JUNIPER NETWORKS

Network Configuration Example

Next-Generation VPLS for Multicast with Multihoming Using BGP

Release

NCE0010

Modified: 2017-01-25

Juniper Networks, Inc.
1133 Innovation Way
Sunnyvale, California 94089
USA
408-745-2000
www.juniper.net

Network Configuration Example Next-Generation VPLS for Multicast with Multihoming Using BGP
NCE0010

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

YEAR 2000 NOTICE

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

END USER LICENSE AGREEMENT

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

Table of Contents

CHAPTER 1

Next-Generation VPLS for Multicast with Multihoming Using BGP

About This Network Configuration Example

This network configuration example provides step-by-step procedures to configure BGP multihoming and verify the BGP multihoming configuration. This document only discusses the multihomed links between the provider edge (PE) routers and the customer edge (CE) routers related to the configuration.

Next-Generation VPLS for Multicast with Multihoming Overview

VPLS emulates the broadcast domain of a LAN across an MPLS network cloud. Traditional MPLS implementations of VPLS require that all participating ingress PE routers make separate copies of each broadcast or multicast packet to send to all other PE routers that are part of the VPLS site for the same extended LAN. In a large virtual private network (VPN), replication overhead can be significant for each ingress router and its attached core-facing links.

Junos OS offers the following VPLS enhancements which provide redundancy for VPLS between PE and CE routers:

- Redundancy using BGP for multihomed links between PE and CE devices— Juniper Networks integrates the local preference and path selection capability of BGP with VPLS to allow a CE Ethernet switch to have a backup path across the network.
- Redundancy using the Spanning Tree Protocol (STP) for multihomed links between PE and CE devices— Various versions of STP can be used in the CE network to avoid loops in a multihoming environment. The provider does not have any control over this customer network configuration. The provider can also implement BGP-based loop avoidance as an additional measure to avoid loops.

The following standardized VPLS implementations are supported by the Internet Engineering Task Force (IETF):

- RFC 4761, *Virtual Private LAN Service (VPLS) Using BGP for Auto-Discovery and Signaling*
- RFC 4762, *Virtual Private LAN Service (VPLS) Using LDP Signaling*

For more information about the basic configuration of next-generation VPLS, see the Network Configuration Example *Next-Generation VPLS Using Point-to-Multipoint LSPs for Unicast and Multicast Forwarding*.

For a detailed technology overview of VPLS, you can refer to *LDP-BGP VPLS Interworking* at the following location:
http://www.juniper.net/us/en/local/pdf/whitepapers/2000282-en.pdf .

Redundancy Using BGP for Multihomed Links between PE and CE Routers

Juniper Networks implements a BGP-based multihoming solution to provide redundancy for VPLS between PE and CE routers.

In this implementation:

- VPLS-enabled PE routers (also called VPLS PE routers) collectively elect one of the VPLS PE routers, to which a site is multihomed, as the designated forwarder of traffic between this site and all other sites.
- All the other VPLS PE routers, to which the same site is connected, do not forward traffic to or from the site.
- Essentially all VPLS PE routers behave as if the site is singlehomed to the VPLS PE router that is the designated forwarder.
- Service providers are able to prevent well-known Layer 2 loops without relying on the customer’s STP configuration.
- Customers can still run STP as a fallback strategy to prevent loops that are formed without the service provider's knowledge.

The benefits of multihoming include:

- Redundancy of the link connecting the PE router and the CE device.
- Redundancy of the directly connected PE routers.
- Faster convergence when there is a link failure between a PE router and CE device.
- The same BGP attributes are used to configure primary and backup links.

Operation of Next-Generation VPLS for Multicast with Multihoming Using BGP

VPLS provides a multipoint-to-multipoint Ethernet service that can span one or more metro areas and multiple sites. VPLS provides connectivity as if these sites are attached to the same Ethernet LAN.

VPLS uses an IP and MPLS service provider infrastructure. From the service provider’s point of view, using IP and MPLS routing protocols and procedures instead of STP, and using MPLS labels instead of VLAN identifiers (IDs), significantly improves the scalability of the VPLS service.

Single CE Site Connected to Multiple VPLS PE Routers

This section describes the process used to elect a single designated forwarder for a multihomed site.

For a multihomed site, all the PE routers in the VPLS instance elect the same designated forwarder PE router using the BGP VPLS multihoming procedure. Only elected designated forwarders forward traffic to and receive traffic from the multihomed site. All other PE routers where this multihomed site is present do not participate in forwarding for that site.

All remote PE routers are aware of the designated forwarder PE router for each multihomed site and do not create a pseudowire to the PE routers that are not the designated forwarder for the multihomed site.

In Figure 1 on page 8:

- The same site ID (sometimes known as a VPLS edge identifier or VE ID) is configured on all VPLS PE routers to which a site is multihomed.
- All PE routers are aware of which sites are multihomed since they see multiple advertisements with the same site ID.
- One of the VPLS PE routers is selected as the designated forwarder for this site by all PE routers based on a deterministic algorithm.
- The algorithm selects the VPLS PE router that originates the best advertisement with a particular site ID as the designated forwarder. There are two possible selection methods:
  - BGP path selection on the route reflector and the PE routers
  - VPLS site selection on the PE router only
- If multiple network layer reachability information (NLRI) advertisements have the same route distinguisher and site ID, the router uses BGP path selection rules to select the best path. The BGP rules are:
  - Always prefer advertisements that do not have the down bit set over ones that do have this bit set.
  - Prefer the advertisement with the higher local preference.
  - Use the configurable per-site site preference to set the BGP local preference in the advertisement and influence the choice of the designated forwarder.
  - Ignore the interior gateway protocol (IGP) metric while doing path selection because the choice of designated forwarder must be the same on all PE routers.
- Among advertisements with the same route distinguisher, apply VPLS site selection rules (a subset of BGP path selection rules) to pick the select advertisement.

Figure 1 on page 8 illustrates the following four-step process to select the designated forwarder and create the pseudowire:

1. Router PE1 and Router PE2 both have the same site ID (Site 1) for Router CE1.

   Router PE1 has a better local preference of 65535 and is configured as the primary router.

2. Router PE3 receives the BGP NLRI advertisement from Router PE1 and Router PE2 with the local preferences of 65535 and 1, respectively.

3. Router PE3 runs the BGP path selection algorithm and selects Router PE1 as the designated forwarder VPLS edge PE router for Site 1.

