



Junos[®] OS

Multiple Instances for Label Distribution Protocol Feature Guide

Release

11.1



Published: 2011-01-19

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 Multiple Instances for Label Distribution Protocol Feature Guide

Release 11.1

Copyright © 2011, Juniper Networks, Inc.

All rights reserved. Printed in USA.

Revision History

January 2011—R1 Junos OS 11.1

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	Multiple Instances for Label Distribution Protocol	
Chapter 1	Multiple Instances for Label Distribution Protocol Concepts and Reference Materials	3
	Multiple Instances for Label Distribution Protocol Overview	3
	Multiple Instances for Label Distribution Protocol System Requirements	4
	Multiple Instances for Label Distribution Protocol Terms and Acronyms	5
Chapter 2	Multiple Instances for Label Distribution Protocol Configuration Example	7
	Example: Configuring Multiple-Instance LDP	7
	Verifying Your Work	26
	Router CE3 Status	27
	Router PE3 Status	27
	Router CE1 Status	29
	Router PE1 Status	30
	Router PE2 Status	32
	Router CE2 Status	37
	Router PE4 Status	38
	Router CE4 Status	40
Part 2	Index	
	Index	45

List of Figures

Part 1	Multiple Instances for Label Distribution Protocol	
Chapter 1	Multiple Instances for Label Distribution Protocol Concepts and Reference Materials	3
	Figure 1: Carrier-of-Carriers Example	4
Chapter 2	Multiple Instances for Label Distribution Protocol Configuration Example	7
	Figure 2: Multiple-Instance LDP Topology Diagram	8

List of Tables

Part 1	Multiple Instances for Label Distribution Protocol	
Chapter 2	Multiple Instances for Label Distribution Protocol Configuration Example	7
	Table 1: Multiple-Instance LDP Example—Routing Protocol Summary	8
	Table 2: Multiple-Instance LDP Example—Loopback Addresses	9

PART 1

Multiple Instances for Label Distribution Protocol

- Multiple Instances for Label Distribution Protocol Concepts and Reference Materials on page 3
- Multiple Instances for Label Distribution Protocol Configuration Example on page 7

CHAPTER 1

Multiple Instances for Label Distribution Protocol Concepts and Reference Materials

This section contains the following topics:

- Multiple Instances for Label Distribution Protocol Overview on page 3
- Multiple Instances for Label Distribution Protocol System Requirements on page 4
- Multiple Instances for Label Distribution Protocol Terms and Acronyms on page 5

Multiple Instances for Label Distribution Protocol Overview

Versions of Junos OS earlier than Release 5.3 support multiple VPN routing and forwarding (VRF) instances of Border Gateway Protocol (BGP), Intermediate System-to-Intermediate System (IS-IS), Open Shortest Path First (OSPF), Protocol Independent Multicast (PIM), and Routing Information Protocol (RIP). Junos OS Release 5.4 and later adds support for multiple instances of the Label Distribution Protocol (LDP).


This support allows LDP to be used to advertise labels in a carrier-of-carriers scenario from a core provider edge (PE) router to a customer carrier edge (CE) router. This is especially useful when the carrier customer is a basic Internet service provider (ISP) and wants to restrict full Internet routes to its PE routers. By using LDP instead of BGP, the carrier customer shields its other internal routers from the Internet at large.

Multiple-instance LDP is also useful when a carrier customer wants to provide Layer 2 VPN or Layer 3 VPN services to its customers.

Using multiple-instance LDP lets you circumvent one of the requirements of RFC 3107: the need to run full-mesh internal BGP (IBGP) within the carrier customer's autonomous system (AS). When you use multiple-instance LDP, full-mesh IBGP is unnecessary.

In Figure 1 on page 4, the customer carrier in AS 21 can configure one instance of LDP for all routers in AS 21 instead of using full-mesh IBGP.

Figure 1: Carrier-of-Carriers Example



In general, if there are a limited number of customer carrier sites and few internal routes in the customer carrier AS, it is simpler and quicker to use LDP than to configure a full IBGP mesh.

An instance of LDP operates essentially in the same way as a master instance. Each instance of LDP must be enabled on all the desired interfaces and a separate set of LDP data structures is maintained for each instance. Instance information includes a set of LDP interfaces, neighbors, sessions, and databases.

For more information about carrier-of-carriers VPNs, see the *Junos VPNs Configuration Guide*.

For more information about LDP, see the *Junos MPLS Applications Configuration Guide*.

Multiple Instances for Label Distribution Protocol System Requirements

To implement the multiple-instance LDP feature, your system must meet these minimum requirements:

- Junos OS Release 8.2 for support on MX Series routers.
- Junos OS Release 5.4 or later for support on M Series and T Series routers.
- Two Juniper Networks M Series, MX Series, or T Series routers for basic multiple-instance LDP; and a minimum of four Juniper Networks routers to act as PE routers in a carrier-of-carriers network.

Multiple Instances for Label Distribution Protocol Terms and Acronyms

C

carrier-of-carriers VPN A VPN that transports data traffic between two or more telecommunications carrier sites across a core provider network. The core provider becomes a carrier for the customer carrier, which, in turn, provides Internet or VPN services to end customers. For more information about carrier-of-carriers VPNs, see the *Junos VPNs Configuration Guide*.

L

Label Distribution Protocol (LDP) A protocol used to distribute labels in an MPLS-enabled network. For more information about LDP, see the *Junos MPLS Applications Configuration Guide*.

V

VPN routing and forwarding (VRF) instance A unique routing table created to maintain VPN routing and forwarding information. One routing table is created per instance, which keeps prefix information and data private from other instances. For more information about VRF instances, see the *Junos VPNs Configuration Guide*.

CHAPTER 2

Multiple Instances for Label Distribution Protocol Configuration Example

This section contains the following topic:

- Example: Configuring Multiple-Instance LDP on page 7

Example: Configuring Multiple-Instance LDP

The master LDP instance is configured at the **[edit protocols]** hierarchy level.

You can configure a specific instance of LDP by using the **ldp** statement at the **[edit routing-instances *routing-instance-name* protocols]** hierarchy level. This creates an instance of LDP for the particular VRF routing instance. You must specify all the required VRF statements and apply export and import policies to your LDP instance for the configuration to commit properly.

Most of the LDP hierarchy levels available in a master instance are also available for specific instances of LDP. However, the **no-forwarding** option does not work in a VRF-based instance of LDP.

Figure 2: Multiple-Instance LDP Topology Diagram




Figure 2 on page 8 shows an example of a carrier-of-carriers network. CE3 and CE4 are end customer CE routers residing in AS 100. The VPN provider in AS 200 has three types of routers: PE3 and PE4 are PE routers that connect to the end customer, CE1 and CE2 act as the intermediate carrier CE routers, and P2 and P3 are internal transit routers. PE1 and PE2 in AS 300 are PE routers servicing the intermediate VPN provider, and P0 and P1 are transit routers for the top-tier carrier.

To make this configuration work, you must complete three major tasks:

1. Configure external BGP between the VPN customer CE and the VPN provider PE.
2. Configure internal BGP using the VPN family between both pairs of PE routers (one IBGP connection between PE1 and PE2 and a second IBGP connection between Router PE3 and Router PE4).
3. Establish LDP and Interior Gateway Protocol (IGP) connections on all remaining links. This example uses OSPF as the IGP, but you can use the IGP of your choice.

Information supporting this carrier-of-carriers multiple-instance LDP example is summarized in Table 1 on page 8 and Table 2 on page 9.

Table 1: Multiple-Instance LDP Example—Routing Protocol Summary

Connection	Protocols
CE3 - PE3	EBGP family inet
PE3 - P2 - CE1	OSPF and LDP

Table 1: Multiple-Instance LDP Example—Routing Protocol Summary (*continued*)

Connection	Protocols
CE1 - PE1	OSPF and LDP
PE1 - P0 - P1 - PE2	OSPF and LDP
PE1 - PE2	IBGP family inet-vpn
PE2 - CE2	OSPF and LDP
CE2 - P3 - PE4	OSPF and LDP
PE4 - CE4	EBGP family inet
PE3 - PE4	IBGP family inet-vpn

Table 2: Multiple-Instance LDP Example—Loopback Addresses

Router	Loopback Address
PE1	10.255.255.171
PE2	10.255.255.172
P0	10.255.255.173
P1	10.255.255.174
P2	10.255.255.175
P3	10.255.255.176
PE3	10.255.255.177
PE4	10.255.255.178
CE1	10.255.255.179
CE2	10.255.255.180
CE3	10.255.255.181 10.49.100.1
CE4	10.255.255.182 10.49.200.1

Your configuration tasks start at Router CE3 and move router by router through the first part of the VPN provider network, into the carrier AS, through the second VPN provider cluster of AS 200, and end at the second VPN customer Router CE4.

