Chapter 7

Configuring VPLS

This chapter describes how to configure the virtual private LAN service (VPLS) on the router, and contains the following sections:

NOTE: Before you configure VPLS, we recommend that you be thoroughly familiar with transparent bridging, Border Gateway Protocol (BGP), Multiprotocol Label Switching (MPLS), Label Distribution Protocol (LDP), BGP/MPLS virtual private networks (VPNs), and layer 2 services over MPLS. For detailed information about these protocols, see the resources listed in *Before You Configure VPLS* on page 538.

Overview

JUNOSe software enables you to configure one or more instances of VPLS, referred to as *VPLS instances*, on the router. VPLS employs an Ethernet-based layer 2 VPN to connect multiple individual LANs across a service provider's MPLS core network. The geographically dispersed multiple LANs function as a single virtual LAN.

VPLS preserves the broadcast and multicast capabilities of the physical LANs. Consequently, any broadcast or multicast traffic from a given customer end station is sent to all sites that participate in the VPLS instance.

You can use either BGP or LDP to provide signaling for VPLS, as follows:

- BGP signaling—VPLS with BGP signaling, which is referred to as BGP-based VPLS, uses BGP as the protocol that signals reachability for the VPLS domain in which the VPLS instance participates. You must configure BGP on each VPLS edge (VE) device in your topology to provide signaling for each VPLS domain. For information about how BGP signaling works, see *BGP Signaling for VPLS* on page 533. For information about configuring BGP signaling, see *Configuration Tasks for VPLS with BGP Signaling* on page 538.

- LDP signaling—VPLS with LDP signaling, which is referred to as LDP-based VPLS, uses LDP as the protocol that signals reachability for the VPLS domain in which the VPLS instance participates. You must configure LDP on each VE device in your topology to provide signaling for each VPLS domain. For information about how LDP signaling works, see *LDP Signaling for VPLS* on page 534. For information about configuring LDP signaling, see *Configuration Tasks for VPLS with LDP Signaling* on page 558.

Figure 120 illustrates an example of a simple VPLS topology. The basic topology of a VPLS network is the same regardless of whether BGP signaling or LDP signaling is used.

Figure 120: VPLS Sample Topology

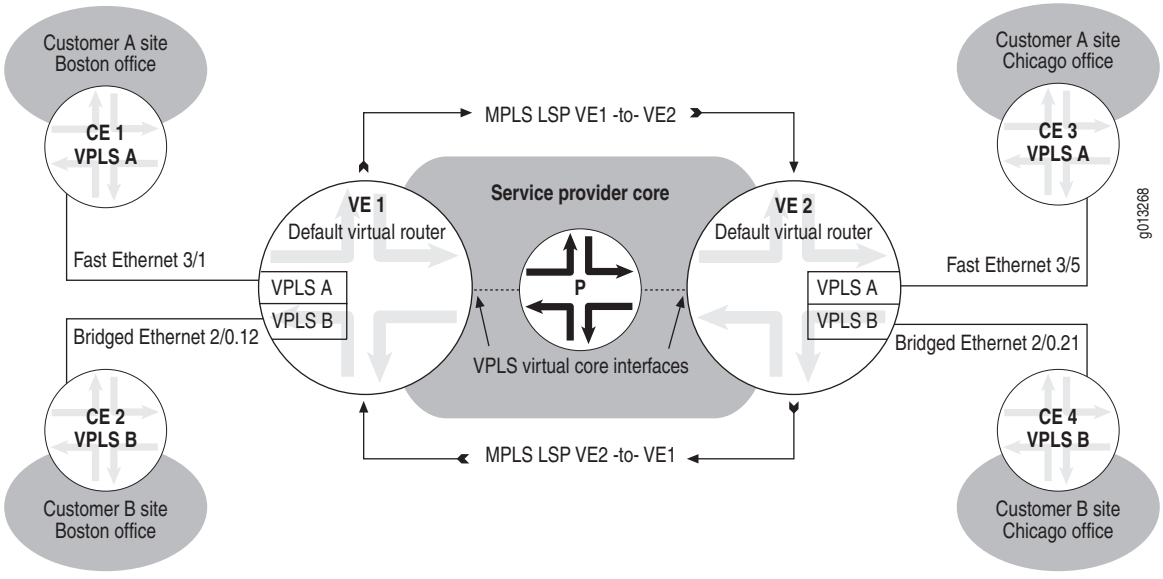

VPLS Components

As illustrated in Figure 120, a typical VPLS topology consists of the following components.

VPLS Domains

Typically, a *VPLS domain* is associated with customers who want to use Ethernet-based layer 2 VPNs to connect geographically dispersed sites in their organization across an MPLS-based service provider core, also known as an MPLS backbone. Each VPLS domain consists of the set of VPLS edge routers running the corresponding VPLS instance that participates in that domain.

Figure 120 depicts two VPLS domains: VPLS A and VPLS B. The VPLS A domain connects Customer A's Boston and Chicago offices, and consists of VPLS edge routers VE 1 and VE 2, each of which runs a VPLS instance named vplsA. Similarly, the VPLS B domain connects Customer B's Boston and Chicago offices, and consists of VPLS edge routers VE 1 and VE 2, each of which also runs a VPLS instance named vplsB.

Customer Edge Devices

Figure 120 on page 530 shows four customer edge (CE) devices: CE 1, CE 2, CE 3, and CE 4. Each CE device is located at the edge of a customer site, and participates in one or more VPLS domains. In the sample topology, CE 1 and CE 3 are members of the VPLS A domain, and CE 2 and CE 4 are members of the VPLS B domain.

A CE device can be a single host, a switch, or, most typically, a router. Each CE device is directly connected to a VPLS edge router by means of an Ethernet or bridged Ethernet network interface, but does not run VPLS. From the perspective of the CE device, the entire VPLS network appears to be a single layer 2 switch that can switch layer 2 packets, learn and filter on media access control (MAC) addresses, and flood packets that have unknown MAC destination addresses (DAs).

VPLS Edge Devices

In a VPLS configuration, E-series routers function as VPLS edge (VE) devices, which are also referred to as VE routers or, simply, VEs. A VE router is analogous to a provider edge (PE) router in BGP, LDP, and MPLS configurations, and performs similar functions.

Figure 120 on page 530 depicts two VE routers: VE 1, which is the local router, and VE 2, which is the remote router located at the other side of the service provider core. Each VE router must have a VPLS instance configured for each VPLS domain in which it participates. Consequently, the sample topology comprises a total of four separate VPLS instances: instances vplsA and vplsB configured on VE 1, and instances vplsA and vplsB configured with matching route target values on VE 2.

Each VPLS instance configured on the router is associated with two types of interfaces, also known as ports. The CE-facing interface is an Ethernet or bridged Ethernet network interface that directly connects the VE router to each CE device. The VPLS virtual core interface, although not an actual physical interface, is automatically generated by the router for each VPLS instance and represents all of the MPLS tunnels from the router to the remote VE devices. The router encapsulates Ethernet frames from the CE device in an MPLS packet and then forwards the encapsulated frames to the service provider core through the provider (P) router. This encapsulation is identical to Martini encapsulation for Ethernet layer 2 services over MPLS.

Each VE router in the sample topology has a total of two network interfaces and two VPLS virtual core interfaces configured, one interface of each type per VPLS instance.

VPLS and Transparent Bridging

A single VPLS instance is analogous to a distributed learning bridge (also known as a bridge group) used for transparent bridging, and performs similar functions. In effect, a VPLS instance is a new or existing bridge group that has additional VPLS attributes configured.

A bridge group is a collection of bridge interfaces stacked on Ethernet layer 2 interfaces to form a broadcast domain. Similarly, a VPLS instance is a collection of network interfaces stacked on Ethernet layer 2 interfaces that transmits packets between the router, or VE device, and the CE device located at the edge of the customer's network. In addition, the VPLS virtual core interface enables a VPLS instance to forward traffic not only between bridge interfaces, like a bridge group, but also between a bridge (network) interface and the service provider core.

Like a bridge group, each VPLS instance maintains its own set of forwarding tables and filters that enables it to learn the network topology by examining the media access control (MAC) source address of every incoming packet. The VPLS instance then creates an entry in its forwarding table that includes the MAC address and associated network interface where the packet was received. For traffic on the VPLS virtual core interface, the VPLS instance captures additional information that includes an outgoing MPLS label used to reach the remote site and an incoming MPLS label used to process traffic received from the remote site.

Table 44 through Table 47 represent the forwarding tables on VE 1 and VE 2 for the sample VPLS topology illustrated in Figure 120 on page 530.

Table 44: VPLS Forwarding Table on VE 1 for VPLS A

Interface	MAC Address	Outgoing Label	Received Label
Fast Ethernet 3/1	1a1a.1a1a.1a1a	–	–
VPLS virtual core interface	3a3a.3a3a.3a3a	18	324

Table 45: VPLS Forwarding Table on VE 1 for VPLS B

Interface	MAC Address	Outgoing Label	Received Label
Bridged Ethernet 2/0.12	2b2b.2b2b.2b2b	–	–
VPLS virtual core interface	4b4b.4b4b.4b4b	25	526

Table 46: VPLS Forwarding Table on VE 2 for VPLS A

Interface	MAC Address	Outgoing Label	Received Label
Fast Ethernet 3/5	3a3a.3a3a.3a3a	–	–
VPLS virtual core interface	1a1a.1a1a.1a1a	42	107

Table 47: VPLS Forwarding Table on VE 2 for VPLS B

Interface	MAC Address	Outgoing Label	Received Label
Bridged Ethernet 2/0.21	4b4b.4b4b.4b4b	–	–
VPLS virtual core interface	2b2b.2b2b.2b2b	63	872

BGP Signaling for VPLS

BGP multiprotocol extensions (MP-BGP) enable BGP to support IPv4 services such as BGP/MPLS VPNs, which are sometimes known as RFC 2547bis VPNs. VPLS with BGP signaling is actually a BGP-MPLS application that has much in common with BGP/MPLS VPNs.

The procedures for configuring BGP signaling for BGP/MPLS VPNs and for VPLS are similar except, for VPLS, you must configure both of the following BGP address families:

- L2VPN—The L2VPN address family enables you to configure the VE router to exchange layer 2 network layer reachability information (NLRI) for all VPLS instances. Optionally, you can use the **signaling** keyword for the L2VPN address family to specify BGP signaling of L2VPN reachability information. Currently, you can omit the **signaling** keyword with no adverse effects.

- VPLS—The VPLS address family enables you to configure the VE router to exchange layer 2 NLRI for a specified VPLS instance.

BGP can exchange information in a VPLS topology within these address families. Specifically, BGP builds a full mesh of label-switched paths (LSPs) among all of the VPLS instances on each of the VPLS edge routers participating in a particular VPLS domain.

Configuring BGP Signaling on page 550 describes one way to configure BGP signaling for VPLS, but does not provide complete details about configuring BGP and BGP/MPLS VPNs. See *Chapter 1, Configuring BGP Routing* for information about configuring BGP, and ***Chapter 3, Configuring BGP-MPLS Applications*** for information about configuring BGP/MPLS VPNs.

LDP Signaling for VPLS

When you configure VPLS with LDP signaling, LDP supports a full mesh of pseudowires among the participating VE routers. This is analogous to BGP signaling, in which BGP builds a full mesh of label-switched paths (LSPs) among all of the VPLS instances on each of the VE routers participating in a particular VPLS domain.

Targeted Sessions

LDP establishes targeted sessions to the remote VEs configured at the edge of the service provider's MPLS core network. The number of targeted sessions supported for a local VE router is equal to the total number of other VE routers that participate in the VPLS instances configured on the local VE. As is the case with Martini encapsulation for Ethernet layer 2 services over MPLS, a targeted session to a remote VE can have many pseudowires that terminate at the same remote VE.

To enable LDP to establish targeted sessions with remote VEs across the MPLS core, you must issue both the **mpls ldp vpls-id** command to configure a VPLS identifier for the VPLS instance, and the **mpls ldp vpls neighbor** command to configure a list of neighbor (peer) addresses to which LDP can send or from which LDP can receive targeted hello messages. For more information about using these commands, see **mpls ldp vpls vpls-id** on page 561 and **mpls ldp vpls neighbor** on page 561.

PWid FEC Element TLV

LDP signaling information for VPLS is carried in a label mapping message. The label mapping message contains the Generic Label type-length-value (TLV), and the pseudowire identifier (PWid) forwarding equivalence class (FEC) element. A FEC is a group of IP packets forwarded over the same path with the same path attributes applied.

The PWid FEC element (FEC Type 128 or 0x80) contains the VPLS identifier information configured for your VPLS instance with the **mpls ldp vpls-id** command. Taken together, the pseudowire type field and the PWid field in the TLV represent a unique VPLS instance. The pseudowire type field is Ethernet to identify the pseudowires that carry Ethernet traffic for multipoint connectivity between the local and remote VEs. The PWid field is a nonzero 32-bit integer that contains the VPLS identifier, which is a globally unique identifier for a VPLS domain. All VEs that participate in the same VPLS domain must use the same VPLS identifier.

Martini encapsulation for Ethernet layer 2 services over MPLS also uses the PWid FEC Element TLV. As a result, the PWid for Martini configurations must not be the same as the VPLS identifier configured for a VPLS instance. To prevent this conflict from occurring, the JUNOSe software displays an error and rejects the configuration if you attempt to configure the same value for the Martini PWid and the VPLS identifier.

Configuration Tasks for VPLS with LDP Signaling on page 558 describes how to configure LDP signaling for VPLS, but does not provide complete details about configuring LDP in an MPLS network. For complete information about using LDP in an MPLS network, see ***Chapter 2, Configuring MPLS***.