4. Router PE3 creates the pseudowire only to Router PE1, which helps to save bandwidth in the network core.

Figure 1: Single CE Site Multihomed with Two PE Routers

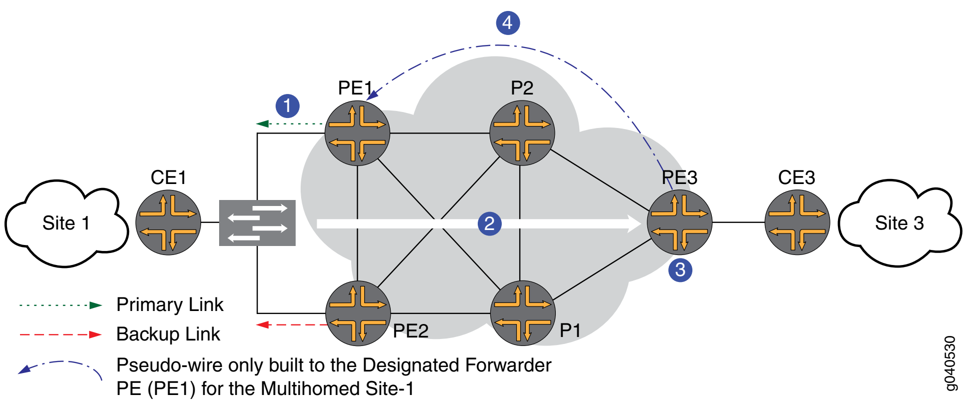

The resulting VPLS PE router roles for Site 1 are:

- Router PE1 is the designated forwarder VPLS edge PE router.
- Router PE2 is the non-designated forwarder VPLS edge PE router.
- Router PE3 is the remote VPLS edge PE router.

All the interfaces linking the CE and PE devices that are connected to the designated forwarder VPLS PE router, are marked **Up** and **forwarding** in **show** command output.

All the interfaces linking the CE and PE devices on the non-designated forwarder VPLS PE router, are marked **vc-down** in **show** command output. The router does not send traffic or forward received traffic on these interfaces.

Remote VPLS PE routers establish pseudowires only to the designated PE router, and tear down any pseudowires to the non-designated PE router.

Multiple CE Sites Connected to a Single VPLS PE Router for Link Redundancy

This section describe some of the operational details of multiple CE sites connected to a single VPLS PE router.

In Figure 2 on page 9:

- Router CE2 is multihomed to Router PE1 and Router PE2.
- Router CE1 is singlehomed to Router PE1.

Figure 2: Two CE Sites Multihomed to a Single PE Router on Different Line Cards

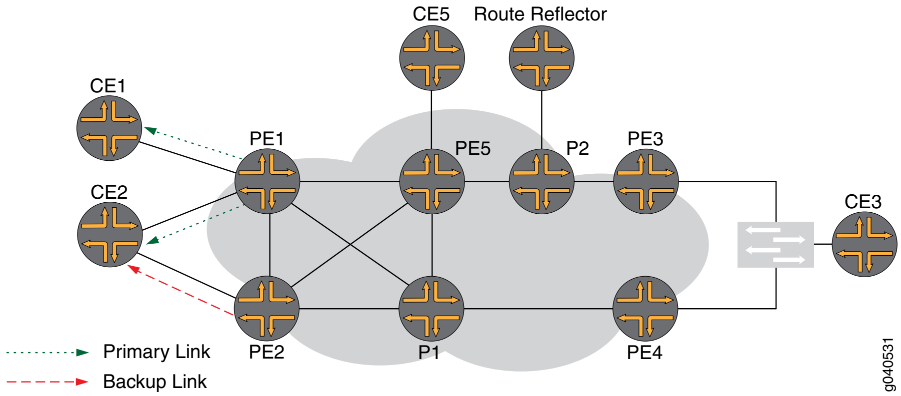

The scenario shown in Figure 2 is common. Your network might have a single PE router in a remote area, but you would like to multihome a Layer 2 network to different Flexible PIC Concentrators (FPCs) on the same PE router. This configuration provides link redundancy on the CE devices and link redundancy on the links between the CE and PE devices, but limited link redundancy on PE devices. In this case, you need the ability to configure a site to use a single active interface for forwarding.

In this scenario:

- Path selection is done per site to determine if a PE router is the designated forwarder for that site or not.
- Only a single pseudowire is established between any two PE routers, even if one or both of them have multiple designated PE routers.

NOTE: **A pseudowire between two PE routers is always established between the designated sites with the minimum site IDs on the two PE routers.**

- Establishing a single pseudowire avoids the need to maintain multiple flooding and media access control (MAC) address tables per instance (one per site) on each PE router.

- The local interfaces are marked **vc-down** in the **show** command output where a site is connected to the non-designated forwarder router.
- When a designated site on a PE router fails, all MAC addresses from this remote PE router have to be learned again, since the router does not know the exact site where the MAC addresses were originally learned from.

Implementation of Redundancy Using VPLS Multihomed Links Between PE and CE Devices

You might need to multihome a CE device to multiple PE routers without causing a Layer 2 forwarding loop. This is not a problem if the CE device is a router, since no Layer 2 loops can form when using a router. However, if the CE device is a Layer 2 device, like a hub or switch, multihoming it to two PE routers can cause a Layer 2 loop.

You can use one of the following methods to prevent the Layer 2 loop:

- BGP-based primary and backup link selection.
- Spanning tree protocol (STP) to prune links to the CE router. However, this method requires the service provider to trust its customer to not cause any Layer 2 loops by misconfiguration.
- Active and standby up link functionality, such as the redundant trunk groups that are supported on Juniper Networks EX Series Ethernet Switches.

The limitations of using STP on the CE site are:

- Backbone and access network bandwidth is not used efficiently.
- PE routers using STP to prevent loops with dual-homed sites receive broadcast traffic unnecessarily because the pseudowire to the standby PE router still exists.
- When the direct link between the CE and PE router fails, multihoming works fine. When a link connected downstream from the CE router fails, multihoming does not work.

The benefits and properties of the BGP-based solution are:

- BGP path selection does not have the limitations of STP.
- A CE device that is multihomed to multiple PE routers is given the same site ID on all the PE routers it is multihomed to.
- The BGP path selection algorithm selects the router that originates the best advertisement as the VPLS PE designated forwarder.
- If desired, you can set the local preference on the PE routers to control BGP path selection.
- BGP path selection occurs on the route reflector and the PE router.
- An IGP metric is not part of the selection process.
- If the route distinguisher is the same on both PE routers, the route reflector selects one PE router as the designated forwarder. If the route distinguishers are different on the PE routers, the route reflector forwards both copies of the route to the remote PE routers.

Related Documentation

- Example: Next-Generation VPLS for Multicast with Multihoming on page 11

- *Example: NG-VPLS Using Point-to-Multipoint LSPs*
- *Next-Generation VPLS Point-to-Multipoint Forwarding Overview*

Example: Next-Generation VPLS for Multicast with Multihoming

This example shows how to configure next-generation VPLS for multicast with multihoming. It is organized in the following sections:

Requirements

The following table lists the hardware and software requirements for this configuration.

