

Chapter 15

VPLS Overview

This chapter provides an overview of virtual private LAN service (VPLS) as it is implemented in the JUNOS software.

For information about virtual private networks (VPNs) and the differences between Layer 2 VPNs, Layer 3 VPNs, and VPLS, see “VPN Overview” on page 3.

This chapter discusses the following topics that provide background information about VPLS:

- VPLS Overview on page 377
- VPLS Standards on page 378
- Supported Platforms and PICs on page 378
- VPLS Routing and Virtual Ports on page 379
- VPLS Multihoming on page 380

VPLS Overview

VPLS is an Ethernet-based point-to-multipoint Layer 2 VPN. It allows you to connect geographically dispersed Ethernet local area networks (LAN) sites to each other across a Multiprotocol Label Switching (MPLS) backbone. For customers who implement VPLS, all sites appear to be in the same Ethernet LAN even though traffic travels across the service provider’s network.

VPLS, in its implementation and configuration, has much in common with a Layer 2 VPN. In a VPLS, a packet originating within a service provider customer’s network is sent first to a customer edge (CE) device (for example, a router or Ethernet switch). It is then sent to a provider edge (PE) router within the service provider’s network. The packet traverses the service provider’s network over a MPLS label-switched path (LSP). It arrives at the egress PE router, which then forwards the traffic to the CE device at the destination customer site.

The difference is that for a VPLS, packets can traverse the service provider's network in point-to-multipoint fashion, meaning that a packet originating from a CE device can be broadcast to all the PE routers participating in a VPLS routing instance. In contrast, a Layer 2 VPN forwards packets in point-to-point fashion only.



NOTE: Graceful Routing Engine switchover is not supported for VPLS.

VPLS Standards

VPLS is described in the Internet draft, draft-ietf-l2vpn-vpls-bgp-08.txt, *Virtual Private LAN Service (VPLS) Using BGP for Auto-discovery and Signaling* (expires December 2006).

You can access Internet RFCs and drafts on the IETF Web site at <http://www.ietf.org>.

Supported Platforms and PICs

VPLS is supported on the following M-series platforms:

- M5
- M7i
- M10
- M10i
- M20
- M40
- M40e
- M320

VPLS is also supported on all T-series platforms.

VPLS is supported on the following PICs:

- All ATM2 IQ PICs
- 4-port Fast Ethernet PIC with 10/100 Base-TX interfaces PIC
- 1-port Gigabit Ethernet PIC
- 1-port 10 Gigabit Ethernet PIC
- 1-port Gigabit Ethernet Intelligent Queuing (IQ) PIC
- 2-port Gigabit Ethernet PIC
- 2-port Gigabit Ethernet IQ PIC

- 4-port, quad-wide Gigabit Ethernet PIC
- 10-port Gigabit Ethernet PIC

To enable VPLS on geographically remote sites of a VPLS domain, the router requires an Adaptive Services PIC, Link Services PIC, or Tunnel Services PIC.

VPLS Routing and Virtual Ports

Because a VPLS carries Ethernet traffic across a service provider network, it must mimic an Ethernet network in some ways. When a PE router configured with a VPLS routing instance receives a packet from a CE device, it first determines whether it knows the destination of the VPLS packet. If it does, it forwards the packet to the appropriate PE router or CE device. If it does not, it broadcasts the packet to all other PE routers and CE devices that are members of that VPLS routing instance. In both cases, the CE device receiving the packet must be different from the one sending the packet.

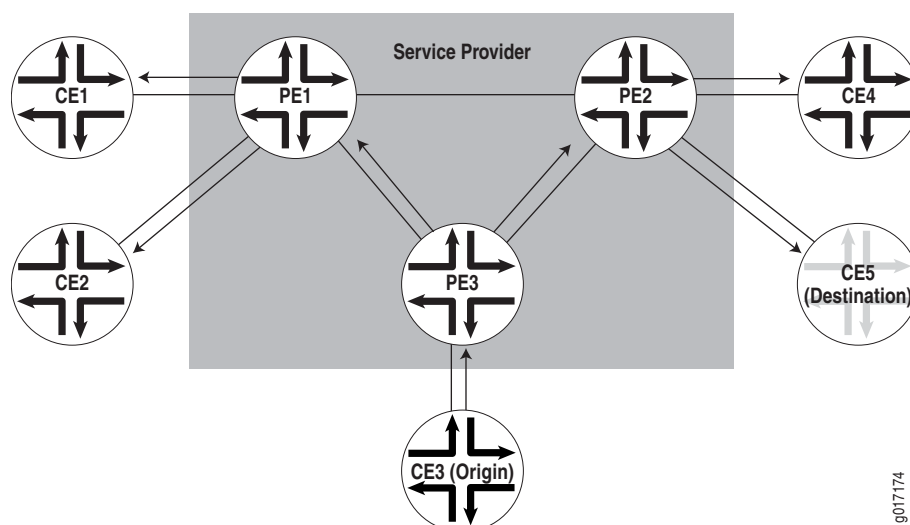
When a PE router receives a packet from another PE router, it first determines whether it knows the destination of the VPLS packet. If the destination is known, the PE router either forwards the packet or drops it depending on whether the destination is a local or remote CE device:

