

# Network Configuration Example

## Configuring RSVP-Signaled Point-to-Multipoint LSPs on Logical Systems

Release  
**13.1**



---

Published: 2013-02-08

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.

*Network Configuration Example Configuring RSVP-Signaled Point-to-Multipoint LSPs on Logical Systems*

Release 13.1

NCE0038

Copyright © 2013, Juniper Networks, Inc.

All rights reserved.

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

## YEAR 2000 NOTICE

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

## END USER LICENSE AGREEMENT

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

# Table of Contents

Introduction .....	1
Point-to-Multipoint LSPs Overview .....	1
Introduction to Logical Systems .....	3
Example: Configuring an RSVP-Signaled Point-to-Multipoint LSP on Logical Systems .....	5



## Introduction

---

This document describes how to configure an RSVP-signaled point-to-multipoint LSP on a Logical System.

## Point-to-Multipoint LSPs Overview


---

A point-to-multipoint MPLS LSP is an LSP with a single source and multiple destinations. By taking advantage of the MPLS packet replication capability of the network, point-to-multipoint LSPs avoid unnecessary packet replication at the ingress router. Packet replication takes place only when packets are forwarded to two or more different destinations requiring different network paths.

This process is illustrated in [Figure 1 on page 2](#). Router PE1 is configured with a point-to-multipoint LSP to Routers PE2, PE3, and PE4. When Router PE1 sends a packet on the point-to-multipoint LSP to Routers P1 and P2, Router P1 replicates the packet and forwards it to Routers PE2 and PE3. Router P2 sends the packet to Router PE4.

This feature is described in detail in the Internet drafts [draft-raggarwa-mpls-p2mp-te-02.txt](#) (expired February 2004), *Establishing Point to Multipoint MPLS TE LSPs*, [draft-ietf-mpls-rsvp-te-p2mp-02.txt](#), *Extensions to Resource Reservation Protocol-Traffic Engineering (RSVP-TE) for Point-to-Multipoint TE Label-Switched Paths (LSPs)*, and [draft-ietf-mpls-ldp-p2mp-10.txt](#), *Label Distribution Protocol Extensions for Point-to-Multipoint and Multipoint-to-Multipoint Label Switched Paths*.

Figure 1: Point-to-Multipoint LSPs



The following are some of the properties of point-to-multipoint LSPs:

- A point-to-multipoint LSP enables you to use MPLS for point-to-multipoint data distribution. This functionality is similar to that provided by IP multicast.
- You can add and remove branch LSPs from a main point-to-multipoint LSP without disrupting traffic. The unaffected parts of the point-to-multipoint LSP continue to function normally.
- You can configure a node to be both a transit and an egress router for different branch LSPs of the same point-to-multipoint LSP.
- You can enable link protection on a point-to-multipoint LSP. Link protection can provide a bypass LSP for each of the branch LSPs that make up the point-to-multipoint LSP. If any of the primary paths fail, traffic can be quickly switched to the bypass.
- You can configure branch LSPs either statically, dynamically, or as a combination of static and dynamic LSPs.
- You can enable graceful Routing Engine switchover (GRES) and graceful restart for point-to-multipoint LSPs at ingress and egress routers. The point-to-multipoint LSPs must be configured using either static routes or circuit cross-connect (CCC). GRES and graceful restart allow the traffic to be forwarded at the Packet Forwarding Engine based on the old state while the control plane recovers. Feature parity for GRES and graceful restart for MPLS point-to-multipoint LSPs on the Junos Trio chipset is supported in Junos OS Releases 11.1R2, 11.2R2, and 11.4.

---

**Related  
Documentation**

- Junos OS High Availability Configuration Guide
- [Example: Configuring an RSVP-Signaled Point-to-Multipoint LSP on Logical Systems on page 5](#)
- Example: NG-VPLS Using Point-to-Multipoint LSPs
- Example: Configuring Point-to-Multipoint LDP LSPs as the Data Plane for Intra-AS MBGP MVPNs

## Introduction to Logical Systems

---

For many years, engineers have combined power supplies, routing hardware and software, forwarding hardware and software, and physical interfaces into a networking device known as a router. Networking vendors have created large routers and small routers, but all routers have been placed into service as individual devices. As a result, the router has been considered a single physical device for most of its history.

The concept of logical systems breaks with this tradition. With the Junos<sup>®</sup> operating system (Junos OS), you can partition a single router into multiple logical devices that perform independent routing tasks. Because logical systems perform a subset of the tasks once handled by the main router, logical systems offer an effective way to maximize the use of a single routing or switching platform.



**NOTE:** Beginning with Junos OS Release 9.3, the logical router feature has been renamed logical system.

All configuration statements, operational commands, show command output, error messages, log messages, and SNMP MIB objects that contain the string logical-router have been changed to logical-system.

Traditionally, service provider network design requires multiple layers of switches and routers. These devices transport packet traffic between customers. As seen on the left side of [Figure 2 on page 4](#), access devices are connected to edge devices, which are in turn connected to core devices.

However, this complexity can lead to challenges in maintenance, configuration, and operation. To reduce such complexity, Juniper Networks supports logical systems. Logical systems perform a subset of the actions of the main router and have their own unique routing tables, interfaces, policies, and routing instances. As shown on the right side of [Figure 2 on page 4](#), a set of logical systems within a single router can handle the functions previously performed by several small routers.


**Figure 2: Logical Systems Concepts**

Figure 3 on page 4 shows the Junos OS architecture without logical systems configured. Figure 4 on page 5 shows the Junos OS architecture when logical systems are configured. Note that each logical system runs its own routing protocol process (rpd).



**Figure 3: Junos OS Without Logical Systems**



Figure 4: Junos OS With Logical Systems



**Related Documentation**

- Logical Systems Operations and Restrictions
- Junos OS Features That Are Supported on Logical Systems

## Example: Configuring an RSVP-Signaled Point-to-Multipoint LSP on Logical Systems

In this example, multiple logical systems in a physical router act as a collection of paths for an RSVP-signaled point-to-multipoint LSP. The logical systems are chained together and connected internally over a series of logical tunnel (**lt**) interfaces.