Supported Features

The JUNOSe implementation of VPLS provides the following features:

- Single-level VPLS hierarchy within a single autonomous system (AS) using MPLS tunneling technology for the core
- Support for the following types of network interfaces between the VE router and the CE device:
  - Bridged Ethernet over ATM1483 subinterfaces
  - Fast Ethernet
  - Gigabit Ethernet
  - 10-Gigabit Ethernet
  - VLAN and S-VLAN subinterfaces over bridged Ethernet, Fast Ethernet, Gigabit Ethernet, or 10-Gigabit Ethernet interfaces
- Autodiscovery of VPLS instance members using MP-BGP
- VPLS signaling using MP-BGP to set up and tear down the pseudowires that constitute a VPLS instance
- VPLS signaling using LDP and the PWid FEC element (FEC Type 128) to set up and tear down the pseudowires that constitute a VPLS instance
- Interworking of the VPLS instance and the VPN routing and forwarding instance (VRF) using an external cable connection
- Class of service (CoS)
- Inter-AS option A, inter-AS option B, and inter-AS option C services
- Minimal filtering and policing support

Platform Considerations

You can configure VPLS on the following E-series routers:

- E120 router
- E320 router
- ERX-1440 router
- ERX-1410 router
- ERX-710 router
- ERX-705 router
- ERX-310 router

Module Requirements

You can configure VPLS network interfaces on all E-series module combinations that support transparent bridging. The VPLS virtual core interface can exist on all E-series module combinations that support MPLS tunnels.

For information about the modules that support VPLS network interfaces and VPLS virtual core interfaces on ERX-14xx models, ERX-7xx models, and ERX-310 routers:

- See *ERX Module Guide, Table 1, Module Combinations* for detailed module specifications.
- See *ERX Module Guide, Appendix A, Module Protocol Support* for information about the modules that support VPLS network interfaces and VPLS virtual core interfaces.

For information about the modules that support VPLS network interfaces and VPLS virtual core interfaces on E120 routers and E320 routers:

- See *E120 and E320 Module Guide, Table 1, Modules and IOAs* for detailed module specifications.
- See *E120 and E320 Module Guide, Appendix A, IOA Protocol Support* for information about the modules that support VPLS network interfaces and VPLS virtual core interfaces.

Interface Specifiers

The configuration task examples in this chapter use the *slot*/*port*[.*subinterface*] format to specify the physical interface on which to configure a VPLS network interface. However, the interface specifier format that you use depends on the router that you are using.

For ERX-7xx models, ERX-14xx models, and ERX-310 routers, use the *slot*/*port*[.*subinterface*] format. For example, the following command specifies Fast Ethernet subinterface 6 on port 2 of the I/O module installed in slot 3 of an ERX-7xx model, ERX-14xx model, or ERX-310 router.

```
host1(config)#interface fastEthernet 3/2.6
```

For E120 and E320 routers, use the *slot*/*adapter*/*port*[.*subinterface*] format, which includes an identifier for the bay in which the I/O adapter (IOA) resides. In the software, the upper IOA bay is identified as adapter 0; the lower IOA bay is identified as adapter 1. For example, the following command specifies Gigabit Ethernet subinterface 20 on port 1 of the IOA installed in the upper adapter bay (adapter 0) of slot 4 in an E320 router.

```
host1(config)#interface gigabitEthernet 4/0/1.20
```

For more information about supported interface types and specifiers on E-series routers, see *Interface Types and Specifiers* in *JUNOSe Command Reference Guide, About This Guide*.

References

For more information about VPLS, consult the following resources:

- *JUNOSe Release Notes*, *Appendix A*, *System Maximums*—Refer to the Release Notes corresponding to your software release for information about the maximum values supported for VPLS configuration.
- RFC 3036—LDP Specification (January 2001)
- RFC 4447—Pseudowire Setup and Maintenance Using the Label Distribution Protocol (LDP) (April 2006)
- RFC 4762—Virtual Private LAN Service (VPLS) Using Label Distribution Protocol (LDP) Signaling (January 2007)
- Virtual Private LAN Service—draft-ietf-l2vpn-vpls-bgp-05.txt (October 2005 expiration)

NOTE: IETF drafts are valid for only 6 months from the date of issuance. They must be considered as works in progress. Please refer to the IETF Web site at http://www.ietf.org for the latest drafts.

Before You Configure VPLS

The JUNOSe implementation of VPLS uses features of transparent bridging, BGP, MPLS, LDP, BGP/MPLS VPNs, and layer 2 services over MPLS. We recommend that you have a thorough understanding of these protocols before you configure and use VPLS in your network.

For more information about configuring transparent bridging, BGP, MPLS, LDP, BGP/MPLS VPNs, and layer 2 services over MPLS, see all of the preceding chapters in this guide, as well as the following chapter:

- *JUNOSe Link Layer Configuration Guide, Chapter 10, Configuring Transparent Bridging*

For more information about configuring the layer 2 interfaces that support VPLS, see the following chapters:

- *JUNOSe Physical Layer Configuration Guide, Chapter 5, Configuring Ethernet Interfaces*
- *JUNOSe Link Layer Configuration Guide, Chapter 9, Configuring Bridged Ethernet*

Configuration Tasks for VPLS with BGP Signaling

To configure VPLS with BGP signaling on the VE router:

1. Configure a single instance of VPLS, known as a VPLS instance, on the VE router for each VPLS domain in which the router participates.
2. (Optional) Configure optional attributes for the VPLS instance.
3. Configure network interfaces to connect the VE router to each CE device.
4. (Optional) Configure nondefault subscriber policies for the VPLS network interface.
5. Configure a loopback interface and assign a router ID that uses the IP address of the loopback interface.
6. Configure MPLS label-switched paths (LSPs) to connect local and remote VE routers.
7. Set up BGP signaling on the autonomous system configured to signal reachability for this VPLS instance.

The following sections describe how to perform each of these tasks. See *VPLS Configuration Example with BGP Signaling* on page 554 for a detailed sample configuration.

NOTE: For information about the maximum values that the router supports for VPLS configuration, see *JUNOSe Release Notes*, *Appendix A*, *System Maximums*.

Configuring VPLS Instances with BGP Signaling

You must configure a VPLS instance for each VPLS domain in which the router participates. From a configuration standpoint, a VPLS instance is simply a new or existing bridge group that you configure with additional VPLS attributes.

Table 48 lists the commands that you use to configure a basic VPLS instance, as described in this section. For more information about the syntax of each command, see the *JUNOSe Command Reference Guide A to M*.

Table 48: Commands to Configure Basic VPLS Instances

bridge vpls rd	**bridge vpls site-range**
bridge vpls route-target	**bridge vpls transport-virtual-routers**
bridge vpls site-name site-id	–

To configure a basic VPLS instance with BGP signaling on the VE router:

1. From Global Configuration mode, create the VPLS instance by specifying the transport virtual router for this instance.

   If the bridge group you specify (customer1 in the example that follows this procedure) already exists on the router, issuing this command causes the bridge group to become a VPLS instance.

NOTE: To configure a VPLS instance, you must issue the **bridge vpls transport-virtual-router** command before you issue any of the other **bridge vpls** commands in this procedure. If the **bridge vpls transport-virtual-router** command is not issued first, the other **bridge vpls** commands fail.

2. Specify the maximum number of customer sites that can participate in the VPLS domain represented by the VPLS instance. (By default, a VPLS domain must consist of at least one site.)

3. Specify a name and unique identifier for the customer site that belongs to the VPLS instance.

   The site ID value must be greater than zero and be unique across the VPLS domain.

4. Specify the unique, two-part route distinguisher (RD) for the VPLS instance.

   Certain rules apply when you configure the route distinguisher for a VPLS instance. For more information, see **bridge vpls rd** on page 540.

   In the example that follows, the first number in the route distinguisher (100) is the number of the AS. The second number in the route distinguisher (11) uniquely identifies the VPLS instance within the AS.

5. Create or add a route target to the import and export lists of VPN extended communities for this VPLS instance.

   The VE router uses the lists of VPN extended communities to determine which routes are imported by this VPLS instance.

In the following example, the first number in the route target (100) is the number of the AS in which the extended community resides. The second number in the route target (1) uniquely identifies the extended community. This example uses the **both** keyword to add the route target to both the import list and the export list for this VPLS instance.

```
! Configure a VPLS instance named customer1.
host1(config)#bridge customer1 vpls transport-virtual-router vr1
host1(config)#bridge customer1 vpls site-range 15
host1(config)#bridge customer1 vpls site-name westford site-id 1
host1(config)#bridge customer1 vpls rd 100:11
host1(config)#bridge customer1 vpls route-target both 100:1
```

bridge vpls rd

- Use to specify a unique two-part route distinguisher to identify a VPLS instance that uses BGP signaling.
- Specify the route distinguisher in the format *number1*:*number2*, where:
  - *number1*—An AS number or an IP address
  - *number2*—A unique integer that is 32 bits if *number1* is an AS number, or 16 bits if *number1* is an IP address
- After you set the route distinguisher for a VPLS instance, you cannot change it for that VPLS instance. To change the route distinguisher, you must either remove the transport virtual router configuration from the VPLS instance or delete the VPLS instance from the router. You can then reconfigure the VPLS instance with a new route distinguisher.
- Multiple VPLS instances that use the same transport virtual router cannot have the same route distinguisher. Conversely, multiple VPLS instances that use different transport virtual routers can have the same route distinguisher.

  For example, the following commands configure the transport virtual router for each of three VPLS instances: vplsA, vplsB, and vplsC. The transport virtual router for both vplsA and vplsC is vr1, and the transport virtual router for vplsB is vr2.

  ```
  host1(config)#bridge vplsA vpls transport-virtual-router vr1
  host1(config)#bridge vplsB vpls transport-virtual-router vr2
  host1(config)#bridge vplsC vpls transport-virtual-router vr1
  ```

  Because vplsA and vplsC use the same transport virtual router, vr1, you cannot assign them the same route distinguisher. Consequently, the following operation fails, and the router displays an error message.

  ```
  host1(config)#bridge vplsA vpls rd 1.1.1.1:10
  host1(config)#bridge vplsC vpls rd 1.1.1.1:10
  % Unable to set VPLS route distinguisher (can't re-use the route-distinguisher)
  ```

However, both vplsA and vplsB can use the same route distinguisher because their transport virtual routers are different. Consequently, the following commands are valid.

```
host1(config)#bridge vplsA vpls rd 1.1.1.1:10
host1(config)#bridge vplsB vpls rd 1.1.1.1:10
```

- Example

  ```
  host1(config)#bridge vplsA vpls rd 100:20
  ```

- Because you cannot change or remove the route distinguisher for a VPLS instance after you set it, issuing the **no** version fails and causes the router to display the following error message:

  ```
  host1(config)#no bridge vplsB vpls rd
  % Unable to set VPLS route distinguisher (can't re-use the route-distinguisher)
  ```

- Because you cannot change or remove the route distinguisher for a VPLS instance after you set it, issuing the **no** version fails, and causes the router to display an error message.

bridge vpls route-target

- Use to create or add a route target to the import list, to the export list, or to both the import and export lists of VPN extended communities for a VPLS instance that uses BGP signaling.
- The VE router uses the lists of VPN extended communities to determine which routes are imported into the BGP address family for the specified VPLS instance. A route is imported when both of the following conditions are met:
  - An update message with a route-target export list advertises a route.
  - That list contains at least one route target that matches a route target in the VPLS instance's route-target import list.
- To add the route target to both the VPLS instance's import list and export list of VPN extended communities, use the **both** keyword. This is the recommended setting for a VPLS instance.
- To add the route target only to the VPLS instance's import list of VPN extended communities, use the **import** keyword.
- To add the route target only to the VPLS instance's export list of VPN extended communities, use the **export** keyword.
- To identify the extended community, specify the route target in the format *number1*:*number2*, where:
  - *number1*—An AS number or an IP address
  - *number2*—A unique integer that is 32 bits if *number1* is an AS number, or 16 bits if *number1* is an IP address
- Example

  ```
  host1(config)#bridge vplsA vpls route-target import 100:1
  ```

- Use the **no** version to remove a route target from the specified VPN extended communities list.

bridge vpls site-name site-id

- Use to configure a site name and a unique site identifier for a VPLS instance that uses BGP signaling.
- You must specify both of the following:
  - A site name of up to 128 alphanumeric characters
  - A site identifier that is an unsigned 16-bit integer greater than zero; the site identifier must be unique across the VPLS domain associated with this VPLS instance
- Example

  host1(config)#**bridge vplsA vpls site-name newyork site-id 5**

- Use the **no** version to remove the site name and the site identifier from the VPLS instance.

bridge vpls site-range

- Use to configure the maximum number of customer sites that can participate in the specified VPLS domain that uses BGP signaling, in the range 1–65534.
- A VPLS domain must consist of at least one site.
- Example

  host1(config)#**bridge vplsA vpls site-range 10**

- Use the **no** version to restore the default site range value, 1.

bridge vpls transport-virtual-router

- Use to configure the transport virtual router for a VPLS instance. The transport virtual router specifies the name of the virtual router on which the BGP instance that signals reachability for this VPLS instance is configured.
- Issuing this command creates a new VPLS instance or causes an existing bridge group configured on the router to become a VPLS instance.
- You must issue the **bridge vpls transport-virtual-router** command before you issue any other **bridge vpls** commands to configure VPLS attributes. If the **bridge vpls transport-virtual-router** command is not issued first, the other **bridge vpls** commands fail.
- Example

  host1(config)#**bridge vplsA vpls transport-virtual-router vr1**

- Use the **no** version to remove the VPLS instance from the router and to clear any attributes for the deleted VPLS instance configured with the **bridge vpls rd**, **bridge vpls route-target**, **bridge vpls site-name site-id**, and **bridge vpls site-range** commands.

Configuring Optional Attributes for VPLS Instances

After you create a basic VPLS instance, you can configure one or more optional attributes to manage the MAC address entries in the VPLS instance's forwarding table, or to enable SNMP link status processing. To configure these attributes, you use the same transparent bridging commands that you use to configure bridge groups that do not function as VPLS instances.

Table 49 lists the optional attributes and associated commands that you can configure for VPLS instances. For more information about using these commands, see *Configuring Optional Bridge Group Attributes* in *JUNOSe Link Layer Configuration Guide, Chapter 10, Configuring Transparent Bridging*.

