



Junos[®] OS

Layer 2 Circuits Feature Guide

Release
11.2



Published: 2011-05-17

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 Feature Guide
Release 11.2
Copyright © 2011, Juniper Networks, Inc.
All rights reserved.

Revision History
April 2011—R1 Junos OS 11.2

The information in this document is current as of the date listed in the revision history.

YEAR 2000 NOTICE

Juniper Networks hardware and software products are Year 2000 compliant. The 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

READ THIS END USER LICENSE AGREEMENT ("AGREEMENT") BEFORE DOWNLOADING, INSTALLING, OR USING THE SOFTWARE. BY DOWNLOADING, INSTALLING, OR USING THE SOFTWARE OR OTHERWISE EXPRESSING YOUR AGREEMENT TO THE TERMS CONTAINED HEREIN, YOU (AS CUSTOMER OR IF YOU ARE NOT THE CUSTOMER, AS A REPRESENTATIVE/AGENT AUTHORIZED TO BIND THE CUSTOMER) CONSENT TO BE BOUND BY THIS AGREEMENT. IF YOU DO NOT OR CANNOT AGREE TO THE TERMS CONTAINED HEREIN, THEN (A) DO NOT DOWNLOAD, INSTALL, OR USE THE SOFTWARE, AND (B) YOU MAY CONTACT JUNIPER NETWORKS REGARDING LICENSE TERMS.

1. **The Parties.** The parties to this Agreement are (i) Juniper Networks, Inc. (if the Customer's principal office is located in the Americas) or Juniper Networks (Cayman) Limited (if the Customer's principal office is located outside the Americas) (such applicable entity being referred to herein as "Juniper"), and (ii) the person or organization that originally purchased from Juniper or an authorized Juniper reseller the applicable license(s) for use of the Software ("Customer") (collectively, the "Parties").

2. **The Software.** In this Agreement, "Software" means the program modules and features of the Juniper or Juniper-supplied software, for which Customer has paid the applicable license or support fees to Juniper or an authorized Juniper reseller, or which was embedded by Juniper in equipment which Customer purchased from Juniper or an authorized Juniper reseller. "Software" also includes updates, upgrades and new releases of such software. "Embedded Software" means Software which Juniper has embedded in or loaded onto the Juniper equipment and any updates, upgrades, additions or replacements which are subsequently embedded in or loaded onto the equipment.

3. **License Grant.** Subject to payment of the applicable fees and the limitations and restrictions set forth herein, Juniper grants to Customer a non-exclusive and non-transferable license, without right to sublicense, to use the Software, in executable form only, subject to the following use restrictions:

- a. Customer shall use Embedded Software solely as embedded in, and for execution on, Juniper equipment originally purchased by Customer from Juniper or an authorized Juniper reseller.
- b. Customer shall use the Software on a single hardware chassis having a single processing unit, or as many chassis or processing units for which Customer has paid the applicable license fees; provided, however, with respect to the Steel-Belted Radius or Odyssey Access Client software only, Customer shall use such Software on a single computer containing a single physical random access memory space and containing any number of processors. Use of the Steel-Belted Radius or IMS AAA software on multiple computers or virtual machines (e.g., Solaris zones) requires multiple licenses, regardless of whether such computers or virtualizations are physically contained on a single chassis.
- c. Product purchase documents, paper or electronic user documentation, and/or the particular licenses purchased by Customer may specify limits to Customer's use of the Software. Such limits may restrict use to a maximum number of seats, registered endpoints, concurrent users, sessions, calls, connections, subscribers, clusters, nodes, realms, devices, links, ports or transactions, or require the purchase of separate licenses to use particular features, functionalities, services, applications, operations, or capabilities, or provide throughput, performance, configuration, bandwidth, interface, processing, temporal, or geographical limits. In addition, such limits may restrict the use of the Software to managing certain kinds of networks or require the Software to be used only in conjunction with other specific Software. Customer's use of the Software shall be subject to all such limitations and purchase of all applicable licenses.
- d. For any trial copy of the Software, Customer's right to use the Software expires 30 days after download, installation or use of the Software. Customer may operate the Software after the 30-day trial period only if Customer pays for a license to do so. Customer may not extend or create an additional trial period by re-installing the Software after the 30-day trial period.
- e. The Global Enterprise Edition of the Steel-Belted Radius software may be used by Customer only to manage access to Customer's enterprise network. Specifically, service provider customers are expressly prohibited from using the Global Enterprise Edition of the Steel-Belted Radius software to support any commercial network access services.

The foregoing license is not transferable or assignable by Customer. No license is granted herein to any user who did not originally purchase the applicable license(s) for the Software from Juniper or an authorized Juniper reseller.

4. **Use Prohibitions.** Notwithstanding the foregoing, the license provided herein does not permit the Customer to, and Customer agrees not to and shall not: (a) modify, unbundle, reverse engineer, or create derivative works based on the Software; (b) make unauthorized copies of the Software (except as necessary for backup purposes); (c) rent, sell, transfer, or grant any rights in and to any copy of the Software, in any form, to any third party; (d) remove any proprietary notices, labels, or marks on or in any copy of the Software or any product in which the Software is embedded; (e) distribute any copy of the Software to any third party, including as may be embedded in Juniper equipment sold in the secondhand market; (f) use any 'locked' or key-restricted feature, function, service, application, operation, or capability without first purchasing the applicable license(s) and obtaining a valid key from Juniper, even if such feature, function, service, application, operation, or capability is enabled without a key; (g) distribute any key for the Software provided by Juniper to any third party; (h) use the

Software in any manner that extends or is broader than the uses purchased by Customer from Juniper or an authorized Juniper reseller; (i) use Embedded Software on non-Juniper equipment; (j) use Embedded Software (or make it available for use) on Juniper equipment that the Customer did not originally purchase from Juniper or an authorized Juniper reseller; (k) disclose the results of testing or benchmarking of the Software to any third party without the prior written consent of Juniper; or (l) use the Software in any manner other than as expressly provided herein.

5. **Audit.** Customer shall maintain accurate records as necessary to verify compliance with this Agreement. Upon request by Juniper, Customer shall furnish such records to Juniper and certify its compliance with this Agreement.

6. **Confidentiality.** The Parties agree that aspects of the Software and associated documentation are the confidential property of Juniper. As such, Customer shall exercise all reasonable commercial efforts to maintain the Software and associated documentation in confidence, which at a minimum includes restricting access to the Software to Customer employees and contractors having a need to use the Software for Customer's internal business purposes.

7. **Ownership.** Juniper and Juniper's licensors, respectively, retain ownership of all right, title, and interest (including copyright) in and to the Software, associated documentation, and all copies of the Software. Nothing in this Agreement constitutes a transfer or conveyance of any right, title, or interest in the Software or associated documentation, or a sale of the Software, associated documentation, or copies of the Software.

8. **Warranty, Limitation of Liability, Disclaimer of Warranty.** The warranty applicable to the Software shall be as set forth in the warranty statement that accompanies the Software (the "Warranty Statement"). Nothing in this Agreement shall give rise to any obligation to support the Software. Support services may be purchased separately. Any such support shall be governed by a separate, written support services agreement. TO THE MAXIMUM EXTENT PERMITTED BY LAW, JUNIPER SHALL NOT BE LIABLE FOR ANY LOST PROFITS, LOSS OF DATA, OR COSTS OR PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES, OR FOR ANY SPECIAL, INDIRECT, OR CONSEQUENTIAL DAMAGES ARISING OUT OF THIS AGREEMENT, THE SOFTWARE, OR ANY JUNIPER OR JUNIPER-SUPPLIED SOFTWARE. IN NO EVENT SHALL JUNIPER BE LIABLE FOR DAMAGES ARISING FROM UNAUTHORIZED OR IMPROPER USE OF ANY JUNIPER OR JUNIPER-SUPPLIED SOFTWARE. EXCEPT AS EXPRESSLY PROVIDED IN THE WARRANTY STATEMENT TO THE EXTENT PERMITTED BY LAW, JUNIPER DISCLAIMS ANY AND ALL WARRANTIES IN AND TO THE SOFTWARE (WHETHER EXPRESS, IMPLIED, STATUTORY, OR OTHERWISE), INCLUDING ANY IMPLIED WARRANTY OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE, OR NONINFRINGEMENT. IN NO EVENT DOES JUNIPER WARRANT THAT THE SOFTWARE, OR ANY EQUIPMENT OR NETWORK RUNNING THE SOFTWARE, WILL OPERATE WITHOUT ERROR OR INTERRUPTION, OR WILL BE FREE OF VULNERABILITY TO INTRUSION OR ATTACK. In no event shall Juniper's or its suppliers' or licensors' liability to Customer, whether in contract, tort (including negligence), breach of warranty, or otherwise, exceed the price paid by Customer for the Software that gave rise to the claim, or if the Software is embedded in another Juniper product, the price paid by Customer for such other product. Customer acknowledges and agrees that Juniper has set its prices and entered into this Agreement in reliance upon the disclaimers of warranty and the limitations of liability set forth herein, that the same reflect an allocation of risk between the Parties (including the risk that a contract remedy may fail of its essential purpose and cause consequential loss), and that the same form an essential basis of the bargain between the Parties.

9. **Termination.** Any breach of this Agreement or failure by Customer to pay any applicable fees due shall result in automatic termination of the license granted herein. Upon such termination, Customer shall destroy or return to Juniper all copies of the Software and related documentation in Customer's possession or control.

10. **Taxes.** All license fees payable under this agreement are exclusive of tax. Customer shall be responsible for paying Taxes arising from the purchase of the license, or importation or use of the Software. If applicable, valid exemption documentation for each taxing jurisdiction shall be provided to Juniper prior to invoicing, and Customer shall promptly notify Juniper if their exemption is revoked or modified. All payments made by Customer shall be net of any applicable withholding tax. Customer will provide reasonable assistance to Juniper in connection with such withholding taxes by promptly: providing Juniper with valid tax receipts and other required documentation showing Customer's payment of any withholding taxes; completing appropriate applications that would reduce the amount of withholding tax to be paid; and notifying and assisting Juniper in any audit or tax proceeding related to transactions hereunder. Customer shall comply with all applicable tax laws and regulations, and Customer will promptly pay or reimburse Juniper for all costs and damages related to any liability incurred by Juniper as a result of Customer's non-compliance or delay with its responsibilities herein. Customer's obligations under this Section shall survive termination or expiration of this Agreement.

11. **Export.** Customer agrees to comply with all applicable export laws and restrictions and regulations of any United States and any applicable foreign agency or authority, and not to export or re-export the Software or any direct product thereof in violation of any such restrictions, laws or regulations, or without all necessary approvals. Customer shall be liable for any such violations. The version of the Software supplied to Customer may contain encryption or other capabilities restricting Customer's ability to export the Software without an export license.

12. **Commercial Computer Software.** The Software is "commercial computer software" and is provided with restricted rights. Use, duplication, or disclosure by the United States government is subject to restrictions set forth in this Agreement and as provided in DFARS 227.7201 through 227.7202-4, FAR 12.212, FAR 27.405(b)(2), FAR 52.227-19, or FAR 52.227-14 (ALT III) as applicable.

13. **Interface Information.** To the extent required by applicable law, and at Customer's written request, Juniper shall provide Customer with the interface information needed to achieve interoperability between the Software and another independently created program, on payment of applicable fee, if any. Customer shall observe strict obligations of confidentiality with respect to such information and shall use such information in compliance with any applicable terms and conditions upon which Juniper makes such information available.

14. **Third Party Software.** Any licensor of Juniper whose software is embedded in the Software and any supplier of Juniper whose products or technology are embedded in (or services are accessed by) the Software shall be a third party beneficiary with respect to this Agreement, and such licensor or vendor shall have the right to enforce this Agreement in its own name as if it were Juniper. In addition, certain third party software may be provided with the Software and is subject to the accompanying license(s), if any, of its respective owner(s). To the extent portions of the Software are distributed under and subject to open source licenses obligating Juniper to make the source code for such portions publicly available (such as the GNU General Public License ("GPL") or the GNU Library General Public License ("LGPL")), Juniper will make such source code portions (including Juniper modifications, as appropriate) available upon request for a period of up to three years from the date of distribution. Such request can be made in writing to Juniper Networks, Inc., 1194 N. Mathilda Ave., Sunnyvale, CA 94089, ATTN: General Counsel. You may obtain a copy of the GPL at <http://www.gnu.org/licenses/gpl.html>, and a copy of the LGPL at <http://www.gnu.org/licenses/lgpl.html>.

15. **Miscellaneous.** This Agreement shall be governed by the laws of the State of California without reference to its conflicts of laws principles. The provisions of the U.N. Convention for the International Sale of Goods shall not apply to this Agreement. For any disputes arising under this Agreement, the Parties hereby consent to the personal and exclusive jurisdiction of, and venue in, the state and federal courts within Santa Clara County, California. This Agreement constitutes the entire and sole agreement between Juniper and the Customer with respect to the Software, and supersedes all prior and contemporaneous agreements relating to the Software, whether oral or written (including any inconsistent terms contained in a purchase order), except that the terms of a separate written agreement executed by an authorized Juniper representative and Customer shall govern to the extent such terms are inconsistent or conflict with terms contained herein. No modification to this Agreement nor any waiver of any rights hereunder shall be effective unless expressly assented to in writing by the party to be charged. If any portion of this Agreement is held invalid, the Parties agree that such invalidity shall not affect the validity of the remainder of this Agreement. This Agreement and associated documentation has been written in the English language, and the Parties agree that the English version will govern. (For Canada: Les parties aux présentes confirment leur volonté que cette convention de même que tous les documents y compris tout avis qui s'y rattache, soient rédigés en langue anglaise. (Translation: The parties confirm that this Agreement and all related documentation is and will be in the English language)).

