CHAPTER 4

Configuring Multiprotocol Label Switching

Multiprotocol label switching (MPLS) combines the performance and capabilities of Layer 2 (data link layer) switching with the proven scalability of Layer 3 (network layer) routing. MPLS enables service providers to meet the challenges of explosive growth in network utilization while providing the opportunity to differentiate services without sacrificing the existing network infrastructure. The MPLS architecture is flexible and can be employed in any combination of Layer 2 technologies. MPLS support is offered for all Layer 3 protocols, and scaling is possible well beyond that typically offered in today's networks.

This chapter describes the following MPLS-related features:

For more information about MPLS, see Chapter 3, "Configuring Remote Access to MPLS VPN" and see the *Multiprotocol Label Switching on Cisco Routers, Release 12.1(3)T feature module*.

BGP Multipath Load Sharing for eBGP and iBGP in an MPLS VPN

Load sharing is a concept that allows the Cisco 10000 series router to take advantage of multiple best paths to a given destination. The paths are derived either statically or with dynamic protocols such as RIP, BGP, OSPF, and IGRP. The best path algorithm decides which is the best path to install in the IP routing table and to use for forwarding traffic.

The BGP Multipath Load Sharing for eBGP and iBGP in an MPLS VPN feature allows you to configure multipath load sharing with both external Border Gateway Protocol (eBGP) and internal BGP (iBGP) paths in BGP networks that are configured to use Multiprotocol Label Switching (MPLS) Virtual Private Networks (VPNs). BGP Multipath Load Sharing provides improved load sharing deployment and service offering capabilities and is useful for multihomed autonomous systems and provider edge (PE) routers that import both eBGP and iBGP paths from multihomed and stub networks.

BGP installs up to the maximum number of paths allowed (configured using the **maximum-paths** command). BGP uses the best path algorithm to select one multipath as the best path, insert the best path into the routing information base (RIB), and advertise the best path to BGP peers. Other multipaths may be inserted into the RIB, but only one path is selected as the best path.

Note The maximum number of configurable paths on the PRE2 is 6.

Cisco Express Forwarding (CEF) uses the multipaths to perform load sharing, which can be performed on a per-packet or per-source/destination pair basis. By default, the BGP Multipath Load Sharing for Both eBGP and iBGP in an MPLS VPN feature performs unequal cost load sharing by selecting BGP paths that do not have an equal cost of the Interior Gateway Protocol (IGP). To enable the feature, configure the router with MPLS VPNs that contain virtual routing and forwarding instances (VRFs) that import both eBGP and iBGP paths. You can configure the number of multipaths separately for each VRF.

Note The BGP Multipath Load Sharing for Both eBGP and iBGP in an MPLS VPN feature operates within the configuration parameters of the existing outbound routing policy.

The BGP Multipath Load Sharing for eBGP and iBGP in an MPLS VPN feature is described in the following topics:

Feature History for BGP Multipath Load Sharing for eBGP and iBGP in an MPLS VPN

Cisco IOS Release	Description	Required PRE
12.3(7)XI1	This feature was introduced on the Cisco 10000 series router.	PRE2
12.2(28)SB	This feature was integrated into Cisco IOS Release 12.2(28)SB.	PRE2
12.2(33)SB	The IGP convergence acceleration feature was added on Cisco 10000 series router.	PRE3 and PRE4
12.2(33)SB3	The IGP convergence acceleration feature was updated to include support for unequal cost paths on Cisco 10000 series router.	PRE3 and PRE4

Restrictions for BGP Multipath Load Sharing for eBGP and iBGP in an MPLS VPN

The BGP Multipath Load Sharing for Both eBGP and iBGP in an MPLS VPN feature has the following restrictions:

- The Cisco 10000 series router supports recursive load sharing, but with the following restriction.

  In recursive load sharing, the information required to forward a packet requires at least 2 lookups. The first lookup determines which provider edge (PE) router is used to reach the final destination. The second lookup determines how to reach the PE router (from first lookup).

  When you configure MPLS VPN, CEF uses recursive load sharing. The first lookup provides the VPN label, the second lookup provides the IGP label. When PXF forwards a packet, it does only 1 lookup which provides both a VPN and an IGP label; 2 lookups in CEF are combined into 1. The restriction for recursive load sharing when PXF forwards a packet is as follows.

  When there are multiple IGP paths between a Cisco 10000 Series PE router to a provider router (P), only per-tag load sharing is supported. That is, PXF is programmed with only one of the paths and this one path is chosen in a round-robin fashion. Because the path is chosen at prefix setup time, it is not possible to predict which path will be selected for which prefix. The path selected depends on the order in which the prefixes are configured in the routing table. The bandwidths of the IGP paths are not considered in the path selection.
- When the routing table contains multiple iBGP paths, a route reflector advertises only one of the paths (one next hop). If a router is behind a route reflector, all routers that are connected to multihomed sites are not advertised unless separate VRFs with different route distinguishers (RDs) are configured for each VRF.
- Each IP routing table entry for a BGP prefix that has multiple iBGP paths uses additional memory. We recommend not using this feature on a router with a low amount of available memory and especially when the router is carrying a full Internet routing table.

Prerequisites for BGP Multipath Load Sharing for eBGP and iBGP in an MPLS VPN

The BGP Multipath Load Sharing for Both eBGP and iBGP in an MPLS VPN feature has the following requirements:

- MPLS VRFs must be configured before load sharing with both eBGP and iBGP routes can be configured.

IGP Convergence Acceleration

From Cisco IOS Release 12.2(33)SB onward, Cisco 10000 series routers support IGP and VPN load balancing for MPLS VPN scenarios on PRE3 and PRE4 engines. This support allows faster failover of IGP routes during load balancing. Therefore, for equal cost paths (load-balanced case), convergence is within an acceptable range.

For unequal cost paths, convergence depends on the number of BGP prefixes; a failover can be more than 30 seconds. From Cisco IOS Release 12.2(33)SB3 onward, Cisco 10000 series routers also support unequal cost paths.

The IGP Convergence Acceleration feature leverages the load balance infrastructure of equal paths for unequal paths. An indirection object is inserted for table output chain building. When the interface goes down, the routing protocols purge their routes and the inplace modifier prevents marking of labels with incomplete adjacency. Therefore, the indirection object can use the new adjacency labels. This update in the IGP Convergence Acceleration feature makes the convergence independent of the number of BGP prefixes, thereby allowing a faster failover.

Configuring IGP Convergence Acceleration

To configure the IGP Convergence Acceleration feature for unequal cost paths, enter the following commands beginning in global configuration mode:

	Command	Purpose
Step 1	`Router(config)# cef table output-chain build favor convergence-speed`	Enables the output chain building to favor convergence.
Step 2	`Router(config)# cef table output-chain build inplace-modify load-sharing`	Enables the output chain building for inplace-modify.
Step 3	`Router(config)# ip routing protocol purge interface`	Allows the routing protocol to purge routes without letting RIB delete BGP prefixes.

Configuring BGP Multipath Load Sharing for eBGP and iBGP in an MPLS VPN

To configure the BGP Multipath Load Sharing for Both eBGP and iBGP in an MPLS VPN feature, perform the following configuration tasks:

- Configuring Multipath Load Sharing for eBGP and iBGP, page 4-5
- Verifying Multipath Load Sharing for eBGP and iBGP, page 4-5

Configuring Multipath Load Sharing for eBGP and iBGP

To configure iBGP and eBGP routes for multipath load sharing, enter the following commands beginning in global configuration mode:

	Command	Purpose
Step 1	`Router(config)# router bgp as-number`	Configures the router to run a BGP process and enters router configuration mode.
Step 2	`Router(config-router)# address-family ipv4 vrf vrf-name`	Configures a VRF instance for an IPv4 session and enters address family configuration mode.
Step 3	`Router(config-router-af)# maximum-paths eibgp number-of-paths`	Configures the number of parallel iBGP and eBGP routes that can be installed into a routing table. The maximum number of configurable paths on the PRE2 is 6. **Note** You must configure the **maximum-paths eibgp** command in address family IPv4 VRF configuration mode. You cannot configure the command in any other address family configuration mode.