Table 1: Hardware and Software Used

Equipment	Components	Software
Four MX Series 3D Universal Edge Routers	DPC40X-1GE -X, DPC 4X-10GE-X, DPC40x-1GE-R, DPC 4X-10GE-R	Junos OS Release 9.3 or later
Two M320 Multiservice Edge Routers and T Series Core Routers	FPC 3, 10GE Xenpak	Junos OS Release 9.3 or later
Five EX Series Ethernet Switches	EX4200, EX3200	Junos OS Release 9.4 or later

NOTE: This configuration example has been tested using the software release listed and is assumed to work on all later releases.

Overview and Topology

Figure 3 on page 12 shows the physical topology used in this next-generation VPLS multihoming example.

Figure 3: Physical Topology of Next-Generation VPLS for Multicast with Multihoming

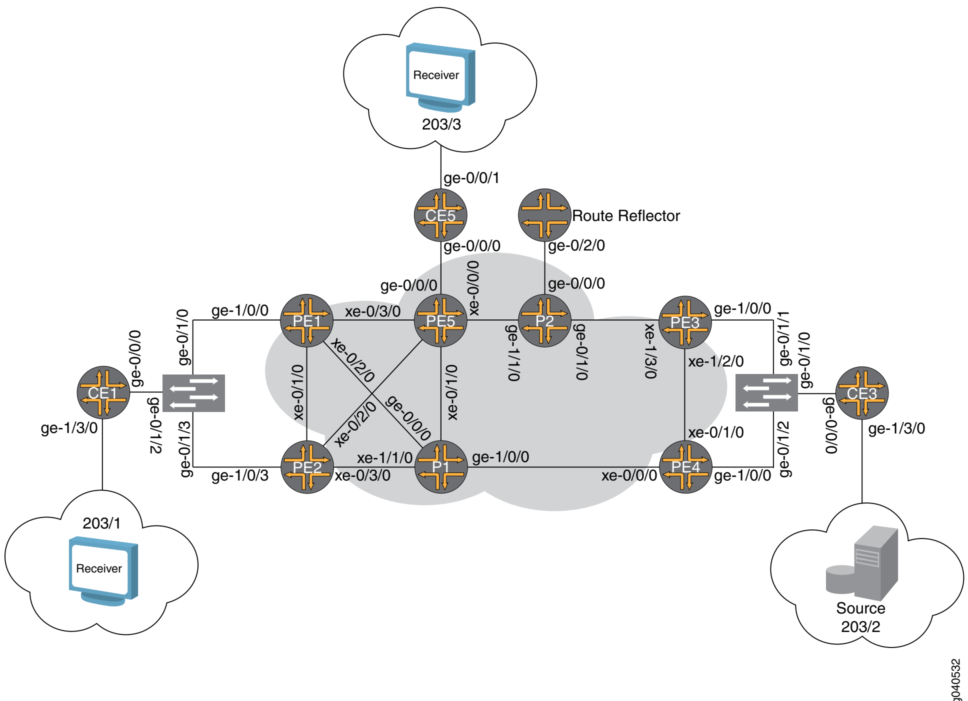

Figure 4 on page 12 show the logical topology of the next-generation VPLS multihoming example.

Figure 4: Logical Topology of Next-Generation VPLS for Multicast with Multihoming

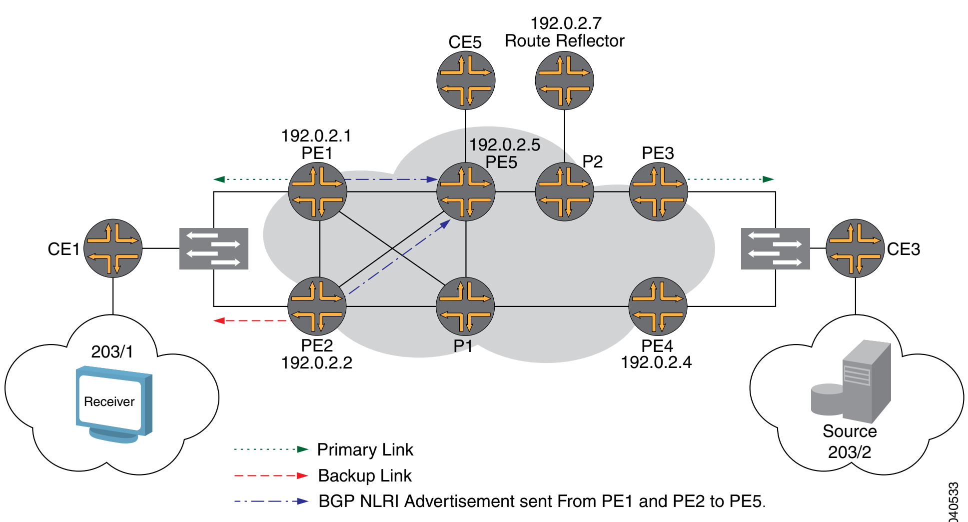

The network state and configuration before the implementation is as follows:

- Five PE routers participating in the next-generation VPLS domain named GOLD.
- OSPF, BGP, and RSVP are configured on the MPLS core interfaces.
- The **no-tunnel-services** statement is included in the VPLS routing instance. This statement supports the use of label-switched interface (LSI) tunnel interfaces for VPLS.
- Router PE1 and Router PE2 are configured with a dynamic point-to-multipoint LSP using the **vpls-GOLD-p2mp-template** template.
- Router PE3 and Router PE4 are configured to use static point-to-multipoint LSPs.

NOTE: **Single-hop point-to-multipoint LSPs are not supported, so single-hop point-to-multipoint LSPs are down.**

- Router CE1 is multihomed to Router PE1 and Router PE2 through an EX4200 Layer 2 switch.
- Router CE3 is multihomed to Router PE3 and Router PE4 through an EX4200 Layer 2 switch.
- Router CE5 is singlehomed to Router PE5.
- The off-path route reflector is configured for BGP. The **family l2vpn** statement is included in the route reflector configuration.
- Router CE3 is connected to test equipment through port 203/2. The test equipment generates multicast traffic to groups 230.1.1.1 through 230.1.1.10 at the rate of 10,000 pps.
- Router CE1 and Router CE5 are configured with static Internet Group Management Protocol (IGMP) joins so they can receive the multicast traffic from Router CE3.
- The Layer 2 switches are configured with truck ports to the PE routers and access ports to the test equipment.