Since Router CE3 is the first customer router, configure EBGP between Router CE3 and the connected VPN provider Router PE3. You must also advertise your loopback address into BGP with a routing policy to allow IP reachability with Router CE4.

```
Router CE3 [edit]
interfaces {
  so-1/2/0 {
    description "to pe3 so-1/2/0";
    unit 0 {
      family inet {
        address 192.255.198.14/30;
      }
    }
  }
  lo0 {
    unit 0 {
      family inet {
        address 10.255.255.181/32;
        address 10.49.100.1/32;
      }
    }
  }
}
routing-options {
  static {
    route 10.49.100.0/24 reject;
    route 10.49.101.0/24 reject;
  }
  autonomous-system 100;
}
protocols {
  bgp {
    group provider {
      type external;
      export static-to-bgp;
      peer-as 200;
      neighbor 192.255.198.13;
    }
  }
}
policy-options {
  policy-statement static-to-bgp {
    term 1 {
      from {
        protocol static;
        route-filter 10.49.100.0/24 exact;
        route-filter 10.49.101.0/24 exact;
      }
      then accept;
    }
    term 2 {
```

```

        from protocol direct;
        then accept;
    }
    term 3 {
        then reject;
    }
}
}

```

On Router PE3, the configuration tasks are more involved. You need to complete the EBGP connection to Router CE3 in a VRF instance, enable MPLS and LDP on the interface pointing toward the VPN provider Router CE1, and configure a master instance of IBGP to reach Router PE4 at the far edge of AS 200.

Finally, set up an outbound VRF policy that places all BGP traffic and directly connected interfaces into a BGP community and an inbound VRF policy that accepts similar BGP community traffic from Router PE4.

```

Router PE3 [edit]
interfaces {
  so-1/2/0 {
    unit 0 {
      family inet {
        address 192.255.198.13/30;
      }
      family mpls;
    }
  }
  so-1/2/1 {
    description "to p2 so-1/2/1";
    unit 0 {
      family inet {
        address 192.255.198.9/30;
      }
      family mpls;
    }
  }
  lo0 {
    unit 0 {
      family inet {
        address 10.255.255.177/32;
      }
    }
  }
}
routing-options {
  autonomous-system 200;
}
protocols {
  mpls {
    interface so-1/2/0.0;
  }
  bgp {
    group internal {
      type internal;
      local-address 10.255.255.177;
    }
  }
}

```

```
    peer-as 200;
    neighbor 10.255.255.178 {
        family inet-vpn {
            unicast;
        }
    }
}
ospf {
    area 0.0.0.0 {
        interface so-1/2/1.0;
        interface lo0.0 {
            passive;
        }
    }
}
ldp {
    interface so-1/2/1.0;
}
}
policy-options {
    policy-statement vpn-customer-import {
        term 1 {
            from {
                protocol bgp;
                community vpn-customer-comm;
            }
            then accept;
        }
        term 2 {
            then reject;
        }
    }
    policy-statement vpn-customer-export {
        term 1 {
            from protocol [bgp direct];
            then {
                community add vpn-customer-comm;
                accept;
            }
        }
        term 2 {
            then reject;
        }
    }
    community vpn-customer-comm members target:200:100;
}
routing-instances {
    vpn-customer {
        instance-type vrf;
        interface so-1/2/0.0;
        route-distinguisher 10.255.255.177:1;
        vrf-import vpn-customer-import;
        vrf-export vpn-customer-export;
        protocols {
            bgp {
```



```

        group customer {
            type external;
            peer-as 100;
            as-override;
            neighbor 192.255.198.14;
        }
    }
}

```

On Router P2, enable LDP and the IGP used for transporting labels (in this case, OSPF). You will repeat these tasks on all transit core routers, both in the VPN provider network and the core carrier network.

```

Router P2 [edit]
interfaces {
    so-1/2/0 {
        description "to ce1 so-1/2/0";
        unit 0 {
            family inet {
                address 192.255.198.2/30;
            }
            family mpls;
        }
    }
    so-1/2/1 {
        description "to pe3 so-1/2/1";
        unit 0 {
            family inet {
                address 192.255.198.10/30;
            }
            family mpls;
        }
    }
    lo0 {
        unit 0 {
            family inet {
                address 10.255.255.175/32;
            }
        }
    }
}
routing-options {
    autonomous-system 200;
}
protocols {
    ospf {
        area 0.0.0.0 {
            interface lo0.0 {
                passive;
            }
            interface so-1/2/0.0;
            interface so-1/2/1.0;
        }
    }
}

```

```
ldp {  
  interface so-1/2/0.0;  
  interface so-1/2/1.0;  
}  
}
```

For Router CE1, configure LDP and OSPF in the same manner that you configured Router P2.

```
Router CE1 [edit]  
interfaces {  
  t3-0/1/0 {  
    description "to pe1 t3-0/2/1";  
    unit 0 {  
      family inet {  
        address 192.255.197.18/30;  
      }  
      family mpls;  
    }  
  }  
  so-1/2/0 {  
    description "to p2 so-1/2/0";  
    unit 0 {  
      family inet {  
        address 192.255.198.1/30;  
      }  
      family mpls;  
    }  
  }  
  lo0 {  
    unit 0 {  
      family inet {  
        address 10.255.255.179/32;  
      }  
    }  
  }  
}  
routing-options {  
  autonomous-system 200;  
}  
protocols {  
  ospf {  
    area 0.0.0.0 {  
      interface so-1/2/0.0;  
      interface lo0.0 {  
        passive;  
      }  
      interface t3-0/1/0.0;  
    }  
  }  
  ldp {  
    interface t3-0/1/0.0;  
    interface so-1/2/0.0;  
  }  
}
```

On core carrier Router PE1, configure a master instance for OSPF, LDP, MPLS, and IBGP (with the **family inet-vpn** option) to connect the router to neighbor Router PE2. Next, implement multiple-instance LDP by establishing a secondary instance. Enable LDP and OSPF in this instance for Router PE1 to communicate with Router CE1. MPLS is not required in the secondary instance.

Finally, set up an outbound VRF policy that places all LDP traffic coming from Router CE1 into a BGP community, an export policy that sends this community traffic to Router PE2, and an inbound VRF policy that accepts similar BGP community traffic from Router PE2. This step tunnels the VPN provider's LDP traffic into the carrier's BGP session.

```
Router PE1 [edit]
interfaces {
  so-0/0/0 {
    description "to p0 so-0/1/0";
    unit 0 {
      family inet {
        address 192.255.197.21/30;
      }
      family mpls;
    }
  }
  t3-0/2/1 {
    description "to ce1 t3-0/1/0";
    unit 0 {
      family inet {
        address 192.255.197.17/30;
      }
      family mpls;
    }
  }
  lo0 {
    unit 0 {
      family inet {
        address 10.255.255.171/32;
      }
    }
  }
}
routing-options {
  autonomous-system 300;
}
protocols {
  mpls {
    interface t3-0/2/1.0;
  }
  bgp {
    group pe {
      type internal;
      local-address 10.255.255.171;
      family inet-vpn {
        unicast;
      }
      peer-as 300;
      neighbor 10.255.255.172;
    }
  }
}
```

```
    }
  }
  ospf {
    area 0.0.0.0 {
      interface lo0.0 {
        passive;
      }
      interface so-0/0/0.0;
    }
  }
  ldp {
    interface so-0/0/0.0;
  }
}
policy-options {
  policy-statement vpn-provider-import {
    term 1 {
      from {
        protocol bgp;
        community vpn-provider-comm;
      }
      then accept;
    }
    term 2 {
      then reject;
    }
  }
  policy-statement vpn-provider-export {
    term 1 {
      from protocol ldp;
      then {
        community add vpn-provider-comm;
        accept;
      }
    }
    term 2 {
      then reject;
    }
  }
  policy-statement bgp-routes-to-export {
    term 1 {
      from {
        protocol bgp;
        community vpn-provider-comm;
      }
      then accept;
    }
    term 2 {
      then reject;
    }
  }
  community vpn-provider-comm members target:300:200;
}
routing-instances {
  vpn-provider {
    instance-type vrf;
```

```

interface t3-0/2/1.0;
route-distinguisher 10.255.255.171:1;
vrf-import vpn-provider-import;
vrf-export vpn-provider-export;
protocols {
  ospf {
    export bgp-routes-to-export;
    area 0.0.0.0 {
      interface t3-0/2/1.0;
    }
  }
  ldp {
    egress-policy bgp-routes-to-export;
    interface t3-0/2/1.0;
  }
}
}