Table 49: Commands to Configure Optional Attributes for VPLS Instances

Attribute	Command
Enable or disable the VPLS instance's ability to acquire dynamically learned MAC addresses	**bridge acquire**
Enable or disable the VPLS instance's ability to filter (forward or discard) frames with a particular MAC source or destination address	**bridge address**
Set the aging time of a dynamic (learned) entry in the VPLS instance's forwarding table	**bridge aging-time**
Set the maximum number of dynamic MAC addresses that a VPLS instance can learn	**bridge learn**
Enable SNMP link status processing for the VPLS instance	**bridge snmp-trap link-status**

bridge acquire

- Use to configure a VPLS instance to acquire dynamically learned MAC addresses.
- Example

  host1(config)#**bridge vplsB acquire**

- Use the **no** version to prevent the VPLS instance from acquiring dynamically learned MAC addresses and to limit forwarding only to those nodes that have a statically configured address entry in the forwarding table.

bridge address

- Use to enable a VPLS instance to filter (forward or discard) frames based on a specific MAC address, and to add static (nonlearned) address entries to the forwarding table.
- You cannot create a static MAC address entry to forward to the VPLS virtual core interface.
- Example 1—Forwards frames destined for the node with MAC address 0090.1a40.4c7c out the specified Gigabit Ethernet interface

  host1(config)#**bridge vplsA address 0090.1a40.4c7c forward gigabitEthernet 3/0.1**

- Example 2—Drops frames sent from or destined for the node with MAC address 1011.22b2.333c

  host1(config)#**bridge vplsB address 1011.22b2.333c discard**

- Use the **no** version to remove the static MAC address entry from the forwarding table.

bridge aging-time

- Use to set the length of time, in the range 1–1000000 seconds, that a dynamic (learned) MAC address entry can remain in the forwarding table of the specified VPLS instance before expiring.
- Example

  host1(config)#**bridge vplsB aging-time 1000**

- Use the **no** version to restore the default aging time, 300 seconds.

bridge learn

- Use to set the maximum number of dynamic MAC address entries that the specified VPLS instance can learn, in the range 0–64000.
- Example

  host1(config)#**bridge vplsB learn 2500**

- Use the **no** version to restore the default value, 0 learned addresses. This default implies that there is no maximum number of learned entries for an individual VPLS instance; that is, an individual VPLS instance can learn an unlimited number of MAC addresses, up to the maximum number that the router supports.

bridge snmp-trap link-status

- Use to enable SNMP link status processing for all network interfaces associated with the specified VPLS instance.
- Example

  host1(config)#**bridge vplsB snmp-trap link-status**

- Use the **no** version to disable SNMP link status processing for all network interfaces associated with the VPLS instance.

Configuring VPLS Network Interfaces

You must configure one of the following types of Ethernet or bridged Ethernet network interfaces to transmit packets between the VE router and each CE device to which the VE is connected:

- Bridged Ethernet over ATM 1483 subinterfaces
- Fast Ethernet
- Gigabit Ethernet

- 10-Gigabit Ethernet
- VLAN and S-VLAN subinterfaces over bridged Ethernet, Fast Ethernet, Gigabit Ethernet, or 10-Gigabit Ethernet interfaces

To configure a network interface for a VPLS instance:

1. From Global Configuration mode, select the interface that you want to assign to the VPLS instance.
2. From Interface Configuration mode or Subinterface Configuration mode, issue the **bridge-group** command to assign the interface to the specified VPLS instance.

   Issuing this command with no optional keywords configures the network interface as a subscriber (client) interface by default.
3. (Optional) Configure one or more the following optional attributes for the network interface. You must issue a separate **bridge-group** command for each attribute.
   - Configure the interface as a trunk (server) interface. For more information about the differences between a subscriber (client) interface and a trunk (server) interface, see *Configuring Subscriber Policies for VPLS Network Interfaces* on page 546.
   - Set the maximum number of MAC addresses that the network interface can learn.
   - Enable SNMP link status processing only for the specified network interface in the VPLS instance.

```
! Configure Gigabit Ethernet 3/0 and assign it to VPLS instance customer1
! as a trunk interface
host1(config)#interface gigabitEthernet 3/0
host1(config-if)#bridge-group customer1
host1(config-if)#bridge-group customer1 subscriber-trunk
host1(config-if)#bridge-group customer1 learn 100
host1(config-if)#bridge-group customer1 snmp-trap link-status
```

bridge-group

- Use to assign a network interface to a specified VPLS instance.
- To assign the network interface to the VPLS instance as a subscriber (client) interface, which is the default, use the **bridge-group** command without any optional keywords.
- To configure the network interface as a trunk (server) interface, use the **subscriber-trunk** keyword.
- To set the maximum number of MAC addresses that the network interface can learn, use the **learn** keyword and specify a value in the range 0–64000. A value of 0 indicates that an individual network interface can learn an unlimited number of MAC addresses, up to the maximum number that the router supports.

- To enable SNMP link status processing for the specified network interface, use the **snmp-trap link-status** keyword.
- Examples

  host1(config-subif)#**bridge-group vpls1 subscriber-trunk**
  host1(config-subif)#**bridge-group vpls1 learn 500**
  host1(config-subif)#**bridge-group vpls1 snmp-trap link-status**

- Use the **no** version to remove the network interface from the VPLS instance and restore the default value for the interface type (subscriber client), maximum number of learned MAC addresses (0), or SNMP link status processing (disabled).

Configuring Subscriber Policies for VPLS Network Interfaces

The router associates a VPLS network interface, as it does a bridge group interface, with a default subscriber policy that enables intelligent flooding of packets within a VPLS domain. This section describes how subscriber policies work and explains some important considerations when you use subscriber policies for VPLS instances.

Network Interface Types

VPLS instances, like bridge groups, support two types of network interfaces:

- Subscriber (client)—A subscriber (client) interface is downstream from the traffic flow; that is, the traffic flow direction is from the server (trunk) to the client (subscriber). This is the default network interface type for both VPLS instances and bridge groups.
- Trunk (server)—A trunk (server) interface is upstream from the traffic flow; that is, the traffic flow direction is from the client (subscriber) to the server (trunk). To configure a trunk interface, you must specify the **subscriber-trunk** keyword as part of the **bridge-group** command. The VPLS virtual core interface always acts as a trunk interface, and cannot be configured as a subscriber interface.

Default Subscriber Policies

Each network interface is associated with a default subscriber policy for that interface type. The subscriber policy is a set of forwarding and filtering rules that defines how the specified interface handles various packet or attribute types, as follows:

- For each packet type listed in Table 50, the subscriber policy specifies whether the network interface permits (forwards) or denies (filters or drops) packets of that type.
- For the relearn attribute, the subscriber policy specifies whether the network interface can relearn a MAC address entry on a different interface from the one initially associated with this entry in the forwarding table. Permit indicates that relearning is allowed; deny indicates that relearning is prohibited.

Table 50 lists the default values for each packet or attribute type defined in the policies for subscriber interfaces and trunk interfaces. The default subscriber policy differs in one way from the default trunk policy: broadcast packets and packets with unknown unicast destination addresses (DAs) are denied in the subscriber policy and permitted in the trunk policy.

Table 50: Default Subscriber Policies for VPLS Network Interfaces

Packet/Attribute Type	Default Subscriber Policy	Default Trunk Policy
ARP	Permit	Permit
Broadcast	Deny	Permit
IP	Permit	Permit
MPLS	Permit	Permit
Multicast	Permit	Permit
PPPoE	Permit	Permit
Relearn	Permit	Permit
Unicast (user-to-user)	Permit	Permit
Unknown unicast DA	Deny	Permit
Unknown protocol	Permit	Permit

Modifying Subscriber Policies

For a network interface configured as a subscriber (client) interface, you can modify the default subscriber policy to change the default permit or deny value for one or more of the packet or attribute types listed in Table 50.

You cannot, however, change the default trunk policy for a network interface configured as a trunk interface or for the VPLS virtual core interface. Trunk interfaces and the VPLS virtual core interface always use the default trunk policy, which forwards packets of all types and permits relearning.

Table 51 lists the commands that you can use to modify subscriber policies for subscriber (client) interfaces associated with either a VPLS instance or a standard bridge group. For information about using these commands, see *Configuring Subscriber Policies* in *JUNOSe Link Layer Configuration Guide, Chapter 10, Configuring Transparent Bridging*.

Table 51: Commands to Configure Subscriber Policies

arp	**pppoe**
bridge subscriber-policy	**relearn**
broadcast	**subscriber-policy**
ip	**unicast**
mpls	**unknown-destination**
multicast	**unknown-protocol**

Considerations for VPLS Network Interfaces

When you configure network interfaces for a VPLS instance, you must ensure that the subscriber policy in effect for the interface is appropriate for your network configuration.

To ensure that the network interface permits relearning and forwards (permits) packets for all of the protocol types listed in Table 50 on page 547, be sure to configure the network interface as a trunk (server) interface so that it always uses the default trunk policy. For example, the following commands associate a 10-Gigabit Ethernet interface with a VPLS instance named vplsBoston, and configure the interface as a trunk.

```
host1(config)#interface tenGigabitEthernet 4/0/1
host1(config-if)#bridge-group vplsBoston subscriber-trunk
```

If you configure a VPLS network interface as a subscriber (client) interface, use care if you modify the default subscriber policy in effect for that interface. For example, if you use the **arp** command to change the default value for ARP packets from permit (forward) to deny (filter or drop), make sure you also use the **bridge address** command to add the appropriate static (nonlearned) ARP entry to the forwarding table. If an ARP entry expires from the forwarding table and the subscriber policy is configured to deny ARP packets, the router cannot properly forward subsequent ARP packets.

Configuring the Loopback Interface and Router ID for BGP Signaling

To establish a BGP session, BGP uses the IP address of the outgoing interface towards the BGP peer as the update source IP address for the TCP connection over which the BGP session runs. Typically, you configure a loopback interface as the update source interface because a loopback interface is inherently stable.

After you configure the loopback interface, you use the **ip router-id** command to assign a router ID to uniquely identify the router within a BGP AS. The router ID is the IP address of the loopback interface.

To configure the loopback interface and router ID on the VE router:

1. Configure a loopback interface on the VE router and assign it an IP address.

2. Assign the router ID using the IP address you configured for the loopback interface.

```
! Configure a loopback interface on the VE and assign it an IP address.
host1(config)#interface loopback 0
host1(config-if)#ip address 10.3.3.3 255.255.255.255
host1(config-if)#exit
! Assign the router ID for the VE using the IP address of the loopback interface.
host1(config)#ip router-id 10.3.3.3
```

interface loopback

- Use to access and configure a loopback interface.
- You can use a loopback interface to provide a stable IP address that can minimize the impact if a physical interface goes down.
- Example

  host1(config)#**interface loopback 0**
  host1(config-if)#**ip address 10.3.3.3 255.255.255.0**

- Use the **no** version to delete the loopback interface.

ip router-id

- Use to assign a router ID, which is a unique identifier that IP routing protocols use to identify the router within an AS.
- Example

  host1(config)#**ip router-id 10.3.3.3**

- Use the **no** version to remove the router ID assignment.

Configuring MPLS LSPs

As part of a VPLS configuration, you must create MPLS label-switched paths (LSPs) to connect the local VE router and the remote VE router.

This section explains one way to create a basic MPLS configuration using the **mpls** and **mpls ldp** commands. For complete information about configuring MPLS LSPs, see *Chapter 2, Configuring MPLS*.

To configure MPLS LSPs on the VE router:

1. Enable MPLS on the virtual router.
2. Configure the core-facing interface on which you want to enable MPLS, Label Distribution Protocol (LDP), and topology-driven LSPs.
3. Enable MPLS on the core-facing interface.
4. Enable LDP and topology-driven LSPs on the core-facing interface.

```
! Enable MPLS on the default virtual router.
host1(config)#mpls
! Configure a core-facing interface between the VE and P routers,
! and assign it an IP address.
host1(config)#interface atm 5/0.100
host1(config-subif)#atm pvc 100 1 100 aal5snap 0 0 0
host1(config-subif)#ip address 192.168.5.5 255.255.255.0
! Enable MPLS on the core-facing interface.
host1(config-subif)#mpls
! Enable LDP and topology-driven LSPs on the core-facing interface.
host1(config-subif)#mpls ldp
host1(config-subif)#exit
```

mpls

- Use from Global Configuration mode to enable MPLS on a virtual router.
- Use from Interface Configuration mode or Subinterface Configuration mode to create an MPLS major interface stacked on the specified layer 2 interface, and to automatically enable MPLS on the current virtual router if it has not already been enabled.
- You cannot enable MPLS on a loopback interface.
- Example

  host1(config-if)#**mpls**

- Use the **no mpls** version from Global Configuration mode to remove MPLS from the virtual router and delete the MPLS configuration.
- Use the **no mpls** version from Interface Configuration mode to remove the MPLS major interface.

mpls ldp

- Use to enable LDP and topology-driven LSPs on an interface, using the default values (that is, using an implicit default profile).
- You cannot enable LDP and topology-driven LSPs on a loopback interface.
- Example

  host1(config-if)#**mpls ldp**

- Use the **no** version to disable LDP on an interface.

Configuring BGP Signaling

This section describes one way to configure BGP signaling for VPLS, but does not provide complete details about configuring BGP and BGP/MPLS VPNs. See *Chapter 1, Configuring BGP Routing* for information about configuring BGP, and *Chapter 3, Configuring BGP-MPLS Applications* for information about configuring BGP/MPLS VPNs.

Table 52 lists the commands discussed in this section to configure BGP signaling for VPLS. For more information about the syntax of each command, see the *JUNOSe Command Reference Guide A to M* and *JUNOSe Command Reference Guide N to Z*.

Table 52: Commands to Configure BGP Signaling for VPLS

address-family l2vpn	neighbor next-hop-self
address-family vpls	neighbor remote-as
exit-address-family	neighbor send-community
ip router-id	neighbor update-source
neighbor activate	router bgp

To configure BGP signaling for VPLS on the VE router:

1. Enable the BGP routing process in the specified AS.

   The AS number identifies the VE router to other BGP routers.

2. Configure the VE-to-VE BGP session. Use **neighbor** commands to specify the peers to which BGP advertises routes.

   In the example that follows this procedure, the BGP peer is a VE router with IP address 10.4.4.4.

   For more information about using **neighbor** commands to configure BGP, see *Chapter 1, Configuring BGP Routing* and *Chapter 3, Configuring BGP-MPLS Applications*.

3. Create the L2VPN address family to configure the router to exchange layer 2 NLRI for all VPLS instances.

   Optionally, you can use the **signaling** keyword for the L2VPN address family to specify BGP signaling of L2VPN reachability information. Currently, you can omit the **signaling** keyword with no adverse effects.

4. Activate the VE-to-VE session in the L2VPN address family. Use **neighbor** commands to configure additional address family parameters for the session.

   For more information about using **neighbor** commands to configure BGP, see *Chapter 1, Configuring BGP Routing* and *Chapter 3, Configuring BGP-MPLS Applications*.