Here is a summary of the steps necessary to complete the configuration successfully:

1. Configure a unique route distinguisher for the VPLS routing instance named GOLD on Router PE1, Router PE2, Router PE3, and Router PE4.
2. Configure the same site ID for the multihomed PE routers. Configure both Router PE1 and Router PE2 with a site ID value of 1. Configure both Router PE3 and Router PE4 with a site ID value of 3.
3. Configure multihoming under the CE1 site configuration.
4. Configure the site-preference **Primary** on Router PE1 and configure the site-preference **Backup** on Router PE2. In this case, Router PE1 has the primary link to Router CE1 and Router PE2 has the backup link to Router CE1.
5. Configure the site preference on Router PE3 and Router PE4. Configure Router PE3 as the primary and Router PE4 as the backup.

Configuration

This section provides a step-by-step procedure to configure next-generation VPLS for multicast with multihoming.

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

This example is organized in the following sections:

Configuring Next-Generation VPLS Multihoming

Step-by-Step Procedure

1. In BGP-based VPLS multihoming, it is recommended that you configure distinct route distinguishers for each multihomed router. Configuring distinct route distinguishers helps with faster convergence when the connection to a primary router goes down. It also requires the other backup PE routers to maintain additional state information for faster convergence.

   There are two levels of path selection:

   - The first is BGP: BGP uses a combination of route distinguisher, site ID, and VE block offset for BGP path selection.
   - The second is in VPLS: VPLS uses the site ID for VPLS path selection.

   By configuring unique route distinguishers, the prefixes for BGP path selection are all unique. Therefore, BGP path selection is skipped and VPLS path selection is used, which only looks at the site ID.

   On Router PE1, Router PE2, Router PE3, and Router PE4 configure a unique router distinguisher for the **GOLD** routing instance.

   ```
   user@PE1# set routing-instance GOLD route-distinguisher 1.1.1.1:1
   ```

   ```
   user@PE2# set routing-instance GOLD route-distinguisher 2.2.2.2:10
   ```

   ```
   user@PE3# set routing-instance GOLD route-distinguisher 3.3.3.3:1
   ```

   ```
   user@PE4# set routing-instance GOLD route-distinguisher 4.4.4.4:10
   ```

2. Configure site ID **1** on Routers PE1 and PE2 for Router CE1. Configure site ID **3** on Routers PE3 and PE4 for Router CE3.

   ```
   user@PE1# set routing-instance GOLD protocols vpls site CE1 site-identifier 1
   ```

   ```
   user@PE2# set routing-instance GOLD protocols vpls site CE1 site-identifier 1
   ```


user@PE3# **set routing-instance GOLD protocols vpls site CE3 site-identifier 3**

user@PE4# **set routing-instance GOLD protocols vpls site CE3 site-identifier 3**

3. Enable multihoming by including the **multi-homing** statement under the multihomed site configuration on Router PE1, Router PE2, Router PE3, and Router PE4.

   user@PE1# **set routing-instance GOLD protocols vpls site CE1 multi-homing**

   user@PE2# **set routing-instance GOLD protocols vpls site CE1 multi-homing**

   user@PE3# **set routing-instance GOLD protocols vpls site CE3 multi-homing**

   user@PE4# **set routing-instance GOLD protocols vpls site CE3 multi-homing**

4. Include the **site-preference primary** statement on Router PE1 and Router PE3, and include the **site-preference backup** statement on Router PE2 and Router PE4. The **site-preference primary** statement sets the local preference to the highest value (65535) and the **site-preference backup** statement sets the BGP local preference to 1. Since the site ID is the same, the routers select the highest local preference value as the designated forwarder.

   user@PE1# **set routing-instance GOLD protocols vpls site CE1 site-preference primary**

   user@PE2# **set routing-instance GOLD protocols vpls site CE1 site-preference backup**

   user@PE3# **set routing-instance GOLD protocols vpls site CE3 site-preference primary**

   user@PE4# **set routing-instance GOLD protocols vpls site CE3 site-preference backup**

Validating the VPLS Control Plane

Step-by-Step Procedure

This section presents show commands that you can use to verify the operation of the example configuration.

In this example the traffic patterns are:

- The source is connected to Router CE3 and sends 10,000 pps for the groups 230.1.1.1 to 230.1.1.10. Router CE3 is configured as a rendezvous point.
- Multicast receivers are connected to both Router CE1 and Router CE5. Protocol Independent Multicast (PIM) join messages are generated by the test equipment.
- The link between Router PE3 and Router CE3 and the link between Router PE1 and Router CE1 are configured as primaries for VPLS multihoming.
- All PE routers have a BGP session with the route reflector.
- All PE routers have a label-switched path (LSP) that is created to the route reflector so that the PE routers have a route to the route reflector in the **inet.3** table for route resolution.

1. On Router PE1, use the **show vpls connections** command to verify that the VPLS connections are **Up** between Router PE1 and Router PE3 and between Router PE1 and PE5. Router PE1 is the primary link selected by the VPLS multihoming configuration.

```
user@PE1# 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 availble
BK -- Backup connection          ST -- Standby connection

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

Instance: GOLD
  Local site: CE1 (1)
    connection-site          Type  St     Time last up          # Up trans
    1                        rmt   RN
    3                        rmt   Up     Nov 16 11:22:44 2009           1
      Remote PE: 3.3.3.3, Negotiated control-word: No
      Incoming label: 262147, Outgoing label: 262145
      Local interface: lsi.1048835, Status: Up, Encapsulation: VPLS
        Description: Intf - vpls GOLD local site 1 remote site 3
    5                        rmt   Up     Nov 16 11:22:46 2009           1
      Remote PE: 5.5.5.5, Negotiated control-word: No
      Incoming label: 262149, Outgoing label: 262161
      Local interface: lsi.1048836, Status: Up, Encapsulation: VPLS
        Description: Intf - vpls GOLD local site 1 remote site 5
```

2. On Router PE2, use the **show vpls connections** command to verify that the VPLS connections to Router PE3 and Router PE5 are in the **LN** state, meaning the local router is not the designated forwarder. Router PE2 is configured to be the backup link for Router CE1.

```
user@PE2# show vpls connections
...

Instance: GOLD
  Local site: CE1 (1)
    connection-site          Type  St     Time last up          # Up trans
    1                        rmt   LN
    3                        rmt   LN
    5                        rmt   LN
```

3. On Router PE3, use the **show vpls connections** command to verify that the VPLS connections to Router PE1 and Router PE5 are **Up**. Router PE3 is configured to be the primary link for Router CE3.

```
user@PE3# show vpls connections
...

Instance: GOLD
  Local site: CE3 (3)
    connection-site           Type  St     Time last up          # Up trans
    1                         rmt   Up     Nov 16 11:22:01 2009           1
      Remote PE: 1.1.1.1, Negotiated control-word: No
      Incoming label: 262145, Outgoing label: 262147
      Local interface: lsi.1048832, Status: Up, Encapsulation: VPLS
        Description: Intf - vpls GOLD local site 3 remote site 1
    3                         rmt   RN
    5                         rmt   Up     Nov 16 11:22:56 2009           1
      Remote PE: 5.5.5.5, Negotiated control-word: No
      Incoming label: 262149, Outgoing label: 262163
      Local interface: lsi.1048834, Status: Up, Encapsulation: VPLS
        Description: Intf - vpls GOLD local site 3 remote site 5
```

