ethernet connectivity-fault-management maintenance-domain
l2circuit-example-md maintenance-association working mep 103 interface
ge-2/0/1.0
set protocols oam ethernet connectivity-fault-management maintenance-domain
l2circuit-example-md maintenance-association working mep 103 interface working
set protocols oam ethernet connectivity-fault-management maintenance-domain
l2circuit-example-md maintenance-association working mep 103 direction down
set protocols oam ethernet connectivity-fault-management maintenance-domain
l2circuit-example-md maintenance-association working mep 103 remote-mep
1000
```

9. Configure OAM on Router PE2 for the protection path.

```
set protocols oam ethernet connectivity-fault-management maintenance-domain
l2circuit-example-md maintenance-association protection continuity-check
interval 100ms
set protocols oam ethernet connectivity-fault-management maintenance-domain
l2circuit-example-md maintenance-association protection mep 104 interface
ge-2/0/1.1
set protocols oam ethernet connectivity-fault-management maintenance-domain
l2circuit-example-md maintenance-association protection mep 104 interface
protect
set protocols oam ethernet connectivity-fault-management maintenance-domain
l2circuit-example-md maintenance-association protection mep 104 direction
down
set protocols oam ethernet connectivity-fault-management maintenance-domain
l2circuit-example-md maintenance-association protection mep 104 remote-mep
1001
```

Results From configuration mode on Router PE1, confirm your configuration by entering the **show protocols l2circuit**, **show policy-options**, **show routing-options**, and **show protocols oam ethernet** commands. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

```
user@host> show protocols l2circuit
local-switching {
  interface ge-2/0/2.0 {
    connection-protection;
  end-interface {
    interface ge-2/0/1.0;
    backup-interface ge-2/0/1.1;
  }
}
}

user@host> show policy-options
policy-statement protection-policy {
  then {
    load-balance per-packet;
  }
}
```

```

user@host> show routing-options
forwarding-table {
  export protection-policy;
}

user@host> show protocols oam ethernet
connectivity-fault-management {
  maintenance-domain l2circuit-example-md {
    level 5;
    maintenance-association working {
      continuity-check {
        interval 100ms;
      }
      mep 1000 {
        interface ge-2/0/1.0 working;
        direction down;
        remote-mep 103;
      }
    }
    maintenance-association protection {
      continuity-check {
        interval 100ms;
      }
      mep 1001 {
        interface ge-2/0/1.1 protect;
        direction down;
        remote-mep 104;
      }
    }
  }
}

```

From configuration mode on Router PE2, confirm your configuration by entering the **show protocols oam ethernet** commands. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

```

connectivity-fault-management {
  maintenance-domain l2circuit-example-md {
    level 5;
    maintenance-association working {
      continuity-check {
        interval 100ms;
      }
      mep 103 {
        interface ge-2/0/1.0 working;
        direction down;
        remote-mep 1000;
      }
    }
    maintenance-association protection {
      continuity-check {
        interval 100ms;
      }
      mep 104 {
        interface ge-2/0/1.1 protect;
        direction down;
      }
    }
  }
}

```

```

        remote-mep 1001;
    }
}
}

```

Configuring Connection Protection Using Another PE Router for the Protection Path

Step-by-Step Procedure

To configure Layer 2 Circuit switching protection as shown in [Figure 6 on page 60](#) on router PE1:

1. Configure the Layer 2 circuit on router PE1.


```

[edit protocols l2circuit]
set protocols l2circuit local-switching interface ge-2/0/2.0 connection-protection
set protocols l2circuit local-switching interface ge-2/0/2.0 backup-neighbor 2.2.2.2
  virtual-circuit-id 2
set protocols l2circuit local-switching interface ge-2/0/2.0 backup-neighbor 2.2.2.2
  community example
set protocols l2circuit local-switching interface ge-2/0/2.0 end-interface interface
  ge-2/0/1.0

```
2. Configure the routing policy on router PE1.


```

[edit policy-options]
set policy-statement load-balance then load-balance per-packet
set policy-statement protection-policy term protect from community example
set policy-statement protection-policy term protect then install-nexthop lsp-regex
  lsp-protect-*

```
3. Configure the routing options on router PE1.


```

[edit routing-options]
set routing-options forwarding-table export load-balance

```
4. Configure OAM on Router PE1 to setup the maintenance domain. OAM is used to monitor the working path between router PE1 and router PE2. In the event of a failure on the working path, traffic is switched automatically to the protection path.


```

[edit protocols oam ethernet]
set connectivity-fault-management maintenance-domain l2circuit-example-md
  level 5

```
5. Configure OAM on Router PE1 for the working path.


```

[edit protocols oam ethernet]
set connectivity-fault-management maintenance-domain l2circuit-example-md
  maintenance-association working mep 1000 interface ge-2/0/1.0
set connectivity-fault-management maintenance-domain l2circuit-example-md
  maintenance-association working mep 1000 direction down
set connectivity-fault-management maintenance-domain l2circuit-example-md
  maintenance-association working mep 1000 remote-mep 103

```
6. Configure OAM on Router PE1 for the protection path.


```

[edit protocols oam ethernet]
set connectivity-fault-management maintenance-domain l2circuit-example-md
  maintenance-association protection mep 1001 interface ge-2/0/0.0

```



```

set connectivity-fault-management maintenance-domain l2circuit-example-md
maintenance-association protection mep 1001 direction down
set connectivity-fault-management maintenance-domain l2circuit-example-md
maintenance-association protection mep 1001 remote-mep 104

```

7. Configure OAM on Router PE2 to setup the maintenance domain.

```

[edit protocols oam ethernet]
set connectivity-fault-management maintenance-domain l2circuit-example-md
level 5

```

8. Configure OAM on Router PE2 for the working path.

```

[edit protocols oam ethernet]
set connectivity-fault-management maintenance-domain l2circuit-example-md
maintenance-association working mep 103 interface ge-2/0/1.0
set connectivity-fault-management maintenance-domain l2circuit-example-md
maintenance-association working mep 103 direction down
set connectivity-fault-management maintenance-domain l2circuit-example-md
maintenance-association working mep 103 remote-mep 1000

```

9. Configure OAM on Router PE2 for the protection path.

```

[edit protocols oam ethernet]
set connectivity-fault-management maintenance-domain l2circuit-example-md
maintenance-association protection mep 104 interface ge-2/0/0.0
set connectivity-fault-management maintenance-domain l2circuit-example-md
maintenance-association protection mep 104 direction down
set connectivity-fault-management maintenance-domain l2circuit-example-md
maintenance-association protection mep 104 remote-mep 1001

```

Results From configuration mode on Router PE1, confirm your configuration by entering the **show protocols l2circuit**, **show policy-options**, **show routing-options**, and **show protocols oam ethernet** commands. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

```

user@host> show protocols l2circuit
local-switching {
  interface ge-2/0/2.0 {
    connection-protection;
    backup-neighbor 2.2.2.2 {
      virtual-circuit-id 2;
      community example;
    }
  }
  end-interface {
    interface ge-2/0/1.0;
  }
}

user@host> show policy-options
policy-statement load-balance {
  then {
    load-balance per-packet;
  }
}
policy-statement protection-policy {
  term protect {

```

```
    from community example;
    then {
        install-nexthop lsp-regex lsp-protect-*;
    }
}

user@host> show routing-options
forwarding-table {
    export load-balance;
}

user@host> show protocols oam ethernet
connectivity-fault-management {
    maintenance-domain l2circuit-example-md {
        level 5;
        maintenance-association working {
            mep 1000 {
                interface ge-2/0/1.0;
                direction down;
                remote-mep 103;
            }
        }
        maintenance-association protection {
            mep 1001 {
                interface ge-2/0/0.0;
                direction down;
                remote-mep 104;
            }
        }
    }
}
```

From configuration mode on Router PE2, confirm your configuration by entering the **show protocols oam ethernet** commands. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

```
connectivity-fault-management {
    maintenance-domain l2circuit-example-md {
        level 5;
        maintenance-association working {
            mep 103 {
                interface ge-2/0/1.0;
                direction down;
                remote-mep 1000;
            }
        }
        maintenance-association protection {
            mep 104 {
                interface ge-2/0/0.0;
                direction down;
                remote-mep 1001;
            }
        }
    }
}
```

Configuring Connection Protection Using an Another PE Router for the Working Path

Step-by-Step Procedure To configure Layer 2 Circuit switching protection as shown in [Figure 7 on page 60](#) on router PE1:

1. Configure the Layer 2 circuit on router PE1.


```
[edit protocols l2circuit]
set protocols l2circuit neighbor 2.2.2.2 interface ge-2/0/2.0 virtual-circuit-id 2
set protocols l2circuit neighbor 2.2.2.2 interface ge-2/0/2.0 community example
set protocols l2circuit neighbor 2.2.2.2 interface ge-2/0/2.0 connection-protection
set protocols l2circuit neighbor 2.2.2.2 interface ge-2/0/2.0 backup-neighbor 3.3.3.3
  virtual-circuit-id 3
set protocols l2circuit neighbor 2.2.2.2 interface ge-2/0/2.0 backup-neighbor 3.3.3.3
  standby
```
2. Configure the policies on router PE1.


```
[edit policy-options]
set policy-options policy-statement load-balance then load-balance per-packet
set policy-options policy-statement protection-policy term protect from community
  example
set policy-options policy-statement protection-policy term protect then
  install-nexthop lsp-regex lsp-primary
```
3.

```
[edit routing-options]
set routing-options forwarding-table export load-balance
```
4. Configure OAM on Router PE1 to setup the maintenance domain. OAM is used to monitor the working path between router PE1 and router PE2. In the event of a failure on the working path, traffic is switched automatically to the protection path.


```
[edit protocols oam ethernet]
set connectivity-fault-management maintenance-domain l2circuit-example-md
  level 5
```
5. Configure OAM on Router PE1 for the working path.


```
[edit protocols oam ethernet]
set connectivity-fault-management maintenance-domain l2circuit-example-md
  maintenance-association working mep 1000 interface ge-2/0/0.0
set connectivity-fault-management maintenance-domain l2circuit-example-md
  maintenance-association working mep 1000 direction down
set connectivity-fault-management maintenance-domain l2circuit-example-md
  maintenance-association working mep 1000 remote-mep 103
```
6. Configure OAM on Router PE1 for the protection path.


```
[edit protocols oam ethernet]
set connectivity-fault-management maintenance-domain l2circuit-example-md
  maintenance-association protection mep 1001 interface ge-2/0/1.0
set connectivity-fault-management maintenance-domain l2circuit-example-md
  maintenance-association protection mep 1001 direction down
set connectivity-fault-management maintenance-domain l2circuit-example-md
  maintenance-association protection mep 1001 remote-mep 104
```
7. Configure OAM on Router PE2 to setup the maintenance domain.

```
[edit protocols oam ethernet]
set connectivity-fault-management maintenance-domain l2circuit-example-md
level 5
```

8. Configure OAM on Router PE2 for the working path.

```
[edit protocols oam ethernet]
set connectivity-fault-management maintenance-domain l2circuit-example-md
maintenance-association working mep 103 interface ge-2/0/0.0
set connectivity-fault-management maintenance-domain l2circuit-example-md
maintenance-association working mep 103 direction down
set connectivity-fault-management maintenance-domain l2circuit-example-md
maintenance-association working mep 103 remote-mep 1000
```

9. Configure OAM on Router PE2 for the protection path.

```
[edit protocols oam ethernet]
set connectivity-fault-management maintenance-domain l2circuit-example-md
maintenance-association protection mep 104 interface ge-2/0/1.0
set connectivity-fault-management maintenance-domain l2circuit-example-md
maintenance-association protection mep 104 direction down
set connectivity-fault-management maintenance-domain l2circuit-example-md
maintenance-association protection mep 104 remote-mep 1001
```

Results From configuration mode on Router PE1, confirm your configuration by entering the **show protocols l2circuit**, **show policy-options**, **show routing-options**, and **show protocols oam ethernet** commands. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

```
user@host> show protocols l2circuit
neighbor 2.2.2.2 {
  interface ge-2/0/2.0 {
    virtual-circuit-id 2;
    community example;
    connection-protection;
    backup-neighbor 3.3.3.3 {
      virtual-circuit-id 3;
      standby;
    }
  }
}

user@host> show policy-options
policy-statement load-balance {
  then {
    load-balance per-packet;
  }
}
policy-statement protection-policy {
  term protect {
    from community example;
    then {
      install-nexthop lsp-regex lsp-primary;
    }
  }
}
```

```

user@host> show routing-options
forwarding-table {
  export load-balance;
}

user@host> show protocols oam ethernet
connectivity-fault-management {
  maintenance-domain l2circuit-example-md {
    level 5;
    maintenance-association working {
      mep 1000 {
        interface ge-2/0/0.0;
        direction down;
        remote-mep 103;
      }
    }
    maintenance-association protection {
      mep 1001 {
        interface ge-2/0/1.0;
        direction down;
        remote-mep 104;
      }
    }
  }
}

```

From configuration mode on Router PE2, confirm your configuration by entering the **show protocols oam ethernet** commands. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

```

connectivity-fault-management {
  maintenance-domain l2circuit-example-md {
    level 5;
    maintenance-association working {
      mep 103 {
        interface ge-2/0/0.0;
        direction down;
        remote-mep 1000;
      }
    }
    maintenance-association protection {
      mep 104 {
        interface ge-2/0/1.0;
        direction down;
        remote-mep 1001;
      }
    }
  }
}

```

Related Documentation

- [Example: Configuring Layer 2 Circuit Protect Interfaces on page 53](#)

Example: Configuring an Egress Protection LSP for a Layer 2 Circuit

This example shows how to configure an egress protection LSP.

- [Requirements on page 70](#)
- [Egress Protection LSP Overview on page 70](#)
- [Egress Protection LSP Configuration on page 71](#)

Requirements

Egress protection LSPs are supported on Juniper Networks MX Series routers only. This requirement applies to the PE routers facilitating the egress protection LSP.

Egress Protection LSP Overview


If there is a link or node failure in the core network, a protection mechanism such as MPLS fast reroute can be triggered on the transport LSPs between the PE routers to repair the connection within tens of milliseconds. An egress protection LSP addresses the problem of when a link failure occurs at the edge of the network (for example, a link failure between a PE router and a CE device). Egress protection LSPs do not address the problem of a node failure at the edge of the network (for example, a failure of a PE router). An egress protection LSP is an RSVP-signaled ultimate hop popping LSP.

This example includes the following configuration concepts and statements that are unique to the configuration of an egress protection LSP:

- **context-identifier**—Specifies an IPv4 address used to define the pair of PE routers participating in the egress protection LSP. The context identifier is used to assign an identifier to the protector PE router. The identifier is propagated to the other PE routers participating in the network, making it possible for the protected egress PE router to signal the egress protection LSP to the protector PE router.
- **egress-protection**—Configures the protector information for the protected Layer 2 circuit and configures the protector Layer 2 circuit at the **[edit protocols l2circuit]** hierarchy level. Configures an LSP as an egress protection LSP at the **[edit protocols mpls label-switched-path lsp-name]** hierarchy level. It also configures the context identifier at the **[edit protocols mpls]** hierarchy level.
- **protected-l2circuit**—Specifies which Layer 2 circuit is to be protected by the egress protect LSP. This statement includes the following sub-statements: **ingress-pe**, **egress-pe**, and **virtual-circuit-id**. These sub-statements specify the address of the PE router at the ingress of the Layer 2 circuit, the address of the PE router at the egress of the Layer 2 circuit, and the Layer 2 circuit's identifier respectively.
- **protector-interface**—Specify the interface used by the egress protection LSP. In the event of a local link failure to a CE device, the egress protect LSP uses the interface specified to communicate with the protector PE router.
- **protector-pe**—Specify the IPv4 address of the protector PE router. The protector PE router must have a connection to the same CE device as the protected PE router for the egress protect LSP to function. This statement includes the following

sub-statements: **context-identifier** and **lsp**. The **lsp** statement specifies the LSP to be used as the actual egress protection LSP.

Figure 8: Egress Protection LSP Configured from Router PE2 to Router PE3



Pseudowires are configured along two paths, one from router PE1 to router PE2 (pseudowire 12) and one from router PE1 to router PE3 (pseudowire 13). In the event of a failure on the link between router PE2 and device CE2, traffic is switched to the egress protection LSP configured between router PE2 and router PE3 (the protector PE router):

- Device CE1—Traffic origin
- Router PE1—Ingress PE router
- Router PE2—Egress PE router
- Router PE3—Protector PE router
- Device CE2—Traffic destination

This example shows how to configure routers PE1, PE2, and PE3.

Egress Protection LSP Configuration

- [Step-by-Step Procedure on page 73](#)
- [Results on page 77](#)

CLI Quick Configuration

To quickly configure an egress protection LSP, copy the following commands into a text file, modify the interface configurations to match your equipment, remove any line breaks, and then paste the commands into the CLI. This group of set commands is for router PE1.

```
set protocols rsvp interface ge-3/1/0.1
set protocols rsvp interface ge-3/1/3.2
set protocols mpls interface ge-3/1/0.1
set protocols mpls interface ge-3/1/3.2
```

```
set protocols ospf traffic-engineering
set protocols ospf area 0.0.0.0 interface ge-3/1/0.1
set protocols ospf area 0.0.0.0 interface ge-3/1/3.2
set protocols ospf area 0.0.0.0 interface lo0.0 passive
set protocols ldp interface ge-3/1/0.1
set protocols ldp interface ge-3/1/3.2
set protocols ldp interface lo0.0
set protocols l2circuit neighbor 1.1.1.3 interface fe-3/3/0.1 virtual-circuit-id 32
set protocols l2circuit neighbor 1.1.1.3 interface fe-3/3/0.1 egress-protection
  protector-interface ge-3/1/2.2
set protocols l2circuit neighbor 1.1.1.4 interface ge-3/1/2.2 virtual-circuit-id 33
set policy-options policy-statement load-balance-example then load-balance per-packet
set routing-options router-id 1.1.1.2
set routing-options forwarding-table export load-balance-example
```

To quickly configure an egress protection LSP, copy the following commands into a text file, modify the interface configurations to match your equipment, remove any line breaks, and then paste the commands into the CLI. This group of set commands is for router PE2.

```
[edit]
set protocols rsvp tunnel-services
set protocols rsvp interface ge-2/3/0.3
set protocols rsvp interface ge-2/3/1.4 link-protection
set protocols ldp interface ge-2/3/0.3
set protocols ldp interface ge-2/3/1.4
set protocols ldp interface lo0.0
set protocols ldp upstream-label-assignment
set protocols mpls label-switched-path protected-lsp to 2.2.3.4
set protocols mpls label-switched-path protected-lsp egress-protection
set protocols mpls interface ge-2/3/0.3
set protocols mpls interface ge-2/3/1.4
set protocols ospf traffic-engineering
set protocols ospf area 0.0.0.0 interface ge-2/3/0.3
set protocols ospf area 0.0.0.0 interface ge-2/3/1.4
set protocols ospf area 0.0.0.0 interface lo0.0 passive
set protocols l2circuit neighbor 1.1.1.2 interface ge-2/0/2.5 virtual-circuit-id 23
set protocols l2circuit neighbor 1.1.1.2 interface ge-2/0/2.5 egress-protection protector-pe
  1.1.1.4
set protocols l2circuit neighbor 1.1.1.2 interface ge-2/0/2.5 egress-protection protector-pe
  context-identifier 2.2.3.4
set policy-options policy-statement load-balance-example then load-balance per-packet
set routing-options router-id 1.1.1.3
set routing-options forwarding-table export load-balance-example
```

To quickly configure an egress protection LSP, copy the following commands into a text file, modify the interface configurations to match your equipment, remove any line breaks, and then paste the commands into the CLI. This group of set commands is for router PE3.