- [Requirements on page 5](#)
- [Overview on page 6](#)
- [Configuration on page 7](#)
- [Verification on page 26](#)

## Requirements

This example uses the following hardware and software components:

- One MX Series router running logical systems. You do not need to use an MX Series router for the logical systems. You can use any Juniper Networks router that supports logical systems.
- On the MX Series router, the logical systems are connected using logical tunnel (**lt**) interfaces. For more information, see [Example: Connecting Logical Systems Within the Same Router Using Logical Tunnel Interfaces on MX Series Routers](#) and [Example: Connecting Logical Systems Within the Same Router Using Logical Tunnel Interfaces](#). An alternative to using **lt** interfaces is to create external back-to-back interconnections between ports on the router.

- Four customer-edge (CE) devices running in separate physical devices. You do not need to use routers for the CE devices. For example, the CE devices can be EX Series Ethernet Switches.
- Junos OS Release 12.1 or later running on the MX Series router.

On M Series Multiservice Edge and T Series Core Routers, you can create an **lt** interface if you have a Tunnel Services PIC installed on an Enhanced FPC in your routing platform.

On M40e routers, you can create an **lt** interface if you have a Tunnel Services PIC. (An Enhanced FPC is not required.)

On an M7i router, **lt** interfaces can be created by using the integrated Adaptive Services Module.

On an MX Series router, as is shown in this example, the master administrator can configure **lt** interfaces by including the **tunnel-services** statement at the **[edit chassis fpc slot-number pic number]** hierarchy level.

## Overview

In this example, the logical systems serve as the transit, branch, and leaf nodes of a single point-to-multipoint LSP. Logical system LS1 is the ingress node. The branches go from LS1 to LS5, LS1 to LS7, and LS1 to LS4. Static unicast routes on the ingress node (LS1) point to the egress nodes.

The following topologies are supported:


- A single logical system in a physical router. The logical system is one node in an RSVP-signaled point-to-multipoint LSP.
- Multiple logical systems in a physical router, with each logical system acting as a label-switched router (LSR). The multiple logical systems can be unconnected, connected to each other internally with **lt** interfaces, or connected to each other externally with back-to-back connections.
- One RSVP-signaled point-to-multipoint LSP, with some nodes being logical systems and other nodes being physical routers.

---

### Topology Diagram

Figure 5 on page 7 shows the topology used in this example.

Figure 5: RSVP-Signaled Point-to-Multipoint LSP on Logical Systems



## Configuration

- [Configuring the MX Series Router to Support Logical Tunnel Interfaces on page 9](#)
- [Configuring the Ingress LSR \(Logical System LS1\) on page 9](#)
- [Configuring the Transit and Egress LSRs \(Logical Systems LS2, LS3, LS4, LS5, LS6, and LS7\) on page 11](#)
- [Configuring Device CE1 on page 22](#)
- [Configuring Device CE2 on page 23](#)
- [Configuring Device CE3 on page 24](#)
- [Configuring Device CE4 on page 25](#)

### CLI Quick Configuration

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

```
Router R1
set logical-systems LS1 interfaces ge-2/0/2 unit 0 description LS1-to-CE1
set logical-systems LS1 interfaces ge-2/0/2 unit 0 family inet address 10.0.244.10/30
set logical-systems LS1 interfaces lt-2/0/10 unit 1 description LS1-to-LS2
set logical-systems LS1 interfaces lt-2/0/10 unit 1 encapsulation ethernet
set logical-systems LS1 interfaces lt-2/0/10 unit 1 peer-unit 2
set logical-systems LS1 interfaces lt-2/0/10 unit 1 family inet address 2.2.2.1/24
set logical-systems LS1 interfaces lt-2/0/10 unit 1 family mpls
set logical-systems LS1 interfaces lt-2/0/10 unit 8 description LS1-to-LS6
set logical-systems LS1 interfaces lt-2/0/10 unit 8 encapsulation ethernet
set logical-systems LS1 interfaces lt-2/0/10 unit 8 peer-unit 6
set logical-systems LS1 interfaces lt-2/0/10 unit 8 family inet address 6.6.6.1/24
set logical-systems LS1 interfaces lt-2/0/10 unit 8 family mpls
set logical-systems LS1 interfaces lt-2/0/10 unit 9 description LS1-to-LS3
```

```
set logical-systems LS1 interfaces lt-2/0/10 unit 9 encapsulation ethernet
set logical-systems LS1 interfaces lt-2/0/10 unit 9 peer-unit 3
set logical-systems LS1 interfaces lt-2/0/10 unit 9 family inet address 3.3.3.1/24
set logical-systems LS1 interfaces lt-2/0/10 unit 9 family mpls
set logical-systems LS1 interfaces lo0 unit 1 family inet address 100.10.10.10/32
set logical-systems LS1 protocols rsvp interface lt-2/0/10.1
set logical-systems LS1 protocols rsvp interface lt-2/0/10.8
set logical-systems LS1 protocols rsvp interface lt-2/0/10.9
set logical-systems LS1 protocols rsvp interface lo0.1
set logical-systems LS1 protocols mpls traffic-engineering bgp-igp
set logical-systems LS1 protocols mpls label-switched-path LS1-LS5 to 100.50.50.50
set logical-systems LS1 protocols mpls label-switched-path LS1-LS5 p2mp p2mp1
set logical-systems LS1 protocols mpls label-switched-path LS1-LS7 to 100.70.70.70
set logical-systems LS1 protocols mpls label-switched-path LS1-LS7 p2mp p2mp1
set logical-systems LS1 protocols mpls label-switched-path LS1-LS4 to 100.40.40.40
set logical-systems LS1 protocols mpls label-switched-path LS1-LS4 p2mp p2mp1
set logical-systems LS1 protocols mpls interface lt-2/0/10.1
set logical-systems LS1 protocols mpls interface lt-2/0/10.8
set logical-systems LS1 protocols mpls interface lt-2/0/10.9
set logical-systems LS1 protocols mpls interface lo0.1
set logical-systems LS1 protocols ospf traffic-engineering
set logical-systems LS1 protocols ospf area 0.0.0.0 interface ge-2/0/2.0
set logical-systems LS1 protocols ospf area 0.0.0.0 interface lt-2/0/10.1
set logical-systems LS1 protocols ospf area 0.0.0.0 interface lt-2/0/10.8
set logical-systems LS1 protocols ospf area 0.0.0.0 interface lt-2/0/10.9
set logical-systems LS1 protocols ospf area 0.0.0.0 interface lo0.1
set logical-systems LS1 routing-options static route 5.5.5.0/24 p2mp-lsp-next-hop p2mp1
set logical-systems LS1 routing-options static route 7.7.7.0/24 p2mp-lsp-next-hop p2mp1
set logical-systems LS1 routing-options static route 4.4.4.0/24 p2mp-lsp-next-hop p2mp1
set logical-systems LS1 routing-options router-id 100.10.10.10
```

