

Chapter 1

VPN Overview

A virtual private network (VPN) consists of two topological areas, the provider's network and the customer's network. The provider's network, which runs across the public Internet infrastructure, consists of routers that provide VPN services to a customer's network as well as routers that provide other services. The customer's network is commonly located at multiple physical sites. The provider's network acts to connect the various customer sites in what appears to the customer and the provider to be a private network.

To ensure that VPNs remain private and isolated from other VPNs and from the public Internet, the provider's network maintains policies that keep routing information from different VPNs separate.

A provider can service multiple VPNs as long as its policies keep routes from different VPNs separate. Similarly, a site can belong to multiple VPNs as long as it keeps routes from the different VPNs separate.

VPN Terminology

VPNs contain the following types of network devices (see Figure 1):

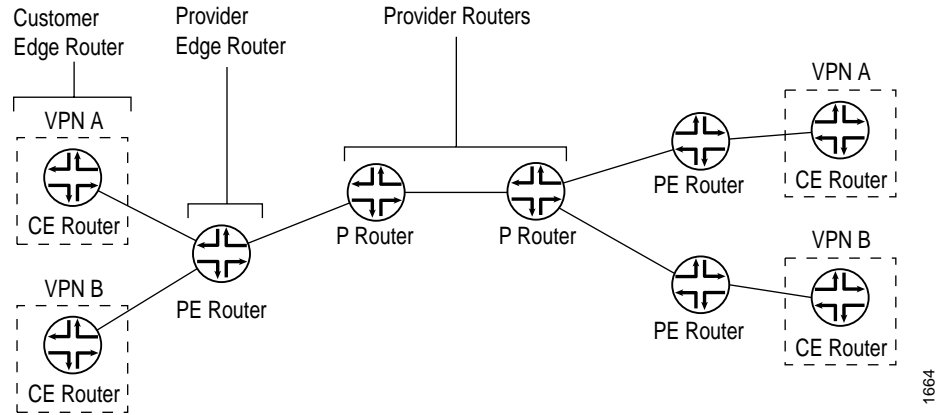
Provider edge (PE) routers—Routers in the provider's network that connect to customer edge devices located at customer sites. PE routers support VPN and label functionality. (The label functionality can be provided either by the Resource Reservation Protocol [RSVP] or Label Distribution Protocol [LDP].) Within a single VPN, pairs of PE routers are connected through a tunnel, which can be either an Multiprotocol Label Switching (MPLS) label-switched path (LSP) or an LDP tunnel.

Provider (P) routers—Routers within the core of the provider's network that are not connected to any routers at a customer site but are part of the tunnel between pairs of PE routers. Provider routers support MPLS LSP or LDP functionality, but do not need to support VPN functionality.

Customer edge (CE) devices—Routers or switches located at the customer's site that connect to the provider's network. CE devices are typically IP routers.

VPN functionality is provided by the PE routers; the provider and CE routers have no special configuration requirements for VPNs.

Figure 1: VPN Router Components



Types of VPNs

The JUNOS Internet software provides several types of VPNs; You can choose the best solution for your network environment. Each of the following VPNs has different capabilities and requires different types of configuration:

Layer 2 VPNs on page 4

Layer 3 VPNs on page 5

VPLS on page 5

Virtual-Router Routing Instances on page 6

Layer 2 VPNs

Implementing a Layer 2 VPN on a router is similar to implementing a VPN using a Layer 2 technology such as Asynchronous Transfer Mode (ATM) or Frame Relay. However, for a Layer 2 VPN on a router, traffic is forwarded to the router in a Layer 2 format. It is carried by MPLS over the service provider's network and then converted back to Layer 2 format at the receiving site. You can configure different Layer 2 formats at the sending and receiving sites. The security and privacy of an MPLS Layer 2 VPN are equal to those of an ATM or Frame Relay VPN.