```
set protocols rsvp tunnel-services
set protocols rsvp interface ge-2/0/0.6
set protocols rsvp interface ge-2/0/1.7
set protocols mpls interface ge-2/0/0.6
set protocols mpls interface ge-2/0/1.7
set protocols mpls egress-protection context-identifier 2.2.3.4 protector
```



```

set protocols ospf traffic-engineering
set protocols ospf area 0.0.0.0 interface ge-2/0/0.6
set protocols ospf area 0.0.0.0 interface ge-2/0/1.7
set protocols ospf area 0.0.0.0 interface lo0.0 passive
set protocols ldp interface ge-2/0/0.6
set protocols ldp interface ge-2/0/1.7
set protocols ldp interface lo0.0
set protocols ldp upstream-label-assignment
set protocols l2circuit neighbor 1.1.1.2 interface ge-2/0/3.21 virtual-circuit-id 42
set protocols l2circuit neighbor 1.1.1.2 interface ge-2/0/3.21 egress-protection
protected-l2circuit PW1
set protocols l2circuit neighbor 1.1.1.2 interface ge-2/0/3.21 egress-protection
protected-l2circuit ingress-pe 1.1.1.2
set protocols l2circuit neighbor 1.1.1.2 interface ge-2/0/3.21 egress-protection
protected-l2circuit egress-pe 1.1.1.3
set protocols l2circuit neighbor 1.1.1.2 interface ge-2/0/3.21 egress-protection
protected-l2circuit virtual-circuit-id 31

```

Step-by-Step Procedure

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 an egress protection LSP, complete the following steps for router PE1:

1. Configure RSVP. Include the interface linked to router PE2 and the interface linked to router PE3.

```

[edit]
user@PE1# edit protocols rsvp
[edit protocols rsvp]
user@PE1# set interface ge-3/1/0.1
[edit protocols rsvp]
user@PE1# set interface ge-3/1/3.2

```

2. Configure LDP. Include the interface linked to router PE2, the interface linked to router PE3, and the loopback interface.

```

[edit]
user@PE1# edit protocols ldp
[edit protocols ldp]
user@PE1# set interface ge-3/1/0.1
[edit protocols ldp]
user@PE1# set interface ge-3/1/3.2
[edit protocols ldp]
user@PE1# set interface lo0.0

```

3. Configure MPLS. Include the interface linked to router PE2 and the interface linked to router PE3.

```

[edit]
user@PE1# edit protocols mpls
[edit protocols mpls]
user@PE1# set interface ge-3/1/0.1
[edit protocols mpls]
user@PE1# set interface ge-3/1/3.2

```

4. Configure OSPF. Include the interface linked to router PE2, the interface linked to router PE3, and the loopback interface in the configuration for the OSPF area.

```
[edit]
user@PE1# edit protocols ospf
[edit protocols ospf]
user@PE1# set interface traffic-engineering
[edit protocols ospf]
user@PE1# set area 0.0.0.0 interface ge-3/1/0.1
[edit protocols ospf]
user@PE1# set area 0.0.0.0 interface ge-3/1/3.2
[edit protocols ospf]
user@PE1# set area 0.0.0.0 interface lo0.0 passive
```

5. Configure Layer 2 circuits to use the egress protection LSP to protect against a link failure to device CE1.

```
[edit]
user@PE1# edit protocols l2circuit
[edit protocols l2circuit]
user@PE1# set neighbor 1.1.1.3 interface fe-3/3/0.1 virtual-circuit-id 32
[edit protocols l2circuit]
user@PE1# edit neighbor 1.1.1.3
[edit protocols l2circuit neighbor 1.1.1.3]
user@PE1# set interface fe-3/3/0.1 egress-protection protector-interface ge-3/1/2.2
[edit protocols l2circuit]
user@PE1# set neighbor 1.1.1.4 interface ge-3/1/2.2 virtual-circuit-id 33
```

6. Configure a load balancing policy.

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

7. Configure the routing options to export routes based on the load balancing policy.

```
[edit]
user@PE1# set routing-options router-id 1.1.1.2
[edit]
user@PE1# set routing-options forwarding-table export load-balance-example
```

8. If you are done configuring the device, commit the configuration.

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 an egress protection LSP, complete the following steps for router PE2:

1. Configure RSVP. Include the interface linked to the ingress PE router and the interface linked to the CE device.

```
[edit]
user@PE2# edit protocols rsvp
[edit protocols rsvp]
user@PE2# set tunnel-services
[edit protocols rsvp]
user@PE2# set interface ge-2/3/0.3
```

```
[edit protocols rsvp]
user@PE2# set interface ge-2/3/1.4 link-protection
```

2. Configure LDP. Include the interface linked to the ingress PE router and the interface linked to the CE device.

```
[edit]
user@PE2# edit protocols ldp
[edit protocols ldp]
user@PE2# set interface ge-2/3/0.3
[edit protocols ldp]
user@PE2# set interface ge-2/3/1.4
[edit protocols ldp]
user@PE2# set interface lo0.0
[edit protocols ldp]
user@PE2# set upstream-label-assignment
```

3. Configure MPLS and the LSP which acts as the egress protection LSP.

```
[edit]
user@PE2# edit protocols mpls
[edit protocols mpls]
user@PE2# set interface ge-2/3/0.3
[edit protocols mpls]
user@PE2# set interface ge-2/3/1.4
[edit protocols mpls]
user@PE2# set label-switched-path protected-lsp to 2.2.3.4
[edit protocols mpls]
user@PE2# set label-switched-path protected-lsp egress-protection
```

4. Configure OSPF.

```
[edit]
user@PE2# edit protocols ospf
[edit protocols ospf]
user@PE2# set interface traffic-engineering
[edit protocols ospf]
user@PE2# set interface area 0.0.0.0 interface ge-2/3/0.3
[edit protocols ospf]
user@PE2# set interface area 0.0.0.0 interface ge-2/3/1.4
[edit protocols ospf]
user@PE2# set interface area 0.0.0.0 interface lo0.0 passive
```

5. Configure the Layer 2 circuit to use the egress protection LSP.

```
[edit]
user@PE2# edit protocols l2circuit
[edit protocols l2circuit]
user@PE2# set neighbor 1.1.1.2 interface ge-2/0/2.5 virtual-circuit-id 23
[edit protocols l2circuit]
user@PE2# edit neighbor 1.1.1.2
[edit protocols l2circuit neighbor 1.1.1.2]
user@PE2# set interface ge-2/0/2.5 egress-protection protector-pe 1.1.1.4
[edit protocols l2circuit neighbor 1.1.1.2]
user@PE2# set interface ge-2/0/2.5 egress-protection protector-pe
context-identifier 2.2.3.4
```

6. Configure a load balancing policy.

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

7. Configure the routing options to export routes based on the load balancing policy.

```
[edit]
user@PE2# set routing-options router-id 1.1.1.3
[edit]
user@PE2# set routing-options forwarding-table export load-balance-example
```

8. If you are done configuring the device, commit the configuration.

Step-by-Step Procedure

To configure an egress protection LSP, complete the following steps for router PE3:

1. Configure RSVP. Include the interface linked to the ingress PE router and the interface linked to the CE device.

```
[edit]
user@PE3# edit protocols rsvp
[edit protocols rsvp]
user@PE3# set tunnel-services
[edit protocols rsvp]
user@PE3# set interface ge-2/0/0.6
[edit protocols rsvp]
user@PE3# set interface ge-2/0/1.7
```

2. Configure LDP. Include the interface linked to the ingress PE router and the interface linked to the CE device.

```
[edit]
user@PE3# edit protocols ldp
[edit protocols ldp]
user@PE3# set interface ge-2/0/0.6
[edit protocols ldp]
user@PE3# set interface ge-2/0/1.7
[edit protocols ldp]
user@PE3# set interface lo0.0
[edit protocols ldp]
user@PE3# set upstream-label-assignment
```

3. Configure MPLS and the LSP which acts as the egress protection LSP.

```
[edit]
user@PE3# edit protocols mpls
[edit protocols mpls]
user@PE3# set interface ge-2/0/0.6
[edit protocols mpls]
user@PE3# set interface ge-2/0/1.7
[edit protocols mpls]
user@PE3# set egress-protection context-identifier 2.2.3.4 protector
```

4. Configure OSPF.

```
[edit]
user@PE3# edit protocols ospf
[edit protocols ospf]
user@PE3# set interface traffic-engineering
[edit protocols ospf]
```

```

user@PE3# set area 0.0.0.0 interface ge-2/0/0.6
[edit protocols ospf]
user@PE3# set area 0.0.0.0 interface ge-2/0/1.7
[edit protocols ospf]
user@PE3# set area 0.0.0.0 interface lo0.0 passive

```

5. Configure the Layer 2 circuit to use the egress protection LSP.

```

[edit]
user@PE3# edit protocols l2circuit
[edit protocols l2circuit]
user@PE3# set neighbor 1.1.1.2 interface ge-2/0/3.21 virtual-circuit-id 42
[edit protocols l2circuit]
user@PE3# edit neighbor 1.1.1.2
[edit protocols l2circuit neighbor 1.1.1.2]
user@PE3# set interface ge-2/0/3.21 egress-protection protected-l2circuit ingress-pe
1.1.1.2
[edit protocols l2circuit neighbor 1.1.1.2]
user@PE3# set interface ge-2/0/3.21 egress-protection protected-l2circuit egress-pe
1.1.1.3
[edit protocols l2circuit neighbor 1.1.1.2]
user@PE3# set interface ge-2/0/3.21 egress-protection
protected-l2circuit virtual-circuit-id 31

```

6. If you are done configuring the device, commit the configuration.

Results

From configuration mode, confirm your configuration on router PE1 by entering the **show protocols**, **show policy-options**, and **show routing-options** commands. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

```

[edit]
user@PE1# show protocols
rsvp {
  interface ge-3/1/0.1;
  interface ge-3/1/3.2;
}
mpls {
  interface ge-3/1/0.1;
  interface ge-3/1/3.2;
}
ospf {
  traffic-engineering;
  area 0.0.0.0 {
    interface ge-3/1/0.1;
    interface ge-3/1/3.2;
    interface lo0.0 {
      passive;
    }
  }
}
ldp {
  interface ge-3/1/0.1;
  interface ge-3/1/3.2;
}

```

```
    interface lo0.0;
  }
  l2circuit {
    neighbor 1.1.1.3 {
      interface fe-3/3/0.1 {
        virtual-circuit-id 32;
        egress-protection {
          protector-interface ge-3/1/2.2;
        }
      }
    }
    neighbor 1.1.1.4 {
      interface ge-3/1/2.2 {
        virtual-circuit-id 33;
      }
    }
  }
}
[edit]
user@PE1# show policy-options
policy-statement load-balance-example {
  then {
    load-balance per-packet;
  }
}
[edit]
user@PE1# show routing-options
router-id 1.1.1.2;
forwarding-table {
  export load-balance-example;
}
```

From configuration mode, confirm your configuration on router PE2 by entering the **show protocols**, **show policy-options**, and **show routing-options** commands. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

```
[edit]
user@PE2# show protocols
rsvp {
  tunnel-services;
  interface ge-2/3/0.3;
  interface ge-2/3/1.4 {
    link-protection;
  }
}
mpls {
  label-switched-path protected-lsp {
    to 2.2.3.4;
    egress-protection;
  }
  interface ge-2/3/0.3;
  interface ge-2/3/1.4;
}
ospf {
  traffic-engineering;
  area 0.0.0.0 {
```

```

        interface ge-2/3/0.3;
        interface ge-2/3/1.4;
        interface lo0.0 {
            passive;
        }
    }
}
ldp {
    interface ge-2/3/0.3;
    interface ge-2/3/1.4;
    interface lo0.0;
    upstream-label-assignment;
}
l2circuit {
    neighbor 1.1.1.2 {
        interface ge-2/0/2.5 {
            virtual-circuit-id 23;
            egress-protection {
                protector-pe 1.1.1.4 context-identifier 2.2.3.4;
            }
        }
    }
}
}

[edit]
user@PE2# show policy-options
policy-options {
    policy-statement load-balance-example {
        then {
            load-balance per-packet;
        }
    }
}

[edit]
user@PE2# show routing-options
routing-options {
    router-id 1.1.1.3;
    forwarding-table {
        export load-balance-example;
    }
}

```

From configuration mode, confirm your configuration on router PE3 by entering the **show protocols** command. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

```

[edit]
user@PE3# show protocols
rsvp {
    tunnel-services;
    interface ge-2/0/0.6;
    interface ge-2/0/1.7;
}
mpls {
    interface ge-2/0/0.6;
    interface ge-2/0/1.7;
}

```

```
    egress-protection {
      context-identifier 2.2.3.4 {
        protector;
      }
    }
  }
}
ospf {
  traffic-engineering;
  area 0.0.0.0 {
    interface ge-2/0/0.6;
    interface ge-2/0/1.7;
    interface lo0.0 {
      passive;
    }
  }
}
ldp {
  interface ge-2/0/0.6;
  interface ge-2/0/1.7;
  interface lo0.0;
  upstream-label-assignment;
}
l2circuit {
  neighbor 1.1.1.2 {
    interface ge-2/0/3.21 {
      virtual-circuit-id 42;
      egress-protection {
        protected-l2circuit PW1 ingress-pe 1.1.1.2 egress-pe 1.1.1.3 virtual-circuit-id 31;
      }
    }
  }
}
}
```

Using the Layer 2 Interworking Interface to Interconnect a Layer 2 Circuit to a Layer 2 VPN

Instead of using a physical Tunnel PIC for looping the packet received from the Layer 2 circuit, the Layer 2 interworking interface uses Junos OS to stitch together both Layer 2 VPN routes.

To configure the interworking interface, include the **iw0** statement. The **iw0** statement is configured at the **[edit interfaces]** hierarchy level. This specifies the peering between two logical interfaces. This configuration is similar to the configuration for a logical tunnel interface. The logical Interfaces must be associated with the endpoints of a Layer 2 circuit and Layer 2 VPN connections.

```
[edit interfaces]
iw0 {
  unit 0 {
    peer-unit 1;
  }
  unit 1 {
    peer-unit 0;
  }
}
```


Configure the Layer 2 circuit protocol by including the **l2circuit** statement at the **[edit protocols]** hierarchy level and specifying the **neighbor** and **iw0** interface.

```
[edit protocols]
l2circuit {
  neighbor 1.2.3.4 {
    interface iw0.0;
  }
}
```

Configure the Layer 2 VPN connection, by including the **routing-instance-name** statement at the **[edit routing-instances]** hierarchy level and specifying the **instance-type l2vpn** option.

```
[edit routing-instances]
routing-instance-name {
  instance-type l2vpn;
  interface iw0.1;
  ...
  protocols {
    l2vpn {
      <l2vpn configuration>;
    }
  }
}
```

In addition to the **iw0** interface configuration, Layer 2 interworking **l2iw** protocols must be enabled. Without the **l2iw** configuration, the **l2iw** routes will not be formed, regardless of whether any **iw** interfaces are present. Within the **l2iw** protocols, only trace options can be configured in the standard fashion. The minimum configuration necessary for the feature to work is shown below:

```
[edit]
protocols {
  l2iw;
}
```

Related Documentation

- [Layer 2 Circuit Overview](#)
- [Layer 2 VPN Overview](#)
- [Example: Interconnecting a Layer 2 Circuit with a Layer 2 VPN on page 81](#)

Example: Interconnecting a Layer 2 Circuit with a Layer 2 VPN

This example provides a step-by-step procedure and commands for configuring and verifying a Layer 2 circuit to a Layer 2 VPN. It contains the following sections:

- [Requirements on page 82](#)
- [Overview and Topology on page 82](#)
- [Configuration on page 83](#)

Requirements

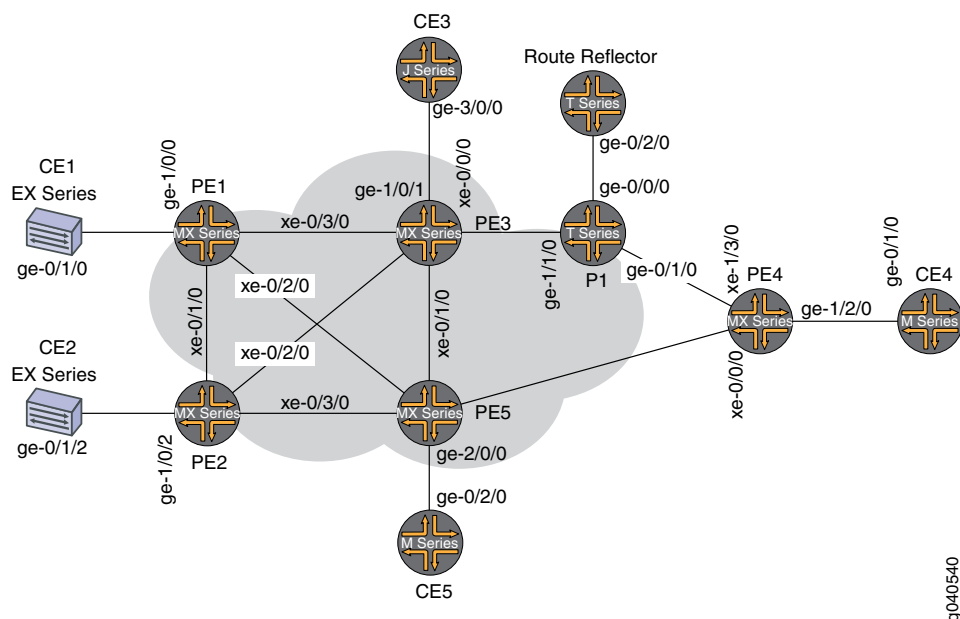
This example uses the following hardware and software components:

- Junos OS Release 9.3 or later
- 2 MX Series 3D Universal Edge Routers
- 2 M Series Multiservice Edge Router
- 1 T Series Core Router
- 1 EX Series Ethernet Switch
- 1 J Series Services Routers

Overview and Topology

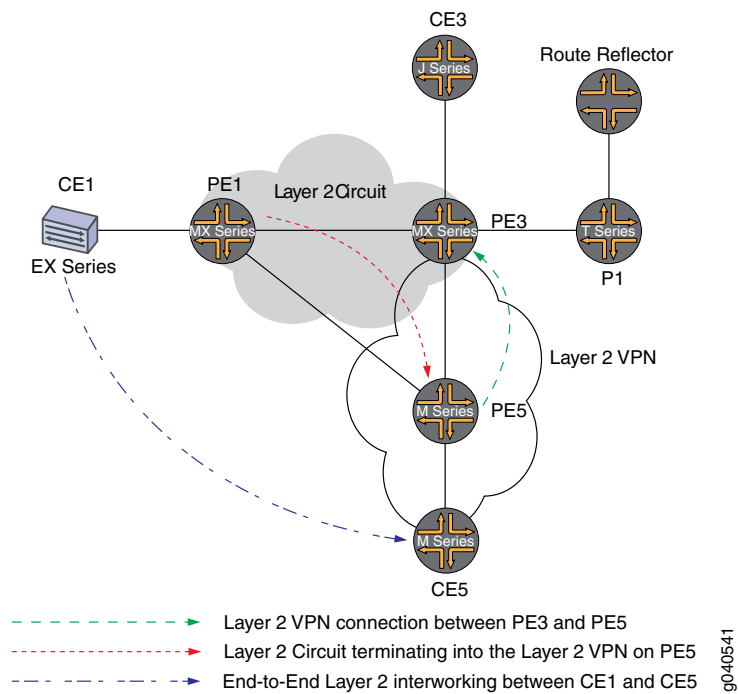
The physical topology of a Layer 2 circuit to a Layer 2 VPN connection is shown in [Figure 9 on page 82](#).

Figure 9: Physical Topology of a Layer 2 Circuit to a Layer 2 VPN Connection



The logical topology of a Layer 2 circuit to a Layer 2 VPN connection is shown in [Figure 10 on page 83](#).

Figure 10: Logical Topology of a Layer 2 Circuit to a Layer 2 VPN Connection



Configuration



NOTE: In any configuration session, it is good practice to verify periodically that the configuration can be committed using the `commit check` command.

In this example, the router being configured is identified using the following command prompts:

- **CE1** identifies the customer edge 1 (CE1) router
- **PE1** identifies the provider edge 1 (PE1) router
- **CE3** identifies the customer edge 3 (CE3) router
- **PE3** identifies the provider edge 3 (PE3) router
- **CE5** identifies the customer edge 5 (CE5) router
- **PE5** identifies the provider edge 5 (PE5) router

This example is organized in the following sections:

- [Configuring Protocols on the PE and P Routers on page 84](#)
- [Verification on page 88](#)

Configuring Protocols on the PE and P Routers

Step-by-Step Procedure In this example, all of the PE routers and P routers are configured with OSPF as the IGP protocol. The MPLS, LDP, and BGP protocols are enabled on all of the interfaces except **fxp0.0**. Core-facing interfaces are enabled with the MPLS address and inet address.

1. Configure all the PE and P routers with OSPF as the IGP. Enable the MPLS, LDP, and BGP protocols on all interfaces except **fxp0.0**. LDP is used as the signaling protocol on Router PE1 for the Layer 2 circuit. The following configuration snippet shows the protocol configuration for Router PE1:

```
[edit]
protocols {
  mpls {
    interface all;
    interface fxp0.0 {
      disable;
    }
  }
  bgp {
    group RR {
      type internal;
      local-address 1.1.1.1;
      family l2vpn {
        signaling;
      }
      neighbor 7.7.7.7;
    }
  }
  ospf {
    traffic-engineering;
    area 0.0.0.0 {
      interface all;
      interface fxp0.0 {
        disable;
      }
    }
  }
  ldp {
    interface all;
    interface fxp0.0 {
      disable;
    }
  }
}
```

2. Configure the PE and P routers with OSPF as the IGP. Enable the MPLS, LDP, and BGP protocols on all interfaces except **fxp0.0**. BGP is used as the signaling protocol on Router PE3 for the Layer 2 VPN. The following configuration snippet shows the protocol configuration for Router PE3:

```
[edit]
protocols {
  mpls {
    interface all;
```

```

interface fxp0.0 {
  disable;
}
}
bgp {
  group RR {
    type internal;
    local-address 3.3.3.3;
    family l2vpn {
      signaling;
    }
    neighbor 7.7.7.7;
  }
}
ospf {
  traffic-engineering;
  area 0.0.0.0 {
    interface all;
    interface fxp0.0 {
      disable;
    }
  }
}
ldp {
  interface all;
  interface fxp0.0 {
    disable;
  }
}
}

```

Step-by-Step Procedure

Configuring Interfaces

1. On Router PE1, configure the **ge-1/0/0** interface encapsulation. To configure the interface encapsulation, include the **encapsulation** statement and specify the **ethernet-ccc** option (vlan-ccc encapsulation is also supported). Configure the **ge-1/0/0.0** logical interface family for circuit cross-connect functionality. To configure the logical interface family, include the **family** statement and specify the **ccc** option. The encapsulation should be configured the same way for all routers in the Layer 2 circuit domain.