Table of Contents

Part 1	Layer 2 Circuits Routing	
Chapter 1	Layer 2 Circuits Concept and Reference Materials	3
	Layer 2 Circuits Overview	3
	Layer 2 Circuits System Requirements	6
	Layer 2 Circuits Terms and Acronyms	7
Chapter 2	Layer 2 Circuits Configuration	9
	Configuring CCC Encapsulation on CE-Facing Ethernet Interfaces	9
	Configuring CCC Encapsulation on CE-Facing SONET/SDH Interfaces	10
	Configuring a CCC Encapsulation and a Layer 2 Circuit Mode on CE-Facing ATM2 IQ Interfaces	11
	Configuring the MPLS Family on Core Interfaces	13
	Configuring the Layer 2 Circuit Neighbor Address and Virtual Circuit Identifier . . .	13
	Configuring LDP and an IGP to Transport Layer 2 Circuits	14
	Option: Applying Traffic Engineering to a Layer 2 Circuit	15
	Option: Mapping Layer 2 Protocol Control Information into a Layer 2 Circuit . . .	16
	Option: Configuring APS for Layer 2 Circuits	16
	Option: Configuring Layer 2 Circuit Trunk Mode on ATM2 IQ Interfaces	17
	Option: Reserving LSP Bandwidth for a Layer 2 Circuit	19
	Option: Selecting an MTU for a Layer 2 Circuit	20
	Option: Configuring Local Interface Switching for a Layer 2 Circuit	21
	Option: Configuring Layer 2 Circuits Simultaneously over RSVP and LDP LSPs	21
Chapter 3	Layer 2 Circuits Configuration Examples	23
	Example: Ethernet-Based Layer 2 Circuit Configuration	23
	Verifying Your Work	27
	Router PE1 Status	28
	Router P0 Status	28
	Router PE2 Status	29
	Example: ATM2 IQ-Based Layer 2 Circuit Configuration	30
	Verifying Your Work	36
	Example: Layer 2 Circuit Traffic Engineering over Multiple LSPs Configuration . . .	38
	Verifying Your Work	47
	Example: APS for a Layer 2 Circuit Configuration	48
	Verifying Your Work	49
	For More Information	51
Part 2	Index	
	Index	55

List of Figures

Part 1	Layer 2 Circuits Routing	
Chapter 1	Layer 2 Circuits Concept and Reference Materials	3
	Figure 1: Layer 2 Circuit Connection	4
	Figure 2: Layer 2 Circuit Concept	4
Chapter 3	Layer 2 Circuits Configuration Examples	23
	Figure 3: Ethernet-Based Layer 2 Circuit Topology Diagram	23
	Figure 4: ATM2 IQ-Based Layer 2 Circuit Topology Diagram	30
	Figure 5: Layer 2 Circuit Traffic Engineering Topology Diagram	38
	Figure 6: APS for a Layer 2 Circuit Topology Diagram	48

PART 1

Layer 2 Circuits Routing

- Layer 2 Circuits Concept and Reference Materials on page 3
- Layer 2 Circuits Configuration on page 9
- Layer 2 Circuits Configuration Examples on page 23

CHAPTER 1

Layer 2 Circuits Concept and Reference Materials

This chapter covers these topics:


- Layer 2 Circuits Overview on page 3
- Layer 2 Circuits System Requirements on page 6
- Layer 2 Circuits Terms and Acronyms on page 7

Layer 2 Circuits Overview

After the introduction and early adoption of Layer 3 virtual private networks (VPNs) based on RFC 4364 (also known as RFC 2547bis), many customers asked their service providers to offer VPNs that would preserve data at Layer 2. One of the Layer 2 VPN options that has emerged is known as a Layer 2 circuit. It is based on a series of Internet Engineering Task Force (IETF) drafts and RFCs authored by Luca Martini. These so-called “Martini-drafts” include Internet draft draft-martini-l2circuit-encap-mpls-11.txt, *Encapsulation Methods for Transport of Layer 2 Frames Over MPLS Networks* (expires August 2006) and Internet draft draft-martini-l2circuit-trans-mpls-19.txt, *Transport of Layer 2 Frames Over MPLS* (expires September 2006), and establish the basis for Juniper Networks implementation of Layer 2 circuits. This guide shows how to implement Layer 2 circuits in a variety of ways.

Layer 2 circuits allow for the creation of point-to-point Layer 2 connections over an IP and Multiprotocol Label Switching (MPLS)-based network. Physical circuits with the same Layer 2 encapsulations can be connected together across such a network. Layer 2 circuits can allow for the replacement of end-to-end Asynchronous Transfer Mode (ATM) networks, Frame Relay networks, and some portions of Time-Division Multiplexing (TDM) networks, with an IP and MPLS-based network.

In Figure 1 on page 4, an OC3/STM1 interface encapsulated with the Point-to-Point Protocol (PPP) on Router PE1 is connected over a Layer 2 circuit to reach an OC3/STM1 interface encapsulated with PPP on Router PE2. To enable the Layer 2 circuits to operate, the provider edge (PE) routers in Figure 1 on page 4 are part of an MPLS network. Routers PE1 and PE2 must also be Label Distribution Protocol (LDP) peers. Additionally, any interface on the PE routers that connects to a customer edge (CE) router must support circuit cross-connect (CCC) interface encapsulations.

Figure 1: Layer 2 Circuit Connection



Layer 2 circuits are very similar to Layer 2 VPNs. However, there are some significant differences:

- You configure Layer 2 VPNs in a routing instance. As a result, Layer 2 VPNs have unique site and VPN identifiers. However, Layer 2 circuits do not require a routing instance configuration and instead use an alternate method of identifying circuits. Layer 2 circuit peer relationships are established with three components: a logical interface on the local PE router, the IP address of the remote PE router neighbor, and a virtual circuit identifier.
- Layer 2 VPNs, like Layer 3 VPNs, require Border Gateway Protocol (BGP) for transport of traffic between PE routers. In contrast, Layer 2 circuits do not require BGP. Instead, Layer 2 circuits rely on LDP and MPLS for their operation. As a result, Layer 2 circuits require less configuration than Layer 2 VPNs.

Layer 2 circuits are configured between two peers. The peers must use the same interior gateway protocol (IGP), such as Open Shortest Path First (OSPF) or Intermediate System-to-Intermediate System (IS-IS). Also, the peers must have a symmetrical Layer 2 configuration and belong to the same routing domain or autonomous system.

The basic building block for a Layer 2 circuit is a virtual circuit (VC). A VC is a point-to-point Layer 2 connection that is transported over MPLS or any other tunneling technology in a service provider network. A VC is similar to a CCC connection except that multiple VCs are transported over a single MPLS label-switched path (LSP) tunnel between two provider edge (PE) routers. In contrast, CCC only supports a single Layer 2 connection over a single LSP.

In Figure 2 on page 4, the basic inner workings of Layer 2 circuits are explained. Two customer edge (CE) router logical interfaces, one local and one remote, are running the same Layer 2 protocol. Packets are sent from the local CE router to the remote CE router over an egress label advertised by the remote PE router. The label is transported over an LDP LSP (or LDP tunneled through RSVP) to the remote PE router that is connected to the remote CE router. Return traffic from the remote CE router is sent over an ingress label advertised by the local PE router. Once again, the label rides over an LDP LSP (or LDP tunneled through RSVP) to the local PE router from the remote PE router.

Figure 2: Layer 2 Circuit Concept


The Layer 2 circuit framework requires LDP to be used as the signaling protocol for advertising ingress labels. In most cases, it is not necessary to transport the Layer 2 encapsulation across the network; rather, the Layer 2 header can be stripped at one PE router, and reproduced at the egress PE router. Such Layer 2 information is carried in a special Layer 2 circuit header called a *control word*.

In the Layer 2 circuit IETF drafts, the control word is optional for most Layer 2 protocols, except Frame Relay and ATM AAL5 where it is required. However, in Junos OS Release 5.6 and later, a control word for all forms of Layer 2 circuits is sent by default. If you are establishing a Layer 2 circuit between a router running Junos OS Release 5.5 or earlier and a router running Junos OS Release 5.6 or later, use of the control word is negotiated automatically.

The Layer 2 protocols that are supported for Layer 2 circuits are:

- ATM cell-relay mode and ATM Adaptation Layer 5 (AAL5) mode on ATM2 intelligent queuing (IQ) interfaces
- Cisco High-Level Data Link Control (HDLC), Frame Relay, and PPP on SONET/SDH-based interfaces
- Ethernet, VLAN, and Extended VLAN on Ethernet-based interfaces

For an Ethernet 802.1q VLAN or simple Ethernet, the entire Ethernet frame without the preamble or frame check sequence (FCS) is transported. For ATM cell-relay mode, ATM cells are transported without a SAR process. For Cisco HDLC, the frame is transported in its entirety except for HDLC flags and the FCS. For PPP, the frame is transported in its entirety except for any media-specific framing information.

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 Layer 2 protocols that map Layer 2 control information into special bit fields in the control word are as follows:

- Frame Relay—This 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 “Option: Mapping Layer 2 Protocol Control Information into a Layer 2 Circuit” on page 16.)
- ATM AAL5 mode—This 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—This 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 MPLS Networks*. The differences are as follows:

- A packet with a sequence number of 0 is treated as out of sequence.
- Any packet which 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.

Layer 2 Circuits System Requirements

To implement Layer 2 circuits, your system must meet these minimum requirements:

- Junos OS Release 9.1 or later for nonstop routing (NSR) support.
- Junos OS Release 8.2 or later for M120 and MX Series routing platform support.
- Junos OS Release 7.3 or later for configuring Layer 2 circuits simultaneously over RSVP and LDP LSPs
- Junos OS Release 7.1 or later for local interface switching for Layer 2 circuits
- Junos OS Release 7.0 or later for specifying a unique maximum transmission unit (MTU) for each Layer 2 circuit
- Junos OS Release 6.4 or later for ATM2 IQ interface-based CoS and bandwidth reservation for trunks
- Junos OS Release 6.2 or later for Layer 2 circuit trunk mode on T Series and M320 routers and bandwidth reservation for Layer 2 circuits
- Junos OS Release 6.1 or later for Automatic Protection Switching (APS) for Layer 2 circuits
- Junos OS Release 6.0 or later for Layer 2 circuit traffic engineering, and Frame Relay or ATM control word mapping
- Junos OS Release 5.7 or later for ATM cell-relay mode or AAL5 Layer 2 circuits
- Junos OS Release 5.6 or later for Frame Relay, HDLC, and PPP-based Layer 2 circuits
- Junos OS Release 5.2 or later for Ethernet-based Layer 2 circuits
- Five Juniper Networks M Series, MX Series, or T Series routers

Layer 2 Circuits Terms and Acronyms

C

circuit cross-connect (CCC)	A Juniper Networks method of exchanging frames between one router interface running a Layer 2 protocol and another router interface using the same Layer 2 protocol. For more information about CCC, see either the <i>Junos Network Interfaces Configuration Guide</i> or the <i>Junos MPLS Applications Configuration Guide</i> .
control word	A 32-bit field used in Layer 2 circuits to transport sequence information, Layer 2 media control information, and padding.

L

Layer 2 circuit	A method of transporting Layer 2 frames between provider edge (PE) routers across an MPLS backbone using LDP signaling.
------------------------	---

CHAPTER 2

Layer 2 Circuits Configuration

When you configure Layer 2 circuits, you can use Ethernet, SONET/SDH, and ATM2 IQ interfaces on a PE router. The specific steps you must take to configure these interface types for Layer 2 circuits are described as follows:

- Configuring CCC Encapsulation on CE-Facing Ethernet Interfaces on page 9
- Configuring CCC Encapsulation on CE-Facing SONET/SDH Interfaces on page 10
- Configuring a CCC Encapsulation and a Layer 2 Circuit Mode on CE-Facing ATM2 IQ Interfaces on page 11
- Configuring the MPLS Family on Core Interfaces on page 13
- Configuring the Layer 2 Circuit Neighbor Address and Virtual Circuit Identifier on page 13
- Configuring LDP and an IGP to Transport Layer 2 Circuits on page 14
- Option: Applying Traffic Engineering to a Layer 2 Circuit on page 15
- Option: Mapping Layer 2 Protocol Control Information into a Layer 2 Circuit on page 16
- Option: Configuring APS for Layer 2 Circuits on page 16
- Option: Configuring Layer 2 Circuit Trunk Mode on ATM2 IQ Interfaces on page 17
- Option: Reserving LSP Bandwidth for a Layer 2 Circuit on page 19
- Option: Selecting an MTU for a Layer 2 Circuit on page 20
- Option: Configuring Local Interface Switching for a Layer 2 Circuit on page 21
- Option: Configuring Layer 2 Circuits Simultaneously over RSVP and LDP LSPs on page 21

Configuring CCC Encapsulation on CE-Facing Ethernet Interfaces

On Ethernet-based CE-facing PE router interfaces, you must configure one of the three Ethernet CCC encapsulation types—Ethernet CCC, VLAN CCC, or Extended VLAN CCC. Use the following guidelines to configure an Ethernet-based interface CCC encapsulation:

- **ethernet-ccc**—Use Ethernet CCC encapsulation on Ethernet interfaces that must accept packets carrying standard Tag Protocol ID (TPID) values.
- **extended-vlan-ccc**—Use extended VLAN CCC encapsulation on Ethernet interfaces that have VLAN 802.1Q tagging and must accept packets carrying TPIDs 0x8100, 0x9100, and 0x9901.

