Chapter 8
Configuring VPLS

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

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.

Related Topics

For more information about configuring BGP, MPLS, LDP, BGP/MPLS VPNs, and layer 2 services over MPLS, see the following chapters:

- *Chapter 1, Configuring BGP Routing*
- *Chapter 2, Configuring MPLS*
- *Chapter 3, Configuring BGP-MPLS Applications*
- *Chapter 4, Layer 2 Services over MPLS Overview*
- *Chapter 5, Configuring Layer 2 Services over MPLS*
- *Chapter 6, Monitoring Layer 2 Services over MPLS*

For more information about configuring transparent bridging, see the following chapter:

- JUNOSe Link Layer Configuration Guide, Chapter 13, 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 12, 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.

Related Topics

- Configuring VPLS Instances with BGP Signaling on page 543
- Configuring Optional Attributes for VPLS Instances on page 545
- Configuring VPLS Network Interfaces on page 547
- Configuring Subscriber Policies for VPLS Network Interfaces on page 548
- Configuring the Loopback Interface and Router ID for VPLS on page 550
- Configuring MPLS LSPs for VPLS on page 551
- Configuring BGP Signaling for VPLS on page 552
- See *VPLS Configuration Example with BGP Signaling* on page 554 for a detailed sample configuration.

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.

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.

   host1(config)#**bridge customer1 vpls transport-virtual-router vr1**

If the bridge group you specify (customer1 in this example) 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.)

   host1(config)#**bridge customer1 vpls site-range 15**

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

   host1(config)#**bridge customer1 vpls site-name westford site-id 1**

   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.

   host1(config)#**bridge customer1 vpls rd 100:11**

   In this example, the first number in the route distinguisher (100) is the number of the AS in which the extended community resides. The second number in the route distinguisher (11) uniquely identifies the VPLS instance within the AS.

TIP: You cannot change or remove the route distinguisher for a VPLS instance after you set it; for this reason the **no bridge vpls rd** command fails. . 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.

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

   host1(config)#**bridge customer1 vpls route-target both 100:1**

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

BEST PRACTICE: We recommend that you add the route target to both the VPLS instance's import list and export list of VPN extended communities. To do so, use the **both** keyword.