```

[edit interfaces]
ge-1/0/0 {
  encapsulation ethernet-ccc;
  unit 0 {
    family ccc;
  }
}
lo0 {
  unit 0 {
    family inet {
      address 1.1.1.1/32;
    }
  }
}
}

```

2. Router PE5 is the router that is *stitching* the Layer 2 circuit to the Layer 2 VPN using the interworking interface. The configuration of the peer unit interfaces is what makes the interconnection.

On Router PE5, configure the **iw0** interface with two logical interfaces. To configure the **iw0** interface, include the **interfaces** statement and specify **iw0** as the interface name. For the unit 0 logical interface, include the **peer-unit** statement and specify the logical interface **unit 1** as the peer interface. For the unit 1 logical interface, include the **peer-unit** statement and specify the logical interface **unit 0** as the peer interface.

```
[edit interfaces]
iw0 {
  unit 0 {
    encapsulation ethernet-ccc;
    peer-unit 1;
  }
  unit 1 {
    encapsulation ethernet-ccc;
    peer-unit 0;
  }
}
```

3. On Router PE5, configure the logical loopback interface. The loopback interface is used to establish the targeted LDP sessions to Routers PE1 and PE5.

```
[edit interfaces]
lo0 {
  unit 0 {
    family inet {
      address 5.5.5.5/32;
    }
  }
}
```

Step-by-Step Procedure

Configuring the Layer 2 circuit protocol

1. On Router PE1, configure the IP address of the remote PE router with the **neighbor** statement. The loopback address and router ID of the PE neighbor is commonly the neighbor's IP address. To allow a Layer 2 circuit to be established even though the maximum transmission unit (MTU) configured on the PE router does not match the MTU configured on the remote PE router, include the **ignore-mtu-mismatch** statement.

```
[edit]
protocols {
  l2circuit {
    neighbor 5.5.5.5 {
      interface ge-1/0/0.0 {
        virtual-circuit-id 100;
        no-control-word;
        ignore-mtu-mismatch;
      }
    }
  }
}
```

- On Router PE5, configure the IP address of the remote PE router. To configure the IP address of the remote PE router, include the **neighbor** statement and specify the IP address of the loopback interface on Router PE1. Configure the virtual circuit ID to be the same as the virtual circuit ID on the neighbor router. To allow a Layer 2 circuit to be established even though the MTU configured on the local PE router does not match the MTU configured on the remote PE router, include the **ignore-mtu-mismatch** statement. Also disable the use of the control word for demultiplexing by including the **no-control-word** statement.

```
[edit protocols]
l2circuit {
  neighbor 1.1.1.1 {
    interface iw0.0 {
      virtual-circuit-id 100;
      no-control-word;
      ignore-mtu-mismatch;
    }
  }
}
```

- On Router PE5, configure the Layer 2 VPN protocols by including the **l2vpn** statement at the **[edit routing-instances routing-instances-name protocols]** hierarchy level. To configure the **iw0** interface, include the **interfaces** statement and specify **iw0** as the interface name. The **iw0** interface is configured under the Layer 2 VPN protocols to receive the looped packet from the **iw0.1** logical interface. The **l2vpn** protocol is configured on Router PE5 with site CE5, which is configured in the BGP L2VPN routing instance. Router CE1 has communication to Router CE5, through the Layer 2 Interworking configuration on Router PE5.

```
[edit]
routing-instances {
  L2VPN {
    instance-type l2vpn;
    interface ge-2/0/0.0;
    interface iw0.1;
    route-distinguisher 65000:5;
    vrf-target target:65000:2;
    protocols {
      l2vpn {
        encapsulation-type ethernet;
        site CE5 {
          site-identifier 5;
          interface ge-2/0/0.0 {
            remote-site-id 3;
          }
        }
        site l2-circuit {
          site-identifier 6;
          interface iw0.1 {
            remote-site-id 3;
          }
        }
      }
    }
  }
}
```

```
}
```

4. In addition to the **iw0** interface configuration, the Layer 2 interworking **l2iw** protocol must be configured. Without the **l2iw** protocol configuration, the Layer 2 interworking routes are not formed, regardless of whether any **iw** interfaces are present.

On Router PE5, configure the **l2iw** protocol. To configure the protocol, include the **l2iw** statement at the **[edit protocols]** hierarchy level.

```
[edit]
protocols {
  l2iw;
}
```

Verification

Step-by-Step Procedure

Verifying the Layer 2 Circuit Connection on Router PE1.

1. On Router PE1, use the **show l2circuit connections** command to verify that the Layer 2 Circuit from Router PE1 to Router PE5 is **Up**.

```
user@PE1> show l2circuit connections
Layer-2 Circuit Connections:
Legend for connection status (St)
EI -- encapsulation invalid      NP -- interface h/w not present
MM -- mtu mismatch              Dn -- down
EM -- encapsulation mismatch    VC-Dn -- Virtual circuit Down
CM -- control-word mismatch     Up -- operational
VM -- vlan id mismatch         CF -- Call admission control failure
OL -- no outgoing label        IB -- TDM incompatible bitrate
NC -- intf encaps not CCC/TCC  TM -- TDM misconfiguration
BK -- Backup Connection        ST -- Standby Connection
CB -- rcvd cell-bundle size bad XX -- unknown
SP -- Static Pseudowire

Legend for interface status
Up -- operational
Dn -- down
Neighbor: 5.5.5.5
Interface              Type St   Time last up   # Up trans
ge-1/0/0.0(vc 100)    rmt  Up    Jan 3 22:00:49 2010    1
Remote PE: 5.5.5.5, Negotiated control-word: No
Incoming label: 301328, Outgoing label: 300192
Local interface: ge-1/0/0.0, Status: Up, Encapsulation: ETHERNET
```

2. On Router PE5, use the **show l2vpn connections** command to verify that the Layer 2 VPN connection is **Up** using the **iw0** peer interface of the Layer 2 circuit.

```
user@PE5> show l2vpn connections
Instance: L2VPN
Local site: CE5 (5)
connection-site        Type St   Time last up   # Up trans
l2-circuit (6)         loc  OR
3                      rmt  Up
Jan 3 22:51:12 2010    1
Remote PE: 3.3.3.3, Negotiated control-word: Yes (Null)
Incoming label: 800258, Outgoing label: 800000
Local interface: ge-2/0/0.0, Status: Up, Encapsulation: ETHERNET
Local site: l2-circuit (6)
connection-site        Type St   Time last up   # Up trans
```



```

CE5 (5)                               loc  OR
3                                     rmt  Up   Jan 3 22:56:38 2010   1
Remote PE: 3.3.3.3, Negotiated control-word: Yes (Null)
Incoming label: 800262, Outgoing label: 800001
Local interface: iw0.1, Status: Up, Encapsulation: ETHERNET

```

Step-by-Step Procedure Verifying that the Layer 2 Circuit is terminating into the Layer 2 VPN connection.

1. On Router PE5, use the **show l2circuit connections** command to verify that the Layer 2 circuit is **Up** using the **iw0** interface. This will be looped through the **iw0.1** interface to the Layer 2 VPN.

```
user@PE5> show l2circuit connections
```

```
Layer-2 Circuit Connections:
```

```
Neighbor: 1.1.1.1
```

```

Interface          Type St   Time last up # Up trans
iw0.0(vc 100)    rmt  Up Jan 3 21:59:07 2010 1
Remote PE: 1.1.1.1, Negotiated control-word: No
Incoming label: 300192, Outgoing label: 301328

```

2. On Router PE 5, use the **show route table mpls.0** command to verify the Layer 2 circuit and Layer 2 VPN routes. In the example below, the Layer 2 circuit is associated with LDP label **301328** and the Layer 2 VPN is associated with LDP label **800001**. Notice the two **iw0** interfaces that are used for the Layer 2 interworking route.

```
user@PE5>show route table mpls.0
```

```

mpls.0: 18 destinations, 20 routes (18 active, 2 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

```

```

0                *[MPLS/0] 5d 20:07:31, metric 1
                  Receive
1                *[MPLS/0] 5d 20:07:31, metric 1
                  Receive
2                *[MPLS/0] 5d 20:07:31, metric 1
                  Receive
299776           *[LDP/9] 2d 03:00:51, metric 1
300048           *[LDP/9] 2d 03:00:49, metric 1
                  > to 10.10.6.1 via xe-0/1/0.0, Pop
300048(S=0)      *[LDP/9] 2d 03:00:49, metric 1
                  > to 10.10.6.1 via xe-0/1/0.0, Pop
300192           *[L2IW/6] 19:11:05, metric2 1
                  > to 10.10.6.1 via xe-0/1/0.0, Swap 800001
                  [L2CKT/7] 20:08:36
                  > via iw0.0, Pop
800258           *[L2VPN/7] 19:16:31
                  > via ge-2/0/0.0, Pop          Offset: 4
800262          *[L2IW/6] 19:11:05, metric2 1          > to 10.10.3.1 via xe-1/1/0.0, Swap
301328          [L2VPN/7] 19:11:05          > via iw0.1, Pop    Offset: 4ge-2/0/0.0
                  *[L2VPN/7] 19:16:31, metric2 1
                  > to 10.10.6.1 via xe-0/1/0.0, Push 800000 Offset: -4
iw0.0           *[L2CKT/7] 20:08:36, metric2 1
                  > to 10.10.3.1 via xe-1/1/0.0, Push 301328
iw0.1           *[L2VPN/7] 19:11:05, metric2 1
                  > to 10.10.6.1 via xe-0/1/0.0, Push 800001 Offset: -4

```

**Related
Documentation**

- [Layer 2 Circuit Overview](#)
- [Layer 2 VPN Overview](#)
- [Layer 2 VPN Applications](#)
- [Using the Layer 2 Interworking Interface to Interconnect a Layer 2 Circuit to a Layer 2 VPN on page 80](#)

Applications for Interconnecting a Layer 2 Circuit with a Layer 2 Circuit

MPLS-based Layer 2 services are growing in demand among enterprise and service providers. This creates new challenges for service providers who want to provide end-to-end value-added services. There are various reasons to stitch different Layer 2 services to one another and to Layer 3 services, for example, to expand the service offerings and to expand geographically. The Junos OS has various features to address the needs of the service provider.

Interconnecting a Layer 2 circuit with a Layer 2 circuit includes the following benefits:

- Interconnecting a Layer 2 circuit with a Layer 2 circuit enables the sharing of a service provider's core network infrastructure between Layer 2 circuit services, reducing the cost of providing those services. A Layer 2 MPLS circuit allows service providers to create a Layer 2 circuit service over an existing IP and MPLS backbone.
- Service providers do not have to invest in separate Layer 2 equipment to provide Layer 2 circuit service. A service provider can configure a provider edge router to run any Layer 2 protocol. Customers who prefer to maintain control over most of the administration of their own networks want Layer 2 circuit connections with their service provider instead of a Layer 3 VPN connection.

**Related
Documentation**

- [Layer 2 Circuit Overview](#)
- [Example: Interconnecting a Layer 2 Circuit with a Layer 2 Circuit on page 90](#)

Example: Interconnecting a Layer 2 Circuit with a Layer 2 Circuit

This example provides a step-by-step procedure and commands for configuring and verifying a Layer 2 circuit to a Layer 2 circuit interconnection. It contains the following sections:

- [Requirements on page 91](#)
- [Overview and Topology on page 91](#)
- [Configuration on page 92](#)

Requirements

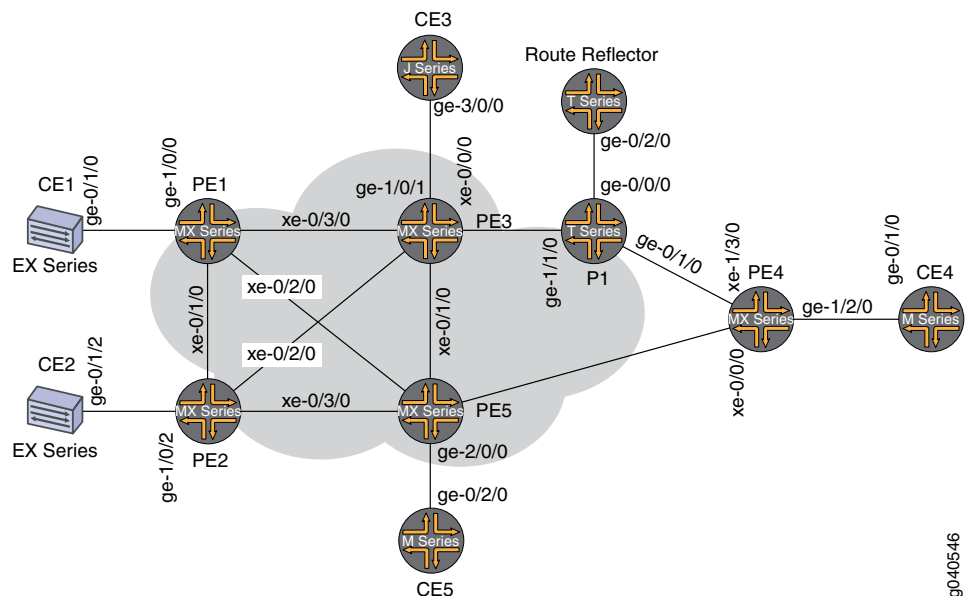
This example uses the following hardware and software components:

- Junos OS Release 9.3 or later
- 2 MX Series routers
- 2 M Series routers
- 1 T Series router
- 1 EX Series router
- 1 J Series router

Overview and Topology


The physical topology of a Layer 2 circuit to Layer 2 circuit interconnection is shown in [Figure 11 on page 91](#)

Figure 11: Physical Topology of a Layer 2 Circuit Terminating into a Layer 2 Circuit



The logical topology of a Layer 2 circuit to Layer 2 circuit interconnection is shown in [Figure 12 on page 92](#)

Figure 12: Logical Topology of a Layer 2 Circuit Terminating into a Layer 2 Circuit



Configuration



NOTE: In any configuration session, it is good practice to verify periodically that the configuration can be committed using the `commit check` command.

In this example, the router being configured is identified using the following command prompts:

- **CE2** identifies the customer edge 2 (CE2) router
- **PE1** identifies the provider edge 1 (PE1) router
- **CE3** identifies the customer edge 3 (CE3) router
- **PE3** identifies the provider edge 3 (PE3) router
- **CE5** identifies the customer edge 5 (CE5) router
- **PE5** identifies the provider edge 5 (PE5) router

This example contains the following procedures:

- [Configuring PE Router Customer-facing and Loopback Interfaces on page 93](#)
- [Configuring Core-facing Interfaces on page 94](#)
- [Configuring Protocols on page 95](#)
- [Configuring the Layer 2 Circuits on page 96](#)
- [Interconnecting the Layer 2 Circuits on page 98](#)

- [Verifying the Layer 2 Circuit to Layer 2 Circuit Interconnection on page 99](#)
- [Results on page 102](#)

Configuring PE Router Customer-facing and Loopback Interfaces

Step-by-Step Procedure

To begin building the interconnection, configure the interfaces on the PE routers. If your network contains provider (P) routers, configure the interfaces on the P routers also. This example shows the configuration for Router PE1 and Router PE5.

1. On Router PE1, configure the **ge-1/0/0** interface encapsulation. To configure the interface encapsulation, include the **encapsulation** statement and specify the **ethernet-ccc** option (vlan-ccc encapsulation is also supported). Configure the **ge-1/0/0.0** logical interface family for circuit cross-connect functionality. To configure the logical interface family, include the **family** statement and specify the **ccc** option.

```
[edit interfaces]
ge-1/0/0 {
  encapsulation ethernet-ccc;
  unit 0 {
    family ccc;
  }
}
lo0 {
  unit 0 {
    family inet {
      address 1.1.1.1/32;
    }
  }
}
```

2. On Router PE5, configure the **ge-2/0/0** interface encapsulation. To configure the interface encapsulation, include the **encapsulation** statement and specify the **ethernet-ccc** option. Configure the **ge-2/0/0.0** logical interface family for circuit cross-connect functionality. To configure the logical interface family, include the **family** statement and specify the **ccc** option

```
[edit interfaces]
ge-2/0/0 {
  encapsulation ethernet-ccc;
  unit 0 {
    family ccc;
  }
}
lo0 {
  unit 0 {
    family inet {
      address 5.5.5.5/32;
    }
  }
}
```

3. On Router PE3, configure the logical loopback interface. The loopback interface is used to establish the targeted LDP sessions to Routers PE1 and PE5.

```
[edit interfaces]
```

```
lo0 {
  unit 0 {
    family inet {
      address 3.3.3.3/32;
    }
  }
}
```

Configuring Core-facing Interfaces

Step-by-Step Procedure

This procedure describes how to configure the core-facing interfaces on the PE routers. This example does not include all the core-facing interfaces shown in the physical topology illustration. Enable the **mpls** and **inet** address families on the core-facing interfaces.

1. On Router PE1, configure the **xe-0/3/0** interface. Include the **family** statement and specify the **inet** address family. Include the **address** statement and specify **10.10.1.1/30** as the interface address. Include the **family** statement and specify the **mpls** address family.

```
[edit interfaces]
xe-0/3/0 {
  unit 0 {
    family inet {
      address 10.10.1.1/30;
    }
    family mpls;
  }
}
```

2. On Router PE3, configure the core-facing interfaces. Include the **family** statement and specify the **inet** address family. Include the **address** statement and specify the IPv4 addresses shown in the example as the interface addresses. Include the **family** statement and specify the **mpls** address family. In the example, the **xe-0/0/0** interface is connected to the route reflector, the **xe-0/1/0** interface is connected to Router PE5, the **xe-0/2/0** interface is connected to Router PE2, and the **xe-0/3/0** interface is connected to Router PE1.

```
[edit interfaces]
xe-0/0/0 {
  unit 0 {
    family inet {
      address 10.10.20.2/30;
    }
    family mpls;
  }
}
xe-0/1/0 {
  unit 0 {
    family inet {
      address 10.10.6.1/30;
    }
    family mpls;
  }
}
```

```

}
xe-0/2/0 {
  unit 0 {
    family inet {
      address 10.10.5.2/30;
    }
    family mpls;
  }
}
xe-0/3/0 {
  unit 0 {
    family inet {
      address 10.10.1.2/30;
    }
    family mpls;
  }
}

```

3. On Router PE5, configure the **xe-0/1/0** interface. Include the **family** statement and specify the **inet** address family. Include the **address** statement and specify **10.10.6.2/30** as the interface address. Include the **family** statement and specify the **mpls** address family.

