Juniper Networks | Engineering Simplicity

DAY ONE: MPLS UP AND RUNNING ON JUNOS

By Nupur Kanoi

DAY ONE: MPLS UP AND RUNNING ON JUNOS

MPLS is an integral part of most Service Provider networks, so any modern book on MPLS needs to have an in-depth understanding of MPLS and traffic engineering (TE). D*ay One: MPLS Up and Running on Junos* covers MPLS in a fast, up and running format, reviewing MPLS basics while moving the reader towards the traffic engineering requirements and configurations that can build a fault tolerant network. Nupur Kanoi has the necessary experience to write this MPLS shorthand and she presents a hands-on lab guide that can help you build a MPLS topology and with an MPLS VPN running end-to-end. Clear definitions and sample configurations make it quick and easy for network administrators to follow along.

"Day One: MPLS Up and Running *on Junos has been a vital learning tool for me, allowing for personal growth and understanding of MPLS and traffic engineering. I highly recommended it for engineers interested in getting started with MPLS and learning traffic engineering techniques."*

Kody Vicknair, Network Engineer, Reserve Telecommunications, (JNCIA-Junos, CCNA, BCNE)

"Have fun learning and labbing MPLS with this awesome guide, and develop the skills you need to deploy and support MPLS services. Everything is covered from basic to advanced topics, and it's all nicely wrapped up with a complete end-to-end topology."

Matt Dinham, Director of Network Operations, Optimity

IT'S DAY ONE AND YOU HAVE A JOB TO DO, SO LEARN HOW TO:

- Understand the basic concepts of MPLS.
- Understand the label distribution protocols, LDP and RSVP, and their differences.
- Design and implement a fault tolerant MPLS network with either LDP or RSVP, or LDP and RSVP.
- Effectively scale that network with the help of MPLS TE.
- Efficiently troubleshoot the core MPLS network.

ISBN 978-1-941441-82-4
52000
9 781941 441824

Juniper Networks Books are focused on network reliability and efficiency. Peruse the complete library at www.juniper.net/books.

JUNIPER NETWORKS

Day One: MPLS Up and Running with Junos®

by Nupur Kanoi

Chapter 1: MPLS Overview 7

Chapter 2: LSP Signaling Protocols 15

Chapter 3: Traffic Engineering 28

Chapter 4: MPLS VPN 45

Chapter 5: Deploying VPN End to End 77

Published by Juniper Networks Books
Author: Nupur Kanoi
Technical Reviewers: Matt Dinham, Martin Brown, and Steve Puluka
Editor in Chief: Patrick Ames
Copyeditor: Nancy Koerbel
Illustrator: Karen Joice

ISBN: 978-1-941441-82-4 (print)
Printed in the USA by Vervante Corporation.

ISBN: 978-1-941441-81-7 (ebook)

Version History: v1, December 2018
2 3 4 5 6 7 8 9 10

http://www.juniper.net/dayone

About the Author
Nupur Kanoi is a Senior Backbone engineer for a global service provider, where she has gained experience in service provider backbone architecture and design. She also holds JNCIE-ENT (#520), JNCIE-SP (#2824), JNCIP-DC, and JNCDS-DC certifications. Nupur can be reached on LinkedIn (linkedin.com/in/nupur-kanoi-520) and on Twitter (@nupur_kanoi).

Author's Acknowledgments
I would like to thank my family and friends, especially my Mom, for being a constant source of inspiration and keeping up with me even when I'm exhausted by my schedule. Also, thanks to my colleagues for their extended support in my learning whenever needed.

Feedback? Comments? Error reports? Email them to dayone@juniper.net.

Welcome to Day One

This book is part of the *Day One* library, produced and published by Juniper Networks Books.

Day One books cover the Junos OS and Juniper Networks networking essentials with straightforward explanations, step-by-step instructions, and practical examples that are easy to follow. You can obtain the books from various sources:

- Download a free PDF edition at http://www.juniper.net/dayone.
- Many of the library's books are available on the Juniper app: Junos Genius.
- Get the ebook edition for iPhones and iPads from the iBooks Store. Search for *Juniper Networks Books* or the title of this book.
- Get the ebook edition for any device that runs the Kindle app (Android, Kindle, iPad, PC, or Mac) by opening your device's Kindle app and going to the Amazon Kindle Store. Search for *Juniper Networks Books* or the title of this book.
- Purchase the paper edition at Vervante Corporation (www.vervante.com) for between $15-$40, depending on page length.
- Note that most mobile devices can also view PDF files.

Key MPLS Resources

There are literally hundreds of MPLS books out there, but not many of them deal with the Junos OS. Here are two key MPLS resources:

The Juniper TechLibrary is stuffed with MPLS information for Junos and for individual hardware platforms. This book is not a substitute for that body of work, so you should take the time to review the documentation: https://www.juniper.net/documentation/en_US/junos/information-products/pathway-pages/config-guide-mpls-applications/config-guide-mpls-applications.html.

MPLS in the SDN Era, by Antonio Sánchez-Monge and Krzysztof Grzegorz Szarkowicz, O'Reilly Media, 2015, shows you how to deploy a large portfolio of multi-vendor MPLS services down to the Junos configuration level. It's available from the publisher and from various bookstores: http://shop.oreilly.com/product/0636920033905.do.

.

What You Need to Know Before Reading This Book

The author has made a few assumptions about the general knowledge level of the reader:

- You are a network engineer who has some experience with managing computer networks.
- You understand typical network management tasks and are willing to see possible benefits to automating those tasks.
- You have a strong understanding of IP networking and prior experience with OSPF, ISIS, and BGP, moreover, you are well versed in the Junos CLI.
- You can build a sample topology using vMX instances.

NOTE There are several books in *Day One* library covering IGP/BGP and the Junos CLI: http://www.juniper.net/dayone.

What You Will Learn by Reading This Book

This book will help you to:

- Understand the basic concepts of MPLS.
- Understand the label distribution protocols, LDP and RSVP, and their differences.
- Design and implement a fault tolerant MPLS network with either LDP or RSVP, or LDP and RSVP.
- Effectively scale that network with the help of MPLS TE.
- Efficiently troubleshoot the core MPLS network.

This Book's Lab

The lab for this book was set up with the help of vMX instances running Junos OS 14.1R1. You can download the vMX 60-day trial version from Juniper Networks (https://www.juniper.net/us/en/dm/free-vmx-trial/). Most of the topics will also work on vSRX in packet-based forwarding mode.

To get help with the vMX installation you can refer to *Day One: vMX Up and Running*: https://www.juniper.net/us/en/training/jnbooks/day-one/automation-series/vmx-up-running/.

Chapter 1

MPLS Overview

This book is written for ease of comprehension and speed of use in the lab. It introduces you to the basics of MPLS that are easy to understand and consume, and then builds on these basics, allowing you to take in subsequent concepts that are slightly more complex.

This chapter begins by covering MPLS basics like *shim header* and *label handling* on a hop-by-hop basis like the *push/pop/swap* operation. It also includes details about the Penultimate-hop popping (PHP) operation with an overview of label-switched paths (LSPs).

Why MPLS?

MPLS is designed to work within a multi-protocol environment and may reduce delay by avoiding the routing table lookup. It is flexible in many ways—such as having separate control and data planes— and it also allows traffic to have an explicit path set through the network by taking advantage of MPLS traffic engineering (TE) for a next hop without making any modifications to an IGP path. You can use MPLS TE to avoid congestion and distribute traffic to non-standard paths making the optimal use of all links. Moreover, MPLS TE can use class of service to prioritize specific types of traffic (say VoIP) over others.

MPLS is more scalable in comparison to IGP routing – you can simply add a new site to the VRF without the need to create complex mesh tunnels.

Understanding the MPLS Header

Before the traffic is sent out of a PE participating in the MPLS backbone towards the next hop, it inserts a 32-bit MPLS header between the Layer 2 header and the Layer 3 IP header – also known as a *shim header*. This header and label value within the header lets the immediate next hop router know what action needs to be taken: for example, *push*, *swap*, or *pop* the label. Because the MPLS shim header sits right between the Layer 3 and Layer 2 header, it's also sometimes known as a *Layer 2.5* protocol, although "2.5" is not defined in any actual standards.

Let's now look at that MPLS shim header in Figure 1.1.

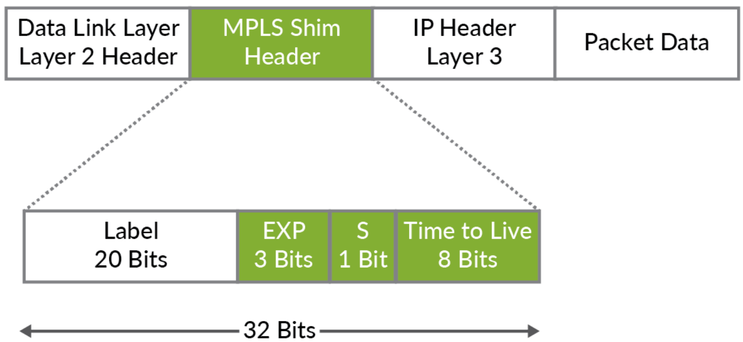

Figure 1.1 *MPLS Shim Header*

Here are the components of the 32-bit header:

Label - This value identifies a LSP and it changes at every hop. Being a 20-bit value, it ranges from 0 to 1,048,575—of which 0-16 are reserved—with the remainder being used for normal packet forwarding. The reserved values have special meanings as listed in Table 1.1.

Table 1.1 *Header Reserved Values*

Label 0	IPv4 Explicit Null
Label 1	Router Alert Label
Label 2	IPv6 Explicit Null
Label 3	Implicit Null
4-15 Label	For Future use

EXP - These three bits are called *experimental bits* and can be used to implement class of service to prioritize the traffic in the MPLS network.

Bottom of stack bit - This 1 bit is used to indicate whether more than one label is stacked on the packet. There can be multiple labels associated with a packet, which in general is pushed by an ingress router. The value 1 in this field tells the egress router that this is the last label associated with the packet, and after popping the label out, what remains will be an unlabeled packet.

TTL - This 8-bit field limits the number of hops a labeled packet can travel. The default functionality is to copy the TTL value from IP header to the MPLS header at the ingress router, and at egress, the MPLS header TTL value is copied back to the IP header by the egress router.

MPLS Terms and Label Handling

Table 1.2 MPLS Terminology

Acronym	Terminology
MPLS	Multiprotocol Label Switching
LSP	Label-switched path
LSR	Label-switching router
LER	Label edge router
PE router	Provider edge router
P router	Provider core router
PHP	Penultimate-hop popping (described later in this chapter)
Push	A MPLS label is added to the top of the stack
Swap	A MPLS label is swapped by an LSR
Pop	The topmost MPLS label is removed from the stack

The key distinctions between LSP, LSR, and LER (see Figure 1.2) are:

- *Label-switched path* (LSP): An LSP is a unidirectional path built solely on labels from an ingress router to an egress router. Hence, for a two-way path (bidirectional path), two LSPs are built – one LSP in each direction.
- *Label-switching router* (LSR): An LSR is a MPLS router that takes forwarding decisions based solely on labels. It swaps or pops (in the case of PHP) the label based on the entries in the *mpls.0* table.
- *Label edge router* (LER): The LER is the router that sits at the edge of the network (say the ingress or egress PE router). The LER/PE will either push the label (ingress router) or pop the label (egress router). The push label operation at LER is performed using the entries in the *inet.3* table and then the packet is forwarded in the MPLS domain.

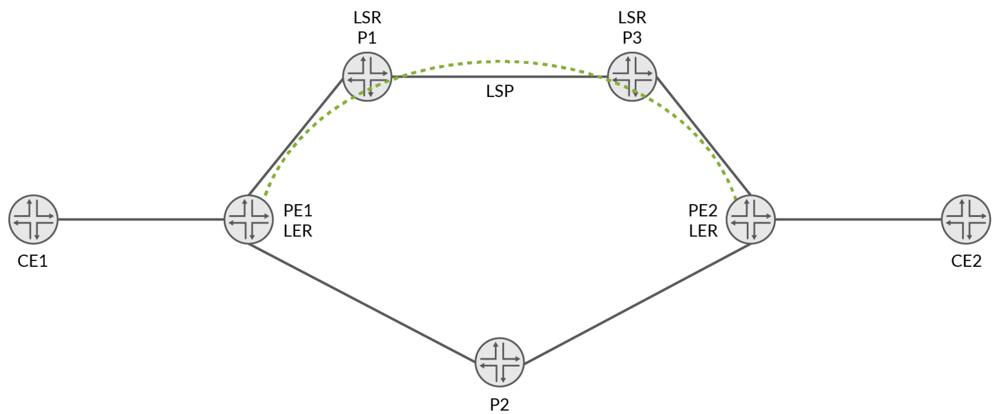

Figure 1.2 LSP Illustration

Figure 1.2 illustrates a unidirectional LSP from PE1 to PE2. PE1 and PE2 are the LERs, precisely, PE1 is the ingress LER and PE2 is the egress LER. The P1 and P3 are the LSR routers performing the label swap operation.

Now let's review how labels are learned. Figure 1.3 illustrates label learning at each hop in a LSP. Label 299776 is installed in the *inet.3* table of PE1, which PE1 learned from P1, stating that if the packet needs to be sent to PE2, it then is encapsulated with label 299776 while sending it towards P1. Similarly, P1 has learned from P3 that the packet destined for PE2 should be encapsulated with label 299792, while sending towards P3. Further, P3 has learned from PE2 that the packet destined for PE2 should have label 100651.

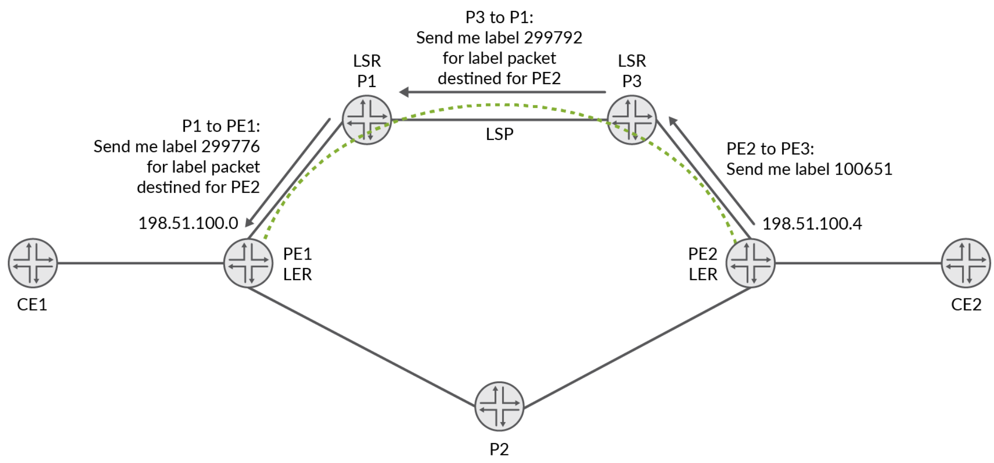

Figure 1.3 Label Distribution

So, the packet enters the LSP at the ingress from PE1. PE1 pushes label 299776, consulting its *inet.3* table, and sends it out towards P1. P1 then swaps the label from 299776 to 299792 referencing the *mpls.0* table and sends it out towards P3. Further, P3 again swaps the label from 299792 to 100651 and sends the packets to PE2. Once the packet reaches PE2, and PE2 sees the label 100651, it knows that the label packet is destined for itself, therefore it pops the label and forwards the packets towards CE2 based on IP routing.

In MPLS-speak, here PE1 and PE2 were LERs, whereas P1 and P3 were LSRs. The following output depicts how to check labels on the ingress LER and LSR.

On the ingress LER (PE1), you can see that label 299776 is pushed:

```
lab@PE1> show route table inet.3 198.51.100.4/32 detail

inet.3: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
198.51.100.4/32 (1 entry, 0 announced)
        State: <FlashAll>
        *RSVP   Preference: 7/1
                Next hop type: Router
                Address: 0x93505b0
                Next-hop reference count: 4
                Next hop: 192.0.2.1 via ge-0/0/2.0 weight 0x1, selected
                Label-switched-path PE1_TO_PE2
                Label operation: Push 299776
                Label TTL action: prop-ttl
                State: <Active Int>
                Age: 14:37      Metric: 3
                Task: RSVP
                AS path: I
```

And on the LSR (P1), you can see the label's swap operation:

```
lab@P1> show route table mpls.0 label 299776

mpls.0: 5 destinations, 5 routes (5 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

299776             *[RSVP/7/1] 00:08:41, metric 1
                    > to 192.0.2.5 via ge-0/0/3.0, label-switched-path PE1_TO_PE2

lab@P1> show route table mpls.0 label 299776 detail

mpls.0: 5 destinations, 5 routes (5 active, 0 holddown, 0 hidden)
299776 (1 entry, 1 announced)
        *RSVP   Preference: 7/1
                Next hop type: Router, Next hop index: 572
                Address: 0x9350688
                Next-hop reference count: 3
                Next hop: 192.0.2.5 via ge-0/0/3.0 weight 0x1, selected
                Label-switched-path PE1_TO_PE2
                Label operation: Swap 299792
                State: <Active Int AckRequest>
                Age: 8:43       Metric: 1
                Task: RSVP
                Announcement bits (1): 0-KRT
                AS path: I
```

Penultimate-hop Popping (PHP)

PHP is taken care of on the penultimate (the second-to-last) router, or the LSR before the LER. It is a process where the outermost label is popped before reaching its destination in the MPLS domain, which means the LSR pops the outer label before the packet is passed to the LER.

For example, in an MPLS L3VPN, this process helps the LER to avoid two-label lookups: the outer label for the LSP and the inner label would be the VRF label. This in turn reduces load on the LER.

The PHP happens with the help of the LER sending of the implicit-null label, which is the reserved label 3, to the upstream router. The implicit label (3) is an indication to the upstream LSR to pop the outer label before passing the label packet to the LER. This distributes the outer label-popping operation across multiple LSRs, therefore it effectively reduces CPU load on an LER.

In the case of MPLS L3VPNs, after the PHP operation, the LER is left with only the inner label on the packet. This virtual routing and forwarding (VRF) label signifies which VRF instance is to be considered for IP routing lookup. The whole process is illustrated in Figure 1.4.

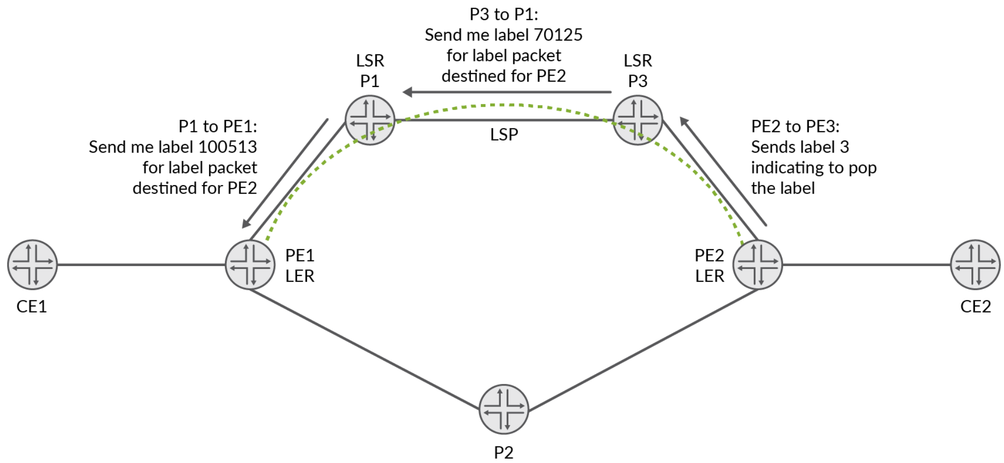

Figure 1.4 *PHP*

As Figure 1.4 illustrates, the egress LER PE2 sends label 3 to inform LSR P3 to pop the outer label before sending the packet to PE2.

And here is the output from P3's *mpls.0* table:

```
lab@P3> show route table mpls.0 label 299792

mpls.0: 6 destinations, 6 routes (6 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

299792             *[RSVP/7/1] 00:11:56, metric 1
                    > to 192.0.2.7 via ge-0/0/3.0, label-switched-path PE1_TO_PE2
299792(S=0)        *[RSVP/7/1] 00:11:56, metric 1
                    > to 192.0.2.7 via ge-0/0/3.0, label-switched-path PE1_TO_PE2

lab@P3> show route table mpls.0 label 299792 detail

mpls.0: 6 destinations, 6 routes (6 active, 0 holddown, 0 hidden)
299792 (1 entry, 1 announced)
        *RSVP   Preference: 7/1
                Next hop type: Router, Next hop index: 572
                Address: 0x9350640
                Next-hop reference count: 3
                Next hop: 192.0.2.7 via ge-0/0/3.0 weight 0x1, selected
                Label-switched-path PE1_TO_PE2
                Label operation: Pop
                State: <Active Int AckRequest>
                Age: 11:59      Metric: 1
                Task: RSVP
                Announcement bits (1): 0-KRT
                AS path: I
```

```
299792(S=0) (1 entry, 1 announced)
        *RSVP   Preference: 7/1
                Next hop type: Router, Next hop index: 573
                Address: 0x9350688
                Next-hop reference count: 2
                Next hop: 192.0.2.7 via ge-0/0/3.0 weight 0x1, selected
                Label-switched-path PE1_TO_PE2
                Label operation: Pop
                State: <Active Int AckRequest>
                Age: 11:59      Metric: 1
                Task: RSVP
                Announcement bits (1): 0-KRT
                AS path: I
```

The *mpls.0* table tells P3 to pop the outer label while indicating that it's not the last label on the stack by S=0.

PHP is the default feature of Junos OS routers and can be altered by enabling `ultimate-hop-popping`. This is achieved on the ingress LER:

```
set protocols mpls label-switched-path <label-switched-path-name> ultimate-hop-popping
```

The knob pops both the outer label and the VRF label on the egress LER.

MORE? For more details on this knob and enabling ultimate hop popping, visit the Juniper TechLibrary: https://www.juniper.net/documentation/en_US/junos/topics/task/configuration/mpls-ultimate-hop-popping-enabling.html.

Summary

This chapter introduced you to the basics of MPLS, including the label handling and various MPLS nomenclature. It should help you to understand the MPLS label packet flow and also the default feature of PHP that reduces the load on the egress router.

In Chapter 2 you will get an MPLS LSP running with the help of some different signaling protocols.

Chapter 2

LSP Signaling Protocols

MPLS signals unidirectional LSPs with the help of distributing labels activated by a signaling protocol. You can have either LDP or RSVP as a signaling protocol to build those LSPs dynamically. This chapter introduces both the signaling protocols by using the reference topology shown in Figure 2.1. Use the topology diagram to get your lab set up.

MORE? Configurations and commands used here were on a vMX running Junos OS 14.1R1.

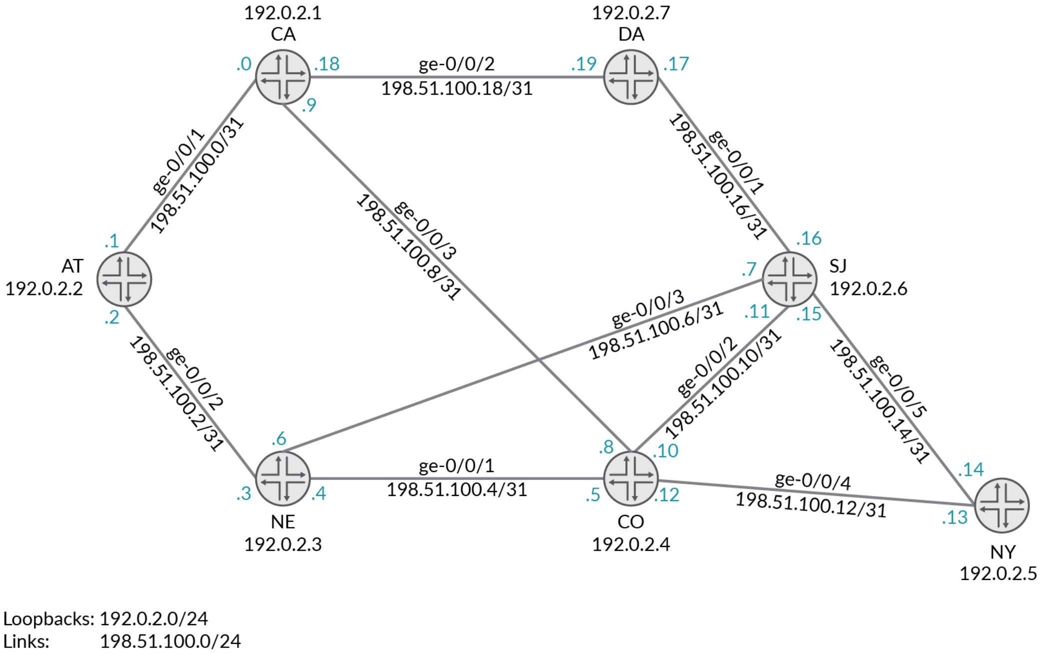

Figure 2.1 Reference Topology

LDP

The LDP runs on top of an IGP. It follows the exact path chosen by the IGP to build an LSP from the ingress LER to the egress LER. The IGP can be OSPF or IS-IS, and it must be stable. The IGP metric/cost decides the path that is followed by the LSP, meaning that to achieve a *specific* path for an LSP, the IGP metric needs to be modified.

LDP is a simple protocol and ideal for networks, which do not require much traffic engineering. It associates a group of destination prefixes to a single LSP. The set of destination prefixes are the *forwarding equivalence class* (FEC), or a set of prefixes to be treated in same way. For example, a set of prefixes for a single VRF is marked as a FEC. A single FEC will share the same LSP and exit node. The LDP configuration is applied here:

```
lab@CA> show configuration protocols ldp
interface ge-0/0/1.0;
interface ge-0/0/2.0;
interface ge-0/0/3.0;
interface lo0.0;

lab@CA> show configuration interfaces ge-0/0/1.0
family inet {
    address 198.51.100.0/31;
}
family mpls;
```

NOTE You do not need to configure the interfaces in the `protocols mpls` hierarchy for LDP to function as expected.

And here, you can see the LDP session information:

```
lab@CA> show ldp session
  Address           State        Connection     Hold time
192.0.2.2           Operational  Open             27
192.0.2.4           Operational  Open             27
192.0.2.7           Operational  Open             27
```

Next, the *input label database* shows received labels and signifies the labels to be pushed when sending traffic to the prefix via that LDP neighbor. The *output label database* shows the labels the router expects to have attached for incoming MPLS-labeled traffic when transiting via, or destined, to the router. For example:

```
lab@CA> show ldp database
Input label database, 192.0.2.1:0--192.0.2.2:0
  Label     Prefix
 299920     192.0.2.1/32
      3     192.0.2.2/32
 299792     192.0.2.3/32
 299872     192.0.2.4/32
 299904     192.0.2.5/32
 299888     192.0.2.6/32
 299936     192.0.2.7/32

Output label database, 192.0.2.1:0--192.0.2.2:0
  Label     Prefix
      3     192.0.2.1/32
 299952     192.0.2.2/32
 299936     192.0.2.3/32
 299888     192.0.2.4/32
 299904     192.0.2.5/32
 299920     192.0.2.6/32
 299872     192.0.2.7/32

Input label database, 192.0.2.1:0--192.0.2.4:0
  Label     Prefix
 299872     192.0.2.1/32
 299904     192.0.2.2/32
 299776     192.0.2.3/32
      3     192.0.2.4/32
 299856     192.0.2.5/32
 299824     192.0.2.6/32
 299888     192.0.2.7/32

Output label database, 192.0.2.1:0--192.0.2.4:0
  Label     Prefix
      3     192.0.2.1/32
 299952     192.0.2.2/32
 299936     192.0.2.3/32
 299888     192.0.2.4/32
 299904     192.0.2.5/32
 299920     192.0.2.6/32
 299872     192.0.2.7/32
```

```
Input label database, 192.0.2.1:0--192.0.2.7:0
  Label     Prefix
 299872     192.0.2.1/32
 299952     192.0.2.2/32
 299936     192.0.2.3/32
 299888     192.0.2.4/32
 299904     192.0.2.5/32
 299920     192.0.2.6/32
      3     192.0.2.7/32

Output label database, 192.0.2.1:0--192.0.2.7:0
  Label     Prefix
      3     192.0.2.1/32
 299952     192.0.2.2/32
 299936     192.0.2.3/32
 299888     192.0.2.4/32
 299904     192.0.2.5/32
 299920     192.0.2.6/32
 299872     192.0.2.7/32

lab@CA> show route 192.0.2.6/32 table inet.3

inet.3: 6 destinations, 6 routes (6 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

192.0.2.6/32       *[LDP/9] 00:40:11, metric 1
                    > to 198.51.100.19 via ge-0/0/2.0, Push 299920
                      to 198.51.100.8 via ge-0/0/3.0, Push 299824

lab@CA> show ldp path 192.0.2.6/32
Output Session (label)          Input Session (label)
192.0.2.4:0(299920)             192.0.2.4:0(299824)
192.0.2.2:0(299920)             192.0.2.7:0(299920)
192.0.2.7:0(299920)
  Attached route:  192.0.2.6/32, Ingress route
```

For those routes having a next hop of 192.0.2.6/32, the output shows two MPLS LDP paths. The route via the DA router is currently active, hence the CA router needs to push MPLS label 299920 (the highlighted label seen in the input label database coming in from DA and CO).

LDP Synchronization

As discussed, LDP relies on the IGP's best route for forwarding MPLS-labeled packets. Under certain scenarios, IGP path selection might not be in synchronization with the LDP label exchange path, leading to MPLS packet loss. For example:

- When an LDP neighborship is not operational due to some error, but the router continues to forward traffic through the IGP link associated with the non-operational LDP neighbor because the IGP neighborship is functional.
- When a new IGP neighborship is established and the shortest path first (SPF) computes the associated link as the best route to a certain destination, the router begins forwarding MPLS-labeled traffic through that link before it completes the LDP label exchange.

LDP synchronization reduces the MPLS packet loss while the LDP and IGP sync is incomplete – during this period the IGP will advertise the concerned link with a maximum cost metric. So, the IGP does not add the link as a transit path for any destination, which in turn allows the router to avoid using that link for MPLS-labeled packets. Once the LDP session is operational and the LDP label exchange is completed, the IGP starts advertising the same link with its expected cost metric. This mechanism ensures that the IGP is in sync with LDP to avoid packet loss.

LDP synchronization is only supported on point-to-point interfaces. To understand why LDP synchronization is not supported on LAN interfaces, you need to get into the basics of the multi-access segment (LAN). Multiple neighbors are expected to be connected on a single LAN interface. Increasing the metric to that LAN interface impacts traffic to other peers connected on the same LAN interface, even when LDP and IGP are in sync with those neighbors.

Here's how to configure LDP synchronization:

```
lab@CA> show configuration protocols ospf
traffic-engineering;
area 0.0.0.0 {
    interface ge-0/0/1.0 {
        interface-type p2p;
        ldp-synchronization;
    }
    interface ge-0/0/2.0 {
        interface-type p2p;
        ldp-synchronization;
    }
    interface ge-0/0/3.0 {
        interface-type p2p;
        ldp-synchronization;
    }
    interface lo0.0 {
        passive;
    }
}
```

And let's verify:

```
lab@CA> show ospf interface extensive
Interface           State   Area            DR ID           BDR ID          Nbrs
ge-0/0/1.0          PtToPt  0.0.0.0         0.0.0.0         0.0.0.0            1
  Type: P2P, Address: 198.51.100.0, Mask: 255.255.255.254, MTU: 1500, Cost: 1
  Adj count: 1
  Hello: 10, Dead: 40, ReXmit: 5, Not Stub
  Auth type: None
  Protection type: None
  Topology default (ID 0) -> Cost: 1
  LDP sync state: in sync, for: 00:00:21, reason: LDP session up
           config holdtime: infinity
ge-0/0/2.0          PtToPt  0.0.0.0         0.0.0.0         0.0.0.0            1
  Type: P2P, Address: 198.51.100.18, Mask: 255.255.255.254, MTU: 1500, Cost: 1
  Adj count: 1
  Hello: 10, Dead: 40, ReXmit: 5, Not Stub
  Auth type: None
  Protection type: None
  Topology default (ID 0) -> Cost: 1
```

```
  LDP sync state: in sync, for: 00:01:13, reason: LDP session up
           config holdtime: infinity
ge-0/0/3.0          PtToPt  0.0.0.0         0.0.0.0         0.0.0.0            1
  Type: P2P, Address: 198.51.100.9, Mask: 255.255.255.254, MTU: 1500, Cost: 1
  Adj count: 1
  Hello: 10, Dead: 40, ReXmit: 5, Not Stub
  Auth type: None
  Protection type: None
  Topology default (ID 0) -> Cost: 1
  LDP sync state: in sync, for: 00:00:46, reason: LDP session up
           config holdtime: infinity
lo0.0               DRother 0.0.0.0         0.0.0.0         0.0.0.0            0
  Type: LAN, Address: 192.0.2.1, Mask: 255.255.255.255, MTU: 65535, Cost: 0
  Adj count: 0, Passive
  Hello: 10, Dead: 40, ReXmit: 5, Not Stub
  Auth type: None
  Protection type: None
  Topology default (ID 0) -> Passive, Cost: 0
```

RSVP

RSVP is the acronym for Resource Reservation Protocol. It was predominantly designed to reserve and to maintain the distributed resource reservation across multiple nodes. Over time, extensions were added so that RSVP could be a signaling protocol for MPLS. RSVP carries opaque objects for building MPLS LSPs that allow for MPLS traffic engineering. The opaque objects in the RSVP extensions do not create any compatibility issue with the existing functions of RSVP.