Device CE1	<pre>set interfaces ge-1/3/2 unit 0 family inet address 10.0.244.9/30 set interfaces ge-1/3/2 unit 0 description CE1-to-LS1 set routing-options static route 10.0.104.8/30 next-hop 10.0.244.10 set routing-options static route 10.0.134.8/30 next-hop 10.0.244.10 set routing-options static route 10.0.224.8/30 next-hop 10.0.244.10</pre>
Device CE2	<pre>set interfaces ge-1/3/3 unit 0 family inet address 10.0.224.9/30 set interfaces ge-1/3/3 unit 0 description CE2-to-LS5 set routing-options static route 10.0.244.8/30 next-hop 10.0.224.10</pre>
Device CE3	<pre>set interfaces ge-2/0/1 unit 0 family inet address 10.0.134.9/30 set interfaces ge-2/0/1 unit 0 description CE3-to-LS7 set routing-options static route 10.0.244.8/30 next-hop 10.0.134.10</pre>
Device CE4	<pre>set interfaces ge-3/1/3 unit 0 family inet address 10.0.104.10/30 set interfaces ge-3/1/3 unit 0 description CE4-to-LS4 set routing-options static route 10.0.244.8/30 next-hop 10.0.104.9</pre>

---

## Configuring the MX Series Router to Support Logical Tunnel Interfaces

### Step-by-Step Procedure

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

This procedure is required for MX Series routers only. If you have an M Series or T Series router, skip this procedure.

To enable **lt** interfaces on the MX Series router:

1. Run the **show chassis fpc** command to verify that the router has a DPC, MPC, or MIC installed and is in the online state.

```
user@host> show chassis fpc
```

Slot	State	Temp (C)	CPU Utilization (%) Total	Interrupt	Memory DRAM (MB)	Utilization (%) Heap	Buffer
0	Empty						
1	Empty						
2	Online	31	3	0	1024	14	21

This output shows that slot 0 and slot 1 are empty. Slot 2 is online.

2. Configure FPC slot 2 to support **lt** interfaces.

```
[edit]
user@host# set chassis fpc 2 pic 0 tunnel-services bandwidth 1g
```

This command creates several tunnel interface types, including **gr**, **ip**, and **lt**. For this example, the important one is the **lt** interface.

3. Commit the configuration.

```
[edit]
user@host# commit
user@host# exit
```

4. Run the **show interfaces terse** command to verify that the router has an **lt** interface.

```
user@host> show interfaces terse
```

Interface	Admin	Link	Proto	Local	Remote
...					
gr-2/0/10	up	up			
ip-2/0/10	up	up			
lt-2/0/10	up	up			
...					

---

## Configuring the Ingress LSR (Logical System LS1)

### Step-by-Step Procedure

To configure Logical System LS1:

1. From the main router, configure the logical system.

```
[edit]
user@R1# set logical-systems LS1
```

2. Commit the configuration.

```
[edit]
```

```
user@R1# commit
user@R1# exit
```

3. Set the CLI to view the logical system.

```
user@R1> set cli logical-system LS1
Logical system: LS1
```

```
user@R1:LS1>
```

4. Configure the interfaces, interface encapsulation, and protocol families.

```
[edit]
user@R1:LS1# edit interfaces
[edit interfaces]
user@R1:LS1# set ge-2/0/2 unit 0 description R1-to-CE1
user@R1:LS1# set ge-2/0/2 unit 0 family inet address 10.0.244.10/30
user@R1:LS1# set lt-2/0/10 unit 1 description LS1-to-LS2
user@R1:LS1# set lt-2/0/10 unit 1 encapsulation ethernet
user@R1:LS1# set lt-2/0/10 unit 1 peer-unit 2
user@R1:LS1# set lt-2/0/10 unit 1 family inet address 2.2.2.1/24
user@R1:LS1# set lt-2/0/10 unit 1 family mpls
user@R1:LS1# set lt-2/0/10 unit 8 description LS1-to-LS6
user@R1:LS1# set lt-2/0/10 unit 8 encapsulation ethernet
user@R1:LS1# set lt-2/0/10 unit 8 peer-unit 6
user@R1:LS1# set lt-2/0/10 unit 8 family inet address 6.6.6.1/24
user@R1:LS1# set lt-2/0/10 unit 8 family mpls
user@R1:LS1# set lt-2/0/10 unit 9 description LS1-to-LS3
user@R1:LS1# set lt-2/0/10 unit 9 encapsulation ethernet
user@R1:LS1# set lt-2/0/10 unit 9 peer-unit 3
user@R1:LS1# set lt-2/0/10 unit 9 family inet address 3.3.3.1/24
user@R1:LS1# set lt-2/0/10 unit 9 family mpls
user@R1:LS1# set lo0 unit 1 family inet address 100.10.10.10/32
user@R1:LS1# exit
```

5. Enable RSVP, MPLS, and OSPF on the interfaces.

```
[edit]
user@R1:LS1# edit protocols
[edit protocols]
user@R1:LS1# set rsvp interface lt-2/0/10.1
user@R1:LS1# set rsvp interface lt-2/0/10.8
user@R1:LS1# set rsvp interface lt-2/0/10.9
user@R1:LS1# set rsvp interface lo0.1
user@R1:LS1# set mpls interface lt-2/0/10.1
user@R1:LS1# set mpls interface lt-2/0/10.8
user@R1:LS1# set mpls interface lt-2/0/10.9
user@R1:LS1# set mpls interface lo0.1
user@R1:LS1# set ospf area 0.0.0.0 interface ge-2/0/2.0
user@R1:LS1# set ospf area 0.0.0.0 interface lt-2/0/10.1
user@R1:LS1# set ospf area 0.0.0.0 interface lt-2/0/10.8
user@R1:LS1# set ospf area 0.0.0.0 interface lt-2/0/10.9
user@R1:LS1# set ospf area 0.0.0.0 interface lo0.1
```