4. On Router PE4, use the **show vpls connections** command to verify that the VPLS connections are in the **LN** state, meaning the local site is not designated. Router PE4 is configured to be the backup link for Router CE3.

```
user@PE4# show vpls connections
...

Instance: GOLD
  Local site: CE3 (3)
    connection-site           Type  St     Time last up          # Up trans
    1                         rmt   LN
    3                         rmt   SC
    5                         rmt   LN
```

5. On Router PE1, use the **show route advertising-protocol** command to verify that Router PE1 (the multihoming primary router) is sending the BGP Layer 2 VPN route advertisement to the route reflector with the local preference value of **65535**. The local preference is used by Router PE3 to select Router PE1 as the designated forwarder, rather than selecting Router PE2 that has a local preference of **1**.

```
user@PE1# show route advertising-protocol bgp 7.7.7.7 extensive
GOLD.l2vpn.0: 5 destinations, 5 routes (5 active, 0 holddown, 0 hidden)
* 1.1.1.1:1:1:1/96 (1 entry, 1 announced)
 BGP group to-RR type Internal
     Route Distinguisher: 1.1.1.1:1
     Label-base: 262145, range: 8
     Nexthop: Self
     Flags: Nexthop Change
     Localpref: 65535
     AS path: [65000] I
     Communities: target:65000:1 Layer2-info: encaps:VPLS, control flags:, mtu: 0, site preference: 65535

     PMSI: Flags 0:RSVP-TE:label[0:0:0]:Session_13[1.1.1.1:0:9519:1.1.1.1]
```

6. On Router PE2, use the **show route advertising-protocol** command to verify that Router PE2 is configured as the multihoming backup with a local preference of **1**.

```
user@PE2# show route advertising-protocol bgp 7.7.7.7 extensive
```



```
GOLD.l2vpn.0: 5 destinations, 5 routes (5 active, 0 holddown, 0 hidden)
* 2.2.2.2:10:1:1/96 (1 entry, 1 announced)
 BGP group to-RR type Internal
     Route Distinguisher: 2.2.2.2:10
     Label-base: 262145, range: 8
     Nexthop: Self
     Flags: Nexthop Change
     Localpref: 1
     AS path: [65000] I
     Communities: target:65000:1 Layer2-info: encaps:VPLS, control flags:, mtu: 0, site preference: 1
```

7. On Router PE3, use the **show route receive-protocol** command to verify that Router PE3 receives the Layer 2 VPN route from the route reflector for Router PE1 and Router PE2 with different local preference values.

   BGP route selection is based on the received **l2vpn** routes for the VPLS site connected to multihomed PE routers. Since the route distinguishers are different on Router PE1 and Router PE2, Router PE3 and Router PE4 consider the received routes from Router PE1 and Router PE2 as different routes. Router PE3 and Router PE4 run the BGP path selection algorithm and select Router PE1, the router advertising the route with the higher local preference value, as the designated forwarder.

```
user@PE3# show route receive-protocol bgp 7.7.7.7
bgp.l2vpn.0: 4 destinations, 4 routes (4 active, 0 holddown, 0 hidden)
  Prefix                  Nexthop              MED     Lclpref    AS path
  1.1.1.1:1:1:1/96
*                         1.1.1.1                      65535      I
  2.2.2.2:10:1:1/96
*                         2.2.2.2                      1          I
  4.4.4.4:10:3:1/96
*                         4.4.4.4                      1          I
  5.5.5.5:10:5:1/96
*                         5.5.5.5                      100        I

GOLD.l2vpn.0: 5 destinations, 5 routes (5 active, 0 holddown, 0 hidden)
  Prefix                  Nexthop              MED     Lclpref    AS path
  1.1.1.1:1:1:1/96
*                         1.1.1.1                      65535      I
  2.2.2.2:10:1:1/96
*                         2.2.2.2                      1          I
  4.4.4.4:10:3:1/96
*                         4.4.4.4                      1          I
  5.5.5.5:10:5:1/96
*                         5.5.5.5                      100        I
```

8. On Router PE3, use the **show route table** command to verify that Router PE3 has selected the static point-to-multipoint LSP from Router PE3 to Router PE1 for forwarding.

   Notice that Router PE2 does not have any provider multicast service interface (PMSI) flags because PMSI attributes are not attached.

```
user@PE3# show route table GOLD.l2vpn.0 extensive
GOLD.l2vpn.0: 5 destinations, 5 routes (5 active, 0 holddown, 0 hidden)
1.1.1.1:1:1:1/96 (1 entry, 1 announced)
        *BGP    Preference: 170/-65536
                Route Distinguisher: 1.1.1.1:1
                PMSI: Flags 0:RSVP-TE:label[0:0:0]:Session_13[1.1.1.1:0:9519:1.1.1.1]
```



```
                Next hop type: Indirect
                Next-hop reference count: 4
                Source: 7.7.7.7
                Protocol next hop: 1.1.1.1
                Indirect next hop: 2 no-forward
                State: <Secondary Active Int Ext>
                Local AS: 65000 Peer AS: 65000
                Age: 2:30:44    Metric2: 1
                Task: BGP_65000.7.7.7.7+179
                Announcement bits (1): 0-GOLD-l2vpn
                AS path: I (Originator) Cluster list:  7.7.7.7
                AS path:  Originator ID: 1.1.1.1
                Communities: target:65000:1 Layer2-info: encaps:VPLS, control flags:, mtu: 0, site
preference: 65535
                Import Accepted
                Label-base: 262145, range: 8
                Localpref: 65535
                Router ID: 7.7.7.7
                Primary Routing Table bgp.l2vpn.0
                Indirect next hops: 1
                        Protocol next hop: 1.1.1.1 Metric: 3
                        Indirect next hop: 2 no-forward
                        Indirect path forwarding next hops: 1
                                Next hop type: Router
                                Next hop: 10.10.8.2 via xe-0/1/0.0 weight 0x1
                        1.1.1.1/32 Originating RIB: inet.3
                          Metric: 3                       Node path count: 1
                          Forwarding nexthops: 1
                                Nexthop: 10.10.8.2 via xe-0/1/0.0

2.2.2.2:10:1:1/96 (1 entry, 1 announced)
        *BGP    Preference: 170/-2
                Route Distinguisher: 2.2.2.2:10
                Next hop type: Indirect
                Next-hop reference count: 3
                Source: 7.7.7.7
                Protocol next hop: 2.2.2.2
                Indirect next hop: 2 no-forward
                State: <Secondary Active Int Ext>
                Local AS: 65000 Peer AS: 65000
                Age: 2:30:44    Metric2: 1
                Task: BGP_65000.7.7.7.7+179
                Announcement bits (1): 0-GOLD-l2vpn
                AS path: I (Originator) Cluster list:  7.7.7.7
                AS path:  Originator ID: 2.2.2.2
                Communities: target:65000:1 Layer2-info: encaps:VPLS, control flags:, mtu: 0, site
preference: 1
                Import Accepted
                Label-base: 262145, range: 8
                Localpref: 1
                Router ID: 7.7.7.7
                Primary Routing Table bgp.l2vpn.0
                Indirect next hops: 1
                        Protocol next hop: 2.2.2.2 Metric: 3
                        Indirect next hop: 2 no-forward
                        Indirect path forwarding next hops: 1
                                Next hop type: Router
                                Next hop: 10.10.8.2 via xe-0/1/0.0 weight 0x1
                        2.2.2.2/32 Originating RIB: inet.3
                          Metric: 3                       Node path count: 1
```

```
Forwarding nexthops: 1
    Nexthop: 10.10.8.2 via xe-0/1/0.0
```