- **vlan-ccc**—Use VLAN CCC encapsulation on Ethernet interfaces with VLAN tagging enabled. VLAN CCC encapsulation supports TPID 0x8100 only. You must configure this encapsulation type on both the physical interface and the logical interface.
- You can configure multiple logical interfaces on Ethernet interfaces configured for VLAN mode. The valid VLAN ID range for the logical interfaces depends upon your PIC type:
 - 1 to 1023—Aggregated Ethernet, 4-port, 8-port, and 12-port Fast Ethernet, and management and internal Ethernet interfaces
 - 1 to 4094—48-port Fast Ethernet, Gigabit Ethernet, and 10-Gigabit Ethernet interfaces
 - VLAN ID 0 is reserved for tagging the priority of frames on all Ethernet PICs.
- For encapsulation type **vlan-ccc**, VLAN IDs from 1 through 511 are reserved for normal VLANs, and VLAN IDs from 512 through 1023 are reserved for CCC VLANs. For encapsulation type **extended-vlan-ccc**, all valid VLAN IDs per PIC type are valid for CCC VLANs.

For more information about VLAN IDs, see “Binding a VLAN ID to a Logical Interface” in the *Junos Network Interfaces Configuration Guide*.

To configure CCC interface encapsulation, include the **encapsulation** statement at the **[edit interfaces ethernet-interface-fpc/pic/port]** hierarchy level and select **ethernet-ccc**, **vlan-ccc**, or **extended-vlan-ccc** as the encapsulation type. If you select the VLAN CCC encapsulation, also include the **vlan-ccc** statement at the **[edit interfaces ethernet-interface-fpc/pic/port unit unit-number encapsulation]** logical interface hierarchy level. When using either VLAN CCC or extended VLAN CCC encapsulations, include the **vlan-tagging** statement at the **[edit interfaces ethernet-interface-fpc/pic/port]** hierarchy level.

```
[edit]
interfaces {
  fe-0/1/0 {
    vlan-tagging;
    encapsulation vlan-ccc;
    unit 0 {
      encapsulation vlan-ccc;
      vlan-id 600;
    }
  }
}
```

Configuring CCC Encapsulation on CE-Facing SONET/SDH Interfaces

On SONET/SDH interfaces in a PE router, you can use Frame Relay CCC, Cisco HDLC CCC, or PPP CCC encapsulation for Layer 2 circuits:

- To configure Frame Relay CCC interface encapsulation, include the **encapsulation** statement at the **[edit interfaces so-fpc/pic/port]** hierarchy level and select **frame-relay-ccc** as the encapsulation type. To enable Frame Relay CCC at the logical

interface level, include the **encapsulation frame-relay-ccc** statement at the **[edit interfaces so-fpc/pic/port unit unit-number]** hierarchy level.

- To configure Cisco HDLC CCC interface encapsulation, include the **encapsulation** statement at the **[edit interfaces so-fpc/pic/port]** hierarchy level and select **cisco-hdlc-ccc** as the encapsulation type.
- To configure PPP CCC interface encapsulation, include the **encapsulation** statement at the **[edit interfaces so-fpc/pic/port]** hierarchy level and select **ppp-ccc** as the encapsulation type.

```
[edit]
interfaces {
  so-0/0/0 {
    encapsulation (frame-relay-ccc | cisco-hdlc-ccc | ppp-ccc);
    unit 0 {
      encapsulation frame-relay-ccc;
    }
  }
}
```

Configuring a CCC Encapsulation and a Layer 2 Circuit Mode on CE-Facing ATM2 IQ Interfaces

On ATM2 IQ interfaces in a PE router, you need to configure two encapsulations to enable Layer 2 circuits: one at the **[edit interfaces at-fpc/pic/port]** hierarchy level and the other at the **[edit chassis fpc fpc-slot pic pic-slot]** hierarchy level. There are two types of ATM2 IQ Layer 2 circuits: cell-relay mode and ATM Adaptation Layer 5 (AAL5) mode. For both modes, you must specify the Physical Interface Card (PIC) type with the **pic-type atm2** statement at the **[edit interfaces at-fpc/pic/port atm-options]** hierarchy level. You can configure only one mode per PIC at a time. If you need to enable both ATM2 IQ Layer 2 circuit modes in the same router, you must configure the different modes on different PICs.

To configure a cell-relay mode Layer 2 circuit, include the **atm-l2circuit-mode cell** statement at the **[edit chassis fpc fpc-slot pic pic-slot]** hierarchy level and the **encapsulation atm-ccc-cell-relay** statement at both the **[edit interfaces at-fpc/pic/port]** physical hierarchy level and the **[edit interfaces at-fpc/pic/port unit unit-number]** logical interface hierarchy level.

```
[edit]
chassis {
  fpc 0 {
    pic 1 {
      atm-l2circuit-mode {
        cell;
      }
    }
  }
}
interfaces {
  at-0/1/0 {
    encapsulation atm-ccc-cell-relay;
```

```
    atm-options {
      cell-bundle-size 4;
      pic-type atm2;
      vpi 0;
    }
    unit 0 {
      encapsulation atm-ccc-cell-relay;
      vci 32;
      cell-bundle-size 10;
    }
  }
}
```

For ATM2 IQ Layer 2 circuit cell-relay mode only, you can adjust the cell bundle size at the physical interface level and the logical interface level. To configure, include the **cell-bundle-size** statement at either the **[edit interfaces at-fpc/pic/port atm-options]** physical interface hierarchy level or the **[edit interfaces at-fpc/pic/port unit unit-number]** logical interface hierarchy level. If the statement is included at both levels, the logical interface setting takes precedence. The default value for cell bundle size is 1 and the maximum value is 190. If you configure the cell bundle size statement, you should configure the same value on all ATM2 IQ neighbors.

To configure an AAL5 mode Layer 2 circuit, include the **atm-l2circuit-mode aal5** statement at the **[edit chassis fpcfpc-slot pic pic-slot]** hierarchy level and the **encapsulation atm-ccc-vc-mux** statement at the **[edit interfaces at-fpc/pic/port]** hierarchy level:

```
[edit]
chassis {
  fpc 1 {
    pic 2 {
      atm-l2circuit-mode {
        aal5;
      }
    }
  }
}
interfaces {
  at-1/2/0 {
    atm-options {
      pic-type atm2;
      vpi 0;
    }
    unit 0 {
      encapsulation atm-ccc-vc-mux;
      vci 32;
    }
  }
}
```

For more information on how to configure interfaces with CCC encapsulation types, see the *Junos MPLS Applications Configuration Guide* or the *Junos Network Interfaces Configuration Guide*.

Configuring the MPLS Family on Core Interfaces

Because LDP is used as the signaling protocol to transport MPLS labels across the core of the network, you must include the **family mpls** statement at the **[edit interfaces interface-name unit logical-unit-number]** hierarchy level. Include the statement on all router interfaces in the path from the local PE router to the remote PE router across the core network that transports the Layer 2 circuit traffic.

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

Configuring the Layer 2 Circuit Neighbor Address and Virtual Circuit Identifier

After you enable the PE router interfaces with the proper encapsulations, you then configure Layer 2 circuits (also referred to as VCs) between two or more PE router neighbors. To configure a Layer 2 circuit, include the **l2circuit** statement at the **[edit protocols]** hierarchy level.

Each Layer 2 circuit is represented by a logical interface on the local PE router, the IP address of the remote PE router neighbor, and a virtual circuit identifier. The logical interface connects the local PE router to the local CE router. The loopback address and router ID of the PE neighbor is commonly the neighbor's IP address. This address is also the destination end-point of the LSP tunnel, which transports the Layer 2 circuit to the neighbor. The virtual circuit ID uniquely identifies the VC to a specific neighbor.

This combination of logical interface, neighbor address, and virtual circuit ID is used to map a particular LDP forwarding equivalence class (FEC) received from a specific neighbor to a local VC. The egress label is added to a table and is used for sending traffic on that VC between the CE routers.

Both ends of a Layer 2 circuit must use the same Layer 2 technology because the Layer 2 encapsulation type is carried in the LDP FEC. The encapsulation type from a received FEC is matched against the local encapsulation type of the VC. If there is a mismatch, the VC is not established.

To add the IP address of the remote PE router neighbor into a Layer 2 circuit, include the **neighbor ip-address** statement at the **[edit protocols l2circuit]** hierarchy level. To map the remote neighbor to the local interface that connects to the CE router, include the **interface** statement at the **[edit protocols l2circuit neighbor ip-address]** hierarchy level. To select the identifier for the virtual circuit, include the **virtual-circuit-identifier** statement at the **[edit protocols l2circuit neighbor ip-address interface interface-name]** hierarchy level. To disable default control word processing, include the **no-control-word** statement at the

[edit protocols l2circuit neighbor *ip-address* interface *interface-name*] hierarchy level. Finally, to assign the Layer 2 circuit to a community, include the **community *community-name*** statement at the [edit protocols l2circuit neighbor *ip-address* interface *interface-name*] hierarchy level.



NOTE: On M Series routers only, if you include the **control-word** statement at the [edit protocols l2circuit neighbor *address* interface *interface-name*] hierarchy level (the default setting for Layer 2 circuits), the software cannot rewrite MPLS EXP bits.

```
[edit]
protocols {
  l2circuit {
    traceoptions {
      file name [replace] [size size] [files files] [nostamp];
      flag (error | topology | nlri | connections | route) [detail];
    }
    neighbor ip-address {
      interface interface-name {
        virtual-circuit-id identifier;
        no-control-word;
        community community-name;
      }
    }
  }
}
```

You do not need to specify the encapsulation type at the [edit protocols l2circuit] hierarchy level because it is already specified in the interface configuration.

Configuring LDP and an IGP to Transport Layer 2 Circuits

LDP is used as the signaling protocol to advertise the ingress MPLS label to the remote PE router. For this purpose, a remote LDP neighbor is established using the extended discovery mechanism described in RFC 5036, *LDP Specification*, and a session is established.

No new configuration is necessary in LDP because the LDP protocol recognizes the Layer 2 circuit configuration and initiates extended neighbor discovery for all Layer 2 circuit neighbors on the remote PE routers. This is very similar to the behavior of LDP when it is tunneled over RSVP. However, you must configure LDP on the **lo0.0** interface for extended neighbor discovery to function correctly.

LDP relies on an underlying IGP, such as OSPF or IS-IS. Therefore, configure LDP and your IGP on all routers in the path from the local PE router to the remote PE router across the service provider backbone.

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

```

        area 0.0.0.0 {
            interface so-0/1/0.0;
            interface lo0.0;
        }
    }
    ldp {
        interface so-0/1/0.0;
        interface lo0.0;
    }
}

```

Option: Applying Traffic Engineering to a Layer 2 Circuit

To traffic engineer Layer 2 circuits over multiple LSPs, you must create a community, assign a set of Layer 2 circuits to that community, define a policy to send the community traffic over a desired LSP, and apply the policy to the forwarding table.