6. Configure the MPLS point-to-multipoint LSPs.

```
[edit protocols]
user@R1:LS1# set mpls label-switched-path LS1-LS5 to 100.50.50.50
user@R1:LS1# set mpls label-switched-path LS1-LS5 p2mp p2mp1
```

---

```
user@R1:LS1# set mpls label-switched-path LS1-LS7 to 100.70.70.70
user@R1:LS1# set mpls label-switched-path LS1-LS7 p2mp p2mp1
user@R1:LS1# set mpls label-switched-path LS1-LS4 to 100.40.40.40
user@R1:LS1# set mpls label-switched-path LS1-LS4 p2mp p2mp1
```

7. Enable MPLS to perform traffic engineering for OSPF.

```
[edit protocols]
user@R1:LS1# set mpls traffic-engineering bgp-igp
user@R1:LS1# exit
```

This causes the ingress routes to be installed in the **inet.0** routing table. By default, MPLS performs traffic engineering for BGP only. You need to enable MPLS traffic engineering on the ingress LSR only.

8. Enable traffic engineering for OSPF.

```
[edit protocols]
user@R1:LS1# set ospf traffic-engineering
user@R1:LS1# exit
```

This causes the shortest-path first (SPF) algorithm to take into account the LSPs configured under MPLS.

9. Configure the router ID.

```
[edit]
user@R1:LS1# edit routing-options
[edit routing-options]
user@R1:LS1# set router-id 100.10.10.10
```

10. Configure static IP unicast routes with the point-to-multipoint LSP name as the next hop for each route.

```
[edit routing-options]
user@R1:LS1# set static route 5.5.5.0/24 p2mp-lsp-next-hop p2mp1
user@R1:LS1# set static route 7.7.7.0/24 p2mp-lsp-next-hop p2mp1
user@R1:LS1# set static route 4.4.4.0/24 p2mp-lsp-next-hop p2mp1
user@R1:LS1# exit
```

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

```
[edit]
user@R1:LS1# commit
```

### Configuring the Transit and Egress LSRs (Logical Systems LS2, LS3, LS4, LS5, LS6, and LS7)

---

#### Step-by-Step Procedure

To configure the transit and egress LSRs:

1. Configure the interfaces, interface encapsulation, and protocol families.

```
[edit]
user@R1# edit logical-systems
[edit logical-systems]
user@R1# set LS2 interfaces lt-2/0/10 unit 2 description LS2-to-LS1
user@R1# set LS2 interfaces lt-2/0/10 unit 2 encapsulation ethernet
user@R1# set LS2 interfaces lt-2/0/10 unit 2 peer-unit 1
```

```
user@R1# set LS2 interfaces lt-2/0/10 unit 2 family inet address 2.2.2.2/24
user@R1# set LS2 interfaces lt-2/0/10 unit 2 family mpls
user@R1# set LS2 interfaces lt-2/0/10 unit 10 description LS2-to-LS5
user@R1# set LS2 interfaces lt-2/0/10 unit 10 encapsulation ethernet
user@R1# set LS2 interfaces lt-2/0/10 unit 10 peer-unit 5
user@R1# set LS2 interfaces lt-2/0/10 unit 10 family inet address 5.5.5.1/24
user@R1# set LS2 interfaces lt-2/0/10 unit 10 family mpls
user@R1# set LS2 interfaces lo0 unit 2 family inet address 100.20.20.20/32
user@R1# set LS3 interfaces lt-2/0/10 unit 3 description LS3-to-LS1
user@R1# set LS3 interfaces lt-2/0/10 unit 3 encapsulation ethernet
user@R1# set LS3 interfaces lt-2/0/10 unit 3 peer-unit 9
user@R1# set LS3 interfaces lt-2/0/10 unit 3 family inet address 3.3.3.2/24
user@R1# set LS3 interfaces lt-2/0/10 unit 3 family mpls
user@R1# set LS3 interfaces lt-2/0/10 unit 12 description LS3-to-LS4
user@R1# set LS3 interfaces lt-2/0/10 unit 12 encapsulation ethernet
user@R1# set LS3 interfaces lt-2/0/10 unit 12 peer-unit 4
user@R1# set LS3 interfaces lt-2/0/10 unit 12 family inet address 4.4.4.1/24
user@R1# set LS3 interfaces lt-2/0/10 unit 12 family mpls
user@R1# set LS3 interfaces lo0 unit 3 family inet address 100.30.30.30/32
user@R1# set LS4 interfaces ge-2/0/0 unit 0 description R1-to-CE4
user@R1# set LS4 interfaces ge-2/0/0 unit 0 family inet address 10.0.104.9/30
user@R1# set LS4 interfaces lt-2/0/10 unit 4 description LS4-to-LS3
user@R1# set LS4 interfaces lt-2/0/10 unit 4 encapsulation ethernet
user@R1# set LS4 interfaces lt-2/0/10 unit 4 peer-unit 12
user@R1# set LS4 interfaces lt-2/0/10 unit 4 family inet address 4.4.4.2/24
user@R1# set LS4 interfaces lt-2/0/10 unit 4 family mpls
user@R1# set LS4 interfaces lo0 unit 4 family inet address 100.40.40.40/32
user@R1# set LS5 interfaces ge-2/0/3 unit 0 description LS1-to-CE1
user@R1# set LS5 interfaces ge-2/0/3 unit 0 family inet address 10.0.224.10/30
user@R1# set LS5 interfaces lt-2/0/10 unit 5 description LS5-to-LS2
user@R1# set LS5 interfaces lt-2/0/10 unit 5 encapsulation ethernet
user@R1# set LS5 interfaces lt-2/0/10 unit 5 peer-unit 10
user@R1# set LS5 interfaces lt-2/0/10 unit 5 family inet address 5.5.5.2/24
user@R1# set LS5 interfaces lt-2/0/10 unit 5 family mpls
user@R1# set LS5 interfaces lo0 unit 5 family inet address 100.50.50.50/32
user@R1# set LS6 interfaces lt-2/0/10 unit 6 description LS6-to-LS1
user@R1# set LS6 interfaces lt-2/0/10 unit 6 encapsulation ethernet
user@R1# set LS6 interfaces lt-2/0/10 unit 6 peer-unit 8
user@R1# set LS6 interfaces lt-2/0/10 unit 6 family inet address 6.6.6.2/24
user@R1# set LS6 interfaces lt-2/0/10 unit 6 family mpls
user@R1# set LS6 interfaces lt-2/0/10 unit 11 description LS6-to-LS7
user@R1# set LS6 interfaces lt-2/0/10 unit 11 encapsulation ethernet
user@R1# set LS6 interfaces lt-2/0/10 unit 11 peer-unit 7
user@R1# set LS6 interfaces lt-2/0/10 unit 11 family inet address 7.7.7.1/24
user@R1# set LS6 interfaces lt-2/0/10 unit 11 family mpls
user@R1# set LS6 interfaces lo0 unit 6 family inet address 100.60.60.60/32
user@R1# set LS7 interfaces ge-2/0/1 unit 0 description R1-to-CE3
user@R1# set LS7 interfaces ge-2/0/1 unit 0 family inet address 10.0.134.10/30
user@R1# set LS7 interfaces lt-2/0/10 unit 7 description LS7-to-LS6
user@R1# set LS7 interfaces lt-2/0/10 unit 7 encapsulation ethernet
user@R1# set LS7 interfaces lt-2/0/10 unit 7 peer-unit 11
user@R1# set LS7 interfaces lt-2/0/10 unit 7 family inet address 7.7.7.2/24
user@R1# set LS7 interfaces lt-2/0/10 unit 7 family mpls
user@R1# set LS7 interfaces lo0 unit 7 family inet address 100.70.70.70/32
```