```

[edit interfaces]
xe-0/1/0 {
  unit 0 {
    family inet {
      address 10.10.6.2/30;
    }
    family mpls;
  }
}

```

Configuring Protocols

Step-by-Step Procedure

This procedure describes how to configure the protocols used in this example. If your network contains P routers, configure the protocols on the P routers also.

Configure all of the PE routers and P routers with OSPF as the IGP protocol. Enable MPLS and LDP protocols on all of the interfaces except **fxp0.0**.

1. On Router PE1, enable OSPF as the IGP. Enable the MPLS and LDP protocols on all interfaces except **fxp0.0**. LDP is used as the signaling protocol on Router PE1 for the Layer 2 circuit. The following configuration snippet shows the protocol configuration for Router PE1:

```

[edit]
protocols {
  mpls {
    interface all;
    interface fxp0.0 {
      disable;
    }
  }
  ospf {
    traffic-engineering;
  }
}

```

```

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

```

2. Configure the PE and P routers with OSPF as the IGP. Enable the MPLS and LDP protocols on all interfaces except **fxp0.0**. The following configuration snippet shows the protocol configuration for Router PE3:

```

[edit]
protocols {
    mpls {
        interface all;
        interface fxp0.0 {
            disable;
        }
    }
    ospf {
        traffic-engineering;
        area 0.0.0.0 {
            interface all;
            interface fxp0.0 {
                disable;
            }
        }
    }
    ldp {
        interface all;
        interface fxp0.0 {
            disable;
        }
    }
}

```

Configuring the Layer 2 Circuits

Step-by-Step Procedure

This procedure describes how to configure the Layer 2 circuits.



NOTE: In this example the **ignore-mtu-mismatch** statement is required for the circuit to come up.

1. On Router PE1, configure the Layer 2 circuit. Include the **l2circuit** statement. Include the **neighbor** statement and specify the loopback IPv4 address of Router PE3 as

the neighbor. Include the interface statement and specify **ge-1/0/0.0** as the logical interface that is participating in the Layer 2 circuit. Include the **virtual-circuit-id** statement and specify **100** as the identifier. Include the **ignore-mtu-mismatch** statement to allow a Layer 2 circuit to be established even though the maximum transmission unit (MTU) configured on the local PE router does not match the MTU configured on the remote PE router.

```
[edit]
protocols {
  l2circuit {
    neighbor 3.3.3.3 {
      interface ge-1/0/0.0 {
        virtual-circuit-id 100;
        ignore-mtu-mismatch;
      }
    }
  }
}
```

2. On Router PE5, configure the Layer 2 circuit. Include the **l2circuit** statement. Include the **neighbor** statement and specify the loopback IPv4 address of Router PE3 as the neighbor. Include the interface statement and specify **ge-2/0/0.0** as the logical interface that is participating in the Layer 2 circuit. Include the **virtual-circuit-id** statement and specify **200** as the identifier. Include the **ignore-mtu-mismatch** statement to allow a Layer 2 circuit to be established even though the MTU configured on the local PE router does not match the MTU configured on the remote PE router.

```
[edit]
protocols {
  l2circuit {
    neighbor 3.3.3.3 {
      interface ge-2/0/0.0 {
        virtual-circuit-id 200;
        ignore-mtu-mismatch;
      }
    }
  }
}
```

3. On Router PE3, configure the Layer 2 circuit to Router PE1. Include the **l2circuit** statement. Include the **neighbor** statement and specify the loopback IPv4 address of Router PE1 as the neighbor. Include the interface statement and specify **iw0.0** as the logical interworking interface that is participating in the Layer 2 circuit. Include the **virtual-circuit-id** statement and specify **100** as the identifier. Include the **ignore-mtu-mismatch** statement to allow a Layer 2 circuit to be established even though the MTU configured on the local PE router does not match the MTU configured on the remote PE router.

On Router PE3, configure the Layer 2 circuit to Router PE5. Include the **l2circuit** statement. Include the **neighbor** statement and specify the loopback IPv4 address of Router PE5 as the neighbor. Include the interface statement and specify **iw0.1** as the logical interworking interface that is participating in the Layer 2 circuit. Include

the **virtual-circuit-id** statement and specify **200** as the identifier. Include the **ignore-mtu-mismatch** statement.

```
[edit protocols]
l2circuit {
  neighbor 1.1.1.1 {
    interface iw0.0 {
      virtual-circuit-id 100;
      ignore-mtu-mismatch;
    }
  }
  neighbor 5.5.5.5 {
    interface iw0.1 {
      virtual-circuit-id 200;
      ignore-mtu-mismatch;
    }
  }
}
```

Interconnecting the Layer 2 Circuits

Step-by-Step Procedure Router PE3 is the router that is *stitching* the Layer 2 circuits together using the interworking interface. The configuration of the peer unit interfaces is what makes the interconnection.

1. On Router PE3, configure the **iw0.0** interface. Include the **encapsulation** statement and specify the **ethernet-ccc** option. Include the **peer-unit** statement and specify the logical interface unit 1 as the peer tunnel interface.

On Router PE3, configure the **iw0.1** interface. Include the **encapsulation** statement and specify the **ethernet-ccc** option. Include the **peer-unit** statement and specify the logical interface unit 0 as the peer tunnel interface.

```
[edit interfaces]
iw0 {
  unit 0 {
    encapsulation ethernet-ccc;
    peer-unit 1;
  }
  unit 1 {
    encapsulation ethernet-ccc;
    peer-unit 0;
  }
}
```

2. On Router PE3, configure the Layer 2 interworking **l2iw** protocol. To configure the Layer 2 interworking protocol, include the **l2iw** statement at the **[edit protocols]** hierarchy level.

```
[edit]
protocols {
  l2iw;
}
```

3. On each router, commit the configuration.

```
user@host> commit check
```

```
configuration check succeeds
user@host> commit
```

Verifying the Layer 2 Circuit to Layer 2 Circuit Interconnection

Step-by-Step Procedure

Verify that the Layer 2 circuit connection on Router PE1 is up, the LDP neighbors are correct, and the MPLS label operations are correct.

1. On Router PE1, use the **show l2circuit connections** command to verify that the Layer 2 circuit from Router PE1 to Router PE3 is **Up**.

```
user@PE1> show l2circuit connections
Layer-2 Circuit Connections:
Legend for connection status (St)
EI -- encapsulation invalid      NP -- interface h/w not present
MM -- mtu mismatch              Dn -- down
EM -- encapsulation mismatch     VC-Dn -- Virtual circuit Down
CM -- control-word mismatch     Up -- operational
VM -- vlan id mismatch          CF -- Call admission control failure
OL -- no outgoing label         IB -- TDM incompatible bitrate
NC -- intf encaps not CCC/TCC   TM -- TDM misconfiguration
BK -- Backup Connection         ST -- Standby Connection
CB -- rcvd cell-bundle size bad XX -- unknown
SP -- Static Pseudowire
```

Legend for interface status

Up -- operational

Dn -- down

Neighbor: 3.3.3.3

```
Interface          Type St   Time last up   # Up trans
ge-1/0/0.0(vc 100) rmt  Up    Jan 5 22:00:49 2010    1
Remote PE: 3.3.3.3, Negotiated control-word: Yes (Null)
Incoming label: 301328, Outgoing label: 314736
Local interface: ge-1/0/0.0, Status: Up, Encapsulation: ETHERNET
```

2. On Router PE1, use the **show ldp neighbor** command to verify that the IPv4 address of Router PE3 is shown as the LDP neighbor.

```
user@PE1> show ldp neighbor
```

```
Address          Interface          Label space ID      Hold time
3.3.3.3          lo0.0              3.3.3.3:0           41
```

3. On Router PE 1, use the **show route table mpls.0** command to verify the Layer 2 circuit is using the LDP label to Router PE3 in both directions (Push and Pop). In the example below, the Layer 2 circuit is associated with LDP label **301328**.

```
user@PE1> show route table mpls.0
```

```
mpls.0: 13 destinations, 13 routes (13 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both
```

```
0          *[MPLS/0] 1w1d 08:25:39, metric 1
            Receive
1          *[MPLS/0] 1w1d 08:25:39, metric 1
            Receive
2          *[MPLS/0] 1w1d 08:25:39, metric 1
            Receive
300432     *[LDP/9] 3d 01:13:57, metric 1
            > to 10.10.2.2 via xe-0/1/0.0, Pop
```

```

300432(S=0)      *[LDP/9] 3d 01:13:57, metric 1
                  > to 10.10.2.2 via xe-0/1/0.0, Pop
300768           *[LDP/9] 3d 01:13:57, metric 1
                  > to 10.10.3.2 via xe-0/2/0.0, Pop
300768(S=0)      *[LDP/9] 3d 01:13:57, metric 1
                  > to 10.10.3.2 via xe-0/2/0.0, Pop
300912           *[LDP/9] 3d 01:13:57, metric 1
                  > to 10.10.3.2 via xe-0/2/0.0, Swap 299856
301264           *[LDP/9] 3d 01:13:53, metric 1
                  > to 10.10.1.2 via xe-0/3/0.0, Swap 308224
301312           *[LDP/9] 3d 01:13:56, metric 1
                  > to 10.10.1.2 via xe-0/3/0.0, Pop
301312(S=0)      *[LDP/9] 3d 01:13:56, metric 1
                  > to 10.10.1.2 via xe-0/3/0.0, Pop
301328           *[L2CKT/7] 02:33:26          > via ge-1/0/0.0, Pop   Offset: 4
ge-1/0/0.0       *[L2CKT/7] 02:33:26, metric2 1          > to 10.10.1.2 via xe-0/3/0.0,
Push 314736 Offset: -4

```

- On Router PE3, use the **show l2circuit connections** command to verify that the Layer 2 circuit from Router PE3 to Router PE5 is **Up**, that the Layer 2 circuit from Router PE3 to Router PE1 is **Up**, that the connections to Router PE1 and Router PE5 use the iw0 interface, and that the status for both local iw0 interfaces is **Up**.

```

user@PE3> show l2circuit connections
Layer-2 Circuit Connections:
Legend for connection status (St)
EI -- encapsulation invalid      NP -- interface h/w not present
MM -- mtu mismatch              Dn -- down
EM -- encapsulation mismatch     VC-Dn -- Virtual circuit Down
CM -- control-word mismatch      Up -- operational
VM -- vlan id mismatch          CF -- Call admission control failure
OL -- no outgoing label         IB -- TDM incompatible bitrate
NC -- intf encaps not CCC/TCC    TM -- TDM misconfiguration
BK -- Backup Connection          ST -- Standby Connection
CB -- rcvd cell-bundle size bad  XX -- unknown
SP -- Static Pseudowire

Legend for interface status
Up -- operational
Dn -- down
Neighbor: 1.1.1.1
  Interface          Type  St    Time last up      # Up trans
  iw0.0(vc 100)      rmt   Up    Jan  5 13:50:14 2010      1

  Remote PE: 1.1.1.1, Negotiated control-word: Yes (Null)
  Incoming label: 314736, Outgoing label: 301328
  Local interface: iw0.0, Status: Up, Encapsulation: ETHERNET
Neighbor: 5.5.5.5
  Interface          Type  St    Time last up      # Up trans
  iw0.1(vc 200)      rmt   Up    Jan  5 13:49:58 2010      1

  Remote PE: 5.5.5.5, Negotiated control-word: Yes (Null)
  Incoming label: 314752, Outgoing label: 300208
  Local interface: iw0.1, Status: Up, Encapsulation: ETHERNET

```

- On Router PE3, use the **show ldp neighbor** command to verify that the correct IPv4 addresses are shown as the LDP neighbor.

```
user@PE3> show ldp neighbor
```

Address	Interface	Label space ID	Hold time
1.1.1.1	lo0.0	1.1.1.1:0	44
2.2.2.2	lo0.0	2.2.2.2:0	42
4.4.4.4	lo0.0	4.4.4.4:0	31
5.5.5.5	lo0.0	5.5.5.5:0	44

6. On Router PE3, use the **show route table mpls.0** command to verify that the **mpls.0** routing table is populated with the Layer 2 interworking routes. Notice that in this example, the router is swapping label **314736** received from Router PE1 on the **iw0.0** to label **301328**.

```

user@PE3> show route table mpls.0
mpls.0: 16 destinations, 18 routes (16 active, 2 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

0          *[MPLS/0] 1w1d 08:28:24, metric 1
            Receive
1          *[MPLS/0] 1w1d 08:28:24, metric 1
            Receive
2          *[MPLS/0] 1w1d 08:28:24, metric 1
            Receive
308160     *[LDP/9] 3d 01:16:55, metric 1
            > to 10.10.1.1 via xe-0/3/0.0, Pop
308160(S=0) *[LDP/9] 3d 01:16:55, metric 1
            > to 10.10.1.1 via xe-0/3/0.0, Pop
308176     *[LDP/9] 3d 01:16:54, metric 1
            > to 10.10.6.2 via xe-0/1/0.0, Pop
308176(S=0) *[LDP/9] 3d 01:16:54, metric 1
            > to 10.10.6.2 via xe-0/1/0.0, Pop
308192     *[LDP/9] 00:21:40, metric 1
            > to 10.10.20.1 via xe-0/0/0.0, Swap 601649
            to 10.10.6.2 via xe-0/1/0.0, Swap 299856
308208     *[LDP/9] 3d 01:16:54, metric 1
            > to 10.10.5.1 via xe-0/2/0.0, Pop
308208(S=0) *[LDP/9] 3d 01:16:54, metric 1
            > to 10.10.5.1 via xe-0/2/0.0, Pop
308224     *[LDP/9] 3d 01:16:52, metric 1
            > to 10.10.20.1 via xe-0/0/0.0, Pop
308224(S=0) *[LDP/9] 3d 01:16:52, metric 1
            > to 10.10.20.1 via xe-0/0/0.0, Pop
314736     *[L2IW/6] 02:35:31, metric2 1
            > to 10.10.6.2 via xe-0/1/0.0, Swap 300208
            [L2CKT/7] 02:35:31
            > via iw0.0, Pop          Offset: 4
314752     *[L2IW/6] 02:35:31, metric2 1
            > to 10.10.1.1 via xe-0/3/0.0, Swap 301328
            [L2CKT/7] 02:35:47
            > via iw0.1, Pop          Offset: 4
iw0.0      *[L2CKT/7] 02:35:31, metric2 1
            > to 10.10.1.1 via xe-0/3/0.0, Push 301328 Offset: -4
iw0.1      *[L2CKT/7] 02:35:47, metric2 1
            > to 10.10.6.2 via xe-0/1/0.0, Push 300208 Offset: -4

```

7. Verify that Router CE1 can send traffic to and receive traffic from Router CE5 across the interconnection, using the **ping** command.

```

user@CE1> ping 40.40.40.11
PING 40.40.40.11 (40.40.40.11): 56 data bytes
64 bytes from 40.40.40.11: icmp_seq=1 ttl=64 time=22.425 ms
64 bytes from 40.40.40.11: icmp_seq=2 ttl=64 time=1.299 ms

```

```
64 bytes from 40.40.40.11: icmp_seq=3 ttl=64 time=1.032 ms
64 bytes from 40.40.40.11: icmp_seq=4 ttl=64 time=1.029 ms
```

8. Verify that Router CE5 can send traffic to and receive traffic from Router CE1 across the interconnection, using the **ping** command.

```
user@CE5>ping 40.40.40.1
PING 40.40.40.1 (40.40.40.1): 56 data bytes
64 bytes from 40.40.40.1: icmp_seq=0 ttl=64 time=1.077 ms
64 bytes from 40.40.40.1: icmp_seq=1 ttl=64 time=0.957 ms
64 bytes from 40.40.40.1: icmp_seq=2 ttl=64 time=1.057 ms 1.017 ms
```

Results

The configuration and verification of this example has been completed. The following section is for your reference.

The relevant sample configuration for Router PE1 follows.

```
Router PE1 [edit]
interfaces {
  xe-0/1/0 {
    unit 0 {
      family inet {
        address 10.10.2.1/30;
      }
      family mpls;
    }
  }
  xe-0/2/0 {
    unit 0 {
      family inet {
        address 10.10.3.1/30;
      }
      family mpls;
    }
  }
  xe-0/3/0 {
    unit 0 {
      family inet {
        address 10.10.1.1/30;
      }
      family mpls;
    }
  }
  ge-1/0/0 {
    encapsulation ethernet-ccc;
    unit 0 {
      family ccc;
    }
  }
  lo0 {
    unit 0 {
      family inet {
        address 1.1.1.1/32;
      }
    }
  }
}
```

```

    }
  }
  forwarding-options {
    hash-key {
      family inet {
        layer-3;
        layer-4;
      }
      family mpls {
        label-1;
        label-2;
      }
    }
  }
}
routing-options {
  static {
    route 172.0.0.0/8 next-hop 172.19.59.1;
  }
  autonomous-system 65000;
}
protocols {
  mpls {
    interface all;
    interface fxp0.0 {
      disable;
    }
  }
  ospf {
    traffic-engineering;
    area 0.0.0.0 {
      interface all;
      interface fxp0.0 {
        disable;
      }
    }
  }
  ldp {
    interface all;
    interface fxp0.0 {
      disable;
    }
  }
  l2circuit {
    neighbor 3.3.3.3 {
      interface ge-1/0/0.0 {
        virtual-circuit-id 100;
        ignore-mtu-mismatch;
      }
    }
  }
}
}

```

The relevant sample configuration for Router PE3 follows.

Router PE3 [edit]
 interfaces {

```
xe-0/0/0 {
  unit 0 {
    family inet {
      address 10.10.20.2/30;
    }
    family mpls;
  }
}
xe-0/1/0 {
  unit 0 {
    family inet {
      address 10.10.6.1/30;
    }
    family mpls;
  }
}
xe-0/2/0 {
  unit 0 {
    family inet {
      address 10.10.5.2/30;
    }
    family mpls;
  }
}
xe-0/3/0 {
  unit 0 {
    family inet {
      address 10.10.1.2/30;
    }
    family mpls;
  }
}
ge-1/0/1 {
  encapsulation ethernet-ccc;
  unit 0 {
    family ccc;
  }
}
iw0 {
  unit 0 {
    encapsulation ethernet-ccc;
    peer-unit 1;
  }
  unit 1 {
    encapsulation ethernet-ccc;
    peer-unit 0;
  }
}
lo0 {
  unit 0 {
    family inet {
      address 3.3.3.3/32;
    }
  }
}
}
```



```

routing-options {
  static {
    route 172.0.0.0/8 next-hop 172.19.59.1;
  }
  autonomous-system 65000;
}
protocols {
  l2iw;
  mpls {
    interface all;
    interface fxp0.0 {
      disable;
    }
  }
  ospf {
    area 0.0.0.0 {
      interface all;
      interface fxp0.0 {
        disable;
      }
    }
  }
  ldp {
    interface all;
    interface fxp0.0 {
      disable;
    }
  }
  l2circuit {
    neighbor 1.1.1.1 {
      interface iw0.0 {
        virtual-circuit-id 100;
        ignore-mtu-mismatch;
      }
    }
    neighbor 5.5.5.5 {
      interface iw0.1 {
        virtual-circuit-id 200;
        ignore-mtu-mismatch;
      }
    }
  }
}

```

**Related
Documentation**

- [Layer 2 Circuit Overview](#)
- [Applications for Interconnecting a Layer 2 Circuit with a Layer 2 Circuit on page 90](#)

Applications for Interconnecting a Layer 2 Circuit with a Layer 3 VPN

MPLS-based Layer 2 services are growing in demand among enterprise and service providers. This creates new challenges related to interoperability between Layer 2 and Layer 3 services for service providers who want to provide end-to-end value-added services. There are various reasons to stitch different Layer 2 services to one another and to Layer 3 services. For example, to expand the service offerings and to expand geographically. The Junos OS has various features to address the needs of the service provider.

Interconnecting a Layer 2 Circuit with a Layer 3 VPN provides the following benefits:

- Interconnecting a Layer 2 Circuit with a Layer 3 VPN enables the sharing of a service provider's core network infrastructure between IP and Layer 2 circuit services, reducing the cost of providing those services. A Layer 2 MPLS circuit allows service providers to create a Layer 2 circuit service over an existing IP and MPLS backbone.
- Service providers do not have to invest in separate Layer 2 equipment to provide Layer 2 circuit service. A service provider can configure a provider edge router to run any Layer 3 protocol in addition to the Layer 2 protocols. Customers who prefer to maintain control over most of the administration of their own networks want Layer 2 circuit connections with their service provider instead of a Layer 3 VPN connection.

Related Documentation

- [Layer 2 Circuit Overview](#)
- [Layer 3 VPN Overview](#)
- [Example: Interconnecting a Layer 2 Circuit with a Layer 3 VPN on page 106](#)

Example: Interconnecting a Layer 2 Circuit with a Layer 3 VPN

This example provides a step-by-step procedure and commands for configuring and verifying a Layer 2 circuit to Layer 3 VPN interconnection. It contains the following sections:

- [Requirements on page 106](#)
- [Overview and Topology on page 107](#)
- [Configuration on page 108](#)
- [Verifying the Layer 2 Circuit to Layer 3 VPN Interconnection on page 118](#)

Requirements

This example uses the following hardware and software components:


- Junos OS Release 9.3 or later
- 3 MX Series 3D Universal Edge Routers
- 1 M Series Multiservice Edge Router
- 1 T Series Core Router

- 1 EX Series Ethernet Switch
- 1 J Series Services Router

Overview and Topology


The physical topology of a Layer 2 circuit to Layer 3 VPN interconnection is shown in [Figure 13 on page 107](#).

Figure 13: Physical Topology of a Layer 2 Circuit to Layer 3 VPN Interconnection



The logical topology of a Layer 2 circuit to Layer 3 VPN interconnection is shown in [Figure 14 on page 108](#).

Figure 14: Logical Topology of a Layer 2 Circuit to Layer 3 VPN Interconnection



Configuration



NOTE: In any configuration session, it is good practice to verify periodically that the configuration can be committed using the `commit check` command.

In this example, the router being configured is identified using the following command prompts:

- **CE2** identifies the customer edge 2 (CE2) router
- **PE1** identifies the provider edge 1 (PE1) router
- **CE3** identifies the customer edge 3 (CE3) router
- **PE3** identifies the provider edge 3 (PE3) router
- **CE5** identifies the customer edge 5 (CE5) router
- **PE5** identifies the provider edge 5 (PE5) router

This example contains the following procedures:

- [Configuring PE Router Customer-facing and Loopback Interfaces on page 109](#)
- [Configuring Core-facing Interfaces on page 110](#)
- [Configuring Protocols on page 112](#)
- [Configuring Routing Instances and Layer 2 Circuits on page 114](#)

- [Configuring the Route Reflector on page 116](#)
- [Interconnecting the Layer 2 Circuit with the Layer 3 VPN on page 117](#)

Configuring PE Router Customer-facing and Loopback Interfaces

Step-by-Step Procedure

To begin building the interconnection, configure the interfaces on the PE routers. If your network contains provider (P) routers, configure the interfaces on the P routers also. This example shows the configuration for Router PE2, Router PE3, and Router PE5.

1. On Router PE2, configure the **ge-1/0/2** interface encapsulation. To configure the interface encapsulation, include the **encapsulation** statement and specify the **ethernet-ccc** option (**vlan-ccc** encapsulation is also supported). Configure the **ge-1/0/2.0** logical interface family for circuit cross-connect functionality. To configure the logical interface family, include the **family** statement and specify the **ccc** option. The encapsulation should be configured the same way for all routers in the Layer 2 circuit domain.