RSVP is used for bandwidth allocation and creates MPLS LSPs with either a manually defined *Explicit Route Object* (ERO), or an ERO computed with the help of the Constrained Shortest Path First (CSPF) algorithm in the MPLS domain. Moreover, an RSVP-signaled LSP leverages traffic engineering and therefore, if desired, does not require the LSP to use the IGP calculated path. RSVP-TE and CSPF will be covered in detail in Chapter 3.

The configuration of RSPV is shown here.

First, let's enable `family mpls` on the interface:

```
lab@AT> show configuration interfaces ge-0/0/2
description NE;
unit 0 {
    family inet {
        address 198.51.100.2/31;
    }
    family mpls;
}

lab@AT> show configuration interfaces ge-0/0/1
description CA;
unit 0 {
    family inet {
        address 198.51.100.1/31;
```

```
    }
    family mpls;
}
```

Now enable the RSVP protocol on the interface:

```
lab@AT> show configuration protocols rsvp
interface ge-0/0/2.0;
interface ge-0/0/1.0;
```

And enable the MPLS protocol on the interface and define the `label-switched-path`:

```
lab@AT> show configuration protocols mpls
label-switched-path AT_SJ {
    to 192.0.2.6;
}
label-switched-path AT_CO {
    to 192.0.2.4;
}
interface ge-0/0/2.0;
interface ge-0/0/1.0;
```

In the above-configured LSP, the AT LER dynamically computes an ERO to SJ and CO. Let's check to see that the two LSPs are operational:

```
lab@AT> show mpls lsp ingress
Ingress LSP: 2 sessions
To              From            State Rt P     ActivePath       LSPname
192.0.2.4       192.0.2.2       Up     0 *                      AT_CO
192.0.2.6       192.0.2.2       Up     0 *                      AT_SJ
Total 2 displayed, Up 2, Down 0
```

And here is an example of the ERO computation for the LSP named `AT_CO`:

```
lab@AT> show mpls lsp ingress name AT_CO extensive
Ingress LSP: 2 sessions

192.0.2.4
  From: 192.0.2.2, State: Up, ActiveRoute: 0, LSPname: AT_CO
  ActivePath:  (primary)
  LSPtype: Static Configured
  LoadBalance: Random
  Encoding type: Packet, Switching type: Packet, GPID: IPv4
 *Primary                    State: Up
    Priorities: 7 0
    SmartOptimizeTimer: 180
    Computed ERO (S [L] denotes strict [loose] hops): (CSPF metric: 2)
 198.51.100.0 S 198.51.100.8 S
    Received RRO (ProtectionFlag 1=Available 2=InUse 4=B/W 8=Node 10=SoftPreempt 20=Node-ID):
          198.51.100.0 198.51.100.8
    5 Jul 29 11:13:10.880 Selected as active path
    4 Jul 29 11:13:10.869 Record Route:  198.51.100.0 198.51.100.8
    3 Jul 29 11:13:10.857 Up
    2 Jul 29 11:13:10.249 Originate Call
    1 Jul 29 11:13:10.244 CSPF: computation result accepted  198.51.100.0 198.51.100.8
  Created: Sun Jul 29 11:12:40 2018
Total 1 displayed, Up 1, Down 0
```

LDP Tunneling

LDP tunneling forms an LDP session neighborship over a RSVP-TE LSP. This helps to utilize the traffic engineering features inline with LDP, instead of simply relying on IGP path selection. LDP tunneling allows the use of a one-way traffic-engineered path, therefore, a *LDP tunneled unidirectional LSP*. To take advantage of bidirectional traffic flow with the LDP label, you need to have LDP tunneling configured on both of the remote ends of the concerned LSP, as illustrated in Figure 2.2.

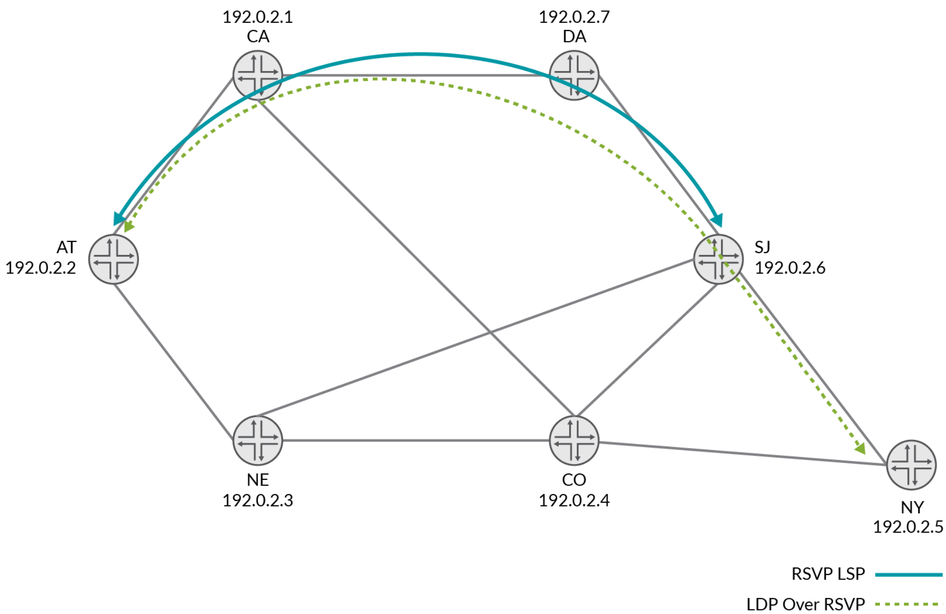

Figure 2.2 LDP Tunneling Between AT and NY

Before LDP tunneling is enabled between AT and SJ, AT and NY don't have labeled path reachability. Furthermore, the NY router is not running RSVP. We already have a LSP in place between AT to SJ, and vice versa:

```
lab@AT> show route table inet.3

inet.3: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

192.0.2.4/32       *[RSVP/7/1] 00:58:42, metric 2
                    > to 198.51.100.0 via ge-0/0/1.0, label-switched-path AT_CO
192.0.2.6/32       *[RSVP/7/1] 00:09:04, metric 2
                    > to 198.51.100.3 via ge-0/0/2.0, label-switched-path AT_SJ
```

```
lab@NY> show route table inet.3

inet.3: 3 destinations, 3 routes (3 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

192.0.2.4/32       *[LDP/9] 01:06:27, metric 1
                    > to 198.51.100.12 via ge-0/0/4.0
192.0.2.6/32       *[LDP/9] 01:08:25, metric 1
                    > to 198.51.100.15 via ge-0/0/5.0
192.0.2.7/32       *[LDP/9] 00:00:28, metric 1
                    > to 198.51.100.15 via ge-0/0/5.0, Push 299888
```

Let's configure LDP tunneling between AT and SJ:

```
lab@AT> show configuration protocols ldp
interface lo0.0;

lab@AT> show configuration protocols mpls
label-switched-path AT_SJ {
    to 192.0.2.6;
    ldp-tunneling;
}
lab@SJ> show configuration protocols ldp
interface ge-0/0/1.0;
interface ge-0/0/2.0;
interface ge-0/0/3.0;
interface ge-0/0/5.0;
interface lo0.0;

lab@SJ> show configuration protocols mpls
label-switched-path SJ_AT {
    to 192.0.2.2;
    ldp-tunneling;
}
```

You can see that the far end routers form an LDP neighborship on the Lo0 interface, hence you need to explicitly enable LDP on Lo0 and then configure `ldp-tunneling` on the required LSP.

Let's verify:

```
lab@AT> show ldp session
  Address           State        Connection     Hold time
192.0.2.6           Operational  Open             27

lab@SJ> show ldp session
  Address           State        Connection     Hold time
192.0.2.2           Operational  Open             26
192.0.2.5           Operational  Open             25
192.0.2.7           Operational  Open             25

lab@AT> show route table inet.3

inet.3: 7 destinations, 11 routes (3 active, 0 holddown, 7 hidden)
+ = Active Route, - = Last Active, * = Both
```

```
192.0.2.4/32       *[RSVP/7/1] 01:21:26, metric 2
                    > to 198.51.100.0 via ge-0/0/1.0, label-switched-path AT_CO
192.0.2.5/32       *[LDP/9] 00:21:50, metric 1
                    > to 198.51.100.3 via ge-0/0/2.0, label-switched-path AT_SJ
192.0.2.6/32       *[RSVP/7/1] 00:31:48, metric 2
                    > to 198.51.100.3 via ge-0/0/2.0, label-switched-path AT_SJ
                    [LDP/9] 00:22:03, metric 1
                    > to 198.51.100.3 via ge-0/0/2.0, label-switched-path AT_SJ

lab@NY> show route table inet.3

inet.3: 4 destinations, 4 routes (4 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

192.0.2.2/32       *[LDP/9] 00:20:52, metric 1
                    > to 198.51.100.15 via ge-0/0/5.0, Push 299936
192.0.2.4/32       *[LDP/9] 01:31:09, metric 1
                    > to 198.51.100.12 via ge-0/0/4.0
192.0.2.6/32       *[LDP/9] 01:33:07, metric 1
                    > to 198.51.100.15 via ge-0/0/5.0
192.0.2.7/32       *[LDP/9] 00:20:56, metric 1
                    > to 198.51.100.15 via ge-0/0/5.0, Push 299888

lab@SJ> show route table inet.3

inet.3: 5 destinations, 7 routes (4 active, 0 holddown, 2 hidden)
+ = Active Route, - = Last Active, * = Both

192.0.2.1/32       *[LDP/9] 00:20:51, metric 1
                    > to 198.51.100.17 via ge-0/0/1.0, Push 299872
192.0.2.2/32       *[RSVP/7/1] 00:20:45, metric 2
                    > to 198.51.100.6 via ge-0/0/3.0, label-switched-path SJ_AT
                    [LDP/9] 00:20:45, metric 1
                    > to 198.51.100.6 via ge-0/0/3.0, label-switched-path SJ_AT
192.0.2.5/32       *[LDP/9] 00:20:49, metric 1
                    > to 198.51.100.14 via ge-0/0/5.0
192.0.2.7/32       *[LDP/9] 00:20:49, metric 1
                    > to 198.51.100.17 via ge-0/0/1.0
```

The configuration of LDP tunneling results in label stacking as shown in Figure 2.3.

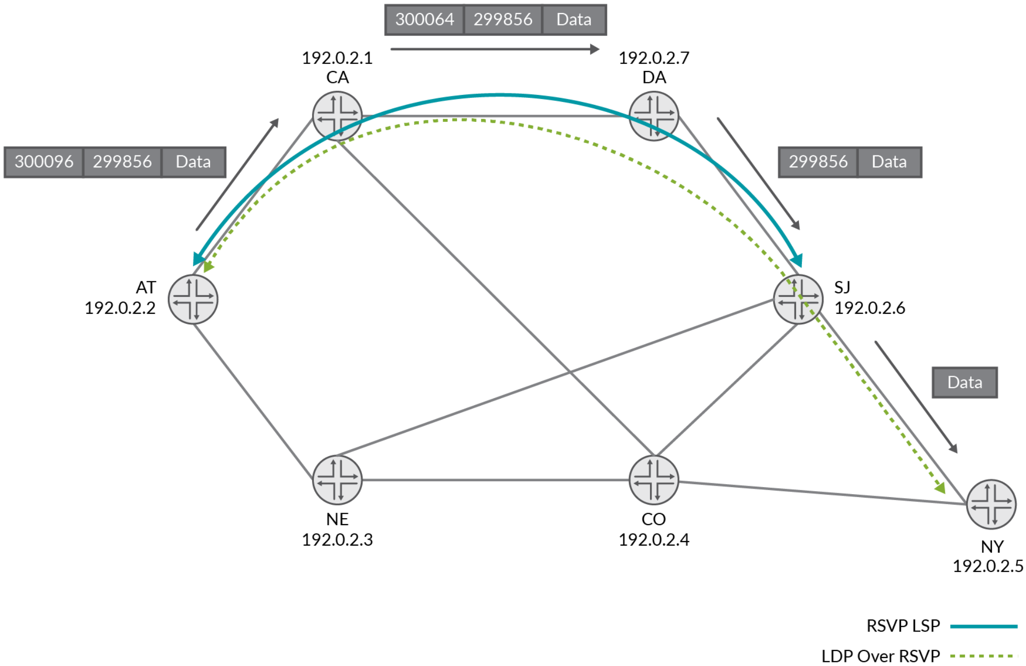

Figure 2.3 Label Stacking in LDP Tunneling

Verify on the ingress router:

```
lab@AT> show rsvp session ingress lsp name AT_SJ
Ingress RSVP: 2 sessions
To              From            State   Rt Style Labelin Labelout LSPname
192.0.2.6       192.0.2.2       Up       0  1 FF       -   300096 AT_SJ
Total 1 displayed, Up 1, Down 0

lab@AT> show ldp database
Input label database, 192.0.2.2:0--192.0.2.6:0
  Label     Prefix
 299872     192.0.2.1/32
 299936     192.0.2.2/32
 299856     192.0.2.5/32
      3     192.0.2.6/32
 299888     192.0.2.7/32

Output label database, 192.0.2.2:0--192.0.2.6:0
  Label     Prefix
      3     192.0.2.2/32
 300048     192.0.2.5/32
 300016     192.0.2.6/32
```

On the transit routers:

```
lab@CA> show rsvp session lsp transit name AT_SJ
Transit RSVP: 3 sessions
```

```
To              From            State   Rt Style Labelin Labelout LSPname
192.0.2.6       192.0.2.2       Up       0  1 FF  300096   300064 AT_SJ
Total 1 displayed, Up 1, Down 0

lab@DA> show rsvp session lsp transit name AT_SJ
Transit RSVP: 2 sessions
To              From            State   Rt Style Labelin Labelout LSPname
192.0.2.6       192.0.2.2       Up       0  1 FF  300064        3 AT_SJ
Total 1 displayed, Up 1, Down 0
```

On the RSVP LSP egress router:

```
lab@SJ> show rsvp session egress lsp name AT_SJ
Egress RSVP: 1 sessions
To              From            State   Rt Style Labelin Labelout LSPname
192.0.2.6       192.0.2.2       Up       0  1 FF       3        - AT_SJ
Total 1 displayed, Up 1, Down 0

lab@SJ> show route table mpls.0 label 299856

mpls.0: 10 destinations, 10 routes (10 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

299856             *[LDP/9] 01:27:22, metric 1
                    > to 198.51.100.14 via ge-0/0/5.0, Pop
299856(S=0)        *[LDP/9] 01:27:22, metric 1
                    > to 198.51.100.14 via ge-0/0/5.0, Pop
```

Label Stacking

By default, a maximum of three labels can be stacked on most Junos OS devices. This default behavior can be changed to the range of 3 to 5 as shown below (please consider hardware dependencies for the following command):

```
lab@NY# set interfaces ge-0/0/4.0 family mpls maximum-labels 5
```

And let's verify the stacking:

```
lab@NY> show interfaces ge-0/0/4.0
  Logical interface ge-0/0/4.0 (Index 333) (SNMP ifIndex 526)
    Description: TOWARDS_R2 ge-0/0/4.0
    Flags: Up SNMP-Traps 0x4000 VLAN-Tag [ 0x8100.123 ]
    Encapsulation: ENET2
    Input packets : 0
    Output packets: 3121
    Protocol inet, MTU: 1500
      Flags: Sendbcast-pkt-to-re
      Addresses, Flags: Is-Preferred Is-Primary
        Destination: 172.30.0.12/30, Local: 172.30.0.14,
        Broadcast: 172.30.0.15
    Protocol iso, MTU: 1497
      Flags: Is-Primary
    Protocol inet6, MTU: 1500
      Flags: Is-Primary
      Addresses, Flags: Is-Preferred
        Destination: fe80::/64, Local: fe80::205:8600:7b71:7104
    Protocol mpls, MTU: 1480, Maximum labels: 5
      Flags: Is-Primary
    Protocol multiservice, MTU: Unlimited
```

In the case of LDP tunneling, there is a RSVP label stacked over an LDP label, which consumes two labels in transit. If an MPLS L3VPN is deployed, between AT and NY there would be one more label in the stack as illustrated in Figure 2.4.

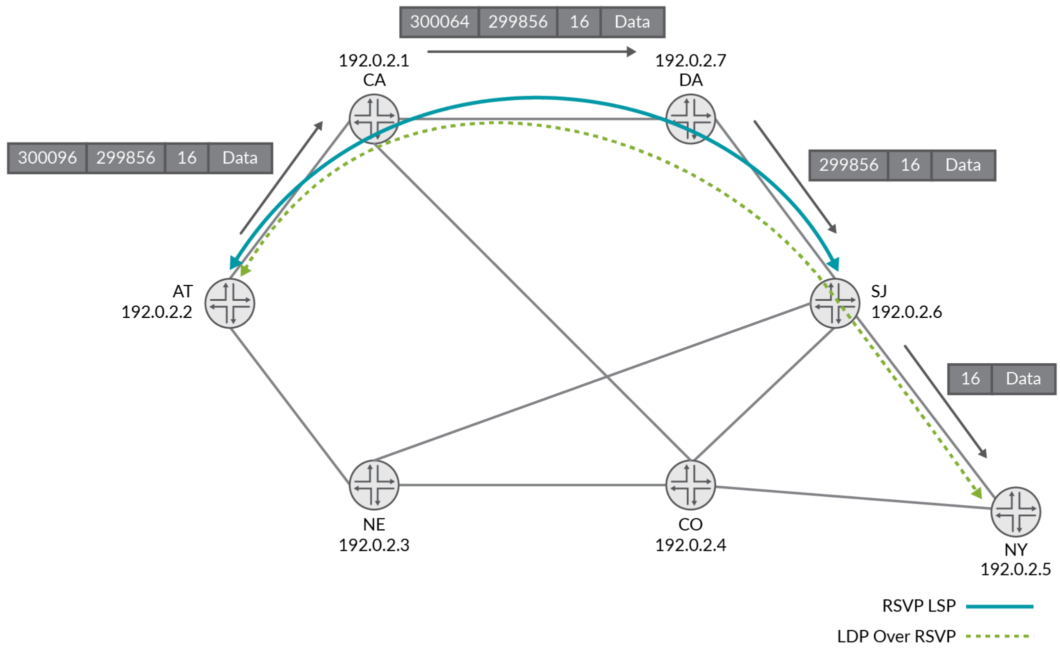

Figure 2.4 *Label Stacking*

Label 16 is the VRF label associated with a VRF present on the NY router that is popped on NY, leaving the data packet.

Summary

The answer to the question *How should we decide on which signaling protocol should be used in a Backbone network* is: *It totally depends on what level traffic engineering is needed*. RSVP signaling offers a lot of flexibility in traffic engineering, but if not much TE is needed then one can simply go ahead with LDP signaling, which follows the IGP's best path. In multi-vendor scenarios where RSVP-TE is not as developed as in the Junos OS, LDP is the preferred signaling protocol.

Chapter 3 discusses traffic engineering with RSVP signaling.

Chapter 3

Traffic Engineering

Resource Reservation Protocol- Traffic Engineering (RSVP-TE) sets up an LSP using an ERO. You can either define an ERO manually or leave this task with the routing protocol to select the optimal ERO to build the LSP. There are constraints to be aware of, however: bandwidth to be reserved, admin groups, strict or loose routing (EROs), and the priority of LSPs.

Introduction to CSPF

By default, CSPF follows the IGP traffic engineering database selected path until the specific constraints have been given for a LSP. Once the constraints are associated with the LSP, the CSPF will try to compute the path in combination with the given parameters and the shortest path selected by the IGP. The traffic engineering database (TED) is constructed with the help of the CSPF algorithm, which is referred to while building the LSPs. All the given constraints are sent via path message to downstream routers. MPLS traffic protection methods, such as fast reroute and link/node protection, rely on the traffic engineering database to set up a protected LSP for the main LSP in the event of failure.

Traffic engineering is enabled by default with the IS-IS protocol, whereas with OSPF, you need to manually enable traffic engineering for the RSVP-TE LSPs to be functional. Here is a configuration example for CSPF.

With ISIS:

```
lab@NY> show configuration protocols isis
interface ge-0/0/4.0;
interface ge-0/0/5.0;
interface lo0.0 {
    passive;
}
lab@NY> show ted protocol
Protocol name        Credibility  Self node
IS-IS(2)             2
IS-IS(1)             1
```

With OSPF:

```
lab@NY> show configuration protocols ospf
traffic-engineering;
area 0.0.0.0 {
    interface ge-0/0/4.0;
    interface ge-0/0/5.0;
    interface lo0.0 {
        passive;
    }
}
lab@NY> show ted protocol
Protocol name        Credibility  Self node
OSPF(0)              0            192.0.2.5
```

And here is the traffic engineering database sample:

```
lab@NY> show ted database
TED database: 0 ISIS nodes 17 INET nodes
ID                            Type Age(s) LnkIn LnkOut Protocol
198.51.100.0-1                Net     362     0      2 OSPF(0.0.0.0)
    To: 192.0.2.1, Local: 0.0.0.0, Remote: 0.0.0.0
      Local interface index: 0, Remote interface index: 0
    To: 192.0.2.2, Local: 0.0.0.0, Remote: 0.0.0.0
      Local interface index: 0, Remote interface index: 0
ID                            Type Age(s) LnkIn LnkOut Protocol
192.0.2.1                     Rtr     362     3      0 OSPF(0.0.0.0)
192.0.2.2                     Rtr     362     2      0 OSPF(0.0.0.0)
198.51.100.2-1                Net     362     0      2 OSPF(0.0.0.0)
    To: 192.0.2.2, Local: 0.0.0.0, Remote: 0.0.0.0
      Local interface index: 0, Remote interface index: 0
    To: 192.0.2.3, Local: 0.0.0.0, Remote: 0.0.0.0
      Local interface index: 0, Remote interface index: 0
ID                            Type Age(s) LnkIn LnkOut Protocol
192.0.2.3                     Rtr     362     3      0 OSPF(0.0.0.0)
192.0.2.4                     Rtr     362     4      0 OSPF(0.0.0.0)
192.0.2.5                     Rtr     363     2      0 OSPF(0.0.0.0)
198.51.100.5-1                Net     362     0      2 OSPF(0.0.0.0)
    To: 192.0.2.3, Local: 0.0.0.0, Remote: 0.0.0.0
      Local interface index: 0, Remote interface index: 0
    To: 192.0.2.4, Local: 0.0.0.0, Remote: 0.0.0.0
      Local interface index: 0, Remote interface index: 0
ID                            Type Age(s) LnkIn LnkOut Protocol
192.0.2.6                     Rtr     362     4      0 OSPF(0.0.0.0)
198.51.100.6-1                Net     362     0      2 OSPF(0.0.0.0)
    To: 192.0.2.3, Local: 0.0.0.0, Remote: 0.0.0.0
```

```
      Local interface index: 0, Remote interface index: 0
    To: 192.0.2.6, Local: 0.0.0.0, Remote: 0.0.0.0
      Local interface index: 0, Remote interface index: 0
ID                            Type Age(s) LnkIn LnkOut Protocol
192.0.2.7                     Rtr     362     2      0 OSPF(0.0.0.0)
198.51.100.9-1                Net     363     0      2 OSPF(0.0.0.0)
    To: 192.0.2.1, Local: 0.0.0.0, Remote: 0.0.0.0
      Local interface index: 0, Remote interface index: 0
    To: 192.0.2.4, Local: 0.0.0.0, Remote: 0.0.0.0
      Local interface index: 0, Remote interface index: 0
ID                            Type Age(s) LnkIn LnkOut Protocol
198.51.100.10-1               Net     363     0      2 OSPF(0.0.0.0)
    To: 192.0.2.4, Local: 0.0.0.0, Remote: 0.0.0.0
      Local interface index: 0, Remote interface index: 0
    To: 192.0.2.6, Local: 0.0.0.0, Remote: 0.0.0.0
      Local interface index: 0, Remote interface index: 0
ID                            Type Age(s) LnkIn LnkOut Protocol
198.51.100.12-1               Net     354     0      2 OSPF(0.0.0.0)
    To: 192.0.2.5, Local: 0.0.0.0, Remote: 0.0.0.0
      Local interface index: 0, Remote interface index: 0
    To: 192.0.2.4, Local: 0.0.0.0, Remote: 0.0.0.0
      Local interface index: 0, Remote interface index: 0
ID                            Type Age(s) LnkIn LnkOut Protocol
198.51.100.15-1               Net     361     0      2 OSPF(0.0.0.0)
    To: 192.0.2.5, Local: 0.0.0.0, Remote: 0.0.0.0
      Local interface index: 0, Remote interface index: 0
    To: 192.0.2.6, Local: 0.0.0.0, Remote: 0.0.0.0
      Local interface index: 0, Remote interface index: 0
ID                            Type Age(s) LnkIn LnkOut Protocol
198.51.100.16-1               Net     363     0      2 OSPF(0.0.0.0)
    To: 192.0.2.6, Local: 0.0.0.0, Remote: 0.0.0.0
      Local interface index: 0, Remote interface index: 0
    To: 192.0.2.7, Local: 0.0.0.0, Remote: 0.0.0.0
      Local interface index: 0, Remote interface index: 0
ID                            Type Age(s) LnkIn LnkOut Protocol
198.51.100.18-1               Net     363     0      2 OSPF(0.0.0.0)
    To: 192.0.2.1, Local: 0.0.0.0, Remote: 0.0.0.0
      Local interface index: 0, Remote interface index: 0
    To: 192.0.2.7, Local: 0.0.0.0, Remote: 0.0.0.0
      Local interface index: 0, Remote interface index: 0
```

RSVP-TE

RSVP-TE helps to avoid overutilization of a link that has been selected as the active route by the IGP. RSVP TE also helps to set up a LSP per the bandwidth reservation that was configured for the LSP, and if the configured bandwidth is not available in any of the transit paths then the LSP fails to come up. Administrative groups, also called *link coloring*, helps to achieve traffic engineering with its 32-named group range, from values 0 through 31, and is carried with the IGP in a 32-bit value. You can either associate a link with a color, leave it uncolored, or assign multiple colors to a single link as shown in Figure 3.1.

NOTE You must set identical administrative groups on all routers within a single MPLS domain.

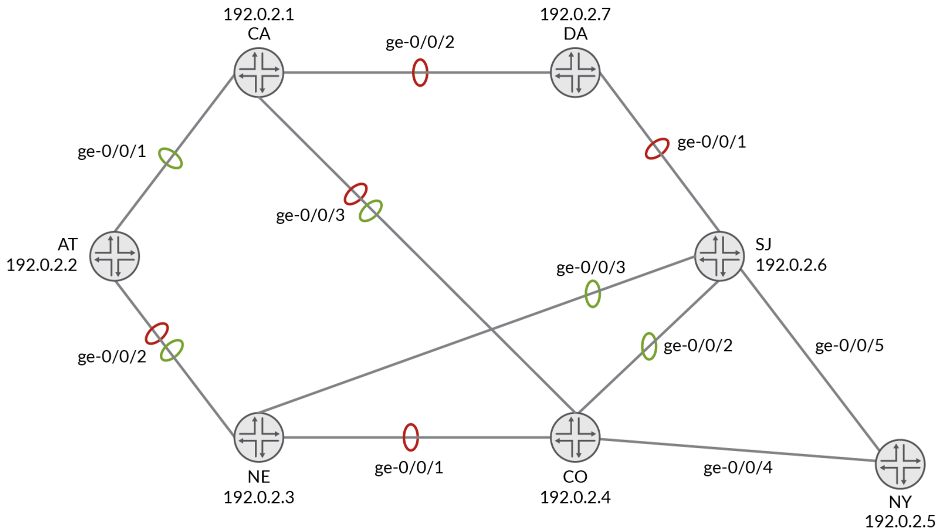

Figure 3.1 Link Coloring

Here's the configuration for Figure 3.1:

```
lab@SJ> show configuration protocols mpls
admin-groups {
    GREEN 0;
    RED 1;
}
label-switched-path SJ_AT {
    to 192.0.2.2;
    ldp-tunneling;
}
interface ge-0/0/1.0 {
    admin-group RED;
}
interface ge-0/0/2.0 {
    admin-group GREEN;
}
interface ge-0/0/3.0 {
    admin-group GREEN;
}
interface ge-0/0/5.0;
```

You can see that the same `admin-groups` are configured on all the routers in this MPLS domain; also, the router interfaces are associated with the colors shown in Figure 3.1.

Let's verify:

```
lab@SJ> show mpls interface
Interface        State      Administrative groups (x: extended)
ge-0/0/1.0       Up         RED
ge-0/0/2.0       Up         GREEN
ge-0/0/3.0       Up         GREEN
ge-0/0/5.0       Up        <none>
```

And now the RSVP-TE LSP configuration. Let's configure SJ_AT with the help of `admin-group` so traffic only follows the RED-colored links, avoiding the GREEN:

```
lab@SJ> show configuration protocols mpls label-switched-path SJ_AT
to 192.0.2.2;
ldp-tunneling;
admin-group include-any RED;
```

And verify:

```
lab@SJ> show mpls lsp ingress name SJ_AT detail
Ingress LSP: 1 sessions

192.0.2.2
  From: 192.0.2.6, State: Up, ActiveRoute: 0, LSPname: SJ_AT
  ActivePath:  (primary)
  LSPtype: Static Configured
  LoadBalance: Random
  Encoding type: Packet, Switching type: Packet, GPID: IPv4
 *Primary                    State: Up
    Priorities: 7 0
    SmartOptimizeTimer: 180
          Include Any: RED
    Computed ERO (S [L] denotes strict [loose] hops): (CSPF metric: 5)
 198.51.100.17 S 198.51.100.18 S 198.51.100.8 S 198.51.100.4 S 198.51.100.2 S
    Received RRO (ProtectionFlag 1=Available 2=InUse 4=B/W 8=Node 10=SoftPreempt 20=Node-ID):
          198.51.100.17 198.51.100.18 198.51.100.8 198.51.100.4 198.51.100.2
Total 1 displayed, Up 1, Down 0
```

The result is depicted in Figure 3.2 on the next page.

You can also use various knobs with `admin-groups` to get more granular control, as listed here:

```
lab@SJ# set protocols mpls label-switched-path SJ_AT admin-group ?
Possible completions:
+ apply-groups         Groups from which to inherit configuration data
+ apply-groups-except  Don't inherit configuration data from these groups
+ exclude              Groups, all of which must be absent
+ include-all          Groups, all of which must be present
+ include-any          Groups, one or more of which must be present
```

The knobs are self-explanatory, but be sure to note something that's not listed here: `exclude` doesn't exclude the uncolored link, but uncolored links are disqualified when used with the `include-any` or `include-all` knob.