To create a community, include the **community *community-name*** statement at the **[edit policy-options]** hierarchy level. To assign a Layer 2 circuit to a community, include the **community *community-name*** statement at the **[edit protocols l2circuit neighbor *ip-address* interface *interface-name*]** hierarchy level. To create a policy that sends community traffic over a specific LSP, include the **community *community-name*** statement at the **[edit policy-options policy-statement *policy-name* term *term-name* from]** hierarchy level and the **install-next-hop lsp *lsp-name*** statement at the **[edit policy-options policy-statement *policy-name* term *term-name* then]** hierarchy level. To apply the policy to the forwarding table, include the **export *policy-name*** statement at the **[edit routing-options forwarding-table]** hierarchy level.

```

[edit]
routing-options {
    forwarding-table {
        export policy-name;
    }
}
protocols {
    l2circuit {
        neighbor ip-address {
            interface interface-name {
                virtual-circuit-id identifier;
                community community-name;
            }
        }
    }
}
policy-options {
    policy-statement policy-name {
        from community community-name;
        then {
            install-nexthop lsp lsp-name;
            accept;
        }
    }
}
community community-name members value;

```

Option: Mapping Layer 2 Protocol Control Information into a Layer 2 Circuit

The control word is defined in Internet draft draft-martini-l2circuit-encap-mpls-07.txt *Encapsulation Methods for Transport of Layer 2 Frames Over IP and MPLS Networks*. It is a set of fields that carry Layer 2 control information across a Layer 2 circuit. The following control word support is available for PE routers:

- Frame Relay—To carry Frame Relay FECN/BECN information in a Layer 2 circuit control word, include the **translate-fecn-and-becn** statement at the **[edit interfaces *interface-name* encapsulation frame-relay-ccc unit *unit-number* encapsulation frame-relay-ccc family *ccc*]** hierarchy level. To carry Frame Relay DE information in a Layer 2 circuit control word, include the **translate-discard-eligible** statement at the **[edit interfaces *interface-name* encapsulation frame-relay-ccc unit *unit-number* encapsulation frame-relay-ccc family *ccc*]** hierarchy level.

```
[edit]
interfaces {
  interface-name {
    encapsulation frame-relay-ccc;
    unit 0 {
      encapsulation frame-relay-ccc;
      point-to-point;
      dlci 512;
      family ccc {
        translate-fecn-and-becn;
        translate-discard-eligible;
      }
    }
  }
}
```

- ATM AAL5 mode—For ATM2 IQ interfaces, the ATM AAL5 control word contains bit fields to carry sequence numbers, ATM cell loss priority (CLP), and explicit forward congestion indication (EFCI) information. When you configure ATM Layer 2 circuits, the control word carries the sequence number, CLP, and EFCI information by default. No additional configuration is necessary.
- ATM cell-relay mode—For ATM2 IQ interfaces, the ATM cell-relay control word supports sequence number processing only. Once you configure a cell-relay mode Layer 2 circuit, the sequence number information is carried by default. No additional configuration is necessary.

Option: Configuring APS for Layer 2 Circuits

To apply Automatic Protection Switching (APS) to a Layer 2 circuit, you must configure an APS working circuit and a protect circuit on PE router interfaces that support SONET options (such as SONET/SDH, ATM, and ATM2 IQ interfaces) and circuit cross-connect (CCC) encapsulation types. Then, you must configure the working circuit as the primary Layer 2 circuit interface and the protect circuit as the protected Layer 2 circuit interface. Note that APS only protects the PE-CE link and not the entire Layer 2 circuit.

To configure an APS working circuit, include the **working-circuit** statement at the **[edit interfaces *interface-name* sonet-options aps]** hierarchy level. To configure an APS protect circuit, include the **protect-circuit** statement at the **[edit interfaces *interface-name* sonet-options aps]** hierarchy level. To configure the primary Layer 2 circuit interface, include the **interface** statement at the **[edit protocols l2circuit neighbor *ip-address*]** hierarchy level. To configure the protected Layer 2 circuit interface, include the **protect-interface** statement at the **[edit protocols l2circuit neighbor *ip-address* interface *interface-name*]** hierarchy level.

```
[edit]
interfaces {
  at-0/0/1 {
    description "APS protect circuit";
    encapsulation CCC-encapsulation-type;
    sonet-options {
      aps {
        protect-circuit name;
      }
    }
  }
  at-1/3/1 {
    description "APS working circuit";
    encapsulation CCC-encapsulation-type;
    sonet-options {
      aps {
        working-circuit name;
      }
    }
  }
}
protocols {
  l2circuit {
    neighbor ip-address {
      interface at-1/3/1.0 {
        protect-interface at-0/0/1.0;
        virtual-circuit-id number;
      }
    }
  }
}
```

Option: Configuring Layer 2 Circuit Trunk Mode on ATM2 IQ Interfaces

When you configure Layer 2 circuits on CE-facing ATM2 IQ interfaces in a PE router that connects to some vendors' ATM switches, you can create a trunk. The trunk bundles several ATM cell streams into one LSP, preserves the cell loss priority (CLP) and class-of-service (CoS) information of the cells within the experimental (EXP) bits of the MPLS header, and provides network-to-network interface (NNI) or user-to-network interface (UNI) information within a proprietary header. A physical interface supports a total of 32 logical trunks in NNI mode and 8 logical trunks when you use the UNI option. To configure a trunk, include the **trunk** statement at the **[edit chassis fpc *fpc-slot* pic *pic-slot* atm-l2circuit-mode]** hierarchy level, select NNI or UNI mode with the **nni** or **uni** statement, and specify a number of bits in the ATM header that will carry an identifier

with the **id-width** statement. You can choose a value from 1 through 8 for the identifier width.

```
[edit]
chassis {
  fpc fpc-slot {
    pic pic-slot {
      atm-l2circuit-mode {
        trunk {
          (nni | uni) {
            id-width number;
          }
        }
      }
    }
  }
}
```

You can also configure several trunk options at the **[edit interfaces at-fpc/pic/port unit unit-number]** hierarchy level:

- To specify an ATM interface as the control channel for a Layer 2 circuit trunk, include the **control-channel** statement at the **[edit interfaces at-fpc/pic/port unit unit-number]** hierarchy level.
- To specify a trunk identifier, include the **trunk-id** statement at the **[edit interfaces at-fpc/pic/port unit unit-number]** hierarchy level. Trunk ID values range from 0 through 31.
- To configure the amount of bandwidth reserved for the trunk, include the **trunk-bandwidth** statement at the **[edit interfaces at-fpc/pic/port unit unit-number]** hierarchy level and specify a value from 1,000,000 bps (1 Mbps) through 542,526,792 bps.
- To apply a CoS scheduler map to the trunk, include the **atm-scheduler-map** statement at the **[edit interfaces at-fpc/pic/port unit unit-number]** hierarchy level. This statement must reference an ATM2 IQ interface-based scheduler map at the **[edit interfaces at-fpc/pic/port atm-options scheduler-maps map-name]** hierarchy level.

```
interfaces {
  at-fpc/pic/port {
    atm-options {
      pic-type atm2;
      scheduler-maps {
        map-name {
          ...
        }
      }
    }
    unit unit-number {
      atm-scheduler-map map-name;
      control-channel;
      trunk-id id-number;
      trunk-bandwidth bandwidth;
    }
  }
}
```

```
}
```

You can configure a variety of CoS-related statements for an ATM2 IQ interface-based scheduler map. To select the CoS mode used for virtual circuits, include the **vc-cos-mode** statement at the **[edit interfaces at-fpc/pic/port atm-options scheduler-maps map-name]** hierarchy level. To specify forwarding class settings, include the **priority**, **transmit-weight**, and **epd-threshold** statements at the **[edit interfaces at-fpc/pic/port atm-options scheduler-maps map-name forwarding-class class-name]** hierarchy level. For more information about CoS, see the *Junos Class of Service Configuration Guide*.

```
[edit]
interfaces {
  at-fpc/pic/port {
    atm-options {
      pic-type atm2;
      scheduler-maps {
        map-name {
          vc-cos-mode (alternate | strict);
          forwarding-class class-name {
            priority (high | low);
            transmit-weight (cells number-of-cells | percent percentage);
            epd-threshold plp0-threshold plp1 plp1-threshold;
          }
        }
      }
    }
  }
}
```

Option: Reserving LSP Bandwidth for a Layer 2 Circuit

You can specify the amount of bandwidth in bytes per second that must be available on an LSP for a specific Layer 2 circuit. By using a bandwidth constraint for a Layer 2 circuit, the router performs a type of call admission control. If an LSP exists that contains the required bandwidth, the Layer 2 circuit is established. If the bandwidth is not available on an LSP, the Layer 2 circuit is not established.

To configure bandwidth requirements for a Layer 2 circuit, include the **bandwidth** statement at the **[edit protocols mpls label-switched-path lsp-name]** hierarchy level and the **[edit protocols l2circuit neighbor ip-address interface interface-name]** hierarchy level.

```
[edit]
protocols {
  mpls {
    label-switched-path lsp-name {
      bandwidth traffic-class bytes-per-second;
    }
  }
  l2circuit {
    neighbor ip-address {
      interface interface-name {
        bandwidth bytes-per-second;
      }
    }
  }
}
```

```

    }
  }

```

You can also assign minimum bandwidth requirements for class-of-service (CoS) queues within a Layer 2 circuit and its corresponding LSP. Class type (CT) queues 0, 1, 2, and 3 in a Layer 2 circuit or LSP correspond to the standard four CoS queues available on M Series, MX Series, and T Series routers. To enable mapping of class type queues to the standard CoS queues, include the **bandwidth-model** statement at the **[edit protocols mpls diffserv-te]** hierarchy level. To assign specific bandwidth requirements to each class type queue, include the **ct0**, **ct1**, **ct2**, and **ct3** statements at the **[edit protocols mpls label-switched-path *lsp-name* bandwidth]** hierarchy level and the **[edit protocols l2circuit neighbor *ip-address* interface *interface-name* bandwidth]** hierarchy level.

```

[edit]
protocols {
  mpls {
    diffserv-te {
      bandwidth-model extended-mam;
    }
    label-switched-path lsp-name {
      bandwidth {
        ct0 100m;
        ct1 100m;
        ct2 50m;
        ct3 5m;
      }
    }
  }
}
l2circuit {
  neighbor ip-address {
    interface interface-name {
      bandwidth {
        ct0 100m;
        ct1 100m;
        ct2 50m;
        ct3 5m;
      }
    }
  }
}
}

```

For more information about class of service, see the *Junos Class of Service Configuration Guide*.

Option: Selecting an MTU for a Layer 2 Circuit

To configure the MTU for each individual Layer 2 circuit, include the **mtu** statement at the **[edit protocols l2circuit neighbor *ip-address* interface *interface-name*]** hierarchy level. If the MTU setting between Layer 2 circuit neighbors does not match, the Layer 2 circuit is torn down.



NOTE: If you configure an MTU value for an ATM cell relay interface on an ATM2 PIC and simultaneously configure an MTU value for a Layer 2 circuit that uses the same ATM2 PIC, the MTU value for the Layer 2 circuit takes precedence when calculating the cell bundle size and is advertised to Layer 2 circuit neighbors.

Option: Configuring Local Interface Switching for a Layer 2 Circuit

You can terminate a Layer 2 circuit locally on an ingress PE router. To configure a locally terminated circuit, include the **local-switching** statement at the **[edit protocols l2circuit]** hierarchy level. Select the Layer 2 circuit interfaces you want to connect locally, specify any APS protect interfaces, and configure an end interface. To select the Layer 2 circuit interfaces that are connected locally, include the **interface** statement at the **[edit protocols l2circuit local-switching]** hierarchy level. To configure an end interface, include the **end-interface** statement at the **[edit protocols l2circuit local-switching interface interface-name]** hierarchy level. To specify APS protect interfaces, include the **protect-interface** statement at the **[edit protocols l2circuit local-switching interface interface-name]** or **[edit protocols l2circuit local-switching interface interface-name end-interface interface-name]** hierarchy levels.

```
[edit]
protocols {
  l2circuit {
    local-switching {
      interface interface1 {
        protect-interface interface2;
        end-interface interface3 {
          protect-interface interface4;
        }
      }
      interface interface5 {
        protect-interface interface6;
        end-interface interface7 {
          protect-interface interface8;
        }
      }
    }
  }
}
```

Option: Configuring Layer 2 Circuits Simultaneously over RSVP and LDP LSPs

You can configure a Layer 2 circuit simultaneously over an RSVP LSP and an LDP LSP between the same two routers. To accomplish this, do the following:

- Configure two loopback addresses—a primary and a secondary loopback.
- Configure the RSVP LSP using the primary loopback address.

- Configure the LDP LSP using the secondary loopback address. You can accomplish this by advertising the secondary loopback address as a forwarding equivalence class (FEC) in LDP.
- Configure the transport tunnel endpoint (also known as a packet-switched network [PSN] tunnel endpoint in the IETF drafts) to be used for transporting the Layer 2 circuit traffic. To configure the tunnel endpoint, include the **psn-tunnel-endpoint** statement at the **[edit protocols l2circuit]** hierarchy level. By default, the tunnel endpoint is the same as the address of the Layer 2 circuit neighbor.

To verify that your configuration is operational, issue the **show l2circuit connections** command. This command has been enhanced to display tunnel endpoints. For more information about configuring Layer 2 circuits simultaneously over RSVP and LDP LSPs, see the *Junos VPNs Configuration Guide*.

CHAPTER 3


Layer 2 Circuits Configuration Examples

This section contains configuration examples and commands you can issue to verify Layer 2 circuit configurations:

- Example: Ethernet-Based Layer 2 Circuit Configuration on page 23
- Example: ATM2 IQ-Based Layer 2 Circuit Configuration on page 30
- Example: Layer 2 Circuit Traffic Engineering over Multiple LSPs Configuration on page 38
- Example: APS for a Layer 2 Circuit Configuration on page 48
- For More Information on page 51

Example: Ethernet-Based Layer 2 Circuit Configuration

Figure 3: Ethernet-Based Layer 2 Circuit Topology Diagram



In Figure 3 on page 23, a Layer 2 circuit is established between routers PE1 and PE2 to deliver Layer 2 traffic between customer routers CE1 and CE2. A Layer 2 circuit VC connection is configured on the PE routers only. No special configuration is required on the CE routers, and the provider core P0 router only requires MPLS and LDP on the appropriate interfaces to enable labels to be shared between the PE routers.