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
```

Related Topics

- Configuration Tasks for VPLS with BGP Signaling on page 542
- Configuration Tasks for VPLS with LDP Signaling on page 558
- **bridge vpls rd** command
- **bridge vpls route-target** command
- **bridge vpls site-name site-id** command
- **bridge vpls site-range** command
- **bridge vpls transport-virtual-router** command

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. The requirements and procedures for optional attributes are the same whether you employ BGP or LDP signaling for VPLS.

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.

1. (Optional) Configure a VPLS instance to acquire dynamically learned MAC addresses.

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

2. (Optional) 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.

   In this example, the VPLS instance 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**

   In this example, the VPLS instance drops frames sent from or destined for the node with MAC address 1011.22b2.333c

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

3. (Optional) Set the length of time that a dynamic (learned) MAC address entry can remain in the forwarding table of the specified VPLS instance before expiring.

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

4. (Optional) Set the maximum number of dynamic MAC address entries that the specified VPLS instance can learn.

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

5. (Optional) Enable SNMP link status processing for all network interfaces associated with the specified VPLS instance.

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

Related Topics

- Configuration Tasks for VPLS with BGP Signaling on page 542
- **bridge acquire** command
- **bridge address** command
- **bridge aging-time** command
- **bridge learn** command
- **bridge snmp-trap link-status** command
- For more information about using these commands, see *Configuring Optional Bridge Group Attributes* in *JUNOSe Link Layer Configuration Guide, Chapter 13, Configuring Transparent Bridging*.

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

The requirements and procedures for network interfaces are the same whether you employ BGP or LDP signaling for VPLS.

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.

   host1(config)#**interface gigabitEthernet 3/0**

2. From Interface Configuration mode or Subinterface Configuration mode, assign the interface to the specified VPLS instance.

   host1(config-if)#**bridge-group customer1**

   Issuing this command with no optional keywords configures the network interface as a subscriber (client) interface by default.

3. (Optional) 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 548.

   host1(config-if)#**bridge-group customer1 subscriber-trunk**

4. (Optional) Set the maximum number of MAC addresses that the network interface can learn.

   host1(config-if)#**bridge-group customer1 learn 100**

5. Enable SNMP link status processing only for the specified network interface in the VPLS instance.

   host1(config-if)#**bridge-group customer1 snmp-trap link-status**

Related Topics

- Configuration Tasks for VPLS with BGP Signaling on page 542
- Configuration Tasks for VPLS with LDP Signaling on page 558
- **bridge-group** command

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. The requirements and procedures for subscriber policies are the same whether you employ BGP or LDP signaling for VPLS.

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 49, 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 49 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 49: 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 49.

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 50 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.

Table 50: 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 49 on page 549, 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.

Related Topics

- Configuration Tasks for VPLS with BGP Signaling on page 542
- Configuration Tasks for VPLS with LDP Signaling on page 558
- For information about using these commands, see *Configuring Subscriber Policies* in *JUNOSe Link Layer Configuration Guide, Chapter 13, Configuring Transparent Bridging*.

Configuring the Loopback Interface and Router ID for VPLS

A loopback interface provides a stable address for BGP or LDP to use so that they can avoid any impact if a physical interface goes down. The loopback interface sends packets back to the router or access server for local processing. Any packets routed from the loopback interface, but not destined to the loopback interface, are dropped.

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.

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 for VPLS* on page 559.

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.

   ```
   host1(config)#interface loopback 0
   host1(config-if)#ip address 10.3.3.3 255.255.255.255
   host1(config-if)#exit
   ```

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

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

Related Topics

- Configuration Tasks for VPLS with BGP Signaling on page 542
- **interface loopback** command
- **ip address** command
- **ip router-id** command

Configuring MPLS LSPs for VPLS

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 through the provider (P) router in the MPLS core. The requirements and procedures for subscriber policies are the same whether you employ BGP or LDP signaling for VPLS.

This section explains one way to create a basic MPLS configuration using the **mpls** and **mpls ldp** commands.

To configure MPLS LSPs on the VE router:

1. Enable MPLS on the default virtual router.

   ```
   host1(config)#mpls
   ```

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

   ```
   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
   ```

3. Create an MPLS major interface stacked on the layer 2 interface Enable MPLS on the core-facing interface.

   ```
   host1(config-subif)#mpls
   ```

4. Enable LDP and topology-driven LSPs on the core-facing interface, using the default values (that is, using an implicit default profile).

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

Related Topics

- Configuration Tasks for VPLS with BGP Signaling on page 542
- Configuration Tasks for VPLS with LDP Signaling on page 558
- **mpls** command
- **mpls ldp** command
- For complete information about configuring MPLS LSPs, see *Chapter 2, Configuring MPLS*.

Configuring BGP Signaling for VPLS

This section describes one way to configure BGP signaling for VPLS, but does not provide complete details about configuring BGP and BGP/MPLS VPNs.

Table 51 lists the commands discussed in this section to configure BGP signaling for VPLS.

Table 51: 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 protocol on the VE router and specify the local AS; that is, the AS to which this BGP speaker belongs.

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

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

2. Configure the VE-to-VE BGP session by first adding an entry to the BGP neighbor table.

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

3. Use **neighbor** commands to specify the peers to which BGP advertises routes.

This example configures only the update-source and next-hop-self attributes. The update-source attribute allows the BGP session to use the IP address of a specific operational interface as the update source address for TCP connections. the next-hop-self attribute forces the BGP speaker to report itself as the next hop for an advertised route that it learned from a neighbor.