- 
2. Enable RSVP, MPLS, and OSPF on the interfaces.

```
[edit]
user@R1# edit logical-systems
[edit logical-systems]
user@R1# set LS2 protocols rsvp interface lt-2/0/10.2
user@R1# set LS2 protocols rsvp interface lt-2/0/10.10
user@R1# set LS2 protocols rsvp interface lo0.2
user@R1# set LS2 protocols mpls interface lt-2/0/10.2
user@R1# set LS2 protocols mpls interface lt-2/0/10.10
user@R1# set LS2 protocols mpls interface lo0.2
user@R1# set LS2 protocols ospf area 0.0.0.0 interface lt-2/0/10.2
user@R1# set LS2 protocols ospf area 0.0.0.0 interface lt-2/0/10.10
user@R1# set LS2 protocols ospf area 0.0.0.0 interface lo0.2
user@R1# set LS3 protocols rsvp interface lt-2/0/10.3
user@R1# set LS3 protocols rsvp interface lt-2/0/10.12
user@R1# set LS3 protocols rsvp interface lo0.3
user@R1# set LS3 protocols mpls interface lt-2/0/10.3
user@R1# set LS3 protocols mpls interface lt-2/0/10.12
user@R1# set LS3 protocols mpls interface lo0.3
user@R1# set LS3 protocols ospf area 0.0.0.0 interface lt-2/0/10.3
user@R1# set LS3 protocols ospf area 0.0.0.0 interface lt-2/0/10.12
user@R1# set LS3 protocols ospf area 0.0.0.0 interface lo0.3
user@R1# set LS4 protocols rsvp interface lt-2/0/10.4
user@R1# set LS4 protocols rsvp interface lo0.4
user@R1# set LS4 protocols mpls interface lt-2/0/10.4
user@R1# set LS4 protocols mpls interface lo0.4
user@R1# set LS4 protocols ospf area 0.0.0.0 interface ge-2/0/0.0
user@R1# set LS4 protocols ospf area 0.0.0.0 interface lt-2/0/10.4
user@R1# set LS4 protocols ospf area 0.0.0.0 interface lo0.4
user@R1# set LS5 protocols rsvp interface lt-2/0/10.5
user@R1# set LS5 protocols rsvp interface lo0.5
user@R1# set LS5 protocols mpls interface lt-2/0/10.5
user@R1# set LS5 protocols mpls interface lo0.5
user@R1# set LS5 protocols ospf area 0.0.0.0 interface ge-2/0/3.0
user@R1# set LS5 protocols ospf area 0.0.0.0 interface lt-2/0/10.5
user@R1# set LS5 protocols ospf area 0.0.0.0 interface lo0.5
user@R1# set LS6 protocols rsvp interface lt-2/0/10.6
user@R1# set LS6 protocols rsvp interface lt-2/0/10.11
user@R1# set LS6 protocols rsvp interface lo0.6
user@R1# set LS6 protocols mpls interface lt-2/0/10.6
user@R1# set LS6 protocols mpls interface lt-2/0/10.11
user@R1# set LS6 protocols mpls interface lo0.6
user@R1# set LS6 protocols ospf area 0.0.0.0 interface lt-2/0/10.6
user@R1# set LS6 protocols ospf area 0.0.0.0 interface lt-2/0/10.11
user@R1# set LS6 protocols ospf area 0.0.0.0 interface lo0.6
user@R1# set LS7 protocols rsvp interface lt-2/0/10.7
user@R1# set LS7 protocols rsvp interface lo0.7
user@R1# set LS7 protocols mpls interface lt-2/0/10.7
user@R1# set LS7 protocols mpls interface lo0.7
user@R1# set LS7 protocols ospf area 0.0.0.0 interface ge-2/0/1.0
user@R1# set LS7 protocols ospf area 0.0.0.0 interface lt-2/0/10.7
user@R1# set LS7 protocols ospf area 0.0.0.0 interface lo0.7
```

3. Enable traffic engineering for OSPF.

```
[edit logical-systems]
user@R1# set LS2 protocols ospf traffic-engineering
user@R1# set LS3 protocols ospf traffic-engineering
user@R1# set LS4 protocols ospf traffic-engineering
user@R1# set LS5 protocols ospf traffic-engineering
user@R1# set LS6 protocols ospf traffic-engineering
user@R1# set LS7 protocols ospf traffic-engineering
```

This causes the SPF algorithm to take into account the LSPs configured under MPLS.