On Router CE1, configure the Fast Ethernet interface to handle VLAN traffic. Be sure to use the same VLAN ID both here and on the Fast Ethernet interface of Router CE2.

```
Router CE1 [edit]
interfaces {
  fe-1/1/2 {
    description "to PE1 fe-1/1/0";
    vlan-tagging;
    unit 600 {
      vlan-id 600; # Be sure this VLAN ID matches the VLAN ID of your CE neighbor.
      family inet {
        address 10.1.1.1/24;
      }
    }
  }
}
```

```

    }
  }
}

```

On Router PE1, configure the Ethernet-based CE-facing interface with the CCC encapsulation type of your choice. If you use VLAN CCC, include the **vlan-tagging** statement at the **[edit interfaces *ethernet-interface-fpc/pic/port*]** hierarchy level. Also, include the **encapsulation vlan-ccc** statement at both the **[edit interfaces *ethernet-interface-fpc/pic/port*]** and **[edit interfaces *ethernet-interface-fpc/pic/port* unit *unit-number*]** hierarchy levels.

Establish your Layer 2 circuit with configuration of the **l2circuit** statement at the **[edit protocols]** hierarchy level. Remember to include in your Layer 2 circuit configuration the IP address of your remote PE neighbor (usually the loopback address of the neighbor), the interface connected to the CE router, and a virtual circuit identifier for this VC. Then, configure MPLS, LDP, and an IGP (such as OSPF) to enable signaling for your Layer 2 circuit.

```

Router PE1 [edit]
interfaces {
  so-0/1/0 {
    description "to P0 so-0/0/0";
    unit 0 {
      family inet {
        address 192.168.245.22/30;
      }
      family mpls; # Include the MPLS family on core-facing interfaces.
    }
  }
  fe-1/1/0 {
    description "to CE1 fe-1/1/2";
    vlan-tagging;
    encapsulation vlan-ccc; # Configure CCC encapsulation on CE-facing interfaces.
    unit 600 {
      encapsulation vlan-ccc; # Enable this encapsulation on the logical interface.
      vlan-id 600;
    }
  }
  lo0 {
    unit 0 {
      family inet {
        address 10.245.14.173/32;
      }
    }
  }
}
protocols {
  ospf {
    traffic-engineering;
    area 0.0.0.0 {
      interface so-0/1/0.0;
      interface lo0.0;
    }
  }
}

```

```

ldp { # LDP is required as the signaling protocol for Layer 2 circuits.
  interface so-0/1/0.0;
  interface lo0.0; # You must include the loopback address in LDP.
}
l2circuit {
  neighbor 10.245.14.179 { # This points to the loopback of the PE neighbor.
    interface fe-1/1/0.600 { # Here you include the local CE-facing interface.
      virtual-circuit-id 5; # Be sure this ID matches the ID of your PE neighbor.
    }
  }
}
}

```

On Router P0, configure LDP, MPLS, and OSPF on the interfaces connected to the PE routers. The core router provides the MPLS backbone needed to tunnel Layer 2 traffic from the ingress PR router to the egress PE router.

```

Router P0 [edit]
interfaces {
  so-0/0/0 {
    description "to PE1 so-0/1/0";
    unit 0 {
      family inet {
        address 192.168.245.21/30;
      }
      family mpls; # Include the MPLS family on core interfaces.
    }
  }
  t3-0/2/1 {
    description "to PE2 t3-0/1/0";
    unit 0 {
      family inet {
        address 192.168.245.17/30;
      }
      family mpls; # Include the MPLS family on core interfaces.
    }
  }
  lo0 {
    unit 0 {
      family inet {
        address 10.245.14.171/32;
      }
    }
  }
}
protocols {
  ospf {
    traffic-engineering;
    area 0.0.0.0 {
      interface so-0/0/0.0;
      interface t3-0/2/1.0;
    }
  }
}
ldp { # LDP is required as the signaling protocol for Layer 2 circuits.
  interface so-0/0/0.0;
  interface t3-0/2/1.0;
}

```

```
}
}
```

On Router PE2, complete the Layer 2 circuit by configuring statements to match those previously set on Router PE1. Configure the Ethernet-based CE-facing interface with CCC encapsulation. Again, you must include the **vlan-tagging** statement at the **[edit interfaces ethernet-interface-fpc/pic/port]** hierarchy level when you use VLAN CCC. Also, include the **encapsulation vlan-ccc** statement at both the **[edit interfaces ethernet-interface-fpc/pic/port]** and **[edit interfaces ethernet-interface-fpc/pic/port unit unit-number]** hierarchy levels.

Establish your Layer 2 circuit with configuration of the **l2circuit** statement at the **[edit protocols]** hierarchy level. Remember to include in your Layer 2 circuit configuration the IP address of your remote PE neighbor (Router PE1), the virtual circuit identifier previously configured on Router PE1, and the interface connected to the CE router. Finally, configure MPLS, LDP, and OSPF to enable signaling for your Layer 2 circuit.

```
Router PE2 [edit]
interfaces {
  t3-0/1/0 {
    description "P0 t3-0/2/1";
    unit 0 {
      family inet {
        address 192.168.245.18/30;
      }
      family mpls; # Include the MPLS family on core-facing interfaces.
    }
  }
  fe-0/2/0 {
    description "to CE2 fe-0/2/0";
    vlan-tagging;
    encapsulation vlan-ccc; # Configure CCC encapsulation on CE-facing interfaces.
    unit 600 {
      encapsulation vlan-ccc; # Enable this encapsulation on the logical interface.
      vlan-id 600;
    }
  }
  lo0 {
    unit 0 {
      family inet {
        address 10.245.14.179/32;
      }
    }
  }
}
protocols {
  ospf {
    traffic-engineering;
    area 0.0.0.0 {
      interface t3-0/1/0.0;
      interface lo0.0;
    }
  }
  ldp { # LDP is required as the signaling protocol for Layer 2 circuits.
    interface t3-0/1/0.0;
```

```

    interface lo0.0; # You must include the loopback address in LDP.
  }
  l2circuit {
    neighbor 10.245.14.173 { # This points to the loopback of the PE neighbor.
      interface fe-0/2/0.600 { # Here you include the local CE-facing interface.
        virtual-circuit-id 5; # Be sure this ID matches the ID of your PE neighbor.
      }
    }
  }
}

```

On Router CE2, configure the Fast Ethernet interface to handle VLAN traffic. Be sure to use the same VLAN ID on this interface as the one seen on the Fast Ethernet interface of Router CE1.

```

Router CE2 [edit]
interfaces {
  fe-0/2/0 {
    description "to PE2 fe-0/2/0";
    vlan-tagging;
    unit 600 {
      vlan-id 600; # Be sure this VLAN ID matches the VLAN ID of your CE neighbor.
      family inet {
        address 10.1.1.2/24;
      }
    }
  }
}

```

Verifying Your Work

To verify proper operation of Layer 2 circuits, use the following commands:

- **ping mpls l2circuit interface***interface-name*
- **ping mpls l2circuit virtual-circuit** *virtual-circuit-id* neighbor *ip-address*
- **show l2circuit connections**
- Options: **[brief]** | **[down]** | **[extensive]** | **[history]** | **[instance]** | **[local-site]** | **[remote-site]** | **[status]** | **[summary]** | **[up]** | **[up-down]**
- **show ldp database**

In addition to displaying bindings for IP prefixes, the **show ldp database** command also displays the bindings for the Layer 2 FECs.

The following sections show the output of these commands used with the configuration example:

- Router PE1 Status on page 28
- Router P0 Status on page 28
- Router PE2 Status on page 29

Router PE1 Status

user@PE1> show l2circuit connections

Layer-2 Circuit Connections:

Legend for connection status (St)	Legend for interface status
EI -- encapsulation invalid	UP -- operational
MM -- mtu mismatch	Dn -- down
EM -- encapsulation mismatch	NP -- no present
OL -- no outgoing label	DS -- disabled
Dn -- down	WE -- wrong encapsulation
VC-Dn -- Virtual circuit Down	UN -- uninitialized
UP -- operational	
XX -- unknown	

Neighbor: 10.245.14.179

Interface	Type	St	Time last up	# Up trans
fe-1/1/0.600 (vc 5)	rmt	Up	Nov 30 00:54:55 2001	1

Local interface: fe-1/1/0.600, Status: Up, Encapsulation: VLAN
 Remote PE: 10.245.14.179, Negotiated control-word: Yes (Null)
 Incoming label: 100007, Outgoing label: 100000

user@PE1> show ldp database

Input label database, 10.245.14.173:0-10.245.14.171:0

Label	Prefix
100019	10.245.14.173/32
100020	10.245.14.179/32
3	10.245.14.171/32

Output label database, 10.245.14.173:0-10.245.14.171:0

Label	Prefix
100009	10.245.14.179/32
3	10.245.14.173/32
100008	10.245.14.171/32

Input label database, 10.245.14.173:0-10.245.14.179:0

Label	Prefix
100001	10.245.14.171/32
100002	10.245.14.173/32
3	10.245.14.179/32
100000	L2CKT VLAN VC 5

Output label database, 10.245.14.173:0-10.245.14.179:0

Label	Prefix
100009	10.245.14.179/32
3	10.245.14.173/32
100008	10.245.14.171/32
100007	L2CKT VLAN VC 5

Router P0 Status

user@P0> show ldp database

Input label database, 10.245.14.171:0-10.245.14.173:0

Label	Prefix
3	10.245.14.173/32
100009	10.245.14.179/32
100008	10.245.14.171/32

Output label database, 10.245.14.171:0-10.245.14.173:0

```

Label      Prefix
100019     10.245.14.173/32
100020     10.245.14.179/32
3          10.245.14.171/32

```

Input label database, 10.245.14.171:0-10.245.14.179:0

```

Label      Prefix
100001     10.245.14.171/32
3          10.245.14.179/32
100002     10.245.14.173/32

```

Output label database, 10.245.14.171:0-10.245.14.179:0

```

Label      Prefix
100019     10.245.14.173/32
100020     10.245.14.179/32
3          10.245.14.171/32

```

Router PE2 Status

user@PE2> show l2circuit connections

Layer-2 Circuit Connections:

Legend for connection status (St)	Legend for interface status
EI -- encapsulation invalid	UP -- operational
MM -- mtu mismatch	Dn -- down
EM -- encapsulation mismatch	NP -- no present
OL -- no outgoing label	DS -- disabled
Dn -- down	WE -- wrong encapsulation
VC-Dn -- Virtual circuit Down	UN -- uninitialized
UP -- operational	
XX -- unknown	

Neighbor: 10.245.14.173

Interface	Type	St	Time last up	# Up trans
fe-0/2/0.600 (vc 5)	rmt	Up	Nov 30 00:54:54 2001	1

Local interface: fe-0/2/0.600, Status: Up, Encapsulation: VLAN
Remote PE: 10.245.14.173, Negotiated control-word: Yes (Null)
Incoming label: 100000, Outgoing label: 100007

user@PE2> show ldp database

Input label database, 10.245.14.179:0-10.245.14.171:0

```

Label      Prefix
100019     10.245.14.173/32
3          10.245.14.171/32
100020     10.245.14.179/32

```

Output label database, 10.245.14.179:0-10.245.14.171:0

```

Label      Prefix
100001     10.245.14.171/32
100002     10.245.14.173/32
3          10.245.14.179/32

```

Input label database, 10.245.14.179:0-10.245.14.173:0

```

Label      Prefix
3          10.245.14.173/32
100008     10.245.14.171/32
100009     10.245.14.179/32
100007 L2CKT VLAN VC 5

```

```

Output label database, 10.245.14.179:0-10.245.14.173:0
Label      Prefix
100001     10.245.14.171/32
100002     10.245.14.173/32
3          10.245.14.179/32
100000     L2CKT VLAN VC 5

```

Example: ATM2 IQ-Based Layer 2 Circuit Configuration

Figure 4: ATM2 IQ-Based Layer 2 Circuit Topology Diagram

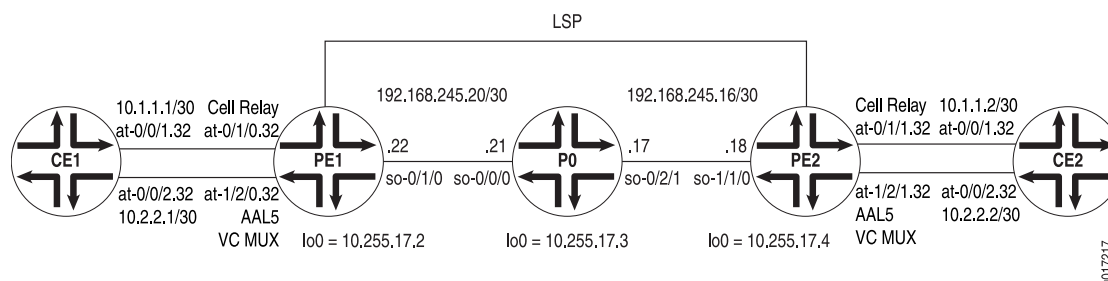


Figure 4 on page 30 shows a similar network topology to our previous two examples. In this example, Routers PE1 and PE2 use ATM cell-relay mode on a CE-facing interface and ATM AAL5 mode on a second CE-facing interface.