On a Layer 2 VPN, routing occurs on the customer's routers, typically on the customer edge (CE) router. The CE router connected to a service provider on a Layer 2 VPN must select the appropriate circuit on which to send traffic. The provider edge (PE) router receiving the traffic sends it across the service provider's network to the PE router connected to the receiving site. PE routers do not need to know the customer's routes or routing topology; they need to know only in which tunnel to send the data.

For a Layer 2 VPN, customers need to configure their own routers to carry all Layer 3 traffic. The service provider needs to know only how much traffic the Layer 2 VPN will need to carry. The service provider's routers carry traffic between the customer's sites using Layer 2 VPN interfaces. The VPN topology is determined by policies configured on the PE routers.

Layer 3 VPNs

In a Layer 3 VPN, the routing occurs on the service provider's routers. Therefore Layer 3 VPNs require more configuration on the part of the service provider, because the service provider's PE routers must know the customer's routes.

In JUNOS, Layer 3 VPNs are based on the Internet draft draft-rosen-rfc2547bis. This Internet draft defines a mechanism by which service providers can use their IP backbones to provide Layer 3 VPN services to their customers. The sites that make up a Layer 3 VPN are connected over a provider's existing public Internet backbone.

VPNs based on draft-rosen-rfc2547bis are also known as BGP/MPLS VPNs because BGP is used to distribute VPN routing information across the provider's backbone, and MPLS is used to forward VPN traffic across the backbone to remote VPN sites.

Customer networks, because they are private, can use either public addresses or private addresses, as defined in RFC 1918. When customer networks that use private addresses connect to the public Internet infrastructure, the private addresses might overlap with the private addresses used by other network users. MPLS/BGP VPNs solve this problem by prefixing a VPN identifier to each address from a particular VPN site, thereby creating an address that is unique both within the VPN and within the public Internet. In addition, each VPN has its own VPN-specific routing table that contains the routing information for that VPN only.

VPLS

VPLS allows you to connect geographically dispersed customer sites as if they were connected to the same LAN. In many ways, it works like a Layer 2 VPN. VPLS and Layer 2 VPNs use the same network topology and function similarly. A packet originating within a service provider customer's network is sent first to a customer edge (CE) device. 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 key difference in VPLS is that packets can traverse the service provider's network in a point-to-multipoint fashion, meaning that a packet originating from a CE device can be broadcast to PE routers in the VPLS. In contrast, a Layer 2 VPN forwards packets in a point-to-point fashion only. The destination of a packet received from a CE device by a PE router must be known for the Layer 2 VPN to function properly.

VPLS is designed to carry Ethernet traffic across an MPLS-enabled service provider network. In certain ways, VPLS mimics the behavior of an Ethernet network. 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 it to the appropriate PE router. If it doesn't know the destination, it broadcasts the packet to all the other PE routers that are members of the same 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.