5. Create the VPLS address family to configure the router to exchange layer 2 NLRI for each VPLS instance configured on the router.

   You must issue the **address-family vpls** command separately for each VPLS instance configured on the router.

   In the following example, two VPLS instances named customer1 and customer2 are configured on the VE router.

After you configure MPLS LSPs and BGP signaling, the router automatically generates a VPLS virtual core interface for each VPLS instance. The VPLS virtual core interface represents all of the MPLS tunnels from the router to the remote VE device.

```
! Enable BGP in AS 100.
host1(config)#router bgp 100
! Configure the VE-to-VE BGP session.
host1(config-router)#neighbor 10.4.4.4 remote-as 100
host1(config-router)#neighbor 10.4.4.4 update-source loopback 0
host1(config-router)#neighbor 10.4.4.4 next-hop-self
! Create the L2VPN address family with BGP signaling for all VPLS instances.
host1(config-router)#address-family l2vpn signaling
! Activate the VE-to-VE session in the L2VPN address family.
host1(config-router-af)#neighbor 10.4.4.4 activate
host1(config-router-af)#neighbor 10.4.4.4 next-hop-self
host1(config-router-af)#exit-address-family
```

```
! Create the VPLS address family for VPLS instance customer1.
host1(config-router)#address-family vpls customer1
host1(config-router-af)#exit-address-family
! Create the VPLS address family for VPLS instance customer2.
host1(config-router)#address-family vpls customer2
host1(config-router-af)#exit-address-family
! Return to Global Configuration mode.
host1(config-router)#exit
```

address-family l2vpn

- Use to create the L2VPN address family, which enables you to configure the router to exchange layer 2 NLRI for all VPLS instances.
- Issuing this command accesses Address Family Configuration mode.
- To specify BGP signaling of L2VPN reachability information, use the optional **signaling** keyword. Currently, you can omit the **signaling** keyword with no adverse effects.
- This command takes effect immediately.
- Example

  host1(config-router)#**address-family l2vpn signaling**

- Use the **no** version to remove the L2VPN address family for all VPLS instances.

address-family vpls

- Use to create the VPLS address family, which enables you to configure the router to exchange layer 2 NLRI only for the specified VPLS instance.
- Issuing this command accesses Address Family Configuration mode.
- This command takes effect immediately.
- Example

  host1(config-router)#**address-family vpls customer1**

- Use the **no** version to remove the VPLS address family for the specified VPLS instance.

exit-address-family

- Use to exit Address Family Configuration mode and access Router Configuration mode.
- Example

  host1(config-router-af)#**exit-address-family**

- There is no **no** version.

neighbor activate

- Use to specify neighbors that exchange routes from within the current address family.
- This command takes effect immediately.
- Example

  host1(config-router)#**neighbor 10.4.4.4 activate**

- Use the **default** version to remove the explicit configuration from the peer or peer group and to reestablish inheritance of the feature configuration.
- Use the **no** version to specify that the router not exchange routes of the current address family with the peer.

neighbor next-hop-self

- Use to force the BGP speaker to report itself as the next hop for an advertised route that it learned from a neighbor.
- Example

  host1(config-router)#**neighbor 10.4.4.4 next-hop-self**

- Use the **default** version to remove the explicit configuration from the peer or peer group and to reestablish inheritance of the feature configuration.
- Use the **no** version to disable this feature, thereby enabling next-hop processing of BGP updates.

neighbor remote-as

- Use to add an entry to the BGP neighbor table.
- Specifying a neighbor with an AS number that matches the AS number specified in the **router bgp** command identifies the neighbor as internal to the local AS. Otherwise, the neighbor is considered external.
- This command takes effect immediately.
- Example

  host1(config-router)#**neighbor 10.4.4.4 remote-as 100**

- Use the **no** version to remove an entry from the BGP neighbor table.

neighbor update-source

- Use to allow the BGP session to use the IP address of a specific operational interface as the update source address for TCP connections.
- This command takes effect immediately and automatically bounces the BGP session.
- If you specify an interface in this command and later remove the interface, this command is also removed from the router configuration.
- Example

  host1(config-router)#**neighbor 10.4.4.4 update-source loopback 0**

- Use the **no** version to restore the interface assignment to the closest interface.

router bgp

- Use to enable the BGP routing protocol on the VE router and to specify the local AS; that is, the AS to which this BGP speaker belongs.
- Issuing this command accesses Router Configuration mode.
- All subsequent BGP configuration commands are placed within the context of this router and AS; you can have only a single BGP instance per virtual router.
- Specify only one BGP AS per virtual router.
- Example

  host1(config)#**router bgp 100**

- Use the **no** version to remove the BGP process.

The following sections describe how to perform each of these tasks. See *VPLS Configuration Example with BGP Signaling* on page 554 for a detailed sample configuration.

NOTE: For information about the maximum values that the router supports for VPLS configuration, see *JUNOSe Release Notes*, *Appendix A*, *System Maximums*.

VPLS Configuration Example with BGP Signaling

The example in this section shows how to configure the VPLS topology illustrated in Figure 121. The example includes the commands for configuring VPLS on both the local E-series router (VE 1) and the remote E-series router (VE 2).

Figure 121: Topology for VPLS Configuration Example with BGP Signaling

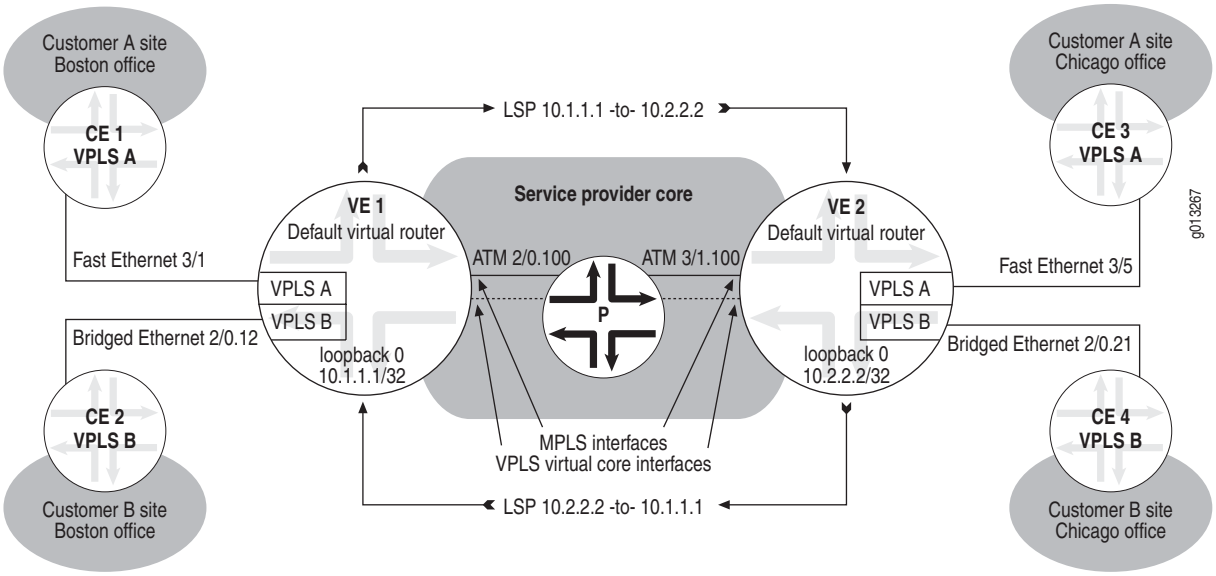

Topology Overview of VPLS with BGP Signaling

The sample topology in Figure 121 includes two VPLS domains, VPLS A and VPLS B. VPLS A connects CE 1, at the edge of Customer A's Boston site, with CE 3, at the edge of Customer A's Chicago site. Similarly, VPLS B connects CE 2, at the edge Customer B's Boston site, with CE 4, at the edge of Customer B's Chicago site.

The E-series routers in the topology, VE 1 and VE 2, each participate in both the VPLS A domain and the VPLS B domain. The example configures a total of four separate VPLS instances, one for each VPLS domain in which the VE router participates. The instances for the VPLS A domain are named vplsA, and the instances for the VPLS B domain are named vplsB.

For each VPLS instance, an Ethernet or bridged Ethernet network interface provides a connection to the associated CE device. Each VPLS instance maintains its own set of forwarding tables and filters to learn the network topology, in a manner that is similar to a bridge group used for transparent bridging.

Each VE router in the sample topology also has an ATM core-facing interface that connects it to the provider (P) router in the service provider core. You must configure MPLS LSPs on the core-facing interfaces to connect VE 1 and VE 2 through the P router across the service provider core. Finally, you must configure BGP on both VE 1 and VE 2 to provide signaling for both VPLS domains.

After you configure the bridging, MPLS, and BGP components of VPLS, the router automatically generates a VPLS virtual core interface for each VPLS instance. The VPLS virtual core interface represents all of the MPLS tunnels from the router to the remote VE device.

Configuration on VE 1 (Local VE Router)

Use the following commands on the local VE router (VE 1) to configure the VPLS topology shown in Figure 121 on page 554.

```
! Configure VPLS instance vplsA.
host1(config)#bridge vplsA vpls transport-virtual-router default
host1(config)#bridge vplsA vpls site-range 10
host1(config)#bridge vplsA vpls site-name boston site-id 1
host1(config)#bridge vplsA vpls rd 100:11
host1(config)#bridge vplsA vpls route-target both 100:1
!
! Configure VPLS instance vplsB.
host1(config)#bridge vplsB vpls transport-virtual-router default
host1(config)#bridge vplsB vpls site-range 20
host1(config)#bridge vplsB vpls site-name boston site-id 1
host1(config)#bridge vplsB vpls rd 100:12
host1(config)#bridge vplsB vpls route-target both 100:2
!
! Configure Fast Ethernet interface 3/0 between VE 1 and CE 1,
! and assign it to vplsA as a trunk interface.
host1(config)#interface fastEthernet 3/1
host1(config-if)#bridge-group vplsA subscriber-trunk
host1(config-if)#exit
!
```

```
! Configure bridged Ethernet interface 2/0.12 between VE 1 and CE 2,
! and assign it to vplsB as a trunk interface.
host1(config)#interface atm 2/0.12 point-to-point
host1(config-subif)#atm pvc 12 0 12 aal5snap 0 0 0
host1(config-subif)#encapsulation bridge1483 mac-address 0090.1a40.9991
host1(config-subif)#bridge-group vplsB subscriber-trunk
host1(config-if)#exit
!
! Configure a loopback interface on VE 1 and assign it an IP address.
host1(config)#interface loopback 0
host1(config-if)#ip address 10.1.1.1 255.255.255.255
host1(config-if)#exit
!
! Assign the router ID for VE 1 using the IP address of the loopback interface.
host1(config)#ip router-id 10.1.1.1
!
! Enable MPLS on the default virtual router.
host1(config)#mpls
!
! Configure ATM core-facing interface 2/0.100 between VE 1 and the P router,
! and assign it an IP address.
host1(config)#interface atm 2/0.100 point-to-point
host1(config-subif)#atm pvc 100 1 100 aal5snap 0 0 0
host1(config-subif)#ip address 192.168.1.1 255.255.255.0
!
! Enable MPLS, LDP, and topology-driven LSPs on the core-facing interface.
host1(config-subif)#mpls
host1(config-subif)#mpls ldp
host1(config-subif)#exit
!
! Configure BGP signaling.
host1(config)#router bgp 100
host1(config-router)#neighbor 10.2.2.2 remote-as 100
host1(config-router)#neighbor 10.2.2.2 update-source loopback 0
host1(config-router)#neighbor 10.2.2.2 next-hop-self
host1(config-router)#address-family l2vpn signaling
host1(config-router-af)#neighbor 10.2.2.2 activate
host1(config-router-af)#neighbor 10.2.2.2 next-hop-self
host1(config-router-af)#exit-address-family
host1(config-router)#address-family vpls vplsA
host1(config-router-af)#exit-address-family
host1(config-router)#address-family vpls vplsB
host1(config-router-af)#exit-address-family
host1(config-router)#exit
```

Configuration on VE 2 (Remote VE Router)

Use the following commands on the remote VE router (VE 2) to configure the VPLS topology shown in Figure 121 on page 554.

```
! Configure VPLS instance vplsA. The route target (100:1)
! matches the route target configured for vplsA on VE 1.
host2(config)#bridge vplsA vpls transport-virtual-router default
host2(config)#bridge vplsA vpls site-range 10
host2(config)#bridge vplsA vpls site-name chicago site-id 2
host2(config)#bridge vplsA vpls rd 100:21
host2(config)#bridge vplsA vpls route-target both 100:1
!
! Configure VPLS instance vplsB. The route target (100:2)
! matches the route target configured for vplsB on VE 1.
host2(config)#bridge vplsB vpls transport-virtual-router default
host2(config)#bridge vplsB vpls site-range 20
host2(config)#bridge vplsB vpls site-name chicago site-id 2
host2(config)#bridge vplsB vpls rd 100:22
host2 (config)#bridge vplsB vpls route-target both 100:2
! Configure Fast Ethernet interface 3/5 between VE 2 and CE 3,
! and assign it to vplsA as a trunk interface.
host2(config)#interface fastEthernet 3/5
host2(config-if)#bridge-group vplsA subscriber-trunk
host2(config-if)#exit
!
! Configure bridged Ethernet interface 2/0.21 between VE 2 and CE 4,
! and assign it to vplsB as a trunk interface.
host2(config)#interface atm 2/0.21 point-to-point
host2(config-subif)#atm pvc 21 0 21 aal5snap 0 0 0
host2(config-subif)#encapsulation bridge1483 mac-address 0090.1a40.9992
host2(config-subif)#bridge-group vplsB subscriber-trunk
host2(config-if)#exit
!
! Configure a loopback interface on VE 2 and assign it an IP address.
host2(config)#interface loopback 0
host2(config-if)#ip address 10.2.2.2 255.255.255.255
host2(config-if)#exit
!
! Assign the router ID for VE 2 using the IP address of the loopback interface.
host2(config)#ip router-id 10.2.2.2
!
! Enable MPLS on the default virtual router.
host2(config)#mpls
!
! Configure ATM core-facing interface 3/1.100 between VE 2 and the P router,
! and assign it an IP address.
host2(config)#interface atm 3/1.100 point-to-point
host2(config-subif)#atm pvc 100 1 100 aal5snap 0 0 0
host2(config-subif)#ip address 192.168.2.2 255.255.255.0
!
! Enable MPLS, LDP, and topology-driven LSPs on the on the core-facing interface.
host2(config-subif)#mpls
host2(config-subif)#mpls ldp
host2(config-subif)#exit
!
```

```
! Configure BGP signaling.
host2(config)#router bgp 100
host2(config-router)#neighbor 10.1.1.1 remote-as 100
host2(config-router)#neighbor 10.1.1.1 update-source loopback 0
host2(config-router)#neighbor 10.1.1.1 next-hop-self
host2(config-router)#address-family l2vpn signaling
host2(config-router-af)#neighbor 10.1.1.1 activate
host2(config-router-af)#neighbor 10.1.1.1 next-hop-self
host2(config-router-af)#exit-address-family
host2(config-router)#address-family vpls vplsA
host2(config-router-af)#exit-address-family
host2(config-router)#address-family vpls vplsB
host2(config-router-af)#exit-address-family
host2(config-router)#exit
```

Configuration Tasks for VPLS with LDP Signaling

To configure VPLS with LDP signaling on the VE router:

1. Configure a single instance of VPLS, known as a VPLS instance, on the VE router for each VPLS domain in which the router participates.
2. (Optional) Configure optional attributes for the VPLS instance.
3. Configure network interfaces to connect the VE router to each CE device.
4. (Optional) Configure nondefault subscriber policies for the VPLS network interface.
5. Set up LDP signaling for this VPLS instance to establish targeted sessions to the remote VE neighbors configured at the edge of the MPLS core network.
6. Configure a loopback interface to be associated with the targeted LDP neighbor, and assign a router ID that uses the IP address of the loopback interface.
7. Configure MPLS LSPs to connect local and remote VE routers.
8. Configure an interior gateway protocol (IGP), such as Open Shortest Path First (OSPF) or Intermediate System–to–Intermediate System (IS-IS), to enable routing within the core network.