Example 4-1 shows how to configure a VRF named *main* for an IPv4 session and configures a router to select 4 eBGP or iBGP paths as multipaths in address family configuration mode.

Example 4-1 Configuring eBGP and iBGP Multipath Load Sharing

```
Router(config)# router bgp 50
Router(config-router)3 address-family ipv4 vrf main
Router(config-router-af)# maximum-paths eibgp 4
```

Verifying Multipath Load Sharing for eBGP and iBGP

To verify that iBGP and eBGP routes are configured for load sharing, enter any of the following commands in privileged EXEC mode:

Command	Purpose
`Router# show ip bgp vpnv4 all`	Displays all available VPNv4 information from the BGP database.
`Router# show ip route vrf vrf-name`	Displays the IP routing table associated with a virtual routing and forwarding (VRF) instance.

Configuration Examples for BGP Multipath Load Sharing for eBGP and iBGP in an MPLS VPN

This section provides the following configuration examples:

- eBGP and iBGP Multipath Load Sharing Configuration Example, page 4-6
- Verifying eBGP and iBGP Multipath Load Sharing, page 4-6

eBGP and iBGP Multipath Load Sharing Configuration Example

Example 4-2 configures a router to select six eBGP or iBGP paths as multipaths in address family configuration mode:

Example 4-2 Configuring eBGP and iBGP Multipath Load Sharing

```
Router(config)# router bgp 100
Router(config-router)# address-family ipv4 vrf vrf-1
Router(config-router-af)# maximum-paths eibgp 6
```

Verifying eBGP and iBGP Multipath Load Sharing

Example 4-3 shows sample output displayed when you enter the **show ip bgp vpnv4** command. The third line of output (`Multipath:eiBGP`) indicates that multipath load sharing is on.

Example 4-3 Verifying eBGP and iBGP Multipath Load Sharing

```
Router# show ip bgp vpnv4 all 10.22.22.0

BGP routing table entry for 10:1:22.22.22.0/24, version 19
Paths: (5 available, best #5)
Multipath:eiBGP
    Advertised to non peer-group peers:
    10.0.0.2 10.0.0.3 10.0.0.4 10.0.0.5
    22
        10.0.0.0 (metric 20) from 10.0.0.4 (10.0.0.4)
            Origin IGP, metric 0, localpref 100, valid, internal, multipath
            Extended Community:0x0:0:0 RT:100:1 0x0:0:0
            Originator:10.0.0.2, Cluster list:10.0.0.4
    22
        10.0.0.2 (metric 20) from 10.0.0.5 (10.0.0.5)
            Origin IGP, metric 0, localpref 100, valid, internal, multipath
            Extended Community:0x0:0:0 RT:100:1 0x0:0:0
            Originator:10.0.0.2, Cluster list:10.0.0.5
    22
        10.0.0.2 (metric 20) from 10.0.0.2 (10.0.0.2)
            Origin IGP, metric 0, localpref 100, valid, internal, multipath
            Extended Community:RT:100:1 0x0:0:0
    22
        10.0.0.2 (metric 20) from 10.0.0.3 (10.0.0.3)
            Origin IGP, metric 0, localpref 100, valid, internal, multipath
            Extended Community:0x0:0:0 RT:100:1 0x0:0:0
            Originator:10.0.0.2, Cluster list:10.0.0.3
    22
        10.1.1.12 from 10.1.1.12 (10.22.22.12)
            Origin IGP, metric 0, localpref 100, valid, internal, multipath, best
            Extended Community:RT:100:1
```

Monitoring and Maintaining BGP Multipath Load Sharing for eBGP and iBGP

To display eBGP and iBGP multipath load sharing information, enter any of the following commands in privileged EXEC mode:

Command	Purpose
`Router# show ip bgp all neighbors`	Displays information about the TCP and BGP connections to neighbors.
`Router# show ip bgp vpnv4 all ip-prefix/length`	Displays attributes and multipaths for a network in an MPLS VPN. The *ip-prefix* option is an IP prefix address (in dotted decimal format) and the *length* option is the length of the mask (0 to 32). You must include the slash mark when you use the *length* option.
`Router# show ip bgp vpnv4 all labels`	Displays incoming and outgoing BGP labels for each NLRI prefix.
`Router# show ip bgp vpnv4 rd route-distinguisher`	Displays Network Layer Reachability Information (NLRI) prefixes that have a matching route distinguisher.
`Router# show ip bgp vpnv4 all summary`	Displays BGP neighbor status.
`Router# show ip bgp vpnv4 vrf vrf-name`	Displays NLRI prefixes associated with the named virtual routing and forwarding instance (VRF).
`Router# show ip route vrf vrf-name ip-prefix`	Displays routing information for a network in an MPLS VPN.

IPv6 VPN over MPLS

Multiprotocol BGP is the centerpiece of the MPLS IPv6 VPN architecture in IPv4 and IPv6. It is used to distribute IPv6 routes over the service provider backbone, with the same set of mechanisms to work with overlapping addresses, redistribution policies, and scalability issues. The 6VPE feature is the IOS implementation as described in RFC 4659.

IPv6 avoids overlapping address space by using either a global IPv6 Unicast prefix—RFC 3587—or Unique IPv6 Local Addressing—RFC 4193. For redistribution, a Network Layer Reachability Information (NLRI) 3-tuple, format—containing length, IPv6 prefix, and label—is defined to distribute routes using multiprotocol BGP. The extended community attribute—the route target—is used to control redistribution of routing information by tagging exported routes and filtering imported ones.

For scalability, route reflectors can be used to concentrate routing paths and avoid a full PE mesh. Similar to IPv4, BGP features in IPv6, such as route refresh, automatic route filtering, and outbound route filtering, help reduce the number of routes held in each PE.

Figure 4-1 illustrates the important aspects of the IPv6 VPN architecture.

Figure 4-1 Simple IPv6 VPN Architecture

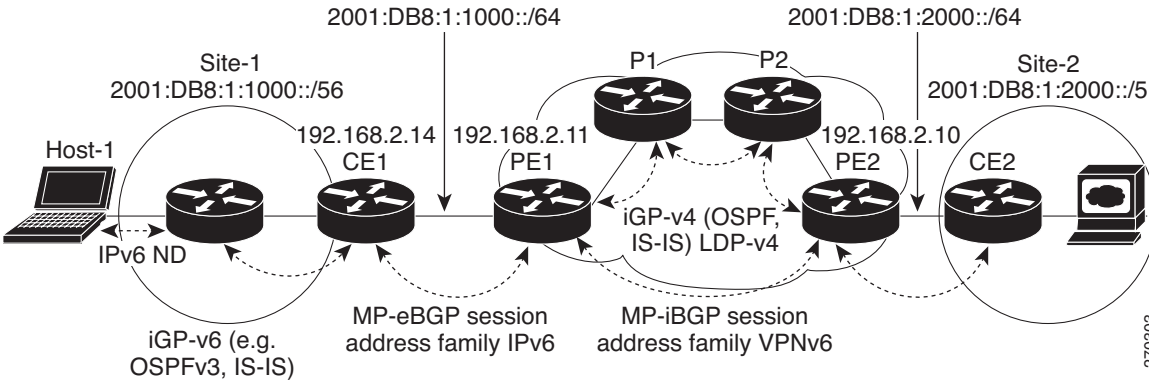

The IPv6 VPN over MPLS (6VPE) feature is described in the following topics:

