Chapter 10
L2VPNs Overview

This chapter describes Layer 2 Virtual Private Networks (L2VPNs), and contains the following sections:

L2VPN Overview

L2VPNs employ layer 2 services over MPLS to build a topology of point-to-point connections that connect end customer sites in a VPN. L2VPNs provide an alternative to private networks that have been provisioned by means of dedicated leased lines or by means of layer 2 virtual circuits that employ ATM or Frame Relay. The service provisioned with L2VPNs is also known as Virtual Private Wire Service (VPWS). You configure an L2VPN *instance* on each associated edge router for each L2VPN.

Traditional VPNs over layer 2 circuits require the provisioning and maintenance of separate networks for IP and for VPN services. In contrast, L2VPNs enable the sharing of a provider's core network infrastructure between IP and L2VPN services, reducing the cost of providing those services.

L2VPNs also use BGP as the signaling protocol, and consequently have a simpler design and require less provisioning overhead than traditional VPNs over layer 2 circuits. BGP signaling also enables autodiscovery of L2VPN peers. L2VPNs are similar to BGP/MPLS VPNs and VPLS in many respects, because all three types of services employ BGP for signaling.

An L2VPN provides the same services as layer 2 over MPLS except for CE-side load-balancing. The main differences between the L2VPNs and L2 over MPLS services are signaling, autodiscovery, and configuration.

L2VPNs can have either a full-mesh or a hub-and-spoke topology. The tunneling mechanism in the core network typically is MPLS. However, L2VPNs can also use other tunneling protocols, such as GRE. L2VPNs are similar to Martini layer 2 services over MPLS, and employ a similar encapsulation scheme for forwarding traffic.

Figure 123 illustrates an example of a simple L2VPN topology.

Figure 123: L2VPN Sample Topology

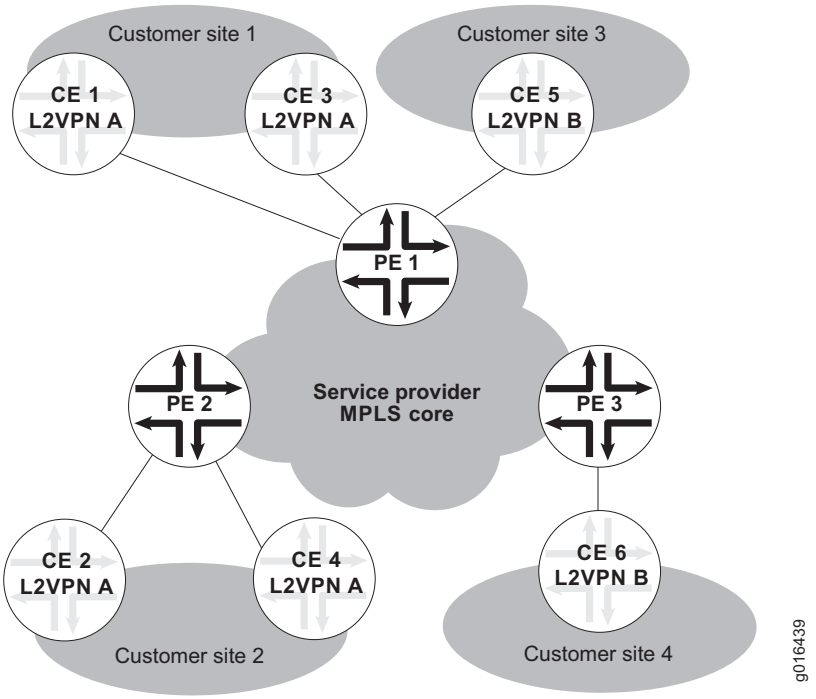

In this example, the service provider offers L2VPN services to Customer A and Customer B. Customer A wants to create a full mesh of point-to-point links between sites 1 and 2. Customer B needs only a single point-to-point link between site 3 and site 4. The service provider uses BGP and MPLS signaling in the core, and creates a set of unidirectional pseudowires at each provider edge (PE) router to separately cross-connect each customer's layer 2 circuits.

In order to provision this service, the provider configures two L2VPNs, L2VPN A and L2VPN B. An encapsulation type is configured for each VPN. All interfaces in a given L2VPN must be configured with the VPN's encapsulation type. The layer 2 interfaces that connect the PE router and CE device pairs are configured to be members of the corresponding L2VPN, L2VPN A or L2VPN B.

Local and remote site information for the interfaces identifies the cross-connect. Local cross-connects are supported when the interfaces that are connected belong to two different sites configured in the same L2VPN instance and on the same PE router.

BGP advertises reachability for the VPNs. The BGP configuration is similar to that used for other VPN services, such as layer 3 VPNs and VPLS. MPLS is configured to set up base LSPs to the remote PE routers similarly to the other VPN services.