```
[edit interfaces]
ge-1/0/2 {
  encapsulation ethernet-ccc;
  unit 0 {
    family ccc;
  }
}
```

2. On Router PE2, configure the **lo0.0** interface. Include the **family** statement and specify the **inet** option. Include the **address** statement and specify **2.2.2.2/32** as the loopback IPv4 address.

```
[edit interfaces]
lo0 {
  unit 0 {
    family inet {
      address 2.2.2.2/32;
    }
  }
}
```

3. On Router PE3, configure the **ge-1/0/1** interface. Include the **family** statement and specify the **inet** option. Include the **address** statement and specify **90.90.90.1/24** as the interface address for this device.

```
[edit interfaces]
ge-1/0/1 {
  unit 0 {
    family inet {
      address 90.90.90.1/24;
    }
  }
}
```

4. On Router PE3, configure the **lo0.0** loopback interface. Include the **family** statement and specify the **inet** option. Include the **address** statement and specify **3.3.3.3/32** as the loopback IPv4 address for this router.

```
[edit interfaces]
```

```

lo0 {
  unit 0 {
    family inet {
      address 3.3.3.3/32;
    }
  }
}

```

- On Router PE5, configure the **ge-2/0/0** interface. Include the **family** statement and specify the **inet** option. Include the **address** statement and specify **80.80.80.1/24** as the interface address.

```

[edit interfaces]
ge-2/0/0 {
  unit 0 {
    family inet {
      address 80.80.80.1/24;
    }
  }
}

```

- On Router PE5, configure the **lo0.0** interface. Include the **family** statement and specify the **inet** option. Include the **address** statement and specify **5.5.5.5/32** as the loopback IPv4 address for this router.

```

[edit interfaces]
lo0 {
  unit 0 {
    family inet {
      address 5.5.5.5/32;
    }
  }
}

```

Configuring Core-facing Interfaces

Step-by-Step Procedure

This procedure describes how to configure the core-facing interfaces on the PE routers. This example does not include all the core-facing interfaces shown in the physical topology illustration. Enable the **mpls** and **inet** address families on the core-facing interfaces.

- On Router PE2, configure the **xe-0/2/0** interface. Include the **family** statement and specify the **inet** address family. Include the **address** statement and specify **10.10.5.1/30** as the interface address. Include the **family** statement and specify the **mpls** address family.

```

[edit interfaces]
xe-0/2/0 {
  unit 0 {
    family inet {
      address 10.10.5.1/30;
    }
    family mpls;
  }
}

```

2. On Router PE3, configure the core-facing interfaces. Include the **family** statement and specify the **inet** address family. Include the **address** statement and specify the IPv4 addresses shown in the example as the interface addresses. Include the **family** statement and specify the **mpls** address family. In the example, the **xe-2/1/0** interface is connected to Router PE5, and the **xe-2/2/0** interface is connected to Router PE2.

```
[edit interfaces]
xe-2/0/0 {
  unit 0 {
    family inet {
      address 10.10.20.2/30;
    }
    family mpls;
  }
}
xe-2/1/0 {
  unit 0 {
    family inet {
      address 10.10.6.1/30;
    }
    family mpls;
  }
}
xe-2/2/0 {
  unit 0 {
    family inet {
      address 10.10.5.2/30;
    }
    family mpls;
  }
}
xe-2/3/0 {
  unit 0 {
    family inet {
      address 10.10.1.2/30;
    }
    family mpls;
  }
}
```

3. On Router PE5, configure the **xe-0/1/0** interface. Include the **family** statement and specify the **inet** address family. Include the **address** statement and specify **10.10.6.2/30** as the interface address. Include the **family** statement and specify the **mpls** address family.

```
[edit interfaces]
xe-0/1/0 {
  unit 0 {
    family inet {
      address 10.10.6.2/30;
    }
    family mpls;
  }
}
```

Configuring Protocols

Step-by-Step Procedure

This procedure describes how to configure the protocols used in this example. If your network contains P routers, configure the interfaces on the P routers also.

1. On Router PE3, enable OSPF as the IGP. Enable the MPLS, LDP, and BGP protocols on all interfaces except **fxp0.0**. LDP is used as the signaling protocol for the Layer 2 circuit to Router PE2. The following configuration snippet shows the protocol configuration for Router PE3:

```
[edit]
protocols {
  rsvp {
    interface all;
    interface fxp0.0 {
      disable;
    }
  }
  mpls {
    label-switched-path to-RR {
      to 7.7.7.7;
    }
    label-switched-path to-PE2 {
      to 2.2.2.2;
    }
    label-switched-path to-PE5 {
      to 5.5.5.5;
    }
    label-switched-path to-PE4 {
      to 4.4.4.4;
    }
    label-switched-path to-PE1 {
      to 1.1.1.1;
    }
    interface all;
    interface fxp0.0 {
      disable;
    }
  }
  bgp {
    group RR {
      type internal;
      local-address 3.3.3.3;
      family inet-vpn {
        unicast;
      }
      family l2vpn {
        signaling;
      }
      neighbor 7.7.7.7;
    }
  }
  ospf {
    traffic-engineering;
    area 0.0.0.0 {
```



```

        interface all;
        interface fxp0.0 {
            disable;
        }
    }
}
ldp {
    interface all;
    interface fxp0.0 {
        disable;
    }
}
}

```

2. On Router PE2, configure the MPLS, OSPF, and LDP protocols.

```

[edit]
protocols {
    mpls {
        interface all;
        interface fxp0.0 {
            disable;
        }
    }
    ospf {
        traffic-engineering;
        area 0.0.0.0 {
            interface all;
            interface fxp0.0 {
                disable;
            }
        }
    }
    ldp {
        interface all;
        interface fxp0.0 {
            disable;
        }
    }
}

```

3. On Router PE5, enable OSPF as the IGP. Enable the MPLS, RSVP, and BGP protocols on all interfaces except **fxp0.0**. Enable core-facing interfaces with the **mpls** and **inet** address families.

```

[edit]
protocols {
    rsvp {
        interface all {
            link-protection;
        }
        interface fxp0.0 {
            disable;
        }
    }
    mpls {
        label-switched-path to-RR {

```

```

        to 7.7.7.7;
    }
    label-switched-path to-PE2 {
        to 2.2.2.2;
    }
    label-switched-path to-PE3 {
        to 3.3.3.3;
    }
    label-switched-path to-PE4 {
        to 4.4.4.4;
    }
    label-switched-path to-PE1 {
        to 1.1.1.1;
    }
    interface all;
    interface fxp0.0 {
        disable;
    }
}
bgp {
    group to-rr {
        type internal;
        local-address 5.5.5.5;
        family inet-vpn {
            unicast;
        }
        family l2vpn {
            signaling;
        }
        neighbor 7.7.7.7;
    }
}
ospf {
    traffic-engineering;
    area 0.0.0.0 {
        interface all;
        interface fxp0.0 {
            disable;
        }
    }
}
}
}

```

Configuring Routing Instances and Layer 2 Circuits

Step-by-Step Procedure

This procedure describes how to configure the Layer 2 circuit and the Layer 3 VPN.

1. On Router PE2, configure the Layer 2 circuit. Include the **l2circuit** statement. Include the **neighbor** statement and specify the loopback IPv4 address of Router PE3 as the neighbor. Include the interface statement and specify **ge-1/0/2.0** as the logical interface that is participating in the Layer 2 circuit. Include the **virtual-circuit-id** statement and specify **100** as the identifier. Include the **no-control-word** statement for equipment that does not support the control word.

[edit]

```

protocols {
  l2circuit {
    neighbor 3.3.3.3 {
      interface ge-1/0/2.0 {
        virtual-circuit-id 100;
        no-control-word;
      }
    }
  }
}

```

2. On Router PE3, configure the Layer 2 circuit to Router PE2. Include the **l2circuit** statement. Include the **neighbor** statement and specify the loopback IPv4 address of Router PE2 as the neighbor. Include the interface statement and specify **lt-1/1/10.0** as the logical tunnel interface that is participating in the Layer 2 circuit. Include the **virtual-circuit-id** statement and specify **100** as the identifier. Include the **no-control-word** statement.

```

[edit ]
protocols {
  l2circuit {
    neighbor 2.2.2.2 {
      interface lt-1/1/10.0 {
        virtual-circuit-id 100;
        no-control-word;
      }
    }
  }
}

```

3. On Router PE3, configure the Layer 3 VPN (**L3VPN**) routing instance to Router PE5 at the **[edit routing-instances]** hierarchy level. Also configure the BGP peer group at the **[edit routing-instances L3VPN protocols]** hierarchy level.

```

[edit ]
routing-instances {
  L3VPN {
    instance-type vrf;
    interface ge-1/0/1.0;
    interface lt-1/1/10.1;
    route-distinguisher 65000:33;
    vrf-target target:65000:2;
    vrf-table-label;
    protocols {
      bgp {
        export direct;
        group ce3 {
          neighbor 90.90.90.2 {
            peer-as 100;
          }
        }
      }
    }
  }
}

```

- On Router PE5, configure the Layer 3 VPN routing instance (**L3VPN**) at the **[edit routing-instances]** hierarchy level. Also configure the BGP peer group at the **[edit routing-instances L3VPN protocols]** hierarchy level.

```
[edit ]
routing-instances {
  L3VPN {
    instance-type vrf;
    interface ge-2/0/0.0;
    route-distinguisher 65000:5;
    vrf-target target:65000:2;
    vrf-table-label;
    protocols {
      bgp {
        group ce5 {
          neighbor 80.80.80.2 {
            peer-as 200;
          }
        }
      }
    }
  }
}
```

Configuring the Route Reflector

Step-by-Step Procedure

Although a route reflector is not required to interconnect a Layer 2 circuit with a Layer 3 VPN, this examples uses a route reflector. This procedure shows the relevant portion of the route reflector configuration.

- Configure the route reflector with RSVP, MPLS, BGP and OSPF. The route reflector is a BGP peer with the PE routers. Notice that the BGP peer group configuration includes the **family** statement and specifies the **inet-vpn** option. The **inet-vpn** option enables BGP to advertise network layer reachability information (NLRI) for the Layer 3 VPN routes. The configuration also includes the **family** statement and specifies the **l2vpn** option. The **l2vpn** option enables BGP to advertise NLRI for the Layer 2 circuit. Layer 2 circuits use the same internal BGP infrastructure as Layer 2 VPNs.

```
[edit ]
protocols {
  rsvp {
    interface all;
    interface fxp0.0 {
      disable;
    }
  }
  mpls {
    label-switched-path to-pe3 {
      to 3.3.3.3;
    }
    label-switched-path to-pe5 {
      to 5.5.5.5;
    }
  }
  interface all;
```

```

interface fxp0.0 {
  disable;
}
}
bgp {
  group RR {
    type internal;
    local-address 7.7.7.7;
    family inet {
      unicast;
    }
    family inet-vpn {
      unicast;
    }
    family l2vpn {
      signaling;
    }
    cluster 7.7.7.7;
    neighbor 1.1.1.1;
    neighbor 2.2.2.2;
    neighbor 4.4.4.4;
    neighbor 5.5.5.5;
    neighbor 3.3.3.3;
  }
}
ospf {
  traffic-engineering;
  area 0.0.0.0 {
    interface all;
    interface fxp0.0 {
      disable;
    }
  }
}
}

```

Interconnecting the Layer 2 Circuit with the Layer 3 VPN

Step-by-Step Procedure

Before you can configure the logical tunnel interface in an MX Series router, you must create the tunnel services interface to be used for tunnel services.

1. Create the tunnel service interface on Router PE3. Include the **bandwidth** statement at the **[edit chassis fpc slot-number pic slot-number tunnel-services]** hierarchy level and specify the amount of bandwidth to reserve for tunnel services in gigabits per second.

```

[edit chassis]
fpc 1 {
  pic 1 {
    tunnel-services {
      bandwidth 1g;
    }
  }
}
}

```

2. On Router PE3, configure the **lt-1/1/10** logical tunnel interface unit 0.

Router PE3 is the router that is *stitching* the Layer 2 circuit to the Layer 3 VPN using the logical tunnel interface. The configuration of the peer unit interfaces is what makes the interconnection.

Include the **encapsulation** statement and specify the **ethernet-ccc** option. Include the **peer-unit** statement and specify the logical interface unit 1 as the peer tunnel interface. Include the **family** statement and specify the **ccc** option.

Configure the **lt-1/1/10** logical interface unit 1 with **ethernet** encapsulation. Include the **peer-unit** statement and specify the logical interface unit 0 as the peer tunnel interface. Include the **family** statement and specify the **inet** option. Also include the **address** statement and specify **70.70.70.1/24** as the IPv4 address of the interface.



NOTE: The peering logical interfaces must belong to the same logical tunnel interface derived from the Tunnel Services PIC.

```
[edit interfaces]
lt-1/1/10 {
  unit 0 {
    encapsulation ethernet-ccc;
    peer-unit 1;
    family ccc;
  }
  unit 1 {
    encapsulation ethernet;
    peer-unit 0;
    family inet {
      address 70.70.70.1/24;
    }
  }
}
```

3. On each router, commit the configuration.

```
user@host> commit check
configuration check succeeds
user@host> commit
```

Verifying the Layer 2 Circuit to Layer 3 VPN Interconnection

To verify that the interconnection is working properly, perform these tasks:

- [Verifying That the Layer 2 Circuit Connection to Router PE3 is Up on page 119](#)
- [Verifying LDP Neighbors and Targeted LDP LSPs on Router PE2 on page 119](#)
- [Verifying the Layer 2 Circuit Routes on Router PE2 on page 120](#)
- [Verifying That the Layer 2 Circuit Connection to Router PE2 is Up on page 121](#)
- [Verifying LDP Neighbors and Targeted LDP LSPs on Router PE3 on page 121](#)
- [Verifying a BGP Peer Session with the Route Reflector on Router PE3 on page 122](#)

- [Verifying the Layer 3 VPN Routes on Router PE3 on page 122](#)
- [Verifying the Layer 2 Circuit Routes on Router PE3 on page 123](#)
- [Verifying the MPLS Routes on Router PE3 on page 123](#)
- [Verifying Traffic Flow Between Router CE2 and Router CE3 on page 124](#)
- [Verifying Traffic Flow Between Router CE2 and Router CE5 on page 124](#)

Verifying That the Layer 2 Circuit Connection to Router PE3 is Up

Purpose To verify that the Layer 2 circuit connection from Router PE2 to Router PE3 is **Up**. To also document the incoming and outgoing LDP labels and the circuit ID used by this Layer 2 circuit connection.

Action Verify that the Layer 2 circuit connection is up, using the **show l2circuit connections** command.

user@PE2> show l2circuit connections

Legend for connection status (St)

EI -- encapsulation invalid	NP -- interface h/w not present
MM -- mtu mismatch	Dn -- down
EM -- encapsulation mismatch	VC-Dn -- Virtual circuit Down
CM -- control-word mismatch	Up -- operational
VM -- vlan id mismatch	CF -- Call admission control failure
OL -- no outgoing label	IB -- TDM incompatible bitrate
NC -- intf encaps not CCC/TCC	TM -- TDM misconfiguration
BK -- Backup Connection	ST -- Standby Connection
CB -- rcvd cell-bundle size bad	SP -- Static Pseudowire
LD -- local site signaled down	RS -- remote site standby
RD -- remote site signaled down	XX -- unknown

Legend for interface status

Up -- operational
Dn -- down

Neighbor: 3.3.3.3

Interface	Type	St	Time last up	# Up trans
ge-1/0/2.0(vc 100)	rmt	Up	Jan 7 02:14:13 2010	1

Remote PE: 3.3.3.3, Negotiated control-word: No
Incoming label: 301488, Outgoing label: 315264
Negotiated PW status TLV: No
Local interface: ge-1/0/2.0, Status: Up, Encapsulation: ETHERNET

Meaning The output shows that the Layer 2 circuit connection from Router PE2 to Router PE3 is **Up** and the connection is using the **ge-1/0/2.0** interface. Note that the outgoing label is **315264** and the incoming label is **301488**, the virtual circuit (VC) identifier is **100** and the encapsulation is **ETHERNET**.

Verifying LDP Neighbors and Targeted LDP LSPs on Router PE2

Purpose To verify that Router PE2 has a targeted LDP LSP to Router PE3 and that Router PE2 and Router PE3 are LDP neighbors.

Action Verify that Router PE2 has a targeted LDP LSP to Router PE3 and that Router PE2 and Router PE3 are LDP neighbors, using the **show ldp neighbor** command.

```
user@PE2> show ldp neighbor
Address          Interface      Label space ID      Hold time
3.3.3.3          lo0.0         3.3.3.3:0          38
```

Meaning The output shows that Router PE2 has an LDP neighbor with the IPv4 address of **3.3.3.3**. Address 3.3.3.3 is the lo0.0 interface address of Router PE3. Notice that Router PE2 uses the local **lo0.0** interface for the LSP.

Verifying that the routers are LDP neighbors also verifies that the targeted LSP is established.

Verifying the Layer 2 Circuit Routes on Router PE2

Purpose To verify that Router PE2 has a route for the Layer 2 circuit and that the route uses the LDP MPLS label to Router PE3.

Action Verify that Router PE2 has a route for the Layer 2 circuit and that the route uses the LDP MPLS label to Router PE3, using the **show route table mpls.0** command.