```

On Router P0, enable LDP and OSPF in the same manner that you configured these protocols on Router P2. You will repeat these tasks on Router P1 and Router P3.

```

Router P0 [edit]
interfaces {
  so-0/1/0 {
    description "to pe1 so-0/0/0";
    unit 0 {
      family inet {
        address 192.255.197.22/30;
      }
      family mpls;
    }
  }
  so-1/0/0 {
    description "to p1 so-1/0/0";
    unit 0 {
      family inet {
        address 192.255.197.85/30;
      }
      family mpls;
    }
  }
  lo0 {
    unit 0 {
      family inet {
        address 10.255.255.173/32;
      }
    }
  }
}
routing-options {
  autonomous-system 300;
}
protocols {
  ospf {
    area 0.0.0.0 {

```

```
        interface so-0/1/0.0;
        interface so-1/0/0.0;
        interface lo0.0 {
            passive;
        }
    }
}
ldp {
    interface so-0/1/0.0;
    interface so-1/0/0.0;
}
}
```

On Router P1, enable LDP and the IGP used for transporting labels (OSPF in this case).

```
Router P1 [edit]
interfaces {
    so-0/0/0 {
        description "to pe2 so-0/2/0";
        unit 0 {
            family inet {
                address 192.255.197.74/30;
            }
            family mpls;
        }
    }
    so-1/0/0 {
        description "to p0 so-1/0/0";
        unit 0 {
            family inet {
                address 192.255.197.86/30;
            }
            family mpls;
        }
    }
    lo0 {
        unit 0 {
            family inet {
                address 10.255.255.174/32;
            }
        }
    }
}
routing-options {
    autonomous-system 300;
}
protocols {
    ospf {
        area 0.0.0.0 {
            interface so-0/0/0.0;
            interface so-1/0/0.0;
            interface lo0.0 {
                passive;
            }
        }
    }
}
```

```

ldp {
  interface so-0/0/0.0;
  interface so-1/0/0.0;
}

```

Core carrier Router PE2 is a mirror image of Router PE1. First, configure a master instance for OSPF, LDP, MPLS, and IBGP (with the **family inet-vpn** option) to connect Router PE2 to neighbor Router PE1. Next, implement multiple-instance LDP by establishing a secondary instance. Enable LDP and OSPF in this instance for Router PE2 to communicate with Router CE2. MPLS is not required in the secondary instance.

Finally, set up an outbound VRF policy that places all LDP traffic coming from Router CE2 into a BGP community, an export policy that sends this community traffic to Router PE1, and an inbound VRF policy that accepts similar BGP community traffic from Router PE1. This step tunnels the VPN provider's LDP traffic into the carrier's BGP session.

```

Router PE2 [edit]
interfaces {
  so-0/2/0 {
    description "to p1 so-0/0/0";
    unit 0 {
      family inet {
        address 192.255.197.73/30;
      }
      family mpls;
    }
  }
  t1-3/0/0 {
    description "to ce2 t1-0/0/0";
    unit 0 {
      family inet {
        address 192.255.197.37/30;
      }
      family mpls;
    }
  }
  lo0 {
    unit 0 {
      family inet {
        address 10.255.255.172/32;
      }
    }
  }
}
routing-options {
  autonomous-system 300;
}
protocols {
  mpls {
    interface t1-3/0/0.0;
  }
  bgp {
    group pe {
      type internal;

```

```
        local-address 10.255.255.172;
        family inet-vpn {
            unicast;
        }
        peer-as 300;
        neighbor 10.255.255.171;
    }
}
ospf {
    area 0.0.0.0 {
        interface so-0/2/0.0;
        interface lo0.0 {
            passive;
        }
    }
}
ldp {
    interface so-0/2/0.0;
}
}
policy-options {
    policy-statement vpn-provider-import {
        term 1 {
            from {
                protocol bgp;
                community vpn-provider-comm;
            }
            then accept;
        }
        term 2 {
            then reject;
        }
    }
    policy-statement vpn-provider-export {
        term 1 {
            from protocol ldp;
            then {
                community add vpn-provider-comm;
                accept;
            }
        }
        term 2 {
            then reject;
        }
    }
    policy-statement bgp-routes-to-export {
        term 1 {
            from {
                protocol bgp;
                community vpn-provider-comm;
            }
            then accept;
        }
        term 2 {
            then reject;
        }
    }
}
```



```

    }
    community vpn-provider-comm members target:300:200;
  }
  routing-instances {
    vpn-provider {
      instance-type vrf;
      interface t1-3/0/0.0;
      route-distinguisher 10.255.255.172:1;
      vrf-import vpn-provider-import;
      vrf-export vpn-provider-export;
      protocols {
        ospf {
          export bgp-routes-to-export;
          area 0.0.0.0 {
            interface t1-3/0/0.0;
          }
        }
        ldp {
          egress-policy bgp-routes-to-export;
          interface t1-3/0/0.0;
        }
      }
    }
  }
}

```

For Router CE2, configure LDP and OSPF as you did on Router CE1 and the transit P routers.

```

Router CE2 [edit]
interfaces {
  t1-0/0/0 {
    description "to pe2 t1-3/0/0";
    unit 0 {
      family inet {
        address 192.255.197.38/30;
      }
      family mpls;
    }
  }
  t3-0/3/3 {
    description "to p3 t3-0/0/3";
    unit 0 {
      family inet {
        address 192.255.198.26/30;
      }
      family mpls;
    }
  }
  lo0 {
    unit 0 {
      family inet {
        address 10.255.255.180/32;
      }
    }
  }
}
routing-options {

```

```

    autonomous-system 200;
  }
  protocols {
    ospf {
      area 0.0.0.0 {
        interface t1-0/0/0.0;
        interface t3-0/3/3.0;
        interface lo0.0 {
          passive;
        }
      }
    }
    ldp {
      interface t1-0/0/0.0;
      interface t3-0/3/3.0;
    }
  }
}

```

Since Router P3 is another core provider router, enable LDP and OSPF on all transit interfaces.

```

Router P3 [edit]
interfaces {
  t3-0/0/3 {
    description "to ce2 t3-0/3/3";
    unit 0 {
      family inet {
        address 192.255.198.25/30;
      }
      family mpls;
    }
  }
  t1-0/1/1 {
    description "to pe4 t1-0/1/1";
    unit 0 {
      family inet {
        address 192.255.198.37/30;
      }
      family mpls;
    }
  }
  lo0 {
    unit 0 {
      family inet {
        address 10.255.255.176/32;
      }
    }
  }
}
routing-options {
  autonomous-system 200;
}
protocols {
  ospf {
    area 0.0.0.0 {

```

```

        interface t3-0/0/3.0;
        interface t1-0/1/1.0;
        interface lo0.0 {
            passive;
        }
    }
    ldp {
        interface t3-0/0/3.0;
        interface t1-0/1/1.0;
    }
}

```

On Router PE4, complete the IBGP connection initiated on Router PE3 to connect the edge routers in AS 200. Also, enable LDP and MPLS on the **t1-0/0/1** interface pointing toward the VPN provider Router CE2 and establish an EBGP connection to Router CE4 through use of a VRF instance.

Finally, set up an outbound VRF policy that places all BGP traffic and directly connected interfaces into a BGP community and an inbound VRF policy that accepts similar BGP community traffic from Router PE3.

```

Router PE4 [edit]
interfaces {
    t3-0/0/3 {
        description to ce4 t3-0/0/3";
        unit 0 {
            family inet {
                address 192.255.198.21/30;
            }
            family mpls;
        }
    }
    t1-0/1/1 {
        unit 0 {
            family inet {
                address 192.255.198.38/30;
            }
            family mpls;
        }
    }
    lo0 {
        unit 0 {
            family inet {
                address 10.255.255.178/32;
            }
        }
    }
}
routing-options {
    autonomous-system 200;
}
protocols {
    mpls {
        interface t3-0/0/3.0;
    }
}

```

```
bgp {
  group internal {
    type internal;
    local-address 10.255.255.178;
    peer-as 200;
    neighbor 10.255.255.177 {
      family inet-vpn {
        unicast;
      }
    }
  }
}
ospf {
  area 0.0.0.0 {
    interface t1-0/1/1.0;
    interface lo0.0 {
      passive;
    }
  }
}
ldp {
  interface t1-0/1/1.0;
}
}
policy-options {
  policy-statement vpn-customer-import {
    term 1 {
      from {
        protocol bgp;
        community vpn-customer-comm;
      }
      then accept;
    }
    term 2 {
      then reject;
    }
  }
  policy-statement vpn-customer-export {
    term 1 {
      from protocol [bgp direct];
      then {
        community add vpn-customer-comm;
        accept;
      }
    }
    term 2 {
      then reject;
    }
  }
  community vpn-customer-comm members target:200:100;
}
routing-instances {
  vpn-customer {
    instance-type vrf;
    interface t3-0/0/3.0;
    route-distinguisher 10.255.255.178:1;
  }
}
```

```

vrf-import vpn-customer-import;
vrf-export vpn-customer-export;
protocols {
  bgp {
    group customer {
      type external;
      peer-as 100;
      as-override;
      neighbor 192.255.198.22;
    }
  }
}