Feature History for IPv6 VPN over MPLS

Cisco IOS Release	Description	Required PRE
12.2(33)SB	This feature was introduced on Cisco 10000 series routers.	PRE2, PRE3, and PRE4
12.2(33)SB2	This feature supports the inter-AS option on Cisco 10000 series routers.	PRE3 and PRE4

Prerequisites for Implementing IPv6 VPN over MPLS

The following Cisco IOS services must be running on the network before you configure IPv6 VPN operation:

- MPLS in provider backbone routers
- MPLS with VPN code in provider routers with VPN PE routers
- BGP in all routers providing a VPN service
- Cisco Express Forwarding switching in every MPLS-enabled router
- The **ipv6 unicast-routing** command enabled on VPN PE routers

Restrictions for Implementing IPv6 VPN over MPLS

The 6VPE feature has the following restrictions:

- 6VPE is supported by an MPLS IPv4-signaled core. An MPLS IPv6-signaled core is not supported.
- The maximum number of IPv6 VRF's that can be supported is 2038, including the global routing instance. However, out of 2038 VRF's, only 1200 eBGP sessions are supported; the remaining VRF's are to be static routed.
- The maximum number of routes supported across all IPv6 VRFs, including the global routing instance, is 50,000, which yields approximately 24 routes/VRF.

Configuration Tasks for Implementing IPv6 VPN over MPLS

Tip The 12.2(33)SRB release introduced the Implementing IPv6 VPN over MPLS (6VPE) feature. As a basic reference to this feature, see the Implementing IPv6 Addressing and Basic Connectivity chapter in the *Cisco IOS IPv6 Configuration* Library at:
http://www.cisco.com/en/US/docs/ios/12_2t/ipv6/v6addres.html

The IPv6VPN over MPLS (6VPE) includes the configuration tasks in the following list. For more information about these tasks, see the Implementing IPv6 VPN over MPLS (6VPE) chapter in the *Cisco IOS IPv6 Configuration Guide, Release 12.2SR* at:

http://www.cisco.com/en/US/docs/ios/ipv6/configuration/guide/ip6-ov_mpls_6vpe.html

- Configuring a Virtual Routing and Forwarding Instance for IPv6
- Binding a VRF to an Interface
- Configuring a Static Route for PE-to-CE-Routing
- Configuring eBGP PE-to-CE Routing Sessions
- Configuring the IPv6 VPN Address Family for iBGP
- Configuring Route Reflectors for Improved Scalability
- Configuring Internet Access

Note Cisco 10000 series routers do not support the **mpls ipv6 vrf** command that has been listed as one of the steps to configure VRF for IPv6.

The IPv6VPN over MPLS (6VPE) feature also supports the configuration of the following features on Cisco 10000 series routers:

- BGP Features, page 4-10
- IPv6 Internet Access, page 4-11
- VRF-Aware Router Applications, page 4-12
- VRF-Lite, page 4-12
- QoS Features, page 4-12

BGP Features

The following features are supported on Cisco 10000 series routers by the IPv6 VPN over MPLS (6VPE) feature:

- Site of Origin (SoO)

  SoO is used to prevent routing loops in the case of a dual-homed CE. The 6VPE feature supports the SoO Attribute for control of IPv6 VPN routes in the same way as it is currently supported for IPv4 VPNs.

  For information on configuring this feature, see the How to Configure EIGRP MPLS VPN PE-CE Site of Origin (SoO) Support section in the *EIGRP MPLS VPN PE-CE Site of Origin (SoO)* guide at:

  http://www.cisco.com/en/US/docs/ios/12_4/ip_route/configuration/guide/h_mvesoo_ps6350_TSD_Products_Configuration_Guide_Chapter.html#wp1048097

- ASN Override

  If a global ASN is specified, BGP automatically replaces the ASN with a unique number to ensure the site is uniquely identified. The 6VPE feature supports the ASN Override BGP feature via the use of the **as-override** keyword in the same way as the feature is currently supported by IPv4 VPNs.

- Allow-AS-in

  A BGP speaker normally ignores a received update that contains its own ASN in the **AS_PATH** attribute. This check must be omitted in the case of hub-and-spoke topology. The 6VPE feature supports the Allow-AS-in BGP feature via the use of the **allowas-in** keyword in the same way as the feature is currently supported by IPv4 VPNs.

- BGP Prefix List Filtering

  The 6VPE feature supports the ability to filter MP-BGP IPv6 advertisements based on configured IPv6 prefixes.

  For information on configuring this feature, see the Configuring BGP Filtering Using Prefix Lists section in the Configuring BGP chapter of the *Cisco IOS IP Configuration Guide, Release 12.2* guide at:

  http://www.cisco.com/en/US/docs/ios/12_2/ip/configuration/guide/1cfbgp.html#wp1001470

- BGP AS Path Filtering

  The 6VPE feature supports the ability to filter MP-BGP IPv6 advertisement based on configured AS paths.

  For information on configuring this feature, see the Configuring BGP Path Filtering by Neighbor section in the Configuring BGP chapter of the *Cisco IOS IP Configuration Guide, Release 12.2* guide at:

  http://www.cisco.com/en/US/docs/ios/12_2/ip/configuration/guide/1cfbgp.html#wp1001644

- BGP Max Prefix

  The 6VPE feature supports an upper limit on the number of BGP routes that have been learned from a given CE. The 6VPE feature also supports the Max Prefix BGP feature via the use of the **maximum-prefix** keyword in the same way as the feature is currently supported by IPv4 VPNs.

  For information on configuring this feature, see the Configuring BGP section in the Configuring BGP chapter of the *Cisco IOS IP Configuration Guide, Release 12.2* guide at:

  http://www.cisco.com/en/US/docs/ios/12_2/ip/configuration/guide/1cfbgp.html

- BGP Route Refresh

Using BGP Route Refresh, an MP-BGP speaker (PE and/or CE) can request another BGP speaker to resend its MP-BGP updates.

For information on configuring this feature, see the Configuring BGP section in the Configuring BGP chapter of the *Cisco IOS IP Configuration Guide, Release 12.2 Guide* at:

http://www.cisco.com/en/US/docs/ios/12_2/ip/configuration/guide/1cfbgp.html

- Route Target Rewrite at AS Boundary

  The 6VPE feature supports the Route Target Rewrite at AS Boundary feature in the same way as the feature is currently supported by IPv4 VPNs.

  For information on configuring this feature, see the Inter-AS RT-Rewrite section in the Spanning Multiple Autonomous Systems chapter of the *Cisco IP Solution Center MPLS VPN User Guide, 5.0* at:

  http://www.cisco.com/en/US/docs/net_mgmt/ip_solution_center/5.0.1/mpls_vpn/user/guide/multauto.html#wp631364

- BGP Multipath

  The 6VPE feature supports eBGP Multipath, iBGP Multipath, eiBGP Multipath and the DMZ-link-bandwidth based load-balancing for IPv6 VPNs in the VPN-IPv6 address family. Load balancing is supported in the same way as they are currently supported by IPv4 VPNs in the VPN-IPv4 address family.

Note The 6VPE feature does not support per-packet load sharing.

For information on configuring this feature, see the How to Configure BGP Multipath Load Sharing for Both eBGP and iBGP in an MPLS-VPN section in the *BGP Multipath Load Sharing for Both eBGP and iBGP in an MPLS-VPN* guide at:

http://www.cisco.com/en/US/docs/ios/iproute/configuration/guide/irp_bgp_ebgp_ibgp.html#wp1054087

- VRF-aware BGP Dampening

  The 6VPE feature supports the same per-VRF BGP dampening mechanism as the one supported for IPv4 VPN, so that BGP Dampening can be separately controlled for each VRF.