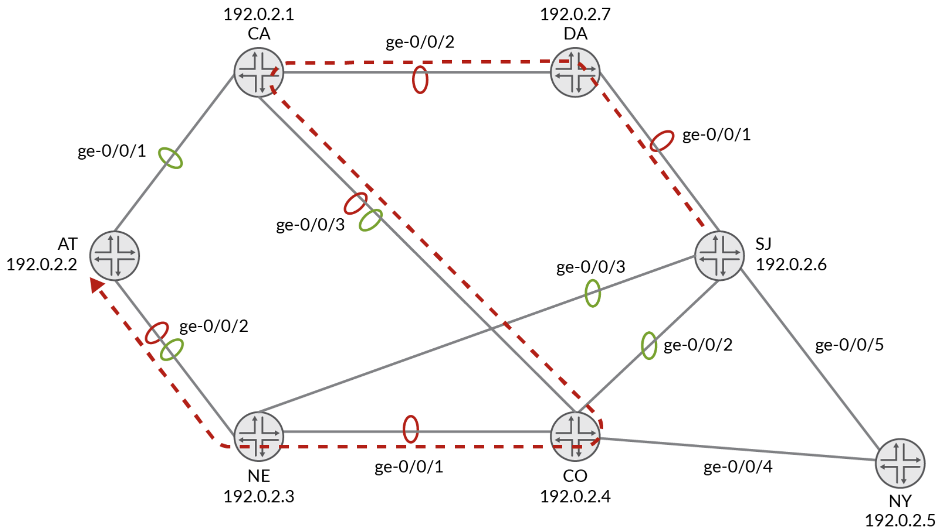

Figure 3.2 RSVP-TE path for LSP ST_AT

Fast Reroute

Fast reroute pre-establishes a temporary backup LSP to avoid real time packet loss in the event of an LSP failure.

A detour object is sent in the RSVP PATH message to the downstream routers while computing ERO for the main LSP. At each hop, the main LSP computes the *detour LSP* or how the packet can reach the egress router if the LSP the traffic is traversing, fails.

Fast reroute plays a vital role in avoiding packet loss in the event of LSP failure. While the reservation TEAR message is sent to the ingress router, which then further sends a PATH message downstream trying to establish a new LSP, the detour path is being used to avoid traffic disruption. Even when a secondary LSP is in a standby state, the already traversing packets will be dropped if the fast reroute is not considered. Once the new LSP, or secondary LSP is up, the traffic will be sent over these new LSPs instead of the detour LSP. Hence, fast reroute detours LSP functions as a one-to-one backup LSP.

Fast reroute builds this detour LSP with the maximum hop-limit of 6. You can customize that with a lower `hop-limit` per your requirements and it inherits the `admin-groups` from the main LSP, if any. Furthermore, when creating a detour LSP it doesn't consider the bandwidth reservation given with the main LSP.

You can see the change in our topology in Figure 3.3.

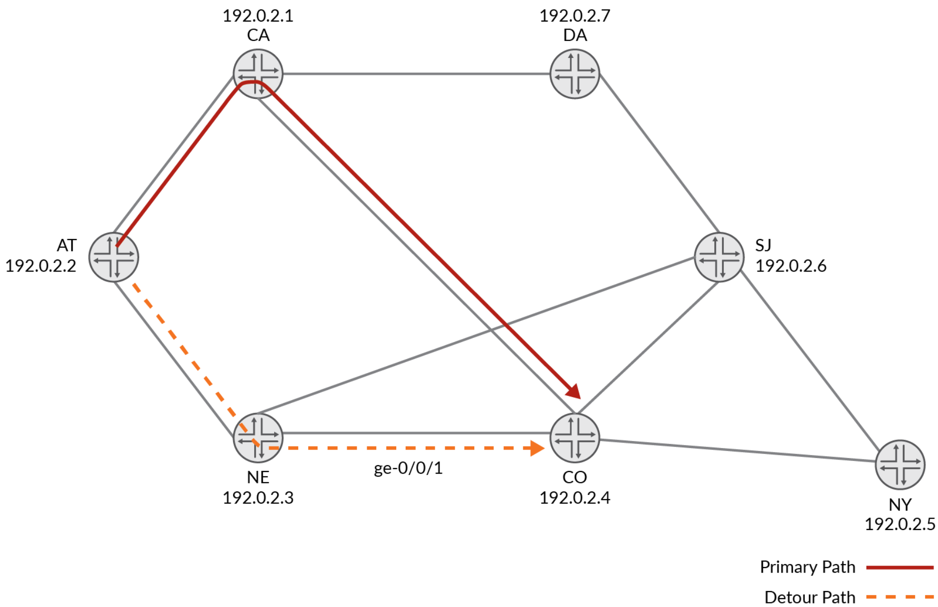

Figure 3.3 Fast Reroute

Configuration of the LSP between AT and CO "fast-reroute" has been added to the LSP configuration:

```
lab@AT> show configuration protocols mpls label-switched-path AT_CO
to 192.0.2.4;
fast-reroute;
```

Verification:

```
lab@AT> show rsvp session ingress lsp name AT_CO extensive
Ingress RSVP: 2 sessions

192.0.2.4
  From: 192.0.2.2, LSPstate: Up, ActiveRoute: 0
  LSPname: AT_CO, LSPpath: Primary
  LSPtype: Static Configured
  Suggested label received: -, Suggested label sent: -
  Recovery label received: -, Recovery label sent: 300160
  Resv style: 1 FF, Label in: -, Label out: 300160
  Time left:    -, Since: Sun Jul 29 14:42:37 2018
  Tspec: rate 0bps size 0bps peak Infbps m 20 M 1500
  Port number: sender 3 receiver 64532 protocol 0
  FastReroute desired
  PATH rcvfrom: localclient
  Adspec: sent MTU 1500
  Path MTU: received 1500
```

```
  PATH sentto: 198.51.100.0 (ge-0/0/1.0) 10 pkts
  RESV rcvfrom: 198.51.100.0 (ge-0/0/1.0) 9 pkts
  Explct route: 198.51.100.0 198.51.100.8
  Record route: <self> 198.51.100.0 198.51.100.8
    Detour is Up
    Detour Tspec: rate 0bps size 0bps peak Infbps m 20 M 1500
    Detour adspec: sent MTU 1500
    Path MTU: received 1500
    Detour PATH sentto: 198.51.100.3 (ge-0/0/2.0) 7 pkts
    Detour RESV rcvfrom: 198.51.100.3 (ge-0/0/2.0) 4 pkts
    Detour Explct route: 198.51.100.3 198.51.100.5
    Detour Record route: <self> 198.51.100.3 198.51.100.5
    Detour Label out: 300032
  Detour branch from 198.51.100.0, to skip 192.0.2.4, Up
    Tspec: rate 0bps size 0bps peak Infbps m 20 M 1500
    Adspec: received MTU 1500
    Path MTU: received 0
    PATH rcvfrom: 198.51.100.0 (ge-0/0/1.0) 9 pkts
    Adspec: received MTU 1500 sent MTU 1500
    PATH sentto: 198.51.100.3 (ge-0/0/2.0) 2 pkts
    RESV rcvfrom: 198.51.100.3 (ge-0/0/2.0) 2 pkts
    Explct route: 198.51.100.3 198.51.100.5
    Record route: 198.51.100.1 198.51.100.0 <self> 198.51.100.3 198.51.100.5
    Label in: 300080, Label out: 300032
Total 1 displayed, Up 1, Down 0
```

When the LSP fails, the routers require time to install the fast reroute detour LSP into the forwarding table. To avoid delays in switching LSPs, in the event of a main RSVP-TE LSP failure you should configure a `load-balancing` policy for the forwarding table to install the fast reroute detour LSP in the PFE ahead of time.

The configuration is here:

```
lab@AT> show configuration policy-options policy-statement LOADBALANCE
term 1 {
    then {
        load-balance per-packet;
    }
}

lab@AT> show configuration routing-options forwarding-table
export LOADBALANCE;
```

Remember, you can always explicitly define bandwidth and other traffic engineering parameters for detour path while in the CLI:

```
[edit]
lab@AT# set protocols mpls label-switched-path AT_CO fast-reroute ?
Possible completions:
  <[Enter]>            Execute this command
+ apply-groups         Groups from which to inherit configuration data
+ apply-groups-except  Don't inherit configuration data from these groups
  bandwidth            Bandwidth to reserve (bps)
  bandwidth-percent    Percentage of main path bandwidth to reserve (1..100)
+ exclude              Groups, all of which must be absent
  hop-limit            Maximum allowed router hops (0..255)
+ include-all          Groups, all of which must be present
```

```
+ include-any          Groups, one or more of which must be present
  no-exclude           Disable exclude checking
  no-include-all       Disable include-all checking
  no-include-any       Disable include-any checking
  |                    Pipe through a command
```

Bypass LSP

Bypass LSPs are built to reduce the impact of any link/node failure in the RSVP-TE path, and they are built around protected interfaces or protected nodes. The difference between fast reroute and bypass LSPs is that fast reroute attempts to protect the entire LSP by installing a detour LSP in advance, whereas the bypass LSP pre-establishes a bypass LSP, protecting a link failure or node failure in the RSVP-TE LSP ERO. Bypass LSPs use CSPF to compute the protected ERO and they also function as “many-to-one” backup LSPs. Link and node protection are discussed in detail in the next sections.

Link Protection

In the case of link protection, bypass LSPs are pre-established against a link failure in RSVP-TE LSP.

When link protection is configured along with the LSP, all the transit LSRs attempt to build an alternate path to reach the next hop in the LSP. This is called *next-hop bypass LSP* and the bypass LSPs are built after the main LSP comes up. The ingress routers flag the main LSP to pre-establish a link-protected bypass LSP and this is again flagged at the ingress router and sent with the path messages. Further, the RSVP link/interface, where the link-protection is explicitly enabled/configured, sees the link-protection flag, and will try to establish a next-hop bypass LSP. It’s an LSP created between two adjacent neighbors.

In the next configuration example, the main LSP from NT_AT takes the NY-SJ-NE-AT path. To better understand the concept of link protection, we have enabled `link-protection` at the ingress and protected the path between SJ-NE. Hence, the next-hop bypass LSP is NY-SJ-CO-NE-AT.

Here, the NE router is called the *merge point* and the SJ router is the *point of local repair* (PLR). Merge point is where the alternate path rejoins the main LSP. The PLR is that node that notices the failure of the downstream link or downstream node. PLR starts forwarding through alternate paths in case of link or node failure. It also informs the ingress router about the failure with a PathErr message. This setup is shown in Figure 3.4.

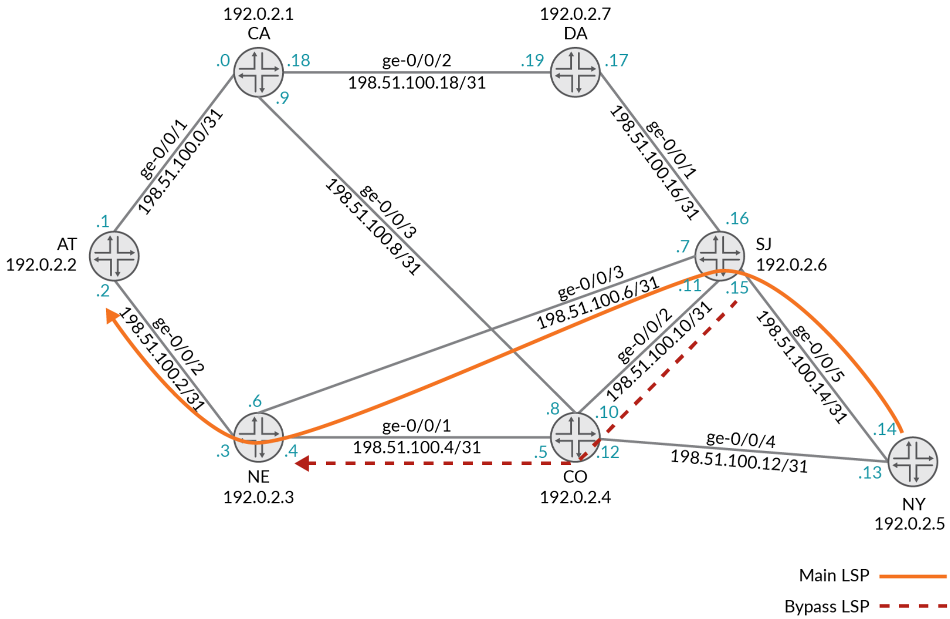

Figure 3.4 *Link Protection*

Here is the link protection configuration:

```
lab@NY> show configuration protocols mpls
label-switched-path NY_AT {
    to 192.0.2.2;
    link-protection;
    primary PATH_AT;
}
path PATH_AT {
    192.0.2.6;
    192.0.2.3;
}
interface ge-0/0/4.0;
interface ge-0/0/5.0;

lab@SJ> show configuration protocols rsvp
interface ge-0/0/1.0;
interface ge-0/0/2.0;
interface ge-0/0/3.0 {
    link-protection;
}
interface ge-0/0/5.0;
```

And to verify:

```
lab@NY> show rsvp session lsp name NY_AT extensive
Ingress RSVP: 1 sessions
```

```
192.0.2.2
  From: 192.0.2.5, LSPstate: Up, ActiveRoute: 0
  LSPname: NY_AT, LSPpath: Primary
  LSPtype: Static Configured
  Suggested label received: -, Suggested label sent: -
  Recovery label received: -, Recovery label sent: 300000
  Resv style: 1 SE, Label in: -, Label out: 300000
  Time left:    -, Since: Sun Jul 29 17:58:56 2018
  Tspec: rate 0bps size 0bps peak Infbps m 20 M 1500
  Port number: sender 1 receiver 4170 protocol 0
  Link protection desired
  Type: Protection down
      1 Jul 29 18:22:21 Link protection disabled on outbound interface[35 times]
  PATH rcvfrom: localclient
  Adspec: sent MTU 1500
  Path MTU: received 1500
  PATH sentto: 198.51.100.15 (ge-0/0/5.0) 34 pkts
  RESV rcvfrom: 198.51.100.15 (ge-0/0/5.0) 34 pkts
  Explct route: 198.51.100.15 198.51.100.6 198.51.100.2
  Record route: <self> 192.0.2.6 (node-id) 198.51.100.15 192.0.2.3 (node-id)
  198.51.100.6 192.0.2.2 (node-id) 198.51.100.2
```

NOTE Link protection on the NY link ge-0/0/5.0 is down since we are only protecting the link between SJ-NE.

```
lab@SJ> show rsvp session transit lsp name NY_AT extensive
Transit RSVP: 1 sessions

192.0.2.2
  From: 192.0.2.5, LSPstate: Up, ActiveRoute: 0
  LSPname: NY_AT, LSPpath: Primary
  Suggested label received: -, Suggested label sent: -
  Recovery label received: -, Recovery label sent: 300096
  Resv style: 1 SE, Label in: 300000, Label out: 300096
  Time left:  137, Since: Sun Jul 29 17:58:55 2018
  Tspec: rate 0bps size 0bps peak Infbps m 20 M 1500
  Port number: sender 1 receiver 4170 protocol 0
  Link protection desired
  Type: Link protected LSP, using Bypass->198.51.100.6
      2 Jul 29 17:58:58 Link protection up, using Bypass->198.51.100.6
      1 Jul 29 17:58:56 New bypass Bypass->198.51.100.6
  PATH rcvfrom: 198.51.100.14 (ge-0/0/5.0) 35 pkts
  Adspec: received MTU 1500 sent MTU 1500
  PATH sentto: 198.51.100.6 (ge-0/0/3.0) 35 pkts
  RESV rcvfrom: 198.51.100.6 (ge-0/0/3.0) 35 pkts
  Explct route: 198.51.100.6 198.51.100.2
  Record route: 198.51.100.14 <self> 192.0.2.3 (node-id) 198.51.100.6
  192.0.2.2 (node-id) 198.51.100.2
Total 1 displayed, Up 1, Down 0

lab@CO> show rsvp session lsp transit name Bypass->198.51.100.6 extensive
Transit RSVP: 2 sessions

192.0.2.3
  From: 192.0.2.6, LSPstate: Up, ActiveRoute: 0
  LSPname: Bypass->198.51.100.6
```

```
  Suggested label received: -, Suggested label sent: -
  Recovery label received: -, Recovery label sent: 3
  Resv style: 1 SE, Label in: 300000, Label out: 3
  Time left:  129, Since: Sun Jul 29 17:58:55 2018
  Tspec: rate 0bps size 0bps peak Infbps m 20 M 1500
  Port number: sender 1 receiver 9313 protocol 0
  PATH rcvfrom: 198.51.100.11 (ge-0/0/2.0) 36 pkts
  Adspec: received MTU 1500 sent MTU 1500
  PATH sentto: 198.51.100.4 (ge-0/0/1.0) 36 pkts
  RESV rcvfrom: 198.51.100.4 (ge-0/0/1.0) 36 pkts
  Explct route: 198.51.100.4
  Record route: 198.51.100.11 <self> 198.51.100.4
Total 1 displayed, Up 1, Down 0
```

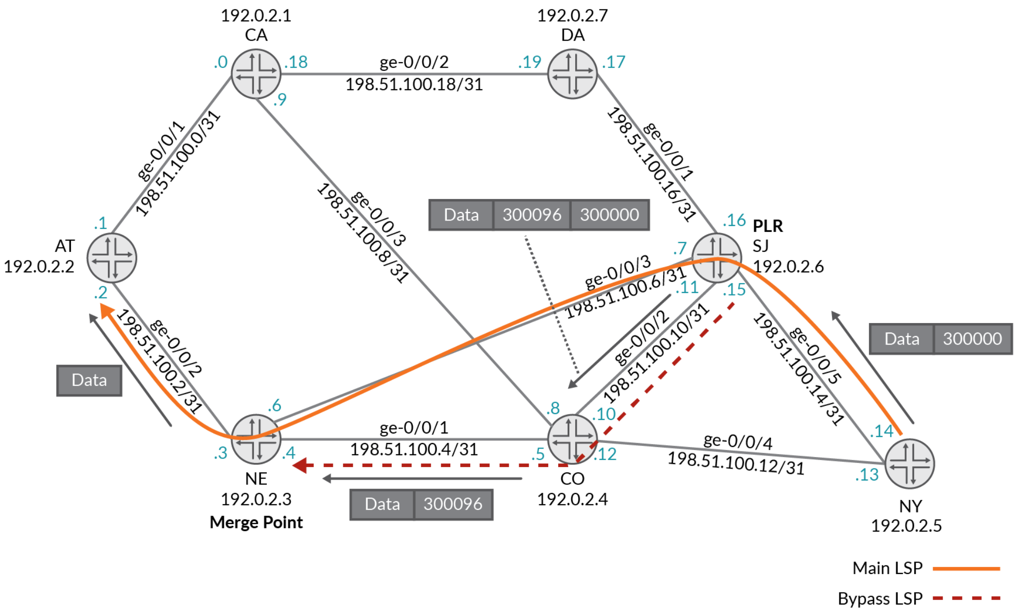

Figure 3.5 Next-hop Bypass LSP Label Stacking

Node-link Protection

In the case of node-link protection, bypass LSPs are pre-established against a node failure in the RSVP-TE LSP.

When node-link protection is configured along the LSP all the transit LSRs attempt to build an alternate path to reach the next hop's next hop in the LSP. This is therefore known as a *next-next-hop bypass LSP*.

The bypass LSPs are built after the main LSP comes up. The ingress routers flag the main LSP to pre-establish a node-link-protected bypass LSP. This is again flagged at the ingress router and signaled with the path messages. Further, seeing the node-link-protection flag, the RSVP link, where link protection is explicitly enabled, will try to establish a *bypass next-next-hop LSP*. You could call it an *LSP created to bypass the immediate downstream node of the main LSP*. If node protection is unavailable for some reason the LSP attempts to use link protection. This is illustrated in Figure 3.6.

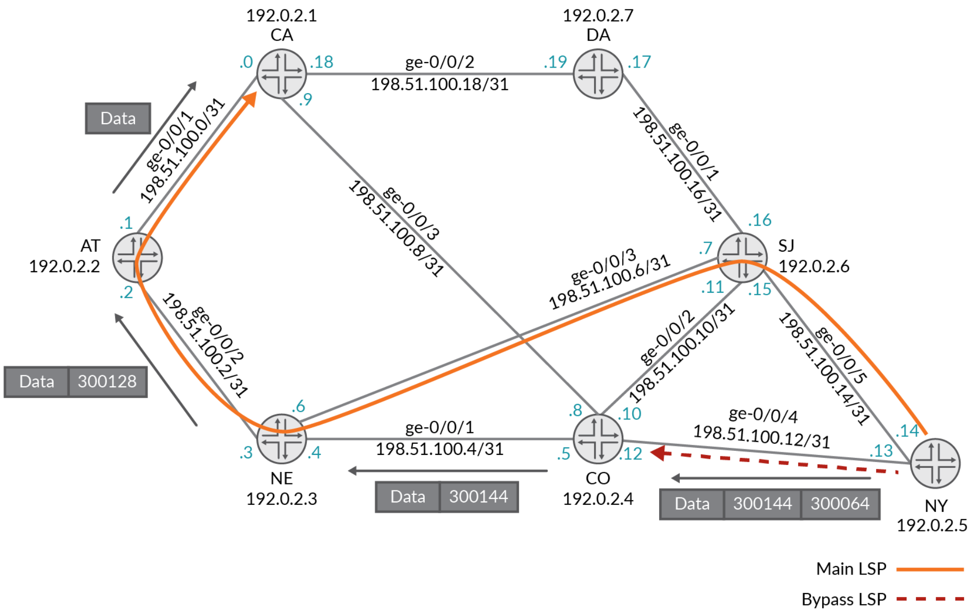

Figure 3.6 Next-Next Hop Bypass LSP Label Stacking

And the configuration:

```
lab@NY> show configuration protocols mpls label-switched-path NY_CA
to 192.0.2.1;
node-link-protection;
primary PATH_AT;

lab@NY> show configuration protocols rsvp
interface ge-0/0/4.0;
interface ge-0/0/5.0 {
    link-protection;
}
```

With verification on the ingress NY router:

```
lab@NY> show rsvp session lsp ingress
Ingress RSVP: 2 sessions
To             From           State   Rt Style Labelin Labelout LSPname
```

```
192.0.2.1      192.0.2.5      Up      0  1 SE      -   300048 NY_CA
192.0.2.3      192.0.2.5      Up      0  1 SE      -   300064 Bypass->198.51.100.15->198.51.100.6
Total 2 displayed, Up 2, Down 0

lab@NY> show rsvp session lsp name NY_CA extensive
Ingress RSVP: 2 sessions

192.0.2.1
  From: 192.0.2.5, LSPstate: Up, ActiveRoute: 0
  LSPname: NY_CA, LSPpath: Primary
  LSPtype: Static Configured
  Suggested label received: -, Suggested label sent: -
  Recovery label received: -, Recovery label sent: 300048
  Resv style: 1 SE, Label in: -, Label out: 300048
  Time left:    -, Since: Sun Jul 29 20:02:33 2018
  Tspec: rate 0bps size 0bps peak Infbps m 20 M 1500
  Port number: sender 3 receiver 4172 protocol 0
  Node/Link protection desired
  Type: Node/Link protected LSP, using Bypass->198.51.100.15->198.51.100.6
      1 Jul 29 20:02:42 Node protection up, using Bypass->198.51.100.15->198.51.100.6
  PATH rcvfrom: localclient
  Adspec: sent MTU 1500
  Path MTU: received 1500
  PATH sentto: 198.51.100.15 (ge-0/0/5.0) 4 pkts
  RESV rcvfrom: 198.51.100.15 (ge-0/0/5.0) 4 pkts
  Explct route: 198.51.100.15 198.51.100.6 198.51.100.2 198.51.100.0
  Record route: <self> 192.0.2.6 (node-id) 198.51.100.15 192.0.2.3 (node-id)
  198.51.100.6 192.0.2.2 (node-id) 198.51.100.2 192.0.2.1 (node-id) 198.51.100.0
Total 1 displayed, Up 1, Down 0

lab@NY> show route table inet.3 192.0.2.1/32 extensive

inet.3: 4 destinations, 4 routes (4 active, 0 holddown, 0 hidden)
192.0.2.1/32 (1 entry, 1 announced)
        State: <FlashAll>
        *RSVP   Preference: 7/1
                Next hop type: Router
                Address: 0x938c0e0
                Next-hop reference count: 2
                Next hop: 198.51.100.15 via ge-0/0/5.0 weight 0x1, selected
                Label-switched-path NY_CA
                Label operation: Push 300048
                Label TTL action: prop-ttl
                Next hop: 198.51.100.12 via ge-0/0/4.0 weight 0x8001
                Label-switched-path Bypass->198.51.100.15->198.51.100.6
                Label operation: Push 300144, Push 300064(top)
                Label TTL action: prop-ttl, prop-ttl(top)
                State: <Active Int>
                Age: 1:42:16    Metric: 2
                Task: RSVP
                Announcement bits (1): 1-Resolve tree 1
                AS path: I
```

This output indicates that for the next-next hop bypass LSP the NY router will push labels 300144 and 300064. Label 300144, the label expected by the NE router for the NY_CA transit LSP, is captured below.

On the node used as the bypass router:

```
lab@CO> show rsvp session transit lsp
Transit RSVP: 2 sessions
To              From            State   Rt Style Labelin Labelout LSPname
192.0.2.2       192.0.2.6       Up       0  1 FF  299968   300048 SJ_AT
192.0.2.3       192.0.2.5       Up       0  1 SE  300064        3 Bypass->198.51.100.1
          5->198.51.100.6
Total 2 displayed, Up 2, Down 0

lab@CO> show rsvp session transit lsp name Bypass->198.51.100.15->198.51.100.6 extensive
Transit RSVP: 2 sessions

192.0.2.3
  From: 192.0.2.5, LSPstate: Up, ActiveRoute: 0
  LSPname: Bypass->198.51.100.15->198.51.100.6
  Suggested label received: -, Suggested label sent: -
  Recovery label received: -, Recovery label sent: 3
  Resv style: 1 SE, Label in: 300064, Label out: 3
  Time left:  134, Since: Sun Jul 29 19:06:10 2018
  Tspec: rate 0bps size 0bps peak Infbps m 20 M 1500
  Port number: sender 1 receiver 4174 protocol 0
  PATH rcvfrom: 198.51.100.13 (ge-0/0/4.0) 118 pkts
  Adspec: received MTU 1500 sent MTU 1500
  PATH sentto: 198.51.100.4 (ge-0/0/1.0) 118 pkts
  RESV rcvfrom: 198.51.100.4 (ge-0/0/1.0) 118 pkts
  Explct route: 198.51.100.4
  Record route: 198.51.100.13 <self> 198.51.100.4
Total 1 displayed, Up 1, Down 0
```

On the next-next hop router of the main LSP:

```
lab@NE> show rsvp session lsp
Ingress RSVP: 0 sessions
Total 0 displayed, Up 0, Down 0

Egress RSVP: 1 sessions
To             From          State   Rt Style Labelin Labelout LSPname
192.0.2.3      192.0.2.5      Up      0  1 SE      3       - Bypass->198.51.100.15->198.51.100.6
Total 1 displayed, Up 1, Down 0

Transit RSVP: 4 sessions
To              From            State   Rt Style Labelin Labelout LSPname
192.0.2.1       192.0.2.5       Up       0  1 SE  300144   300128 NY_CA
192.0.2.2       192.0.2.6       Up       0  1 FF  300048        3 SJ_AT
192.0.2.4       192.0.2.2       Up       0  1 FF  300032        3 AT_CO
192.0.2.6       192.0.2.2       Up       0  1 FF  300080        3 AT_SJ
Total 4 displayed, Up 4, Down 0

lab@NE> show rsvp session lsp name Bypass->198.51.100.15->198.51.100.6 extensive
Ingress RSVP: 0 sessions
Total 0 displayed, Up 0, Down 0

Egress RSVP: 1 sessions
```

```
192.0.2.3
  From: 192.0.2.5, LSPstate: Up, ActiveRoute: 0
  LSPname: Bypass->198.51.100.15->198.51.100.6
  Suggested label received: -, Suggested label sent: -
  Recovery label received: -, Recovery label sent: -
  Resv style: 1 SE, Label in: 3, Label out: -
  Time left:  155, Since: Sun Jul 29 19:06:10 2018
  Tspec: rate 0bps size 0bps peak Infbps m 20 M 1500
  Port number: sender 1 receiver 4174 protocol 0
  PATH rcvfrom: 198.51.100.5 (ge-0/0/1.0) 146 pkts
  Adspec: received MTU 1500
  PATH sentto: localclient
  RESV rcvfrom: localclient
  Record route: 198.51.100.13 198.51.100.5 <self>
Total 1 displayed, Up 1, Down 0

Transit RSVP: 4 sessions
Total 0 displayed, Up 0, Down 0

lab@NE> show rsvp session lsp name NY_CA extensive
Ingress RSVP: 0 sessions
Total 0 displayed, Up 0, Down 0

Egress RSVP: 1 sessions
Total 0 displayed, Up 0, Down 0

Transit RSVP: 4 sessions

192.0.2.1
  From: 192.0.2.5, LSPstate: Up, ActiveRoute: 0
  LSPname: NY_CA, LSPpath: Primary
  Suggested label received: -, Suggested label sent: -
  Recovery label received: -, Recovery label sent: 300128
  Resv style: 1 SE, Label in: 300144, Label out: 300128
  Time left:  138, Since: Sun Jul 29 20:02:31 2018
  Tspec: rate 0bps size 0bps peak Infbps m 20 M 1500
  Port number: sender 3 receiver 4172 protocol 0
  Node/Link protection desired
  Type: Protection down
      1 Jul 29 20:53:47 Link protection disabled on outbound interface[71 times]
  PATH rcvfrom: 198.51.100.7 (ge-0/0/3.0) 71 pkts
  Adspec: received MTU 1500 sent MTU 1500
  PATH sentto: 198.51.100.2 (ge-0/0/2.0) 71 pkts
  RESV rcvfrom: 198.51.100.2 (ge-0/0/2.0) 71 pkts
  Explct route: 198.51.100.2 198.51.100.0
  Record route: 198.51.100.14 198.51.100.7 <self> 192.0.2.2 (node-id) 198.51.100.2 192.0.2.1 (node-
id) 198.51.100.0
Total 1 displayed, Up 1, Down 0
```

Here, the bypass LSP was built around CO to bypass the SJ node since the main LSP was NY-SJ-NE-AT-CA. The next-next hop bypass LSP was built to reach NE. The NE router is the merge point.

NOTE All protection LSPs (fast reroute or bypass LSP) share bandwidth with the main LSP while under reservation.

Choosing between multiple protection mechanisms to protect an LSP can be a task. During implementation of either the bypass LSP or fast reroute protection mechanism, the size, vendor interoperability, and scaling dimensions of the MPLS network should also be considered. Chapter 5 drops a hint about how you can select the suited mechanism for your MPLS network.

Summary

With RSVP-TE the network design should not be an issue until all the bandwidth constraints and link colorings are placed/designed as expected. An example being within a large network with hundreds of LSPs – not all the LSPs might need the detailed traffic engineering mechanism and knobs to be deployed, and at times only a few of these LSPs might need traffic engineering mechanisms (like link coloring, bandwidth control, etc.).

These few LSPs can be between the data centers and the LSPs, with the rest of the network solely dependent on CSPF EROs without any link coloring or bandwidth control, hence making the task easier and more convenient for an engineer, while adding a new PE router in the network.

Chapter 4 covers how to leverage the MPLS backbone with the help of VPNs.

Chapter 4

MPLS VPNs

VPN is the acronym for *Virtual Private Network*, a private network on a shared infrastructure built with the help of VRFs. Each user or site will have a separate VRF on a single router and work simultaneously with traffic in isolation from other VRFs. Each VRF creates a separate routing and forwarding table, which is independent of the other VRFs. This allows overlapping of subnets, which cater hosting services for multiple customers on a single PE router. Each customer is free to use any private IP address schema without giving a thought to overlapping IP address subnets with other users on the provider backbone. To differentiate each site and VRF IP address pool, route-distinguisher and route-targets are used.

VPN Routing and Forwarding

The routing and forwarding of VPN routes happens with the help of *route-distinguisher* and *route-target* and the upcoming sections cover them in detail. Moreover, `family inet-vpn unicast` needs to be enabled between the BGP neighbors, so VPN routes can be exchanged between them.

Route Distinguisher

While the routes are propagated to the remote end PE, the routes are differentiated with the help of the route distinguisher (RD) identifier used to create a VPNv4 prefix. Each VRF will have a unique RD in a MPLS domain that can be manually configured or dynamically generated by the originating PE.

The three types of RD format are shown in Figure 4.1.

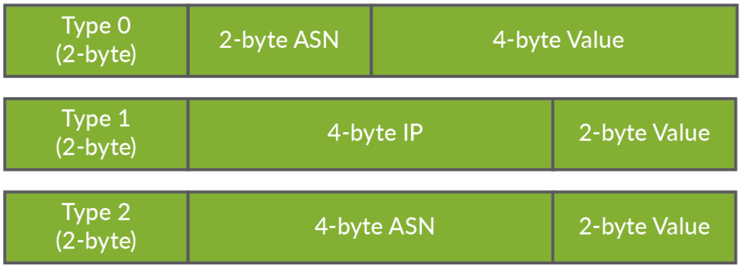

Figure 4.1 RD Format

The VPNv4 prefix is a VPN IP prefix with an RD prefix attached. It is 96 bits and is also called *RD prefix*.

