

Chapter 14

VPLS Overview

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

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

This chapter includes the following sections:

VPLS Overview on page 279

VPLS Routing and Virtual Ports on page 280

VPLS Standards on page 281

Supported Platforms and PICs on page 281

VPLS Overview

VPLS is an Ethernet-based point-to-multipoint Layer 2 VPN. It allows you to connect geographically dispersed Ethernet LAN sites to each other across an 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 Multiprotocol Label Switching (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.

VPLS Routing and Virtual Ports

Since 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. If it doesn't, it broadcasts the packet to all the other PE routers that are members of that particular VPLS routing instance. The PE routers forward the packet to their CE devices. The CE device that is the intended recipient of the packet forwards it to its final destination. The other CE devices discard it. This process is illustrated in Figure 40.

Figure 40: Flooding a Packet with an Unknown Destination to All PE Routers in the VPLS Instance



A VPLS can be directly connected to an Ethernet switch. Layer 2 information gathered by an Ethernet switch (for example, 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 label-switched path (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 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 PIC when you configure VPLS on the router (a Tunnel PIC is required on each VPLS router).

One restriction on flooding behavior in VPLS is that traffic received from remote provider edge (PE) routers is never forwarded to other PE routers. This restriction helps prevent loops in the core network. However, if a customer edge (CE) Ethernet switch has two connections or more to the same PE router, you must enable the Spanning Tree Protocol on the CE switch to prevent loops. Spanning Tree Protocol (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 instances. However, customer edge (CE) Ethernet switches that generate proprietary BPDU frames might not be able to run STP across Juniper Networks routers configured for these emulated Layer 2 connections.

VPLS Standards

VPLS is described in the Internet draft draft-kompella-ppvnp-vpls-01.txt, *Virtual Private LAN Service*.

Supported Platforms and PICs

VPLS is supported on the following M-series platforms:

M5

M10

M20

M40

M40e

VPLS is supported on the following PICs:

Four-port Fast Ethernet PIC with 10/100 Base-TX interfaces

One-port Gigabit Ethernet PIC

Two-port Gigabit Ethernet PIC

Four-port, quad-wide Gigabit Ethernet PIC

To enable VPLS on geographically remote sites of a VPLS domain, a Tunnel Services PIC is required.