```
user@PE2> show route table mpls.0
mpls.0: 13 destinations, 13 routes (13 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

0          *[MPLS/0] 1w3d 05:24:11, metric 1
           Receive
1          *[MPLS/0] 1w3d 05:24:11, metric 1
           Receive
2          *[MPLS/0] 1w3d 05:24:11, metric 1
           Receive
300560     *[LDP/9] 16:12:23, metric 1
           > to 10.10.2.1 via xe-0/1/0.0, Pop
300560(S=0) *[LDP/9] 16:12:23, metric 1
           > to 10.10.2.1 via xe-0/1/0.0, Pop
301008     *[LDP/9] 16:12:23, metric 1
           > to 10.10.4.2 via xe-0/3/0.0, Swap 299856
301488     *[L2CKT/7] 11:07:28
           > via ge-1/0/2.0, Pop
301536     *[LDP/9] 16:12:23, metric 1
           > to 10.10.4.2 via xe-0/3/0.0, Pop
301536(S=0) *[LDP/9] 16:12:23, metric 1
           > to 10.10.4.2 via xe-0/3/0.0, Pop
301712     *[LDP/9] 12:41:22, metric 1
           > to 10.10.5.2 via xe-0/2/0.0, Swap 315184
301728     *[LDP/9] 12:41:22, metric 1
           > to 10.10.5.2 via xe-0/2/0.0, Pop
301728(S=0) *[LDP/9] 12:41:22, metric 1
           > to 10.10.5.2 via xe-0/2/0.0, Pop
ge-1/0/2.0 *[L2CKT/7] 11:07:28, metric 1
           > to 10.10.5.2 via xe-0/2/0.0, Push 315264
```

Meaning The output shows that Router PE2 pushes the **315264** outgoing label on the **L2CKT** route going out interface **ge-1/0/2.0**. The output also shows that Router PE2 pops the **301488** incoming label on the **L2CKT** coming from interface **ge-1/0/2.0**

Verifying That the Layer 2 Circuit Connection to Router PE2 is Up

Purpose To verify that the Layer 2 circuit connection from Router PE3 to Router PE2 is **Up**. To also document the incoming and outgoing LDP labels and the circuit ID used by this Layer 2 circuit connection.

Action Verify that the Layer 2 circuit connection is up, using the **show l2circuit connections** command.

```
user@PE3> show l2circuit connections
```

Layer-2 Circuit Connections:

Legend for connection status (St)

EI -- encapsulation invalid	NP -- interface h/w not present
MM -- mtu mismatch	Dn -- down
EM -- encapsulation mismatch	VC-Dn -- Virtual circuit Down
CM -- control-word mismatch	Up -- operational
VM -- vlan id mismatch	CF -- Call admission control failure
OL -- no outgoing label	IB -- TDM incompatible bitrate
NC -- intf encaps not CCC/TCC	TM -- TDM misconfiguration
BK -- Backup Connection	ST -- Standby Connection
CB -- rcvd cell-bundle size bad	XX -- unknown

Legend for interface status

Up -- operational

Dn -- down

Neighbor: 2.2.2.2

Interface	Type	St	Time last up	# Up trans
lt-1/1/10.0(vc 100)	rmt	Up	Jan 7 02:15:03 2010	1

Remote PE: 2.2.2.2, Negotiated control-word: No
Incoming label: 315264, Outgoing label: 301488
Local interface: lt-1/1/10.0, Status: Up, Encapsulation: ETHERNET

Meaning The output shows that the Layer 2 circuit connection from Router PE3 to Router PE2 is **Up** and the connection is using the logical tunnel (lt) interface. Note that the incoming label is **315264** and the outgoing label is **301488**, the virtual circuit (VC) identifier is **100**, and that the encapsulation is **ETHERNET**.

Verifying LDP Neighbors and Targeted LDP LSPs on Router PE3

Purpose To verify that Router PE3 has a targeted LDP LSP to Router PE2 and that Router PE3 and Router PE2 are LDP neighbors.

Action Verify that Router PE2 has a targeted LDP LSP to Router PE3 and that Router PE2 and Router PE3 are LDP neighbors, using the **show ldp neighbor** command.

```
user@PE2> show ldp neighbor
```

Address	Interface	Label space ID	Hold time
2.2.2.2	lo0.0	2.2.2.2:0	43
4.4.4.4	lo0.0	4.4.4.4:0	33

Meaning The output shows that Router PE3 has an LDP neighbor with the IPv4 address of **2.2.2.2**. Address 2.2.2.2 is the lo0.0 interface address of Router PE2. The output also shows that the interface used on Router PE3 for the LSP is **lo0.0**. Verifying that the routers are LDP neighbors also verifies that the targeted LSP is established.

Verifying a BGP Peer Session with the Route Reflector on Router PE3

Purpose To verify that Router PE3 has a peer session established with the route reflector.

Action Verify that Router PE3 has a peer session established with the route reflector, using the **show bgp summary** command.

```
user@PE2> show bgp summary
```

```
Groups: 2 Peers: 2 Down peers: 0
Table      Tot Paths  Act Paths Suppressed  History  Damp State   Pending
bgp.13vpn.0      1          1          0          0        0      0        0
Peer          AS      InPkt   OutPkt   OutQ   Flaps  Last Up/Dwn
State|#Active/Received/Accepted/Damped...
7.7.7.7        65000      1597     1612        0        1   12:03:21 Establ
  bgp.12vpn.0: 0/0/0/0
  bgp.13vpn.0: 1/1/1/0
  L3VPN.inet.0: 1/1/1/0
```

Meaning The output shows that Router PE3 has a peer session with the router with the IPv4 address of **7.7.7.7**. Address 7.7.7.7 is the lo0.0 interface address of the route reflector. The output also shows that the peer session state is **Establ**, meaning that the session is established.

Verifying the Layer 3 VPN Routes on Router PE3

Purpose To verify that Router PE3 has Layer 3 VPN routes to Router CE2, Router CE3, and Router CE5.

Action Verify that Router PE3 has routes to Router CE2, Router CE3, and Router CE5 in the Layer 3 VPN route table, using the **show route table L3VPN.inet.0** command. In this example, **L3VPN** is the name configured for the routing instance.

```
user@PE3> show route table L3VPN.inet.0
L3VPN.inet.0: 5 destinations, 5 routes (5 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

70.70.70.0/24    *[Direct/0] 11:13:59
                 > via lt-1/1/10.1
70.70.70.1/32    *[Local/0] 11:13:59
                 Local via lt-1/1/10.1
80.80.80.0/24    *[BGP/170] 11:00:41, localpref 100, from 7.7.7.7
                 AS path: I
                 > to 10.10.6.2 via xe-2/1/0.0, Push 16
90.90.90.0/24    *[Direct/0] 11:54:41
                 > via ge-1/0/1.0
90.90.90.1/32    *[Local/0] 11:54:41
                 Local via ge-1/0/1.0
```

Meaning The output shows that Router PE3 has a route to the IPv4 subnetwork address of **70.70.70.0**. Address 70.70.70.2 is the interface address of Router CE2. The output shows that Router PE3 has a route to the IPv4 subnetwork address of **80.80.80.0**. Address 80.80.80.2 is the interface address of Router CE5. The output shows that Router PE3 has a route to the IPv4 subnetwork address of **90.90.90.0**. Address 90.90.90.2 is the interface address of Router CE3.

Verifying the Layer 2 Circuit Routes on Router PE3

Purpose To verify that Router PE3 has a route to Router PE2 in the Layer 2 circuit route table.

Action Verify that Router PE3 has a route to Router PE2 in the Layer 2 circuit route table, using the **show route table l2circuit.0** command.

```
user@PE3> show route table l2circuit.0
2.2.2.2:NoCtrlWord:5:100:Local/96 (1 entry, 1 announced)
    *L2CKT Preference: 7
        Next hop type: Indirect
        Next-hop reference count: 1
        Next hop type: Router
        Next hop: 10.10.5.1 via xe-2/2/0.0, selected
        Protocol next hop: 2.2.2.2
        Indirect next hop: 8cae0a0 -
        State: <Active Int>
        Local AS: 65000
        Age: 11:16:50 Metric2: 1
        Task: 12 circuit
        Announcement bits (1): 0-LDP
        AS path: I
        VC Label 315264, MTU 1500
```

Meaning The output shows that Router PE3 has a route to the IPv4 address of **2.2.2.2**. Address 2.2.2.2 is the lo0.0 interface address of Router PE2. Note that the VC label is **315264**. This label is the same as the incoming MPLS label displayed using the **show l2circuit connections** command.

Verifying the MPLS Routes on Router PE3

Purpose To verify that Router PE3 has a route to Router PE2 in the MPLS route table.

Action Verify Router PE3 has a route to Router PE2 in the MPLS route table, using the **show route table mpls.0** command.

```
user@PE3> show route table mpls.0
mpls.0: 21 destinations, 21 routes (21 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

0          *[MPLS/0] 1w3d 05:29:02, metric 1
            Receive
1          *[MPLS/0] 1w3d 05:29:02, metric 1
            Receive
2          *[MPLS/0] 1w3d 05:29:02, metric 1
            Receive
16         *[VPN/0] 12:22:45
            to table L3VPN.inet.0, Pop
315184     *[LDP/9] 12:45:14, metric 1
            > to 10.10.20.1 via xe-2/0/0.0, Pop
315184(S=0) *[LDP/9] 12:45:14, metric 1
            > to 10.10.20.1 via xe-2/0/0.0, Pop
315200     *[LDP/9] 00:03:53, metric 1
            > to 10.10.20.1 via xe-2/0/0.0, Swap 625297
            to 10.10.6.2 via xe-2/1/0.0, Swap 299856
315216     *[LDP/9] 12:45:14, metric 1
            > to 10.10.6.2 via xe-2/1/0.0, Pop
```

```

315216(S=0)      *[LDP/9] 12:45:14, metric 1
                  > to 10.10.6.2 via xe-2/1/0.0, Pop
315232           *[LDP/9] 12:45:06, metric 1
                  > to 10.10.1.1 via xe-2/3/0.0, Pop
315232(S=0)      *[LDP/9] 12:45:06, metric 1
                  > to 10.10.1.1 via xe-2/3/0.0, Pop
315248           *[LDP/9] 12:45:14, metric 1
                  > to 10.10.5.1 via xe-2/2/0.0, Pop
315248(S=0)      *[LDP/9] 12:45:14, metric 1
                  > to 10.10.5.1 via xe-2/2/0.0, Pop
315264           *[L2CKT/7] 11:11:20
                  > via lt-1/1/10.0, Pop
315312           *[RSVP/7] 11:26:01, metric 1
                  > to 10.10.6.2 via xe-2/1/0.0, label-switched-path to-pe5
315312(S=0)      *[RSVP/7] 11:26:01, metric 1
                  > to 10.10.6.2 via xe-2/1/0.0, label-switched-path to-pe5
315328           *[RSVP/7] 11:26:01, metric 1
                  > to 10.10.20.1 via xe-2/0/0.0, label-switched-path to-RR
315360           *[RSVP/7] 11:26:01, metric 1
                  > to 10.10.20.1 via xe-2/0/0.0, label-switched-path to-RR
316208           *[RSVP/7] 00:03:32, metric 1
                  > to 10.10.6.2 via xe-2/1/0.0, label-switched-path
Bypass->10.10.9.1
316208(S=0)      *[RSVP/7] 00:03:32, metric 1
                  > to 10.10.6.2 via xe-2/1/0.0, label-switched-path
Bypass->10.10.9.1
lt-1/1/10.0      *[L2CKT/7] 11:11:20, metric2 1
                  > to 10.10.5.1 via xe-2/2/0.0, Push 301488

```

Meaning The output shows that Router PE3 has a route for the Layer 2 circuit and that the route uses the LDP MPLS label to Router PE2. Notice that the **301488** label is the same as the outgoing label displayed on Router PE2 using the **show l2circuit connections** command.

Verifying Traffic Flow Between Router CE2 and Router CE3

Purpose To verify that the CE routers can send and receive traffic across the interconnection.

Action Verify that Router CE2 can send traffic to and receive traffic from Router CE3 across the interconnection, using the **ping** command.

```

user@CE2>ping 90.90.90.2
PING 90.90.90.2 (90.90.90.2): 56 data bytes
64 bytes from 90.90.90.2: icmp_seq=0 ttl=63 time=0.708 ms
64 bytes from 90.90.90.2: icmp_seq=1 ttl=63 time=0.610 ms

```

Meaning The output shows that Router CE2 can send an ICMP request to and receive a response from Router CE3 across the interconnection.

Verifying Traffic Flow Between Router CE2 and Router CE5

Purpose To verify that the CE routers can send and receive traffic across the interconnection.

Action Verify that Router CE2 can send traffic to and receive traffic from Router CE5 across the interconnection, using the **ping** command.

```

user@CE2>ping 80.80.80.2

```

```
PING 80.80.80.2 (80.80.80.2): 56 data bytes
64 bytes from 80.80.80.2: icmp_seq=0 ttl=62 time=0.995 ms
64 bytes from 80.80.80.2: icmp_seq=1 ttl=62 time=1.005 ms
```

Meaning The output shows that Router CE2 can send an ICMP request to and receive a response from Router CE5 across the interconnection.

- Related Documentation**
- Layer 2 Circuit Overview
 - Layer 3 VPN Overview
 - [Applications for Interconnecting a Layer 2 Circuit with a Layer 3 VPN on page 106](#)

PART 3

Administration

- [Layer 2 Circuit Reference on page 129](#)
- [Summary of Layer 2 Circuit Configuration Statements on page 131](#)

CHAPTER 5

Layer 2 Circuit Reference

- [Supported Layer 2 Circuit Standards on page 129](#)

Supported Layer 2 Circuit Standards

Junos OS substantially supports the following RFCs, which define standards for Layer 2 circuits.

- RFC 4447, *Pseudowire Setup and Maintenance Using the Label Distribution Protocol (LDP)*

Junos OS does not support Section 5.3, "The Generalized PWid FEC Element."

- RFC 4448, *Encapsulation Methods for Transport of Ethernet over MPLS Networks*

The following Internet drafts do not define standards, but provide information about Layer 2 technologies. The IETF classifies them as "Historic."

- Internet draft draft-martini-l2circuit-encap-mpls-11.txt, *Encapsulation Methods for Transport of Layer 2 Frames Over IP and MPLS Networks*

Junos OS differs from the Internet draft in the following ways:

- A packet with a sequence number of 0 (zero) is treated as out of sequence.
- Any packet that does not have the next incremental sequence number is considered out of sequence.
- When out-of-sequence packets arrive, the expected sequence number for the neighbor is set to the sequence number in the Layer 2 circuit control word.
- Internet draft draft-martini-l2circuit-trans-mpls-19.txt, *Transport of Layer 2 Frames Over MPLS*

Related Documentation

- [Supported Carrier-of-Carriers and Interprovider VPN Standards](#)
- [Supported Layer 2 VPN Standard](#)
- [Supported Layer 3 VPN Standards](#)
- [Supported Multicast VPN Standards](#)
- [Supported VPLS Standards](#)

- Accessing Standards Documents on the Internet

CHAPTER 6

Summary of Layer 2 Circuit Configuration Statements


bandwidth (Protocols Layer 2 Circuit)

Syntax	<code>bandwidth (<i>bandwidth</i> <i>ctnumber bandwidth</i>);</code>
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols l2circuit neighbor <i>address</i> interface <i>interface-name</i>], [edit protocols l2circuit neighbor <i>address</i> interface <i>interface-name</i>]
Release Information	Statement introduced before Junos OS Release 7.4.
Description	Specify bandwidth allocation for a Layer 2 circuit or for the class types of a Layer 2 circuit.
Options	<p><i>bandwidth</i>—Configure the bandwidth in bits per second for the Layer 2 circuit. You cannot configure the bandwidth for the Layer 2 circuit and for the class types at the same time.</p> <p><i>ctnumber bandwidth</i>—Configure the bandwidth in bits per second for a class type on the Layer 2 circuit. You can configure bandwidth for up to four class types (ct0, ct1, ct2, ct3) per Layer 2 circuit. If you configure the class types, you must configure them in order, starting with class type ct0.</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">• Configuring Bandwidth Allocation and Call Admission Control in Layer 2 Circuits on page 48

backup-interface (Layer 2 Circuits)

Syntax	<code>backup-interface <i>interface-name</i>;</code>
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols l2circuit local-switching interface <i>interface-name</i> end-interface interface <i>interface-name</i>], [edit protocols l2circuit local-switching interface <i>interface-name</i> end-interface interface <i>interface-name</i>]
Release Information	Statement introduced in Junos OS Release 12.3.
Description	Specify the interface to be used by the protection pseduowire in connection protection configurations for Layer 2 circuits.
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">• Example: Configuring Layer 2 Circuit Switching Protection on page 58

backup-neighbor

Syntax	<pre> backup-neighbor address { community name; mtu number; psn-tunnel-endpoint address; standby; static; virtual-circuit-id number; } </pre>
Hierarchy Level	<pre> [edit logical-systems logical-system-name protocols l2circuit local-switching interface interface-name], [edit logical-systems logical-system-name protocols l2circuit neighbor address interface interface-name], [edit logical-systems logical-system-name routing-instances routing-instance-name protocols vpls neighbor address], [edit protocols l2circuit local-switching interface interface-name], [edit protocols l2circuit neighbor address interface interface-name], [edit routing-instances routing-instance-name protocols vpls neighbor address] </pre>
Release Information	<p>Statement introduced in Junos OS Release 9.2.</p> <p>Statement introduced in Junos OS Release 12.2 for ACX Series routers.</p> <p>Statement introduced in Junos OS Release 12.3 at the [edit protocols l2circuit local-switching interface interface-name] hierarchy level.</p>
Description	<p>Configure pseudowire redundancy for Layer 2 circuits and VPLS. A redundant pseudowire can act as a backup connection and can be configured between a PE router and a CE device or between PE routers, maintaining Layer 2 circuit and VPLS services after certain types of failures. This feature can help improve the reliability of certain types of networks where a single point of failure could interrupt service for customers.</p>
	<div>  <p>NOTE: VPLS is not supported on ACX Series routers. The psn-tunnel-endpoint statement is not supported at the [edit protocols l2circuit local-switching interface interface-name end-interface interface interface-name] hierarchy level.</p> </div>
Options	<p>address—Specifies the address for the backup neighbor.</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"> Configuring Redundant Pseudowires for Layer 2 Circuits and VPLS Example: Configuring Layer 2 Circuit Switching Protection on page 58 community (Protocols Layer 2 Circuit) on page 134 psn-tunnel-endpoint on page 163

- [virtual-circuit-id on page 168](#)

community (Protocols Layer 2 Circuit)

Syntax	<pre>community <i>community-name</i> { invert-match; members <i>community-members</i>; }</pre>
Hierarchy Level	<pre>[edit logical-systems <i>logical-system-name</i> policy-options], [edit logical-systems <i>logical-system-name</i> protocols l2circuit local-switching interface <i>interface-name</i> backup-neighbor <i>address</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> protocols l2circuit neighbor <i>address</i> interface <i>interface-name</i> backup-neighbor <i>address</i>], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols vpls neighbor <i>address</i> backup-neighbor <i>address</i>], [edit policy-options], [edit protocols l2circuit local-switching interface <i>interface-name</i> backup-neighbor <i>address</i>], [edit protocols l2circuit neighbor <i>address</i> interface <i>interface-name</i>], [edit protocols l2circuit neighbor <i>address</i> interface <i>interface-name</i> backup-neighbor <i>address</i>], [edit routing-instances <i>routing-instance-name</i> protocols vpls neighbor <i>address</i> backup-neighbor <i>address</i>]</pre>
Release Information	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Hierarchy levels associated with the backup-neighbor statement (pseudowire redundancy) added in Junos OS Release 9.2.</p> <p>Statement introduced in Junos OS Release 12.3 at the [edit protocols l2circuit local-switching interface <i>interface-name</i> backup-neighbor <i>address</i>] hierarchy level.</p>
Description	Specify the community for the Layer 2 circuit.
Options	<p><i>community-name</i>—Name of the Layer 2 circuit community.</p> <p><i>invert-match</i>—Invert the results of the community expression match.</p> <p><i>members community-members</i>—Specify the members of the community.</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 the Layer 2 Circuit Community on page 41• Configuring Redundant Pseudowires for Layer 2 Circuits and VPLS• Example: Configuring Layer 2 Circuit Switching Protection on page 58

connection-protection

Syntax	connection-protection;
Hierarchy Level	[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], [edit protocols l2circuit local-switching interface <i>interface-name</i>], [edit protocols l2circuit neighbor <i>address</i> interface <i>interface-name</i>]
Release Information	Statement introduced in Junos OS Release 12.3.
Description	Enable connection protection on the Layer 2 circuit. Connection protection enables you to configure a redundant pseudowire to act as a backup connection and can be configured between PE routers, maintaining Layer 2 circuit and VPLS services after certain types of failures. This feature helps to improve the reliability of networks where a single point of failure could interrupt service for customers.
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"> • Example: Configuring Layer 2 Circuit Switching Protection on page 58

control-word (Protocols Layer 2 Circuit Neighbor)

Syntax	(control-word no-control-word);
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols l2circuit neighbor <i>address</i> interface <i>interface-name</i>], [edit protocols l2circuit neighbor <i>address</i> interface <i>interface-name</i>]
Release Information	Statement introduced before Junos OS Release 7.4.
Description	Specify the control word. The control word is four bytes long and is inserted between the Layer 2 protocol data unit (PDU) being transported and the virtual circuit (VC) label that is used for demultiplexing.
Options	<p>control-word—Enable the use of the control word.</p> <p>Default: A null control word is enabled by default. You can also configure the control word explicitly using the control-word statement.</p> <p>no-control-word—Disable the use of the control word.</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"> • Configuring the Control Word for Frame Relay Interfaces on page 34

description (Protocols Layer 2 Circuit Neighbor)

Syntax	<code>description text;</code>
Hierarchy Level	<code>[edit logical-systems <i>logical-system-name</i> protocols l2circuit neighbor <i>address</i> interface <i>interface-name</i>],</code> <code>[edit protocols l2circuit neighbor <i>address</i> interface <i>interface-name</i>]</code>
Release Information	Statement introduced before Junos OS Release 7.4.
Description	Provide a text description for the Layer 2 circuit. If the text includes one or more spaces, enclose the entire text string in quotation marks (" ").
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">• Configuring the Description on page 19

egress-protection (Layer 2 circuit)

Syntax	<pre>egress-protection { protected-l2circuit { egress-pe <i>address</i>; ingress-pe <i>address</i>; virtual-circuit-id <i>identifier</i>; } protector-interface <i>interface-name</i>; protector-pe <i>address</i> { context-identifier <i>identifier</i>; lsp <i>lsp-name</i>; } }</pre>
Hierarchy Level	<code>[edit logical-systems <i>logical-system-name</i> protocols l2circuit neighbor <i>address</i> interface <i>interface-name</i>],</code> <code>[edit protocols l2circuit neighbor <i>address</i> interface <i>interface-name</i>]</code>
Release Information	Statement introduced in Junos OS release 10.4.
Description	Configures an egress protection virtual circuit (EPVC).
Options	The other 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">• Example: Configuring an Egress Protection LSP for a Layer 2 Circuit on page 70

egress-protection (MPLS)

Syntax	<pre>egress-protection { context-identifier <i>context-id</i> { primary protector; metric <i>igp-metric-value</i>; } }</pre>
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols mpls], [edit logical-systems <i>logical-system-name</i> protocols mpls label-switched-path <i>lsp-name</i>], [edit protocols mpls], [edit protocols mpls label-switched-path <i>lsp-name</i>]
Release Information	Statement introduced in Junos OS Release 10.4. Options primary , protector , and metric introduced in Junos OS Release 11.4R3.
Description	Enables an Edge Protection Virtual Circuit (EPVC) for the MPLS protocol.
Options	<p>context-identifier <i>context-id-ip-address</i>—(Optional) The context identifier IPv4 address.</p> <p>metric <i>igp-metric-value</i>—(Optional) The IGP metric value ranging from 2 through 16777215.</p> <p>(primary protector)—On the primary PE router, configure as type primary. On the protector PE router, configure as type protector.</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"> • • Example: Configuring Egress Protection for Layer 3 VPN Services

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 Switches (flexible-ethernet-services , ethernet-ccc , and ethernet-tcc options only).
Description	Specify the physical link-layer encapsulation type. Not all encapsulation types are supported on the switches. See the switch CLI.
Default	ppp —Use serial PPP encapsulation.
Options	<p>atm-ccc-cell-relay—Use ATM cell-relay encapsulation.</p> <p>atm-pvc—Use ATM PVC encapsulation.</p> <p>cisco-hdlc—Use Cisco-compatible High-Level Data Link Control (HDLC) framing.</p> <p>cisco-hdlc-ccc—Use Cisco-compatible HDLC framing on CCC circuits.</p> <p>cisco-hdlc-tcc—Use Cisco-compatible HDLC framing on TCC circuits for connecting different media.</p> <p>ethernet-bridge—Use Ethernet bridge encapsulation on Ethernet interfaces that have bridging enabled and that must accept all packets.</p> <p>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.</p> <p>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, <i>Multiprotocol Encapsulation over ATM Adaptation Layer 5</i>, 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</p>

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.

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. 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), 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.

frame-relay-ccc—Use Frame Relay encapsulation on CCC circuits.

frame-relay-ether-type—Use Frame Relay ether type encapsulation for compatibility with the Cisco Frame Relay.

frame-relay-ether-type-tcc—Use Frame Relay ether type TCC for Cisco-compatible Frame Relay on TCC circuits to connect different media.

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. When you use this encapsulation type, you can configure the **ccc** family only.

frame-relay-tcc—Use Frame Relay encapsulation on TCC circuits to connect different media.

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.

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-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.

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 Encapsulation
- Configuring Extended VLAN Encapsulation
- Configuring Encapsulation for Layer 2 Wholesale VLAN Interfaces
- [Configuring Interfaces for Layer 2 Circuits on page 32](#)
- Configuring Interface Encapsulation on PTX Series Packet Transport Switches
- Configuring an MPLS-Based Layer 2 VPN (CLI Procedure)
- 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
- Understanding Encapsulation on an Interface

encapsulation-type (Layer 2 Circuits)

Syntax	encapsulation-type (atm-aal5 atm-cell atm-cell-port-mode atm-cell-vc-mode atm-cell-vp-mode cesop cisco-hdlc ethernet ethernet-vlan frame-relay frame-relay-port-mode interworking ppp satop-e1 satop-e3 satop-t1 satop-t3);
Hierarchy Level	[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 protocols l2circuit local-switching interface <i>interface-name</i>], [edit protocols l2circuit neighbor <i>address</i> interface <i>interface-name</i>]
Release Information	Statement introduced in Junos OS Release 9.2.
Description	Specify the type of Layer 2 traffic transiting the Layer 2 circuit.
Options	<p>atm-aal5—ATM Adaptation Layer (AAL/5)</p> <p>atm-cell—ATM cell relay</p> <p>atm-cell-port-mode—ATM cell relay port promiscuous mode</p> <p>atm-cell-vc-mode—ATM VC cell relay nonpromiscuous mode</p> <p>atm-cell-vp-mode—ATM virtual path (VP) cell relay promiscuous mode</p> <p>cesop—CESOP-based Layer 2 circuit</p> <p>cisco-hdlc—Cisco Systems-compatible HDLC</p> <p>ethernet—Ethernet</p> <p>ethernet-vlan—Ethernet VLAN</p> <p>frame-relay—Frame Relay</p> <p>frame-relay-port-mode—Frame Relay port mode</p> <p>interworking—Layer 2.5 interworking</p> <p>ppp—PPP</p> <p>satop-e1—SATSOP-E1-based Layer 2 circuit</p> <p>satop-e3—SATSOP-E3-based Layer 2 circuit</p> <p>satop-t1—SATSOP-T1-based Layer 2 circuit</p> <p>satop-t3—SATSOP-T3-based Layer 2 circuit</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**
- [Configuring the Encapsulation Type for the Layer 2 Circuit Neighbor Interface on page 35](#)

end-interface

Syntax