```

Router CE4 is the destination VPN customer router. Configure EBGp between Router CE4 and the connected VPN provider Router PE4 to complete the configuration. Remember to advertise the loopback address into BGP by using a routing policy to allow IP reachability with Router CE3.

```

Router CE4 [edit]
interfaces {
  t3-0/0/3 {
    description "to pe4 t3-0/0/3";
    unit 0 {
      family inet {
        address 192.255.198.22/30;
      }
    }
  }
  lo0 {
    unit 0 {
      family inet {
        address 10.255.255.182/32;
        address 10.49.200.1/32;
      }
    }
  }
}
routing-options {
  static {
    route 10.49.200.0/24 reject;
    route 10.49.201.0/24 reject;
  }
  autonomous-system 100;
}
protocols {
  bgp {
    group provider {
      type external;
      export static-to-bgp;
      peer-as 200;
      neighbor 192.255.198.21;
    }
  }
}

```

```
policy-options {
  policy-statement static-to-bgp {
    term 1 {
      from {
        protocol static;
        route-filter 10.49.200.0/24 exact;
        route-filter 10.49.201.0/24 exact;
      }
      then accept;
    }
    term 2 {
      from protocol direct;
      then accept;
    }
    term 3 {
      then reject;
    }
  }
}
```

Verifying Your Work

To verify the proper operation of your multiple-instance LDP configuration, use the following commands:

- **show ldp database**
- **show ldp interface**
- **show ldp neighbor**
- **show ldp path**
- **show ldp route**
- **show ldp session**
- **show ldp statistics**

The display output for these commands is the same as in previous Junos OS Releases, except for one difference. An instance name can now be used as an argument.

If you include an instance name with these commands, you display information for the specified LDP instance. For example, the command **show ldp neighbor instance crockett** shows all the LDP neighbors for a VRF instance named **crockett**. Conversely, **show ldp neighbor** without an instance name displays the LDP neighbors associated with the master instance.

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

- Router CE3 Status on page 27
- Router PE3 Status on page 27
- Router CE1 Status on page 29
- Router PE1 Status on page 30

- Router PE2 Status on page 32
- Router CE2 Status on page 37
- Router PE4 Status on page 38
- Router CE4 Status on page 40

Router CE3 Status

user@CE3> show bgp summary

```
Groups: 1 Peers: 1 Down peers: 0
Table      Tot Paths  Act Paths Suppressed  History  Damp State  Pending
inet.0      10          5          0          0        0        0        0
Peer      AS      InPkt    OutPkt    OutQ    Flaps  Last Up/DwnState|#Active/Received/Damped...
192.255.198.13  200      440      433      0        0      3:34:34 5/10/0 0/0/0
```

user@CE3> show route protocol bgp

```
inet.0: 23 destinations, 28 routes (22 active, 0 holddown, 6 hidden)
+ = Active Route, - = Last Active, * = Both
10.49.200.0/24  * [BGP/170] 00:19:20, localpref 100
                  AS path: 200 200 I
                  > to 192.255.198.13 via so-1/2/0.0
10.49.200.1/32  * [BGP/170] 00:19:20, localpref 100
                  AS path: 200 200 I
                  > to 192.255.198.13 via so-1/2/0.0
10.49.201.0/24  * [BGP/170] 00:19:20, localpref 100
                  AS path: 200 200 I
                  > to 192.255.198.13 via so-1/2/0.0
10.255.255.182/32 * [BGP/170] 00:19:20, localpref 100
                  AS path: 200 200 I
                  > to 192.255.198.13 via so-1/2/0.0
192.255.198.20/30 * [BGP/170] 00:19:20, localpref 100
                  AS path: 200 I
                  > to 192.255.198.13 via so-1/2/0.0
```

Router PE3 Status

user@PE3> show bgp summary

```
Groups: 2 Peers: 2 Down peers: 0
Table      Tot Paths  Act Paths Suppressed  History  Damp State  Pending
bgp.l3vpn.0  6          6          0          0        0        0        0
Peer      AS      InPkt    OutPkt    OutQ    Flaps  Last Up/DwnState|#Active/Received/Damped...
192.255.198.14  100      432      441      0        0      3:34:55 Establ
  vpn-customer.inet.0: 5/6/0
10.255.255.178  200      62       63      0        2      27:23 Establ
  bgp.l3vpn.0: 6/6/0
  vpn-customer.inet.0: 5/6/0
```

user@PE3> show route protocol ldp

```
inet.0: 19 destinations, 20 routes (18 active, 0 holddown, 1 hidden)
inet.3: 5 destinations, 5 routes (5 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both
10.255.255.175/32 * [LDP/9] 03:35:45, metric 1
                  > via so-1/2/1.0
10.255.255.176/32 * [LDP/9] 00:29:32, metric 1
                  > via so-1/2/1.0, Push 100007
10.255.255.178/32 * [LDP/9] 00:29:32, metric 1
                  > via so-1/2/1.0, Push 100008
10.255.255.179/32 * [LDP/9] 03:34:39, metric 1
                  > via so-1/2/1.0, Push 100001
```

```

10.255.255.180/32 *[LDP/9] 03:31:15, metric 1
                  > via so-1/2/1.0, Push 100002
vpn-customer.inet.0: 12 destinations, 14 routes (12 active, 0 holddown, 0 hidden)
mpls.0: 13 destinations, 13 routes (13 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both
100000             *[LDP/9] 03:35:45, metric 1
                  > via so-1/2/1.0, Pop
100000(S=0)        *[LDP/9] 03:35:45, metric 1
                  > via so-1/2/1.0, Pop
100001             *[LDP/9] 03:34:39, metric 1
                  > via so-1/2/1.0, Swap 100001
100002             *[LDP/9] 03:31:15, metric 1
                  > via so-1/2/1.0, Swap 100002
100011             *[LDP/9] 00:29:32, metric 1
                  > via so-1/2/1.0, Swap 100007
100012             *[LDP/9] 00:29:32, metric 1
                  > via so-1/2/1.0, Swap 100008
bgp.l3vpn.0: 6 destinations, 6 routes (6 active, 0 holddown, 0 hidden)

user@PE3> show route protocol bgp
inet.0: 19 destinations, 20 routes (18 active, 0 holddown, 1 hidden)
inet.3: 5 destinations, 5 routes (5 active, 0 holddown, 0 hidden)
vpn-customer.inet.0: 12 destinations, 14 routes (12 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both
10.49.100.0/24     *[BGP/170] 03:34:59, MED 0, localpref 100
                  AS path: 100 I
                  > to 192.255.198.14 via so-1/2/0.0
10.49.100.1/32     *[BGP/170] 03:34:59, localpref 100
                  AS path: 100 I
                  > to 192.255.198.14 via so-1/2/0.0
10.49.101.0/24     *[BGP/170] 03:34:59, MED 0, localpref 100
                  AS path: 100 I
                  > to 192.255.198.14 via so-1/2/0.0
10.49.200.0/24     *[BGP/170] 00:26:39, MED 0, localpref 100, from 10.255.255.178
                  AS path: 100 I
                  > via so-1/2/1.0, Push 100019, Push 100008(top)
10.49.200.1/32     *[BGP/170] 00:26:39, localpref 100, from 10.255.255.178
                  AS path: 100 I
                  > via so-1/2/1.0, Push 100019, Push 100008(top)
10.49.201.0/24     *[BGP/170] 00:26:39, MED 0, localpref 100, from 10.255.255.178
                  AS path: 100 I
                  > via so-1/2/1.0, Push 100019, Push 100008(top)
10.255.255.181/32  *[BGP/170] 03:34:59, localpref 100
                  AS path: 100 I
                  > to 192.255.198.14 via so-1/2/0.0
10.255.255.182/32  *[BGP/170] 00:26:39, localpref 100, from 10.255.255.178
                  AS path: 100 I
                  > via so-1/2/1.0, Push 100019, Push 100008(top)
192.255.14.0/24    *[BGP/170] 03:34:59, localpref 100
                  AS path: 100 I
                  > to 192.255.198.14 via so-1/2/0.0
                  [BGP/170] 00:26:39, localpref 100, from 10.255.255.178
                  AS path: 100 I
                  > via so-1/2/1.0, Push 100019, Push 100008(top)
192.255.198.12/30  [BGP/170] 03:34:59, localpref 100
                  AS path: 100 I
                  > to 192.255.198.14 via so-1/2/0.0
192.255.198.20/30  *[BGP/170] 00:26:39, localpref 100, from 10.255.255.178
                  AS path: I
                  > via so-1/2/1.0, Push 100020, Push 100008(top)
mpls.0: 13 destinations, 13 routes (13 active, 0 holddown, 0 hidden)