  For information on configuring this feature, see the Configuring Route Dampening section in the Configuring Internal BGP Features chapter of the *Cisco IOS IP Routing Protocols Configuration Guide, Release 12.4* at:

  http://www.cisco.com/en/US/docs/ios/12_4/ip_route/configuration/guide/1cfbgph.html#wp1002395

IPv6 Internet Access

Most VPN sites require access to the Internet. The following Internet access models are supported by IPv6 VPNs for enabling VPN access to the Internet:

- Model 1: Non-VRF Internet Access case.

  In some VPNs, one or more of the sites can obtain Internet access using an Internet gateway, such as a firewall, attached to a non-VRF interface to an Internet service provider (ISP). The ISP may or may not be the same organization as the service provider (SP) that is providing the VPN service. Traffic to or from the Internet gateway can be routed according to the PE router's default forwarding table.

- Model 2: Some VPNs may obtain Internet access via an VRF interface.

  If a packet is received by a PE over a VRF interface and the packet's destination address does not match any route in the VRF, the packet can be matched against the PE's default forwarding table. If the packet matches the PE's default forwarding table, the packet can be forwarded natively through the backbone to the Internet instead of being forwarded by MPLS. In this model, the default forwarding table might have the full set of Internet routes, or it might have just a single default route leading to another router that does have the full set of Internet routes in its default forwarding table.

- Model 3: Using static routes in VRF that can be resolved using the IPv6 Global Table.

  The static routes in VRFs can be resolved in the IPv6 Global Table in the same way they are currently supported for IPv4 VPN. Therefore, in the IPv6 VRF, the network administrator can add static routes as a default route, that points to an IPv6 Internet Gateway for outbound traffic from CE to Internet.

- Model 4: All Internet routes in VRF.

  You can obtain Internet access via a VRF interface by having the VRF include the Internet routes. This model involves redistributing the Internet routes into the VRF.

VRF-Aware Router Applications

The following features are supported on Cisco 10000 series routers by the IPv6VPN over MPLS (6VPE) feature:

- VRF-aware Ping

  The VRF-aware Ping **ping vrf** [*VRF name*] [*IPv6-address*] command is supported.

- VRF-aware Traceroute

  The VRF-aware Traceroute **traceroute vrf** [*VRF name*] [*IPv6-address*] command is supported.

- VRF-aware Telnet

  The VRF-aware Telnet **telnet vrf** [*VRF name*] [*IPv6-address*] command is supported.

VRF-Lite

VRF-lite, also known as Multi-VRF CE, is an extension of IP routing with multiple routing instances on a CE router. The VRF-lite feature performs the following functions:

- Enables the creation of a Layer 3 VPN service by keeping separate IP routing and forwarding tables for each VPN customer.
- Uses input interfaces to distinguish routes for different VPNs.
- Forms virtual packet-forwarding tables by associating one or more interfaces with each VRF. An interface cannot belong to more than one VRF at any time.
- Supports overlapping unicast IP addresses across different VRFs.

VRF-lite is typically deployed along with Multiprotocol Label Switching (MPLS) VPN at the customer edge to support multiple customers on a single switch. The 6VPE feature supports the VRF-Lite feature in the same way as the feature is currently supported by IPv4 VPNs.

QoS Features

The following features are supported on Cisco 10000 series routers by the IPv6VPN over MPLS (6VPE) feature:

- Diff-Serv on Ingress PE

  The 6VPE feature supports the same QoS mechanisms for IPv6 VPNs that is currently supported for IPv4 VPNs on the Ingress PE.

- Diff-Serv on Egress PE

  The 6VPE feature supports the same QoS mechanisms for IPv6 VPNs that is currently supported for IPv4 VPNs on the Egress PE.

- FRF.12

  The 6VPE feature supports FRF.12 fragmentation and interleaving for IPv6 traffic, in addition to IPv4 traffic, on PE to CE Frame Relay connections.

  For configuration tasks, see the FRF.12 Fragmentation section in the Fragmenting and Interleaving Real-Time and Nonreal-Time Packets chapter of the *Cisco 10000 Series Router Quality of Service Configuration Guide* at:

  http://www.cisco.com/en/US/docs/routers/10000/10008/configuration/guides/qos/10qlfi.html#wp1043021

Note When an IPv6 packet arrives on an input interface configured for IPv6, either the packet has a Differentiated Services Code Point (DSCP) value set or an IPv6 QoS setup is done on the router to mark the DSCP value. This packet sent over a MPLS output interface receives the DSCP value that is mapped to the MPLS Experimental (EXP) bits. The mapping propagates the IPv6 QoS value to its MPLS equivalent.

Tip See the *Configuring a Basic MPLS VPN* document for setting up an IPv4 MPLS core network at: http://www.cisco.com/en/US/tech/tk436/tk428/technologies_configuration_example09186a00800a6c11.shtml

Configuration Example for Implementing IPv6 VPN over MPLS

Figure 4-2 illustrates the Example 4-4, which follows the figure.