The following sections describe how to perform each of these tasks. See *VPLS Configuration Example with LDP Signaling* on page 564 for a detailed sample configuration.

NOTE: For information about the maximum values that the router supports for VPLS configuration, see *JUNOSe Release Notes*, *Appendix A*, *System Maximums*.

Configuring VPLS Instances for LDP Signaling

As is the case with BGP signaling, when you use LDP signaling you must configure a VPLS instance for each VPLS domain in which the router participates. Unlike BGP signaling, however, configuring a VPLS instance for LDP signaling requires only that you specify the transport virtual router for this instance by issuing the **bridge vpls transport-virtual-router** command.

To configure a basic VPLS instance with LDP signaling on the VE router:

- From Global Configuration mode, create the VPLS instance by specifying the transport virtual router for this instance.

  ```
  ! Configure a VPLS instance named customer3.
  host1(config)#bridge customer3 vpls transport-virtual-router vr1
  ```

  If the bridge group you specify (customer3 in this example) already exists on the router, issuing this command causes the bridge group to become a VPLS instance.

bridge vpls transport-virtual-router

- Use to configure the transport virtual router for a VPLS instance. The transport virtual router specifies the name of the virtual router on which the LDP instance that signals reachability for this VPLS instance is configured.
- Issuing this command creates a new VPLS instance or causes an existing bridge group configured on the router to become a VPLS instance.
- Example

  ```
  host1(config)#bridge vplsC vpls transport-virtual-router vr1
  ```

- Use the **no** version to remove the VPLS instance from the router and to clear any attributes configured for the deleted VPLS instance.

Configuring Optional Attributes for VPLS Instances

After you create a basic VPLS instance for LDP signaling, you can configure one or more optional attributes for this instance that provide transparent bridging functions such as managing MAC address entries and enabling SNMP link status processing. To configure these attributes, you use the same transparent bridging commands that you use to configure VPLS instances with BGP signaling.

For instructions, see *Configuring Optional Attributes for VPLS Instances* on page 543.

Configuring VPLS Network Interfaces

VPLS instances with LDP signaling, like VPLS instances with BGP signaling, use Ethernet or bridged Ethernet network interfaces to transmit packets between the VE router and each CE device to which the VE is connected. To configure network interfaces for LDP signaling, you use the same commands and procedure that you use to configure network interfaces for BGP signaling.

For instructions, see *Configuring VPLS Network Interfaces* on page 544.

Configuring Subscriber Policies for VPLS Network Interfaces

Network interfaces for VPLS instances with LDP signaling, like network interfaces for VPLS instances with BGP signaling, are associated with two default subscriber policies, subscriber and trunk, to enable intelligent flooding of packets within a VPLS domain. To configure and use nondefault subscriber policies for LDP signaling, you use the same commands and procedures that you use to configure nondefault subscriber policies for BGP signaling.

For instructions, see *Configuring Subscriber Policies for VPLS Network Interfaces* on page 546.

Configuring LDP Signaling

LDP signaling establishes targeted sessions to the remote VEs configured at the edge of the service provider's MPLS core network. To enable LDP to establish these targeted sessions, you issue the **mpls ldp vpls-id** command to configure a VPLS identifier for the VPLS instance, and the **mpls ldp vpls neighbor** command to configure a list of neighbor (peer) addresses to which LDP can send or from which LDP can receive targeted hello messages.

This section describes how to configure LDP signaling for a VPLS network, but does not provide complete details about configuring LDP on E-series routers. For more information about LDP, see *Chapter 2, Configuring MPLS*.

Table 53 lists the commands discussed in this section to configure LDP signaling for VPLS. For more information about the syntax of each command, see the *JUNOSe Command Reference Guide A to M*.

Table 53: Commands to Configure LDP Signaling for VPLS

mpls ldp vpls neighbor	**mpls ldp vpls vpls-id**

To configure LDP signaling for VPLS on the VE router:

1. Configure the VPLS identifier, which is a globally unique identifier for each VPLS domain.

2. Configure a list of neighbor (peer) addresses to which LDP can send or from which LDP can receive targeted hello messages.

The following example configures LDP signaling for two VPLS instances named customer3 and customer4 on the VE router.

```
! Enable LDP signaling for customer3.
host1(config)#mpls ldp vpls customer3 vpls-id 3
host1(config)#mpls ldp vpls customer3 neighbor 10.3.3.3
! Enable LDP signaling for customer4.
host1(config)#mpls ldp vpls customer4 vpls-id 4
host1(config)#mpls ldp vpls customer3 neighbor 10.4.4.4
```

mpls ldp vpls neighbor

- Use to enable LDP signaling for a VPLS instance by configuring the remote VE device address of a neighbor in the VPLS domain in which this instance participates.
- If either or both LDP or MPLS are not configured on the current virtual router, issuing the **mpls ldp vpls neighbor** command creates the LDP and MPLS configurations.
- Example

  host1(config)#**mpls ldp vpls vplsC neighbor 10.1.1.1**

- Use the **no** version to delete the neighbor from the VPLS domain.

mpls ldp vpls vpls-id

- Use to configure the VPLS identifier of a VPLS instance that uses LDP as the signaling protocol.
- The VPLS identifier is a globally unique identifier for the VPLS domain, in the range 1–4294967295, that must meet the following requirements:
  - All VEs that participate in the same VPLS domain must use the same VPLS identifier.
  - The VPLS identifier configured for a VPLS instance must not be the same as the PWid for Martini configurations for Ethernet layer 2 services over MPLS.
- Example

  host1(config)#**mpls ldp vpls vplsC vpls-id 1**

- Use the **no** version to delete the VPLS identifier from the VPLS instance.

Configuring the Loopback Interface and Router ID for LDP Signaling

VPLS with LDP signaling, like VPLS with BGP signaling, requires configuration of a loopback interface. You can use a loopback interface to provide a stable IP address that can minimize the impact if a physical interface goes down. LDP uses the loopback interface as the associated interface for the targeted neighbors configured with the **mpls ldp vpls neighbor** command, as described in *Configuring LDP Signaling* on page 560.

After you configure the loopback interface, you use the **ip router-id** command to assign a router ID to uniquely identify the router within the VPLS domain. The router ID is the IP address of the loopback interface.

To configure the loopback interface and router ID on the VE router:

1. Configure a loopback interface on the VE router and assign it an IP address.
2. Assign the router ID using the IP address you configured for the loopback interface.

```
! Configure a loopback interface on the VE and assign it an IP address.
host1(config)#interface loopback 0
host1(config-if)#ip address 10.1.1.1 255.255.255.255
host1(config-if)#exit
! Assign the router ID for the VE using the IP address of the loopback interface.
host1(config)#ip router-id 10.1.1.1
```

interface loopback

- Use to access and configure a loopback interface.
- LDP uses the loopback interface as the associated interface for the targeted remote neighbors.
- Example

  ```
  host1(config)#interface loopback 0
  host1(config-if)#ip address 10.3.3.3 255.255.255.0
  ```

- Use the **no** version to delete the loopback interface.

ip router-id

- Use to assign a router ID, which is a unique identifier that IP routing protocols use to identify the router within the VPLS domain.
- Example

  ```
  host1(config)#ip router-id 10.3.3.3
  ```

- Use the **no** version to remove the router ID assignment.

Configuring MPLS LSPs

VPLS with LDP signaling, like VPLS with BGP signaling, requires configuration of MPLS LSPs to connect the local VE router and the remote VE router through the provider (P) router in the MPLS core. To configure MPLS LSPs for LDP signaling, you can use the same commands and procedure that you use to configure MPLS LSPs for BGP signaling.

For instructions, see *Configuring MPLS LSPs* on page 549.

Configuring Routing in the Core Network

After you configure the transparent bridging, LDP, and MPLS components of the VPLS network, you must configure an IGP, such as OSPF or IS-IS, on the VE to set up routing within the core MPLS network.

This section explains one way to configure OSPF to enable routing in the core network. For complete information about configuring and using OSPF, see *JUNOSe IP, IPv6, and IGP Configuration Guide, Chapter 5, Configuring OSPF*. For complete information about configuring and using IS-IS, see *JUNOSe IP, IPv6, and IGP Configuration Guide, Chapter 6, Configuring IS-IS*.

Table 54 lists the commands discussed in this section to configure OSPF. For more information about the syntax of each command, see the *JUNOSe Command Reference Guide N to Z*.

Table 54: Commands to Configure OSPF for a VPLS Network

network area	router ospf

To configure the VE to set up OSPF routing for the core MPLS network:

1. Create the OSPF routing process.
2. Create the range of IP addresses associated with the routing process and the corresponding OSPF interfaces.
3. Assign an area ID associated with each range of IP addresses.

This example configures an OSPF routing process with process ID 1, and creates two OSPF interfaces in the backbone area (area 0.0.0.0): one using IP address 1.1.1.1, and one using IP address 10.10.10.0. The **network area** commands also create the two OSPF areas if they do not already exist.

```
host1(config)#router ospf 1
host1(config-router)#network 1.1.1.1 0.0.0.0 area 0.0.0.0
host1(config-router)#network 10.10.10.0 0.0.0.255 area 0.0.0.0
```

network area

- Use to configure a range of OSPFv2 interfaces and their related area.
- If the specified range matches one or more of the IP addresses configured for IP interfaces, one or more corresponding OSPF interfaces are created and placed in the specified area.
- Create address ranges that do not overlap; you can attach only the same range of interfaces to a single area.
- Example—shows the creation of one OSPF interface in the backbone area

  ```
  host1(config-if)#ip address 2.2.2.1 255.255.0.0
  host1(config)#router ospf 2
  host1(config-router)#network 2.2.2.0 0.0.0.255 area 0
  ```

- Use the **no** version to delete OSPF interfaces, ranges, and areas.

router ospf

- Use to set a process ID for an OSPF routing process.
- The process ID can be any positive integer in the range 1–65535.
- You must assign a unique ID for each OSPF routing process.
- Example

  ```
  host1(config)#router ospf 5
  ```

- Use the **no** version to end the designated OSPF routing process.

VPLS Configuration Example with LDP Signaling

The example in this section shows how to configure the VPLS topology illustrated in Figure 122. The example includes the commands for configuring VPLS on both the local E-series router (VE 1) and the remote E-series router (VE 2).

Figure 122: Topology for VPLS Configuration Example with LDP Signaling

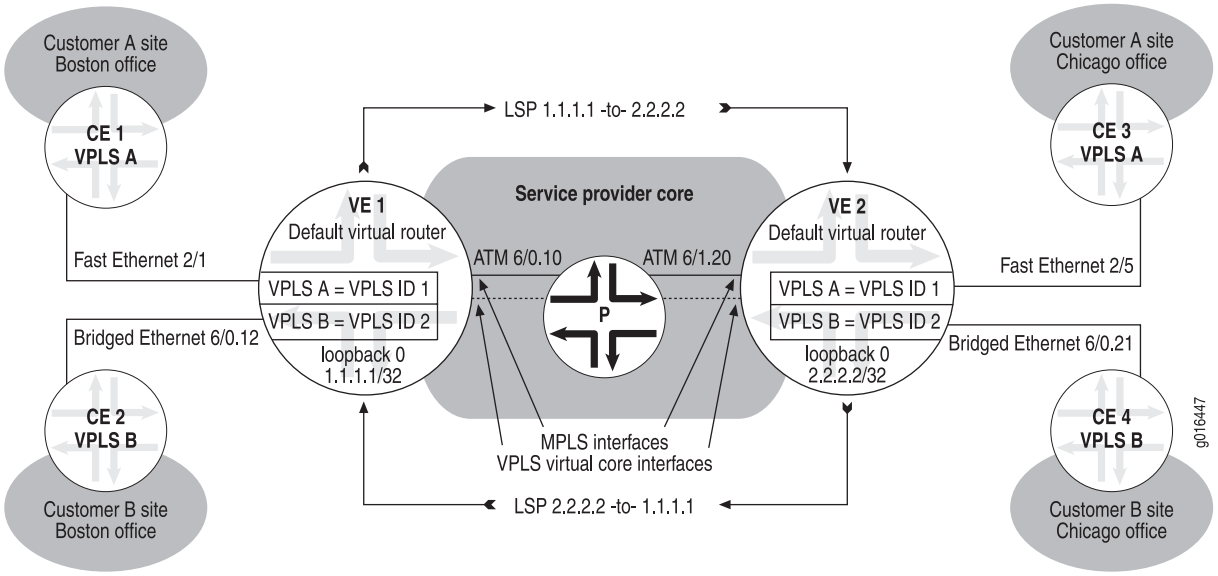

Topology Overview of VPLS with LDP Signaling

Because the basic components of a VPLS network are the same regardless of whether BGP signaling or LDP signaling is used, the sample topology shown for LDP signaling in Figure 122 on page 564 is almost identical to the sample topology shown for BGP signaling in Figure 121 on page 554. Figure 122 on page 564 includes two VPLS domains: VPLS A, which connects CE 1 and CE 3, and VPLS B, which connects CE 2 and CE 4. The local VE router, VE 1, and the remote VE router, VE 2, each participate in both the VPLS A domain and the VPLS B domain, and have one VPLS instance associated with each domain configured on each router.