On Router CE1, configure the ATM2 IQ interfaces to handle ATM traffic. Interface **at-0/0/1** handles standard ATM traffic while interface **at-0/0/2** handles AAL5 traffic.

```

Router CE1 [edit]
interfaces {
  at-0/0/1 {
    description "to PE1 at-0/1/0";
    atm-options {
      pic-type atm2; # Layer 2 circuits are compatible with
      vpi 0; # ATM2 IQ interfaces.
    }
    unit 0 {
      vci 32;
      family inet {
        address 10.1.1.1/30;
      }
    }
  }
  at-0/0/2 {
    description "to PE1 at-1/2/0";
    atm-options {
      pic-type atm2; # Layer 2 circuits are compatible with
      vpi 0; # ATM2 IQ interfaces.
    }
    unit 0 {
      encapsulation atm-vc-mux;
      vci 32;
      family inet {
        address 10.2.2.1/30;
      }
    }
  }
}

```

```

    }
  }
}

```

On Router PE1, configure the ATM2 IQ-based CE-facing interfaces: one with ATM cell-relay mode CCC encapsulation and the other with ATM VC multiplexing CCC encapsulation. Also enable the corresponding Layer 2 circuit modes at the **[edit chassis]** hierarchy level. In this case, you must configure cell-relay mode on Physical Interface Card (PIC) 1 in Flexible PIC Concentrator (FPC) 0 and AAL5 mode on PIC 2 in FPC 1.

Establish your Layer 2 circuit with configuration of the **l2circuit** statement at the **[edit protocols]** hierarchy level. Remember to include in your Layer 2 circuit configuration the IP address of your remote PE neighbor (usually the loopback address of the neighbor), the interfaces connected to the CE router, and a virtual circuit identifier for each VC. In this case, you will establish one VC for cell-relay mode traffic and a second VC for AAL5 traffic. Then, configure MPLS, LDP, and an IGP (such as OSPF) to enable signaling for your Layer 2 circuit.

```

Router PE1 [edit]
chassis {
  fpc 0 {
    pic 1 {
      atm-l2circuit-mode {
        cell; # This dedicates FPC 0 PIC 1 to cell-relay mode.
      }
    }
  }
  fpc 1 {
    pic 2 {
      atm-l2circuit-mode {
        aal5; # This dedicates FPC 1 PIC 2 to AAL5 mode.
      }
    }
  }
}
interfaces {
  at-0/1/0 {
    description "to CE1 at-0/0/1";
    encapsulation atm-ccc-cell-relay; # Cell-relay requires cell-relay encapsulation.
    atm-options {
      cell-bundle-size 4; # This sets the cell bundle size for the interface.
      pic-type atm2; # Layer 2 circuits are compatible with
      vpi 0; # ATM2 IQ interfaces.
    }
    unit 0 {
      encapsulation atm-ccc-cell-relay; # Encapsulation for the logical interface.
      vci 32;
      cell-bundle-size 10; # The cell bundle size for the logical interface overrides
    } # the physical interface setting.
  }
  at-1/2/0 {
    description "to CE1 at-0/0/2";
    atm-options {
      pic-type atm2; # Layer 2 circuits are compatible with
      vpi 0; # ATM2 IQ interfaces.
    }
  }
}

```

```

    }
    unit 0 {
        encapsulation atm-ccc-vc-mux; # AAL5 requires CCC VC MUX encapsulation.
        vci 32;
    }
}
so-0/1/0 {
    description "to P0 so-0/0/0";
    unit 0 {
        family inet {
            address 192.168.245.22/30;
        }
        family mpls; # Include the MPLS family on core-facing interfaces.
    }
}
lo0 {
    unit 0 {
        family inet {
            address 10.255.17.2/32;
        }
    }
}
}
protocols {
    ospf {
        traffic-engineering;
        area 0.0.0.0 {
            interface so-0/1/0.0;
            interface lo0.0;
        }
    }
    ldp { # LDP is required as the signaling protocol for Layer 2 circuits.
        interface so-0/1/0.0;
        interface lo0.0;
    }
    l2circuit {
        neighbor 10.255.17.4 { # This points to the loopback of the PE neighbor.
            interface at-0/1/0.32 { # Here you include the local CE-facing interface.
                virtual-circuit-id 1; # Be sure this ID matches the ID of your PE neighbor.
            }
        }
        neighbor 10.255.17.4 { # This points to the loopback of the PE neighbor.
            interface at-1/2/0.32 { # Here you include the local CE-facing interface.
                virtual-circuit-id 2; # Be sure this ID matches the ID of your PE neighbor.
            }
        }
    }
}
}

```

On Router P0, configure LDP, MPLS, and OSPF on the interfaces connected to the PE routers. The core router provides the MPLS backbone needed to tunnel Layer 2 traffic from the ingress PR router to the egress PE router.

Router P0 [edit]
 interfaces {
 so-0/0/0 {

```

description "to PE1 so-0/1/0";
unit 0 {
    family inet {
        address 192.168.245.21/30;
    }
    family mpls; # Include the MPLS family on core interfaces.
}
}
so-0/2/1 {
    description "to PE2 so-1/1/0";
    unit 0 {
        family inet {
            address 192.168.245.17/30;
        }
        family mpls; # Include the MPLS family on core interfaces.
    }
}
}
lo0 {
    unit 0 {
        family inet {
            address 10.255.17.3/32;
        }
    }
}
}
protocols {
    ospf {
        traffic-engineering;
        area 0.0.0.0 {
            interface so-0/0/0.0;
            interface so-0/2/1.0;
        }
    }
    ldp { # LDP is required as the signaling protocol for Layer 2 circuits.
        interface so-0/0/0.0;
        interface so-0/2/1.0;
    }
}
}

```

On Router PE2, complete the Layer 2 circuit by configuring statements to match those previously set on Router PE1.

Configure the ATM2 IQ-based CE-facing interfaces: one with ATM cell-relay mode CCC encapsulation and the other with ATM VC multiplexing CCC encapsulation. Also enable the corresponding Layer 2 circuit modes at the **[edit chassis]** hierarchy level. In this case, you must configure cell-relay mode on PIC 1 in FPC 0 and AAL5 mode on PIC 2 in FPC 1.

Complete your Layer 2 circuit with configuration of the **l2circuit** statement at the **[edit protocols]** hierarchy level. Remember to include in your Layer 2 circuit configuration the IP address of your remote PE neighbor (Router PE1), the interfaces connected to the CE router, and a virtual circuit identifier for each VC. In this case, you will establish one VC for cell-relay mode traffic and a second VC for AAL5 traffic. Then, configure MPLS, LDP, and an IGP (such as OSPF) to enable signaling for your Layer 2 circuit.

Router PE2 [edit]

```
chassis {
  fpc 0 {
    pic 1 {
      atm-l2circuit-mode {
        cell; # This dedicates FPC 0 PIC 1 to cell-relay mode.
      }
    }
  }
  fpc 1 {
    pic 2 {
      atm-l2circuit-mode {
        aal5; # This dedicates FPC 1 PIC 2 to AAL5 mode.
      }
    }
  }
}
interfaces {
  at-0/1/1 {
    description "to CE2 at-1/0/1";
    encapsulation atm-ccc-cell-relay; # Cell-relay requires cell-relay encapsulation.
    atm-options {
      cell-bundle-size 4; # This sets the cell bundle size for the physical interface.
      pic-type atm2; # Layer 2 circuits are compatible with
      vpi 0; # ATM2 IQ interfaces.
    }
    unit 0 {
      encapsulation atm-ccc-cell-relay; # Also configure the encapsulation
      vci 32; # on the logical interface.
      cell-bundle-size 10; # The cell bundle size for the logical interface overrides
    } # the physical interface setting.
  }
  at-1/2/1 {
    description "to CE2 at-1/0/2";
    atm-options {
      pic-type atm2; # Layer 2 circuits are compatible with
      vpi 0; # ATM2 IQ interfaces.
    }
    unit 0 {
      encapsulation atm-ccc-vc-mux; # AAL5 requires CCC VC MUX encapsulation.
      vci 32;
    }
  }
  so-1/1/0 {
    description "to P0 so-0/2/1";
    unit 0 {
      family inet {
        address 192.168.245.18/30;
      }
      family mpls; # Include the MPLS family on core-facing interfaces.
    }
  }
  lo0 {
    unit 0 {
      family inet {
        address 10.255.17.4/32;
      }
    }
  }
}
```

```

    }
  }
}
protocols {
  ospf {
    traffic-engineering;
    area 0.0.0.0 {
      interface so-1/1/0.0;
      interface lo0.0;
    }
  }
}
ldp { # LDP is required as the signaling protocol for Layer 2 circuits.
  interface so-1/1/0.0;
  interface lo0.0;
}
l2circuit {
  neighbor 10.255.17.2 { # This points to the loopback of the PE neighbor.
    interface at-0/1/1.32 { # Here you include the local CE-facing interface.
      virtual-circuit-id 1; # Be sure this ID matches the ID of your PE neighbor.
    }
  }
  neighbor 10.255.17.2 { # This points to the loopback of the PE neighbor.
    interface at-1/2/1.32 { # Here you include the local CE-facing interface.
      virtual-circuit-id 2; # Be sure this ID matches the ID of your PE neighbor.
    }
  }
}
}
}

```

On Router CE2, configure the ATM2 IQ interfaces to handle ATM traffic. Interface **at-1/0/1** handles standard ATM traffic while interface **at-1/0/2** handles AAL5 traffic.

```

Router CE2 [edit]
interfaces {
  at-1/0/1 {
    description "to PE2 at-0/1/1";
    atm-options {
      pic-type atm2; # Layer 2 circuits are compatible with
      vpi 0; # ATM2 IQ interfaces.
    }
    unit 0 {
      vci 32;
      family inet {
        address 10.1.1.2/30;
      }
    }
  }
  at-1/0/2 {
    description "to PE2 at-1/2/1";
    atm-options {
      pic-type atm2; # Layer 2 circuits are compatible with
      vpi 0; # ATM2 IQ interfaces.
    }
    unit 0 {
      encapsulation atm-vc-mux;
      vci 32;
    }
  }
}

```

```

        family inet {
            address 10.2.2.2/30;
        }
    }
}

```

Verifying Your Work

To verify proper operation of Layer 2 circuits, use the following commands:

- **ping mpls l2circuit interface *interface-name***
- **ping mpls l2circuit virtual-circuit *virtual-circuit-id* neighbor *ip-address***
- **show l2circuit connections**
- Options: **[brief] | [down] | [extensive] | [history] | [instance] | [local-site] | [remote-site] | [status] | [summary] | [up] | [up-down]**
- **show interfaces**
- **show route table l2circuit.0**
- **show ldp database l2circuit detail**

This is what the operational command output looks like for cell-relay mode on Router PE1:

```

user@PE1> show l2circuit connections
Layer-2 Circuit Connections:

```

Legend for connection status (St)

EI -- encapsulation invalid	NP -- interface not present
MM -- mtu mismatch	Dn -- down
EM -- encapsulation mismatch	VC-Dn -- Virtual circuit Down
CM -- control-word mismatch	Up -- operational
OL -- no outgoing label	XX -- unknown
NC -- intf encaps not CCC/TCC	

Legend for interface status

Up -- operational
Dn -- down

Neighbor: 10.255.17.4

Interface	Type	St	Time last up	# Up trans
at-0/1/0.0 (vc 32)	rmt	Up	Jan 22 15:15:52 2003	1

Local interface: at-0/1/0.0, Status: Up, Encapsulation: ATM CELL (VC Mode)

Remote PE: 10.255.17.4, Negotiated control-word: Yes (Non-null)
Incoming label: 100000, Outgoing label: 100000

```

user@PE1> show route table l2circuit.0 detail

```

```

l2circuit.0: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)
10.255.17.4:CtrlWord:9:32:Local/96 (1 entry, 1 announced)
  *L2CKT Preference: 7
    Next hop: via so-0/2/0.0 weight 1, selected
    Label-switched-path PE1-PE2
    Protocol next hop: 10.255.17.4 Indirect next hop: 85135e8 367
    State: <Active Int>

```

```

Local AS: 69
Age: 2:34
Task: l2 circuit
Announcement bits (1): 0-LDP
AS path: I
VC Label 100000, MTU 0, cell-bundle size 80

10.255.17.4:CtrlWord:9:32:Remote/96 (1 entry, 1 announced)
*LDP Preference: 9
Next hop type: Discard
State: <Active Int>
Local AS: 69
Age: 28:11
Task: LDP
Announcement bits (1): 1-l2 circuit
AS path: I
VC Label 100000, MTU 0, cell-bundle size 80

user@PE1> show interfaces at-0/1/0.0 extensive
Logical interface at-0/1/0.0 (Index 66) (SNMP ifIndex 40) (Generation 4)
Flags: Point-To-Point SNMP-Traps Encapsulation: ATM-CCC-Cell-Relay
L2 circuit cell bundle size: 10 , bundle timeout: 125 usec, timeout count:
0
L2 circuit out-of-sequence count: 0
Traffic statistics:
[...]

user@PE1> show interfaces media at-0/1/0
Physical interface: at-0/1/0, Enabled, Physical link is Up
Interface index: 154, SNMP ifIndex: 50
Link-level type: ATM-CCC-Cell-Relay, MTU: 4482, Clocking: Internal, SONET mode,

L2 circuit mode: Cell, Speed: OC12, Loopback: None
[...]

user@PE1> show ldp database l2circuit detail
Input label database, 10.255.17.2:0--10.255.17.4:0
Label Prefix
100000 L2CKT CtrlWord ATM CELL (VC Mode) VC 32
Cell bundle size: 80
State: Active
Age: 9:48

Output label database, 10.255.17.2:0--10.255.17.4:0
Label Prefix
100000 L2CKT CtrlWord ATM CELL (VC Mode) VC 32
Cell bundle size: 80
State: Active
Age: 9:48