```



```

bgp.l3vpn.0: 6 destinations, 6 routes (6 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both
10.255.255.178:1:10.49.200.0/24
    *[BGP/170] 00:27:27, MED 0, localpref 100, from 10.255.255.178
        AS path: 100 I
        > via so-1/2/1.0, Push 100019, Push 100008(top)
10.255.255.178:1:10.49.200.1/32
    *[BGP/170] 00:27:27, localpref 100, from 10.255.255.178
        AS path: 100 I
        > via so-1/2/1.0, Push 100019, Push 100008(top)
10.255.255.178:1:10.49.201.0/24
    *[BGP/170] 00:27:27, MED 0, localpref 100, from 10.255.255.178
        AS path: 100 I
        > via so-1/2/1.0, Push 100019, Push 100008(top)
10.255.255.178:1:10.255.255.182/32
    *[BGP/170] 00:27:27, localpref 100, from 10.255.255.178
        AS path: 100 I
        > via so-1/2/1.0, Push 100019, Push 100008(top)
10.255.255.178:1:192.255.14.0/24
    *[BGP/170] 00:27:27, localpref 100, from 10.255.255.178
        AS path: 100 I
        > via so-1/2/1.0, Push 100019, Push 100008(top)
10.255.255.178:1:192.255.198.20/30
    *[BGP/170] 00:27:27, localpref 100, from 10.255.255.178
        AS path: I
        > via so-1/2/1.0, Push 100020, Push 100008(top)

```

Router CE1 Status

```

user@CE1> show ldp neighbor

```

Address	Interface	Label space ID	Hold time
192.255.197.17	t3-0/1/0.0	192.255.197.17:0	11
192.255.198.2	so-1/2/0.0	10.255.255.175:0	14

```

user@CE1> show route

```

```

inet.0: 21 destinations, 23 routes (20 active, 0 holddown, 1 hidden)
+ = Active Route, - = Last Active, * = Both
0.0.0.0/0      *[Static/5] 07:53:10, metric 0
                Discard
10.255.255.175/32 *[OSPF/10] 00:31:44, metric 1
                > via so-1/2/0.0
10.255.255.176/32 *[OSPF/150] 00:31:44, metric 1, tag 3489661228
                > via t3-0/1/0.0
10.255.255.177/32 *[OSPF/10] 00:31:44, metric 2
                > via so-1/2/0.0
10.255.255.178/32 *[OSPF/150] 00:31:44, metric 1, tag 3489661228
                > via t3-0/1/0.0
10.255.255.179/32 *[Direct/0] 07:53:10
                > via lo0.0
10.255.255.180/32 *[OSPF/150] 00:31:44, metric 1, tag 3489661228
                > via t3-0/1/0.0
172.16.0.0/12   *[Static/5] 07:53:10
                > to 192.255.14.254 via fxp0.0
192.255.0.0/18  *[Static/5] 07:53:10
                > to 192.255.14.254 via fxp0.0
192.255.14.0/24 *[Direct/0] 07:53:10
                > via fxp0.0
192.255.14.179/32 *[Local/0] 07:53:10
                Local via fxp0.0
192.255.40.0/22 *[Static/5] 03:38:37

```

```

> to 192.255.14.254 via fxp0.0
192.255.64.0/18 * [Static/5] 03:38:37
> to 192.255.14.254 via fxp0.0
192.255.197.16/30 * [Direct/0] 03:37:42
> via t3-0/1/0.0
[OSPF/10] 00:31:44, metric 2
> via t3-0/1/0.0
192.255.197.18/32 * [Local/0] 07:52:01
Local via t3-0/1/0.0
192.255.198.0/30 * [Direct/0] 07:51:18
> via so-1/2/0.0
[OSPF/10] 00:31:44, metric 1
> via so-1/2/0.0
192.255.198.1/32 * [Local/0] 07:51:59
Local via so-1/2/0.0
192.255.198.8/30 * [OSPF/10] 00:31:44, metric 2
> via so-1/2/0.0
207.17.136.192/32 * [Static/5] 07:53:10
> to 192.255.14.254 via fxp0.0
224.0.0.5/32 * [OSPF/10] 07:53:14, metric 1
MultiRecv
inet.3: 5 destinations, 5 routes (5 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both
10.255.255.175/32 * [LDP/9] 01:00:52, metric 1
> via so-1/2/0.0
10.255.255.176/32 * [LDP/9] 00:33:24, metric 1
> via t3-0/1/0.0, Push 100020
10.255.255.177/32 * [LDP/9] 01:00:52, metric 1
> via so-1/2/0.0, Push 100000
10.255.255.178/32 * [LDP/9] 00:33:24, metric 1
> via t3-0/1/0.0, Push 100021
10.255.255.180/32 * [LDP/9] 01:00:52, metric 1
> via t3-0/1/0.0, Push 100015
mpls.0: 6 destinations, 6 routes (6 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both
100000 * [LDP/9] 03:38:31, metric 1
> via so-1/2/0.0, Pop
100000(S=0) * [LDP/9] 03:38:31, metric 1
> via so-1/2/0.0, Pop
100001 * [LDP/9] 03:38:31, metric 1
> via so-1/2/0.0, Swap 100000
100002 * [LDP/9] 03:35:06, metric 1
> via t3-0/1/0.0, Swap 100015
100007 * [LDP/9] 00:33:24, metric 1
> via t3-0/1/0.0, Swap 100020
100008 * [LDP/9] 00:33:24, metric 1
> via t3-0/1/0.0, Swap 100021

```

Router PE1 Status

```
user@PE1> show ldp neighbor instance vpn-provider
```

Address	Interface	Label space ID	Hold time
192.255.197.18	t3-0/2/1.0	10.255.255.179:0	11

```
user@PE1> show ldp database instance vpn-provider
```

```
Input label database, 192.255.197.17:0--10.255.255.179:0
Label Prefix
3 10.255.255.179/32
100002 10.255.255.180/32
100007 10.255.255.176/32
100001 10.255.255.177/32
```

```

100008      10.255.255.178/32
100000      10.255.255.175/32
Output label database, 192.255.197.17:0--10.255.255.179:0
Label      Prefix
100007      10.255.255.175/32
100020      10.255.255.176/32
100008      10.255.255.177/32
100021      10.255.255.178/32
100006      10.255.255.179/32
100015      10.255.255.180/32

```

```
user@PE1> show ldp interface instance vpn-provider
```

```

Interface      Label space ID      Nbr count      Next hello
t3-0/2/1.0      192.255.197.17:0      1              0

```

```
user@PE1> show ldp path instance vpn-provider
```

```

Output Session (label)      Input Session (label)
10.255.255.179:0(100006)(      ) 10.255.255.179:0(3)( )
10.255.255.179:0(100007)      10.255.255.179:0(100000)
10.255.255.179:0(100008)      10.255.255.179:0(100001)
10.255.255.179:0(100015)      ( )
10.255.255.179:0(100020)      ( )
10.255.255.179:0(100021)      ( )

```

```
user@PE1> show ldp route instance vpn-provider
```

```

Destination      Next-hop intf/lsp      Next-hop address
10.255.255.175/32  t3-0/2/1.0
10.255.255.176/32  so-0/0/0.0
10.255.255.177/32  t3-0/2/1.0
10.255.255.178/32  so-0/0/0.0
10.255.255.179/32  t3-0/2/1.0
10.255.255.180/32  so-0/0/0.0
192.255.197.16/30  t3-0/2/1.0
192.255.197.17/32
192.255.198.0/30   t3-0/2/1.0
192.255.198.8/30   t3-0/2/1.0
224.0.0.5/32

```

```
user@PE1> show ldp session instance vpn-provider
```

```

Address      State      Connection      Hold time
10.255.255.179      Operational      Open      24

```

```
user@PE1> show ldp statistics instance vpn-provider
```

```

Message type      Total      Last 5 seconds
Sent      Received      Sent      Received
Hello      2838      2839      1      2
Initialization      1      1      0      0
Keepalive      1240      1239      0      0
Notification      0      0      0      0
Address      1      1      0      0
Address withdraw      0      0      0      0
Label mapping      10      10      0      0
Label request      0      0      0      0
Label withdraw      4      4      0      0
Label release      4      4      0      0
Label abort      0      0      0      0
All UDP      2837      2839      1      2
All TCP      1258      1251      0      0
Event type      Total      Last 5 seconds
Sessions opened      1      0
Sessions closed      0      0
Topology changes      21      0

```

No router id	0	0
No address	0	0
No interface	0	0
No session	0	0
No adjacency	0	0
Unknown version	0	0
Malformed PDU	0	0
Malformed message	0	0
Unknown message type	0	0
Inappropriate message	0	0
Malformed TLV	0	0
Bad TLV value	0	0
Missing TLV	0	0
PDU too large	0	0
PDU too small	0	0

user@PE1> show ldp traffic-statistics instance vpn-provider

FEC	Type	Packets	Bytes	Shared
10.255.255.175/32	Transit	0	0	No
10.255.255.175/32	Ingress	0	0	No
10.255.255.176/32	Transit	0	0	No
10.255.255.177/32	Transit	2798	241984	No
10.255.255.177/32	Ingress	0	0	No
10.255.255.178/32	Transit	1365	125580	No
10.255.255.179/32	Transit	0	0	No
10.255.255.179/32	Ingress	2427	149076	No
10.255.255.180/32	Transit	0	0	No

user@PE1> show bgp summary

Groups: 1 Peers: 1 Down peers: 0

Table	Tot Paths	Act Paths	Suppressed	History	Damp	State	Pending
bgp.13vpn.0	3	3	0	0		0	0