Unlike a VPLS configuration with BGP signaling, a VPLS configuration with LDP signaling requires that you configure a VPLS ID for each VPLS instance to uniquely identify each VPLS domain. In the sample topology in Figure 122, instance vplsA is assigned VPLS ID 1, and instance vplsB is assigned VPLS ID 2 on both the local VE and the remote VE. You must also configure a list of remote neighbor (peer) addresses to which LDP can send or from which LDP can receive targeted hello messages. In the sample topology, the remote neighbor configured for VE 1 is VE 2 with IP address 2.2.2.2, and the remote neighbor configured for VE 2 is VE 1 with IP address 1.1.1.1.

The Ethernet and bridged Ethernet network interfaces, ATM core-facing interfaces, VPLS virtual core interfaces, and MPLS LSPs play the same role in a VPLS topology with LDP signaling as they do in a VPLS topology with BGP signaling. For more information about these components, see *Topology Overview of VPLS with BGP Signaling* on page 555.

Configuration on VE 1 (Local VE Router)

Use the following commands on the local VE router (VE 1) to configure the VPLS topology shown in Figure 122 on page 564.

```
! Configure VPLS instance vplsA.
host1(config)#bridge vplsA vpls transport-virtual-router default
!
! Configure VPLS instance vplsB.
host1(config)#bridge vplsB vpls transport-virtual-router default
!
! Configure Fast Ethernet interface 2/1 between VE 1 and CE 1,
! and assign it to vplsA as a trunk interface.
host1(config)#interface fastEthernet 2/1
host1(config-if)#bridge-group vplsA subscriber-trunk
host1(config-if)#exit
!
! Configure bridged Ethernet interface 6/0.12 between VE 1 and CE 2,
! and assign it to vplsB as a trunk interface.
host1(config)#interface atm 6/0.12 point-to-point
host1(config-subif)#atm pvc 12 0 12 aal5snap 0 0 0
host1(config-subif)#encapsulation bridge1483 mac-address 0090.1a40.9991
host1(config-subif)#bridge-group vplsB subscriber-trunk
host1(config-if)#exit
!
! Configure LDP signaling for vplsA.
host1(config)#mpls ldp vpls vplsA vpls-id 1
host1(config)#mpls ldp vpls vplsA neighbor 2.2.2.2
!
! Configure LDP signaling for vplsB.
host1(config)#mpls ldp vpls vplsB vpls-id 2
host1(config)#mpls ldp vpls vplsB neighbor 2.2.2.2
!
! Configure a loopback interface on VE 1 and assign it an IP address.
host1(config)#interface loopback 0
host1(config-if)#ip address 1.1.1.1 255.255.255.255
host1(config-if)#exit
!
! Assign the router ID for VE 1 using the IP address of the loopback interface.
host1(config)#ip router-id 1.1.1.1
!
! Configure ATM core-facing interface 6/0.10 between VE 1 and the P router,
! and assign it an IP address.
host1(config)#interface atm 6/0.10 point-to-point
host1(config-subif)#atm pvc 10 0 10 aal5snap 0 0 0
host1(config-subif)#ip address 10.10.10.1 255.255.255.0
!
```

```
! Enable MPLS, LDP, and topology-driven LSPs on the core-facing interface.
host1(config-subif)#mpls
host1(config-subif)#mpls ldp
host1(config-subif)#exit
!
! Configure OSPF routing in the core MPLS network.
host1(config)#router ospf 1
host1(config-router)#network 1.1.1.1 0.0.0.0 area 0.0.0.0
host1(config-router)#network 10.10.10.0 0.0.0.255 area 0.0.0.0
host1(config-router)#exit
```

Configuration on VE 2 (Remote VE Router)

Use the following commands on the remote VE router (VE 2) to configure the VPLS topology shown in Figure 122 on page 564.

```
! Configure VPLS instance vplsA.
host2(config)#bridge vplsA vpls transport-virtual-router default
!
! Configure VPLS instance vplsB.
host2(config)#bridge vplsB vpls transport-virtual-router default
!
! Configure Fast Ethernet interface 2/5 between VE 2 and CE 3,
! and assign it to vplsA as a trunk interface.
host2(config)#interface fastEthernet 2/5
host2(config-if)#bridge-group vplsA subscriber-trunk
host2(config-if)#exit
!
! Configure bridged Ethernet interface 6/0.21 between VE 2 and CE 4,
! and assign it to vplsB as a trunk interface.
host2(config)#interface atm 6/0.21 point-to-point
host2(config-subif)#atm pvc 21 0 21 aal5snap 0 0 0
host2(config-subif)#encapsulation bridge1483 mac-address 0090.1a40.9992
host2(config-subif)#bridge-group vplsB subscriber-trunk
host2(config-if)#exit
!
! Configure LDP signaling for vplsA.
host2(config)#mpls ldp vpls vplsA vpls-id 1
host2(config)#mpls ldp vpls vplsA neighbor 1.1.1.1
!
! Configure LDP signaling for vplsB.
host2(config)#mpls ldp vpls vplsB vpls-id 2
host2(config)#mpls ldp vpls vplsB neighbor 1.1.1.1
!
! Configure a loopback interface on VE 2 and assign it an IP address.
host2(config)#interface loopback 0
host2(config-if)#ip address 2.2.2.2 255.255.255.255
host2(config-if)#exit
!
! Assign the router ID for VE 2 using the IP address of the loopback interface.
host2(config)#ip router-id 2.2.2.2
!
```

```
! Configure ATM core-facing interface 6/1.20 between VE 2 and the P router,
! and assign it an IP address.
host2(config)#interface atm 6/1.20 point-to-point
host2(config-subif)#atm pvc 20 0 20 aal5snap 0 0 0
host2(config-subif)#ip address 20.20.20.2 255.255.255.0
!
! Enable MPLS, LDP, and topology-driven LSPs on the core-facing interface.
host2(config-subif)#mpls
host2(config-subif)#mpls ldp
host2(config-subif)#exit
!
! Configure OSPF routing in the core MPLS network.
host2(config)#router ospf 1
host2(config-router)#network 2.2.2.2 0.0.0.0 area 0.0.0.0
host2(config-router)#network 20.20.20.0 0.0.0.255 area 0.0.0.0
host2(config-router)#exit
```

Monitoring VPLS

This section describes the following tasks to manage and monitor a VPLS configuration:

NOTE: The E120 router and E320 router output for **monitor** and **show** commands is identical to output from other E-series routers, except that the E120 and E320 router output also includes information about the adapter identifier in the interface specifier (*slot/adapter/port*).

Setting Statistics Baselines

You can use the following **baseline** commands to set a statistics baseline for a VPLS instance, for a network interface associated with a VPLS instance, or for the VPLS virtual core interface associated with a VPLS instance. The router implements the baseline by reading and storing the statistics at the time the baseline is set and then subtracting this baseline whenever baseline-relative statistics are retrieved.

baseline bridge

- Use to set a statistics baseline for a specified VPLS instance.
- Example

  host1#**baseline bridge vplsA**

- There is no **no** version.

baseline bridge interface

- Use to set a statistics baseline for a particular network interface associated with a VPLS instance.
- You must specify the following:
  - *interfaceType*—One of the following interface types listed in *Interface Types and Specifiers* in *JUNOSe Command Reference Guide, About This Guide*:
    - **atm**
    - **fastEthernet**
    - **gigabitEthernet**
    - **tenGigabitEthernet**
  - *interfaceSpecifier*—Particular interface; format varies according to interface type; see *Interface Types and Specifiers* in *JUNOSe Command Reference Guide, About This Guide* for information
- Example

  host1#**baseline bridge interface gigabitEthernet 4/1**

- There is no **no** version.

baseline bridge interface vpls

- Use to set a statistics baseline for the VPLS virtual core interface associated with a VPLS instance.
- Example

  host1#**baseline bridge interface vpls vplsA**

- There is no **no** version.

Clearing the VPLS Forwarding Table

You can use the following **clear** commands to remove all dynamic (learned) MAC address entries or a specific MAC address entry from the forwarding table for a VPLS instance.

clear bridge

- Use to remove from the forwarding table all dynamic MAC address entries for the specified VPLS instance.
- Example

  host1#**clear bridge vplsB**

- There is no **no** version.

clear bridge address

- Use to remove from the forwarding table a specific dynamic MAC address entry for the specified VPLS instance.
- Example

  host1#**clear bridge vplsB address 0090.1a40.9992**

- There is no **no** version.

clear bridge interface

- Use to remove from the forwarding table all dynamic MAC address entries for a network interface associated with a VPLS instance.
- You must specify the following:
  - *interfaceType*—One of the following interface types listed in *Interface Types and Specifiers* in *JUNOSe Command Reference Guide, About This Guide*:
    - **atm**
    - **fastEthernet**
    - **gigabitEthernet**
    - **tenGigabitEthernet**
  - *interfaceSpecifier*—Particular interface; format varies according to interface type; see *Interface Types and Specifiers* in *JUNOSe Command Reference Guide, About This Guide* for information
- Example

  host1#**clear bridge interface atm 3/3.2**

- There is no **no** version.

clear bridge interface vpls

- Use to remove from the forwarding table all dynamic MAC address entries for the VPLS virtual core interface associated with a VPLS instance.
- Example

  host1#**clear bridge interface vpls vplsA**

- There is no **no** version.

Clearing BGP Attributes

You can use the **clear ip bgp** commands listed in Table 55 to clear BGP attributes for the L2VPN address family and, in one case, for the VPLS address family associated with a specific VPLS instance.

For more information about using these commands, see *Chapter 1, Configuring BGP Routing*. For information about the syntax of each command, see the *JUNOSe Command Reference Guide A to M*.

Table 55: Commands for Clearing BGP Attributes

Attribute	Command	Examples
Clears reachability information for the L2VPN address family	**clear ip bgp**	host1#**clear ip bgp l2vpn soft in**
Clears route flap dampening information for the L2VPN address family, or for the VPLS address family associated with the specified VPLS instance	**clear ip bgp dampening**	host1#**clear ip bgp l2vpn dampening** host1#**clear ip bgp vpls vplsA dampening**
Clears the wait to receive an End-of-RIB marker from the peer for the L2VPN address family	**clear ip bgp wait-end-of-rib**	host1#**clear ip bgp l2vpn wait-end-of-rib**

Monitoring Bridging-Related Settings for VPLS

You can use the **show** commands listed in Table 56 to display VPLS settings related to transparent bridging. Except for the **show bridge interface vpls** command, which is only for VPLS instances, you can use these commands to monitor both VPLS instances and standard bridge groups for transparent bridging. The examples in this section show those portions of the command display relevant to a VPLS configuration.

For more information about using these commands, see *Monitoring Transparent Bridging* in *JUNOSe Link Layer Configuration Guide, Chapter 10, Configuring Transparent Bridging*. For information about the syntax of each command, see the *JUNOSe Command Reference Guide N to Z*.

Table 56: Commands for Monitoring VPLS Bridging Settings

show bridge	**show bridge port**
show bridge groups	**show bridge table**
show bridge interface	**show subscriber-policy**
show bridge interface vpls	

show bridge

- Use to display configuration and statistics information for the specified VPLS instance.
- To display address table and statistics information for all network interfaces associated with the VPLS instance, use the **all** keyword.
- Field descriptions
  - BridgeGroup—Name of the VPLS instance for which information is displayed
  - Bridge Mode—Bridging capability currently enabled; for a VPLS instance, this field always displays default
  - Aging Time—Length of time, in seconds, that a MAC address entry can remain in the forwarding table before expiring
  - Learning—Whether acquisition of dynamically learned MAC addresses is enabled or disabled
  - Max Learn—Maximum number of dynamic MAC addresses that the VPLS instance can learn
  - Link Status Snmp Traps—Whether SNMP link status processing is enabled or disabled
  - Subscriber Policy—Name of the subscriber policy currently in effect
  - Port Count—Number of ports currently configured for the VPLS instance, including network interfaces and the VPLS virtual core interface
  - Interface Count—Number of network interfaces currently configured for the VPLS instance
  - Transport Virtual Rtr—Name of the transport virtual router configured for the VPLS instance
  - Route Distinguisher—Unique route distinguisher configured for the VPLS instance
  - SiteName—Site name configured for the VPLS instance
  - SiteId—Numerical site identifier configured for the VPLS instance
  - SiteRange—Maximum number of sites that can participate in the VPLS domain associated with the VPLS instance
  - VPLS Route Targets—Extended community identifiers, also known as route targets, for each VPLS instance configured on the router
  - Flood Next Hop—Index of the MPLS next hop to which the router floods packets with unknown destination addresses. For more information about displaying MPLS next hops and any available next-hop statistics, see **show mpls next-hop** on page 343.