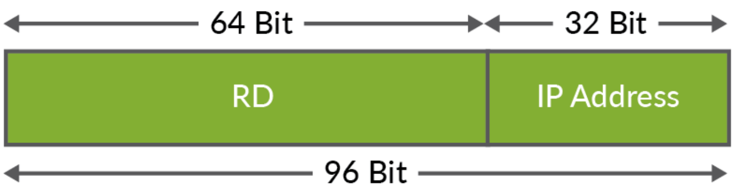

Figure 4.2 Bits of a VPNv4 Prefix or RD Prefix

In the next example, 100.64.11.0/24 is the VPN route where 192.0.2.7:4 is the RD:

```
lab@NE> show route table bgp.l3vpn.0

bgp.l3vpn.0: 10 destinations, 10 routes (10 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

192.0.2.7:4:100.64.11.0/24
                   *[BGP/170] 00:18:16, localpref 100, from 192.0.2.7
                      AS path: 64510 I
                    > to 198.51.100.7 via ge-0/0/3.0, Push 299856, Push 299824(top)
```

Route Target

The route target (RT) is a type of BGP extended community. RD and route target are two different concepts and serve two different roles, but RD and route target share a similar format. RD is used to differentiate a route in a MPLS domain, whereas route target is used to import and export routes between VRFs on single or multiple VRFs. This example shows the presence of route target in a `rd-prefix`:

```
lab@NE> show route table bgp.l3vpn.0 rd-prefix 192.0.2.7:4:100.64.11.0/24 extensive

bgp.l3vpn.0: 10 destinations, 10 routes (10 active, 0 holddown, 0 hidden)
192.0.2.7:4:100.64.11.0/24 (1 entry, 0 announced)
        *BGP    Preference: 170/-101
                Route Distinguisher: 192.0.2.7:4
                Next hop type: Indirect
                Address: 0x9351498
                Next-hop reference count: 15
                Source: 192.0.2.7
                Next hop type: Router, Next hop index: 610
                Next hop: 198.51.100.7 via ge-0/0/3.0, selected
                Label operation: Push 299856, Push 299824(top)
                Label TTL action: prop-ttl, prop-ttl(top)
                Protocol next hop: 192.0.2.7
                Push 299856
                Indirect next hop: 9475c18 262143
                State: <Active Int Ext>
                Local AS: 64500 Peer AS: 64500
                Age: 27:05      Metric2: 1
                Task: BGP_64500.192.0.2.7+179
                AS path: 64510 I
                Communities: target:64510:100
                Import Accepted
                VPN Label: 299856
                Localpref: 100
                Router ID: 192.0.2.7
                Secondary Tables: CUST_SITE2.inet.0
                Indirect next hops: 1
                        Protocol next hop: 192.0.2.7 Metric: 1
                        Push 299856
                        Indirect next hop: 9475c18 262143
                        Indirect path forwarding next hops: 1
                                Next hop type: Router
                                Next hop: 198.51.100.7 via ge-0/0/3.0
                        192.0.2.7/32 Originating RIB: inet.3
                          Metric: 1                       Node path count: 1
                          Forwarding nexthops: 1
                                Nexthop: 198.51.100.7 via ge-0/0/3.0
```

Importing and Exporting VPN Routes

There are two ways you can import and export routes in a VRF. One option is to configure the `vrf-target` knob, and the second is configuring import and export policies. The `vrf-target` knob automatically exports routes from VRFs adding a target community, and imports routes in the VRF with the same target community, allowing easy configuration of simple VPN topologies. The use of import and export policies can allow configuration of far more complex VPN topologies.

The next example configuration shows both ways to configure import and export.

Using `vrf-target`:

```
lab@DA# show routing-instances CUST_SITE1
instance-type vrf;
interface ge-0/0/9.200;
vrf-target target:64510:100;
protocols {
    bgp {
        group PE-CE {
            type external;
            peer-as 64510;
            neighbor 203.0.113.3;
        }
    }
}
```

Using import and export policies:

```
lab@NE> show configuration routing-instances CUST_SITE2
instance-type vrf;
interface ge-0/0/9.300;
vrf-import CUST-IMPORT;
vrf-export CUST-EXPORT;
protocols {
    bgp {
        group PE-CE {
            type external;
            peer-as 64509;
            neighbor 203.0.113.1;
        }
    }
}

lab@NE> show configuration policy-options community CUST-TARGET
members target:64510:100;

lab@NE> show configuration policy-options policy-statement CUST-IMPORT
term 1 {
    from community CUST-TARGET;
    then accept;
}
term 2 {
    then reject;
}

lab@NE> show configuration policy-options policy-statement CUST-EXPORT
term 1 {
    then {
        community add CUST-TARGET;
        accept;
    }
}
```

VPN Label

A PE router assigns a VPN label based on the next hop in the VRF and advertises that VPN label to the remote PE router. For example, Figure 4.3's topology there is a customer VRF on the DA router and also on the NE router. When the NE router sends packets to the customer VPN site on DA, it first pushes the VPN label associated with that next hop and then the LSP label.

By default, penultimate hop popping happens one hop before the egress router, which leaves the packet with only the VPN label when it arrives at the DA egress router. The DA router then identifies the concerned VRF with the help of the VPN label, and pops the VPN label before sending it to the destination.

To avoid use of multiple labels per VRF that depend on the next hops, one single label can be advertised per VRF with the help of the `vrf-table-label` command:

```
lab@DA# set routing-instances CUST_SITE1 vrf-table-label
```

NOTE The `vrf-table-label` command also aids the export of direct routes present in the VRF.

L3 VPN Control and Data Plane

This section covers the implementation of the topology shown in Figure 4.3. The LSP between the PE routers NE, DA, and SJ can be achieved either by RSVP LSPs or by LDP LSP.

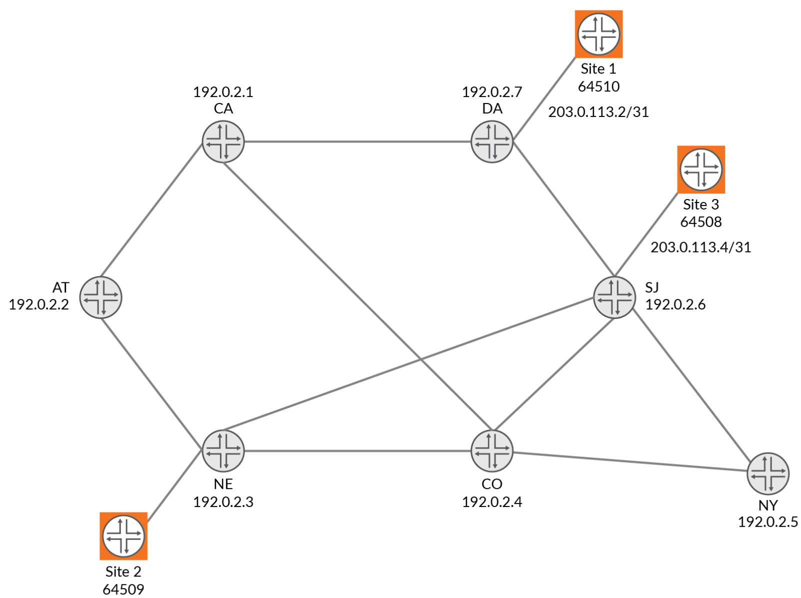

Figure 4.3 MPLS Layer 3 VPN

NOTE We already configured LSPs using the LDP protocol before starting the L3 VPN configuration.

Configuration

On DA for Site 1

```
lab@DA> show configuration protocols bgp
group MP-BGP {
    type internal;
    local-address 192.0.2.7;
    family inet {
        unicast;
    }
    family inet-vpn {
        unicast;
    }
    neighbor 192.0.2.3 {
        description NE;
    }
    neighbor 192.0.2.6 {
        description SJ;
    }
}
```

The `family inet-vpn unicast` enables exchange of VPN routes between the PEs, while `family inet unicast` also aids the exchange of IPv4 routes:

```
lab@DA> show configuration routing-options route-distinguisher-id
route-distinguisher-id 192.0.2.7;

lab@DA> show configuration routing-instances CUST_SITE1
instance-type vrf;
interface ge-0/0/9.200;
vrf-target target:64510:100;
protocols {
    bgp {
        group PE-CE {
            type external;
            peer-as 64510;
            neighbor 203.0.113.3;
        }
    }
}
```

You can see the `vrf-target` knob on the DA router that enables import and export of VPN routes. Moreover, it's using the dynamic method for generating the RD for VRFs.

On NE for Site 2

```
lab@NE> show configuration protocols bgp
group MP-BGP {
    type internal;
```

```
    local-address 192.0.2.3;
    family inet {
        unicast;
    }
    family inet-vpn {
        unicast;
    }
    neighbor 192.0.2.6 {
        description SJ;
    }
    neighbor 192.0.2.7 {
        description DA;
    }
}

lab@NE> show configuration routing-options route-distinguisher-id
route-distinguisher-id 192.0.2.3;

lab@NE> show configuration routing-instances CUST_SITE2
instance-type vrf;
interface ge-0/0/9.300;
vrf-import CUST-IMPORT;
vrf-export CUST-EXPORT;
protocols {
    bgp {
        group PE-CE {
            type external;
            peer-as 64509;
            neighbor 203.0.113.1;
        }
    }
}

lab@NE> show configuration policy-options policy-statement CUST-IMPORT
term 1 {
    from community CUST-TARGET;
    then accept;
}
term 2 {
    then reject;
}

lab@NE> show configuration policy-options policy-statement CUST-EXPORT
term 1 {
    then {
        community add CUST-TARGET;
        accept;
    }
}

lab@NE> show configuration policy-options community CUST-TARGET
members target:64510:100;
```

The configuration on the NE router uses VRF import and export policies to assist route advertisement and installation.

On SJ for Site 3

```
lab@SJ> show configuration protocols bgp
group MP-BGP {
    type internal;
    local-address 192.0.2.6;
    family inet {
        unicast;
    }
    family inet-vpn {
        unicast;
    }
    neighbor 192.0.2.3 {
        description NE;
    }
    neighbor 192.0.2.7 {
        description DA;
    }
}

lab@SJ> show configuration routing-instances CUST_SITE3
instance-type vrf;
interface ge-0/0/9.400;
route-distinguisher 65000:65000;
vrf-target target:64510:100;
protocols {
    bgp {
        group PE-CE {
            type external;
            peer-as 64508;
            neighbor 203.0.113.5;
        }
    }
}
```

Verification

Let's verify multi-procotol BGP (MP-BGP) and the customer's BGP neighborship.

On DA PE

```
lab@DA> show bgp summary
Groups: 2 Peers: 3 Down peers: 0
Table          Tot Paths  Act Paths Suppressed    History Damp State    Pending
inet.0                0          0          0          0          0          0
bgp.l3vpn.0          10         10          0          0          0          0
Peer                     AS      InPkt     OutPkt    OutQ   Flaps Last Up/Dwn State|#Active/
Received/Accepted/Damped...
192.0.2.3             64500        631        636       0       0     4:44:51 Establ
  inet.0: 0/0/0/0
  bgp.l3vpn.0: 5/5/5/0
  CUST_SITE1.inet.0: 5/5/5/0
192.0.2.6             64500        635        634       0       0     4:44:47 Establ
```

```
  inet.0: 0/0/0/0
  bgp.l3vpn.0: 5/5/5/0
  CUST_SITE1.inet.0: 5/5/5/0
203.0.113.3           64510        125        126       0       0       54:22 Establ
  CUST_SITE1.inet.0: 4/4/4/0
```

You can see that the MP-BGP neighborships are up and exchanging L3VPN routes with other PEs. The customer peer 203.0.113.3 is up and is advertising four routes to the PE router.

On NE PE

```
lab@NE> show bgp summary
Groups: 2 Peers: 3 Down peers: 0
Table          Tot Paths  Act Paths Suppressed    History Damp State    Pending
inet.0                 0          0          0          0          0          0
bgp.l3vpn.0           10         10          0          0          0          0
Peer                     AS      InPkt     OutPkt    OutQ   Flaps Last Up/Dwn State|#Active/
Received/Accepted/Damped...
192.0.2.6             64500        697        692       0       0     5:11:29 Establ
  CUST_SITE2.inet.0: 5/5/5/0
  inet.0: 0/0/0/0
  bgp.l3vpn.0: 5/5/5/0
192.0.2.7             64500        650        645       0       0     4:50:40 Establ
  CUST_SITE2.inet.0: 5/5/5/0
  inet.0: 0/0/0/0
  bgp.l3vpn.0: 5/5/5/0
203.0.113.1           64509        138        140       0       0     1:00:25 Establ
  CUST_SITE2.inet.0: 4/4/4/0
```

On SJ PE

```
lab@SJ> show bgp summary
Groups: 2 Peers: 3 Down peers: 0
Table          Tot Paths  Act Paths Suppressed    History Damp State    Pending
inet.0                 0          0          0          0          0          0
bgp.l3vpn.0           10         10          0          0          0          0
Peer                     AS      InPkt     OutPkt    OutQ   Flaps Last Up/Dwn State|#Active/
Received/Accepted/Damped...
192.0.2.3             64500        742        747       0       0     5:34:17 Establ
  inet.0: 0/0/0/0
  bgp.l3vpn.0: 5/5/5/0
  CUST_SITE3.inet.0: 5/5/5/0
192.0.2.7             64500        699        699       0       0     5:13:25 Establ
  inet.0: 0/0/0/0
  bgp.l3vpn.0: 5/5/5/0
  CUST_SITE3.inet.0: 5/5/5/0
203.0.113.5           64508        188        190       0       0     1:22:56 Establ
  CUST_SITE3.inet.0: 4/4/4/0
```

Verifying Customer Routes

On DA Router

```
lab@DA> show route receive-protocol bgp 203.0.113.3

inet.0: 22 destinations, 23 routes (22 active, 0 holddown, 0 hidden)

inet.3: 6 destinations, 6 routes (6 active, 0 holddown, 0 hidden)

CUST_SITE1.inet.0: 16 destinations, 16 routes (16 active, 0 holddown, 0 hidden)
  Prefix                  Nexthop              MED     Lclpref    AS path
* 100.64.11.0/24          203.0.113.3                             64510 I
* 100.64.12.0/24          203.0.113.3                             64510 I
* 100.64.13.0/24          203.0.113.3                             64510 I
* 100.64.14.0/24          203.0.113.3                             64510 I

iso.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)

mpls.0: 14 destinations, 14 routes (14 active, 0 holddown, 0 hidden)

bgp.l3vpn.0: 10 destinations, 10 routes (10 active, 0 holddown, 0 hidden)
```

On NE Router

```
lab@NE> show route receive-protocol bgp 203.0.113.1

inet.0: 23 destinations, 24 routes (23 active, 0 holddown, 0 hidden)

inet.3: 6 destinations, 6 routes (6 active, 0 holddown, 0 hidden)

CUST_SITE2.inet.0: 16 destinations, 16 routes (16 active, 0 holddown, 0 hidden)
  Prefix                  Nexthop              MED     Lclpref    AS path
* 100.64.21.0/24          203.0.113.1                             64509 I
* 100.64.22.0/24          203.0.113.1                             64509 I
* 100.64.23.0/24          203.0.113.1                             64509 I
* 100.64.24.0/24          203.0.113.1                             64509 I

iso.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)

mpls.0: 21 destinations, 21 routes (21 active, 0 holddown, 0 hidden)

bgp.l3vpn.0: 10 destinations, 10 routes (10 active, 0 holddown, 0 hidden)
```

On SJ Router

```
lab@SJ> show route receive-protocol bgp 203.0.113.5

inet.0: 24 destinations, 25 routes (24 active, 0 holddown, 0 hidden)

inet.3: 6 destinations, 7 routes (6 active, 0 holddown, 0 hidden)

CUST_SITE3.inet.0: 16 destinations, 16 routes (16 active, 0 holddown, 0 hidden)
  Prefix                  Nexthop              MED     Lclpref    AS path
* 100.64.31.0/24          203.0.113.5                             64508 I
```

```
* 100.64.32.0/24         203.0.113.5                              64508 I
* 100.64.33.0/24         203.0.113.5                              64508 I
* 100.64.34.0/24         203.0.113.5                              64508 I

iso.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)

mpls.0: 16 destinations, 16 routes (16 active, 0 holddown, 0 hidden)

bgp.l3vpn.0: 10 destinations, 10 routes (10 active, 0 holddown, 0 hidden)
```

Verifying MP-BGP Learned Routes

```
lab@DA> show route table bgp.l3vpn.0 terse

bgp.l3vpn.0: 10 destinations, 10 routes (10 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

A Destination        P Prf   Metric 1   Metric 2  Next hop         AS path
  192.0.2.3:5:100.64.21.0/24
*                    B 170       100             >198.51.100.16    64509 I
  192.0.2.3:5:100.64.22.0/24
*                    B 170       100             >198.51.100.16    64509 I
  192.0.2.3:5:100.64.23.0/24
*                    B 170       100             >198.51.100.16    64509 I
  192.0.2.3:5:100.64.24.0/24
*                    B 170       100             >198.51.100.16    64509 I
  192.0.2.3:5:203.0.113.0/31
*                    B 170       100             >198.51.100.16    I
  65000:65000:100.64.31.0/24
*                    B 170       100             >198.51.100.16    64508 I
  65000:65000:100.64.32.0/24
*                    B 170       100             >198.51.100.16    64508 I
  65000:65000:100.64.33.0/24
*                    B 170       100             >198.51.100.16    64508 I
  65000:65000:100.64.34.0/24
*                    B 170       100             >198.51.100.16    64508 I
  65000:65000:203.0.113.4/31
*                    B 170       100             >198.51.100.16    I
```

As seen here, L3VPN routes from other MP-BGP peers are being learned and installed in the `bgp.l3vpn.0 table`. The `bgp.l3vpn.0` table hosts all of the L3VPN routes learned from other MP-BGP peers. The routes are differentiated with the help of the RD. After getting installed in the `bgp.l3vpn.0` table, the VPNv4 routes learned from other MP-BGP peers will then be installed in the respective customer VRF routing table with the help of `route-target`.

```
lab@DA> show route table bgp.l3vpn.0 rd-prefix 192.0.2.3:5:100.64.21.0/24 extensive

bgp.l3vpn.0: 10 destinations, 10 routes (10 active, 0 holddown, 0 hidden)
192.0.2.3:5:100.64.21.0/24 (1 entry, 0 announced)
        *BGP    Preference: 170/-101
                Route Distinguisher: 192.0.2.3:5
                Next hop type: Indirect
                Address: 0x9351528
                Next-hop reference count: 15
```

```
                Source: 192.0.2.3
                Next hop type: Router, Next hop index: 599
                Next hop: 198.51.100.16 via ge-0/0/1.0, selected
                Label operation: Push 300032, Push 299856(top)
                Label TTL action: prop-ttl, prop-ttl(top)
                Protocol next hop: 192.0.2.3
                Push 300032
                Indirect next hop: 9445ed0 262142
                State: <Active Int Ext>
                Local AS: 64500 Peer AS: 64500
                Age: 1:54:30    Metric2: 1
                Task: BGP_64500.192.0.2.3+179
                AS path: 64509 I
                Communities: target:64510:100
                Import Accepted
                VPN Label: 300032
                Localpref: 100
                Router ID: 192.0.2.3
                Secondary Tables: CUST_SITE1.inet.0
                Indirect next hops: 1
                        Protocol next hop: 192.0.2.3 Metric: 1
                        Push 300032
                        Indirect next hop: 9445ed0 262142
                        Indirect path forwarding next hops: 1
                                Next hop type: Router
                                Next hop: 198.51.100.16 via ge-0/0/1.0
                        192.0.2.3/32 Originating RIB: inet.3
                          Metric: 1                       Node path count: 1
                          Forwarding nexthops: 1
                                Nexthop: 198.51.100.16 via ge-0/0/1.0

lab@DA> show route table bgp.l3vpn.0 rd-prefix 65000:65000:100.64.32.0/24 extensive

bgp.l3vpn.0: 10 destinations, 10 routes (10 active, 0 holddown, 0 hidden)
65000:65000:100.64.32.0/24 (1 entry, 0 announced)
        *BGP    Preference: 170/-101
                Route Distinguisher: 65000:65000
                Next hop type: Indirect
                Address: 0x9351600
                Next-hop reference count: 15
                Source: 192.0.2.6
                Next hop type: Router, Next hop index: 600
                Next hop: 198.51.100.16 via ge-0/0/1.0, selected
                Label operation: Push 299888
                Label TTL action: prop-ttl
                Protocol next hop: 192.0.2.6
                Push 299888
                Indirect next hop: 9445fb8 262143
                State: <Active Int Ext>
                Local AS: 64500 Peer AS: 64500
                Age: 1:55:28    Metric2: 1
                Task: BGP_64500.192.0.2.6+179
                AS path: 64508 I
                Communities: target:64510:100
                Import Accepted
                VPN Label: 299888
                Localpref: 100
                Router ID: 192.0.2.6
```

```
Secondary Tables: CUST_SITE1.inet.0
Indirect next hops: 1
        Protocol next hop: 192.0.2.6 Metric: 1
        Push 299888
        Indirect next hop: 9445fb8 262143
        Indirect path forwarding next hops: 1
                Next hop type: Router
                Next hop: 198.51.100.16 via ge-0/0/1.0
        192.0.2.6/32 Originating RIB: inet.3
          Metric: 1                     Node path count: 1
          Forwarding nexthops: 1
                Nexthop: 198.51.100.16 via ge-0/0/1.0
```

Both NE and SJ are sending routes with CUST_SITE1 `route-target` attached of CUST_SITE1, hence this output indicates "`Secondary Tables`" as "`CUST_SITE1.inet.0`".

The next output shows the `CUST_SITE1.inet.0` table has routes from customer site 2 and site 3:

```
lab@DA> show route table CUST_SITE1.inet.0

CUST_SITE1.inet.0: 16 destinations, 16 routes (16 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

100.64.11.0/24     *[BGP/170] 02:01:39, localpref 100
                      AS path: 64510 I
                    > to 203.0.113.3 via ge-0/0/9.200
100.64.12.0/24     *[BGP/170] 02:01:39, localpref 100
                      AS path: 64510 I
                    > to 203.0.113.3 via ge-0/0/9.200
100.64.13.0/24     *[BGP/170] 02:01:39, localpref 100
                      AS path: 64510 I
                    > to 203.0.113.3 via ge-0/0/9.200
100.64.14.0/24     *[BGP/170] 02:01:39, localpref 100
                      AS path: 64510 I
                    > to 203.0.113.3 via ge-0/0/9.200
100.64.21.0/24     *[BGP/170] 02:01:39, localpref 100, from 192.0.2.3
                      AS path: 64509 I
                    > to 198.51.100.16 via ge-0/0/1.0, Push 300032, Push 299856(top)
100.64.22.0/24     *[BGP/170] 02:01:39, localpref 100, from 192.0.2.3
                      AS path: 64509 I
                    > to 198.51.100.16 via ge-0/0/1.0, Push 300032, Push 299856(top)
100.64.23.0/24     *[BGP/170] 02:01:39, localpref 100, from 192.0.2.3
                      AS path: 64509 I
                    > to 198.51.100.16 via ge-0/0/1.0, Push 300032, Push 299856(top)
100.64.24.0/24     *[BGP/170] 02:01:39, localpref 100, from 192.0.2.3
                      AS path: 64509 I
                    > to 198.51.100.16 via ge-0/0/1.0, Push 300032, Push 299856(top)
100.64.31.0/24     *[BGP/170] 02:01:39, localpref 100, from 192.0.2.6
                      AS path: 64508 I
                    > to 198.51.100.16 via ge-0/0/1.0, Push 299888
100.64.32.0/24     *[BGP/170] 02:01:39, localpref 100, from 192.0.2.6
                      AS path: 64508 I
                    > to 198.51.100.16 via ge-0/0/1.0, Push 299888
100.64.33.0/24     *[BGP/170] 02:01:39, localpref 100, from 192.0.2.6
                      AS path: 64508 I
```

```
                      > to 198.51.100.16 via ge-0/0/1.0, Push 299888
100.64.34.0/24     *[BGP/170] 02:01:39, localpref 100, from 192.0.2.6
                      AS path: 64508 I
                    > to 198.51.100.16 via ge-0/0/1.0, Push 299888
203.0.113.0/31     *[BGP/170] 02:01:39, localpref 100, from 192.0.2.3
                      AS path: I
                    > to 198.51.100.16 via ge-0/0/1.0, Push 300032, Push 299856(top)
203.0.113.2/31     *[Direct/0] 05:57:30
                    > via ge-0/0/9.200
203.0.113.2/32     *[Local/0] 05:57:30
                      Local via ge-0/0/9.200
203.0.113.4/31     *[BGP/170] 02:01:39, localpref 100, from 192.0.2.6
                      AS path: I
                    > to 198.51.100.16 via ge-0/0/1.0, Push 299888
```

Similarly, the routing tables for Site 2 and Site 3 on NE and SJ have routes for all the sites.

On NE Router

```
lab@NE> show route table CUST_SITE2.inet.0

CUST_SITE2.inet.0: 16 destinations, 16 routes (16 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

100.64.11.0/24     *[BGP/170] 02:06:23, localpref 100, from 192.0.2.7
                      AS path: 64510 I
                    > to 198.51.100.7 via ge-0/0/3.0, Push 299936, Push 299776(top)
100.64.12.0/24     *[BGP/170] 02:06:23, localpref 100, from 192.0.2.7
                      AS path: 64510 I
                    > to 198.51.100.7 via ge-0/0/3.0, Push 299936, Push 299776(top)
100.64.13.0/24     *[BGP/170] 02:06:23, localpref 100, from 192.0.2.7
                      AS path: 64510 I
                    > to 198.51.100.7 via ge-0/0/3.0, Push 299936, Push 299776(top)
100.64.14.0/24     *[BGP/170] 02:06:23, localpref 100, from 192.0.2.7
                      AS path: 64510 I
                    > to 198.51.100.7 via ge-0/0/3.0, Push 299936, Push 299776(top)
100.64.21.0/24     *[BGP/170] 02:06:23, localpref 100
                      AS path: 64509 I
                    > to 203.0.113.1 via ge-0/0/9.300
100.64.22.0/24     *[BGP/170] 02:06:23, localpref 100
                      AS path: 64509 I
                    > to 203.0.113.1 via ge-0/0/9.300
100.64.23.0/24     *[BGP/170] 02:06:23, localpref 100
                      AS path: 64509 I
                    > to 203.0.113.1 via ge-0/0/9.300
100.64.24.0/24     *[BGP/170] 02:06:23, localpref 100
                      AS path: 64509 I
                    > to 203.0.113.1 via ge-0/0/9.300
100.64.31.0/24     *[BGP/170] 02:06:23, localpref 100, from 192.0.2.6
                      AS path: 64508 I
                    > to 198.51.100.7 via ge-0/0/3.0, Push 299888
100.64.32.0/24     *[BGP/170] 02:06:23, localpref 100, from 192.0.2.6
                      AS path: 64508 I
                    > to 198.51.100.7 via ge-0/0/3.0, Push 299888
100.64.33.0/24     *[BGP/170] 02:06:23, localpref 100, from 192.0.2.6
                      AS path: 64508 I
```

```
                    > to 198.51.100.7 via ge-0/0/3.0, Push 299888
100.64.34.0/24     *[BGP/170] 02:06:23, localpref 100, from 192.0.2.6
                      AS path: 64508 I
                    > to 198.51.100.7 via ge-0/0/3.0, Push 299888
203.0.113.0/31     *[Direct/0] 06:27:50
                    > via ge-0/0/9.300
203.0.113.0/32     *[Local/0] 06:27:50
                      Local via ge-0/0/9.300
203.0.113.2/31     *[BGP/170] 02:06:23, localpref 100, from 192.0.2.7
                      AS path: I
                    > to 198.51.100.7 via ge-0/0/3.0, Push 299936, Push 299776(top)
203.0.113.4/31     *[BGP/170] 02:06:23, localpref 100, from 192.0.2.6
                      AS path: I
                    > to 198.51.100.7 via ge-0/0/3.0, Push 299888
```

On SJ Router

```
lab@SJ> show route table CUST_SITE3.inet.0

CUST_SITE3.inet.0: 16 destinations, 16 routes (16 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

100.64.11.0/24     *[BGP/170] 02:07:21, localpref 100, from 192.0.2.7
                      AS path: 64510 I
                    > to 198.51.100.17 via ge-0/0/1.0, Push 299936
100.64.12.0/24     *[BGP/170] 02:07:21, localpref 100, from 192.0.2.7
                      AS path: 64510 I
                    > to 198.51.100.17 via ge-0/0/1.0, Push 299936
100.64.13.0/24     *[BGP/170] 02:07:21, localpref 100, from 192.0.2.7
                      AS path: 64510 I
                    > to 198.51.100.17 via ge-0/0/1.0, Push 299936
100.64.14.0/24     *[BGP/170] 02:07:21, localpref 100, from 192.0.2.7
                      AS path: 64510 I
                    > to 198.51.100.17 via ge-0/0/1.0, Push 299936
100.64.21.0/24     *[BGP/170] 02:07:21, localpref 100, from 192.0.2.3
                      AS path: 64509 I
                    > to 198.51.100.6 via ge-0/0/3.0, Push 300032
100.64.22.0/24     *[BGP/170] 02:07:21, localpref 100, from 192.0.2.3
                      AS path: 64509 I
                    > to 198.51.100.6 via ge-0/0/3.0, Push 300032
100.64.23.0/24     *[BGP/170] 02:07:21, localpref 100, from 192.0.2.3
                      AS path: 64509 I
                    > to 198.51.100.6 via ge-0/0/3.0, Push 300032
100.64.24.0/24     *[BGP/170] 02:07:21, localpref 100, from 192.0.2.3
                      AS path: 64509 I
                    > to 198.51.100.6 via ge-0/0/3.0, Push 300032
100.64.31.0/24     *[BGP/170] 02:07:21, localpref 100
                      AS path: 64508 I
                    > to 203.0.113.5 via ge-0/0/9.400
100.64.32.0/24     *[BGP/170] 02:07:21, localpref 100
                      AS path: 64508 I
                    > to 203.0.113.5 via ge-0/0/9.400
100.64.33.0/24     *[BGP/170] 02:07:21, localpref 100
                      AS path: 64508 I
                    > to 203.0.113.5 via ge-0/0/9.400
100.64.34.0/24     *[BGP/170] 02:07:21, localpref 100
                      AS path: 64508 I
```

```
                    > to 203.0.113.5 via ge-0/0/9.400
203.0.113.0/31     *[BGP/170] 02:07:21, localpref 100, from 192.0.2.3
                      AS path: I
                    > to 198.51.100.6 via ge-0/0/3.0, Push 300032
203.0.113.2/31     *[BGP/170] 02:07:21, localpref 100, from 192.0.2.7
                      AS path: I
                    > to 198.51.100.17 via ge-0/0/1.0, Push 299936
203.0.113.4/31     *[Direct/0] 05:21:32
                    > via ge-0/0/9.400
203.0.113.4/32     *[Local/0] 06:23:54
                      Local via ge-0/0/9.400
```

Let's verify the advertisement of routes to the customer's router.