```
end-interface {  
  interface interface-name;  
  no-revert;  
  protect-interface interface-name;  
}
```

Hierarchy Level [edit logical-systems *logical-system-name* protocols l2circuit local-switching interface *interface-name*],
[edit protocols l2circuit local-switching interface *interface-name*]

Release Information Statement introduced before Junos OS Release 7.4.

Description Specify the end interface for a local interface switch.



.....

NOTE: The protect interface must be configured prior to configuring the no-revert statement.

.....

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 Local Interface Switching in Layer 2 Circuits on page 30](#)

fast-aps-switch

Syntax	fast-aps-switch;
Hierarchy Level	[edit interfaces <i>interface-name</i> sonet-options aps]
Release Information	Statement introduced in Junos OS Release 12.1.
Description	(M320 routers with Channelized OC3/STM1 Circuit Emulation PIC with SFP only) Reduce the Automatic Protection Switching (APS) switchover time in Layer 2 circuits.



NOTE:

- Configuring this statement reduces the APS switchover time only when the Layer 2 circuit encapsulation type for the interface receiving traffic from a Layer 2 circuit neighbor is SAToP.
- When the fast-aps-switch statement is configured in revertive APS mode, you must configure an appropriate value for revert time to achieve reduction in APS switchover time.
- To prevent the logical interfaces in the data path from being shut down, configure appropriate hold-time values on all the interfaces in the data path that support TDM.
- The fast-aps-switch statement cannot be configured when the APS annex-b option is configured.
- The interfaces that have the fast-aps-switch statement configured cannot be used in virtual private LAN service (VPLS) environments.

Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • Reducing APS Switchover Time in Layer 2 Circuits on page 49

ignore-encapsulation-mismatch

Syntax	ignore-encapsulation-mismatch;
Hierarchy Level	[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 protocols l2circuit local-switching interface <i>interface-name</i>], [edit protocols l2circuit neighbor <i>address</i> interface <i>interface-name</i>]
Release Information	Statement introduced in Junos OS Release 9.2. Statement extended to support local switching in Junos OS Release 10.4.
Description	Allow a Layer 2 circuit to be established even though the encapsulation configured on the CE device interface does not match the encapsulation configured on the Layer 2 circuit interface.
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">• Enabling the Layer 2 Circuit When the Encapsulation Does Not Match on page 35

ignore-mtu-mismatch

Syntax	ignore-mtu-mismatch;
Hierarchy Level	[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 protocols l2circuit local-switching interface <i>interface-name</i>], [edit protocols l2circuit neighbor <i>address</i> interface <i>interface-name</i>]
Release Information	Statement introduced in Junos OS Release 8.5. Support for remote PE routers added in Junos OS Release 9.2.
Description	Ignore the MTU configuration set for the physical interface associated with the local switching interface or with the remote PE router. This allows a Layer 2 circuit to be brought up between two logical interfaces that are defined on physical interfaces with different MTU values.
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">• Enabling the Layer 2 Circuit When the MTU Does Not Match on page 36• Enabling Local Interface Switching When the MTU Does Not Match on page 31

install-nexthop

Syntax	<code>install-nexthop (except lsp <i>lsp-name</i> lsp-regex <i>lsp-regular-expression</i>);</code>
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> policy-options policy-statement <i>policy-name</i> term <i>term-name</i> then], [edit policy-options policy-statement <i>policy-name</i> term <i>term-name</i> then]
Release Information	Statement introduced before Junos OS Release 7.4.
Description	Select a specific label-switched path (LSP), or select an LSP from a set of similarly named LSPs as the traffic destination for the configured community. Also can prevent the installation of any matching next hops.
Options	<p>except—Prevent the installation of any matching next hops.</p> <p>lsp <i>lsp-name</i>—Configure a specific LSP.</p> <p>lsp-regex <i>lsp-regular-expression</i>—Configure a range of similarly named LSPs. You can use the following wildcard characters when configuring an LSP regular expression:</p> <ul style="list-style-type: none">• Asterisk (*)—Match any characters.• Period (.)—Match any single digit.
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">• Configuring the Policy Statement for the Layer 2 Circuit Community on page 42

interface (Protocols Layer 2 Circuit)

Syntax interface *interface-name* {
 backup-neighbor *address* {
 community *name*;
 psn-tunnel-endpoint *address*;
 standby;
 virtual-circuit-id *number*;
 }
 bandwidth (*bandwidth* | *ctnumber bandwidth*);
 community *community-name*;
 (control-word | no-control-word);
 description *text*;
 egress-protection {
 protected-l2circuit {
 egress-pe *address*;
 ingress-pe *address*;
 virtual-circuit-id *identifier*;
 }
 protector-interface *interface-name*;
 protector-pe *address* {
 context-identifier *identifier*;
 lsp *lsp-name*;
 }
 }
 encapsulation-type *type*;
 ignore-encapsulation-mismatch;
 ignore-mtu-mismatch;
 mtu *mtu-number*;
 no-revert;
 protect-interface *interface-name*;
 pseudowire-status-tlv;
 psn-tunnel-endpoint *address*;
 revert-time *seconds*;
 switchover-delay *milliseconds*;
 virtual-circuit-id *identifier*;
 }

Hierarchy Level [edit logical-systems *logical-system-name* protocols l2circuit local-switching],
 [edit logical-systems *logical-system-name* protocols l2circuit neighbor *address*],
 [edit protocols l2circuit local-switching],
 [edit protocols l2circuit neighbor *address*]

Release Information Statement introduced before Junos OS Release 7.4.

Description Interface over which Layer 2 circuit traffic travels.

Options *interface-name*—Name of the interface to configure.

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">• Configuring the Neighbor Interface for the Layer 2 Circuit on page 32

l2circuit

Syntax	<pre> l2circuit { local-switching { interface interface-name { description text; } end-interface { interface interface-name; protect-interface interface-name; } ignore-mtu-mismatch; protect-interface interface-name; } } neighbor address { interface interface-name { bandwidth (bandwidth ctnumber bandwidth); community community-name; (control-word no-control-word); description text; encapsulation-type type; ignore-encapsulation-mismatch; ignore-mtu-mismatch; mtu mtu-number; protect-interface interface-name; pseudowire-status-tlv; psn-tunnel-endpoint address; virtual-circuit-id identifier; } } traceoptions { file filename <files number> <size size> <world-readable no-world-readable>; flag flag <flag-modifier> <disable>; } </pre>
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols], [edit protocols]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 11.1 for EX Series switches.
Description	Enables a Layer 2 circuit. 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"> • Configuring ATM Trunking on Layer 2 Circuits on page 47 • Configuring Bandwidth Allocation and Call Admission Control in Layer 2 Circuits on page 48

- [Configuring Interfaces for Layer 2 Circuits on page 32](#)
- [Configuring LDP for Layer 2 Circuits on page 9](#)
- [Configuring Policies for Layer 2 Circuits on page 41](#)
- [Configuring Static Layer 2 Circuits on page 40](#)
- [Tracing Layer 2 Circuit Operations on page 171](#)

l2vpn-use-bgp-rules

Syntax	l2vpn-use-bgp-rules;
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols bgp path-selection], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp path-selection], [edit protocols bgp path-selection], [edit routing-instances <i>routing-instance-name</i> protocols bgp path-selection]
Release Information	Statement introduced in Junos OS Release 12.3. Statement introduced in Junos OS Release 12.3 for ACX Series routers.
Description	Enable routers to use both the BGP path selection algorithm and the designated forwarder path selection algorithm when selecting the preferred path to each destination in a Layer 2 VPN or VPLS routing instance. The BGP path selection algorithm is used by all of the Provider routers participating in the routing instance. The designated forwarder path selection algorithm is used by the PE router participating in the routing instance.
Default	By default, the designated forwarder path selection algorithm is used to select the best path to reach each destination within Layer 2 VPN and VPLS routing instances.
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"> • Enabling BGP Path Selection for Layer 2 VPNs and VPLS on page 45 • route-distinguisher

local-switching (Layer 2 Circuits)

Syntax