Figure 4-2 IPv6 VPN over MPLS

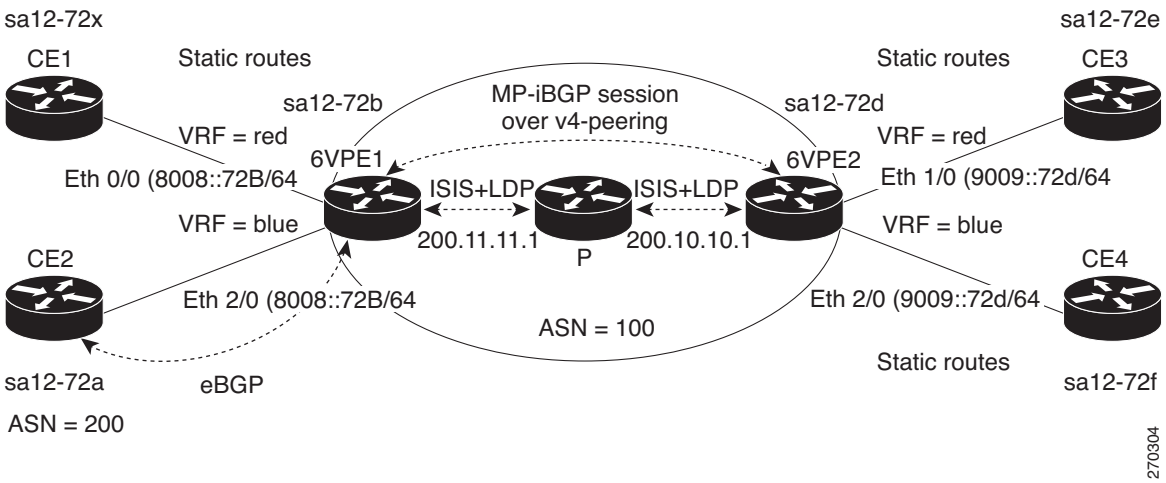

Example 4-4 Configuring IPv6 VPN over MPLS

```
version 12.2
no service pad
service timestamps debug uptime
service timestamps log uptime
no service password-encryption
!
hostname sa14-72b
!
logging snmp-authfail
logging queue-limit 100
!
clock timezone GMT 0
ip subnet-zero
ip cef
!
ipv6 unicast-routing
vrf definition blue
 rd 200:1
 address-family ipv6
 route-target export 200:1
 route-target import 200:1
 exit-address-family
!
vrf definition red
 rd 100:1
 address-family ipv6
 route-target export 100:1
 route-target import 100:1
 exit-address-family
!
ipv6 cef
mpls ldp logging neighbor-changes
mpls ldp router-id Loopback0
!
!
interface Loopback0
 ip address 200.11.11.1 255.255.255.255
 ipv6 address BEEF:11::1/64
 ipv6 nd prefix default 0 0 off-link no-autoconfig
 no ipv6 mfib fast
!
interface Ethernet0/0
 vrf forwarding red
 ip address 50.1.1.2 255.255.255.0
 no ip route-cache
 no ip mroute-cache
 ipv6 address 4000::72B/64
 ipv6 address 8008::72B/64
 ipv6 nd prefix default infinite infinite
 no ipv6 mfib fast
!
interface Ethernet1/0
 ip address 40.1.1.2 255.255.255.0
 ip router isis
 no ip mroute-cache
 mpls ip
!
interface Ethernet2/0
 vrf forwarding blue
 ip address 90.1.1.2 255.255.255.0
 ipv6 address 8008::72B/64
 no ipv6 mfib fast
```

```
!
router isis
 net 49.0000.0000.0002.00
 redistribute connected metric 50
 passive-interface Loopback0
!
router bgp 100
 bgp log-neighbor-changes
 neighbor 200.10.10.1 remote-as 100
 neighbor 200.10.10.1 update-source Loopback0
 neighbor 8008::72a remote-as 200
 !
 address-family ipv4
 neighbor 200.10.10.1 activate
 no auto-summary
 no synchronization
 exit-address-family
 !
 address-family ipv4 multicast
 no auto-summary
 exit-address-family
 !
 address-family vpnv6
 neighbor 200.10.10.1 activate
 neighbor 200.10.10.1 send-community extended
 exit-address-family
 !
 address-family ipv6 vrf red
 no synchronization
 redistribute connected
 exit-address-family
 !
 address-family ipv6 vrf blue
 neighbor 8008::72a activate
 no synchronization
 redistribute connected
 exit-address-family
!
ip classless
no ip http server
!
end
```

Monitoring and Maintaining IPv6 VPN over MPLS

For information on monitoring and maintaining IPv6 VPN over MPLS, see the Verifying and Troubleshooting IPv6 VPN section in the Implementing IPv6 VPN over MPLS (6VPE) chapter of the *Cisco IOS IPv6 Configuration Guide, Release 12.4T* guide at:

http://www.cisco.com/en/US/docs/ios/12_2t/ipv6/SA_vpnv6_ps6441_TSD_Products_Configuration_Guide_Chapter.html#wp1078529

Session Limit Per VRF

The session limit Per VRF feature enables you to limit the number of sessions that can be established for VPDN groups associated with a specific VPDN template. Previously, you associated all VPDN groups configured on the router with a single template. By using the session limit Per VRF feature, you can create, define, and name multiple VPDN templates, including a default VPDN template. You can

then associate a VPDN group with a specific VPDN template. By configuring a group session limit for a VPDN template, you can limit the maximum number of concurrent sessions allowed for all VPDN groups associated with the VPDN template.

If you configure a group session limit for the default VPDN template (the unnamed VPDN template), that session limit is the same for all VPDN groups not associated with a named VPDN template. If you configure a group session limit that is less than the number of current active sessions, no sessions are terminated and no new sessions can start. For example, if you configure a group session limit of 30 and 50 sessions are active, the router does not terminate any active sessions and it does not allow any new sessions to start.

If a VPDN group is associated with a VPDN template and the VPDN group has a session limit configured, the value of the VPDN group session limit takes precedence over the VPDN template session limit if the VPDN group value is less than the VPDN template value.

You can associate a VPDN group with only one VPDN template at a time. If you associate a VPDN group with a named VPDN template and then with a second VPDN template, the VPDN group is detached from the first VPDN template and associated with the second.

If you attempt to associate a VPDN group with a named VPDN template that you have not configured, the VPDN group uses the system defaults.

The **session-limit** global configuration command takes precedence over the **group session-limit** VPDN template configuration command. The **session-limit** command limits the number of VPDN sessions and the **group session-limit** command specifies the maximum concurrent sessions allowed across all VPDN groups associated with a particular VPDN template.

The session limit Per VRF feature is described in the following topics:

Application of VPDN Parameters to VPDN Groups

By default, the router applies VPDN parameters to a VPDN group in the following way:

- VPDN parameters configured for the individual VPDN group are always applied to that VPDN group.
- VPDN parameters configured in the VPDN template are applied for any setting not specified in the individual VPDN group configuration.
- System default settings for VPDN parameters are applied for any setting not configured in the individual VPDN group or VPDN template.

When you detach a VPDN group from a VPDN template by using the **no source vpdn-template** command, the router applies VPDN parameters to that VPDN group in the following way:

- VPDN parameters configured for the individual VPDN group are always applied to that VPDN group.
- System default settings for VPDN parameters are applied for any settings not configured in the individual VPDN group or VPDN template.

VPDN Template Configuration

Not all commands that are available for configuring a VPDN group can be used to configure a VPDN template. For a list of the commands available for VPDN template configuration, see the **vpdn-template** command reference page in the *Session Limit Per VRF*, Release 12.2(4)B feature module.

Feature History for Session Limit Per VRF

Cisco IOS Release	Description	Required PRE
12.2(15)BX	This feature was integrated into Cisco IOS Release 12.2(15)BX.	PRE2
12.3(7)XI1	This feature was integrated into Cisco IOS Release 12.3(7)XI1.	PRE2
12.2(28)SB	This feature was integrated into Cisco IOS Release 12.2(28)SB.	PRE2

Restrictions for Session Limit Per VRF

The session limit Per VRF feature has the following restrictions:

- Nesting of VPDN templates is not supported. You can associate a VPDN group with only one VPDN template at a time. If you associate a VPDN group with a named VPDN template and then with a second VPDN template, the VPDN group is detached from the first VPDN template and associated with the second template.
- If you attempt to associate a VPDN group with a named VPDN template that you have not configured, the VPDN group uses the system defaults.
- The **session-limit** global configuration command takes precedence over the **group session-limit** VPDN template configuration command.
- If the VPDN group value is less than the VPDN template value, the VPDN group session limit takes precedence over the VPDN template session limit.

Prerequisites for Session Limit Per VRF

The session limit Per VRF feature has the following requirements:

- You must have a VPDN enabled on the router and at least one VPDN group configured. The router must make an L2TP connection before VPDN configurations can be established.

Configuring Session Limit Per VRF

To configure the session limit Per VRF feature on the Cisco 10000 series router, enter the following commands beginning in global configuration mode:

	Command	Purpose
Step 1	`Router(config)# vpdn enable`	Enables virtual private dialup networking (VPDN) on the router and informs the router to look for tunnel definitions in a local database and on a remote authorization server if one is present.
Step 2	`Router(config)# vpdn session-limit` *sessions*	Limits the number of simultaneous VPN sessions that can be established on a router. The *sessions* option is the maximum number of simultaneous VPN sessions that you want to allow on a router. Valid values are 1 to 10,000.
Step 3	`Router(config)# vpdn-template` *template-name*	Configures a VPDN template and enters VPDN group configuration mode. The *template-name* option is the name of the VPDN template
Step 4	`Router(config-vpdn)# group session-limit` *number*	Specifies the maximum concurrent sessions allowed across all VPDN groups associated with the VPDN template you specified in step 3. The *number* option is a value from 1 to 32,767.
Step 5	Repeat steps 2 and 3 to configure additional named VPDN templates.	
Step 6	`Router(config-vpdn)# exit`	Exits VPDN group configuration mode.
Step 7	`Router(config)# vpdn-group` *tag*	Associates a VPDN group to a customer or VPDN profile. The *tag* option is the name of the VPDN group.