```
host1(config-router)#neighbor 10.4.4.4 update-source loopback 0
host1(config-router)#neighbor 10.4.4.4 next-hop-self
```

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

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

5. Activate the VE-to-VE session in the L2VPN address family by specifying neighbors that exchange routes from within the current address family.

   ```
   host1(config-router-af)#neighbor 10.4.4.4 activate
   ```

6. Use **neighbor** commands to configure additional address family parameters for the session, then exit the address family.

   This example configures only the next-hop-self attribute, forcing the BGP speaker to report itself as the next hop for an advertised route that it learned from a neighbor.

   ```
   host1(config-router-af)#neighbor 10.4.4.4 next-hop-self
   host1(config-router-af)#exit-address-family
   ```

7. 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.

   ```
   host1(config-router)#address-family vpls customer1
   host1(config-router-af)#exit-address-family
   host1(config-router)#address-family vpls customer2
   ```

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.

Related Topics

- Configuration Tasks for VPLS with BGP Signaling on page 542
- **address-family l2vpn** command
- **address-family vpls** command
- **exit-address-family** command

- **neighbor activate** command
- **neighbor next-hop-self** command
- **neighbor remote-as** command
- **router bgp** command
- See *Chapter 1, Configuring BGP Routing* for information about configuring BGP.
- See *Chapter 3, Configuring BGP-MPLS Applications* for information about configuring BGP/MPLS VPNs.

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

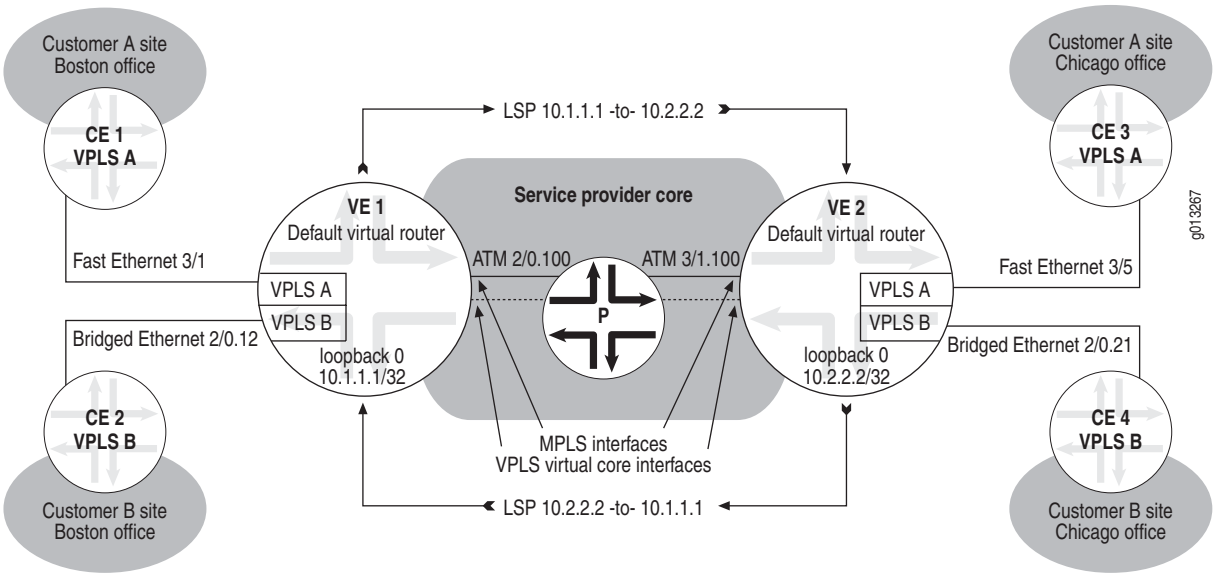

Topology Overview of VPLS with BGP Signaling

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
```

Related Topics

- Configuration Tasks for VPLS with BGP Signaling on page 542

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.

   For instructions, see *Configuring Optional Attributes for VPLS Instances* on page 545.
3. Configure network interfaces to connect the VE router to each CE device.

   For instructions, see *Configuring VPLS Network Interfaces* on page 547.
4. (Optional) Configure nondefault subscriber policies for the VPLS network interface.

   For instructions, see *Configuring Subscriber Policies for VPLS Network Interfaces* on page 548.
5. 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.

   For instructions, see *Configuring the Loopback Interface and Router ID for VPLS* on page 550.
6. Configure MPLS LSPs to connect local and remote VE routers.

   For instructions, see *Configuring MPLS LSPs for VPLS* on page 551.
7. 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.
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.

Related Topics

- Configuring VPLS Instances with LDP Signaling on page 559
- Configuring Optional Attributes for VPLS Instances on page 545
- Configuring VPLS Network Interfaces on page 547
- Configuring Subscriber Policies for VPLS Network Interfaces on page 548
- Configuring the Loopback Interface and Router ID for VPLS on page 550
- Configuring MPLS LSPs for VPLS on page 551
- Configuring LDP Signaling for VPLS on page 559

- Configuring Routing in the Core Network for VPLS on page 560
- See *VPLS Configuration Example with LDP Signaling* on page 561 for a detailed sample configuration.