```
local-switching {  
  interface interface-name {  
    description text;  
    encapsulation-type;  
    end-interface {  
      interface interface-name;  
      no-revert;  
      protect-interface interface-name;  
    }  
    ignore-encapsulation-mismatch;  
    ignore-mtu-mismatch;  
    no-revert;  
    protect-interface interface-name;  
  }  
}
```

Hierarchy Level [edit logical-systems *logical-system-name* protocols l2circuit],
[edit protocols l2circuit]

Release Information Statement introduced before Junos OS Release 7.4.

Description Configure a local switching interface. A local switching interface allows you to terminate a virtual circuit on the local router.

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 Local Interface Switching in Layer 2 Circuits on page 30](#)

mtu

Syntax	<code>mtu bytes;</code>
Hierarchy Level	<pre>[edit interfaces <i>interface-name</i>], [edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family <i>family</i>], [edit interfaces <i>interface-range name</i>], [edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family <i>family</i>], [edit logical-systems <i>logical-system-name</i> protocols l2circuit local-switching interface <i>interface-name</i> backup-neighbor <i>address</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> protocols l2circuit neighbor <i>address</i> interface <i>interface-name</i> backup-neighbor <i>address</i>], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols vpls neighbor <i>address</i> backup-neighbor <i>address</i>], [edit protocols l2circuit local-switching interface <i>interface-name</i> backup-neighbor <i>address</i>], [edit protocols l2circuit neighbor <i>address</i> interface <i>interface-name</i>], [edit protocols l2circuit neighbor <i>address</i> interface <i>interface-name</i> backup-neighbor <i>address</i>], [edit routing-instances <i>routing-instance-name</i> protocols vpls neighbor <i>address</i> backup-neighbor <i>address</i>]</pre>
Release Information	<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 12.1X48 for PTX Series Packet Transport Switches.</p> <p>Statement introduced in Junos OS Release 12.2 for ACX Series Universal Access Routers.</p>
Description	<p>Specify the maximum transmission unit (MTU) size for the media or protocol. The default MTU size depends on the device type. Changing the media MTU or protocol MTU causes an interface to be deleted and added again.</p> <p>To route jumbo data packets on the routed VLAN interface (RVI) on EX Series switches, you must configure the jumbo MTU size on the member physical interfaces and also on the RVI itself (the vlan interface).</p>



CAUTION: For EX Series switches, setting or deleting the jumbo MTU size on the RVI (the **vlan** interface) while the switch is transmitting packets might cause packets to be dropped.



NOTE: If a packet whose size is larger than the configured MTU size is received on the receiving interface, the packet is eventually dropped. The value considered for MRU (maximum receive unit) size is also the same as the MTU size configured on that interface.



NOTE: Not all devices allow you to set an MTU value, and some devices have restrictions on the range of allowable MTU values. You cannot configure an MTU for management Ethernet interfaces (fxp0, em0, or me0) or for loopback, multilink, and multicast tunnel devices.

For more information about configuring MTU for specific interfaces and router or switch combinations, see [Configuring the Media MTU](#).

Options *bytes*—MTU size.

Range: 256 through 9192 bytes, 256 through 9500 bytes (Junos OS 12.1X48R2 for PTX Series systems)

Default: 1500 bytes (INET, INET6, and ISO families), 1448 bytes (MPLS), 1514 bytes (EX Series switch interfaces)

Required Privilege interface—To view this statement in the configuration.

Level interface-control—To add this statement to the configuration.

**Related
Documentation**

- [Configuring Gigabit Ethernet Interfaces \(CLI Procedure\)](#)
- [Configuring the MTU Advertised for a Layer 2 Circuit on page 36](#)
- [Configuring the Media MTU](#)
- [Configuring Routed VLAN Interfaces \(CLI Procedure\)](#)
- [Setting the Protocol MTU](#)

neighbor (Protocols Layer 2 Circuit)

Syntax	<pre> neighbor address { interface interface-name { backup-neighbor address { community name; psn-tunnel-endpoint address; standby; static; virtual-circuit-id number; } bandwidth (bandwidth ctnumber bandwidth); community community-name; (control-word no-control-word); description text; ignore-encapsulation-mismatch; ignore-mtu-mismatch; mtu mtu-number; no-revert; protect-interface interface-name; pseudowire-status-tlv; psn-tunnel-endpoint address; revert-time seconds; static { incoming-label label; outgoing-label label; } switchover-delay milliseconds; virtual-circuit-id identifier; } } </pre>
Hierarchy Level	<p>[edit logical-systems <i>logical-system-name</i> protocols l2circuit], [edit protocols l2circuit]</p>
Release Information	<p>Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 11.1 for EX Series switches.</p>
Description	<p>Each Layer 2 circuit is represented by the logical interface connecting the local provider edge (PE) router or switch to the local customer edge (CE) router or switch. All the Layer 2 circuits using a particular remote PE router or switch designated for remote CE routers or switches are listed under the neighbor statement (neighbor designates the PE router or switch). Each neighbor is identified by its IP address and is usually the end-point destination for the LSP tunnel (transporting the Layer 2 circuit).</p>
Options	<p>address—IP address of a neighboring router or switch.</p> <p>The remaining statements are explained separately.</p>
Required Privilege Level	<p>routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.</p>

- Related Documentation**
- [Configuring the Neighbor Interface for the Layer 2 Circuit on page 32](#)

no-revert (Local Switching)

Syntax	no-revert;
Hierarchy Level	[edit protocols l2circuit local-switching interfaces <i>interface-name</i>] [edit protocols l2circuit local-switching interfaces <i>interface-name</i> end-interface <i>interface-name</i>]
Release Information	Statement introduced in Junos OS Release 11.4.
Description	(Optional) Prevent the local switching interface from reverting to the primary interface.



NOTE: The protect interface must be configured prior to configuring the **no-revert** statement.

Typically, when the primary interface goes down, the pseudowire starts using the protect interface. By default, when the primary interface comes back online, the interface is switched-over back from the protect interface to the primary interface. To prevent the switchover back to the primary interface, unless the primary interface goes down, include the **no-revert** statement. This prevents loss of traffic during the switchover.



NOTE: If the protect interface fails, the interface is switched-over back to the primary interface, irrespective of whether or not the **no-revert** statement is included in the configuration.

This statement can be configured both for the starting interface and the ending interface.

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

- Related Documentation**
- [Configuring Local Interface Switching in Layer 2 Circuits on page 30](#)

no-revert (Neighbor Interface)

Syntax	no-revert;
Hierarchy Level	[edit protocols l2circuit neighbor <i>address</i> interfaces <i>interface-name</i>], [edit logical-systems <i>logical-system-name</i> protocols l2circuit neighbor <i>address</i> interfaces <i>interface-name</i>]
Release Information	Statement introduced in Junos OS Release 11.3.
Description	(Optional) Prevent the protect interface from reverting to the primary interface.



NOTE: The protect interface must be configured prior to configuring the **no-revert** statement.

Typically, when the primary interface goes down, the pseudowire starts using the protect interface. By default, when the primary interface comes back online, the interface is switched-over back from the protect interface to the primary interface. To prevent the switchover back to the primary interface, unless the protect interface goes down, include the **no-revert** statement. This prevents loss of traffic during the switchover.



NOTE: If the protect interface fails, the interface is switched-over back to the primary interface, irrespective of whether or not the **no-revert** statement is included in the configuration.

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"> • Configuring Interfaces for Layer 2 Circuits on page 32

ping-interval

Syntax	ping-interval;
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols l2circuit neighbor <i>address</i> interface <i>interface-name</i> oam], [edit logical-systems <i>logical-system-name</i> routing-instances <i>instance-name</i> protocols l2vpn oam], [edit logical-systems <i>logical-system-name</i> routing-instances <i>instance-name</i> protocols vpls neighbor <i>address</i> oam], [edit logical-systems <i>logical-system-name</i> routing-instances <i>instance-name</i> protocols vpls mesh-group <i>mesh-group-name</i> neighbor <i>address</i> oam], [edit logical-systems <i>logical-system-name</i> routing-instances <i>instance-name</i> protocols vpls oam], [edit protocols l2circuit neighbor <i>address</i> interface <i>interface-name</i> oam], [edit routing-instances <i>instance-name</i> protocols l2vpn oam], [edit routing-instances <i>instance-name</i> protocols vpls neighbor <i>address</i> oam], [edit routing-instances <i>instance-name</i> protocols vpls mesh-group <i>mesh-group-name</i> neighbor <i>address</i> oam], [edit routing-instances <i>instance-name</i> protocols vpls oam]
Release Information	Statement introduced in Junos OS Release 10.0. Support for FEC 129 VPLS added in Junos OS Release 12.2.
Description	Configure the time interval between ping messages for bidirectional forwarding detection (BFD) sessions enabled over pseudowires inside a VPN.
Options	<i>seconds</i> —Time interval between ping messages. Range: 30 through 3600
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">Configuring BFD for VCCV for Layer 2 VPNs, Layer 2 Circuits, and VPLS in the Junos OS VPNs Configuration Guide

protect-interface

Syntax	<code>protect-interface <i>interface-name</i>;</code>
Hierarchy Level	<p>[edit logical-systems <i>logical-system-name</i> protocols l2circuit neighbor <i>address</i> interface <i>interface-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols l2circuit local-switching interface <i>interface-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols l2circuit local-switching interface <i>interface-name</i> end-interface],</p> <p>[edit protocols l2circuit local-switching interface <i>interface-name</i>],</p> <p>[edit protocols l2circuit neighbor <i>address</i> interface <i>interface-name</i>],</p> <p>[edit protocols l2circuit local-switching interface <i>interface-name</i> end-interface]</p>
Release Information	Statement introduced before Junos OS Release 7.4.
Description	Provide a backup for the protected interface in case of failure. Network traffic uses the primary interface only, as long as the primary interface functions.
Options	<i>interface-name</i> —Name of the protect interface to configure.
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 the Protect Interface on page 37

protected-l2circuit

Syntax	<pre>protected-l2circuit { egress-pe <i>address</i>; ingress-pe <i>address</i>; virtual-circuit-id <i>identifier</i>; }</pre>
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols l2circuit neighbor <i>address</i> interface <i>interface-name</i> egress-protection], [edit protocols l2circuit neighbor <i>address</i> interface <i>interface-name</i> egress-protection]
Release Information	Statement introduced in Junos OS release 10.4.
Description	Configures the protected Layer 2 circuit as part of an egress protection virtual circuit (EPVC).
Options	<p>egress-pe <i>address</i>—Specify the address of the egress PE router for the protected Layer 2 circuit.</p> <p>ingress-pe <i>address</i>—Specify the address of the ingress PE router for the protected Layer 2 circuit.</p> <p>virtual-circuit-id <i>identifier</i>—Specify the virtual circuit identifier for the protected Layer 2 circuit.</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">• Example: Configuring an Egress Protection LSP for a Layer 2 Circuit on page 70

protector-interface

Syntax	<code>protector-interface <i>interface-name</i>;</code>
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols l2circuit neighbor <i>address</i> interface <i>interface-name</i> egress-protection], [edit protocols l2circuit neighbor <i>address</i> interface <i>interface-name</i> egress-protection]
Release Information	Statement introduced in Junos OS release 10.4.
Description	Configures the protector interface for an egress protection LSP.
Options	<i>interface-name</i> —Name of the interface used to protect traffic for an egress protection LSP.
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"> • Example: Configuring an Egress Protection LSP for a Layer 2 Circuit on page 70

protector-pe

Syntax	<code>protector-pe <i>address</i> { context-identifier <i>identifier</i>; lsp <i>lsp-name</i>; }</code>
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols l2circuit neighbor <i>address</i> interface <i>interface-name</i> egress-protection], [edit protocols l2circuit neighbor <i>address</i> interface <i>interface-name</i> egress-protection]
Release Information	Statement introduced in Junos OS release 10.4.
Description	Configures the protector PE router for an egress protection LSP. Test.
Options	<i>address</i> —IPv4 address for the protector PE router. <i>context-identifier identifier</i> —Identifies the context for the egress protection LSP. <i>lsp lsp-name</i> —Specifies the LSP for the egress protection LSP.
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"> • Example: Configuring an Egress Protection LSP for a Layer 2 Circuit on page 70

pseudowire-status-tlv

Syntax	pseudowire-status-tlv;
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols l2circuit neighbor <i>address</i> 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 vpls mesh-group <i>mesh-group-name</i> neighbor <i>address</i>], [edit routing-instances <i>routing-instance-name</i> protocols vpls neighbor <i>address</i>]
Release Information	Statement introduced in Junos OS Release 10.0.
Description	Enables the pseudowire type length variable (TLV). The pseudowire status TLV is used to communicate the status of a pseudowire back and forth between two PE routers. The pseudowire status TLV is configurable for each pseudowire connection and is disabled by default.
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">• Configuring the Pseudowire Status TLV on page 38

psn-tunnel-endpoint

Syntax	<code>psn-tunnel-endpoint <i>address</i>;</code>
Hierarchy Level	<p>[edit logical-systems <i>logical-system-name</i> protocols l2circuit neighbor <i>address</i> interface <i>interface-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols l2circuit neighbor <i>address</i> interface <i>interface-name</i> backup-neighbor <i>address</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols vpls neighbor <i>address</i> backup-neighbor <i>address</i>],</p> <p>[edit protocols l2circuit neighbor <i>address</i> interface <i>interface-name</i>],</p> <p>[edit protocols l2circuit neighbor <i>address</i> interface <i>interface-name</i> backup-neighbor <i>address</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols vpls neighbor <i>address</i> backup-neighbor <i>address</i>]</p>
Release Information	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Hierarchy levels associated with the backup-neighbor statement added in Junos OS Release 9.2.</p>
Description	Specify the endpoint of the packet switched network (PSN) tunnel on the remote PE router.
Options	<i>address</i> —Address for the tunnel endpoint.
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 Layer 2 Circuits over Both RSVP and LDP LSPs on page 38 • Configuring Redundant Pseudowires for Layer 2 Circuits and VPLS

standby (Protocols Layer 2 Circuit)

Syntax	standby;
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols l2circuit local-switching interface <i>interface-name</i> end-interface interface <i>interface-name</i> backup-neighbor <i>address</i>], [edit logical-systems <i>logical-system-name</i> protocols l2circuit neighbor <i>address</i> interface <i>interface-name</i> backup-neighbor <i>address</i>], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols vpls neighbor <i>address</i> backup-neighbor <i>address</i>], [edit protocols l2circuit neighbor <i>address</i> interface <i>interface-name</i> backup-neighbor <i>address</i>], [edit routing-instances <i>routing-instance-name</i> protocols vpls neighbor <i>address</i> backup-neighbor <i>address</i>]
Release Information	Statement introduced before Junos OS Release 7.4. Hierarchy levels associated with the backup-neighbor statement added in Junos OS Release 9.2.
Description	Configure the pseudowire to the specified backup neighbor as the standby. When you configure this statement, traffic flows over both the active and standby pseudowires to the backup device (either a CE device or PE router). The backup device drops the traffic from the standby pseudowire, unless the active pseudowire fails. If the active pseudowire fails, the backup device automatically switches to the standby pseudowire.
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">• Configuring Layer 2 Circuits over Both RSVP and LDP LSPs on page 38• Configuring Redundant Pseudowires for Layer 2 Circuits and VPLS• Example: Configuring Layer 2 Circuit Switching Protection on page 58

static (Protocols Layer 2 Circuit)

Syntax	<pre>static { incoming-label <i>label</i>; outgoing-label <i>label</i>; send-oam }</pre>
Hierarchy Level	<pre>[edit logical-systems <i>logical-system-name</i> l2circuit neighbor <i>address</i> interface <i>interface-name</i>], [edit logical-systems <i>logical-system-name</i> l2circuit neighbor <i>address</i> interface <i>interface-name</i> backup-neighbor <i>neighbor</i>], [edit protocols l2circuit neighbor <i>address</i> interface <i>interface-name</i>], [edit protocols l2circuit neighbor <i>address</i> interface <i>interface-name</i> backup-neighbor <i>neighbor</i>]</pre>
Release Information	<p>Statement introduced in Junos OS Release 9.5.</p> <p>Statement introduced in Junos OS Release 12.1X48 for PTX Series switches.</p> <p>Statement introduced in Junos OS Release 12.2 for ACX Series routers.</p>
Description	<p>Configures static Layer 2 circuit pseudowires. Static pseudowires are designed for networks that do not support LDP or do not have LDP enabled. You configure a static pseudowire by configuring static values for the in and out labels needed to enable a pseudowire connection.</p>
Options	<p>incoming-label—(Optional for PTX Series Packet Transport Switches only) Configure the Layer 2 circuit incoming static pseudowire label.</p> <p>Range: 1000000 through 1048575</p> <p>outgoing-label—(Optional for PTX Series Packet Transport Switches only) Configure the Layer 2 circuit outgoing static pseudowire label.</p> <p>Range: 16 through 1048575</p> <p>The remaining statement is 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"> • Configuring Static Layer 2 Circuits on page 40

traceoptions (Protocols Layer 2 Circuit)

Syntax	<pre>traceoptions { file <i>filename</i> <files <i>number</i>> <size <i>size</i>> <world-readable no-world-readable>; flag <i>flag</i> <<i>flag-modifier</i>> <disable>; }</pre>
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols l2circuit], [edit protocols l2circuit]
Release Information	Statement introduced before Junos OS Release 7.4.
Description	Trace traffic flowing through a Layer 2 circuit.
Options	<p>disable—(Optional) Disable the tracing operation. You can use this option to disable a single operation when you have defined a broad group of tracing operations, such as all.</p> <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 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.</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.</p> <ul style="list-style-type: none">• connections—Layer 2 circuit connections (events and state changes)• error—Error conditions• fec—Layer 2 circuit advertisements received or sent by means of LDP• topology—Layer 2 circuit topology changes caused by reconfiguration or advertisements received from other PE routers <p>flag-modifier—(Optional) Modifier for the tracing flag. You can specify the detail modifier if you want to provide detailed trace information.</p> <p>no-world-readable—(Optional) Do not allow any user to read the log file.</p> <p>size <i>size</i>—(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 the trace-file again reaches its maximum size, trace-file.0 is</p>

renamed *trace-file.1* and *trace-file* is renamed *trace-file.0*. This renaming scheme continues until the maximum number of trace files 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—(Optional) 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 Layer 2 Circuit Operations on page 171

virtual-circuit-id

Syntax	<code>virtual-circuit-id <i>identifier</i>;</code>
Hierarchy Level	<code>[edit logical-systems <i>logical-system-name</i> protocols l2circuit local-switching interface <i>interface-name</i> backup-neighbor <i>address</i>],</code> <code>[edit logical-systems <i>logical-system-name</i> protocols l2circuit neighbor <i>address</i> interface <i>interface-name</i>],</code> <code>[edit logical-systems <i>logical-system-name</i> protocols l2circuit neighbor <i>address</i> interface <i>interface-name</i> backup-neighbor <i>address</i>],</code> <code>[edit protocols l2circuit local-switching interface <i>interface-name</i> backup-neighbor <i>address</i>],</code> <code>[edit protocols l2circuit neighbor <i>address</i> interface <i>interface-name</i>],</code> <code>[edit protocols l2circuit neighbor <i>address</i> interface <i>interface-name</i> backup-neighbor <i>address</i>]</code>
Release Information	Statement introduced before Junos OS Release 7.4. Hierarchy levels for backup-neighbor (pseudowire redundancy) added in Junos OS Release 9.2. Statement introduced in Junos OS Release 11.1 for EX Series switches. Statement introduced in Junos OS Release 12.3 at the <code>[edit protocols l2circuit local-switching interface <i>interface-name</i> backup-neighbor <i>address</i>]</code> hierarchy level.
Description	Uniquely identify a Layer 2 circuit for either a standard pseudowire or a redundant pseudowire.
Options	<i>identifier</i> —1 through 4,294,967,295
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">• Configuring the Virtual Circuit ID on page 39• Configuring Redundant Pseudowires for Layer 2 Circuits and VPLS• Example: Configuring Layer 2 Circuit Switching Protection on page 58

PART 4

Troubleshooting

- [Troubleshooting Layer 2 Circuits on page 171](#)

Troubleshooting Layer 2 Circuits

- [Tracing Layer 2 Circuit Operations on page 171](#)

Tracing Layer 2 Circuit Operations

To trace the creation of and changes to Layer 2 circuits, include the **traceoptions** statement:

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

You can include this statement at the following hierarchy levels:

- **[edit protocols l2circuit]**
- **[edit logical-systems *logical-system-name* protocols l2circuit]**

Specify the following flags to trace the indicated operations on Layer 2 circuits:

- **connections**—Layer 2 circuit connections (events and state changes)
- **error**—Error conditions
- **FEC**—Layer 2 circuit advertisements received or sent using LDP
- **topology**—Layer 2 circuit topology changes caused by reconfiguration or advertisements received from other PE routers

PART 5

Index

- [Index on page 175](#)

Index

Symbols

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

A

ATM trunking.....	7
Layer 2 circuits.....	47
automatic route distinguisher.....	23
autonomous system number	
route distinguisher.....	23

B

backup-interface.....	132
backup-neighbor statement.....	133
bandwidth accounting.....	4
bandwidth statement.....	131
Layer 2 circuits	
usage guidelines.....	48
braces, in configuration statements.....	xiv
brackets	
angle, in syntax descriptions.....	xiv
square, in configuration statements.....	xiv

C

CAC.....	4, 48
fast reroute.....	6
link and node protection.....	6
LSP path protection.....	5
secondary paths.....	6
call admission control See CAC	
CCC	
Frame Relay, control word.....	34
comments, in configuration statements.....	xiv
communities	
regular expressions.....	26

community statement.....	134
Layer 2 circuits	
usage guidelines.....	33, 41
connection-protection.....	135
control word, Frame Relay.....	34
control-word statement	
Layer 2 circuits.....	135
usage guidelines.....	35
conventions	
text and syntax.....	xiii
curly braces, in configuration statements.....	xiv
customer support.....	xv
contacting JTAC.....	xv

D

description statement.....	136
documentation	
comments on.....	xv

E

egress protection LSP, configuring.....	70
egress-protection statement	
Layer 2 circuits.....	136
MPLS.....	137
encapsulation mismatch	
Layer 2 circuits.....	35
encapsulation statement	
physical interface.....	138
encapsulation-type statement	
Layer 2 circuits.....	143
usage guidelines.....	35
end-interface statement.....	144
usage guidelines.....	30
example	
egress protection LSP.....	70
export policy, VRF.....	27

F

family inet-vpn.....	14
family inet6-vpn.....	14
family l2vpn.....	14
fast reroute, CAC.....	6
fast-aps-switch	
configuring.....	49
fast-aps-switch statement.....	145
font conventions.....	xiii

I

ignore-encapsulation-mismatch statement.....	146
usage guidelines.....	35
ignore-mtu-mismatch statement.....	146
usage guidelines.....	31, 36
import communities	
regular expressions.....	26
import policy, VRF.....	25
install-nexthop statement.....	147
usage guidelines.....	42
instance-type statement	
usage guidelines.....	20
interface statement	
Layer 2 circuits.....	148
usage guidelines.....	32
VPNs	
usage guidelines.....	20

L

l2circuit statement.....	150
l2vpn-use-bgp-rules statement.....	151
Layer 2 circuits	
CAC.....	4
egress protection LSP.....	70
supported software standards.....	129
Layer 2 VPNs	
path selection.....	45
route distinguisher.....	22
Layer 2 VPNs, multihoming.....	45
Layer 2 circuit	
MTU.....	36
Layer 2 circuits	
ATM trunking.....	7, 47
bandwidth accounting.....	4
call admission control.....	4
encapsulation mismatch.....	35
local interface switching.....	30
static pseudowires.....	40
trunk mode.....	7
link and node protection, CAC.....	6
local interface switching.....	30
local switching interface	
MTU.....	31
local-switching statement	
Layer 2 circuits.....	152
usage guidelines.....	30

M

manuals	
comments on.....	xv
MTU	
Layer 2 circuit.....	36
local switching interface.....	31
mtu statement.....	153
usage guidelines.....	36
multihoming	
path selection.....	45

N

neighbor statement.....	155
no-control-word statement	
Layer 2 circuits	
usage guidelines.....	35

P

parentheses, in syntax descriptions.....	xiv
path selection, Layer 2 VPNs and VPLS.....	45
ping-interval statement.....	158
protect-interface statement.....	159
protected-l2circuit statement.....	160
protector-interface statement.....	161
protector-pe statement.....	161
pseudowire-status-tlv statement.....	162
usage guidelines.....	38
pseudowires	
static.....	40
psn-tunnel-endpoint statement.....	163
usage guidelines.....	38

R

reducing aps switchover time.....	49
layer 2 circuits	
configuring.....	49
regular expressions, import communities.....	26
route distinguisher.....	22
automatic.....	23
autonomous system number.....	23
route-distinguisher statement	
usage guidelines.....	22
route-distinguisher-id statement	
usage guidelines.....	23
routing engine, sampling.....	20
routing instance name.....	19
routing instance type.....	20

S

standby statement.....	164
static pseudowires	
Layer 2 circuits.....	40
static statement	
Layer 2 circuits.....	165
usage guidelines.....	40
support, technical See technical support	
syntax conventions.....	xiii

T

technical support	
contacting JTAC.....	xv
traceoptions statement	
protocols Layer 2 circuit.....	166
trunk mode.....	7

V

virtual-circuit-id statement.....	168
VPLS	
path selection.....	45
route distinguisher.....	22
VPLS, multihoming.....	45
VPNs	
export policy.....	27
import policy.....	25
interfaces.....	20
route distinguisher.....	22
automatic.....	23
routing instance name.....	19
unicast RPF.....	22
VRF	
export policy.....	27
import policy.....	25
regular expressions.....	26