For Site 1

```
lab@DA> show route advertising-protocol bgp 203.0.113.3

CUST_SITE1.inet.0: 16 destinations, 16 routes (16 active, 0 holddown, 0 hidden)
  Prefix                  Nexthop              MED     Lclpref    AS path
* 100.64.21.0/24          Self                                    64509 I
* 100.64.22.0/24          Self                                    64509 I
* 100.64.23.0/24          Self                                    64509 I
* 100.64.24.0/24          Self                                    64509 I
* 100.64.31.0/24          Self                                    64508 I
* 100.64.32.0/24          Self                                    64508 I
* 100.64.33.0/24          Self                                    64508 I
* 100.64.34.0/24          Self                                    64508 I
* 203.0.113.0/31          Self                                    I
* 203.0.113.4/31          Self                                    I
```

For Site 2

```
lab@NE> show route advertising-protocol bgp 203.0.113.1

CUST_SITE2.inet.0: 16 destinations, 16 routes (16 active, 0 holddown, 0 hidden)
  Prefix                  Nexthop              MED     Lclpref    AS path
* 100.64.11.0/24          Self                                    64510 I
* 100.64.12.0/24          Self                                    64510 I
* 100.64.13.0/24          Self                                    64510 I
* 100.64.14.0/24          Self                                    64510 I
* 100.64.31.0/24          Self                                    64508 I
* 100.64.32.0/24          Self                                    64508 I
* 100.64.33.0/24          Self                                    64508 I
* 100.64.34.0/24          Self                                    64508 I
* 203.0.113.2/31          Self                                    I
* 203.0.113.4/31          Self                                    I
```

For Site 3

```
lab@SJ> show route advertising-protocol bgp 203.0.113.5

CUST_SITE3.inet.0: 16 destinations, 16 routes (16 active, 0 holddown, 0 hidden)
  Prefix                  Nexthop              MED     Lclpref    AS path
* 100.64.11.0/24          Self                                    64510 I
```

```
* 100.64.12.0/24         Self                                    64510 I
* 100.64.13.0/24         Self                                    64510 I
* 100.64.14.0/24         Self                                    64510 I
* 100.64.21.0/24         Self                                    64509 I
* 100.64.22.0/24         Self                                    64509 I
* 100.64.23.0/24         Self                                    64509 I
* 100.64.24.0/24         Self                                    64509 I
* 203.0.113.0/31         Self                                    I
* 203.0.113.2/31         Self                                    I
```

And let's verify reachability between customer sites.

Between Site 1 and Site 2

```
lab@Site1> ping 100.64.21.1 rapid count 10
PING 100.64.21.1 (100.64.21.1): 56 data bytes
!!!!!!!!!!
--- 100.64.21.1 ping statistics ---
10 packets transmitted, 10 packets received, 0% packet loss
round-trip min/avg/max/stddev = 19.877/27.432/30.888/3.377 ms
```

Between Site 1 and Site 3

```
lab@Site1> ping 100.64.31.1 rapid count 10
PING 100.64.31.1 (100.64.31.1): 56 data bytes
!!!!!!!!!!
--- 100.64.31.1 ping statistics ---
10 packets transmitted, 10 packets received, 0% packet loss
round-trip min/avg/max/stddev = 15.262/21.361/25.063/2.990 ms
```

Between Site 2 and Site 3

```
lab@Site2> ping 100.64.32.1 rapid count 10
PING 100.64.32.1 (100.64.32.1): 56 data bytes
!!!!!!!!!!
--- 100.64.32.1 ping statistics ---
10 packets transmitted, 10 packets received, 0% packet loss
round-trip min/avg/max/stddev = 19.053/20.367/23.931/1.264 ms
```

For traffic to reach the destination from the customer site router, it first goes to the PE router, and the PE router looks up the concerned VRF table and follows the LSP to reach *protocol next hop* or the egress router. Once it reaches the egress router, the egress router looks up the VPN label to figure out the associated FEC, and sends the packet to the concerned destination, popping the VPN label. If `vrf-table-label` is being used, then the egress router pops the VPN label and sends the packet to the VRF, and once the packet is in the VRF it is then sent out to the customer site by referring to the VRF routing table. The same approach follows when traffic traverses the path in the other direction.

VPLS

There are cases when you need a Layer 2 service that is point-to-multipoint or a LAN segment across multiple sites. Virtual private LAN service (VPLS) is a service that can cater to that requirement – appearing to be a single LAN segment, like a learning bridge. When the destination MAC address is known traffic is sent directly to the site, and when it's unknown an Ethernet frame is broadcast to all the remote sites. VPLS can be configured using BGP or LDP, as discussed in the next walk through.

VPLS Using BGP

The VPLS control plane is used for auto discovery and signaling. BGP allows auto discovery and signaling between all the PE devices that participate in a particular VPLS instance and all the PEs participating must be logically in full mesh.

After auto discovery, each pair of PEs participating in a particular VPLS instance creates a bidirectional pseudowire, which is achieved by signaling. Within the signaling, the PE router sends pseudowire characteristics and demultiplexors in a single BGP update message to all the remote PEs participating in that VPLS instance, which aids in building the pseudowires. *Bidirectional* pseudowires benefit by sending packets directly from an ingress PE to an egress PE without any need for an intermediate PE.

Network layer reachability information (NLRI) exchanges label blocks between PEs, which are used to compute the label for a particular site. The address family identifier (AFI) of VPLS BGP NLRI is AFI (25) and subsequent AFI is SAFI (65). The label block requires the site ID, the label base, and the label base offset. Each label block is carried in a separate NLRI update.

The transmit (TX) and receive (RX) labels are computed using label blocks (local and learned) through the following method:

TX label = Remote base + local site ID - remote offset

RX label = Local base + remote site ID - local offset

Layer 2 information with the extended community is used to signal control information about pseudowires, which is used to set up pseudowires for a particular VPLS. The L2 information includes community type (0x800A), encapsulation type (set to VPLS), control flags, MTU, and the preference to be used on pseudowires. The MTU values should be made the same for the local CE interface for all PE routers, since the L2 pseudowire doesn't support fragmentation.

The VPN routing and forwarding tables are created for each VPLS instance in respective PEs (see Figure 4.4). MAC learning is performed in the data plane. For each remote site a pseudowire is dynamically created and a receive label for each remote site is learned for that pseudowire.

The information in the VRF table for the local CE device is:

- Local site ID
- Logical interface connected to the local CE device
- Layer 2 encapsulation for the site
- The label used to forward traffic to one of the logical interfaces

The VRF routing and forwarding table also entails information received from remote PEs, like:

- Remote site ID
- Layer 2 encapsulation
- Label base

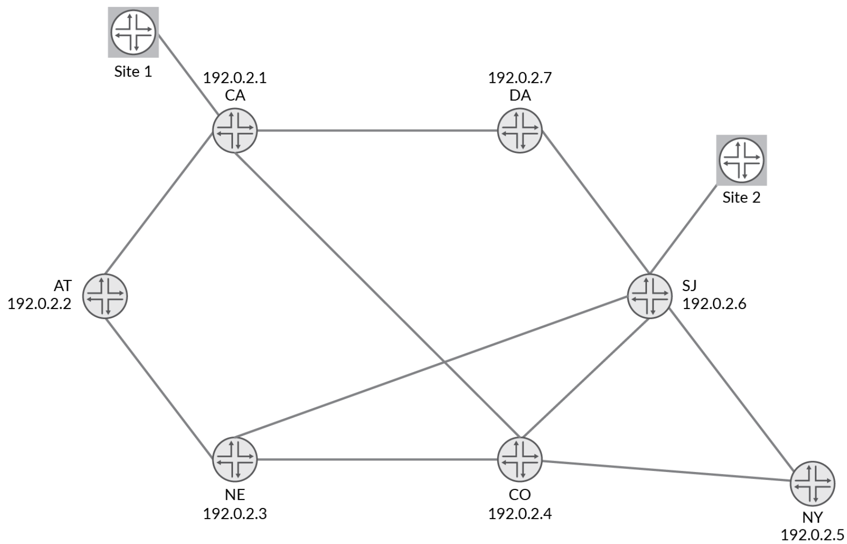

Figure 4.4 *MPLS VPLS using BGP*

Configuration of VPLS Using BGP

In Figure 4.4, the LSPs are already in place between PEs. You can use RSVP or LDP to bring up LSPs. Site 1, Site 2, and Site 3 are VPLS customer sites.

To configure, enable `family l2vpn` with MP-BGP peers to enable VPLS auto discovery and signaling between the VPLS PEs:

```
lab@CA> show configuration protocols bgp
group MP-BGP {
    type internal;
    local-address 192.0.2.1;
    family inet {
        unicast;
    }
    family inet-vpn {
        unicast;
    }
    family l2vpn {
        signaling;
    }
    neighbor 192.0.2.2 {
        description AT;
    }
    neighbor 192.0.2.3 {
        description NE;
    }
    neighbor 192.0.2.4 {
        description CO;
    }
    neighbor 192.0.2.5 {
        description NY;
    }
    neighbor 192.0.2.6 {
        description SJ;
    }
    neighbor 192.0.2.7 {
        description DA;
    }
}
lab@SJ> show configuration protocols bgp
group MP-BGP {
    type internal;
    local-address 192.0.2.6;
    family inet {
        unicast;
    }
    family inet-vpn {
        unicast;
    }
    family l2vpn {
        signaling;
    }
    neighbor 192.0.2.1 {
        description CA;
    }
    neighbor 192.0.2.2 {
        description AT;
    }
    neighbor 192.0.2.3 {
        description NE;
    }
    neighbor 192.0.2.4 {
        description CO;
    }
    neighbor 192.0.2.5 {
```

```
        description NY;
    }
    neighbor 192.0.2.7 {
        description DA;
    }
}
```

Interface Configuration

The VPLS configuration requires `encapsulation ethernet-vpls` and `family VPLS`.

On CA PE

```
lab@CA> show configuration interfaces ge-0/0/7
description CUST_VPLS_SITE1;
encapsulation ethernet-vpls;
unit 0 {
    family vpls;
}
```

On SJ PE

```
lab@SJ> show configuration interfaces ge-0/0/8
description CUST_VPLS_SITE2;
encapsulation ethernet-vpls;
unit 0 {
    family vpls;
}
```

In some topologies, you need to have the VPLS service enabled only on a single unit of a physical interface; in such cases this sample configuration can help:

```
lab@SJ# show interfaces ge-0/0/6
vlan-tagging;
encapsulation flexible-ethernet-services;    >>> This helps in assigning per-
unit Ethernet encapsulation configuration
unit 251 {
    encapsulation vlan-vpls;
    vlan-id 251;
    family vpls;
}
unit 500 {
    vlan-id 500;
    family inet {
        address 100.64.5.1/24;
    }
}
```

Configuring the VPLS Routing-instance On Both PEs

On CA PE

```
lab@CA> show configuration routing-instances CUST_SITE1
instance-type vpls;
interface ge-0/0/7.0;
```

```
vrf-target target:65000:111;
protocols {
    vpls {
        no-tunnel-services;
        site 1 {
            site-identifier 1;
        }
    }
}
```

On SJ PE

```
lab@SJ> show configuration routing-instances
CUST_SITE2 {
    instance-type vpls;
    interface ge-0/0/8.0;
    vrf-target target:65000:111;
    protocols {
        vpls {
            site-range 8;
            no-tunnel-services;
            site 2 {
                site-identifier 2;
            }
        }
    }
}
```

A few explanations:

- The site identifier needs to be unique number.
- The site range (optional) - Signifies the number of CE devices (remote device) it can connect (by default it's also 8).
- The `no-tunnel-services` knob is needed if the chassis isn't enabled with `tunnel-services` in the global configuration. If you enable `tunnel-services` in the global with the following command, you can skip configuring this knob:

```
lab@SJ# set chassis fpc 0 pic 0 tunnel-services bandwidth 1g
```

Verification

Let's verify the MP-BGP peering.

On CA PE

```
lab@CA> show bgp summary | find 192.0.2.6
192.0.2.6             64500       1916       1928       0       1    14:37:31 Establ
  inet.0: 0/0/0/0
  bgp.l3vpn.0: 0/0/0/0
  bgp.l2vpn.0: 1/1/1/0
  CUST_SITE1.l2vpn.0: 1/1/1/0
```

On SJ PE

```
lab@SJ> show bgp summary | find 192.0.2.5
192.0.2.5             64500       4619       4655       0       0 1d 11:41:29 Establ
  inet.0: 0/0/0/0
  bgp.l3vpn.0: 0/0/0/0
  bgp.l2vpn.0: 1/1/1/0
  CUST_SITE2.l2vpn.0: 1/1/1/0
```

You can see that the `bgp.l2vpn.0` routes are being learned on each PE and then installed in the respective customer VRF table. Now let's verify the global `bgp.l2vpn.0` table.

On CA PE

```
lab@CA> show route table bgp.l2vpn.0

bgp.l2vpn.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

192.0.2.6:5:2:1/96
                   *[BGP/170] 00:12:37, localpref 100, from 192.0.2.6
                      AS path: I
                    > to 198.51.100.19 via ge-0/0/2.0, label-switched-path CA_SJ
                      to 198.51.100.1 via ge-0/0/1.0, label-switched-path Bypass->198.51.100.19
```

On SJ PE

```
lab@SJ> show route table bgp.l2vpn.0

bgp.l2vpn.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

192.0.2.1:7:1:1/96
                   *[BGP/170] 00:12:01, localpref 100, from 192.0.2.1
                      AS path: I
                    > to 198.51.100.17 via ge-0/0/1.0, label-switched-path SJ_CA
                      to 198.51.100.10 via ge-0/0/2.0, label-switched-path Bypass->198.51.100.17
```

And now verify the VRF l2vpn table.

On CA Router

```
lab@CA> show route table CUST_SITE1.l2vpn.0

CUST_SITE1.l2vpn.0: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

192.0.2.1:7:1:1/96
                   *[L2VPN/170/-101] 00:13:50, metric2 1
                      Indirect
192.0.2.6:5:2:1/96
                   *[BGP/170] 00:13:52, localpref 100, from 192.0.2.6
                      AS path: I
                    > to 198.51.100.19 via ge-0/0/2.0, label-switched-path CA_SJ
                      to 198.51.100.1 via ge-0/0/1.0, label-switched-path Bypass->198.51.100.19
```

On SJ Router

```
lab@SJ> show route table CUST_SITE2.l2vpn.0

CUST_SITE2.l2vpn.0: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

192.0.2.1:7:1:1/96
                   *[BGP/170] 00:14:11, localpref 100, from 192.0.2.1
                      AS path: I
                    > to 198.51.100.17 via ge-0/0/1.0, label-switched-path SJ_CA
                      to 198.51.100.10 via ge-0/0/2.0, label-switched-path Bypass->198.51.100.17
192.0.2.6:5:2:1/96
                   *[L2VPN/170/-101] 1d 11:54:59, metric2 1
                      Indirect
```

The L2VPN route format is “RD:Site ID:Offset”.

Verifying VPLS Connections

On CA Router

```
lab@CA> show vpls connections
Layer-2 VPN connections:

Legend for connection status (St)
EI -- encapsulation invalid      NC -- interface encapsulation not CCC/TCC/VPLS
EM -- encapsulation mismatch     WE -- interface and instance encaps not same
VC-Dn -- Virtual circuit down    NP -- interface hardware not present
CM -- control-word mismatch      -> -- only outbound connection is up
CN -- circuit not provisioned    <- -- only inbound connection is up
OR -- out of range               Up -- operational
OL -- no outgoing label          Dn -- down
LD -- local site signaled down   CF -- call admission control failure
RD -- remote site signaled down  SC -- local and remote site ID collision
LN -- local site not designated  LM -- local site ID not minimum designated
RN -- remote site not designated RM -- remote site ID not minimum designated
XX -- unknown connection status  IL -- no incoming label
MM -- MTU mismatch               MI -- Mesh-Group ID not available
BK -- Backup connection          ST -- Standby connection
PF -- Profile parse failure      PB -- Profile busy
RS -- remote site standby        SN -- Static Neighbor
VM -- VLAN ID mismatch

Legend for interface status
Up -- operational
Dn -- down

Instance: CUST_SITE1
  Local site: 1 (1)
    connection-site           Type  St     Time last up          # Up trans
    2                         rmt   Up     Aug 21 23:37:13 2018           1
      Remote PE: 192.0.2.6, Negotiated control-word: No
      Incoming label: 262154, Outgoing label: 262145
      Local interface: lsi.1048581, Status: Up, Encapsulation: VPLS
        Description: Intf - vpls CUST_SITE1 local site 1 remote site 2
```

On SJ PE

```
lab@SJ> show vpls connections
Layer-2 VPN connections:

Legend for connection status (St)
EI -- encapsulation invalid      NC -- interface encapsulation not CCC/TCC/VPLS
EM -- encapsulation mismatch     WE -- interface and instance encaps not same
VC-Dn -- Virtual circuit down    NP -- interface hardware not present
CM -- control-word mismatch      -> -- only outbound connection is up
CN -- circuit not provisioned    <- -- only inbound connection is up
OR -- out of range               Up -- operational
OL -- no outgoing label          Dn -- down
LD -- local site signaled down   CF -- call admission control failure
RD -- remote site signaled down  SC -- local and remote site ID collision
LN -- local site not designated  LM -- local site ID not minimum designated
RN -- remote site not designated RM -- remote site ID not minimum designated
XX -- unknown connection status  IL -- no incoming label
MM -- MTU mismatch               MI -- Mesh-Group ID not available
BK -- Backup connection          ST -- Standby connection
PF -- Profile parse failure      PB -- Profile busy
RS -- remote site standby        SN -- Static Neighbor
VM -- VLAN ID mismatch

Legend for interface status
Up -- operational
Dn -- down

Instance: CUST_SITE2
  Local site: 2 (2)
    connection-site           Type  St     Time last up          # Up trans
    1                         rmt   Up     Aug 21 23:38:30 2018           1
      Remote PE: 192.0.2.1, Negotiated control-word: No
      Incoming label: 262145, Outgoing label: 262154
      Local interface: lsi.1048580, Status: Up, Encapsulation: VPLS
        Description: Intf - vpls CUST_SITE2 local site 2 remote site 1
```

You can see that the VPLS connections are up and the incoming and outgoing label values can be seen in the outputs.

Customer Site 1 and Site 2 have reachability to each other and act as a direct connection, as shown here:

```
lab@SITE1> show ospf neighbor
Address          Interface              State     ID               Pri  Dead
100.64.65.2      ge-0/0/7.0             Full      100.64.65.2      128    34

lab@SITE2> show ospf neighbor
Address          Interface              State     ID               Pri  Dead
100.64.65.1      ge-0/0/8.0             Full      100.64.65.1      128    37
```

The VPLS forwarding table displays the MAC addresses learned, interface details, and the label used to forward frames to remote sites:

```
lab@CA> show route forwarding-table family vpls
Routing table: CUST_SITE1.vpls
VPLS:
Destination        Type RtRef Next hop           Type Index NhRef Netif
default            perm     0                    rjct   545     1
ge-0/0/7.0         user     0                    comp   569     2
lsi.1048581        user     0                    comp   622     2
00:0c:29:3f:c6:a5/48 dynm     0                  ucst   567     5 ge-0/0/7.0
00:0c:29:3f:c6:af/48 dynm     0                  indr 262144     4
                              198.51.100.19     Push 262145, Push 300480(top)   585     1 ge-0/0/2.0
00:0c:29:4f:ef:27/48 dynm     0                  ucst   567     5 ge-0/0/7.0
00:0c:29:94:67:c5/48 dynm     0                  ucst   567     5 ge-0/0/7.0
lab@SJ> show route forwarding-table family vpls
Routing table: CUST_SITE2.vpls
VPLS:
Destination        Type RtRef Next hop           Type Index NhRef Netif
default            perm     0                    rjct   544     1
ge-0/0/8.0         user     0                    comp   571     2
lsi.1048580        user     0                    comp   584     2
00:0c:29:3f:c6:a5/48 dynm     0                  indr 262143     6
                              198.51.100.17     Push 262154, Push 300336(top)   581     1 ge-0/0/1.0
00:0c:29:3f:c6:af/48 dynm     0                  ucst   569     3 ge-0/0/8.0
00:0c:29:4f:ef:27/48 dynm     0                  indr 262143     6
                              198.51.100.17     Push 262154, Push 300336(top)   581     1 ge-0/0/1.0
00:0c:29:94:67:c5/48 dynm     0                  indr 262143     6
                              198.51.100.17     Push 262154, Push 300336(top)   581     1 ge-0/0/1.0
```

VPLS Using LDP

VPLS can be configured using the label distribution protocol (LDP) as the signaling protocol. LDP signaling doesn't allow auto discovery of remote PEs, instead you must explicitly refer neighbor IP addresses. The PE routers use LDP signaling to distribute the VPLS instance to label mapping (refer to Figure 4.5 for further implementation).

For this configuration you need to have an LDP LSP in place between the PE routers having VPLS customers, or you can make use of LDP tunneling. The main goal is to have an LDP label in the database for each VPLS PE.

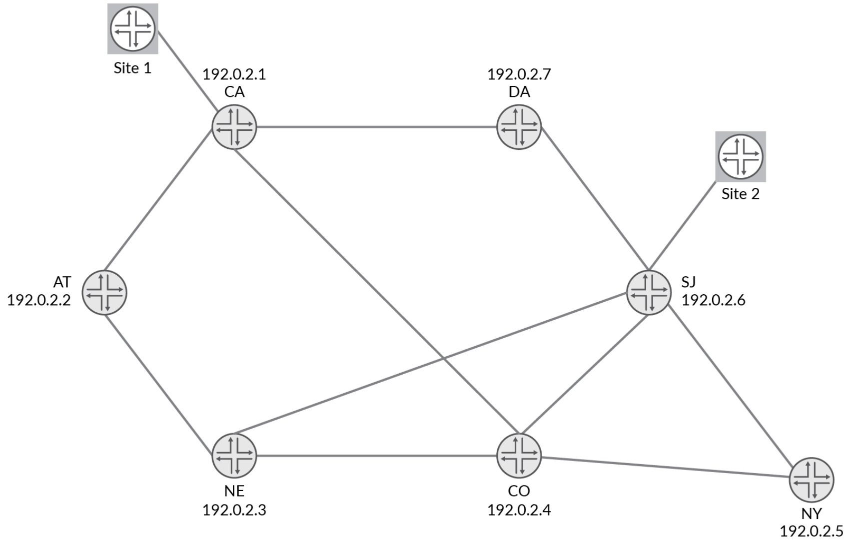

Figure 4.5 MPLS VPLS using LDP

In Figure 4.5, the required MPLS backbone is up and running with LDP label exchange. Site 1 and Site 2 are the VPLS customer sites.

There is always a Virtual Circuit(VC) label associated with each VPLS instance, which is exchanged between PEs. Along with the VC label, a FEC element is also associated for each VPLS instance. The FEC element supported by the Junos OS is FEC 128 (see Figure 4.6). This FEC element is used to brief the parameters of a PE router to the remote PE router.

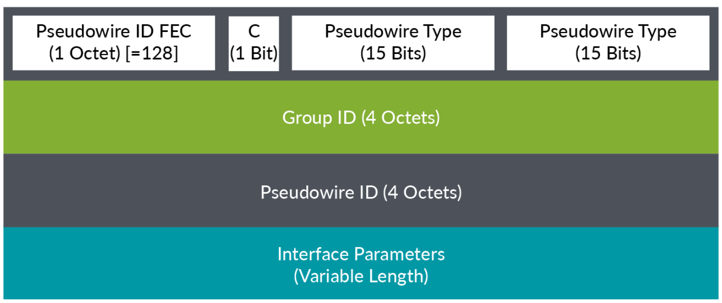

Figure 4.6 FEC 128

A few explanations for FEC 128:

- Pseudowire ID FEC: Set to 0x80 (128 in decimal).
- C Bit: This indicates whether or not the Martini control word is present; by default it is set to 0 in Junos (hence no control word is present).
- Pseudowire type: Describes the Layer 2 encapsulation type on the VPN interface. (Ethernet is only supported in Junos.).
- Pseudowire info length: Length of pseudowire ID field and Interface Parameters in octets.
- Group ID: Identifies a group of pseudowires.
- PseudowireID: A 32-bit value that together with the pseudowire type identifies the pseudowire. Both of these values should match on each endpoint of a pseudowire.
- Interface Parameter Sub-TLV: This variable length TLV is used to provide interface-specific parameters like MTU.

Using the PW ID FEC, each endpoint of the pseudowire (VC) initiates a unidirectional LSP. Once both endpoints set up their outgoing and incoming LSPs, they are bound together, provided that their PW ID and PW Type match.

Configure Interfaces

On CA PE

```
lab@CA> show configuration interfaces ge-0/0/7
description CUST_VPLS_SITE1;
encapsulation ethernet-vpls;
unit 0 {
    family vpls;
}
```

On SJ PE

```
lab@SJ> show configuration interfaces ge-0/0/8
description CUST_VPLS_SITE2;
encapsulation ethernet-vpls;
unit 0 {
    family vpls;
}
```

Configure the VPLS Instance

On CA PE

```
lab@CA>show configuration routing-instances CUST_VPLS_LDP_SITE1
instance-type vpls;
interface ge-0/0/7.0;
```

```
protocols {
    vpls {
        no-tunnel-services;
        vpls-id 200;
        neighbor 192.2.0.6;
    }
}
```

On SJ PE

```
lab@SJ>show configuration routing-instances CUST_VPLS_LDP_SITE2
instance-type vpls;
interface ge-0/0/8.0;
protocols {
    vpls {
        no-tunnel-services;
        vpls-id 200;
        neighbor 192.2.0.1;
    }
}
```

You need to explicitly configure the remote VPLS PE IP address since autodiscovery is not supported by LDP. The VPLS-ID value should be unique within the domain and should be the same on both ends.

You can confirm from the next round of output that no control word is present with the associated VC label for VPLS instance.

Verification

Let's verify the VPLS connection.

On CA PE

```
lab@CA>show vpls connections
Layer-2 VPN connections:

Legend for connection status (St)
EI -- encapsulation invalid      NC -- interface encapsulation not CCC/TCC/VPLS
EM -- encapsulation mismatch     WE -- interface and instance encaps not same
VC-Dn -- Virtual circuit down    NP -- interface hardware not present
CM -- control-word mismatch      -> -- only outbound connection is up
CN -- circuit not provisioned    <- -- only inbound connection is up
OR -- out of range               Up -- operational
OL -- no outgoing label          Dn -- down
LD -- local site signaled down   CF -- call admission control failure
RD -- remote site signaled down  SC -- local and remote site ID collision
LN -- local site not designated  LM -- local site ID not minimum designated
RN -- remote site not designated RM -- remote site ID not minimum designated
XX -- unknown connection status  IL -- no incoming label
MM -- MTU mismatch               MI -- Mesh-Group ID not available
BK -- Backup connection          ST -- Standby connection
PF -- Profile parse failure      PB -- Profile busy
RS -- remote site standby        SN -- Static Neighbor
LB -- Local site not best-site   RB -- Remote site not best-site
VM -- VLAN ID mismatch           HS -- Hot-standby Connection
```

```
Legend for interface status
Up -- operational
Dn -- down

Instance: CUST_VPLS_LDP_SITE1
  VPLS-id: 200
    Neighbor                  Type  St     Time last up          # Up trans
    192.2.0.6(vpls-id 200)    rmt   Up     Aug 30 15:54:23 2018           1
      Remote PE: 192.2.0.6, Negotiated control-word: No
      Incoming label: 262145, Outgoing label: 262146
      Negotiated PW status TLV: No
      Local interface: lsi.1048576, Status: Up, Encapsulation: ETHERNET
        Description: Intf - vpls CUST_VPLS_LDP_SITE1 neighbor 192.2.0.6 vpls-id 200
      Flow Label Transmit: No, Flow Label Receive: No
```

On SJ PE

```
lab@SJ>show vpls connections
Layer-2 VPN connections:

Legend for connection status (St)
EI -- encapsulation invalid      NC -- interface encapsulation not CCC/TCC/VPLS
EM -- encapsulation mismatch     WE -- interface and instance encaps not same
VC-Dn -- Virtual circuit down    NP -- interface hardware not present
CM -- control-word mismatch      -> -- only outbound connection is up
CN -- circuit not provisioned    <- -- only inbound connection is up
OR -- out of range               Up -- operational
OL -- no outgoing label          Dn -- down
LD -- local site signaled down   CF -- call admission control failure
RD -- remote site signaled down  SC -- local and remote site ID collision
LN -- local site not designated  LM -- local site ID not minimum designated
RN -- remote site not designated RM -- remote site ID not minimum designated
XX -- unknown connection status  IL -- no incoming label
MM -- MTU mismatch               MI -- Mesh-Group ID not available
BK -- Backup connection          ST -- Standby connection
PF -- Profile parse failure      PB -- Profile busy
RS -- remote site standby        SN -- Static Neighbor
LB -- Local site not best-site   RB -- Remote site not best-site
VM -- VLAN ID mismatch

Legend for interface status
Up -- operational
Dn -- down

Instance: CUST_VPLS_LDP_SITE2
  VPLS-id: 200
    Neighbor                  Type  St     Time last up          # Up trans
    192.2.0.1(vpls-id 200) rmt Up    Aug 30 15:59:08 2018           1
      Remote PE: 192.2.0.1, Negotiated control-word: No
      Incoming label: 262146, Outgoing label: 262145
      Negotiated PW status TLV: No
      Local interface: lsi.1048576, Status: Up, Encapsulation: ETHERNET
        Description: Intf - vpls CUST_VPLS_LDP_SITE2 neighbor 192.2.0.1 vpls-id 200
      Flow Label Transmit: No, Flow Label Receive: No
```

And the MAC address verification.

On CA PE

```
lab@CA> show route forwarding-table family vpls
Routing table: CUST_SITE1.vpls
VPLS:
Destination        Type RtRef Next hop           Type Index NhRef Netif
default            perm     0                    rjct   545     1
ge-0/0/7.0         user     0                    comp   569     2
lsi.1048581        user     0                    comp   622     2
00:0c:29:3f:c6:a5/48 dynm     0                  ucst   567     5 ge-0/0/7.0
00:0c:29:3f:c6:af/48 dynm     0                  indr 262144     4
                              198.51.100.19     Push 262146, Push 22457(top)   585     1 ge-0/0/2.0
```

On SJ PE

```
lab@SJ> show route forwarding-table family vpls
Routing table: CUST_SITE2.vpls
VPLS:
Destination        Type RtRef Next hop           Type Index NhRef Netif
default            perm     0                    rjct   544     1
ge-0/0/8.0         user     0                    comp   571     2
lsi.1048580        user     0                    comp   584     2
00:0c:29:3f:c6:a5/48 dynm     0                  indr 262143     6
                              198.51.100.17     Push 262145, Push 19639(top)   581     1 ge-0/0/1.0
00:0c:29:3f:c6:af/48 dynm     0                  ucst   569     3 ge-0/0/8.0
```

Verifying the L2 VPN routes.

```
lab@CA>show route table l2circuit.0

l2circuit.0: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

192.2.0.6:NoCtrlWord:5:200:Local/96
                   *[VPLS/7] 02:30:44, metric2 1
                    > to 20.50.40.4 via ae0.1
192.2.0.6:NoCtrlWord:5:200:Remote/96
                   *[LDP/9] 02:33:07
                      Discard

lab@SJ>show route table l2circuit.0

l2circuit.0: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

192.2.0.1:NoCtrlWord:5:200:Local/96
                   *[VPLS/7] 02:32:52, metric2 1
                    > to 20.50.40.5 via ae0.1
192.2.0.1:NoCtrlWord:5:200:Remote/96
                   *[LDP/9] 02:30:28
                      Discard
```

Customer Site 1 and Site 2 have reachability to each other and act as a direct connection, as shown here:

```
lab@SITE1> show ospf neighbor
Address          Interface              State     ID               Pri  Dead
100.64.65.2      ge-0/0/7.0             Full      100.64.65.2      128    34

lab@SITE2> show ospf neighbor
Address          Interface              State     ID               Pri  Dead
100.64.65.1      ge-0/0/8.0             Full      100.64.65.1      128    37
```

Summary

MPLS VPNs are a popular solution since they provide customers with transparency between two or more geographically diverse sites without exposing their data to the Internet or experiencing the Internet's unpredictable latency or packet loss.

This chapter shared some major concepts of L3VPNs as well as L2VPNs, including their control and data plane traffic flow. It's recommended to lab the scenarios and verify the working of this book's concepts in order to aid in learning the process of getting MPLS up and running.

In Chapter 5 you will deploy the VPN end to end, starting off with building a MPLS backbone network.

Chapter 5

Deploying VPN End to End

This chapter deploys a VPN network end to end, including L3VPN and VPLS. It covers multiple customers over a MPLS backbone, leveraging Layer 3 VPN services and Layer 2 VPN services. You will set up a MPLS network with customers utilizing the services configured on it, implementing the knowledge you have gained from previous chapters.

Network Setup

Figure 5.1 is the referenced topology for this chapter. Customer A and Customer B are Layer 3 VPN customers, and have multiple respective sites, whereas Customer C is a VPLS customer with three sites. Figure 5.1 gives a clear picture of the respective site locations, as well as the routes that are being exported from those sites.