Step 8	`Router(config-vpdn)# accept-dialin`
	or
	`Router(config-vpdn)# request-dialout`
		Enables the router to accept dialin requests and enters VPDN accept-dialin group configuration mode.
		Enables the router to send L2TP dialout requests and enters VPDN request-dialout group configuration mode.

Step 9	`Router(config-vpdn-acc-in)# protocol` *protocol* or `Router(config-vpdn-req-out)# protocol` *protocol*	Specifies the tunneling protocol to be used.

Step 10	`Router(config-vpdn-acc-in)# exit`
	or
	`Router(config-vpdn-req-out)# exit`
		Exits VPDN accept-dialin or VPDN request-dialout group configuration mode.

Step 11	`Router(config-vpdn)# source vpdn-template` *template-name*	Configures the VPDN group to use the VPDN template settings for all unspecified parameters. The *template-name* option is the name of the VPDN template to be associated with a VPDN group.

	Command	Purpose
Step 12	`Router(config-vpdn)# session-limit session-number`	Limits the number of sessions allowed on the VPDN group. The *session-number* option is the maximum number of sessions allowed on the specified VPDN group. Valid values are from 0 to 32,767.
Step 13	Repeat steps 7 through 12 to configure session limiting on additional VPDN groups.	

Verifying a Session Limit Per VRF Configuration

To verify the configuration of the session limit Per VRF feature, enter the following commands in privileged EXEC mode:

Command	Purpose
`Router# show running-config`	Displays the current configuration of the router. Check the output of this command to confirm the configuration of a VPDN template group.
`Router# show vpdn session`	Displays the status of all active tunnels.

Configuration Examples for Session Limit Per VRF

Example 4-5 creates three VPDN groups named *group1*, *group2*, and *group3*. VPDN group1 and group2 are attached to the default VPDN template, which has a session limit of 10. VPDN group1 and group2 can have no more than a combined total of 10 concurrent sessions. For example, if group1 has three sessions, group2 can only have seven sessions.

In Example 4-5, using the **session-limit 5** command allows VPDN group1 to have no more than 5 sessions. Using the **session-limit 20** command allows VPDN group2 to have no more than 20 sessions. However, as previously indicated, the default VPDN template has a session limit of 10. Therefore, the combined number of sessions for VPDN group1 and group2 cannot exceed 10 sessions. If group1 has 5 sessions, group2 can only have 5 sessions. If group1 does not have any active sessions, group2 can have a maximum of 10 sessions, even though group2 is configured with the **session-limit 20** command.

In Example 4-5, VPDN group3 does not have a session limit configured. Using the **no source vpdn-template** command detaches group3 from the default VPDN template.

Example 4-5 Configuring Session Limit Per VRF

```
vpdn-template
 group session-limit 10
 exit

vpdn-group group2
 accept-dialin
  protocol any
  exit
 session-limit 20
 exit

vpdn-group group1
 accept-dialin
  protocol any
```

```
  exit
 session-limit 5

vpdn-group group3
 accept-dialin
  protocol any
  exit
 no source vpdn-template
```

Example 4-6 creates a default VPDN template and three VPDN groups named groupA, groupB, and groupC. As indicated in the default VPDN template configuration, the maximum combined number of sessions allowed for all VPDN groups associated with the default template is 10 sessions. The local name of the default VPDN template is local-name. Example 4-6 also creates an additional VPDN template named templateA, which has a session limit of 50. The combined number of sessions for all VPDN groups associated with templateA cannot exceed 50 concurrent sessions, regardless of the session limits set for the individual VPDN groups. VPDN groupA and groupB are attached to VPDN templateA and each group has an individual session limit of 30 sessions. Because groupA and groupB are attached to VPDN templateA, they use the hostname host1 as their local name.

In Example 4-6, the **source vpdn-template** command is not used to associate VPDN groupC with a specific VPDN template. Therefore, by default, VPDN groupC is attached to the default VPDN template, which has a group session limit of 10. VPDN groupC inherits the local name local-name from the default VPDN template.

Example 4-6 Configuring Session Limit Per VRF

```
hostname host1
vpdn-template
 group session-limit 10
 local name local-name
 exit

vpdn-template templateA
 group session-limit 50
 exit

vpdn-group groupA
 accept-dialin
  protocol any
  exit
 source vpdn-template templateA
 session-limit 30
 exit

vpdn-group groupB
 accept-dialin
  protocol any
  exit
 source vpdn-template templateA
 session-limit 30
 exit

vpdn-group groupC
 accept-dialin
  protocol any
```

Monitoring and Maintaining Session Limit Per VRF

To monitor and maintain the session limit Per VRF feature, enter the following commands in privileged EXEC mode:

Command	Purpose
`Router# show vpdn session [all [interface | tunnel | username] | packets | sequence | state | timers | window]`	Displays VPDN session information including interface, tunnel, username, packets, status, and window statistics. The options are: • **all**—All session information for active sessions • **all interface**—Interface associated to a specific session • **all tunnel**—Tunnel attribute filter • **all username**—Username filter • **packets**—Packet and byte count • **sequence**—Sequence numbers • **state**—State of each session • **timers**—Timer information • **window**—Window information
`Router# show vpdn`	Displays a summary of all active VPDN tunnels.
`Router# show vpdn group name`	Displays the session limit set and the number of active sessions and tunnels on the VPDN group you specify.
`Router# show vpdn history failure`	Displays information about VPDN user failures.

Half-Duplex VRF

The Half-Duplex VRF (HDVRF) feature provides scalable hub and spoke connectivity for subscribers of a multiprotocol label switching-based virtual private network (MPLS VPN) service. These subscribers connect to the provider edge (PE) router of the wholesale service provider, and they use the same or different services (for example, the same or different VRFs). The HDVRF feature prevents local connectivity between subscribers at the spoke PE router and ensures that a hub site provides subscriber connectivity. Any sites that connect to the same PE router must forward intersite traffic using the hub site. This ensures that the routing done at the spoke site is always access side interface to network side interface, or network side interface to access side interface, and never access side to access side.

In hub and spoke topologies in which multiple-spoke customer edge (CE) routers, also referred to as spokes, connect to the same PE router, the PE router locally switches the spokes without passing the traffic through the upstream Internet service provider (ISP). In releases earlier than Cisco IOS Release 12.2(16)BX2, when spokes connect to the same PE router, it was necessary to configure each spoke in a separate VRF to ensure that the traffic between the spokes always traverses the central link between the wholesale service provider and the ISP. However, this solution is manageable only if the number of spokes is relatively small. When a large number of spokes are connected to the same PE router, configuring a single VRF for each spoke can become quite complex and can greatly increase memory usage. This is true especially in large-scale wholesale service provider environments that support high-density remote access to Layer 3 VPNs.

The HDVRF feature addresses the limitations previously imposed on hub and spoke topologies by removing the requirement of one VRF per spoke and ensuring that subscriber traffic always traverses the central link between the wholesale service provider and the ISP, whether the subscriber traffic is being routed to a remote network by way of the upstream ISP or to another locally or remotely connected subscriber.

Figure 4-3 shows a sample hub and spoke topology for HDVRF.