```

This is what the operational command output looks like on Router PE1 if AAL5 mode is used:

```

user@PE1> show l2circuit connections
Layer-2 Circuit Connections:

Legend for connection status (St)
EI -- encapsulation invalid      NP -- interface not present
MM -- mtu mismatch              Dn -- down
EM -- encapsulation mismatch     VC-Dn -- Virtual circuit Down
CM -- control-word mismatch      Up -- operational

```

OL -- no outgoing label XX -- unknown
NC -- intf encaps not CCC/TCC

Legend for interface status

Up -- operational

Dn -- down

Neighbor: 10.255.17.4

Interface	Type	St	Time last up	# Up trans
at-1/2/0.0 (vc 32)	rmt	Up	Feb 18 18:00:00 2003	1

Local interface: at-1/2/0.0, Status: Up, Encapsulation: ATM AAL5

Remote PE: 10.255.17.4, Negotiated control-word: Yes (Non-null)

Incoming label: 100016, Outgoing label: 100032

user@PE1> show interfaces media at-0/1/0

Physical interface: at-0/1/0, Enabled, Physical link is Up

Interface index: 154, SNMP ifIndex: 50

Link-level type: ATM-PVC, MTU: 4482, Clocking: Internal, SONET mode,

L2 circuit mode: AAL5, Speed: OC12, Loopback: None, Payload scrambler: Enabled

[...]

user@PE1> show interfaces at-1/2/0.0 extensive

Logical interface at-1/2/0.0 (Index 68) (SNMP ifIndex 40) (Generation 38)

Flags: Point-To-Point SNMP-Traps Encapsulation: ATM-CCC-VCMUX

L2 circuit out-of-sequence count: 0

Traffic statistics:[...]

Example: Layer 2 Circuit Traffic Engineering over Multiple LSPs Configuration

Figure 5: Layer 2 Circuit Traffic Engineering Topology Diagram




Figure 5 on page 38 shows a network topology designed to traffic engineer different Layer 2 circuits over select LSPs. Across provider edge routers PE1 and PE2, an ATM AAL5 mode Layer 2 circuit connects customer edge routers CE1 and CE2, and a Frame Relay Layer 2 circuit connects routers CE3 and CE4. To maintain traffic separation, the ATM traffic is mapped onto LSP1 with a community named **gold**, and the Frame Relay traffic is mapped onto LSP2 with a community named **silver**. LSP1 takes the direct route between routers PE1 and PE2, while LSP2 travels from Router PE1 to PE2 through Router P0.

In addition to traffic engineering, you can send Layer 2 control information in the control word of a Layer 2 circuit. In this case, Frame Relay discard eligible (DE), forward explicit congestion notification (FECN), and backward explicit congestion notification (BECN) information is mapped into the control word. Likewise, ATM cell loss priority (CLP) and explicit forward congestion indicator (EFCI) information is mapped into the control word.

To traffic engineer Layer 2 circuits over multiple LSPs, you assign a set of Layer 2 circuits to a community and then apply a policy to send the community traffic over a desired LSP. To create communities, include the **community *community-name*** statement at the **[edit policy-options]** hierarchy level. To assign a Layer 2 circuit to a community, include the **community *community-name*** statement at the **[edit protocols l2circuit neighbor *neighbor-id* interface *interface-name*]** hierarchy level. To send community traffic over a specific LSP, include the **community *community-name*** statement at the **[edit policy-options policy-statement *policy-name* term *term-name* from]** hierarchy level and the **install-next-hop lsp *lsp-name*** statement at the **[edit policy-options policy-statement *policy-name* term *term-name* then]** hierarchy level.

On Router CE1, configure the ATM2 IQ interface **at-0/0/1.32** to handle ATM AAL5 traffic:

```
Router CE1 [edit]
            interfaces {
            at-0/0/1 {
                description "to PE1 at-1/1/0";
                atm-options {
                    pic-type atm2; # Layer 2 circuits are compatible with
                    vpi 0; # ATM2 IQ interfaces.
                }
                unit 0 {
                    encapsulation atm-vc-mux; # Use ATM VC MUX encapsulation on the CE.
                    point-to-point;
                    vci 0.32;
                    family inet {
                        address 10.0.0.1/30;
                    }
                }
            }
        }
```

On Router CE3, configure the SONET/SDH interface at **so-0/0/1** to handle Frame Relay traffic:

```
Router CE3 [edit]
            interfaces
            so-0/0/1 {
                description "to PE1 so-1/1/1"
                encapsulation frame-relay; # Use Frame Relay encapsulation on the CE router.
                unit 0 {
                    encapsulation frame-relay;
                    point-to-point;
                    dlci 512;
                    family inet {
                        address 10.10.0.1/30;
                    }
                }
            }
```

```
}
```

On Router PE1, configure the ATM2 IQ-based CE1-facing interface **at-1/1/0** with ATM VC multiplexing CCC encapsulation on the logical interface. Also enable the corresponding Layer 2 circuit modes at the **[edit chassis]** hierarchy level. In this case, you must configure AAL5 mode on PIC 1 in FPC 1. Once you configure the ATM2 IQ-based Layer 2 circuit, the CLP and EFCI bits are mapped to the control word by default.

Next, configure the Frame Relay interface **so-1/1/1** with Frame Relay CCC encapsulation on both the physical and logical interface. Map the DE, FECN, and BECN bits to the control word with the **translate-fecn-and-becn** and **translate-discard-eligible** statements at the **[edit interfaces so-fpc/pic/port unit unit-number family ccc]** hierarchy level.

Establish your Layer 2 circuits with configuration of the **l2circuit** statement at the **[edit protocols]** hierarchy level. Remember to include in your Layer 2 circuit configuration the IP address of your remote PE neighbor (usually the loopback address of the neighbor), the interfaces connected to the CE router, and a virtual circuit identifier for each VC. In this case, you will establish one VC for ATM AAL5 traffic and a second VC for Frame Relay traffic. Then, configure MPLS, LDP, and an IGP (such as OSPF) to enable signaling for your Layer 2 circuit. Two LSPs are established for the ATM and Frame Relay traffic: LSP1 for ATM traffic going directly to Router PE2 and LSP 2 for Frame Relay traffic going through Router P0 before going on to Router PE2.

Finally, configure a community for traffic separation for the ATM and Frame Relay Layer 2 circuits. Assign community **gold** to the ATM VC and community **silver** to the Frame Relay VC. Remember to give the communities numerical values and configure a routing policy to match the communities to specific LSPs. This policy is applied as an **export** policy for the forwarding table at the **[edit routing-options]** hierarchy level.

```
Router PE1 [edit]
            chassis {
              fpc 1 {
                pic 1 {
                  atm-l2circuit-mode {
                    aal5; # This dedicates FPC 1 PIC 1 to AAL5 mode.
                  }
                }
              }
            }
            interfaces {
              at-1/1/0 {
                description "to CE1 at-0/0/2";
                atm-options {
                  pic-type atm2; # Layer 2 circuits are compatible with
                  vpi 0; # ATM2 IQ interfaces.
                }
                unit 0 {
                  encapsulation atm-ccc-vc-mux; # CLP/EFCI bits are mapped to control word.
                  vci 0.32;
                }
              }
              so-0/1/0 {
                description "to P0 so-0/0/0";
                unit 0 {
```

```

        family inet {
            address 192.168.27.17/30;
        }
        family mpls; # Include the MPLS family on core-facing interfaces.
    }
}
so-1/1/0 {
    description "to PE2 so-1/0/1";
    unit 0 {
        family inet {
            address 192.168.27.10/30;
        }
        family mpls; # Include the MPLS family on core-facing interfaces.
    }
}
so-1/1/1 {
    description "to CE3 so-0/0/1";
    dce;
    encapsulation frame-relay-ccc;
    unit 0 {
        encapsulation frame-relay-ccc;
        point-to-point;
        dlci 512;
        family ccc {
            translate-fecn-and-becn; # Option to map FECN/BECN bits to control word.
            translate-discard-eligible; # Option to map DE bit to control word.
        }
    }
}
lo0 {
    unit 0 {
        family inet {
            address 10.255.17.2/32;
        }
    }
}
}
routing-options {
    forwarding-table {
        export layer2communities; # This applies communities to the Layer 2 circuits.
    }
}
protocols {
    mpls {
        label-switched-path lsp1 { # ATM LSP 1 goes directly to PE2.
            to 10.255.17.4;
            primary direct;
        }
        label-switched-path lsp2 { # Frame Relay LSP 2 goes through P0.
            to 10.255.17.4;
            primary thruP0;
        }
        path direct {
            192.168.27.9 strict;
        }
        path thruP0 {

```

```

        192.168.27.18 strict;
        192.168.27.25 strict;
    }
    interface so-0/0/1.0;
    interface so-1/1/0.0;
}
ospf {
    traffic-engineering;
    area 0.0.0.0 {
        interface so-0/0/1.0;
        interface so-1/1/0.0;
        interface lo0.0;
    }
}
ldp { # LDP is required as the signaling protocol for Layer 2 circuits.
    interface so-0/0/1.0;
    interface so-1/1/0.0;
    interface lo0.0;
}
l2circuit {
    neighbor 10.255.17.4 {# This points to the loopback of the PE neighbor.
        interface at-1/1/0.32 { # Here you include the local CE-facing interface.
            virtual-circuit-id 1; # Be sure this ID matches the ID of your PE neighbor.
            community gold;# Assigns the ATM Layer 2 circuit to the gold community.
        }
    }
    neighbor 10.255.17.4 {# This points to the loopback of the PE neighbor.
        interface so-1/1/1.512 { # Here you include the local CE-facing interface.
            virtual-circuit-id 2; # Be sure this ID matches the ID of your PE neighbor.
            community silver; # Assigns the Frame Relay Layer 2 circuit to silver.
        }
    }
}
}
policy-options {
    policy-statement layer2communities { # Here you map the communities to LSPs.
        term 10 {
            from community gold; # Apply community gold to LSP 1.
            then {
                install-nexthop lsp lsp1;
                accept;
            }
        }
        term 20 {
            from community silver; # Apply community silver to LSP 2.
            then {
                install-nexthop lsp lsp2;
                accept;
            }
            community gold members 103:1; # Assign numerical value to community gold.
            community silver members 103:2; # Assign numerical value to community silver.
        }
    }
}
}
}

```

On Router P0, configure LDP, MPLS, and OSPF on the interfaces connected to the PE routers. The core router provides the MPLS backbone needed to tunnel Layer 2 traffic from the ingress PR router to the egress PE router. Only LSP 2 for Frame Relay passes through Router P0.

```
Router P0 [edit]
interfaces {
  so-1/1/1 {
    description "to PE1 so-0/0/1";
    unit 0 {
      family inet {
        address 192.168.27.18/30;
      }
      family mpls; # Include the MPLS family on core interfaces.
    }
  }
  so-1/1/2 {
    description "to PE2 so-0/1/1";
    unit 0 {
      family inet {
        address 192.168.27.26/30;
      }
      family mpls; # Include the MPLS family on core interfaces.
    }
  }
  lo0 {
    unit 0 {
      family inet {
        address 10.255.17.3/32;
      }
    }
  }
}
protocols {
  mpls {
    interface so-1/1/1.0;
    interface so-1/1/2.0;
  }
  ospf {
    traffic-engineering;
    area 0.0.0.0 {
      interface so-1/1/1.0;
      interface so-1/1/2.0;
    }
  }
  ldp {# LDP is required as the signaling protocol for Layer 2 circuits.
    interface so-1/1/1.0;
    interface so-1/1/2.0;
  }
}
```

On Router PE2, complete the Layer 2 circuit by configuring statements to match those previously set on Router PE1.

Establish your Layer 2 circuits with configuration of the **l2circuit** statement at the **[edit protocols]** hierarchy level. Remember to include in your Layer 2 circuit configuration the IP address of your remote PE neighbor (usually the loopback address of the neighbor), the interfaces connected to the CE router, and a virtual circuit identifier for each VC. In this case, you will establish one VC for ATM AAL5 traffic and a second VC for Frame Relay traffic. Then, configure MPLS, LDP, and an IGP (such as OSPF) to enable signaling for your Layer 2 circuit. Two LSPs are established for the ATM and Frame Relay traffic: LSP1 for ATM traffic going directly to Router PE2 and LSP 2 for Frame Relay traffic going through Router P0 before going on to Router PE2.