- Example

```
host1#show bridge vplsA

  BridgeGroup: vplsA(vpls)

        Bridge Mode:            default
        Aging Time:             300 secs
        Learning:               Enabled
        Max Learn:              Unlimited
        Link Status Snmp Traps: Disabled
        Subscriber Policy:      default Subscriber
        Port Count:             2
        Interface Count:        1
        Transport Virtual Rtr:  default
        Route Distinguisher:    1.1.1.1:10
        SiteName:               boston
        SiteId:                 1
        SiteRange:              10
        VPLS Route Targets
          Route Target: RT:100:1 (both)
          Route Target: RT:100:2 (both)
        Flood Next Hop: Index 1048577
```

show bridge groups

- Use to display configuration and statistics information for all VPLS instances configured on the router.
- To display only the names of the VPLS instances configured on the router, use the command with no keywords.
- To display configuration settings for all VPLS instances on the router, use the **details** keyword.
- The **show bridge groups details** command displays the same fields for each VPLS instance as the **show bridge** command. For information, see the field descriptions for **show bridge** on page 571.
- Example 1

```
host1#show bridge groups

  BridgeGroup: vplsA(vpls)

  BridgeGroup: vplsB(vpls)
```

- Example 2

```
host1#show bridge groups details

  BridgeGroup: vplsA(vpls)

        Bridge Mode:            default
        Aging Time:             300 secs
        Learning:               Enabled
        Max Learn:              Unlimited
        Link Status Snmp Traps: Disabled
        Subscriber Policy:      default Subscriber
        Port Count:             2
        Interface Count:        1
        Transport Virtual Rtr:  default
        Route Distinguisher:    1.1.1.1:10
        SiteName:               boston
```

```
      SiteId:                 1
      SiteRange:              10
      VPLS Route Targets
        Route Target: RT:100:1 (both)
        Route Target: RT:100:2 (both)
      Flood Next Hop: Index 1048577

BridgeGroup: vplsB(vpls)

      Bridge Mode:            default
      Aging Time:             300 secs
      Learning:               Enabled
      Max Learn:              Unlimited
      Link Status Snmp Traps: Disabled
      Subscriber Policy:      default Subscriber
      Port Count:             2
      Interface Count:        1
      Transport Virtual Rtr:  default
      Route Distinguisher:    1.1.1.1:11
      SiteName:               boston
      SiteId:                 1
      SiteRange:              20
      VPLS Route Targets
        No Route Targets configured
      Flood Next Hop: Index 1048578
```

show bridge interface

- Use to display configuration, statistics, and status information for a specified network interface or for all interfaces assigned to a VPLS instance.
- To display information about all interfaces that belong to the VPLS instance, including the VPLS virtual core interface, specify the command with the name of the VPLS instance.
- To display only the interface type and specifier, associated port number, and operational status for each interface in a VPLS instance, specify the command with the name of the VPLS instance and the **brief** keyword.
- Field descriptions
  - BridgeGroup—Name of the VPLS instance to which the interface belongs
  - Port Number—Port number on which this interface resides
  - Operational Status—Operational status of the physical interface: Up, Down, LowerLayerDown, NotPresent
  - Admin Status—State of the physical interface: Up, Down
  - Snmp Link Status Trap—Whether SNMP link status processing is enabled or disabled for the specified interface
  - Max Learn—Maximum number of dynamic MAC addresses that the interface can learn
  - Subscriber Policy—Name of the subscriber policy currently in effect for the interface

- Statistics—Displays statistics information for the specified port
  - In Octets—Number of octets received on this interface
  - In Frames—Number of frames received on this interface
  - In Discards—Number of incoming packets discarded on this interface
  - In Errors—Number of incoming errors received on this interface
  - Out Octets—Number of octets transmitted on this interface
  - Out Frames—Number of frames transmitted on this interface
  - Out Discards—Number of outgoing packets discarded on this interface
  - Out Errors—Number of outgoing errors on this interface
- Time since counters last reset—Elapsed time since statistics counters were last reset
- queue—Hardware packet queue associated with the specified traffic class and interface
  - Queue length—Length of the queue, in bytes
  - Forwarded packets, bytes—Number of packets and bytes forwarded on this queue
  - Dropped committed packets, bytes—Number of committed packets and bytes that were dropped
  - Dropped conformed packets, bytes—Number of conformed packets and bytes that were dropped
  - Dropped exceeded packets, bytes—Number of exceeded packets and bytes that were dropped
- vpls *vplsName*—Identifies the VPLS virtual core interface for the VPLS instance
- Using the **brief** keyword displays only the following fields:
  - Interface—Interface type and specifier associated with the port
  - Port—Port number on which this interface resides
  - Status—Operational status of the physical interface: Up, Down, LowerLayerDown, NotPresent

- Example 1—Displays information about a specified network interface in a VPLS instance

```
host1#show bridge interface atm 3/1.10

    atm3/1.10
      BridgeGroup: vplsB
      Port Number: 1
      Operational Status: Up
      Admin Status: Up
      Snmp Link Status Trap: Disabled
      Max Learn: Unlimited
      Subscriber Policy: default Trunk
```

```
  Statistics:
     In Octets:    1958
     In Frames:    14
     In Discards:  1
     In Errors:    0
     Out Octets:   1930
     Out Frames:   14
     Out Discards: 1
     Out Errors:   0
  Time since counters last reset: 00:14:32

  queue 0: traffic class best-effort, bound to bridge ATM3/1.10
    Queue length 0 bytes
    Forwarded packets 14, bytes 2238
    Dropped committed packets 0, bytes 0
    Dropped conformed packets 0, bytes 0
    Dropped exceeded packets 0, bytes 0
```

- Example 2—Displays information about all interfaces in a VPLS instance, including the VPLS virtual core interface (vplsB)

```
host1#show bridge vplsB interface

  FastEthernet1/1.1
    Port Number: 1
    Operational Status: Up
    Admin Status: Up
    Snmp Link Status Trap: Disabled
    Max Learn: Unlimited
    Subscriber Policy: samplepolicy
    Statistics:
       In Octets:    3770
       In Frames:    27
       In Discards:  0
       In Errors:    0
       Out Octets:   3682
       Out Frames:   27
       Out Discards: 0
       Out Errors:   0
    Time since counters last reset: 01:07:08

    queue 0: traffic class best-effort, bound to bridge FastEthernet1/1.1
      Queue length 0 bytes
      Forwarded packets 27, bytes 3898
      Dropped committed packets 0, bytes 0
      Dropped conformed packets 0, bytes 0
      Dropped exceeded packets 0, bytes 0

  vpls vplsB
    Port Number: 2
    Operational Status: Down
    Admin Status: Up
    Snmp Link Status Trap: Disabled
    Max Learn: Unlimited
    Subscriber Policy: default Trunk
```

```
  Statistics:
     In Octets:    0
     In Frames:    0
     In Discards:  0
     In Errors:    0
     Out Octets:   0
     Out Frames:   0
     Out Discards: 40
     Out Errors:   0
  Time since counters last reset: 01:04:10
```

- Example 3—Displays a summary of all interfaces configured for the specified VPLS instance

```
host1#show bridge vplsB interface brief
        Interface            Port       Status
-------------------------   ------   ------------
FastEthernet1/1.1           1        Up
ATM10/1.1.1                 2        Up
vpls vplsB                  3        Up
```

show bridge interface vpls

- Use to display configuration, statistics, and status information for the VPLS virtual core interface associated with a VPLS instance.
- The **show bridge interface vpls** command displays the same fields for the VPLS virtual core interface as the **show bridge interface** command. For information, see the field descriptions for **show bridge interface** on page 573.
- Example

```
host1#show bridge interface vpls vplsB

  vpls vplsB
    BridgeGroup: vplsB
    Port Number: 2
    Operational Status: Up
    Admin Status: Up
    Snmp Link Status Trap: Disabled
    Max Learn: Unlimited
    Subscriber Policy: default Trunk
    Statistics:
       In Octets:    0
       In Frames:    0
       In Discards:  0
       In Errors:    0
       Out Octets:   0
       Out Frames:   0
       Out Discards: 0
       Out Errors:   0
    Time since counters last reset: 00:12:53
```

show bridge port

- Use to display configuration, statistics, and status information for ports (interfaces) associated with a VPLS instance.
- The **show bridge port** command displays the same fields for each port as the **show bridge interface** command. For information, see the field descriptions for **show bridge interface** on page 573.
- Example 1

```
host1#show bridge vplsC port

    FastEthernet1/1
      Port Number: 1
      Operational Status: Up
      Admin Status: Up
      Snmp Link Status Trap: Disabled
      Max Learn: Unlimited
      Subscriber Policy: samplepolicy
      Statistics:
         In Octets:    2018
         In Frames:    15
         In Discards:  0
         In Errors:    0
         Out Octets:   1930
         Out Frames:   14
         Out Discards: 0
         Out Errors:   0
      Time since counters last reset: 00:10:55

      queue 0: traffic class best-effort, bound to bridge FastEthernet1/1.1
        Queue length 0 bytes
        Forwarded packets 14, bytes 2042
        Dropped committed packets 0, bytes 0
        Dropped conformed packets 0, bytes 0
        Dropped exceeded packets 0, bytes 0

    vpls vplsC
      Port Number: 2
      Operational Status: Up
      Admin Status: Up
      Snmp Link Status Trap: Disabled
      Max Learn: Unlimited
      Subscriber Policy: default Trunk
      Statistics:
         In Octets:    0
         In Frames:    0
         In Discards:  0
         In Errors:    0
         Out Octets:   0
         Out Frames:   0
         Out Discards: 0
         Out Errors:   0
      Time since counters last reset: 00:07:07
```

- Example 2

```
host1#show bridge vplsTest port brief
 Port            Interface              Status
------   ------------------------   ------------
1        FastEthernet1/1.1          Up
2        ATM10/1.1.1                Up
3        vpls vplsTest              Up
```

show bridge table

- Use to display information about the MAC address entries in the forwarding table for the specified VPLS instance.
- To display information only for static (nonlearned) entries, use the **static** keyword.
- To display information only for dynamic (learned) entries, use the **dynamic** keyword.
- To display information for both static and dynamic entries, use the command with no keywords.
- Field descriptions
  - Bridge—Name of the VPLS instance for which the MAC address table is displayed
  - Address—MAC address of the entry
  - Action—Specifies how the VPLS instance handles this entry: forward or discard
  - Interface—Interface type and specifier on which the entry is forwarded; this value does not appear for entries that are discarded; vpls identifies the VPLS virtual core interface
  - Age—Length of time that a dynamic entry has been in the forwarding table; this value does not appear for static entries
- Example

```
host1#show bridge vpls1 table
  Bridge: vpls1 MAC Address Table

   Address            Action                Interface              Age
--------------  --------------------  -------------------------  -------
0009.01a0.002e  forward               ATM10/1.1.1                0
0090.1a41.3aca  forward               vpls (10)                  0
```

show subscriber-policy

- Use to display the set of forwarding and filtering rules for all subscriber policies configured on the router, or for a specified subscriber policy.
- To display information about all default and nondefault subscriber policies configured on the router, issue the command without specifying a subscriber policy name.
- You can modify the default Subscriber policy for a network interface configured as a subscriber (client) interface. You cannot change the default Trunk policy for a network interface configured as a trunk (server) interface, or for the VPLS virtual core interface, which always uses the default Trunk policy.

- Field descriptions
  - Subscriber—Name of the subscriber policy
  - Permit—Indicates that the subscriber interface forwards packets of the specified type. For the relearn attribute, specifies that relearning a MAC address entry on a different interface from the one initially associated with this entry in the forwarding table is allowed on this interface
  - Deny—Indicates that the subscriber interface filters packets of the specified type. For the relearn attribute, specifies that relearning is prohibited on this interface
- Example 1—Displays the rules for all default and nondefault subscriber policies configured on the router

```
host1#show subscriber-policy

Subscriber: default Subscriber
  ARP                 : Permit
  Broadcast           : Deny
  Multicast           : Permit
  Unknown Destination : Deny
  IP                  : Permit
  Unknown Protocol    : Permit
  Unicast             : Permit
  PPPoE               : Permit
  Relearn             : Permit
  Mpls                : Permit

Subscriber: default Trunk
  ARP                 : Permit
  Broadcast           : Permit
  Multicast           : Permit
  Unknown Destination : Permit
  IP                  : Permit
  Unknown Protocol    : Permit
  Unicast             : Permit
  PPPoE               : Permit
  Relearn             : Permit
  Mpls                : Permit

  Subscriber: client01
  ARP                 : Permit
  Broadcast           : Permit
  Multicast           : Deny
  Unknown Destination : Deny
  IP                  : Permit
  Unknown Protocol    : Permit
  Unicast             : Permit
  PPPoE               : Permit
  Relearn             : Deny
  Mpls                : Permit
```

- Example 2—Displays the rules for a specified subscriber policy

```
host1#show subscriber-policy client01
  Subscriber: client01
  ARP                 : Permit
  Broadcast           : Permit
  Multicast           : Deny
  Unknown Destination : Deny
  IP                  : Permit
  Unknown Protocol    : Permit
  Unicast             : Permit
  PPPoE               : Permit
  Relearn             : Deny
  Mpls                : Permit
```

Monitoring BGP-Related Settings for VPLS

You can use the **show ip bgp** commands listed in Table 57 to display BGP-related settings for VPLS instances in the L2VPN address family or in the VPLS address family.

For more information about using these commands, see *Chapter 1, Configuring BGP Routing* and *Chapter 3, Configuring BGP-MPLS Applications*. For information about the syntax of each command, see the *JUNOSe Command Reference Guide N to Z*.

Table 57: Commands for Monitoring VPLS BGP Settings

show ip bgp	**show ip bgp neighbors paths**
show ip bgp advertised-routes	**show ip bgp neighbors received-routes**
show ip bgp community	**show ip bgp neighbors routes**
show ip bgp community-list	**show ip bgp network**
show ip bgp dampened-paths	**show ip bgp next-hops**
show ip bgp filter-list	**show ip bgp paths**
show ip bgp flap-statistics	**show ip bgp peer-group**
show ip bgp l2vpn	**show ip bgp quote-regexp**
show ip bgp l2vpn vpls	**show ip bgp regexp**
show ip bgp neighbors	**show ip bgp summary**
show ip bgp neighbors dampened-routes	–

This section provides examples of some of the **show ip bgp** commands that you can use to monitor VPLS configurations.

show ip bgp l2vpn
show ip bgp l2vpn vpls

- Use to display layer 2 NLRI for all VPLS instances in the L2VPN address family, for a particular VPLS instance in the L2VPN address family, or for a particular VPLS instance in the VPLS address family.
- To display layer 2 NLRI for the route that matches a specified prefix (site ID and block offset) in the L2VPN address family or in the VPLS address family, use the **site-id** and **block-offset** keywords.