Peer	AS	InPkt	OutPkt	OutQ	Flaps	Last
10.255.255.172	300	428	422	0	0	3:28:37 Establ

Up/DwnState|#Active/Received/Damped...

bgp.13vpn.0: 3/3/0

vpn-provider.inet.0: 3/3/0

Router PE2 Status

user@PE2> show ldp neighbor instance vpn-provider

Address	Interface	Label space ID	Hold time
192.255.197.38	tl-3/0/0.0	10.255.255.180:0	11

user@PE2> show ldp database instance vpn-provider

Input label database, 192.255.197.37:0--10.255.255.180:0

Label	Prefix
3	10.255.255.180/32
100003	10.255.255.177/32
100010	10.255.255.178/32
100009	10.255.255.176/32
100002	10.255.255.175/32
100004	10.255.255.179/32

Output label database, 192.255.197.37:0--10.255.255.180:0

Label	Prefix
100026	10.255.255.175/32
100028	10.255.255.179/32
100027	10.255.255.177/32
100021	10.255.255.180/32
100039	10.255.255.178/32
100037	10.255.255.176/32

```
user@PE2> show ldp interface instance vpn-provider
```

Interface	Label space ID	Nbr count	Next hello
t1-3/0/0.0	192.255.197.37:0		
1	1		

```
user@PE2> show ldp path instance vpn-provider
```

Output Session (label)	Input Session (label)
10.255.255.180:0(100021)() 10.255.255.180:0(3)()
10.255.255.180:0(100026)	()
10.255.255.180:0(100027)	()
10.255.255.180:0(100028)	()
10.255.255.180:0(100037)	10.255.255.180:0(100009)
10.255.255.180:0(100039)	10.255.255.180:0(100010)

```
user@PE2> show ldp route instance vpn-provider
```

Destination	Next-hop intf/lsp	Next-hop address
10.255.255.175/32	so-0/2/0.0	
10.255.255.176/32	t1-3/0/0.0	
10.255.255.177/32	so-0/2/0.0	
10.255.255.178/32	t1-3/0/0.0	
10.255.255.179/32	so-0/2/0.0	
10.255.255.180/32	t1-3/0/0.0	
192.255.197.36/30	t1-3/0/0.0	
192.255.197.37/32		
192.255.198.24/30	t1-3/0/0.0	
192.255.198.36/30	t1-3/0/0.0	
224.0.0.5/32		

```
user@PE2> show ldp session instance vpn-provider
```

Address	State	Connection	Hold time
10.255.255.180	Operational	Open	29

```
user@PE2> show ldp statistics instance vpn-provider
```

Message type	Total		Last 5 seconds	
	Sent	Received	Sent	Received
Hello	2948	2939	1	1
Initialization	1	1	0	0
Keepalive	1285	1285	0	0
Notification	0	0	0	0
Address	1	1	0	0
Address withdraw	0	0	0	0
Label mapping	10	10	0	0
Label request	0	0	0	0
Label withdraw	4	4	0	0
Label release	4	4	0	0
Label abort	0	0	0	0
All UDP	2947	2939	1	1
All TCP	1297	1299	0	0
Event type	Total		Last 5 seconds	
Sessions opened	1		0	
Sessions closed	0		0	
Topology changes	33		0	
No router id	0		0	
No address	0		0	
No interface	0		0	
No session	0		0	
No adjacency	0		0	
Unknown version	0		0	
Malformed PDU	0		0	
Malformed message	0		0	

Unknown message type	0	0
Inappropriate message	0	0
Malformed TLV	0	0
Bad TLV value	0	0
Missing TLV	0	0
PDU too large	0	0
PDU too small	0	0

```
user@PE2> show ldp traffic-statistics instance vpn-provider
```

FEC	Type	Packets	Bytes	Shared
10.255.255.175/32	Transit	0	0	No
10.255.255.176/32	Transit	0	0	No
10.255.255.176/32	Ingress	0	0	No
10.255.255.177/32	Transit	3131	274830	No
10.255.255.178/32	Transit	1966	178256	No
10.255.255.178/32	Ingress	0	0	No
10.255.255.179/32	Transit	1	44	No
10.255.255.180/32	Transit	0	0	No
10.255.255.180/32	Ingress	2330	144838	No

```
user@PE2> show bgp summary
```

```
Groups: 1 Peers: 1 Down peers: 0
Table      Tot Paths  Act Paths Suppressed  History  Damp State  Pending
inet.0      0          0          0          0          0      0      0
bgp.l3vpn.0 3          3          0          0          0      0      0
Peer      AS      InPkt  OutPkt  OutQ  Flaps Last
Up/DwnState|#Active/Received/Damped...
10.255.255.171 300      429      438      0      0      3:33:32 Establ
  bgp.l3vpn.0: 3/3/0
  vpn-provider.inet.0: 3/3/0
```

```
user@PE2> show route protocol bgp
```

```
inet.0: 18 destinations, 19 routes (17 active, 0 holddown, 1 hidden)
inet.3: 3 destinations, 3 routes (3 active, 0 holddown, 0 hidden)
vpn-provider.inet.0: 11 destinations, 15 routes (11 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both
10.255.255.175/32 *[BGP/170] 00:27:59, MED 1, localpref 100, from 10.255.255.171
```

```
AS path: I
```

```
> via so-0/2/0.0, Push 100012, Push 100028(top)
```

```
10.255.255.177/32 *[BGP/170] 00:27:59, MED 1, localpref 100, from 10.255.255.171
```

```
AS path: I
```

```
> via so-0/2/0.0, Push 100013, Push 100028(top)
```

```
10.255.255.179/32 *[BGP/170] 00:27:59, MED 1, localpref 100, from 10.255.255.171
```

```
AS path: I
```

```
> via so-0/2/0.0, Push 100014, Push 100028(top)
```

```
vpn-provider.inet.3: 3 destinations, 3 routes (3 active, 0 holddown, 0 hidden)
```

```
mpls.0: 11 destinations, 11 routes (11 active, 0 holddown, 0 hidden)
```

```
vpn-provider.mpls.0: 7 destinations, 7 routes (7 active, 0 holddown, 0 hidden)
```

```
bgp.l3vpn.0: 3 destinations, 3 routes (3 active, 0 holddown, 0 hidden)
```

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

```
10.255.255.171:1:10.255.255.175/32
```

```
*[BGP/170] 03:33:34, MED 1, localpref 100, from 10.255.255.171
```

```
AS path: I
```

```
> via so-0/2/0.0, Push 100012, Push 100028(top)
```

```
10.255.255.171:1:10.255.255.177/32
```

```
*[BGP/170] 03:33:34, MED 1, localpref 100, from 10.255.255.171
```

```

AS path: I
> via so-0/2/0.0, Push 100013, Push 100028(top)
10.255.255.171:1:10.255.255.179/32
*[BGP/170] 03:33:34, MED 1, localpref 100, from 10.255.255.171

```

```

AS path: I
> via so-0/2/0.0, Push 100014, Push 100028(top)
Address          Interface      Label space ID      Hold time
192.255.197.38   t1-3/0/0.0    10.255.255.180:0    11

```

```

user@PE2> show ldp database instance vpn-provider
Input label database, 192.255.197.37:0--10.255.255.180:0
  Label      Prefix
    3        10.255.255.180/32
 100003      10.255.255.177/32
 100010      10.255.255.178/32
 100009      10.255.255.176/32
 100002      10.255.255.175/32
 100004      10.255.255.179/32
Output label database, 192.255.197.37:0--10.255.255.180:0
  Label      Prefix
 100026      10.255.255.175/32
 100028      10.255.255.179/32
 100027      10.255.255.177/32
 100021      10.255.255.180/32
 100039      10.255.255.178/32
 100037      10.255.255.176/32

```

```

user@PE2> show ldp interface instance vpn-provider
Interface      Label space ID      Nbr count      Next hello
t1-3/0/0.0     192.255.197.37:0    1               1

```

```

user@PE2> show ldp path instance vpn-provider
Output Session (label)      Input Session (label)
10.255.255.180:0(100021)(   ) 10.255.255.180:0(3)( )
10.255.255.180:0(100026)      ( )
10.255.255.180:0(100027)      ( )
10.255.255.180:0(100028)      ( )
10.255.255.180:0(100037)      10.255.255.180:0(100009)
10.255.255.180:0(100039)      10.255.255.180:0(100010)