4. Configure the router IDs.

```
[edit logical-systems]
user@R1# set LS2 routing-options router-id 100.20.20.20
user@R1# set LS3 routing-options router-id 100.30.30.30
user@R1# set LS4 routing-options router-id 100.40.40.40
user@R1# set LS5 routing-options router-id 100.50.50.50
user@R1# set LS6 routing-options router-id 100.60.60.60
user@R1# set LS7 routing-options router-id 100.70.70.70
```

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

```
[edit logical-systems]
user@R1# commit
```

**Results** From configuration mode, confirm your configuration by entering the **show logical-systems** command. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
user@R1# show logical-systems
LS1 {
  interfaces {
    ge-2/0/2 {
      unit 0 {
        description R1-to-CE1;
        family inet {
          address 10.0.244.10/30;
        }
      }
    }
  }
  lt-2/0/10 {
    unit 1 {
      description LS1-to-LS2;
      encapsulation ethernet;
      peer-unit 2;
      family inet {
        address 2.2.2.1/24;
      }
      family mpls;
    }
    unit 8 {
      description LS1-to-LS6;
      encapsulation ethernet;
```

```

        peer-unit 6;
        family inet {
            address 6.6.6.1/24;
        }
        family mpls;
    }
    unit 9 {
        description LS1-to-LS3;
        encapsulation ethernet;
        peer-unit 3;
        family inet {
            address 3.3.3.1/24;
        }
        family mpls;
    }
}
lo0 {
    unit 1 {
        family inet {
            address 100.10.10.10/32;
        }
    }
}
}
protocols {
    rsvp {
        interface lt-2/0/10.1;
        interface lt-2/0/10.8;
        interface lt-2/0/10.9;
        interface lo0.1;
    }
    mpls {
        traffic-engineering bgp-igp;
        label-switched-path LS1-to-LS5 {
            to 100.50.50.50;
            p2mp p2mpl;
        }
        label-switched-path LS1-to-LS7 {
            to 100.70.70.70;
            p2mp p2mpl;
        }
        label-switched-path LS1-to-LS4 {
            to 100.40.40.40;
            p2mp p2mpl;
        }
        interface lt-2/0/10.1;
        interface lt-2/0/10.8;
        interface lt-2/0/10.9;
        interface lo0.1;
    }
}
ospf {
    traffic-engineering;
    area 0.0.0.0 {
        interface ge-2/0/2.0;
        interface lt-2/0/10.1;
        interface lt-2/0/10.8;

```

```
        interface lt-2/0/10.9;
        interface lo0.1;
    }
}
routing-options {
    static {
        route 5.5.5.0/24 {
            p2mp-lsp-next-hop p2mp1;
        }
        route 7.7.7.0/24 {
            p2mp-lsp-next-hop p2mp1;
        }
        route 4.4.4.0/24 {
            p2mp-lsp-next-hop p2mp1;
        }
    }
    router-id 100.10.10.10;
}
LS2 {
    interfaces {
        lt-2/0/10 {
            unit 2 {
                description LS2-to-LS1;
                encapsulation ethernet;
                peer-unit 1;
                family inet {
                    address 2.2.2.2/24;
                }
                family mpls;
            }
            unit 10 {
                description LS2-to-LS5;
                encapsulation ethernet;
                peer-unit 5;
                family inet {
                    address 5.5.5.1/24;
                }
                family mpls;
            }
        }
        lo0 {
            unit 2 {
                family inet {
                    address 100.20.20.20/32;
                }
            }
        }
    }
    protocols {
        rsvp {
            interface lt-2/0/10.2;
            interface lt-2/0/10.10;
            interface lo0.2;
        }
    }
}
```

```

mpls {
    interface lt-2/0/10.2;
    interface lt-2/0/10.10;
    interface lo0.2;
}
ospf {
    traffic-engineering;
    area 0.0.0.0 {
        interface lt-2/0/10.2;
        interface lt-2/0/10.10;
        interface lo0.2;
    }
}
}
routing-options {
    router-id 100.20.20.20;
}
}
LS3 {
    interfaces {
        lt-2/0/10 {
            unit 3 {
                description LS3-to-LS1;
                encapsulation ethernet;
                peer-unit 9;
                family inet {
                    address 3.3.3.2/24;
                }
                family mpls;
            }
            unit 12 {
                description LS3-to-LS4;
                encapsulation ethernet;
                peer-unit 4;
                family inet {
                    address 4.4.4.1/24;
                }
                family mpls;
            }
        }
    }
    lo0 {
        unit 3 {
            family inet {
                address 100.30.30.30/32;
            }
        }
    }
}
}
protocols {
    rsvp {
        interface lt-2/0/10.3;
        interface lt-2/0/10.12;
        interface lo0.3;
    }
    mpls {
        interface lt-2/0/10.3;
    }
}

```

```
        interface lt-2/0/10.12;
        interface lo0.3;
    }
    ospf {
        traffic-engineering;
        area 0.0.0.0 {
            interface lt-2/0/10.3;
            interface lt-2/0/10.12;
            interface lo0.3;
        }
    }
}
routing-options {
    router-id 100.30.30.30;
}
}
LS4 {
    interfaces {
        ge-2/0/0 {
            unit 0 {
                description R1-to-CE4;
                family inet {
                    address 10.0.104.9/30;
                }
            }
        }
    }
    lt-2/0/10 {
        unit 4 {
            description LS4-to-LS3;
            encapsulation ethernet;
            peer-unit 12;
            family inet {
                address 4.4.4.2/24;
            }
            family mpls;
        }
    }
    lo0 {
        unit 4 {
            family inet {
                address 100.40.40.40/32;
            }
        }
    }
}
protocols {
    rsvp {
        interface lt-2/0/10.4;
        interface lo0.4;
    }
    mpls {
        interface lt-2/0/10.4;
        interface lo0.4;
    }
    ospf {
        traffic-engineering;
    }
}
```