BGP Signaling for L2VPNs

When you configure the L2VPN service at a given PE router for a given L2VPN customer, BGP signals reachability for all sites that belong to that L2VPN. This signaling is identical to the signaling used for BGP/MPLS VPNs and VPLS. The network layer reachability information (NLRI) for both services are encoded in a similar manner.

A new NLRI format carries the individual L2VPN information listed in Table 70. One or more of these NLRIs is carried in the MP_REACH_NLRI and MP_UNREACH_NLRI BGP attributes.

Table 70: Components of L2VPN NLRI

NLRI value	Size in octets
Length	2
Route Distinguisher	8
CE-ID	2
Label-block Offset	2
Label Base	3
Variable TLVs	0–*n*

The local PE router selects a contiguous label block to cover all the remote sites for a given L2VPN instance. The local PE router then advertises that label block as part of the reachability information for a given customer site in a particular L2VPN instance. This label block represents the set of demultiplexers that are used to cross-connect incoming MPLS traffic to a specific local interface in the L2VPN instance.

The local PE router also processes advertisements from all remote PE routers and for each local interface in an L2VPN instance. The local PE router selects a demultiplexer label from a label block received from the remote PE router associated with each remote site in the L2VPN instance. Traffic coming into the local interface from the CE device is cross-connected to an MPLS next hop that corresponds to the demultiplexer. Traffic is then encapsulated in MPLS and sent across the MPLS core to the remote PE router in the L2VPN.

A new address family identifier (AFI) and a new subsequent address family identifier (SAFI) are used in the NLRI for L2VPNs. The identifier values have yet to be assigned by IANA.

The L2VPN NLRIs must be accompanied by a route-target extended community. PE routers that receive VPN information can filter route advertisements with the route target import lists and export lists. This route filtering enables the PE routers to control CE-to-CE connectivity.

An L2VPN NLRI is uniquely identified by the route distinguisher, CE ID, and the label block offset.

In addition to the site ID and label block information, BGP also signals control flags that indicate whether a control word is included in the encapsulation and whether packets have a sequence number. If a control word mismatch occurs, the pseudowire remains in a down state with a status of control word mismatch.

A control status vector is sent along with the other NLRI information. This vector carries the operational state of the local layer 2 interfaces between the PE router and CE device for a given L2VPN instance. A TLV type of 1 is used currently to interoperate with JUNOS software.

Related Topics

- *Chapter 3, Configuring BGP-MPLS Applications*
- *Chapter 8, Configuring VPLS*

L2VPN Components

Figure 124 shows the components of a typical L2VPN topology.

Figure 124: L2VPN Components

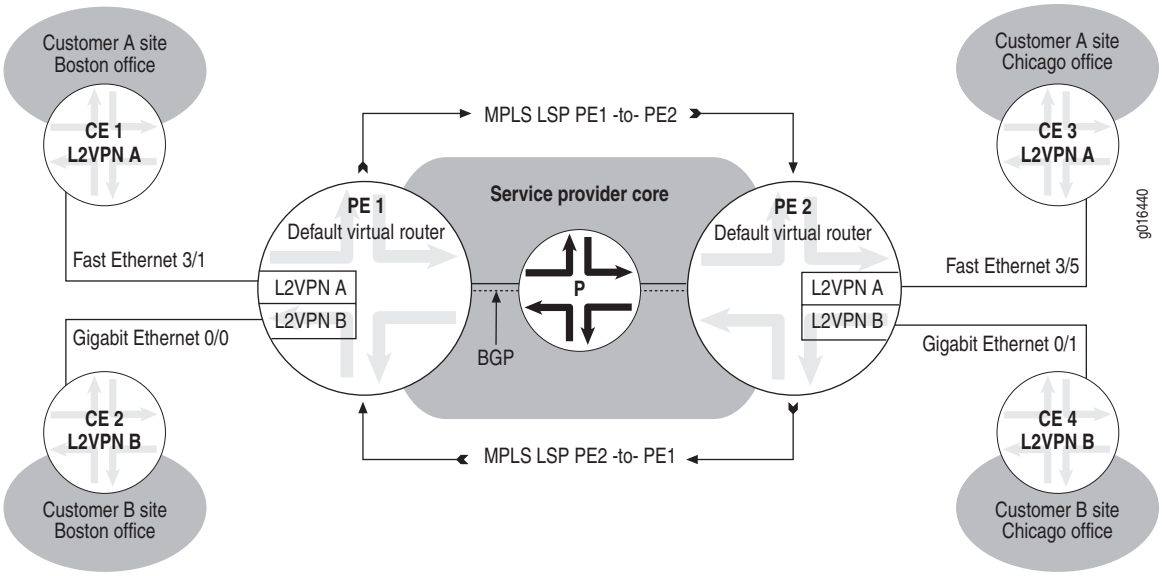

L2VPN Instances

Typically, an L2VPN is associated with customers who want to use L2VPNs to connect geographically dispersed sites in their organization across an MPLS-based service provider core, also known as an MPLS backbone. Each L2VPN consists of the set of provider edge routers running the corresponding L2VPN instance. To provide connectivity for the L2VPN, BGP builds a full mesh of pseudowires among all of the L2VPN instances on each of the provider edge routers participating in a particular L2VPN.