9. On Router PE3, use the **show vpls connections** command to verify that the VPLS connection is in the **Up** state.

   Notice the display also shows the local interface and the incoming and outgoing label values used.

```
user@PE3# show vpls connections extensive
...

Instance: GOLD
  Local site: CE3 (3)
    Number of local interfaces: 1
    Number of local interfaces up: 1
    IRB interface present: no
    ge-1/0/0.1
    lsi.1048832        1         Intf - vpls GOLD local site 3 remote site 1
    lsi.1048833        2         Intf - vpls GOLD local site 3 remote site 2
        Interface flags: VC-Down
    lsi.1048834        5         Intf - vpls GOLD local site 3 remote site 5
        Interface flags: VC-Down
    Label-base        Offset     Range     Preference
    262145            1          8         65535
    connection-site           Type  St     Time last up          # Up trans
    1                         rmt   Up     Nov 16 11:22:01 2009           1
      Remote PE: 1.1.1.1, Negotiated control-word: No
      Incoming label: 262145, Outgoing label: 262147
      Local interface: lsi.1048832, Status: Up, Encapsulation: VPLS
        Description: Intf - vpls GOLD local site 3 remote site 1
      RSVP-TE P2MP lsp:
        Egress branch LSP:  3.3.3.3:1.1.1.1:1:vpls:GOLD, State: Up
    Connection History:
        Nov 16 11:22:54 2009  PE route changed
        Nov 16 11:22:01 2009  status update timer
        Nov 16 11:22:01 2009  PE route changed
        Nov 16 11:22:01 2009  Out lbl Update                    262147
        Nov 16 11:22:01 2009  In lbl Update                     262145
        Nov 16 11:22:01 2009  loc intf up                  lsi.1048832
    3                         rmt   RN
    5                         rmt   RD
  Ingress RSVP-TE P2MP LSP: vpls-GOLD, Flood next-hop ID: 616
```

Verifying the VPLS Data Plane

Step-by-Step Procedure

After the control plane is verified using the previous steps, you can verify the data plane. The data plane operation in the VPLS multihoming scenario is the same as the regular next-generation VPLS operation. This section describes the **show** command outputs that you can use to validate the data plane.

1. On Router PE3, use the **show mpls lsp** command to verify the state of the static LSPs and sub-LSPs.

   Router PE2 is configured with static point-to-multipoint LSPs and sub-LSPs with link protection. Point to multipoint LSPs are not supported for single-hop LSPs. In the following output notice that the single-hop point-to-multipoint LSP from Router PE3 to Router PE4 is **down**.

```
user@PE3# show mpls lsp p2mp ingress
Ingress LSP: 1 sessions
P2MP name: vpls-GOLD, P2MP branch count: 4
To              From            State Rt P     ActivePath       LSPname
5.5.5.5         3.3.3.3         Up      0 *                     to-pe5
1.1.1.1         3.3.3.3         Up      0 *                     to-pe1
4.4.4.4         3.3.3.3         Dn      0 *                     to-pe4
2.2.2.2         3.3.3.3         Up      0 *                     to-pe2
Total 4 displayed, Up 3, Down 1
```

2. On Router PE1, use the **show mpls lsp** command to verify the state of the dynamic LSPs.

   Router PE1 is using a dynamic point-to-multipoint LSP template configured with link protection. Notice that the LSP state is **Up** and that link protection is **desired**.

```
user@PE1# show mpls lsp p2mp ingress extensive
Ingress LSP: 1 sessions
P2MP name: 1.1.1.1:1:vpls:GOLD, P2MP branch count: 1

3.3.3.3
  From: 1.1.1.1, State: Up, ActiveRoute: 0, LSPname: 3.3.3.3:1.1.1.1:1:vpls:GOLD
  ActivePath:  (primary)
  P2MP name: 1.1.1.1:1:vpls:GOLD
  Link protection desired
  LoadBalance: Random
  Encoding type: Packet, Switching type: Packet, GPID: IPv4
 *Primary                    State: Up
    Priorities: 7 0
    OptimizeTimer: 50
    SmartOptimizeTimer: 180
    Reoptimization in 45 second(s).
    Computed ERO (S [L] denotes strict [loose] hops): (CSPF metric: 3)
 10.10.3.2 S 10.10.9.2 S 10.10.8.1 S
    Received RRO (ProtectionFlag 1=Available 2=InUse 4=B/W 8=Node 10=SoftPreempt):
         10.10.3.2(Label=488645) 4.4.4.4(flag=0x21) 10.10.9.2(flag=1 Label=299936) 10.10.8.1(Label=262145)

   12 Nov 16 15:38:08.116 CSPF: computation result ignored[314 times]
   11 Nov 16 11:23:44.856 Link-protection Up
   10 Nov 16 11:23:32.696 CSPF: computation result ignored[3 times]
    9 Nov 16 11:22:47.859 Record Route:  10.10.3.2(Label=488645) 4.4.4.4(flag=0x21) 10.10.9.2(flag=1
Label=299936) 10.10.8.1(Label=262145)
    8 Nov 16 11:22:44.910 Record Route:  10.10.3.2(Label=488645) 4.4.4.4(flag=0x20) 10.10.9.2(Label=299936)
 10.10.8.1(Label=262145)
```



```
    7 Nov 16 11:22:44.910 Up
    6 Nov 16 11:22:44.910 10.10.3.1: Down
    5 Nov 16 11:22:44.866 Selected as active path
    4 Nov 16 11:22:44.864 Record Route:  10.10.3.2(Label=488629) 4.4.4.4(flag=0x20) 10.10.9.2(Label=299920)
 10.10.8.1(Label=3)
    3 Nov 16 11:22:44.864 Up
    2 Nov 16 11:22:44.852 Originate Call
    1 Nov 16 11:22:44.852 CSPF: computation result accepted  10.10.3.2 10.10.9.2 10.10.8.1
  Created: Mon Nov 16 11:22:45 2009
Total 1 displayed, Up 1, Down 0
```

3. On Router PE3, use the **monitor interface traffic** command to verify the multicast replication behavior for the point-to-multipoint LSP on the designated forwarder Router PE3.