```

```

user@PE2> show ldp route instance vpn-provider
Destination      Next-hop intf/lsp      Next-hop address
10.255.255.175/32 so-0/2/0.0
10.255.255.176/32 t1-3/0/0.0
10.255.255.177/32 so-0/2/0.0
10.255.255.178/32 t1-3/0/0.0
10.255.255.179/32 so-0/2/0.0
10.255.255.180/32 t1-3/0/0.0
192.255.197.36/30 t1-3/0/0.0
192.255.197.37/32
192.255.198.24/30 t1-3/0/0.0
192.255.198.36/30 t1-3/0/0.0
224.0.0.5/32

```

```

user@PE2> show ldp session instance vpn-provider
Address      State      Connection      Hold time
10.255.255.180 Operational Open            29

```

```
user@PE2> show ldp statistics instance vpn-provider
```

Message type	Total		Last 5 seconds	
	Sent	Received	Sent	Received
Hello	2948	2939	1	1
Initialization	1	1	0	0
Keepalive	1285	1285	0	0
Notification	0	0	0	0
Address	1	1	0	0
Address withdraw	0	0	0	0
Label mapping	10	10	0	0
Label request	0	0	0	0
Label withdraw	4	4	0	0
Label release	4	4	0	0
Label abort	0	0	0	0
All UDP	2947	2939	1	1
All TCP	1297	1299	0	0
Event type	Total		Last 5 seconds	
Sessions opened	1		0	
Sessions closed	0		0	
Topology changes	33		0	
No router id	0		0	
No address	0		0	
No interface	0		0	
No session	0		0	
No adjacency	0		0	
Unknown version	0		0	
Malformed PDU	0		0	
Malformed message	0		0	
Unknown message type	0		0	
Inappropriate message	0		0	
Malformed TLV	0		0	
Bad TLV value	0		0	
Missing TLV	0		0	
PDU too large	0		0	
PDU too small	0		0	

```
user@PE2> show ldp traffic-statistics instance vpn-provider
```

FEC	Type	Packets	Bytes	Shared
10.255.255.175/32	Transit	0	0	No
10.255.255.176/32	Transit	0	0	No
10.255.255.176/32	Ingress	0	0	No
10.255.255.177/32	Transit	3131	274830	No
10.255.255.178/32	Transit	1966	178256	No
10.255.255.178/32	Ingress	0	0	No
10.255.255.179/32	Transit	1	44	No
10.255.255.180/32	Transit	0	0	No
10.255.255.180/32	Ingress	2330	144838	No

```
user@PE2> show bgp summary
```

```
Groups: 1 Peers: 1 Down peers: 0
```

Table	Tot Paths	Act Paths	Suppressed	History	Damp	State	Pending
inet.0	0	0	0	0	0	0	0
bgp.13vpn.0	3	3	0	0	0	0	0
Peer	AS	InPkt	OutPkt	OutQ	Flaps	Last	
Up/DwnState #Active/Received/Damped...							
10.255.255.171	300	429	438	0	0	3:33:32	Establ
bgp.13vpn.0: 3/3/0							
vpn-provider.inet.0: 3/3/0							

```
user@PE2> show route protocol bgp
```

```
inet.0: 18 destinations, 19 routes (17 active, 0 holddown, 1 hidden)
```



```

inet.3: 3 destinations, 3 routes (3 active, 0 holddown, 0 hidden)
vpn-provider.inet.0: 11 destinations, 15 routes (11 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both
10.255.255.175/32 * [BGP/170] 00:27:59, MED 1, localpref 100, from 10.255.255.171

    AS path: I
    > via so-0/2/0.0, Push 100012, Push 100028(top)
10.255.255.177/32 * [BGP/170] 00:27:59, MED 1, localpref 100, from 10.255.255.171

    AS path: I
    > via so-0/2/0.0, Push 100013, Push 100028(top)
10.255.255.179/32 * [BGP/170] 00:27:59, MED 1, localpref 100, from 10.255.255.171

    AS path: I
    > via so-0/2/0.0, Push 100014, Push 100028(top)
vpn-provider.inet.3: 3 destinations, 3 routes (3 active, 0 holddown, 0 hidden)
mpls.0: 11 destinations, 11 routes (11 active, 0 holddown, 0 hidden)
vpn-provider.mpls.0: 7 destinations, 7 routes (7 active, 0 holddown, 0 hidden)
bgp.l3vpn.0: 3 destinations, 3 routes (3 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both
10.255.255.171:1:10.255.255.175/32
    * [BGP/170] 03:33:34, MED 1, localpref 100, from 10.255.255.171

    AS path: I
    > via so-0/2/0.0, Push 100012, Push 100028(top)
10.255.255.171:1:10.255.255.177/32
    * [BGP/170] 03:33:34, MED 1, localpref 100, from 10.255.255.171

    AS path: I
    > via so-0/2/0.0, Push 100013, Push 100028(top)
10.255.255.171:1:10.255.255.179/32
    * [BGP/170] 03:33:34, MED 1, localpref 100, from 10.255.255.171

    AS path: I
    > via so-0/2/0.0, Push 100014, Push 100028(top)

```

Router CE2 Status

```

user@CE2> show ldp neighbor
Address          Interface      Label space ID      Hold time
192.255.197.37   t1-0/0/0.0    192.255.197.37:0    12
192.255.198.25   t3-0/3/3.0    10.255.255.176:0    13

user@CE2> show route
inet.0: 21 destinations, 23 routes (20 active, 0 holddown, 1 hidden)
+ = Active Route, - = Last Active, * = Both
0.0.0.0/0        * [Static/5] 07:53:49, metric 0
                  Discard
10.255.255.175/32 * [OSPF/150] 00:29:56, metric 1, tag 3489661228
                  > via t1-0/0/0.0
10.255.255.176/32 * [OSPF/10] 00:29:56, metric 2
                  > via t3-0/3/3.0
10.255.255.177/32 * [OSPF/150] 00:29:56, metric 1, tag 3489661228
                  > via t1-0/0/0.0
10.255.255.178/32 * [OSPF/10] 00:29:56, metric 67
                  > via t3-0/3/3.0
10.255.255.179/32 * [OSPF/150] 00:29:56, metric 1, tag 3489661228
                  > via t1-0/0/0.0
10.255.255.180/32 * [Direct/0] 07:53:49
                  > via lo0.0
172.16.0.0/12    * [Static/5] 07:53:49

```

```

> to 192.255.14.254 via fxp0.0
192.255.0.0/18 * [Static/5] 07:53:49
> to 192.255.14.254 via fxp0.0
192.255.14.0/24 * [Direct/0] 07:53:49
> via fxp0.0
192.255.14.180/32 * [Local/0] 07:53:49
Local via fxp0.0
192.255.40.0/22 * [Static/5] 06:07:28
> to 192.255.14.254 via fxp0.0
192.255.64.0/18 * [Static/5] 07:49:39
> to 192.255.14.254 via fxp0.0
192.255.197.36/30 * [Direct/0] 03:38:03
> via t1-0/0/0.0
[OSPF/10] 00:29:56, metric 65
> via t1-0/0/0.0
192.255.197.38/32 * [Local/0] 07:52:52
Local via t1-0/0/0.0
192.255.198.24/30 * [Direct/0] 03:33:17
> via t3-0/3/3.0
[OSPF/10] 00:29:56, metric 2
> via t3-0/3/3.0
192.255.198.26/32 * [Local/0] 07:52:49
Local via t3-0/3/3.0
192.255.198.36/30 * [OSPF/10] 00:29:56, metric 67
> via t3-0/3/3.0
207.17.136.192/32 * [Static/5] 07:53:49
> to 192.255.14.254 via fxp0.0
224.0.0.5/32 * [OSPF/10] 03:38:55, metric 1
MultiRecv
inet.3: 5 destinations, 5 routes (5 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both
10.255.255.175/32 * [LDP/9] 03:35:53, metric 1
> via t1-0/0/0.0, Push 100026
10.255.255.176/32 * [LDP/9] 00:34:13, metric 1
> via t3-0/3/3.0
10.255.255.177/32 * [LDP/9] 03:35:53, metric 1
> via t1-0/0/0.0, Push 100027
10.255.255.178/32 * [LDP/9] 00:34:13, metric 1
> via t3-0/3/3.0, Push 100014
10.255.255.179/32 * [LDP/9] 03:35:53, metric 1
> via t1-0/0/0.0, Push 100028
mpls.0: 6 destinations, 6 routes (6 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both
100002 * [LDP/9] 03:35:53, metric 1
> via t1-0/0/0.0, Swap 100026
100003 * [LDP/9] 03:35:53, metric 1
> via t1-0/0/0.0, Swap 100027
100004 * [LDP/9] 03:35:53, metric 1
> via t1-0/0/0.0, Swap 100028
100009 * [LDP/9] 00:34:13, metric 1
> via t3-0/3/3.0, Pop
100009(S=0) * [LDP/9] 00:34:13, metric 1
> via t3-0/3/3.0, Pop
100010 * [LDP/9] 00:34:13, metric 1
> via t3-0/3/3.0, Swap 100014