Configuring VPLS Instances with 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:

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

   host1(config)#**bridge customer3 vpls transport-virtual-router vr1**

   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. 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.

Related Topics

- Configuration Tasks for VPLS with LDP Signaling on page 558
- **bridge vpls transport-virtual-router** command

Configuring LDP Signaling for VPLS

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.

Table 52 lists the commands discussed in this section to configure LDP signaling for VPLS.

Table 52: 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.

   host1(config)#**mpls ldp vpls customer3 vpls-id 3**

2. Enable LDP signaling for a VPLS instance by configuring a list of neighbor (peer) addresses on remote VE devices in the VPLS domain to which LDP can send or from which LDP can receive targeted hello messages.

   host1(config)#**mpls ldp vpls customer3 neighbor 10.3.3.3**
   host1(config)#**mpls ldp vpls customer3 neighbor 10.3.10.100**

3. Repeat these steps for each VPLs instance on the VE router.

Related Topics

- Configuration Tasks for VPLS with LDP Signaling on page 558
- **mpls ldp vpls neighbor** command
- **mpls ldp vpls vpls-id** command
- For more information about LDP, see *Chapter 2, Configuring MPLS*.

Configuring Routing in the Core Network for VPLS

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.

Table 53 lists the commands discussed in this section to configure OSPF.

Table 53: 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.

   host1(config)#**router ospf 1**

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

   host1(config-router)#**network 10.1.1.1 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**

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 10.1.1.1, and one using IP address 10.10.10.0. The **network area** commands create the two OSPF areas if they do not already exist.

Related Topics

- Configuration Tasks for VPLS with LDP Signaling on page 558
- **network area** command
- **router ospf** command
- 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*.

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

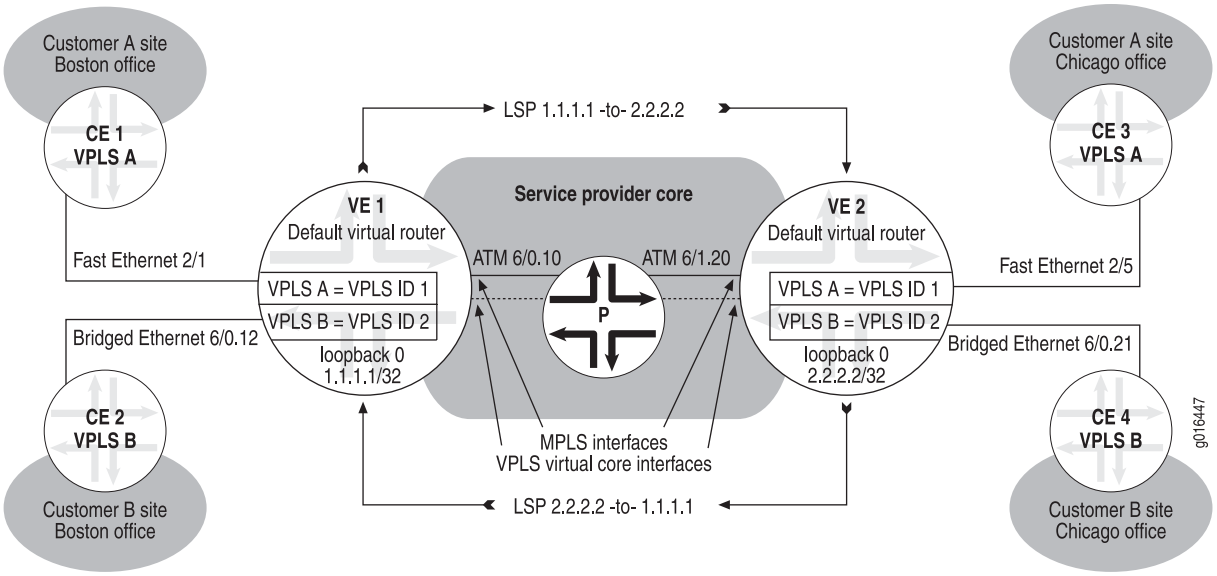

Topology Overview of VPLS with LDP Signaling

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 561 is almost identical to the sample topology shown for BGP signaling in Figure 121 on page 554. Figure 122 on page 561 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 554.

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 561.

```
! 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 561.

```
! 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
```

Related Topics

- Configuration Tasks for VPLS with LDP Signaling on page 558