All the routers run the LDP protocol except the AT and DA LERs. All the routers also run the RSVP protocol, except NY. We will build RSVP-TE LSPs between all the PE routers except NY, since it has only LDP enabled. *Not enabling* RSVP on NY can be an example of a multi-vendor implementation, where NY is a non-Junos router that doesn't support RSVP extensions to the level that Junos supports. Moreover, not enabling LDP on the DA and AT LERs is simply for the book's learning purposes. All the RSVP-TEs will be protected by bypass LSPs.

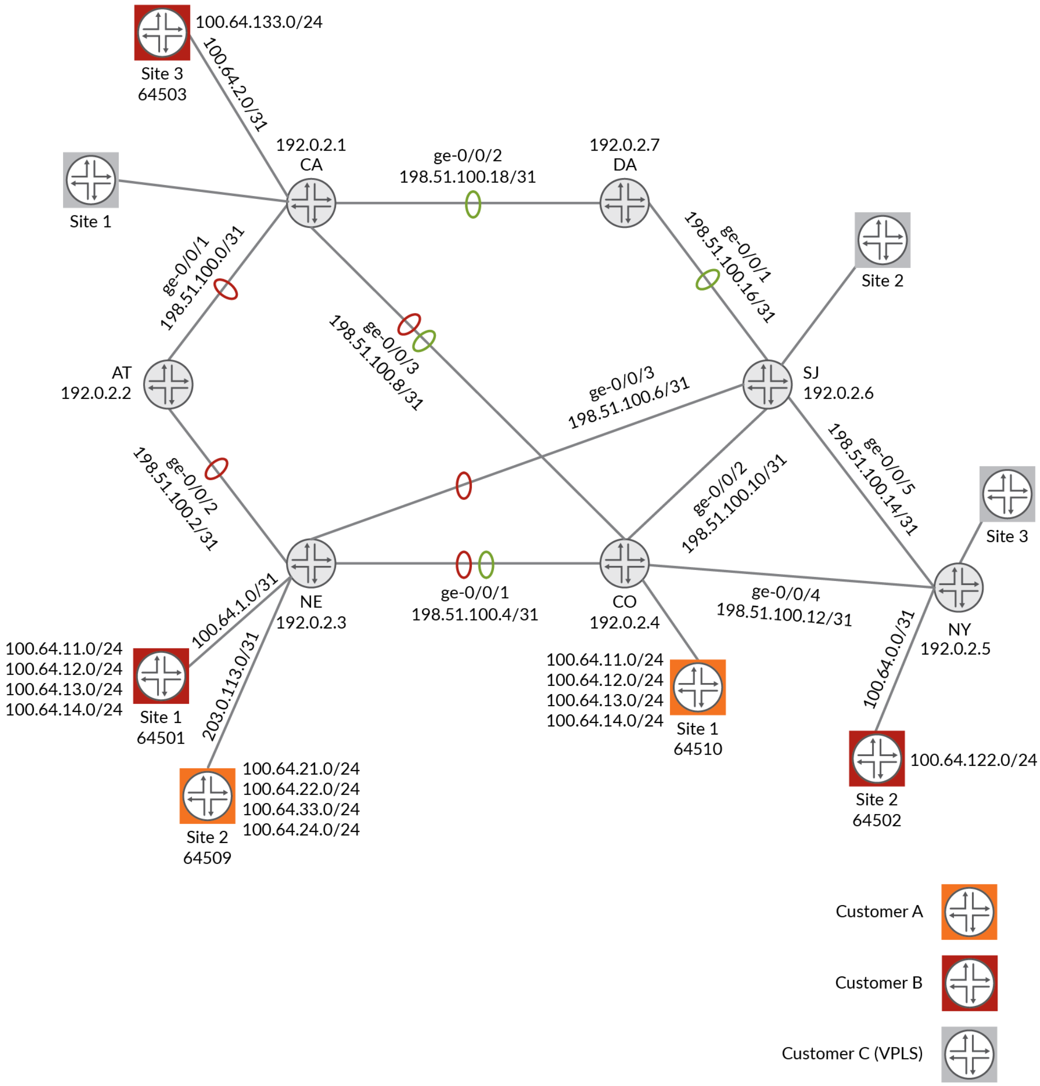

Figure 5.1 MPLS Backbone

You protect the RSVP-TE LSPs with bypass LSPs, because they are shared between multiple main LSPs and hence the router has to maintain fewer states, whereas fast reroute is a one-to-one backup LSP which leads to maintaining multiple states in all transit, as well as LER, routers. For a small MPLS network, protecting RSVP-TE LSP with fast reroute should not be a problem, but for larger MPLS networks, protection with bypass LSPs is always a better, more scalable option.

Configuration Steps

1. Enable traffic engineering with IGP.
2. Enable RSVP and MPLS on all MPLS backbone routers except NY.
3. Set up RSVP-TE LSPs between all LER routers except NY.
4. Enable LDP on routers except AT and DA.
5. Label exchange between NY and CA LERs.
6. Implement Customer A, Customer B, and Customer C VPNs.
7. Verify the reachability between customer sites.

Configuration

First let's combine steps 1 and 2 to enable traffic engineering with IGP and then enable RSVP and MPLS on all MPLS backbone routers except NY:

```
lab@AT> show configuration protocols | display inheritance no-comments
rsvp {
    interface ge-0/0/1.0;
    interface ge-0/0/2.0;
}
mpls {
    interface ge-0/0/1.0;
    interface ge-0/0/2.0;
}
ospf {
    traffic-engineering;
    area 0.0.0.0 {
        interface ge-0/0/1.0 {
            interface-type p2p;
        }
        interface ge-0/0/2.0 {
            interface-type p2p;
        }
        interface lo0.0 {
            interface-type p2p;
            passive;
        }
    }
}
```

A similar configuration will be on all backbone routers.

Here's the traffic engineering database:

```
lab@AT> show ted database
```

```
TED database: 0 ISIS nodes 7 INET nodes
ID                            Type Age(s) LnkIn LnkOut Protocol
192.0.2.1                     Rtr     955     3      2 OSPF(0.0.0.0)
    To: 192.0.2.2, Local: 198.51.100.0, Remote: 198.51.100.1
      Local interface index: 71, Remote interface index: 0
    To: 192.0.2.7, Local: 198.51.100.18, Remote: 198.51.100.19
      Local interface index: 72, Remote interface index: 0
ID                            Type Age(s) LnkIn LnkOut Protocol
192.0.2.2                     Rtr     988     2      2 OSPF(0.0.0.0)
    To: 192.0.2.3, Local: 198.51.100.2, Remote: 198.51.100.3
      Local interface index: 72, Remote interface index: 0
    To: 192.0.2.1, Local: 198.51.100.1, Remote: 198.51.100.0
      Local interface index: 71, Remote interface index: 0
ID                            Type Age(s) LnkIn LnkOut Protocol
192.0.2.3                     Rtr     878     3      2 OSPF(0.0.0.0)
    To: 192.0.2.2, Local: 198.51.100.3, Remote: 198.51.100.2
      Local interface index: 74, Remote interface index: 0
    To: 192.0.2.4, Local: 198.51.100.4, Remote: 198.51.100.5
      Local interface index: 73, Remote interface index: 0
ID                            Type Age(s) LnkIn LnkOut Protocol
192.0.2.4                     Rtr     922     2      3 OSPF(0.0.0.0)
    To: 192.0.2.3, Local: 198.51.100.5, Remote: 198.51.100.4
      Local interface index: 71, Remote interface index: 0
    To: 192.0.2.1, Local: 198.51.100.8, Remote: 198.51.100.9
      Local interface index: 73, Remote interface index: 0
    To: 192.0.2.6, Local: 198.51.100.10, Remote: 198.51.100.11
      Local interface index: 72, Remote interface index: 0
ID                            Type Age(s) LnkIn LnkOut Protocol
192.0.2.5                     Rtr    2540     0      0 OSPF(0.0.0.0)
192.0.2.6                     Rtr     803     2      3 OSPF(0.0.0.0)
    To: 192.0.2.3, Local: 198.51.100.7, Remote: 198.51.100.6
      Local interface index: 84, Remote interface index: 0
    To: 192.0.2.7, Local: 198.51.100.16, Remote: 198.51.100.17
      Local interface index: 82, Remote interface index: 0
    To: 192.0.2.4, Local: 198.51.100.11, Remote: 198.51.100.10
      Local interface index: 83, Remote interface index: 0
ID                            Type Age(s) LnkIn LnkOut Protocol
192.0.2.7                     Rtr     905     2      2 OSPF(0.0.0.0)
    To: 192.0.2.1, Local: 198.51.100.19, Remote: 198.51.100.18
      Local interface index: 73, Remote interface index: 0
    To: 192.0.2.6, Local: 198.51.100.17, Remote: 198.51.100.16
      Local interface index: 72, Remote interface index: 0
```

As you can see, the NY router traffic engineering database is not available since it's not running RSVP.

Step 3: Set up RSVP-TE LSPs between all LER routers except NY.

You want CA_CO and CA_SJ to take different EROs to manage bandwidth utilization. Hence, we will make use of `admin-groups` to manipulate the EROs:

```
lab@CA> show configuration interfaces
ge-0/0/1 {
    description AT;
    unit 0 {
        family inet {
            address 198.51.100.0/31;
```

```
        }
        family mpls;
    }
}
ge-0/0/2 {
    description DA;
    unit 0 {
        family inet {
            address 198.51.100.18/31;
        }
        family mpls;
    }
}
ge-0/0/3 {
    description CO;
    unit 0 {
        family inet {
            address 198.51.100.9/31;
        }
        family mpls;
    }
}
lo0 {
    unit 0 {
        family inet {
            address 192.0.2.1/32;
        }
        family mpls;
    }
}
lab@CA> show configuration protocols rsvp
interface ge-0/0/2.0 {
    link-protection;
}
interface ge-0/0/1.0 {
    link-protection;
}
lab@CA> show configuration protocols mpls
admin-groups {
    red 0;
    green 1;
}
label-switched-path CA_AT {
    to 192.0.2.2;
    link-protection;
}
label-switched-path CA_DA {
    to 192.0.2.7;
    link-protection;
}
label-switched-path CA_SJ {
    to 192.0.2.6;
    admin-group include-any green;
    link-protection;
}
label-switched-path CA_CO {
    to 192.0.2.4;
    bandwidth 100;
```

```
    admin-group include-any red;
    link-protection;
}
label-switched-path CA_NE {
    to 192.0.2.3;
    link-protection;
}
interface ge-0/0/1.0 {
    admin-group red;
}
interface ge-0/0/2.0 {
    admin-group green;
}
interface ge-0/0/3.0 {
    admin-group [ red green ];
}
```

Similar to the configuration on CA LER, all routers in the backbone are configured and the LSPs are up as shown below. Moreover, the bypass LSPs are up protecting the main LSPs.

On AT

```
lab@AT> show mpls lsp
Ingress LSP: 5 sessions
To              From            State Rt P     ActivePath       LSPname
192.0.2.1       192.0.2.2       Up     0 *                      AT_CA
192.0.2.3       192.0.2.2       Up     0 *                      AT_NE
192.0.2.4       192.0.2.2       Up     0 *                      AT_CO
192.0.2.6       192.0.2.2       Up     0 *                      AT_SJ
192.0.2.7       192.0.2.2       Up     0 *                      AT_DA
Total 5 displayed, Up 5, Down 0

Egress LSP: 7 sessions
To              From            State   Rt Style Labelin Labelout LSPname
192.0.2.2       192.0.2.6       Up       0  1 SE       3        - SJ_AT
192.0.2.2       192.0.2.4       Up       0  1 SE       3        - CO_AT
192.0.2.2       192.0.2.7       Up       0  1 SE       3        - DA_AT
192.0.2.2       192.0.2.1       Up       0  1 SE       3        - CA_AT
192.0.2.2       192.0.2.1       Up       0  1 SE       3        - Bypass->198.51.100.1
192.0.2.2       192.0.2.3       Up       0  1 SE       3        - NE_AT
192.0.2.2       192.0.2.3       Up       0  1 SE       3        - Bypass->198.51.100.2
Total 7 displayed, Up 7, Down 0

Transit LSP: 11 sessions
To              From            State   Rt Style Labelin Labelout LSPname
192.0.2.1       192.0.2.7       Up       0  1 SE  300320        3 Bypass->198.51.100.18
192.0.2.1       192.0.2.3       Up       0  1 SE  300400        3 NE_CA
192.0.2.3       192.0.2.4       Up       0  1 SE  300304        3 Bypass->198.51.100.4
192.0.2.3       192.0.2.1       Up       0  1 SE  300256        3 CA_NE
192.0.2.4       192.0.2.6       Up       0  1 SE  300368   300640 Bypass->198.51.100.10
192.0.2.4       192.0.2.1       Up       0  1 SE  300272   300480 CA_CO
192.0.2.4       192.0.2.3       Up       0  1 SE  300416   300640 Bypass->198.51.100.5
192.0.2.6       192.0.2.4       Up       0  1 SE  300352   300560 Bypass->198.51.100.11
192.0.2.6       192.0.2.7       Up       0  1 SE  300336   300576 Bypass->198.51.100.16
192.0.2.7       192.0.2.6       Up       0  1 SE  300288   300480 Bypass->198.51.100.17
```

```
192.0.2.7       192.0.2.1       Up       0  1 SE  300384   300672 Bypass->198.51.100.19
Total 11 displayed, Up 11, Down 0
```

On CA Router

```
lab@CA> show mpls lsp
Ingress LSP: 5 sessions
To              From            State Rt P     ActivePath       LSPname
192.0.2.2       192.0.2.1       Up     0 *                      CA_AT
192.0.2.3       192.0.2.1       Up     0 *                      CA_NE
192.0.2.4       192.0.2.1       Up     0 *                      CA_CO
192.0.2.6       192.0.2.1       Up     0 *                      CA_SJ
192.0.2.7       192.0.2.1       Up     0 *                      CA_DA
Total 5 displayed, Up 5, Down 0

Egress LSP: 7 sessions
To              From            State   Rt Style Labelin Labelout LSPname
192.0.2.1       192.0.2.6       Up       0  1 SE       3        - SJ_CA
192.0.2.1       192.0.2.4       Up       0  1 SE       3        - CO_CA
192.0.2.1       192.0.2.7       Up       0  1 SE       3        - DA_CA
192.0.2.1       192.0.2.7       Up       0  1 SE       3        - Bypass->198.51.100.18
192.0.2.1       192.0.2.2       Up       0  1 SE       3        - AT_CA
192.0.2.1       192.0.2.2       Up       0  1 SE       3        - Bypass->198.51.100.0
192.0.2.1       192.0.2.3       Up       0  1 SE       3        - NE_CA
Total 7 displayed, Up 7, Down 0

Transit LSP: 8 sessions
To              From            State   Rt Style Labelin Labelout LSPname
192.0.2.2       192.0.2.3       Up       0  1 SE  300624        3 Bypass->198.51.100.2
192.0.2.3       192.0.2.4       Up       0  1 SE  300544   300304 Bypass->198.51.100.4
192.0.2.3       192.0.2.2       Up       0  1 SE  300576   300432 Bypass->198.51.100.3
192.0.2.4       192.0.2.6       Up       0  1 SE  300608   300368 Bypass->198.51.100.10
192.0.2.4       192.0.2.3       Up       0  1 SE  300640   300512 Bypass->198.51.100.5
192.0.2.6       192.0.2.4       Up       0  1 SE  300560   300400 Bypass->198.51.100.11
192.0.2.6       192.0.2.7       Up       0  1 SE  300592   300336 Bypass->198.51.100.16
192.0.2.7       192.0.2.6       Up       0  1 SE  300480        3 Bypass->198.51.100.17
Total 8 displayed, Up 8, Down 0
```

On DA Router

```
lab@DA> show mpls lsp
Ingress LSP: 5 sessions
To              From            State Rt P     ActivePath       LSPname
192.0.2.1       192.0.2.7       Up     0 *                      DA_CA
192.0.2.2       192.0.2.7       Up     0 *                      DA_AT
192.0.2.3       192.0.2.7       Up     0 *                      DA_NE
192.0.2.4       192.0.2.7       Up     0 *                      DA_CO
192.0.2.6       192.0.2.7       Up     0 *                      DA_SJ
Total 5 displayed, Up 5, Down 0

Egress LSP: 7 sessions
To              From            State   Rt Style Labelin Labelout LSPname
192.0.2.7       192.0.2.6       Up       0  1 SE       3        - SJ_DA
192.0.2.7       192.0.2.6       Up       0  1 SE       3        - Bypass->198.51.100.17
192.0.2.7       192.0.2.4       Up       0  1 SE       3        - CO_DA
192.0.2.7       192.0.2.2       Up       0  1 SE       3        - AT_DA
```

```
192.0.2.7       192.0.2.1       Up       0  1 SE       3        - CA_DA
192.0.2.7       192.0.2.1       Up       0  1 SE       3        - Bypass->198.51.100.19
192.0.2.7       192.0.2.3       Up       0  1 SE       3        - NE_DA
Total 7 displayed, Up 7, Down 0

Transit LSP: 11 sessions
To              From            State   Rt Style Labelin Labelout LSPname
192.0.2.1       192.0.2.6       Up       0  1 SE  300336        3 SJ_CA
192.0.2.1       192.0.2.4       Up       0  1 SE  300352        3 CO_CA
192.0.2.1       192.0.2.2       Up       0  1 SE  300384        3 Bypass->198.51.100.0
192.0.2.2       192.0.2.1       Up       0  1 SE  300464   300528 Bypass->198.51.100.1
192.0.2.2       192.0.2.3       Up       0  1 SE  300496   300624 Bypass->198.51.100.2
192.0.2.3       192.0.2.4       Up       0  1 SE  300416   300544 Bypass->198.51.100.4
192.0.2.3       192.0.2.2       Up       0  1 SE  300432   300448 Bypass->198.51.100.3
192.0.2.4       192.0.2.6       Up       0  1 SE  300448   300608 Bypass->198.51.100.10
192.0.2.4       192.0.2.3       Up       0  1 SE  300512   300560 Bypass->198.51.100.5
192.0.2.6       192.0.2.4       Up       0  1 SE  300400        3 Bypass->198.51.100.11
192.0.2.6       192.0.2.1       Up       0  1 SE  300480        3 CA_SJ
Total 11 displayed, Up 11, Down 0
```

On SJ LER

```
lab@SJ> show mpls lsp
Ingress LSP: 5 sessions
To              From            State Rt P     ActivePath       LSPname
192.0.2.1       192.0.2.6       Up     0 *                      SJ_CA
192.0.2.2       192.0.2.6       Up     0 *                      SJ_AT
192.0.2.3       192.0.2.6       Up     0 *                      SJ_NE
192.0.2.4       192.0.2.6       Up     0 *                      SJ_CO
192.0.2.7       192.0.2.6       Up     0 *                      SJ_DA
Total 5 displayed, Up 5, Down 0

Egress LSP: 7 sessions
To              From            State   Rt Style Labelin Labelout LSPname
192.0.2.6       192.0.2.4       Up       0  1 SE       3        - CO_SJ
192.0.2.6       192.0.2.4       Up       0  1 SE       3        - Bypass->198.51.100.11
192.0.2.6       192.0.2.7       Up       0  1 SE       3        - DA_SJ
192.0.2.6       192.0.2.7       Up       0  1 SE       3        - Bypass->198.51.100.16
192.0.2.6       192.0.2.2       Up       0  1 SE       3        - AT_SJ
192.0.2.6       192.0.2.1       Up       0  1 SE       3        - CA_SJ
192.0.2.6       192.0.2.3       Up       0  1 SE       3        - NE_SJ
Total 7 displayed, Up 7, Down 0

Transit LSP: 15 sessions
To              From            State   Rt Style Labelin Labelout LSPname
192.0.2.1       192.0.2.4       Up       0  1 SE  300400   300352 CO_CA
192.0.2.1       192.0.2.7       Up       0  1 SE  300496   300464 Bypass->198.51.100.18
192.0.2.1       192.0.2.2       Up       0  1 SE  300464   300384 Bypass->198.51.100.0
192.0.2.2       192.0.2.7       Up       0  1 SE  300240   300288 DA_AT
192.0.2.2       192.0.2.1       Up       0  1 SE  300528   300496 Bypass->198.51.100.1
192.0.2.2       192.0.2.3       Up       0  1 SE  300544   300496 Bypass->198.51.100.2
192.0.2.3       192.0.2.4       Up       0  1 SE  300480   300416 Bypass->198.51.100.4
192.0.2.3       192.0.2.7       Up       0  1 SE  300224   300272 DA_NE
192.0.2.3       192.0.2.2       Up       0  1 SE  300448   300432 Bypass->198.51.100.3
192.0.2.4       192.0.2.7       Up       0  1 SE  299888        3 DA_CO
192.0.2.4       192.0.2.3       Up       0  1 SE  300560        3 Bypass->198.51.100.5
192.0.2.7       192.0.2.4       Up       0  1 SE  300000        3 CO_DA
```

```
192.0.2.7       192.0.2.2       Up       0  1 SE  300288        3 AT_DA
192.0.2.7       192.0.2.1       Up       0  1 SE  300512        3 Bypass->198.51.100.19
192.0.2.7       192.0.2.3       Up       0  1 SE  300256        3 NE_DA
Total 15 displayed, Up 15, Down 0
```

On CO LER

```
lab@CO> show mpls lsp
Ingress LSP: 5 sessions
To              From            State Rt P     ActivePath       LSPname
192.0.2.1       192.0.2.4       Up     0 *                      CO_CA
192.0.2.2       192.0.2.4       Up     0 *                      CO_AT
192.0.2.3       192.0.2.4       Up     0 *                      CO_NE
192.0.2.6       192.0.2.4       Up     0 *                      CO_SJ
192.0.2.7       192.0.2.4       Up     0 *                      CO_DA
Total 5 displayed, Up 5, Down 0

Egress LSP: 7 sessions
To              From            State   Rt Style Labelin Labelout LSPname
192.0.2.4       192.0.2.6       Up       0  1 SE       3        - SJ_CO
192.0.2.4       192.0.2.6       Up       0  1 SE       3        - Bypass->198.51.100.10
192.0.2.4       192.0.2.7       Up       0  1 SE       3        - DA_CO
192.0.2.4       192.0.2.2       Up       0  1 SE       3        - AT_CO
192.0.2.4       192.0.2.1       Up       0  1 SE       3        - CA_CO
192.0.2.4       192.0.2.3       Up       0  1 SE       3        - NE_CO
192.0.2.4       192.0.2.3       Up       0  1 SE       3        - Bypass->198.51.100.5
Total 7 displayed, Up 7, Down 0

Transit LSP: 16 sessions
To              From            State   Rt Style Labelin Labelout LSPname
192.0.2.1       192.0.2.7       Up       0  1 SE  300464   300592 Bypass->198.51.100.18
192.0.2.1       192.0.2.2       Up       0  1 SE  300448   300464 Bypass->198.51.100.0
192.0.2.2       192.0.2.6       Up       0  1 SE  300112   300256 SJ_AT
192.0.2.2       192.0.2.7       Up       0  1 SE  300288   300416 DA_AT
192.0.2.2       192.0.2.1       Up       0  1 SE  300496   300656 Bypass->198.51.100.1
192.0.2.2       192.0.2.3       Up       0  1 SE  300512   300544 Bypass->198.51.100.2
192.0.2.3       192.0.2.6       Up       0  1 SE  300256        3 SJ_NE
192.0.2.3       192.0.2.7       Up       0  1 SE  300272        3 DA_NE
192.0.2.3       192.0.2.2       Up       0  1 SE  300432        3 Bypass->198.51.100.3
192.0.2.6       192.0.2.7       Up       0  1 SE  300416        3 Bypass->198.51.100.16
192.0.2.6       192.0.2.2       Up       0  1 SE  300336        3 AT_SJ
192.0.2.6       192.0.2.3       Up       0  1 SE  300304        3 NE_SJ
192.0.2.7       192.0.2.6       Up       0  1 SE  300384   300496 Bypass->198.51.100.17
192.0.2.7       192.0.2.2       Up       0  1 SE  300352   300288 AT_DA
192.0.2.7       192.0.2.1       Up       0  1 SE  300480   300512 Bypass->198.51.100.19
192.0.2.7       192.0.2.3       Up       0  1 SE  300320   300256 NE_DA
Total 16 displayed, Up 16, Down 0
```

On NE LER

```
lab@NE> show mpls lsp
Ingress LSP: 5 sessions
To              From            State Rt P     ActivePath       LSPname
192.0.2.1       192.0.2.3       Up     0 *                      NE_CA
192.0.2.2       192.0.2.3       Up     0 *                      NE_AT
192.0.2.4       192.0.2.3       Up     0 *                      NE_CO
```

```
192.0.2.6       192.0.2.3       Up     0 *                      NE_SJ
192.0.2.7       192.0.2.3       Up     0 *                      NE_DA
Total 5 displayed, Up 5, Down 0

Egress LSP: 7 sessions
To              From            State   Rt Style Labelin Labelout LSPname
192.0.2.3       192.0.2.6       Up       0  1 SE       3        - SJ_NE
192.0.2.3       192.0.2.4       Up       0  1 SE       3        - CO_NE
192.0.2.3       192.0.2.4       Up       0  1 SE       3        - Bypass->198.51.100.4
192.0.2.3       192.0.2.7       Up       0  1 SE       3        - DA_NE
192.0.2.3       192.0.2.2       Up       0  1 SE       3        - AT_NE
192.0.2.3       192.0.2.2       Up       0  1 SE       3        - Bypass->198.51.100.3
192.0.2.3       192.0.2.1       Up       0  1 SE       3        - CA_NE
Total 7 displayed, Up 7, Down 0

Transit LSP: 15 sessions
To              From            State   Rt Style Labelin Labelout LSPname
192.0.2.1       192.0.2.7       Up       0  1 SE  300592   300320 Bypass->198.51.100.18
192.0.2.1       192.0.2.2       Up       0  1 SE  300608   300448 Bypass->198.51.100.0
192.0.2.2       192.0.2.6       Up       0  1 SE  300256        3 SJ_AT
192.0.2.2       192.0.2.4       Up       0  1 SE  300432        3 CO_AT
192.0.2.2       192.0.2.7       Up       0  1 SE  300416        3 DA_AT
192.0.2.2       192.0.2.1       Up       0  1 SE  300656        3 Bypass->198.51.100.1
192.0.2.4       192.0.2.6       Up       0  1 SE  300640        3 Bypass->198.51.100.10
192.0.2.4       192.0.2.2       Up       0  1 SE  300064        3 AT_CO
192.0.2.4       192.0.2.1       Up       0  1 SE  300480        3 CA_CO
192.0.2.6       192.0.2.4       Up       0  1 SE  300624   300352 Bypass->198.51.100.11
192.0.2.6       192.0.2.7       Up       0  1 SE  300576   300416 Bypass->198.51.100.16
192.0.2.6       192.0.2.2       Up       0  1 SE  300448   300336 AT_SJ
192.0.2.7       192.0.2.6       Up       0  1 SE  300496   300288 Bypass->198.51.100.17
192.0.2.7       192.0.2.2       Up       0  1 SE  300464   300352 AT_DA
192.0.2.7       192.0.2.1       Up       0  1 SE  300672   300480 Bypass->198.51.100.19
Total 15 displayed, Up 15, Down 0
```

Step 4: Enable LDP on routers except AT and DA.

On CA LER

```
lab@CA> show configuration protocols ldp
interface lo0.0;
```

On DA LER

```
lab@DA> show configuration protocols ldp

lab@DA>
```

On SJ LER

```
lab@SJ> show configuration protocols ldp
interface ge-0/0/2.0;
interface ge-0/0/3.0;
interface ge-0/0/5.0;
interface lo0.0;
```

On CO LER

```
lab@CO> show configuration protocols ldp
interface ge-0/0/1.0;
interface ge-0/0/2.0;
interface ge-0/0/4.0;
interface lo0.0;
```

On NY LER

```
lab@NY> show configuration protocols ldp
interface ge-0/0/4.0;
interface ge-0/0/5.0;
interface lo0.0;
```

On NE LER

```
lab@NE> show configuration protocols ldp
interface ge-0/0/1.0;
interface ge-0/0/3.0;
```

On AT LER

```
lab@AT> show configuration protocols ldp

lab@AT>
```

Step 5: Label exchange between NY and CA LERs.

The extended LDP sessions between CA < > SJ and CA < > CO happen over loopback, hence the above configuration has enabled LDP on the loopbacks on CA, SJ, and CO.