```

Router PE4 Status

```
user@PE4> show bgp summary
```

```

Groups: 2 Peers: 2 Down peers: 0
Table      Tot Paths  Act Paths Suppressed  History Damp State   Pending
bgp.l3vpn.0      6          6          0          0          0          0
inet.0          12         10          0          0          0          0
Peer
192.255.198.22   100        420        429         0          0      3:28:57 Establ
  vpn-customer.inet.0: 5/6/0
10.255.255.177   200        394        406         0          2      28:35 Establ
  bgp.l3vpn.0: 6/6/0
  vpn-customer.inet.0: 5/6/0

```

```
user@PE4> show route protocol bgp
```

```

inet.0: 20 destinations, 21 routes (19 active, 0 holddown, 1 hidden)
inet.3: 5 destinations, 5 routes (5 active, 0 holddown, 0 hidden)
vpn-customer.inet.0: 12 destinations, 14 routes (12 active, 0 holddown,
0 hidden)
+ = Active Route, - = Last Active, * = Both
10.49.100.0/24    * [BGP/170] 00:23:27, MED 0, localpref 100, from 10.255.255.177
                  AS path: 100 I
                  > via t1-0/1/1.0, Push 100013, Push 100012(top)
10.49.100.1/32    * [BGP/170] 00:23:27, localpref 100, from 10.255.255.177
                  AS path: 100 I
                  > via t1-0/1/1.0, Push 100013, Push 100012(top)
10.49.101.0/24    * [BGP/170] 00:23:27, MED 0, localpref 100, from 10.255.255.177
                  AS path: 100 I
                  > via t1-0/1/1.0, Push 100013, Push 100012(top)
10.49.200.0/24    * [BGP/170] 03:29:00, MED 0, localpref 100
                  AS path: 100 I
                  > to 192.255.198.22 via t3-0/0/3.0
10.49.200.1/32    * [BGP/170] 03:29:00, localpref 100
                  AS path: 100 I
                  > to 192.255.198.22 via t3-0/0/3.0
10.49.201.0/24    * [BGP/170] 03:29:00, MED 0, localpref 100
                  AS path: 100 I
                  > to 192.255.198.22 via t3-0/0/3.0
10.255.255.181/32 * [BGP/170] 00:23:27, localpref 100, from 10.255.255.177
                  AS path: 100 I
                  > via t1-0/1/1.0, Push 100013, Push 100012(top)
10.255.255.182/32 * [BGP/170] 03:29:00, localpref 100
                  AS path: 100 I
                  > to 192.255.198.22 via t3-0/0/3.0
192.255.14.0/24   * [BGP/170] 03:29:00, localpref 100
                  AS path: 100 I
                  > to 192.255.198.22 via t3-0/0/3.0
                  [BGP/170] 00:23:27, localpref 100, from 10.255.255.177
                  AS path: 100 I
                  > via t1-0/1/1.0, Push 100013, Push 100012(top)
192.255.198.12/30 * [BGP/170] 00:23:27, localpref 100, from 10.255.255.177
                  AS path: I
                  > via t1-0/1/1.0, Push 100014, Push 100012(top)
192.255.198.20/30 [BGP/170] 03:29:00, localpref 100
                  AS path: 100 I
                  > to 192.255.198.22 via t3-0/0/3.0
mpls.0: 13 destinations, 13 routes (13 active, 0 holddown, 0 hidden)
bgp.l3vpn.0: 6 destinations, 6 routes (6 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both
10.255.255.177:1:10.49.100.0/24
                  * [BGP/170] 00:28:38, MED 0, localpref 100, from 10.255.255.177
                  AS path: 100 I
                  > via t1-0/1/1.0, Push 100013, Push 100012(top)
10.255.255.177:1:10.49.100.1/32

```

```

* [BGP/170] 00:28:38, localpref 100, from 10.255.255.177
  AS path: 100 I
  > via t1-0/1/1.0, Push 100013, Push 100012(top)
10.255.255.177:1:10.49.101.0/24
* [BGP/170] 00:28:38, MED 0, localpref 100, from 10.255.255.177
  AS path: 100 I
  > via t1-0/1/1.0, Push 100013, Push 100012(top)
10.255.255.177:1:10.255.255.181/32
* [BGP/170] 00:28:38, localpref 100, from 10.255.255.177
  AS path: 100 I
  > via t1-0/1/1.0, Push 100013, Push 100012(top)
10.255.255.177:1:192.255.14.0/24
* [BGP/170] 00:28:38, localpref 100, from 10.255.255.177
  AS path: 100 I
  > via t1-0/1/1.0, Push 100013, Push 100012(top)
10.255.255.177:1:192.255.198.12/30
* [BGP/170] 00:28:38, localpref 100, from 10.255.255.177
  AS path: I
  > via t1-0/1/1.0, Push 100014, Push 100012(top)

```

user@PE4> show route protocol ldp

```

inet.0: 20 destinations, 21 routes (19 active, 0 holddown, 1 hidden)
inet.3: 5 destinations, 5 routes (5 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both
10.255.255.175/32 * [LDP/9] 00:29:08, metric 1
  > via t1-0/1/1.0, Push 100011
10.255.255.176/32 * [LDP/9] 00:29:08, metric 1
  > via t1-0/1/1.0
10.255.255.177/32 * [LDP/9] 00:29:08, metric 1
  > via t1-0/1/1.0, Push 100012
10.255.255.179/32 * [LDP/9] 00:29:08, metric 1
  > via t1-0/1/1.0, Push 100013
10.255.255.180/32 * [LDP/9] 00:29:08, metric 1
  > via t1-0/1/1.0, Push 100010
vpn-customer.inet.0: 12 destinations, 14 routes (12 active, 0 holddown, 0 hidden)
mpls.0: 13 destinations, 13 routes (13 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both
100014 * [LDP/9] 00:29:08, metric 1
  > via t1-0/1/1.0, Pop
100014(S=0) * [LDP/9] 00:29:08, metric 1
  > via t1-0/1/1.0, Pop
100015 * [LDP/9] 00:29:08, metric 1
  > via t1-0/1/1.0, Swap 100010
100016 * [LDP/9] 00:29:08, metric 1
  > via t1-0/1/1.0, Swap 100011
100017 * [LDP/9] 00:29:08, metric 1
  > via t1-0/1/1.0, Swap 100012
100018 * [LDP/9] 00:29:08, metric 1
  > via t1-0/1/1.0, Swap 100013
bgp.l3vpn.0: 6 destinations, 6 routes (6 active, 0 holddown, 0 hidden)

```

Router CE4 Status

user@CE4> show route protocol bgp

```

inet.0: 20 destinations, 25 routes (19 active, 0 holddown, 6 hidden)
+ = Active Route, - = Last Active, * = Both
10.49.100.0/24 * [BGP/170] 00:28:00, localpref 100
  AS path: 200 200 I
  > to 192.255.198.21 via t3-0/0/3.0
10.49.100.1/32 * [BGP/170] 00:28:00, localpref 100
  AS path: 200 200 I

```

```

> to 192.255.198.21 via t3-0/0/3.0
10.49.101.0/24 * [BGP/170] 00:28:00, localpref 100
                AS path: 200 200 I
> to 192.255.198.21 via t3-0/0/3.0
10.255.255.181/32 * [BGP/170] 00:28:00, localpref 100
                AS path: 200 200 I
> to 192.255.198.21 via t3-0/0/3.0
192.255.198.12/30 * [BGP/170] 00:28:00, localpref 100
                AS path: 200 I
> to 192.255.198.21 via t3-0/0/3.0

```

```
user@CE4> show bgp summary
```

```
Groups: 1 Peers: 1 Down peers: 0
```

Table	Tot Paths	Act Paths	Suppressed	History	Damp	State	Pending
inet.0	0	0	0	0	0	0	0

Peer	AS	InPkt	OutPkt	OutQ	Flaps	Last
Up/DwnState #Active/Received/Damped...						
192.255.198.21	200	426	421	0	0	3:28:20 5/10/0
	0/0/0					

Related Documentation

- *Junos VPNs Configuration Guide*
- *Junos MPLS Applications Configuration Guide*
- *Junos Policy Framework Configuration Guide*
- RFC 3017, *Carrying Label Information in BGP-4*
- RFC 3036, *LDP Specification*

PART 2

Index

- Index on page 45

Index

L

LDP

multiple instances

operational mode commands.....26

overview.....3

sample configuration.....7

system requirements.....4

S

system requirements

multiple instance LDP.....4