   The output shows that **10,000** pps are received on interface **ge-1/0/0** from Router CE3. The traffic has been forwarded to the provider (P) Router P2 and Router PE4 through xe-0/0/0 and xe-0/1/0, respectively. Based on the output, you can determine that a single copy of the packet is being sent to Router P2 and Router PE4.

```
user@PE3> monitor interface traffic
PE3                        Seconds: 8                  Time: 11:58:40

Interface    Link  Input packets        (pps)     Output packets        (pps)
 lc-0/0/0     Up               0                               0
 xe-0/0/0     Up        13570505          (0)         4507338866      (10000)
 lc-0/1/0     Up               0                               0
 xe-0/1/0     Up          292843          (1)          628972219      (10000)
 lc-0/2/0     Up               0                               0
 xe-0/2/0     Up          343292          (0)             206808          (1)
 lc-0/3/0     Up               0                               0
 xe-0/3/0   Down               0          (0)                  0          (0)
 ge-1/0/0     Up      2703709733       (9999)           13203544          (1)
 lc-1/0/0     Up               0                               0
 ge-1/0/1   Down     50380341937          (0)        60024542111          (0)
 ge-1/0/2   Down     60652323068          (0)        84480825838          (0)
 ge-1/0/3   Down     81219536264          (0)        84614255165          (0)
 ge-1/0/4   Down     54379241112          (0)        83656815208          (0)
```

4. On Router P2, use the **monitor interface traffic** command to verify that the multicast packet replication happens close to the PE routers connected to the receivers.

   Router PE1 and Router PE5 are connected to receivers that have joined this multicast group. Notice that incoming multicast packets from Router PE3 on the **ge-0/1/0** interface are replicated twice and sent out on the **ge-1/1/0** interface.

```
user@P2> monitor interface traffic
P2                         Seconds: 6                  Time: 12:07:58

Interface    Link  Input packets        (pps)     Output packets        (pps)
 ge-0/1/0     Up       661459806      (10000)             116236          (0)
 ge-1/1/0     Up          115956          (0)         1322690473      (20000)
 gr-2/1/0     Up               0          (0)                  0          (0)
 ip-2/1/0     Up               0          (0)                  0          (0)
```

5. On Router PE3, use the **show vpls flood** command to verify information about the flood next-hop route.

Junos OS Release 9.0 and later identifies the flood next-hop route as a composite next hop. Notice that the interface is **ge-1/0/0.1**, the next-hop type is **composite**, and that the flood composition is **flood-to-all**. This means the traffic is flooded to all the PE routers.

```
user@PE3# show vpls flood extensive
Name: GOLD
CEs: 1
VEs: 1
  Flood route prefix: 0x30002/51
  Flood route type: FLOOD_GRP_COMP_NH
  Flood route owner: __ves__
  Flood group name: __ves__
  Flood group index: 0
  Nexthop type: comp
  Nexthop index: 606
    Flooding to:
    Name             Type          NhType          Index
    __all_ces__      Group          comp            603
        Composition: split-horizon
        Flooding to:
        Name             Type          NhType          Index
        ge-1/0/0.1       CE             ucst            578

  Flood route prefix: 0x30003/51
  Flood route type: FLOOD_GRP_COMP_NH
  Flood route owner: __all_ces__
  Flood group name: __all_ces__
  Flood group index: 1
  Nexthop type: comp
  Nexthop index: 611
    Flooding to:
    Name             Type          NhType          Index
    __ves__          Group          comp            594
        Composition: flood-to-all
    Component p2mp NH (for all core facing interfaces):
    Index
    616
    Flooding to:
    Name             Type          NhType          Index
    __all_ces__      Group          comp            603
        Composition: split-horizon
        Flooding to:
        Name             Type          NhType          Index
        ge-1/0/0.1       CE             ucst            578

  Flood route prefix: 0x30001/51
  Flood route type: FLOOD_GRP_COMP_NH
  Flood route owner: __re_flood__
  Flood group name: __re_flood__
  Flood group index: 65534
  Nexthop type: comp
  Nexthop index: 598
    Flooding to:
    Name             Type          NhType          Index
    __ves__          Group          comp            594
        Composition: flood-to-all
    Component p2mp NH (for all core facing interfaces):
    Index
    616
```



```
    Flooding to:
    Name              Type          NhType          Index
    __all_ces__       Group          comp            603
        Composition: split-horizon
        Flooding to:
        Name              Type          NhType          Index
        ge-1/0/0.1        CE             ucst            578
Name: __juniper_private1__
CEs: 0
VEs: 0
```

6. On Router PE3, use the **show vpls mac-table** command to verify that the MAC address of the PE router at the remote end of the VPLS has been learned and added to the MAC address table.

   Notice that the MAC address is learned on the **ge-1/0/0.1** interface.

```
user@PE3# show vpls mac-table
MAC flags (S -static MAC, D -dynamic MAC,
           SE -Statistics enabled, NM -Non configured MAC)

Routing instance : GOLD
 Bridging domain : __GOLD__, VLAN : NA
   MAC                   MAC      Logical
   address               flags    interface
   00:14:f6:75:78:00  D     ge-1/0/0.1
```

7. On Router PE3, use the **show route forwarding-table** command to verify that the forwarding table has the required entries with two labels: one for the VPLS service and the other for the next-hop interface.

```
user@PE3> show route forwarding-table family vpls vpn GOLD
Routing table: GOLD.vpls
VPLS:
Destination        Type RtRef Next hop           Type Index NhRef Netif
default            perm     0                    dscd   574     1
lsi.1048832        intf     0                    indr 1048575     4
                          10.10.7.1             Push 262147, Push 309680(top)   596     2 xe-0/0/0.0
lsi.1048836        intf     0                    indr 1048574     4
                          10.10.7.1             Push 262179, Push 299856(top)   589     2 xe-0/0/0.0
00:10:db:e9:4e:b6/48
                   user     0                    indr 1048574     4
                          10.10.7.1             Push 262179, Push 299856(top)   589     2 xe-0/0/0.0
00:12:1e:c6:98:00/48
                   user     0                    indr 1048575     4
                          10.10.7.1             Push 262147, Push 309680(top)   596     2 xe-0/0/0.0
00:14:f6:75:78:00/48
                   user     0                    ucst   578     4 ge-1/0/0.1
0x30002/51         user     0                    comp   606     2
ge-1/0/0.1         intf     0                    ucst   578     4 ge-1/0/0.1
0x30003/51         user     0                    comp   611     2
0x30001/51         user     0                    comp   598     2
```