- Field descriptions
  - Local BGP identifier—IP address of the local VE router
  - local AS—Autonomous system number
  - Local-RIB version—Version number of the local routing information base
  - FIB version—Version number of the forwarding information base
  - Status codes—Status codes for the route
  - Prefix—Route prefix in the format *siteID*:*blockOffset*
  - Peer—IP address of the peer from which the route was learned
  - Next-hop (or Next hop IP address)—IP address of the next router that is used when a packet is forwarded to the destination network
  - MED—Multiexit discriminator for the route
  - LocPrf—Local preference for the route
  - Weight—Weight of the route
  - Origin—Origin of the route
  - AS path—AS path through which this route has been advertised
  - Extended communities—Route targets of the communities associated with this route
- Example 1—Displays information for all VPLS instances in the L2VPN address family

```
host1#show ip bgp l2vpn all

Local BGP identifier 1.1.1.1, local AS 100
  4 routes (264 bytes)
  4 destinations (288 bytes) of which 4 have a route
  0 routes selected for route tables installation
  0 unicast/multicast routes selected for route table installation
  0 unicast/multicast tunnel-usable routes selected for route table installation
  0 tunnel-only routes selected for tunnel-route table installation
  4 path attribute entries (608 bytes)
  Local-RIB version 11. FIB version 11.

Status codes: > best, * invalid, s suppressed, d dampened, r rejected,
              a auto-summarized

Prefix                      Peer           Next-hop       MED LocPrf Weight Origin
> 1:1                       0.0.0.0        self                        0 IGP
> 1:1                       0.0.0.0        self                        0 IGP
> 2:1                       2.2.2.2        2.2.2.2              100      0 IGP
> 2:1                       2.2.2.2        2.2.2.2              100      0 IGP
```

- Example 2—Displays summary information for the L2VPN address family

```
host1#show ip bgp l2vpn all summary

Display summary information for the l2vpn address-family
Local router ID 1.1.1.1, local AS 100
  Administrative state is Start
  BGP Operational state is Up
  Shutdown in overload state is disabled
  Default local preference is 100
  IGP synchronization is enabled
```

```
  Default originate is disabled
  Always compare MED is disabled
  Compare MED within confederation is disabled
  Advertise inactive routes is disabled
  Advertise best external route to internal peers is disabled
  Enforce first AS is disabled
  Missing MED as worst is disabled
  Route flap dampening is disabled
  Log neighbor changes is disabled
  Fast External Fallover is disabled
  No maximum received AS-path length
  BGP administrative distances are 20 (ext), 200 (int), and 200 (local)
  Client-to-client reflection is enabled
  Cluster ID is 1.1.1.1
  Route-target filter is enabled
  Default IPv4-unicast is enabled
  Check next-hops of vpn routes is disabled
  Redistribution of iBGP routes is disabled
  Graceful restart is globally disabled
  Global graceful-restart restart time is 120 seconds
  Global graceful-restart stale paths time is 360 seconds
  Graceful-restart path selection defer time is 360 seconds
  Graceful-restart is not ready to switch to the standby SRP
  The last restart was not graceful
  Local-RIB version 11. FIB version 11.

                                                 Messages  Messages  Prefixes
Neighbor           AS State       Up/down time      Sent  Received  Received
2.2.2.2           100 Established     00:30:35        65        65         2
```

- Example 3—Displays information for the route that matches the specified prefix (2:1) for a VPLS instance named customer1 in the VPLS address family

```
host1#show ip bgp l2vpn vpls customer1 site-id 2 block-offset 1

BGP route information for prefix 2:1
  Received route learned from internal peer 2.2.2.2 (best route)
    Leaked route
    Route placed in IP forwarding table
    Best to advertise to external peers
    Address Family Identifier (AFI) is layer2
    Subsequent Address Family Identifier (SAFI) is unicast
    Route Distinguisher (RD) is 100:11
    Original Route Distinguisher (RD) is 100:21
    MPLS in-label is none
    MPLS in-label block size is 0
    MPLS out-label is 46
    MPLS out-label block size is 20
    Next hop IP address is 2.2.2.2 (metric 3)
    Multi-exit discriminator is not present
    Local preference is 100
    Weight is 0
    Origin is IGP
    AS path is empty
    Extended communities RT:100:1 Layer 2:19:00:0
```

show ip bgp next-hops

- Use to display information about BGP next hops in the L2VPN address family or in the VPLS address family.
- Field descriptions
  - Indirect next-hop—BGP next-hop attribute received in the BGP update message
  - Resolution—Describes where the indirect next hop is resolved (the IP routing table, the IP tunnel routing table, or both) and whether this is in a VR or VRF
  - IP indirect next-hop index—Index number of the IP indirect next hop that corresponds to the BGP indirect next hop and its resolution
  - MPLS indirect next-hop index—Index number of the MPLS indirect next hop that corresponds to the BGP indirect next hop and its resolution
  - Reachable—Indicates whether or not the indirect next hop is reachable. For more information about the reachability rules that apply for various route types, see the command description for **show ip bgp next-hops** on page 469.
  - metric—Metric number of the BGP indirect next hop
  - Number of direct next-hops—Number of the equal-cost legs of direct next hops to which this indirect next hop resolves
  - Direct next-hop—MPLS next-hop index that resolves the MPLS indirect next hop
  - Reference count—Number of label mappings of BGP routes that use this next hop
- Example—Displays next hop information that matches the specified indirect next-hop address (2.2.2.2) in the L2VPN address family

```
host1#show ip bgp l2vpn all next-hops 2.2.2.2
Indirect next-hop 2.2.2.2
  Resolution in IP route table of VR
       IP indirect next-hop index 2
       Reachable (metric 3)
       Number of direct next-hops is 1
              Direct next-hop ATM2/0.10 (10.10.10.2)
  Resolution in IP tunnel-route table of VR
       MPLS indirect next-hop index 19
       Reachable (metric 3)
       Number of direct next-hops is 1
              Direct next-hop 0000000c
  Reference count is 2
```

Monitoring LDP-Related Settings for VPLS

You can use the **show ldp vpls** command to display MPLS configuration information for a VPLS instance that uses LDP as the signaling protocol.

show ldp vpls

- Use to display MPLS configuration information for a VPLS instance that uses LDP as the signaling protocol.
- To display information for a specific VPLS instance, specify the name of the designated VPLS instance.
- To display information for a specific neighbor adress, use the command with the **neighbor** keyword.
- To display information for all VPLS instances configured on the virtual router, use the command with no keywords.
- The **mpls** keyword is optional and is provided for compatibility with non–E-series implementations.
- Field descriptions
  - Vpls Instance—Name of the VPLS instance for which the configuration information is displayed
  - Vpls Id—Globally unique identifier for the VPLS domain
  - Remote PE—IP address of the remote VE (also known as the PE) router
  - In-label—Incoming MPLS label from the remote site
  - Out-label—Outgoing MPLS label used to reach the remote site
- Example

```
host1:ve1#show ldp vpls
  Vpls     Vpls   Remote
Instance    Id      PE      In-label   Out-label
--------   ----   -------   --------   ---------
vpls1         1   2.2.2.2         25          27
vpls2         2   2.2.2.2         26          28
```

Monitoring MPLS-Related Settings for VPLS

You can use the **show mpls forwarding** command to display MPLS-related settings for VPLS instances.

show mpls forwarding

- Use to display information for MPLS labels being used for forwarding.
- Field descriptions
  - In label—Label sent to upstream neighbor for route
  - Out label—Label received from downstream neighbor for route
  - Label space—Label space in which the label is assigned
  - Owner—Signaling protocol that placed the label in the forwarding table: BGP, LDP, or RSVP-TE

  - Spoof check—Type and location of spoof checking performed on the MPLS packet, router, or interface
  - Action—Action taken for MPLS packets arriving with that label
  - in pkts—Number of packets sent with the label
  - in Octets—Number of octets sent with the label
  - in errors—Number of packets that are dropped for some reason before being sent
  - in discardPkts—Number of packets that are discarded due to lack of buffer space before being sent

- Example 1

```
host1:ve1#show mpls forwarding brief

In-label   Owner    Action
--------   -----    ----------------------------------
17         bgp      Forward to bridge-group customer1
27         bgp      Forward to bridge-group customer2
```

- Example 2

```
host1:ve1#show mpls forwarding label 17
In label: 17
  Label space: platform label space
  Owner: bgp
  Spoof check: router pe1
  Action:
    MPLS next-hop: 3, Forward to bridge-group customer1
  Statistics:
    0 in pkts
    0 in Octets
    0 in errors
    0 in discard pkts
```

Monitoring VPLS-Specific Settings

You can use the **show vpls connections** command to display configuration and status information for VPLS connections configured on the router.

show vpls connections

- Use to display connection information for a specified VPLS instance configured on the router, or for all VPLS instances configured on the router.
- To display detailed configuration and status information for VPLS connections, use the **details** keyword.
- To display information only for operational (up) VPLS connections, use the **state up** keywords.
- To display information only for nonoperational (down) VPLS connections, use the **state down** keywords.
- To display information only for a specified VPLS instance, use the **bridge-group** keyword and specify the name of the VPLS instance.

- To display information only for the VPLS connection with the specified remote site identifier, use the **remote-site** keyword and specify the site identifier value. (For more information about configuring the site identifier for a VPLS instance, see **bridge vpls site-name site-id** on page 542.)
- Field descriptions
  - BridgeGroup—Name of the VPLS instance for which information is displayed
  - Bridge Mode—Bridging capability currently enabled; for a VPLS instance, this field always displays default
  - Aging Time—Length of time, in seconds, that a MAC address entry can remain in the forwarding table before expiring
  - Learning—Whether acquisition of dynamically learned MAC addresses is enabled or disabled
  - Max Learn—Maximum number of dynamic MAC addresses that the VPLS instance can learn
  - Link Status Snmp Traps—Whether SNMP link status processing is enabled or disabled
  - Subscriber Policy—Name of the subscriber policy currently in effect
  - Port Count—Number of ports currently configured for the VPLS instance, including network interfaces and the VPLS virtual core interface
  - Interface Count—Number of network interfaces currently configured for the VPLS instance
  - Transport Virtual Rtr—Name of the transport virtual router configured for the VPLS instance
  - Route Distinguisher—Unique route distinguisher configured for the VPLS instance
  - SiteName—Site name configured for the VPLS instance
  - SiteId—Numerical site identifier configured for the VPLS instance
  - SiteRange—Maximum number of sites that can participate in the VPLS domain associated with the VPLS instance
  - VPLS Route Targets—Extended community identifiers, also known as route targets, for each VPLS instance configured on the router
  - Flood Next Hop—Index number of the MPLS next hop to which the router floods packets with unknown destination addresses. For more information about displaying MPLS next hops and any available next-hop statistics, see the **show mpls next-hop** command description in *Chapter 2, Configuring MPLS*.
  - Interface—Type and specifier of the network interfaces and VPLS virtual core interface associated with the VPLS instance; vpls *vplsName* in this field identifies the VPLS virtual core interface
  - Port—Port number of the module on which the network interface or VPLS virtual core interface resides
  - Status—Operational status of the physical interface: Up, Down, LowerLayerDown, NotPresent

- Connections status code—Possible status codes for the VPLS connection that appear in the State field
- Site—Remote site identifier
- State—Status of the connection with the remote VPLS instance; possible values for this field appear in the Connections status code legend in the command display
- Remote PE—IP address of the remote VPLS edge (VE) router, which is analogous to the remote provider edge (PE) router in a BGP/MPLS VPN configuration
- In-label—Incoming MPLS label from the remote site
- Out-label—Outgoing MPLS label used to reach the remote site
- MPLS NH Idx—MPLS next-hop index number that corresponds to the outgoing MPLS label
- Up-down Time—Time since the last state change for this VPLS connection

- Example

```
host1#show vpls connections details

  BridgeGroup: vpls1(vpls)

        Bridge Mode:            default
        Aging Time:             300 secs
        Learning:               Enabled
        Max Learn:              Unlimited
        Link Status Snmp Traps: Disabled
        Subscriber Policy:      default Subscriber
        Port Count:             2
        Interface Count:        1
        Transport Virtual Rtr:  pe1
          Route Distinguisher:    1.1.1.1:10
        SiteName:               westford
        SiteId:                 1
        SiteRange:              10
        VPLS Route Targets
          Route Target: RT:100:1 (both)
        Flood Next Hop: Index 1048577
          MPLS next-hop: 20, label 46, resolved by MPLS next-hop 19
            MPLS next-hop: 19, resolved by MPLS next-hop 17, peer 2.2.2.2
              MPLS next-hop: 17, label 82 on ATM2/0.10, nbr 10.10.10.2

        Interface            Port       Status
------------------------   ------   ------------
FastEthernet3/1            1        Up
vpls vpls1                 2        Up

Connections status code:
UP = Operational
SC = Local and Remote Site Identifier Collision
EM = Encapsulation Mismatch
OR = Out of Range
DN = VC Down because Remote PE Unreachable
LD = Local Site Down
RD = Remote Site Down
AS = Max BGP AS path length exceeded
OL = No Out Label
```

```
Site   State     Remote PE     In-label  Out-label  MPLS NH Idx  Up-down Time
-----  -----  ---------------  --------  ---------  -----------  ------------
2      UP     2.2.2.2          17        46         20           00:02:56

  BridgeGroup: vpls2(vpls)

        Bridge Mode:            default
        Aging Time:             300 secs
        Learning:               Enabled
        Max Learn:              Unlimited
        Link Status Snmp Traps: Disabled
        Subscriber Policy:      default Subscriber
        Port Count:             2
        Interface Count:        1
        Transport Virtual Rtr:  pe1
          Route Distinguisher:    1.1.1.1:10
        SiteName:               westford
        SiteId:                 1
        SiteRange:              20
        VPLS Route Targets
          Route Target: RT:100:2 (both)
        Flood Next Hop: Index 1048578
          MPLS next-hop: 21, label 56, resolved by MPLS next-hop 19
            MPLS next-hop: 19, resolved by MPLS next-hop 17, peer 2.2.2.2
              MPLS next-hop: 17, label 82 on ATM2/0.10, nbr 10.10.10.2

        Interface            Port       Status
-------------------------   ------   ------------
ATM2/0.12                   1        Up
vpls vpls2                  2        Up

Connections status code:
UP = Operational
SC = Local and Remote Site Identifier Collision
EM = Encapsulation Mismatch
OR = Out of Range
DN = VC Down because Remote PE Unreachable
LD = Local Site Down
RD = Remote Site Down
AS = Max BGP AS path length exceeded
OL = No Out Label

Site   State     Remote PE     In-label  Out-label  MPLS NH Idx  Up-down Time
-----  -----  ---------------  --------  ---------  -----------  ------------
2      UP     2.2.2.2          27        56         21           00:02:56
```