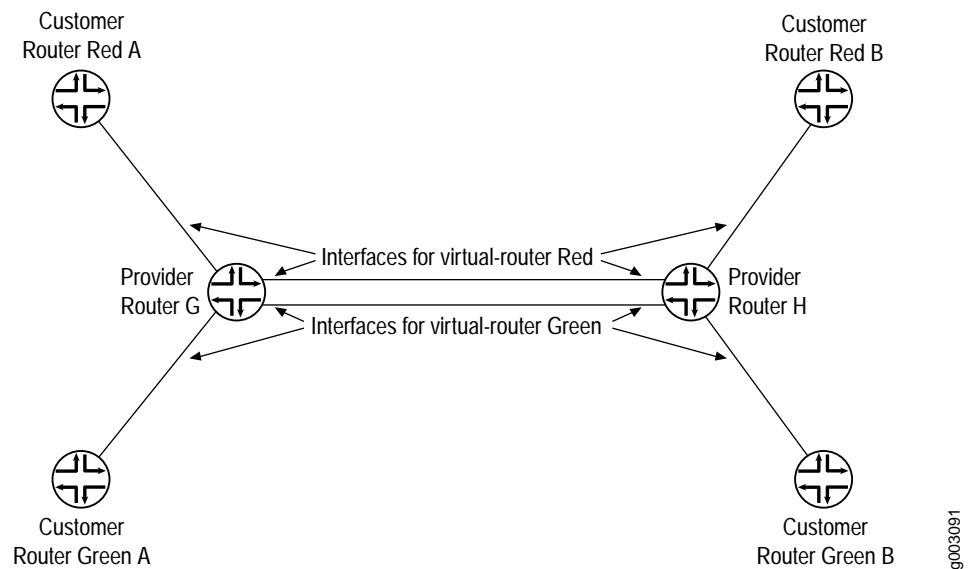
Virtual-Router Routing Instances

A virtual-router routing instance, like a VRF (Layer 3 VPN) routing instance, maintains separate routing and forwarding tables for each instance. However, many configuration steps required for VRF routing instances are not required for virtual-router routing instances. Specifically, you do not need to configure a route distinguisher, a routing table policy (the `vrf-export`, `vrf-import`, and `route-distinguisher` statements), or MPLS between the provider routers.

However, you need to configure separate logical interfaces between each of the service provider routers participating in a virtual-router routing instance. You also need to configure separate logical interfaces between the service provider routers and the customer routers participating in each routing instance. Each virtual-router instance requires its own unique set of logical interfaces to all participating routers.

Figure 2 shows how this works. The provider routers G and H are configured for virtual-router routing instances Red and Green. Each provider router is directly connected to two local customer routers, one in each routing instance. The provider routers are also connected to each other over the service provider network. These routers need four logical interfaces: a logical interface to each of the locally connected customer routers and a logical interface to carry traffic between the two provider routers for each virtual-router instance.

Figure 2: One Logical Interface Configured for Each Router in Each Virtual-Router Routing Instance



Layer 3 VPNs do not have this configuration requirement. If you configure several Layer 3 VPN routing instances on a PE router, all the instances can use the same logical interface to reach another PE router. This is possible because Layer 3 VPNs use MPLS (VPN) labels that differentiate traffic going to and from various routing instances. Without MPLS and VPN labels, as in a virtual-router routing instance, you need separate logical interfaces to separate traffic from different instances.

One method of providing this logical interface between the provider routers is by configuring tunnels between them. You can configure IPSec, GRE, or IP-IP tunnels between the provider routers, terminating the tunnels at the virtual-router instance.

VPN Graceful Restart

VPN graceful restart allows a router whose VPN control plane is undergoing a restart to continue to forward traffic while recovering its state from neighboring routers. Without graceful restart, a control plane restart disrupts any Layer 2 or Layer 3 VPN services provided by the router.

For VPN graceful restart to function properly, the following needs to be configured on the PE router:

- BGP graceful restart must be active on the PE-to-PE sessions carrying any service signaling data in the session's network layer reachability information (NLRI).

- OSPF, ISIS, LDP, and RSVP graceful restart must be active, because routes added by these protocols are used to resolve Layer 2 and Layer 3 VPN NLRIs.

- For other protocols (static, RIP, OSPF, LDP and so on), graceful restart functionality must also be active when these protocols are run between the PE and CE routers. Layer 2 VPNs do not rely on this because protocols are not configured between the PE and CE routers.

In VPN graceful restart, a restarting router does the following:

- Waits for all the BGP NLRI information from other PE routers before it starts advertising routes to its CE routers.

- Waits for all protocols in all routing instances to converge (or finish graceful restart) before sending CE router information to the other PE routers.

- Waits for all routing instance information (whether it is local configuration or advertisements from a remote peer router) to be processed before sending it to the other PE routers.

- Preserves all forwarding state information in the MPLS routing tables until new labels and transit routes are allocated and then advertises them to other PE routers (and CE routers in carrier-of-carriers VPNs).