Finally, configure a community for traffic separation for the ATM and Frame Relay Layer 2 circuits. The ATM VC has community **gold** and the Frame Relay VC has community **silver**. Remember to give the communities numerical values and configure a routing policy to match the communities to specific LSPs. This policy is applied as an **export** policy for the forwarding table at the **[edit routing-options]** hierarchy level.

```
Router PE2 [edit]
            chassis {
              fpc 1 {
                pic 0 {
                  atm-l2circuit-mode {
                    aal5; # This dedicates FPC 1 PIC 0 to AAL5 mode.
                  }
                }
              }
            }
            interfaces {
              at-1/0/1 {
                description "to CE2 at-1/1/2";
                atm-options {
                  pic-type atm2; # Layer 2 circuits are compatible with
                  vpi 0; # ATM2 IQ interfaces.
                }
                unit 0 {
                  encapsulation atm-ccc-vc-mux; # CLP and EFCI appear in the control word.
                  vci 0.32;
                }
              }
              so-0/1/1 {
                description "to P0 so-1/1/2";
                unit 0 {
                  family inet {
                    address 192.168.27.25/30;
                  }
                  family mpls; # Include the MPLS family on core-facing interfaces.
                }
              }
              so-1/0/1 {
                description "to PE1 so-1/1/0";
                unit 0 {
                  family inet {
                    address 192.168.27.9/30;
                  }
                  family mpls; # Include the MPLS family on core-facing interfaces.
                }
              }
            }
          }
```

```

    }
  }
  so-1/1/1 {
    description "to CE4 so-1/1/2";
    dce;
    encapsulation frame-relay-ccc;
    unit 0 {
      encapsulation frame-relay-ccc;
      point-to-point;
      dlci 512;
      family ccc {
        translate-fecn-and-becn; # Option to map FECN/BECN bits to control word.
        translate-discard-eligible; # Option to map DE bit to control word.
      }
    }
  }
  lo0 {
    unit 0 {
      family inet {
        address 10.255.17.4/32;
      }
    }
  }
}
routing-options {
  forwarding-table {
    export layer2communities; # This maps communities to the Layer 2 circuits.
  }
}
protocols {
  mpls {
    label-switched-path lsp1 { # ATM LSP 1 goes directly to Router PE2.
      to 10.255.17.2;
      primary direct;
    }
    label-switched-path lsp2 { # Frame Relay LSP 2 goes through Router PO.
      to 10.255.17.2;
      primary thruPO;
    }
    path direct {
      192.168.27.10 strict;
    }
    path thruPO {
      192.168.27.26 strict;
      192.168.27.17 strict;
    }
    interface so-0/1/1.0;
    interface so-1/0/1.0;
  }
  ospf {
    traffic-engineering;
    area 0.0.0.0 {
      interface so-0/1/1.0;
      interface so-1/0/1.0;
      interface lo0.0;
    }
  }
}

```

```

}
ldp { # LDP is required as the signaling protocol for Layer 2 circuits.
    interface so-0/1/1.0;
    interface so-1/0/1.0;
    interface lo0.0;
}
l2circuit {
    neighbor 10.255.17.2 {# This points to the loopback of the PE neighbor.
        interface at-1/0/1.32 { # Here you include the local CE-facing interface.
            virtual-circuit-id 1; # Be sure this ID matches the ID of your PE neighbor.
            community gold;# Assigns the ATM Layer 2 circuit to the gold community.
        }
    }
    neighbor 10.255.17.2 {# This points to the loopback of the PE neighbor.
        interface so-1/1/1.512 { # Here you include the local CE-facing interface.
            virtual-circuit-id 2; # Be sure this ID matches the ID of your PE neighbor.
            community silver; # Assigns the Frame Relay Layer 2 circuit to silver.
        }
    }
}
}
policy-options {
    policy-statement layer2communities { # Here you map communities to LSPs.
        term 10 {
            from community gold; # Apply community gold to LSP 1.
            then {
                install-nexthop lsp lsp1;
                accept;
            }
        }
        term 20 {
            from community silver; # Apply community silver to LSP 2.
            then {
                install-nexthop lsp lsp2;
                accept;
            }
            community gold members 103:1; # Assign numerical value to community gold.
            community silver members 103:2; # Assign numerical value to community silver.
        }
    }
}
}

```

On Router CE2, configure the ATM2 IQ interfaces to handle ATM traffic. Interface **at-1/0/1** handles AAL5 traffic.

```

Router CE2 [edit]
interfaces {
    at-1/1/2 {
        description "to PE2 at-1/0/1";
        atm-options {
            pic-type atm2; # Layer 2 circuits are compatible with
            vpi 0; # ATM2 IQ interfaces.
        }
        unit 0 {
            encapsulation atm-vc-mux; # Use ATM VC MUX encapsulation on the CE.
            point-to-point;
        }
    }
}

```

```

        vci 0.32;
        family inet {
            address 10.0.0.2/30;
        }
    }
}

```

On Router CE4, configure the SONET/SDH interface at **so-1/1/2** to handle Frame Relay traffic:

```

Router CE3 [edit]
interfaces {
    so-1/1/2 {
        description " to PE2 so-1/1/1";
        encapsulation frame-relay-ccc; # Use Frame Relay encapsulation on the CE.
        unit 0 {
            encapsulation frame-relay-ccc;
            point-to-point;
            dlci 512;
            family inet {
                address 10.10.0.2/30;
            }
        }
    }
}

```

Verifying Your Work

To verify proper operation of traffic engineered Layer 2 circuits, use the following command:

```
show route table mpls.0 detail
```

On Router PE1, you can see that ATM traffic is part of the **gold** community that has a value of **103:1** and is associated with LSP 1. Likewise, Frame Relay traffic is part of the **silver** community that has a value of **103:2** and is associated with LSP 2:

```

user@PE1> show route table mpls.0 detail
mpls.0: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)
at-1/1/0.32 (1 entry, 1 announced)
    *L2CKT Preference: 7
    Next hop: 192.168.27.9 via so-1/1/0.0 weight 1, selected
    Label-switched-path lsp1
        Label operation: Push 100032 Offset: -4
        Next hop: via so-0/0/1.0 weight 1
        Label-switched-path lsp2
        Label operation: Push 100032 Offset: -4
        Protocol next hop: 10.255.17.4
        Push 100032 Offset: -4
        Indirect next hop: 8576bd0 300
        State: <Active Int>
        Age: 7:18
        Task: Common L2 VC
        Announcement bits (2): 0-KRT 1-Common L2 VC
        AS path: I
        Communities: 103:1 # This is the gold community.

```

```

so-1/1/1.512 (1 entry, 1 announced)
  *L2CKT Preference: 7
    Next hop: 192.168.27.9 via so-1/1/0.0 weight 1
    Label-switched-path lsp1
    Label operation: Push 100048 Offset: -4
    Next hop: via so-0/0/1.0 weight 1, selected
  Label-switched-path lsp2
    Label operation: Push 100048 Offset: -4
    Protocol next hop: 10.255.17.4
    Push 100048 Offset: -4
    Indirect next hop: 860f1f8 293
    State: <Active Int>
    Age: 5:15
    Task: Common L2 VC
    Announcement bits (2): 0-KRT 1-Common L2 VC
    AS path: I
    Communities: 103:2 # This is the silver community.

```

Example: APS for a Layer 2 Circuit Configuration

Figure 6: APS for a Layer 2 Circuit Topology Diagram




Figure 6 on page 48 shows that APS is configured on a PE router to protect a PE-CE link in a Layer 2 circuit. This example shows only the PE router configuration and assumes that you have preconfigured a full Layer 2 circuit topology. For more information about configuring Layer 2 circuits, see *Junos Layer 2 Circuits Feature Guide*.

On Router PE1, configure ATM2 IQ interface **at-0/0/1** as an APS protect circuit and ATM2 IQ interface **at-1/3/1** as a working circuit. Also, configure the working circuit interface as the primary interface for your Layer 2 circuit and configure the protect circuit interface as the protected interface for your Layer 2 circuit.

```

Router PE1 [edit]
chassis {
  fpc 0 {
    pic 0 {
      atm-l2circuit-mode {
        cell; # This dedicates FPC 0 PIC 0 to cell-relay mode.
      }
    }
  }
  fpc 1 {
    pic 3 {
      atm-l2circuit-mode {
        cell; # This dedicates FPC 1 PIC 3 to cell-relay mode.
      }
    }
  }
}

```

```

    }
  }
}
interfaces {
  at-0/0/1 {
    description "To CE1 at-0/1/2";
    encapsulation atm-ccc-cell-relay;
    sonet-options {
      aps {
        protect-circuit TEST; # This interface is the APS protect circuit.
      }
    }
    atm-options {
      pic-type atm2;
      promiscuous-mode;
    }
    unit 0 {
      allow-any-vci;
    }
  }
  at-1/3/1 {
    description "To CE1 at-0/1/1";
    encapsulation atm-ccc-cell-relay;
    sonet-options {
      aps {
        working-circuit TEST; # This interface is the APS working circuit.
      }
    }
    atm-options {
      pic-type atm2;
      promiscuous-mode;
    }
    unit 0 {
      allow-any-vci;
    }
  }
}
protocols {
  l2circuit {
    neighbor 10.255.17.4 {
      interface at-1/3/1.0 { # The Layer 2 circuit interface is the working circuit.
        protect-interface at-0/0/1.0; # The protect-interface is the protect circuit.
        virtual-circuit-id 100;
      }
    }
  }
}
}

```

Verifying Your Work

To verify proper operation of APS for Layer 2 circuits, use the following command:

```
show l2circuit connections
```

After you configure the Layer 2 circuit and the APS working and protect circuits, you can see which APS circuit is active for the Layer 2 circuit with the **show l2circuit connections**

command. The first local interface that is displayed is always the active circuit. If the second local interface field indicates **Protect-Inactive**, the working circuit is active, as shown in this output sample.

```
user@PE0> 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
OL -- no outgoing label        XX -- unknown
NC -- intf encaps not CCC/TCC
CB -- rcvd cell-bundle size bad
```

Legend for interface status

```
Up -- operational
Dn -- down
```

Neighbor: 10.255.17.4

```
Interface      Type  St      Time last up      # Up trans
at-1/3/1.0(vc 100)  rmt  Up      Sep 3 17:48:25 2003      1
  Local interface: at-1/3/1.0, Status: Up, Encapsulation: ATM CELL (PORT Mode)
  Remote PE: 10.255.17.4, Negotiated control-word: Yes (Null)
  Incoming label: 100368, Outgoing label: 100112
  Local interface: at-0/0/1.0, Status: Dn, Encapsulation: ATM CELL (PORT Mode),
  Protect-Inactive
```

Conversely, if the first local interface is marked with the **Protect-Active** indicator, and the second local interface indicates **Primary-Inactive**, the protect circuit is active, as shown here:

```
user@PE0> 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
OL -- no outgoing label        XX -- unknown
NC -- intf encaps not CCC/TCC
CB -- rcvd cell-bundle size bad
```

Legend for interface status

```
Up -- operational
Dn -- down
```

Neighbor: 10.255.17.4

```
Interface      Type  St      Time last up      # Up trans
at-1/3/1.0(vc 100)  rmt  Up      Sep 3 17:51:06 2003      2
  Local interface: at-0/0/1.0, Status: Up, Encapsulation: ATM CELL (PORT Mode),
  Protect-Active
  Remote PE: 10.255.17.4, Negotiated control-word: No
  Incoming label: 100368, Outgoing label: 100112
  Local interface: at-1/3/1.0, Status: Dn, Encapsulation: ATM CELL (PORT Mode),
  Primary-Inactive
```

For More Information

For additional information about Layer 2 circuits, see the following:

- *Junos Network Interfaces Configuration Guide*
- *Junos Class of Service Configuration Guide*
- *Junos VPNs Configuration Guide*
- RFC 3036, *LDP Specification*
- RFC 4447, *Pseudowire Setup and Maintenance Using the Label Distribution Protocol (LDP)* (except section 5.3)
- RFC 4448, *Encapsulation Methods for Transport of Ethernet over MPLS Networks*
- Internet draft draft-martini-atm-encap-mpls-01.txt, *Encapsulation Methods for Transport of ATM Cells/Frame Over IP and MPLS Networks* (expires December 2002)
- Internet draft draft-martini-l2circuit-encap-mpls-11.txt, *Encapsulation Methods for Transport of Layer 2 Frames over MPLS Networks* (expires August 2006)
- Internet draft draft-martini-l2circuit-trans-mpls-19.txt, *Transport of Layer 2 Frames Over MPLS* (expires September 2006)
- Internet draft draft-kompella-ppvpn-l2vpn-03.txt, *Layer 2 VPNs Over Tunnels* (expires October 2003)

PART 2

Index

- Index on page 55

Index

A

APS

Layer 2 circuits.....16

L

Layer 2 circuits

configuration procedure.....9

example configuration

APS.....48

ATM2 IQ.....30

traffic engineering.....38

operational mode commands

ATM2 IQ.....36

Ethernet.....27

traffic engineering.....47

options

APS.....16

control word.....16

local interface switching.....21

MTU.....20

reserving LSP bandwidth.....19

simultaneous RSVP and LDP LSPs.....21

traffic engineering.....15

trunk mode for ATM2 IQ interfaces.....17

overview.....3

system requirements.....6

Layer 2 circuits

example configuration

Ethernet.....23

operational mode commands

APS.....49

S

system requirements

Layer 2 circuits.....6

T

traffic engineering

Layer 2 circuits.....15