On CA LER

```
lab@CA> show configuration protocols mpls label-switched-path CA_CO
to 192.0.2.4;
ldp-tunneling;
bandwidth 100;
admin-group include-any red;

lab@CA> show configuration protocols mpls label-switched-path CA_SJ
to 192.0.2.6;
ldp-tunneling;
admin-group include-any green;
```

On SJ LER

```
lab@SJ> show configuration protocols mpls label-switched-path SJ_CA
to 192.0.2.1;
ldp-tunneling;
```

On CO LER

```
lab@CO> show configuration protocols mpls label-switched-path CO_CA
to 192.0.2.1;
ldp-tunneling;
```

Verification

On CA LER

```
lab@CA> show ldp session
  Address           State        Connection     Hold time
192.0.2.4           Operational  Open             28
192.0.2.6           Operational  Open             28
lab@CA> show ldp database
Input label database, 192.0.2.1:0--192.0.2.4:0
  Label     Prefix
 300400     192.0.2.1/32
 299808     192.0.2.3/32
      3     192.0.2.4/32
 299776     192.0.2.5/32
 299792     192.0.2.6/32

Output label database, 192.0.2.1:0--192.0.2.4:0
  Label     Prefix
      3     192.0.2.1/32
 300496     192.0.2.4/32
 300512     192.0.2.5/32
 300528     192.0.2.6/32

Input label database, 192.0.2.1:0--192.0.2.6:0
  Label     Prefix
 300432     192.0.2.1/32
 299776     192.0.2.3/32
 299808     192.0.2.4/32
 299792     192.0.2.5/32
      3     192.0.2.6/32

Output label database, 192.0.2.1:0--192.0.2.6:0
  Label     Prefix
      3     192.0.2.1/32
 300496     192.0.2.4/32
 300512     192.0.2.5/32
 300528     192.0.2.6/32
lab@CA> show route table inet.3

inet.3: 13 destinations, 21 routes (6 active, 0 holddown, 13 hidden)
+ = Active Route, - = Last Active, * = Both

192.0.2.2/32       *[RSVP/7/1] 02:47:53, metric 1
                    > to 198.51.100.1 via ge-0/0/1.0, label-switched-path CA_AT
                      to 198.51.100.19 via ge-0/0/2.0, label-switched-path Bypass->198.51.100.1
192.0.2.3/32       *[RSVP/7/1] 02:47:53, metric 2
                    > to 198.51.100.1 via ge-0/0/1.0, label-switched-path CA_NE
                      to 198.51.100.19 via ge-0/0/2.0, label-switched-path Bypass->198.51.100.1
192.0.2.4/32       *[RSVP/7/1] 00:44:40, metric 1
                    > to 198.51.100.1 via ge-0/0/1.0, label-switched-path CA_CO
                      to 198.51.100.19 via ge-0/0/2.0, label-switched-path Bypass->198.51.100.1
                    [LDP/9] 00:44:36, metric 1
                    > to 198.51.100.1 via ge-0/0/1.0, label-switched-path CA_CO
                      to 198.51.100.19 via ge-0/0/2.0, label-switched-path CA_CO
192.0.2.5/32       *[LDP/9] 00:44:36, metric 1
                    > to 198.51.100.1 via ge-0/0/1.0, label-switched-path CA_CO
```

```
                     to 198.51.100.19 via ge-0/0/2.0, label-switched-path CA_CO
192.0.2.6/32       *[RSVP/7/1] 00:44:40, metric 2
                    > to 198.51.100.19 via ge-0/0/2.0, label-switched-path CA_SJ
                      to 198.51.100.1 via ge-0/0/1.0, label-switched-path Bypass->198.51.100.19
                    [LDP/9] 00:20:31, metric 1
                    > to 198.51.100.19 via ge-0/0/2.0, label-switched-path CA_SJ
                      to 198.51.100.1 via ge-0/0/1.0, label-switched-path CA_SJ
192.0.2.7/32       *[RSVP/7/1] 02:47:53, metric 1
                    > to 198.51.100.19 via ge-0/0/2.0, label-switched-path CA_DA
                      to 198.51.100.1 via ge-0/0/1.0, label-switched-path Bypass->198.51.100.19
lab@CA> show route protocol ldp

inet.0: 25 destinations, 25 routes (23 active, 0 holddown, 2 hidden)

inet.3: 13 destinations, 21 routes (6 active, 0 holddown, 13 hidden)
+ = Active Route, - = Last Active, * = Both

192.0.2.4/32        [LDP/9] 04:46:47, metric 1
                    > to 198.51.100.1 via ge-0/0/1.0, label-switched-path CA_CO
                      to 198.51.100.19 via ge-0/0/2.0, label-switched-path CA_CO
192.0.2.5/32       *[LDP/9] 04:46:47, metric 1
                    > to 198.51.100.1 via ge-0/0/1.0, label-switched-path CA_CO
                      to 198.51.100.19 via ge-0/0/2.0, label-switched-path CA_CO
192.0.2.6/32        [LDP/9] 04:46:39, metric 1
                    > to 198.51.100.19 via ge-0/0/2.0, label-switched-path CA_SJ
                      to 198.51.100.1 via ge-0/0/1.0, label-switched-path CA_SJ

CUST-B_SITE3.inet.0: 10 destinations, 10 routes (10 active, 0 holddown, 0 hidden)

iso.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)

mpls.0: 22 destinations, 22 routes (22 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

300496             *[LDP/9] 04:46:47, metric 1
                    > to 198.51.100.1 via ge-0/0/1.0, label-switched-path CA_CO
                      to 198.51.100.19 via ge-0/0/2.0, label-switched-path CA_CO
300512             *[LDP/9] 04:46:47, metric 1
                    > to 198.51.100.1 via ge-0/0/1.0, label-switched-path CA_CO
                      to 198.51.100.19 via ge-0/0/2.0, label-switched-path CA_CO
300528             *[LDP/9] 04:46:39, metric 1
                    > to 198.51.100.19 via ge-0/0/2.0, label-switched-path CA_SJ
                      to 198.51.100.1 via ge-0/0/1.0, label-switched-path CA_SJ

bgp.l3vpn.0: 7 destinations, 7 routes (7 active, 0 holddown, 0 hidden)

bgp.l2vpn.0: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)

CUST3_SITE1.l2vpn.0: 3 destinations, 3 routes (3 active, 0 holddown, 0 hidden)
```

On SJ LER

```
lab@SJ> show ldp session
  Address           State        Connection     Hold time
192.0.2.1           Operational  Open             28
192.0.2.3           Operational  Open             26
192.0.2.4           Operational  Open             26
```

```
192.0.2.5           Operational  Open             22
lab@SJ> show route table inet.3

inet.3: 12 destinations, 20 routes (6 active, 0 holddown, 11 hidden)
+ = Active Route, - = Last Active, * = Both

192.0.2.1/32       *[RSVP/7/1] 02:38:28, metric 2
                    > to 198.51.100.17 via ge-0/0/1.0, label-switched-path SJ_CA
                      to 198.51.100.10 via ge-0/0/2.0, label-switched-path Bypass->198.51.100.17
                    [LDP/9] 00:01:24, metric 1
                    > to 198.51.100.17 via ge-0/0/1.0, label-switched-path SJ_CA
                      to 198.51.100.10 via ge-0/0/2.0, label-switched-path SJ_CA
192.0.2.2/32       *[RSVP/7/1] 02:38:29, metric 2
                    > to 198.51.100.10 via ge-0/0/2.0, label-switched-path SJ_AT
                      to 198.51.100.17 via ge-0/0/1.0, label-switched-path Bypass->198.51.100.10
192.0.2.3/32       *[RSVP/7/1] 02:38:28, metric 1
                    > to 198.51.100.10 via ge-0/0/2.0, label-switched-path SJ_NE
                      to 198.51.100.17 via ge-0/0/1.0, label-switched-path Bypass->198.51.100.10
                    [LDP/9] 00:04:14, metric 1
                    > to 198.51.100.6 via ge-0/0/3.0
192.0.2.4/32       *[RSVP/7/1] 02:38:27, metric 1
                    > to 198.51.100.10 via ge-0/0/2.0, label-switched-path SJ_CO
                      to 198.51.100.17 via ge-0/0/1.0, label-switched-path Bypass->198.51.100.10
                    [LDP/9] 00:15:16, metric 1
                    > to 198.51.100.10 via ge-0/0/2.0
192.0.2.5/32       *[LDP/9] 00:04:13, metric 1
                    > to 198.51.100.14 via ge-0/0/5.0
192.0.2.7/32       *[RSVP/7/1] 02:38:28, metric 1
                    > to 198.51.100.17 via ge-0/0/1.0, label-switched-path SJ_DA
                      to 198.51.100.10 via ge-0/0/2.0, label-switched-path Bypass->198.51.100.17
```

On CO LER

```
lab@CO> show ldp session
  Address           State        Connection     Hold time
192.0.2.1           Operational  Open             27
192.0.2.3           Operational  Open             25
192.0.2.5           Operational  Open             21
192.0.2.6           Operational  Open             21
lab@CO> show route table inet.3

inet.3: 12 destinations, 20 routes (6 active, 0 holddown, 11 hidden)
+ = Active Route, - = Last Active, * = Both

192.0.2.1/32       *[RSVP/7/1] 00:02:38, metric 1
                    > to 198.51.100.11 via ge-0/0/2.0, label-switched-path CO_CA
                      to 198.51.100.4 via ge-0/0/1.0, label-switched-path Bypass->198.51.100.11
                    [LDP/9] 00:02:37, metric 1
                    > to 198.51.100.11 via ge-0/0/2.0, label-switched-path CO_CA
                      to 198.51.100.4 via ge-0/0/1.0, label-switched-path CO_CA
192.0.2.2/32       *[RSVP/7/1] 06:08:28, metric 2
                    > to 198.51.100.4 via ge-0/0/1.0, label-switched-path CO_AT
                      to 198.51.100.11 via ge-0/0/2.0, label-switched-path Bypass->198.51.100.4
192.0.2.3/32       *[RSVP/7/1] 04:34:49, metric 1
                    > to 198.51.100.4 via ge-0/0/1.0, label-switched-path CO_NE
                      to 198.51.100.11 via ge-0/0/2.0, label-switched-path Bypass->198.51.100.4
                    [LDP/9] 00:05:38, metric 1
```

```
                    > to 198.51.100.4 via ge-0/0/1.0
192.0.2.5/32       *[LDP/9] 00:05:38, metric 1
                    > to 198.51.100.13 via ge-0/0/4.0
192.0.2.6/32       *[RSVP/7/1] 02:39:58, metric 1
                    > to 198.51.100.11 via ge-0/0/2.0, label-switched-path CO_SJ
                      to 198.51.100.4 via ge-0/0/1.0, label-switched-path Bypass->198.51.100.11
                    [LDP/9] 00:05:38, metric 1
                    > to 198.51.100.11 via ge-0/0/2.0
192.0.2.7/32       *[RSVP/7/1] 02:44:16, metric 2
                    > to 198.51.100.4 via ge-0/0/1.0, label-switched-path CO_DA
                      to 198.51.100.11 via ge-0/0/2.0, label-switched-path Bypass->198.51.100.4
```

On NY LER

```
lab@NY> show ldp session
  Address           State        Connection     Hold time
192.0.2.4           Operational  Open             27
192.0.2.6           Operational  Open             27
lab@NY> show route table inet.3


inet.3: 4 destinations, 4 routes (4 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

192.0.2.1/32       *[LDP/9] 04:49:10, metric 1
                    > to 198.51.100.12 via ge-0/0/4.0, Push 300400
192.0.2.3/32       *[LDP/9] 04:49:10, metric 1
                    > to 198.51.100.12 via ge-0/0/4.0, Push 299808
                      to 198.51.100.15 via ge-0/0/5.0, Push 299776
192.0.2.4/32       *[LDP/9] 04:49:10, metric 1
                    > to 198.51.100.12 via ge-0/0/4.0
192.0.2.6/32       *[LDP/9] 04:49:11, metric 1
                    > to 198.51.100.15 via ge-0/0/5.0
```

On NE LER

```
lab@NE> show ldp session
  Address           State        Connection     Hold time
192.0.2.4           Operational  Open             26
192.0.2.6           Operational  Open             21
lab@NE> show route table inet.3


inet.3: 6 destinations, 9 routes (6 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

192.0.2.1/32       *[RSVP/7/1] 05:21:20, metric 2
                    > to 198.51.100.2 via ge-0/0/2.0, label-switched-path NE_CA
                      to 198.51.100.5 via ge-0/0/1.0, label-switched-path Bypass->198.51.100.2
                    [LDP/9] 05:21:34, metric 1
                    > to 198.51.100.5 via ge-0/0/1.0, Push 300400
192.0.2.2/32       *[RSVP/7/1] 1d 02:22:54, metric 1
                    > to 198.51.100.2 via ge-0/0/2.0, label-switched-path NE_AT
                      to 198.51.100.5 via ge-0/0/1.0, label-switched-path Bypass->198.51.100.2
192.0.2.4/32       *[RSVP/7/1] 05:22:13, metric 1
                    > to 198.51.100.5 via ge-0/0/1.0, label-switched-path NE_CO
                      to 198.51.100.2 via ge-0/0/2.0, label-switched-path Bypass->198.51.100.5
                    [LDP/9] 05:21:35, metric 1
```

```
                    > to 198.51.100.5 via ge-0/0/1.0
192.0.2.5/32       *[LDP/9] 05:21:35, metric 1
                    > to 198.51.100.7 via ge-0/0/3.0, Push 299792
                      to 198.51.100.5 via ge-0/0/1.0, Push 299776
192.0.2.6/32       *[RSVP/7/1] 05:22:12, metric 1
                    > to 198.51.100.5 via ge-0/0/1.0, label-switched-path NE_SJ
                      to 198.51.100.2 via ge-0/0/2.0, label-switched-path Bypass->198.51.100.5
                    [LDP/9] 05:21:36, metric 1
                    > to 198.51.100.7 via ge-0/0/3.0
192.0.2.7/32       *[RSVP/7/1] 05:22:12, metric 2
                    > to 198.51.100.5 via ge-0/0/1.0, label-switched-path NE_DA
                      to 198.51.100.2 via ge-0/0/2.0, label-switched-path Bypass->198.51.100.5
```

Customer VRF Implementations

On CA PE

```
lab@CA> show configuration interfaces ge-0/0/7
description “VPLS SITE 1 CUSTOMER C”;
encapsulation ethernet-vpls;
unit 0 {
    family vpls;
}

lab@CA> show configuration interfaces ge-0/0/9.600
description "CUST B SITE 3";
vlan-id 600;
family inet {
    address 100.64.2.0/31;
}

lab@CA> show configuration routing-instances
CUST-B_SITE3 {
    instance-type vrf;
    interface ge-0/0/9.600;
    vrf-target target:100.64.0.0:100;
    protocols {
        bgp {
            group PE-CE {
                type external;
                peer-as 64503;
                neighbor 100.64.2.1;
            }
        }
    }
}
CUST3_SITE1 {
    instance-type vpls;
    interface ge-0/0/7.0;
    vrf-target target:65000:111;
    protocols {
        vpls {
            no-tunnel-services;
            site 1 {
                site-identifier 1;
            }
        }
```

```
    }
}
```

On SJ PE

```
lab@SJ> show configuration interfaces ge-0/0/8
description “VPLS SITE2 CUSTOMER C”;
encapsulation ethernet-vpls;
unit 0 {
    family vpls;
}
lab@SJ> show configuration routing-instances
CUST3_SITE2 {
    instance-type vpls;
    interface ge-0/0/8.0;
    vrf-target target:65000:111;
    protocols {
        vpls {
            no-tunnel-services;
            site 2 {
                site-identifier 2;
            }
        }
    }
}
```

On NY PE

```
lab@NY> show configuration interfaces ge-0/0/6
description "VPLS SITE3 CUSTOMER C";
encapsulation ethernet-vpls;
unit 0 {
    family vpls;
}
lab@NY> show configuration interfaces ge-0/0/9.400
description "CUST B SITE 2";
vlan-id 400;
family inet {
    address 100.64.0.0/31;
}
lab@NY> show configuration routing-instances
CUST-B_SITE2 {
    instance-type vrf;
    interface ge-0/0/9.400;
    vrf-target target:100.64.0.0:100;
    protocols {
        bgp {
            group PE-CE {
                type external;
                peer-as 64502;
                neighbor 100.64.0.1;
            }
        }
    }
}
CUST3_SITE3 {
```

```
    instance-type vpls;
    interface ge-0/0/6.0;
    vrf-target target:65000:111;
    protocols {
        vpls {
            no-tunnel-services;
            site 3 {
                site-identifier 3;
            }
        }
    }
}
```

On CO PE

```
lab@CO> show configuration interfaces ge-0/0/9.200
description "CUST A SITE 1";
vlan-id 200;
family inet {
    address 203.0.113.2/31;
}
lab@CO> show configuration routing-instances
CUST-A_SITE1 {
    instance-type vrf;
    interface ge-0/0/9.200;
    vrf-target target:64510:100;
    protocols {
        bgp {
            group PE-CE {
                type external;
                peer-as 64510;
                neighbor 203.0.113.3;
            }
        }
    }
}
```

On NE PE

```
lab@NE> show configuration interfaces ge-0/0/9
vlan-tagging;
unit 300 {
    description "CUST A SITE2";
    vlan-id 300;
    family inet {
        address 203.0.113.0/31;
    }
}
unit 500 {
    description "CUST B SITE1";
    vlan-id 500;
    family inet {
        address 100.64.1.0/31;
    }
}
lab@NE> show configuration routing-instances
```

```
CUST-A_SITE2 {
    instance-type vrf;
    interface ge-0/0/9.300;
    vrf-import CUST_IMPORT;
    vrf-export CUST_EXPORT;
    protocols {
        bgp {
            group PE-CE {
                type external;
                peer-as 64509;
                neighbor 203.0.113.1;
            }
        }
    }
}
CUST-B_SITE1 {
    instance-type vrf;
    interface ge-0/0/9.500;
    vrf-target target:100.64.0.0:100;
    protocols {
        bgp {
            group PE-CE {
                type external;
                peer-as 64501;
                neighbor 100.64.1.1;
            }
        }
    }
}
```

Verification of VPN label allocations.

On CA PE

```
lab@CA> show route table mpls.0

mpls.0: 22 destinations, 22 routes (22 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

0                  *[MPLS/0] 1d 03:46:01, metric 1
                      Receive
1                  *[MPLS/0] 1d 03:46:01, metric 1
                      Receive
2                  *[MPLS/0] 1d 03:46:01, metric 1
                      Receive
13                 *[MPLS/0] 1d 03:46:01, metric 1
                      Receive
262146             *[VPLS/7] 06:17:52
                    > via lsi.1048579, Pop
262147             *[VPLS/7] 06:17:56
                    > via lsi.1048578, Pop
300480             *[RSVP/7/1] 06:18:21, metric 1
                    > to 198.51.100.19 via ge-0/0/2.0, label-switched-path Bypass->198.51.100.17
300480(S=0)        *[RSVP/7/1] 06:18:21, metric 1
                    > to 198.51.100.19 via ge-0/0/2.0, label-switched-path Bypass->198.51.100.17
300496             *[LDP/9] 06:18:06, metric 1
                    > to 198.51.100.1 via ge-0/0/1.0, label-switched-path CA_CO
```

```
                     to 198.51.100.19 via ge-0/0/2.0, label-switched-path CA_CO
300512            *[LDP/9] 06:18:06, metric 1
                   > to 198.51.100.1 via ge-0/0/1.0, label-switched-path CA_CO
                     to 198.51.100.19 via ge-0/0/2.0, label-switched-path CA_CO
300528            *[LDP/9] 06:17:58, metric 1
                   > to 198.51.100.19 via ge-0/0/2.0, label-switched-path CA_SJ
                     to 198.51.100.1 via ge-0/0/1.0, label-switched-path CA_SJ
300544            *[RSVP/7/1] 06:18:16, metric 1
                   > to 198.51.100.1 via ge-0/0/1.0, label-switched-path Bypass->198.51.100.4
300560            *[RSVP/7/1] 06:18:16, metric 1
                   > to 198.51.100.19 via ge-0/0/2.0, label-switched-path Bypass->198.51.100.11
300576            *[RSVP/7/1] 06:18:16, metric 1
                   > to 198.51.100.19 via ge-0/0/2.0, label-switched-path Bypass->198.51.100.3
300592            *[RSVP/7/1] 06:18:16, metric 1
                   > to 198.51.100.1 via ge-0/0/1.0, label-switched-path Bypass->198.51.100.16
300608            *[RSVP/7/1] 06:18:15, metric 1
                   > to 198.51.100.1 via ge-0/0/1.0, label-switched-path Bypass->198.51.100.10
300624            *[RSVP/7/1] 06:17:55, metric 1
                   > to 198.51.100.1 via ge-0/0/1.0, label-switched-path Bypass->198.51.100.2
300624(S=0)       *[RSVP/7/1] 06:17:55, metric 1
                   > to 198.51.100.1 via ge-0/0/1.0, label-switched-path Bypass->198.51.100.2
300640            *[RSVP/7/1] 06:17:52, metric 1
                   > to 198.51.100.19 via ge-0/0/2.0, label-switched-path Bypass->198.51.100.5
300656            *[VPN/170] 02:09:28
                   > to 100.64.2.1 via ge-0/0/9.600, Pop
lsi.1048578       *[VPLS/7] 06:17:56, metric2 1
                   > to 198.51.100.1 via ge-0/0/1.0, label-switched-path CA_CO
lsi.1048579       *[VPLS/7] 06:17:52, metric2 2
                   > to 198.51.100.19 via ge-0/0/2.0, label-switched-path CA_SJ
                     to 198.51.100.1 via ge-0/0/1.0, label-switched-path Bypass->198.51.100.19
```

On SJ PE

```
lab@SJ> show route table mpls.0

mpls.0: 36 destinations, 36 routes (36 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

0                 *[MPLS/0] 1d 03:29:16, metric 1
                     Receive
1                 *[MPLS/0] 1d 03:29:16, metric 1
                     Receive
2                 *[MPLS/0] 1d 03:29:16, metric 1
                     Receive
13                *[MPLS/0] 1d 03:29:16, metric 1
                     Receive
262145            *[VPLS/7] 06:21:15
                   > via lsi.1048579, Pop
262147            *[VPLS/7] 1d 03:23:52
                   > via lsi.1048577, Pop
299776            *[LDP/9] 06:21:14, metric 1
                   > to 198.51.100.6 via ge-0/0/3.0, Pop
299776(S=0)       *[LDP/9] 06:21:14, metric 1
                   > to 198.51.100.6 via ge-0/0/3.0, Pop
299792            *[LDP/9] 06:21:23, metric 1
                   > to 198.51.100.14 via ge-0/0/5.0, Pop
299792(S=0)       *[LDP/9] 06:21:23, metric 1
                   > to 198.51.100.14 via ge-0/0/5.0, Pop
```

```
299808              *[LDP/9] 06:21:07, metric 1
                     > to 198.51.100.10 via ge-0/0/2.0, Pop
299808(S=0)         *[LDP/9] 06:21:07, metric 1
                     > to 198.51.100.10 via ge-0/0/2.0, Pop
299888              *[RSVP/7/1] 1d 03:23:50, metric 1
                     > to 198.51.100.10 via ge-0/0/2.0, label-switched-path DA_CO
                       to 198.51.100.17 via ge-0/0/1.0, label-switched-path Bypass->198.51.100.10
299888(S=0)         *[RSVP/7/1] 06:21:14, metric 1
                     > to 198.51.100.10 via ge-0/0/2.0, label-switched-path DA_CO
                       to 198.51.100.17 via ge-0/0/1.0, label-switched-path Bypass->198.51.100.10
300000              *[RSVP/7/1] 07:31:07, metric 1
                     > to 198.51.100.17 via ge-0/0/1.0, label-switched-path CO_DA
                       to 198.51.100.10 via ge-0/0/2.0, label-switched-path Bypass->198.51.100.17
300000(S=0)         *[RSVP/7/1] 06:21:38, metric 1
                     > to 198.51.100.17 via ge-0/0/1.0, label-switched-path CO_DA
                       to 198.51.100.10 via ge-0/0/2.0, label-switched-path Bypass->198.51.100.17
300224              *[RSVP/7/1] 06:22:44, metric 1
                     > to 198.51.100.10 via ge-0/0/2.0, label-switched-path DA_NE
                       to 198.51.100.17 via ge-0/0/1.0, label-switched-path Bypass->198.51.100.10
300240              *[RSVP/7/1] 06:22:44, metric 1
                     > to 198.51.100.10 via ge-0/0/2.0, label-switched-path DA_AT
                       to 198.51.100.17 via ge-0/0/1.0, label-switched-path Bypass->198.51.100.10
300256              *[RSVP/7/1] 06:22:43, metric 1
                     > to 198.51.100.17 via ge-0/0/1.0, label-switched-path NE_DA
                       to 198.51.100.10 via ge-0/0/2.0, label-switched-path Bypass->198.51.100.17
300256(S=0)         *[RSVP/7/1] 06:21:37, metric 1
                     > to 198.51.100.17 via ge-0/0/1.0, label-switched-path NE_DA
                       to 198.51.100.10 via ge-0/0/2.0, label-switched-path Bypass->198.51.100.17
300288              *[RSVP/7/1] 06:22:17, metric 1
                     > to 198.51.100.17 via ge-0/0/1.0, label-switched-path AT_DA
                       to 198.51.100.10 via ge-0/0/2.0, label-switched-path Bypass->198.51.100.17
300288(S=0)         *[RSVP/7/1] 06:21:22, metric 1
                     > to 198.51.100.17 via ge-0/0/1.0, label-switched-path AT_DA
                       to 198.51.100.10 via ge-0/0/2.0, label-switched-path Bypass->198.51.100.17
300400              *[RSVP/7/1] 06:21:45, metric 1
                     > to 198.51.100.17 via ge-0/0/1.0, label-switched-path CO_CA
                       to 198.51.100.10 via ge-0/0/2.0, label-switched-path Bypass->198.51.100.17
300432              *[LDP/9] 06:21:40, metric 1
                     > to 198.51.100.17 via ge-0/0/1.0, label-switched-path SJ_CA
                       to 198.51.100.10 via ge-0/0/2.0, label-switched-path SJ_CA
300448              *[RSVP/7/1] 06:21:39, metric 1
                     > to 198.51.100.10 via ge-0/0/2.0, label-switched-path Bypass->198.51.100.3
300464              *[RSVP/7/1] 06:21:39, metric 1
                     > to 198.51.100.17 via ge-0/0/1.0, label-switched-path Bypass->198.51.100.0
300480              *[RSVP/7/1] 06:21:39, metric 1
                     > to 198.51.100.17 via ge-0/0/1.0, label-switched-path Bypass->198.51.100.4
300496              *[RSVP/7/1] 06:21:39, metric 1
                     > to 198.51.100.10 via ge-0/0/2.0, label-switched-path Bypass->198.51.100.18
300512              *[RSVP/7/1] 06:21:31, metric 1
                     > to 198.51.100.17 via ge-0/0/1.0, label-switched-path Bypass->198.51.100.19
300512(S=0)         *[RSVP/7/1] 06:21:31, metric 1
                     > to 198.51.100.17 via ge-0/0/1.0, label-switched-path Bypass->198.51.100.19
300528              *[RSVP/7/1] 06:21:31, metric 1
                     > to 198.51.100.10 via ge-0/0/2.0, label-switched-path Bypass->198.51.100.1
300544              *[RSVP/7/1] 06:21:17, metric 1
                     > to 198.51.100.17 via ge-0/0/1.0, label-switched-path Bypass->198.51.100.2
300560              *[RSVP/7/1] 06:21:16, metric 1
                     > to 198.51.100.10 via ge-0/0/2.0, label-switched-path Bypass->198.51.100.5
```

```
300560(S=0)        *[RSVP/7/1] 06:21:16, metric 1
                    > to 198.51.100.10 via ge-0/0/2.0, label-switched-path Bypass->198.51.100.5
lsi.1048579        *[VPLS/7] 06:21:15, metric2 2
                    > to 198.51.100.17 via ge-0/0/1.0, label-switched-path SJ_CA
                      to 198.51.100.10 via ge-0/0/2.0, label-switched-path Bypass->198.51.100.17
lsi.1048577        *[VPLS/7] 1d 03:23:52, metric2 1
                    > to 198.51.100.14 via ge-0/0/5.0, Push 262146
```

On NY PE

```
lab@NY> show route table mpls.0

mpls.0: 15 destinations, 15 routes (15 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

0                  *[MPLS/0] 1d 03:29:55, metric 1
                      Receive
1                  *[MPLS/0] 1d 03:29:55, metric 1
                      Receive
2                  *[MPLS/0] 1d 03:29:55, metric 1
                      Receive
13                 *[MPLS/0] 1d 03:29:55, metric 1
                      Receive
262145             *[VPLS/7] 06:23:09
                    > via lsi.1048578, Pop
262146             *[VPLS/7] 1d 03:25:30
                    > via lsi.1048577, Pop
299872             *[LDP/9] 06:23:14, metric 1
                    > to 198.51.100.15 via ge-0/0/5.0, Pop
299872(S=0)        *[LDP/9] 06:23:14, metric 1
                    > to 198.51.100.15 via ge-0/0/5.0, Pop
299888             *[LDP/9] 06:23:13, metric 1
                    > to 198.51.100.12 via ge-0/0/4.0, Swap 299808
                      to 198.51.100.15 via ge-0/0/5.0, Swap 299776
299904             *[LDP/9] 06:23:13, metric 1
                    > to 198.51.100.12 via ge-0/0/4.0, Pop
299904(S=0)        *[LDP/9] 06:23:13, metric 1
                    > to 198.51.100.12 via ge-0/0/4.0, Pop
299920             *[LDP/9] 06:23:13, metric 1
                    > to 198.51.100.12 via ge-0/0/4.0, Swap 300400
299936             *[VPN/170] 02:17:06
                    > to 100.64.0.1 via ge-0/0/9.400, Pop
lsi.1048578        *[VPLS/7] 06:23:09, metric2 1
                    > to 198.51.100.12 via ge-0/0/4.0, Push 262147, Push 300400(top)
lsi.1048577        *[VPLS/7] 1d 03:25:30, metric2 1
                    > to 198.51.100.15 via ge-0/0/5.0, Push 262147
```

On CO PE

```
lab@CO> show route table mpls.0

mpls.0: 34 destinations, 34 routes (34 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

0                  *[MPLS/0] 1d 03:31:30, metric 1
                      Receive
```

```
1                  *[MPLS/0] 1d 03:31:30, metric 1
                      Receive
2                  *[MPLS/0] 1d 03:31:30, metric 1
                      Receive
13                 *[MPLS/0] 1d 03:31:30, metric 1
                      Receive
299776             *[LDP/9] 06:26:45, metric 1
                    > to 198.51.100.13 via ge-0/0/4.0, Pop
299776(S=0)        *[LDP/9] 06:26:45, metric 1
                    > to 198.51.100.13 via ge-0/0/4.0, Pop
299792             *[LDP/9] 06:26:27, metric 1
                    > to 198.51.100.11 via ge-0/0/2.0, Pop
299792(S=0)        *[LDP/9] 06:26:27, metric 1
                    > to 198.51.100.11 via ge-0/0/2.0, Pop
299808             *[LDP/9] 06:26:27, metric 1
                    > to 198.51.100.4 via ge-0/0/1.0, Pop
299808(S=0)        *[LDP/9] 06:26:27, metric 1
                    > to 198.51.100.4 via ge-0/0/1.0, Pop
300112             *[RSVP/7/1] 07:36:20, metric 1
                    > to 198.51.100.4 via ge-0/0/1.0, label-switched-path SJ_AT
                      to 198.51.100.11 via ge-0/0/2.0, label-switched-path Bypass->198.51.100.4
300256             *[RSVP/7/1] 06:28:08, metric 1
                    > to 198.51.100.4 via ge-0/0/1.0, label-switched-path SJ_NE
                      to 198.51.100.11 via ge-0/0/2.0, label-switched-path Bypass->198.51.100.4
300256(S=0)        *[RSVP/7/1] 06:26:37, metric 1
                    > to 198.51.100.4 via ge-0/0/1.0, label-switched-path SJ_NE
                      to 198.51.100.11 via ge-0/0/2.0, label-switched-path Bypass->198.51.100.4
300272             *[RSVP/7/1] 06:28:08, metric 1
                    > to 198.51.100.4 via ge-0/0/1.0, label-switched-path DA_NE
                      to 198.51.100.11 via ge-0/0/2.0, label-switched-path Bypass->198.51.100.4
300272(S=0)        *[RSVP/7/1] 06:26:32, metric 1
                    > to 198.51.100.4 via ge-0/0/1.0, label-switched-path DA_NE
                      to 198.51.100.11 via ge-0/0/2.0, label-switched-path Bypass->198.51.100.4
300288             *[RSVP/7/1] 06:28:08, metric 1
                    > to 198.51.100.4 via ge-0/0/1.0, label-switched-path DA_AT
                      to 198.51.100.11 via ge-0/0/2.0, label-switched-path Bypass->198.51.100.4
300304             *[RSVP/7/1] 06:28:06, metric 1
                    > to 198.51.100.11 via ge-0/0/2.0, label-switched-path NE_SJ
                      to 198.51.100.4 via ge-0/0/1.0, label-switched-path Bypass->198.51.100.11
300304(S=0)        *[RSVP/7/1] 06:26:27, metric 1
                    > to 198.51.100.11 via ge-0/0/2.0, label-switched-path NE_SJ
                      to 198.51.100.4 via ge-0/0/1.0, label-switched-path Bypass->198.51.100.11
300320             *[RSVP/7/1] 06:28:06, metric 1
                    > to 198.51.100.11 via ge-0/0/2.0, label-switched-path NE_DA
                      to 198.51.100.4 via ge-0/0/1.0, label-switched-path Bypass->198.51.100.11
300336             *[RSVP/7/1] 06:27:45, metric 1
                    > to 198.51.100.11 via ge-0/0/2.0, label-switched-path AT_SJ
                      to 198.51.100.4 via ge-0/0/1.0, label-switched-path Bypass->198.51.100.11
300336(S=0)        *[RSVP/7/1] 06:26:59, metric 1
                    > to 198.51.100.11 via ge-0/0/2.0, label-switched-path AT_SJ
                      to 198.51.100.4 via ge-0/0/1.0, label-switched-path Bypass->198.51.100.11
300352             *[RSVP/7/1] 06:27:39, metric 1
                    > to 198.51.100.11 via ge-0/0/2.0, label-switched-path AT_DA
                      to 198.51.100.4 via ge-0/0/1.0, label-switched-path Bypass->198.51.100.11
300384             *[RSVP/7/1] 06:27:07, metric 1
                    > to 198.51.100.4 via ge-0/0/1.0, label-switched-path Bypass->198.51.100.17
300400             *[LDP/9] 06:27:00, metric 1
                    > to 198.51.100.11 via ge-0/0/2.0, label-switched-path CO_CA
```

```
                     to 198.51.100.4 via ge-0/0/1.0, label-switched-path CO_CA
300416             *[RSVP/7/1] 06:27:03, metric 1
                    > to 198.51.100.11 via ge-0/0/2.0, label-switched-path Bypass->198.51.100.16
300416(S=0)        *[RSVP/7/1] 06:27:03, metric 1
                    > to 198.51.100.11 via ge-0/0/2.0, label-switched-path Bypass->198.51.100.16
300432             *[RSVP/7/1] 06:27:03, metric 1
                    > to 198.51.100.4 via ge-0/0/1.0, label-switched-path Bypass->198.51.100.3
300432(S=0)        *[RSVP/7/1] 06:27:03, metric 1
                    > to 198.51.100.4 via ge-0/0/1.0, label-switched-path Bypass->198.51.100.3
300448             *[RSVP/7/1] 06:27:02, metric 1
                    > to 198.51.100.11 via ge-0/0/2.0, label-switched-path Bypass->198.51.100.0
300464             *[RSVP/7/1] 06:27:02, metric 1
                    > to 198.51.100.4 via ge-0/0/1.0, label-switched-path Bypass->198.51.100.18
300480             *[RSVP/7/1] 06:26:55, metric 1
                    > to 198.51.100.11 via ge-0/0/2.0, label-switched-path Bypass->198.51.100.19
300496             *[RSVP/7/1] 06:26:55, metric 1
                    > to 198.51.100.4 via ge-0/0/1.0, label-switched-path Bypass->198.51.100.1
300512             *[RSVP/7/1] 06:26:40, metric 1
                    > to 198.51.100.11 via ge-0/0/2.0, label-switched-path Bypass->198.51.100.2
300528             *[VPN/170] 01:59:57
                    > to 203.0.113.3 via ge-0/0/9.200, Pop
```