Results

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

The relevant sample configuration for Router PE1 follows:

Router PE1

```
chassis {
   dump-on-panic;
   fpc 1 {
      pic 3 {
         tunnel-services {
            bandwidth 1g;
         }
      }
   }
   network-services ethernet;
}
interfaces {
   xe-0/1/0 {
      unit 0 {
         family inet {
            address 10.10.2.1/30;
         }
         family mpls;
      }
   }
   xe-0/2/0 {
      unit 0 {
         family inet {
            address 10.10.3.1/30;
         }
         family mpls;
      }
   }
   xe-0/3/0 {
      unit 0 {
         family inet {
            address 10.10.1.1/30;
         }
         family mpls;
      }
   }
   ge-1/0/0 {
      vlan-tagging;
      encapsulation vlan-vpls;
      unit 1 {
         encapsulation vlan-vpls;
         vlan-id 1000;
         family vpls;
      }
   }
   lo0 {
      unit 0 {
         family inet {
            address 1.1.1.1/32;
         }
      }
   }
}
routing-options {
   static {
      route 172.0.0.0/8 next-hop 172.19.59.1;
```

```
    }
    autonomous-system 65000;
}
protocols {
    rsvp {
        interface all;
        interface fxp0.0 {
            disable;
        }
        interface xe-0/3/0.0 {
            link-protection;
        }
        interface xe-0/2/0.0 {
            link-protection;
        }
        interface xe-0/1/0.0 {
            link-protection;
        }
    }
    mpls {
        label-switched-path to-RR {
            to 7.7.7.7;
        }
        label-switched-path vpls-GOLD-p2mp-template {
            template;
            optimize-timer 50;
            link-protection;
            p2mp;
        }
        label-switched-path to-PE2 {
            to 2.2.2.2;
        }
        label-switched-path to-PE3 {
            to 3.3.3.3;
        }
        label-switched-path to-PE4 {
            to 4.4.4.4;
        }
        label-switched-path to-PE5 {
            to 5.5.5.5;
        }
        interface all;
        interface fxp0.0 {
            disable;
        }
    }
    bgp {
        group to-RR {
            type internal;
            local-address 1.1.1.1;
            family l2vpn {
                signaling;
            }
            neighbor 7.7.7.7;
        }
    }
```

```
  ospf {
    traffic-engineering;
    area 0.0.0.0 {
      interface all;
      interface fxp0.0 {
        disable;
      }
    }
  }
}
routing-instances {
  GOLD {
    instance-type vpls;
    interface ge-1/0/0.1;
    route-distinguisher 1.1.1.1:1;
    provider-tunnel {
      rsvp-te {
        label-switched-path-template {
          vpls-GOLD-p2mp-template;
        }
      }
    }
    vrf-target target:65000:1;
    protocols {
      vpls {
        site-range 8;
        no-tunnel-services;
        site CE1 {
          site-identifier 1;
          multi-homing;
          site-preference primary;
          interface ge-1/0/0.1;
        }
      }
    }
  }
}
```

The relevant sample configuration for Router PE2 follows.

PE2 Router

```
chassis {
  dump-on-panic;
  fpc 1 {
    pic 3 {
      tunnel-services {
        bandwidth 1g;
      }
    }
  }
  network-services ethernet;
}
interfaces {
  xe-0/1/0 {
    unit 0 {
      family inet {
        address 10.10.2.2/30;
```

```
            }
            family mpls;
        }
    }
    xe-0/2/0 {
        unit 0 {
            family inet {
                address 10.10.5.1/30;
            }
            family mpls;
        }
    }
    xe-0/3/0 {
        unit 0 {
            family inet {
                address 10.10.4.1/30;
            }
            family mpls;
        }
    }
    ge-1/0/1 {
        vlan-tagging;
        encapsulation vlan-vpls;
    }
    ge-1/0/3 {
        vlan-tagging;
        encapsulation vlan-vpls;
        unit 1 {
            encapsulation vlan-vpls;
            vlan-id 1000;
            family vpls;
        }
    }
    fxp0 {
        apply-groups [ re0 re1 ];
    }
    lo0 {
        unit 0 {
            family inet {
                address 2.2.2.2/32;
            }
        }
    }
}
routing-options {
    static {
        route 172.0.0.0/8 next-hop 172.19.59.1;
    }
    autonomous-system 65000;
}
protocols {
    rsvp {
        interface all;
        interface fxp0.0 {
            disable;
        }
```

```
    }
    mpls {
      label-switched-path to-RR {
        to 7.7.7.7;
      }
      label-switched-path vpls-GOLD-p2mp-template {
        template;
        optimize-timer 50;
        link-protection;
        p2mp;
      }
      label-switched-path to-PE1 {
        to 1.1.1.1;
      }
      label-switched-path to-PE3 {
        to 3.3.3.3;
      }
      label-switched-path to-PE4 {
        to 4.4.4.4;
      }
      label-switched-path to-PE5 {
        to 5.5.5.5;
      }
      interface all;
      interface fxp0.0 {
        disable;
      }
    }
    bgp {
      group to-RR {
        type internal;
        local-address 2.2.2.2;
        family l2vpn {
          signaling;
        }
        neighbor 7.7.7.7;
      }
    }
    ospf {
      traffic-engineering;
      area 0.0.0.0 {
        interface all;
        interface fxp0.0 {
          disable;
        }
      }
    }
  }
  routing-instances {
    GOLD {
      instance-type vpls;
      interface ge-1/0/3.1;
      route-distinguisher 2.2.2.2:10;
      provider-tunnel {
        rsvp-te {
          label-switched-path-template {
```



```
                vpls-GOLD-p2mp-template;
              }
            }
          }
          vrf-target target:65000:1;
          protocols {
            vpls {
              site-range 8;
              no-tunnel-services;
              site CE1 {
                site-identifier 1;
                multi-homing;
                site-preference backup;
                interface ge-1/0/3.1;
              }
            }
          }
        }
```

Related Documentation

- *Example: NG-VPLS Using Point-to-Multipoint LSPs*
- Next-Generation VPLS for Multicast with Multihoming Overview on page 5
- *Next-Generation VPLS Point-to-Multipoint Forwarding Overview*