```

        area 0.0.0.0 {
            interface ge-2/0/0.0;
            interface lt-2/0/10.4;
            interface lo0.4;
        }
    }
}
routing-options {
    router-id 100.40.40.40;
}
}
LS5 {
    interfaces {
        ge-2/0/3 {
            unit 0 {
                description LS1-to-CE1;
                family inet {
                    address 10.0.224.10/30;
                }
            }
        }
        lt-2/0/10 {
            unit 5 {
                description LS5-to-LS2;
                encapsulation ethernet;
                peer-unit 10;
                family inet {
                    address 5.5.5.2/24;
                }
                family mpls;
            }
        }
        lo0 {
            unit 5 {
                family inet {
                    address 100.50.50.50/32;
                }
            }
        }
    }
}
protocols {
    rsvp {
        interface lt-2/0/10.5;
        interface lo0.5;
    }
    mpls {
        interface lt-2/0/10.5;
        interface lo0.5;
    }
    ospf {
        traffic-engineering;
        area 0.0.0.0 {
            interface ge-2/0/3.0;
            interface lt-2/0/10.5;
            interface lo0.5;
        }
    }
}

```

```
    }  
  }  
  routing-options {  
    router-id 100.50.50.50;  
  }  
}  
LS6 {  
  interfaces {  
    lt-2/0/10 {  
      unit 6 {  
        description LS6-to-LS1;  
        encapsulation ethernet;  
        peer-unit 8;  
        family inet {  
          address 6.6.6.2/24;  
        }  
        family mpls;  
      }  
      unit 11 {  
        description LS6-to-LS7;  
        encapsulation ethernet;  
        peer-unit 7;  
        family inet {  
          address 7.7.7.1/24;  
        }  
        family mpls;  
      }  
    }  
  }  
  lo0 {  
    unit 6 {  
      family inet {  
        address 100.60.60.60/32;  
      }  
    }  
  }  
}  
protocols {  
  rsvp {  
    interface lt-2/0/10.6;  
    interface lt-2/0/10.11;  
    interface lo0.6;  
  }  
  mpls {  
    interface lt-2/0/10.6;  
    interface lt-2/0/10.11;  
    interface lo0.6;  
  }  
  ospf {  
    traffic-engineering;  
    area 0.0.0.0 {  
      interface lt-2/0/10.6;  
      interface lt-2/0/10.11;  
      interface lo0.6;  
    }  
  }  
}
```



```

routing-options {
  router-id 100.60.60.60;
}
}
LS7 {
  interfaces {
    ge-2/0/1 {
      unit 0 {
        description R1-to-CE3;
        family inet {
          address 10.0.134.10/30;
        }
      }
    }
    lt-2/0/10 {
      unit 7 {
        description LS7-to-LS6;
        encapsulation ethernet;
        peer-unit 11;
        family inet {
          address 7.7.7.2/24;
        }
        family mpls;
      }
    }
    lo0 {
      unit 7 {
        family inet {
          address 100.70.70.70/32;
        }
      }
    }
  }
  protocols {
    rsvp {
      interface lt-2/0/10.7;
      interface lo0.7;
    }
    mpls {
      interface lt-2/0/10.7;
      interface lo0.7;
    }
    ospf {
      traffic-engineering;
      area 0.0.0.0 {
        interface ge-2/0/1.0;
        interface lt-2/0/10.7;
        interface lo0.7;
      }
    }
  }
  routing-options {
    router-id 100.70.70.70;
  }
}

```

### Configuring Device CE1

---

#### Step-by-Step Procedure

To configure Device CE1:

1. Configure an interface to Logical System LS1.

```
[edit]
user@CE1# edit interfaces
[edit interfaces]
user@CE1# set ge-1/3/2 unit 0 family inet address 10.0.244.9/30
user@CE1# set ge-1/3/2 unit 0 description CE1-to-LS1
user@CE1# exit
```

2. Configure static routes from Device CE1 to the three other customer networks, with Logical System LS1 as the next hop.

```
[edit]
user@CE1# edit routing-options
[edit routing-options]
set static route 10.0.104.8/30 next-hop 10.0.244.10
set static route 10.0.134.8/30 next-hop 10.0.244.10
set static route 10.0.224.8/30 next-hop 10.0.244.10
user@CE1# exit
```

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

```
[edit]
user@CE1# commit
```

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

```
user@CE1# show interfaces
interfaces {
  ge-1/3/2 {
    unit 0 {
      family inet {
        address 10.0.244.9/30;
        description CE1-to-LS1;
      }
    }
  }
}

user@CE1# show routing-options
routing-options {
  static {
    route 10.0.104.8/30 next-hop 10.0.244.10;
    route 10.0.134.8/30 next-hop 10.0.244.10;
    route 10.0.224.8/30 next-hop 10.0.244.10;
  }
}
```

---

## Configuring Device CE2

---

### Step-by-Step Procedure

To configure Device CE2:

1. Configure an interface to Logical System LS5.  

```
[edit]
user@CE2# edit interfaces
[edit interfaces]
user@CE2# set ge-1/3/3 unit 0 family inet address 10.0.224.9/30
user@CE2# set ge-1/3/3 unit 0 description CE2-to-LS5
user@CE2# exit
```
2. Configure a static route from Device CE2 to CE1, with Logical System LS5 as the next hop.  

```
[edit]
user@CE2# edit routing-options
[edit routing-options]
user@CE2# set static route 10.0.244.8/30 next-hop 10.0.224.10
user@CE2# exit
```
3. If you are done configuring the device, commit the configuration.  

```
[edit]
user@CE2# commit
```

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

```
user@CE2# show interfaces
interfaces {
  ge-1/3/3 {
    unit 0 {
      family inet {
        address 10.0.224.9/30;
        description CE2-to-LS5;
      }
    }
  }
}

user@CE2# show routing-options
routing-options {
  static {
    route 10.0.244.8/30 next-hop 10.0.224.10;
  }
}
```

### Configuring Device CE3

---

#### Step-by-Step Procedure

To configure Device CE3:

1. Configure an interface to Logical System LS7.  