On NE PE

```
lab@NE> show route table mpls.0

mpls.0: 34 destinations, 34 routes (34 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

0                  *[MPLS/0] 1d 03:36:24, metric 1
                      Receive
1                  *[MPLS/0] 1d 03:36:24, metric 1
                      Receive
2                  *[MPLS/0] 1d 03:36:24, metric 1
                      Receive
13                 *[MPLS/0] 1d 03:36:24, metric 1
                      Receive
299840             *[VPN/170] 1d 03:30:05
                    > to 203.0.113.1 via ge-0/0/9.300, Pop
300064             *[RSVP/7/1] 1d 03:29:41, metric 1
                    > to 198.51.100.5 via ge-0/0/1.0, label-switched-path AT_CO
                      to 198.51.100.2 via ge-0/0/2.0, label-switched-path Bypass->198.51.100.5
300064(S=0)        *[RSVP/7/1] 06:27:36, metric 1
                    > to 198.51.100.5 via ge-0/0/1.0, label-switched-path AT_CO
                      to 198.51.100.2 via ge-0/0/2.0, label-switched-path Bypass->198.51.100.5
300256             *[RSVP/7/1] 07:37:20, metric 1
                    > to 198.51.100.2 via ge-0/0/2.0, label-switched-path SJ_AT
                      to 198.51.100.5 via ge-0/0/1.0, label-switched-path Bypass->198.51.100.2
300256(S=0)        *[RSVP/7/1] 06:27:56, metric 1
                    > to 198.51.100.2 via ge-0/0/2.0, label-switched-path SJ_AT
                      to 198.51.100.5 via ge-0/0/1.0, label-switched-path Bypass->198.51.100.2
300416             *[RSVP/7/1] 06:29:12, metric 1
                    > to 198.51.100.2 via ge-0/0/2.0, label-switched-path DA_AT
                      to 198.51.100.5 via ge-0/0/1.0, label-switched-path Bypass->198.51.100.2
300416(S=0)        *[RSVP/7/1] 06:27:46, metric 1
                    > to 198.51.100.2 via ge-0/0/2.0, label-switched-path DA_AT
                      to 198.51.100.5 via ge-0/0/1.0, label-switched-path Bypass->198.51.100.2
```

```
300432              *[RSVP/7/1] 06:29:12, metric 1
                     > to 198.51.100.2 via ge-0/0/2.0, label-switched-path CO_AT
                       to 198.51.100.5 via ge-0/0/1.0, label-switched-path Bypass->198.51.100.2
300432(S=0)         *[RSVP/7/1] 06:27:42, metric 1
                     > to 198.51.100.2 via ge-0/0/2.0, label-switched-path CO_AT
                       to 198.51.100.5 via ge-0/0/1.0, label-switched-path Bypass->198.51.100.2
300448              *[RSVP/7/1] 06:28:49, metric 1
                     > to 198.51.100.5 via ge-0/0/1.0, label-switched-path AT_SJ
                       to 198.51.100.2 via ge-0/0/2.0, label-switched-path Bypass->198.51.100.5
300464              *[RSVP/7/1] 06:28:43, metric 1
                     > to 198.51.100.5 via ge-0/0/1.0, label-switched-path AT_DA
                       to 198.51.100.2 via ge-0/0/2.0, label-switched-path Bypass->198.51.100.5
300480              *[RSVP/7/1] 06:28:34, metric 1
                     > to 198.51.100.5 via ge-0/0/1.0, label-switched-path CA_CO
                       to 198.51.100.2 via ge-0/0/2.0, label-switched-path Bypass->198.51.100.5
300480(S=0)         *[RSVP/7/1] 06:27:40, metric 1
                     > to 198.51.100.5 via ge-0/0/1.0, label-switched-path CA_CO
                       to 198.51.100.2 via ge-0/0/2.0, label-switched-path Bypass->198.51.100.5
300496              *[RSVP/7/1] 06:28:34, metric 1
                     > to 198.51.100.2 via ge-0/0/2.0, label-switched-path Bypass->198.51.100.17
300512              *[LDP/9] 06:28:33, metric 1
                     > to 198.51.100.7 via ge-0/0/3.0, Pop
300512(S=0)         *[LDP/9] 06:28:33, metric 1
                     > to 198.51.100.7 via ge-0/0/3.0, Pop
300528              *[LDP/9] 06:28:32, metric 1
                     > to 198.51.100.7 via ge-0/0/3.0, Swap 299792
                       to 198.51.100.5 via ge-0/0/1.0, Swap 299776
300544              *[LDP/9] 06:28:32, metric 1
                     > to 198.51.100.5 via ge-0/0/1.0, Pop
300544(S=0)         *[LDP/9] 06:28:32, metric 1
                     > to 198.51.100.5 via ge-0/0/1.0, Pop
300560              *[LDP/9] 06:28:31, metric 1
                     > to 198.51.100.5 via ge-0/0/1.0, Swap 300400
300576              *[RSVP/7/1] 06:28:30, metric 1
                     > to 198.51.100.5 via ge-0/0/1.0, label-switched-path Bypass->198.51.100.16
300592              *[RSVP/7/1] 06:28:30, metric 1
                     > to 198.51.100.2 via ge-0/0/2.0, label-switched-path Bypass->198.51.100.18
300608              *[RSVP/7/1] 06:28:30, metric 1
                     > to 198.51.100.5 via ge-0/0/1.0, label-switched-path Bypass->198.51.100.0
300624              *[RSVP/7/1] 06:28:30, metric 1
                     > to 198.51.100.2 via ge-0/0/2.0, label-switched-path Bypass->198.51.100.11
300640              *[RSVP/7/1] 06:28:29, metric 1
                     > to 198.51.100.5 via ge-0/0/1.0, label-switched-path Bypass->198.51.100.10
300640(S=0)         *[RSVP/7/1] 06:28:29, metric 1
                     > to 198.51.100.5 via ge-0/0/1.0, label-switched-path Bypass->198.51.100.10
300656              *[RSVP/7/1] 06:28:22, metric 1
                     > to 198.51.100.2 via ge-0/0/2.0, label-switched-path Bypass->198.51.100.1
300656(S=0)         *[RSVP/7/1] 06:28:22, metric 1
                     > to 198.51.100.2 via ge-0/0/2.0, label-switched-path Bypass->198.51.100.1
300672              *[RSVP/7/1] 06:28:22, metric 1
                     > to 198.51.100.5 via ge-0/0/1.0, label-switched-path Bypass->198.51.100.19
300688              *[VPN/170] 02:22:15
                     > to 100.64.1.1 via ge-0/0/9.500, Pop
```

Verification of Customer Site Reachability

Customer A

On CO PE

```
lab@CO> show route table CUST-A_SITE1.inet.0

CUST-A_SITE1.inet.0: 11 destinations, 11 routes (11 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

100.64.11.0/24     *[BGP/170] 02:03:10, localpref 100
                      AS path: 64510 I
                    > to 203.0.113.3 via ge-0/0/9.200
100.64.12.0/24     *[BGP/170] 02:03:10, localpref 100
                      AS path: 64510 I
                    > to 203.0.113.3 via ge-0/0/9.200
100.64.13.0/24     *[BGP/170] 02:03:10, localpref 100
                      AS path: 64510 I
                    > to 203.0.113.3 via ge-0/0/9.200
100.64.14.0/24     *[BGP/170] 02:03:10, localpref 100
                      AS path: 64510 I
                    > to 203.0.113.3 via ge-0/0/9.200
100.64.21.0/24     *[BGP/170] 02:03:20, localpref 100, from 192.0.2.3
                      AS path: 64509 I
                    > to 198.51.100.4 via ge-0/0/1.0, label-switched-path CO_NE
                      to 198.51.100.11 via ge-0/0/2.0, label-switched-path Bypass->198.51.100.4
100.64.22.0/24     *[BGP/170] 02:03:20, localpref 100, from 192.0.2.3
                      AS path: 64509 I
                    > to 198.51.100.4 via ge-0/0/1.0, label-switched-path CO_NE
                      to 198.51.100.11 via ge-0/0/2.0, label-switched-path Bypass->198.51.100.4
100.64.23.0/24     *[BGP/170] 02:03:20, localpref 100, from 192.0.2.3
                      AS path: 64509 I
                    > to 198.51.100.4 via ge-0/0/1.0, label-switched-path CO_NE
                      to 198.51.100.11 via ge-0/0/2.0, label-switched-path Bypass->198.51.100.4
100.64.24.0/24     *[BGP/170] 02:03:20, localpref 100, from 192.0.2.3
                      AS path: 64509 I
                    > to 198.51.100.4 via ge-0/0/1.0, label-switched-path CO_NE
                      to 198.51.100.11 via ge-0/0/2.0, label-switched-path Bypass->198.51.100.4
203.0.113.0/31     *[BGP/170] 02:03:19, localpref 100, from 192.0.2.3
                      AS path: I
                    > to 198.51.100.4 via ge-0/0/1.0, label-switched-path CO_NE
                      to 198.51.100.11 via ge-0/0/2.0, label-switched-path Bypass->198.51.100.4
203.0.113.2/31     *[Direct/0] 02:03:14
                    > via ge-0/0/9.200
203.0.113.2/32     *[Local/0] 02:03:14
                      Local via ge-0/0/9.200
```

On NE PE

```
lab@NE> show route table CUST-A_SITE2.inet.0

CUST-A_SITE2.inet.0: 11 destinations, 11 routes (11 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both
```

```
100.64.11.0/24     *[BGP/170] 02:04:41, localpref 100, from 192.0.2.4
                      AS path: 64510 I
                    > to 198.51.100.5 via ge-0/0/1.0, label-switched-path NE_CO
                      to 198.51.100.2 via ge-0/0/2.0, label-switched-path Bypass->198.51.100.5
100.64.12.0/24     *[BGP/170] 02:04:41, localpref 100, from 192.0.2.4
                      AS path: 64510 I
                    > to 198.51.100.5 via ge-0/0/1.0, label-switched-path NE_CO
                      to 198.51.100.2 via ge-0/0/2.0, label-switched-path Bypass->198.51.100.5
100.64.13.0/24     *[BGP/170] 02:04:41, localpref 100, from 192.0.2.4
                      AS path: 64510 I
                    > to 198.51.100.5 via ge-0/0/1.0, label-switched-path NE_CO
                      to 198.51.100.2 via ge-0/0/2.0, label-switched-path Bypass->198.51.100.5
100.64.14.0/24     *[BGP/170] 02:04:41, localpref 100, from 192.0.2.4
                      AS path: 64510 I
                    > to 198.51.100.5 via ge-0/0/1.0, label-switched-path NE_CO
                      to 198.51.100.2 via ge-0/0/2.0, label-switched-path Bypass->198.51.100.5
100.64.21.0/24     *[BGP/170] 1d 03:33:52, localpref 100
                      AS path: 64509 I
                    > to 203.0.113.1 via ge-0/0/9.300
100.64.22.0/24     *[BGP/170] 1d 03:33:52, localpref 100
                      AS path: 64509 I
                    > to 203.0.113.1 via ge-0/0/9.300
100.64.23.0/24     *[BGP/170] 1d 03:33:52, localpref 100
                      AS path: 64509 I
                    > to 203.0.113.1 via ge-0/0/9.300
100.64.24.0/24     *[BGP/170] 1d 03:33:52, localpref 100
                      AS path: 64509 I
                    > to 203.0.113.1 via ge-0/0/9.300
203.0.113.0/31     *[Direct/0] 1d 03:37:00
                    > via ge-0/0/9.300
203.0.113.0/32     *[Local/0] 1d 03:37:27
                      Local via ge-0/0/9.300
203.0.113.2/31     *[BGP/170] 02:04:41, localpref 100, from 192.0.2.4
                      AS path: I
                    > to 198.51.100.5 via ge-0/0/1.0, label-switched-path NE_CO
                      to 198.51.100.2 via ge-0/0/2.0, label-switched-path Bypass->198.51.100.5
```

Customer Router

```
lab@CUSTA_SITE1>ping 100.64.24.1 rapid count 5
PING 100.64.24.1 (100.64.24.1): 56 data bytes
!!!!!
--- 100.64.24.1 ping statistics ---
5 packets transmitted, 5 packets received, 0% packet loss
round-trip min/avg/max/stddev = 39.115/44.197/48.891/4.050 ms

lab@CUSTA_SITE1>ping 100.64.21.1 rapid count 5
PING 100.64.21.1 (100.64.21.1): 56 data bytes
!!!!!
--- 100.64.21.1 ping statistics ---
5 packets transmitted, 5 packets received, 0% packet loss
round-trip min/avg/max/stddev = 39.394/51.435/73.669/11.661 ms
```

Customer B

On NE PE

```
lab@NE> show route table CUST-B_SITE1.inet.0

CUST-B_SITE1.inet.0: 10 destinations, 10 routes (10 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

100.64.0.0/31      *[BGP/170] 02:34:01, localpref 100, from 192.0.2.5
                      AS path: I
                    > to 198.51.100.7 via ge-0/0/3.0, Push 299936, Push 299792(top)
                      to 198.51.100.5 via ge-0/0/1.0, Push 299936, Push 299776(top)
100.64.1.0/31      *[Direct/0] 1d 03:45:34
                    > via ge-0/0/9.500
100.64.1.0/32      *[Local/0] 1d 03:46:01
                      Local via ge-0/0/9.500
100.64.2.0/31      *[BGP/170] 02:31:33, localpref 100, from 192.0.2.1
                      AS path: I
                    > to 198.51.100.2 via ge-0/0/2.0, label-switched-path NE_CA
                      to 198.51.100.5 via ge-0/0/1.0, label-switched-path Bypass->198.51.100.2
100.64.11.0/24     *[BGP/170] 02:34:06, localpref 100
                      AS path: 64501 I
                    > to 100.64.1.1 via ge-0/0/9.500
100.64.12.0/24     *[BGP/170] 02:34:06, localpref 100
                      AS path: 64501 I
                    > to 100.64.1.1 via ge-0/0/9.500
100.64.13.0/24     *[BGP/170] 02:34:06, localpref 100
                      AS path: 64501 I
                    > to 100.64.1.1 via ge-0/0/9.500
100.64.14.0/24     *[BGP/170] 02:34:06, localpref 100
                      AS path: 64501 I
                    > to 100.64.1.1 via ge-0/0/9.500
100.64.122.0/24    *[BGP/170] 02:34:01, localpref 100, from 192.0.2.5
                      AS path: 64502 I
                      to 198.51.100.7 via ge-0/0/3.0, Push 299936, Push 299792(top)
                    > to 198.51.100.5 via ge-0/0/1.0, Push 299936, Push 299776(top)
100.64.133.0/24    *[BGP/170] 02:31:33, localpref 100, from 192.0.2.1
                      AS path: 64503 I
                    > to 198.51.100.2 via ge-0/0/2.0, label-switched-path NE_CA
                      to 198.51.100.5 via ge-0/0/1.0, label-switched-path Bypass->198.51.100.2
```

On NY PE

```
lab@NY> show route table CUST-B_SITE2.inet.0

CUST-B_SITE2.inet.0: 10 destinations, 10 routes (10 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

100.64.0.0/31      *[Direct/0] 1d 03:44:02
                    > via ge-0/0/9.400
100.64.0.0/32      *[Local/0] 1d 03:44:14
                      Local via ge-0/0/9.400
100.64.1.0/31      *[BGP/170] 02:34:38, localpref 100, from 192.0.2.3
                      AS path: I
                    > to 198.51.100.12 via ge-0/0/4.0, Push 300688, Push 299808(top)
```

```
                    to 198.51.100.15 via ge-0/0/5.0, Push 300688, Push 299776(top)
100.64.2.0/31      *[BGP/170] 02:32:05, localpref 100, from 192.0.2.1
                      AS path: I
                    > to 198.51.100.12 via ge-0/0/4.0, Push 300656, Push 300400(top)
100.64.11.0/24     *[BGP/170] 02:34:39, localpref 100, from 192.0.2.3
                      AS path: 64501 I
                      to 198.51.100.12 via ge-0/0/4.0, Push 300688, Push 299808(top)
                    > to 198.51.100.15 via ge-0/0/5.0, Push 300688, Push 299776(top)
100.64.12.0/24     *[BGP/170] 02:34:39, localpref 100, from 192.0.2.3
                      AS path: 64501 I
                      to 198.51.100.12 via ge-0/0/4.0, Push 300688, Push 299808(top)
                    > to 198.51.100.15 via ge-0/0/5.0, Push 300688, Push 299776(top)
100.64.13.0/24     *[BGP/170] 02:34:39, localpref 100, from 192.0.2.3
                      AS path: 64501 I
                      to 198.51.100.12 via ge-0/0/4.0, Push 300688, Push 299808(top)
                    > to 198.51.100.15 via ge-0/0/5.0, Push 300688, Push 299776(top)
100.64.14.0/24     *[BGP/170] 02:34:39, localpref 100, from 192.0.2.3
                      AS path: 64501 I
                      to 198.51.100.12 via ge-0/0/4.0, Push 300688, Push 299808(top)
                    > to 198.51.100.15 via ge-0/0/5.0, Push 300688, Push 299776(top)
100.64.122.0/24    *[BGP/170] 02:34:34, localpref 100
                      AS path: 64502 I
                    > to 100.64.0.1 via ge-0/0/9.400
100.64.133.0/24    *[BGP/170] 02:32:06, localpref 100, from 192.0.2.1
                      AS path: 64503 I
                    > to 198.51.100.12 via ge-0/0/4.0, Push 300656, Push 300400(top)
```

On CA PE

```
lab@CA> show route table CUST-B_SITE3.inet.0

CUST-B_SITE3.inet.0: 10 destinations, 10 routes (10 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

100.64.0.0/31      *[BGP/170] 02:35:01, localpref 100, from 192.0.2.5
                      AS path: I
                    > to 198.51.100.1 via ge-0/0/1.0, label-switched-path CA_CO
100.64.1.0/31      *[BGP/170] 02:35:06, localpref 100, from 192.0.2.3
                      AS path: I
                    > to 198.51.100.1 via ge-0/0/1.0, label-switched-path CA_NE
                      to 198.51.100.19 via ge-0/0/2.0, label-switched-path Bypass->198.51.100.1
100.64.2.0/31      *[Direct/0] 02:32:43
                    > via ge-0/0/9.600
100.64.2.0/32      *[Local/0] 02:38:04
                      Local via ge-0/0/9.600
100.64.11.0/24     *[BGP/170] 02:35:06, localpref 100, from 192.0.2.3
                      AS path: 64501 I
                    > to 198.51.100.1 via ge-0/0/1.0, label-switched-path CA_NE
                      to 198.51.100.19 via ge-0/0/2.0, label-switched-path Bypass->198.51.100.1
100.64.12.0/24     *[BGP/170] 02:35:06, localpref 100, from 192.0.2.3
                      AS path: 64501 I
                    > to 198.51.100.1 via ge-0/0/1.0, label-switched-path CA_NE
                      to 198.51.100.19 via ge-0/0/2.0, label-switched-path Bypass->198.51.100.1
100.64.13.0/24     *[BGP/170] 02:35:06, localpref 100, from 192.0.2.3
                      AS path: 64501 I
                    > to 198.51.100.1 via ge-0/0/1.0, label-switched-path CA_NE
                      to 198.51.100.19 via ge-0/0/2.0, label-switched-path Bypass->198.51.100.1
```

```
100.64.14.0/24     *[BGP/170] 02:35:06, localpref 100, from 192.0.2.3
                      AS path: 64501 I
                    > to 198.51.100.1 via ge-0/0/1.0, label-switched-path CA_NE
                      to 198.51.100.19 via ge-0/0/2.0, label-switched-path Bypass->198.51.100.1
100.64.122.0/24    *[BGP/170] 02:35:01, localpref 100, from 192.0.2.5
                      AS path: 64502 I
                    > to 198.51.100.1 via ge-0/0/1.0, label-switched-path CA_CO
100.64.133.0/24    *[BGP/170] 02:32:33, localpref 100
                      AS path: 64503 I
                    > to 100.64.2.1 via ge-0/0/9.600
```

Customer Router

Site 1

```
lab@CUSTB_SITE1>ping 100.64.122.1 count 5 rapid
PING 100.64.122.1 (100.64.122.1): 56 data bytes
!!!!!
--- 100.64.122.1 ping statistics ---
5 packets transmitted, 5 packets received, 0% packet loss
round-trip min/avg/max/stddev = 49.282/82.647/156.908/38.201 ms

lab@CUSTB_SITE1>ping 100.64.133.1 count 5 rapid
PING 100.64.133.1 (100.64.133.1): 56 data bytes
!!!!!
--- 100.64.133.1 ping statistics ---
5 packets transmitted, 5 packets received, 0% packet loss
round-trip min/avg/max/stddev = 33.338/53.855/90.055/19.180 ms
```

Site 2

```
lab@CUSTB_SITE2>ping 100.64.11.1 count 5 rapid
PING 100.64.11.1 (100.64.11.1): 56 data bytes
!!!!!
--- 100.64.11.1 ping statistics ---
5 packets transmitted, 5 packets received, 0% packet loss
round-trip min/avg/max/stddev = 48.359/66.086/87.872/13.350 ms

lab@CUSTB_SITE2>ping 100.64.133.1 count 5 rapid
PING 100.64.133.1 (100.64.133.1): 56 data bytes
!!!!!
--- 100.64.133.1 ping statistics ---
5 packets transmitted, 5 packets received, 0% packet loss
round-trip min/avg/max/stddev = 67.801/80.094/94.643/10.377 ms
```

Customer C

ON CA PE

```
lab@CA> show vpls connections extensive
Layer-2 VPN connections:

Legend for connection status (St)
EI -- encapsulation invalid      NC -- interface encapsulation not CCC/TCC/VPLS
```

```
EM -- encapsulation mismatch     WE -- interface and instance encaps not same
VC-Dn -- Virtual circuit down    NP -- interface hardware not present
CM -- control-word mismatch      -> -- only outbound connection is up
CN -- circuit not provisioned    <- -- only inbound connection is up
OR -- out of range               Up -- operational
OL -- no outgoing label          Dn -- down
LD -- local site signaled down   CF -- call admission control failure
RD -- remote site signaled down  SC -- local and remote site ID collision
LN -- local site not designated  LM -- local site ID not minimum designated
RN -- remote site not designated RM -- remote site ID not minimum designated
XX -- unknown connection status  IL -- no incoming label
MM -- MTU mismatch               MI -- Mesh-Group ID not available
BK -- Backup connection          ST -- Standby connection
PF -- Profile parse failure      PB -- Profile busy
RS -- remote site standby        SN -- Static Neighbor
VM -- VLAN ID mismatch

Legend for interface status
Up -- operational
Dn -- down

Instance: CUST3_SITE1
  Local site: 1 (1)
    Number of local interfaces: 1
    Number of local interfaces up: 1
    IRB interface present: no
    ge-0/0/7.0
    lsi.1048579         2         Intf - vpls CUST3_SITE1 local site 1 remote site 2
    lsi.1048578         3         Intf - vpls CUST3_SITE1 local site 1 remote site 3
    Label-base        Offset     Size  Range     Preference
    262145            1          8      8         100
    connection-site           Type  St     Time last up          # Up trans
    2                         rmt   Up     Aug 21 09:04:10 2018           1
      Remote PE: 192.0.2.6, Negotiated control-word: No
      Incoming label: 262146, Outgoing label: 262145
      Local interface: lsi.1048579, Status: Up, Encapsulation: VPLS
        Description: Intf - vpls CUST3_SITE1 local site 1 remote site 2
    Connection History:
        Aug 21 09:04:10 2018  status update timer
        Aug 21 09:04:10 2018  loc intf up                  lsi.1048579
        Aug 21 09:04:10 2018  PE route changed
        Aug 21 09:04:10 2018  Out lbl Update                    262145
        Aug 21 09:04:10 2018  In lbl Update                     262146
        Aug 21 09:04:10 2018  loc intf down
    3                         rmt   Up     Aug 21 09:04:06 2018           1
      Remote PE: 192.0.2.5, Negotiated control-word: No
      Incoming label: 262147, Outgoing label: 262145
      Local interface: lsi.1048578, Status: Up, Encapsulation: VPLS
        Description: Intf - vpls CUST3_SITE1 local site 1 remote site 3
    Connection History:
        Aug 21 09:04:06 2018  status update timer
        Aug 21 09:04:06 2018  loc intf up                  lsi.1048578
        Aug 21 09:04:05 2018  PE route changed
        Aug 21 09:04:05 2018  Out lbl Update                    262145
        Aug 21 09:04:05 2018  In lbl Update                     262147
        Aug 21 09:04:05 2018  loc intf down
```

On SJ PE

```
lab@SJ> show vpls connections extensive
Layer-2 VPN connections:

Legend for connection status (St)
EI -- encapsulation invalid      NC -- interface encapsulation not CCC/TCC/VPLS
EM -- encapsulation mismatch     WE -- interface and instance encaps not same
VC-Dn -- Virtual circuit down    NP -- interface hardware not present
CM -- control-word mismatch      -> -- only outbound connection is up
CN -- circuit not provisioned    <- -- only inbound connection is up
OR -- out of range               Up -- operational
OL -- no outgoing label          Dn -- down
LD -- local site signaled down   CF -- call admission control failure
RD -- remote site signaled down  SC -- local and remote site ID collision
LN -- local site not designated  LM -- local site ID not minimum designated
RN -- remote site not designated RM -- remote site ID not minimum designated
XX -- unknown connection status  IL -- no incoming label
MM -- MTU mismatch               MI -- Mesh-Group ID not available
BK -- Backup connection          ST -- Standby connection
PF -- Profile parse failure      PB -- Profile busy
RS -- remote site standby        SN -- Static Neighbor
VM -- VLAN ID mismatch

Legend for interface status
Up -- operational
Dn -- down

Instance: CUST3_SITE2
  Local site: 2 (2)
    Number of local interfaces: 1
    Number of local interfaces up: 1
    IRB interface present: no
    ge-0/0/8.0
    lsi.1048579         1         Intf - vpls CUST3_SITE2 local site 2 remote site 1
    lsi.1048577         3         Intf - vpls CUST3_SITE2 local site 2 remote site 3
    Label-base        Offset     Size  Range     Preference
    262145            1          8      8         100
    connection-site           Type  St     Time last up          # Up trans
    1                         rmt   Up     Aug 21 09:05:17 2018           1
      Remote PE: 192.0.2.1, Negotiated control-word: No
      Incoming label: 262145, Outgoing label: 262146
      Local interface: lsi.1048579, Status: Up, Encapsulation: VPLS
        Description: Intf - vpls CUST3_SITE2 local site 2 remote site 1
    Connection History:
        Aug 21 09:05:17 2018  status update timer
        Aug 21 09:05:17 2018  loc intf up                  lsi.1048579
        Aug 21 09:05:17 2018  PE route changed
        Aug 21 09:05:17 2018  Out lbl Update                    262146
        Aug 21 09:05:17 2018  In lbl Update                     262145
        Aug 21 09:05:17 2018  loc intf down
    3                         rmt   Up     Aug 20 12:02:40 2018           1
      Remote PE: 192.0.2.5, Negotiated control-word: No
      Incoming label: 262147, Outgoing label: 262146
      Local interface: lsi.1048577, Status: Up, Encapsulation: VPLS
        Description: Intf - vpls CUST3_SITE2 local site 2 remote site 3
    Connection History:
        Aug 20 12:02:40 2018  status update timer
```

```
        Aug 20 12:02:40 2018  loc intf up                  lsi.1048577
        Aug 20 12:02:40 2018  PE route changed
        Aug 20 12:02:40 2018  Out lbl Update                    262146
        Aug 20 12:02:40 2018  In lbl Update                     262147
        Aug 20 12:02:40 2018  loc intf down
```

On NY PE

```
lab@NY> show vpls connections extensive
Layer-2 VPN connections:

Legend for connection status (St)
EI -- encapsulation invalid      NC -- interface encapsulation not CCC/TCC/VPLS
EM -- encapsulation mismatch     WE -- interface and instance encaps not same
VC-Dn -- Virtual circuit down    NP -- interface hardware not present
CM -- control-word mismatch      -> -- only outbound connection is up
CN -- circuit not provisioned    <- -- only inbound connection is up
OR -- out of range               Up -- operational
OL -- no outgoing label          Dn -- down
LD -- local site signaled down   CF -- call admission control failure
RD -- remote site signaled down  SC -- local and remote site ID collision
LN -- local site not designated  LM -- local site ID not minimum designated
RN -- remote site not designated RM -- remote site ID not minimum designated
XX -- unknown connection status  IL -- no incoming label
MM -- MTU mismatch               MI -- Mesh-Group ID not available
BK -- Backup connection          ST -- Standby connection
PF -- Profile parse failure      PB -- Profile busy
RS -- remote site standby        SN -- Static Neighbor
VM -- VLAN ID mismatch

Legend for interface status
Up -- operational
Dn -- down

Instance: CUST3_SITE3
  Local site: 3 (3)
    Number of local interfaces: 1
    Number of local interfaces up: 1
    IRB interface present: no
    ge-0/0/6.0
    lsi.1048578         1         Intf - vpls CUST3_SITE3 local site 3 remote site 1
    lsi.1048577         2         Intf - vpls CUST3_SITE3 local site 3 remote site 2
    Label-base        Offset     Size  Range     Preference
    262145            1          8      8         100
    connection-site           Type  St     Time last up          # Up trans
    1                         rmt   Up     Aug 21 09:05:01 2018           1
      Remote PE: 192.0.2.1, Negotiated control-word: No
      Incoming label: 262145, Outgoing label: 262147
      Local interface: lsi.1048578, Status: Up, Encapsulation: VPLS
        Description: Intf - vpls CUST3_SITE3 local site 3 remote site 1
    Connection History:
        Aug 21 09:05:01 2018  status update timer
        Aug 21 09:05:00 2018  loc intf up                  lsi.1048578
        Aug 21 09:04:59 2018  PE route changed
        Aug 21 09:04:59 2018  Out lbl Update                    262147
        Aug 21 09:04:59 2018  In lbl Update                     262145
        Aug 21 09:04:59 2018  loc intf down
    2                         rmt   Up     Aug 20 12:02:40 2018           1
```

```
      Remote PE: 192.0.2.6, Negotiated control-word: No
      Incoming label: 262146, Outgoing label: 262147
      Local interface: lsi.1048577, Status: Up, Encapsulation: VPLS
        Description: Intf - vpls CUST3_SITE3 local site 3 remote site 2
    Connection History:
        Aug 20 12:02:40 2018  status update timer
        Aug 20 12:02:40 2018  loc intf up                  lsi.1048577
        Aug 20 12:02:40 2018  PE route changed
        Aug 20 12:02:40 2018  Out lbl Update                    262147
        Aug 20 12:02:40 2018  In lbl Update                     262146
        Aug 20 12:02:40 2018  loc intf down
```

Customer Router

```
lab@CUSTC_SITE1>show ospf neighbor instance VPLS_SITE1
Address          Interface              State     ID               Pri  Dead
100.64.65.3      ge-0/0/7.0             Full      100.64.65.3      128    33
100.64.65.2      ge-0/0/7.0             Full      100.64.65.2      128    31
```

So according to this example device output, Customer A, Customer B, and Customer C's sites are all reachable.

Summary

While this final chapter covers topics mentioned in previous chapters, it should reinforce your MPLS familiarity and help you to fully grasp how flexible and powerful MPLS is as a protocol on the Junos OS.

Where to Go Next

Try these resources to continue your learning curve:

Metro Ethernet Design Guide (Best practices)
https://www.juniper.net/documentation/en_US/release-independent/solutions/information-products/pathway-pages/solutions/metro-ethernet-dg.pdf

MPLS Applications Feature Guide (detailed configuration examples)
https://www.juniper.net/documentation/en_US/junos/information-products/pathway-pages/config-guide-mpls-applications/config-guide-mpls-applications.pdf

The J-Net Communities forum is dedicated to sharing information, best practices, and questions about Juniper products, technologies, and solutions. Please register to participate in this free forum and ask your next MPLS question.
http://forums.juniper.net/jnet

If you need more on Junos, try starting here:
http://www.juniper.net/techpubs/software/junos
or, the *Day One* book library also has several books on the Junos OS:
http://www.juniper.net/dayone
and visit the Learning Portal at Juniper Networks:
https://learningportal.juniper.net/juniper/user_courses.aspx

If you want or need an app for your devices use the Junos Genius app:
https://www.juniper.net/us/en/training/junos-genius/