Figure 4-3 Hub and Spoke Topology for Half-Duplex VRF

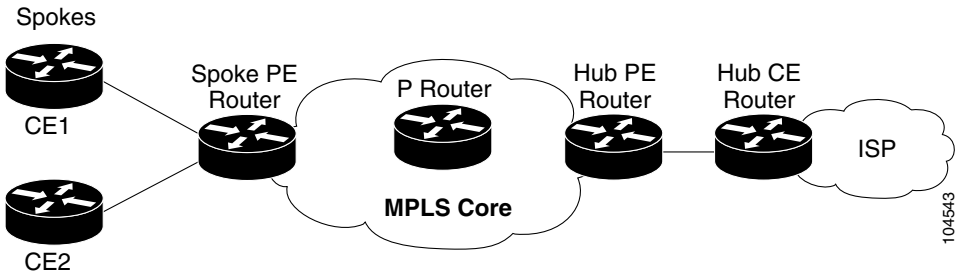

The Half-Duplex VRF feature is described in the following topics:

Upstream and Downstream VRFs

HDVRF uses two unidirectional VRFs, called upstream VRF and downstream VRF, to forward IP traffic between the spokes and the hub PE router.

The upstream VRF is used to forward the IP traffic from the spokes toward the MPLS VPN backbone. This VRF typically contains only a default route; but, depending on the configuration, it might also contain such information as summary routes and multiple default routes. The default route points to the interface on the hub PE router that connects to the upstream ISP. The Cisco 10000 series router dynamically learns about the default route from the routing updates that the hub PE router or home gateway sends. The upstream VRF also contains the virtual access interfaces that connect the spokes, but it contains no other local interfaces.

The downstream VRF is used to forward the traffic from the MPLS core back to the spokes. This VRF contains PPP peer routes for the spokes and per-user static routes imported from the authentication, authorization, and accounting (AAA) server. It also contains the routes imported from the hub PE router. These routes are the dynamically allocated virtual access interfaces of the subscribers associated with a particular service.

The Cisco 10000 series router redistributes routes from the downstream VRF into Multiprotocol Border Gateway Protocol (MP-BGP). The spoke PE router typically advertises a summary route across the MPLS core for the connected spokes. The upstream VRF configured on the hub PE router imports the advertised summary route.

Reverse Path Forwarding Check Support

Reverse Path Forwarding (RPF) check ensures that an IP packet entered the router using the correct inbound interface. The HDVRF feature supports unicast RPF check on the spoke-side interfaces. Because different VRFs are used for downstream and upstream forwarding, HDVRF extends the RPF mechanism to ensure that source address checks occur in the downstream VRF.

Feature History for Half-Duplex VRF

Cisco IOS Release	Description	Required PRE
12.2(16)BX2	This feature was introduced on the Cisco 10000 series router.	PRE2
12.3(7)XI1	This feature was integrated into Cisco IOS Release 12.3(7)XI1.	PRE2
12.2(28)SB	This feature was integrated into Cisco IOS Release 12.2(28)SB.	PRE2

Restrictions for Half-Duplex VRF

The Half-Duplex VRF feature has the following restrictions:

- In both the upstream and downstream VRFs, routing protocols are not supported on interfaces configured for half-duplex VRFs.
- Half-duplex VRFs apply only to virtual access interfaces (VAIs) and virtual template interfaces. Only IP unnumbered interfaces are supported.
- It is not supported with Routing with Bridged Encapsulation (RBE).

Prerequisites for Half-Duplex VRF

The Half-Duplex VRF feature has the following requirements:

- The spoke PE routers must be running Cisco IOS Release 12.2(16)BX2, Cisco IOS Release 12.3(7)XI1, or a later release.
- The performance routing engine (PRE), part number ESR-PRE2, must be installed in the router's chassis.

Configuration Tasks for Half-Duplex VRF

To configure the Half-Duplex VRF feature, perform the following configuration tasks:

- Configuring Upstream and Downstream VRFs on the L2TP Access Concentrator and PE Router, page 4-24
- Associating VRFs, page 4-25
- Configuring RADIUS, page 4-26

Configuring Upstream and Downstream VRFs on the L2TP Access Concentrator and PE Router

To configure the upstream and downstream VRFs on the PE router, enter the following commands beginning in global configuration mode:

	Command	Purpose
Step 1	`Router(config)# ip vrf vrf-name`	Enters VRF configuration mode and defines the VRF instance by assigning a VRF name.
Step 2	`Router(config-vrf)# rd route-distinguisher`	Creates routing and forwarding tables.
Step 3	`Router(config-vrf)# route-target {import | export | both} route-target-ext-community`	Creates a list of import and export route target communities for the specified VRF. The **import** keyword is required to create an upstream VRF. The upstream VRF is used to import the default route from the hub PE router. The **export** keyword is required to create a downstream VRF. The downstream VRF is used to export the routes of all subscribers of a given service that the VRF serves.

Example 4-7 shows how to configure a downstream VRF named D.

Example 4-7 Configuring the Downstream VRF

```
Router(config)# ip vrf D
Router(config-vrf)# description Downstream VRF - to subscribers
Router(config-vrf)# rd 1:8
Router(config-vrf)# route-target export 1:100
```

Example 4-8 shows how to configure an upstream VRF named U.

Example 4-8 Configuring the Upstream VRF

```
Router(config)# ip vrf U
Router(config-vrf)# description Upstream VRF - to hub PE
Router(config-vrf)# rd 1:0
Router(config-vrf)# route-target import 1:0
```

Associating VRFs

After you define and configure the VRFs on the PE routers, associate each VRF with:

- An interface or subinterface, or
- A virtual template interface

The virtual template interface is used to create and configure a virtual access interface (VAI). For information about configuring a virtual template interface, see the "Configuring a Virtual Template Interface" section on page 3-17.

To associate a VRF, enter the following commands on the PE router beginning in interface configuration mode:

	Command	Purpose
Step 1	`Router(config-if)# ip vrf forwarding vrf-name`	Associates an interface with the VRF you specify. *vrf-name* is the name of the VRF associated with the interface.
Step 2	`Router(config-if)# ip unnumbered type number`	Enables IP processing on an interface without assigning an explicit IP address to the interface. The *type* and *number* arguments are the type and number of another interface on which the router has an assigned IP address. It cannot be another unnumbered interface. **Note** The Cisco 10000 series router supports only unnumbered interfaces for the Half-Duplex VRF feature.
Step 3	`Router(config-if)# exit`	Returns to global configuration mode.
Step 4	`Router(config)# interface virtual-template number`	Creates a virtual template interface and enters interface configuration mode.
Step 5	`Router(config-if)# ip vrf forwarding vrf-name1 [downstream vrf-name2]`	Associates a virtual template interface with the VRF you specify. The *vrf-name1* argument is the name of the VRF associated with the virtual template interface. The *vrf-name2* argument is the name of the downstream VRF into which the PPP peer route and all of the per-user routes from the AAA server are installed. If a AAA server is used, the AAA server provides the VRF membership; you do not need to configure the VRF members on the virtual templates.

Example 4-9 associates the VRF named vpn1 with the Virtual-Template1 interface and specifies the downstream VRF named D.

Example 4-9 Associating a Downstream VRF with a Virtual Template Interface

```
Router(config)# interface Virtual-Template1
Router(config-if)# ip vrf forwarding vpn1 downstream D
Router(config-if)# ip unnumbered Loopback1
Router(config-if)# no peer default ip address
Router(config-if)# ppp authentication chap vpn1
Router(config-if)# ppp authorization vpn1
Router(config-if)# ppp accounting vpn1
```

Configuring RADIUS

To configure the downstream VRF for an AAA server, enter the following Cisco attribute value:

```
cisco-avpair = "ip:vrf-id=vrf-name1 downstream vrf-name2"
```

where:

The *vrf-name1* argument is the name of the VRF associated with the subinterface or virtual template interface.

The *vrf-name2* argument is the name of the downstream VRF into which all of the subscriber routes from the AAA server are installed.

Note Instead of using the **lcp:interface-config** RADIUS attribute, we recommend that you use the **ip:vrf-id** RADIUS attribute when supported in Cisco IOS software. Unlike the **lcp:interface-config** attribute, which causes full virtual interfaces to be used, the **ip:vrf-id** attribute causes virtual subinterfaces to be used, which significantly improves scalability.