```
[edit]
user@CE3# edit interfaces
[edit interfaces]
user@CE3# set ge-2/0/1 unit 0 family inet address 10.0.134.9/30
user@CE3# set ge-2/0/1 unit 0 description CE3-to-LS7
user@CE3# exit
```
2. Configure a static route from Device CE3 to CE1, with Logical System LS7 as the next hop.  

```
[edit]
user@CE3# edit routing-options
[edit routing-options]
user@CE3# set static route 10.0.244.8/30 next-hop 10.0.134.10
user@CE3# exit
```
3. If you are done configuring the device, commit the configuration.  

```
[edit]
user@CE3# commit
```

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

```
user@CE3# show interfaces
interfaces {
  ge-2/0/1 {
    unit 0 {
      family inet {
        address 10.0.134.9/30;
        description CE3-to-LS7;
      }
    }
  }
}

user@CE3# show routing-options
routing-options {
  static {
    route 10.0.244.8/30 next-hop 10.0.134.10;
  }
}
```

---

## Configuring Device CE4

---

### Step-by-Step Procedure

To configure Device CE4:

1. Configure an interface to Logical System LS4.

```
[edit]
user@CE4# edit interfaces
[edit interfaces]
user@CE4# set ge-3/1/3 unit 0 family inet address 10.0.104.10/30
user@CE4# set ge-3/1/3 unit 0 description CE4-to-LS4
```

2. Configure a static route from Device CE4 to CE1, with Logical System LS4 as the next hop.

```
[edit]
user@CE4# edit routing-options
[edit routing-options]
user@CE4# set static route 10.0.244.8/30 next-hop 10.0.104.9
user@CE4# exit
```

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

```
[edit]
user@CE4# commit
```

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

```
user@CE4# show interfaces
interfaces {
  ge-3/1/3 {
    unit 0 {
      family inet {
        address 10.0.104.10/30;
        description CE4-to-LS4;
      }
    }
  }
}

user@CE4# show routing-options
routing-options {
  static {
    route 10.0.244.8/30 next-hop 10.0.104.9;
  }
}
```

## Verification

Confirm that the configuration is working properly.

- [Verifying Connectivity on page 26](#)
- [Verifying the State of the Point-to-Multipoint LSP on page 26](#)
- [Checking the Forwarding Table on page 27](#)

---

### Verifying Connectivity

**Purpose** Make sure that the devices can ping each other.

**Action** Run the **ping** command from CE1 to the interface on CE2 connecting to LS5.

```
user@CE1> ping 10.0.224.9
PING 10.0.224.9 (10.0.224.9): 56 data bytes
64 bytes from 10.0.224.9: icmp_seq=0 ttl=61 time=1.387 ms
64 bytes from 10.0.224.9: icmp_seq=1 ttl=61 time=1.394 ms
64 bytes from 10.0.224.9: icmp_seq=2 ttl=61 time=1.506 ms
^C
--- 10.0.224.9 ping statistics ---
3 packets transmitted, 3 packets received, 0% packet loss
round-trip min/avg/max/stddev = 1.387/1.429/1.506/0.055 ms
```

Run the **ping** command from CE1 to the interface on CE3 connecting to LS7.

```
user@CE1> ping 10.0.134.9
PING 10.0.134.9 (10.0.134.9): 56 data bytes
64 bytes from 10.0.134.9: icmp_seq=0 ttl=61 time=1.068 ms
64 bytes from 10.0.134.9: icmp_seq=1 ttl=61 time=1.062 ms
64 bytes from 10.0.134.9: icmp_seq=2 ttl=61 time=1.053 ms
^C
--- 10.0.134.9 ping statistics ---
3 packets transmitted, 3 packets received, 0% packet loss
round-trip min/avg/max/stddev = 1.053/1.061/1.068/0.006 ms
```

Run the **ping** command from CE1 to the interface on CE4 connecting to LS4.

```
user@CE1> ping 10.0.104.10
PING 10.0.104.10 (10.0.104.10): 56 data bytes
64 bytes from 10.0.104.10: icmp_seq=0 ttl=61 time=1.079 ms
64 bytes from 10.0.104.10: icmp_seq=1 ttl=61 time=1.048 ms
64 bytes from 10.0.104.10: icmp_seq=2 ttl=61 time=1.070 ms
^C
--- 10.0.104.10 ping statistics ---
3 packets transmitted, 3 packets received, 0% packet loss
round-trip min/avg/max/stddev = 1.048/1.066/1.079/0.013 ms
```

---

### Verifying the State of the Point-to-Multipoint LSP

**Purpose** Make sure that the ingress, transit, and egress LSRs are in the Up state.



**NOTE:** For this example, the **show rsvp session** command displays the same output as the **show mpls lsp p2mp** command.

---

**Action** Run the `show mpls lsp p2mp` command on all of the LSRs. Only the ingress LSR is shown here.

```
user@R1> set cli logical-system LS1
Logical system: LS1
```

```
user@R1:LS1> show mpls lsp p2mp
Ingress LSP: 1 sessions
P2MP name: p2mp1, P2MP branch count: 3
To          From          State Rt P    ActivePath    LSPname
100.40.40.40 100.10.10.10 Up    0 *           LS1-LS4
100.70.70.70 100.10.10.10 Up    0 *           LS1-LS7
100.50.50.50 100.10.10.10 Up    0 *           LS1-LS5
Total 3 displayed, Up 3, Down 0
...
```

---

### Checking the Forwarding Table

**Purpose** Make sure that the routes are set up as expected by running the `show route forwarding-table` command. Only the routes to the remote customer networks are shown here.

**Action** `user@R1:LS1> show route forwarding-table`

```
Routing table: default.inet
Internet:
Destination          Type RtRef Next hop          Type Index NhRef Netif
...
10.0.104.8/30         user    0 3.3.3.2          ucst 1006    6 1t-2/0/10.9
10.0.134.8/30         user    0 6.6.6.2          ucst 1010    6 1t-2/0/10.8
10.0.224.8/30         user    0 2.2.2.2          ucst 1008    6 1t-2/0/10.1
...
```

**Related Documentation**

- [Point-to-Multipoint LSPs Overview on page 1](#) in the Junos OS MPLS Applications Configuration Guide