Figure 124 on page 594 depicts two L2VPNs: L2VPN A and L2VPN B. L2VPN A connects Customer A's Boston and Chicago offices, and consists of provider edge routers PE 1 and PE 2, each of which runs an L2VPN instance named l2vpnA. Similarly, L2VPN B connects Customer B's Boston and Chicago offices, and consists of provider edge routers PE 1 and PE 2, each of which also runs an L2VPN instance named l2vpnB.

Customer Edge Devices

Figure 124 on page 594 shows four customer edge devices: CE 1, CE 2, CE 3, and CE 4. Each CE device is located at the edge of a customer site. In the sample topology, CE 1 and CE 3 are members of L2VPN A, and CE 2 and CE 4 are members of L2VPN B.

A CE device can be a single host, a switch, or, most typically, a router. Each CE device is directly connected to an L2VPN provider edge router by means of a layer 2 interface.

L2VPN Provider Edge Devices

In an L2VPN configuration, E-series routers function as provider edge devices, which are also referred to as PE routers. These PE routers perform a similar function to PE routers in a BGP/MPLS VPN configuration.

Figure 124 on page 594 depicts two PE routers: PE 1, which is the local router, and PE 2, which is the remote router located at the other side of the service provider core. Each PE router must have an L2VPN instance configured for each L2VPN in which it participates. Consequently, the sample topology comprises four separate L2VPN instances: instances l2vpnA and l2vpnB configured on PE 1, and instances l2vpnA and l2vpnB configured with matching route target values on PE 2.

Each L2VPN instance configured on the router is associated with two types of interfaces. The CE-facing or customer-facing interface is a layer 2 interface that directly connects the PE router to a local CE device. The router encapsulates layer 2 frames from the CE device in an MPLS packet and then forwards the encapsulated frames to the service provider core from an MPLS interface through the provider (P) router. This encapsulation is identical to Martini encapsulation for layer 2 services over MPLS.

L2VPNs and BGP/MPLS VPNs

BGP multiprotocol extensions (MP-BGP) enable BGP to support IPv4 services such as BGP/MPLS VPNs, which are sometimes known as RFC 2547bis VPNs. An L2VPN 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 L2VPNs are similar except that for L2VPNs you must configure both of the following address families:

- L2VPN—The L2VPN address family enables you to configure the PE router (L2VPNs) or VE router (VPLS) to exchange layer 2 NLRI for all L2VPN or VPLS instances. Optionally, you can use the **signaling** keyword with the **address-family** command for the L2VPN address family to specify BGP signaling of L2VPN reachability information. Currently, you can omit the **signaling** keyword with no adverse effects.

- VPWS—The VPWS address family enables you to configure the PE router to exchange layer 2 NLRI for a specified VPWS (L2VPN) instance.

BGP can exchange information in an L2VPN topology within these address families.

Related Topics

- *Chapter 1, Configuring BGP Routing*
- *Chapter 3, Configuring BGP-MPLS Applications*

L2VPN Supported Features

The JUNOSe software implementation of L2VPNs provides the following features:

- Support for the following types of network interfaces between the PE router and the CE device:
  - ATM with ATM Adaptation Layer 5 (AAL5) encapsulation
  - ATM with virtual channel connection (VCC) cell relay encapsulation
  - Cisco HDLC
  - Ethernet (Fast Ethernet, Gigabit Ethernet, 10-Gigabit Ethernet, Ethernet/VLAN)
  - Frame Relay
  - PPP
- Autodiscovery of L2VPN instance members using MP-BGP
- L2VPN signaling using MP-BGP to set up and tear down the pseudowires that constitute an L2VPN instance
- Inter-AS option A, inter-AS option B, and inter-AS option C services

As with VPLS, L2VPNs do not support BGP multipaths.

L2VPN Platform Considerations

L2VPNs are supported on all E-series routers:

Module Requirements

You can configure L2VPNs on all E-series module combinations that support MPLS tunnels.

For information about the modules that support L2VPNs on ERX-14xx models, ERX-7xx models, and ERX-310 routers:

- See *ERX Module Guide, Chapter 1, Module Combinations* for detailed module specifications.
- See *ERX Module Guide, Appendix A, Module Protocol Support* for information about the modules that support L2VPNs.

For information about the modules that support L2VPNs 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 an L2VPN 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, adapter 0 identifies the right IOA bay (E120 router) and the upper IOA bay (E320 router); adapter 1 identifies the left IOA bay (E120 router) and the lower IOA bay (E320 router). 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
```

Related Topics

- 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*.

L2VPN References

For more information about L2VPNs, consult the following resources:

- Layer 2 VPNs over Tunnels—draft-kompella-l2vpn-l2vpn-01.txt (July 2006 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.