- If the destination is a local CE device, the PE router forwards the packet to it.
- If the destination is a remote CE device (connected to another PE router), the PE router discards the packet.

If the PE router cannot determine the destination of the VPLS packet, it floods the packet to all attached CE devices.

This process is illustrated in Figure 45.

Figure 45: Flooding a Packet with an Unknown Destination to All PE Routers in the VPLS Instance



g017174

A VPLS can be directly connected to an Ethernet switch. Layer 2 information gathered by an Ethernet switch (for example, media access control [MAC] addresses and interface ports) is included in the VPLS routing instance table. However, instead of all VPLS interfaces being physical switch ports, the router allows remote traffic for a VPLS instance to be delivered across an MPLS LSP and arrive on a virtual port. The virtual port emulates a local, physical port. Traffic can be learned, forwarded, or flooded to the virtual port in almost the same way as traffic is sent to a local port.

The VPLS routing table learns MAC address and interface information for both physical and virtual ports. The main difference between a physical port and a virtual port is that the router captures additional information from the virtual port—an outgoing MPLS label used to reach the remote site and an incoming MPLS label for VPLS traffic received from the remote site. The virtual port is generated dynamically on a Tunnel Services Physical Interface Card (PIC) when you configure VPLS on the router. A Tunnel Services PIC is required on each VPLS router.

One restriction on flooding behavior in VPLS is that traffic received from remote PE routers is never forwarded to other PE routers. This restriction helps prevent loops in the core network. However, if a CE Ethernet switch has two or more connections to the same PE router, you must enable the Spanning Tree Protocol (STP) on the CE switch to prevent loops. STP is not supported directly on M-series routers.



NOTE: The JUNOS software allows standard Bridge Protocol Data Unit (BPDU) frames to pass through emulated Layer 2 connections, such as those configured with Layer 2 VPNs, Layer 2 circuits, and VPLS routing instances. However, CE Ethernet switches that generate proprietary BPDU frames might not be able to run STP across Juniper Networks routing platforms configured for these emulated Layer 2 connections.

VPLS Multihoming

VPLS multihoming allows you to connect a customer site to multiple PE routers to provide redundant connectivity while preventing the formation of Layer 2 loops in the service provider's network. A VPLS site multihomed to two or more PE routers provides redundant connectivity in the event of a PE router-to-CE device link failure or the failure of a PE router.

When multihoming a VPLS site (potentially in different autonomous systems [Ass]), the PE routers connected to the same site can either be configured with the same VPLS edge (VE) device identifier or with different VE device identifiers. In the latter case, you must run STP on the CE device, and possibly on the PE routers, to construct a loop-free VPLS topology.

If the PE routers are connected to the same site and assigned the same VE device identifier, a loop-free topology is constructed using a routing mechanism such as BGP path selection. When a BGP speaker receives two equivalent network layer reachability information (NLRIs), it applies standard path selection criteria such as local preference and AS path length to determine which NLRI to choose; it selects only one.

Since a PE router would just pick one of the received NLRI advertisements with a particular VE device identifier, it would establish pseudowires to only one of the remote PE routers, the PE router that originated the winning advertisement. This prevents multiple paths from being created in the network between sites, preventing the formation of Layer 2 loops in the network. If the selected PE router fails, all PE routers in the network automatically switch to the backup PE router and establish pseudowires through the backup PE router.

Two VPLS NLRIs are considered equivalent from a path selection perspective if the following are the same:

- Route distinguisher
- VE device identifier
- VE block offset

If two PE routers are assigned the same VE device identifier in a given VPLS, they must also use the same route distinguisher, and they should advertise the same VE block size for a given VE offset.



NOTE: Traffic loss can occur when the old pseudowires are brought down and new ones established.