Example 4-10 shows how to configure a downstream VRF named D on a AAA server:

Example 4-10 Configuring the Downstream VRF on RADIUS

```
cisco-avpair = "ip:vrf-id=U downstream D"
```

Configuration Examples for Half-Duplex VRF

This section provides the following configuration examples. These examples use the hub and spoke topology shown in Figure 4-4.

- Hub and Spoke Sample Configuration with Half-Duplex VRFs, page 4-27
- RADIUS Sample Configuration, page 4-28

Figure 4-4 Sample Topology for Half-Duplex Configuration

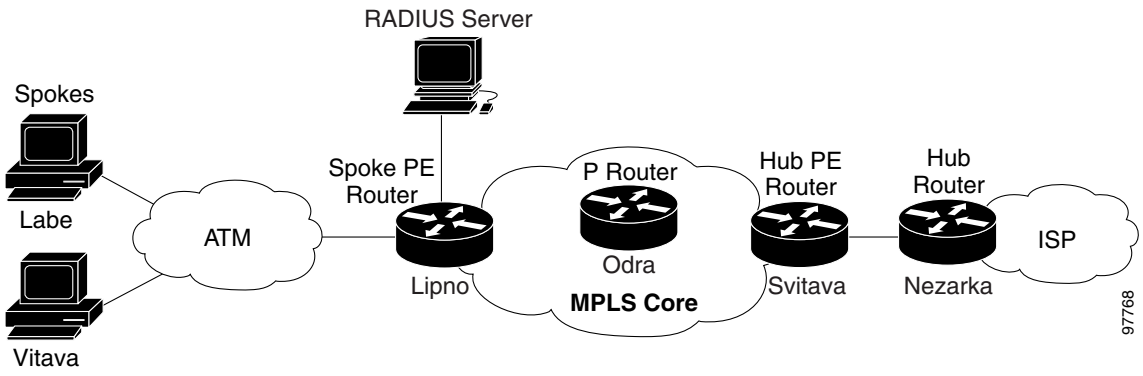

Hub and Spoke Sample Configuration with Half-Duplex VRFs

Example 4-11 shows how to connect two PPPoE clients to a single VRF pair on the spoke PE router named *Lipno*. Although both PPPoE clients are configured in the same VRF, all communication occurs using the hub PE router. Half-duplex VRFs are configured on the spoke PE. The client configuration is downloaded to the spoke PE from the RADIUS server.

Example 4-11 Configuring the Spoke PE Router

```
aaa new-model
!
aaa group server radius R
 server 22.0.20.26 auth-port 1812 acct-port 1813
!
aaa authentication ppp default group radius
aaa authorization network default group radius
!
ip vrf D
 description Downstream VRF - to spokes
 rd 1:8
 route-target export 1:100
!
ip vrf U
 description Upstream VRF - to hub
 rd 1:0
 route-target import 1:0
!
vpdn enable
!
vpdn-group U
 accept-dialin
  protocol pppoe
  virtual-template 1
!
interface Loopback0
ip address 100.0.0.8 255.255.255.255
!
interface Loopback2
ip unnumbered Loopback2
 ip vrf forwarding U
 ip address 2.0.0.8 255.255.255.255
!
interface ATM2/0
 pvc 3/100
  protocol pppoe
 !
 pvc 3/101
  protocol pppoe
 !
interface Virtual-Template1
 no ip address
 ppp authentication chap
!
router bgp 1
 no synchronization
 neighbor 100.0.0.34 remote-as 1
 neighbor 100.0.0.34 update-source Loopback0
 no auto-summary
 !
address-family vpnv4
  neighbor 100.0.0.34 activate
  neighbor 100.0.0.34 send-community extended
```

```
  no auto-summary
  exit-address-family
 !
 address-family ipv4 vrf U
  no auto-summary
  no synchronization
  exit-address-family
 !
 address-family ipv4 vrf D
  redistribute static
  no auto-summary
  no synchronization
  exit-address-family
!
ip local pool U-pool 2.8.1.1 2.8.1.100
!
radius-server host 22.0.20.26 auth-port 1812 acct-port 1813
radius-server key cisco
```

RADIUS Sample Configuration

Example 4-12 shows how to configure the RADIUS server for HDVRF support. In this example, the spokes inherit the default configuration. Static routes per spoke are defined to demonstrate that HDVRF supports per-user static routes. The functionality of the HDVRF feature does not require that you define static routes per spoke. This configuration was tested on FreeRADIUS 0.8.1.

Example 4-12 Configuring RADIUS for Half-Duplex VRFs

```
DEFAULT Service-Type == Framed-User
        Framed-Protocol = PPP,
        cisco-avpair = "ip:vrf-id=U downstream D",
        cisco-avpair = "ip:ip-unnumbered=Loopback 2",
        cisco-avpair = "ip:addr-pool=U-pool",
        Fall-Through = Yes

labe    Auth-Type := Local, User-Password == "labe"
        cisco-avpair = "ip:route=2.0.0.5 255.255.255.255"

vltava  Auth-Type := Local, User-Password == "vltava"
        cisco-avpair = "ip:route=2.0.0.2 255.255.255.255"
```

Note Instead of using the **lcp:interface-config** RADIUS attribute, we recommend that you use the **ip:vrf-id** RADIUS attribute when supported in Cisco IOS software. Unlike the **lcp:interface-config** attribute, which causes full virtual interfaces to be used, the **ip:vrf-id** attribute causes virtual subinterfaces to be used, which significantly improves scalability.

Monitoring and Maintaining Half-Duplex VRF

To monitor and maintain upstream and downstream VRFs, enter any of the following commands in privileged EXEC mode:

Command	Purpose
`Router# show cef interface virtual-interface number internal`	Displays internal information about the virtual access interface (VAI) you specify, including the downstream VRF associated with the VAI.
`Router# show ip interface virtual-interface number`	Displays information about the VAI you specify, including the downstream VRF associated with the VAI.
`Router# show ip route vrf vrf-name`	Displays the IP routing table for the VRF you specify. Use this command to display information about the per-user static routes installed in the downstream VRF.
`Router# show ip vrf`	Displays information about all of the VRFs configured on the router, including the downstream VRF for each associated VAI.
`Router# show ip vrf detail vrf-name`	Displays detailed information about the VRF you specify, including all of the VAIs associated with the VRF. If you do not specify a value for *vrf-name*, detailed information about all of the VRFs configured on the router appears, including all of the VAIs associated with each VRF.
`Router# show running-config interface type number`	Displays information about the virtual access interface you specify, including information about the upstream and downstream VRFs.

Example 4-13 shows how to display information about the interface named virtual-access 3.

Example 4-13 show running-config interface—virtual-access 3

```
Lipno# show running-config interface virtual-access 3

Building configuration...

Current configuration : 92 bytes
!
interface Virtual-Access3
 ip vrf forwarding U downstream D
 ip unnumbered Loopback2
end
```

Example 4-14 shows how to display information about the interface named virtual-access 4.

Example 4-14 show running-config interface—virtual-access 4

```
Lipno# show running-config interface virtual-access 4

Building configuration...

Current configuration : 92 bytes
!
interface Virtual-Access4
 ip vrf forwarding U downstream D
 ip unnumbered Loopback2
end
```

Example 4-15 shows how to display the routing table for the downstream VRF named D.

Example 4-15 show ip route vrf—Downstream

```
Lipno# show ip route vrf D

Routing Table: D
Codes: C - connected, S - static, R - RIP, M - mobile, B - BGP
       D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area
       N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
       E1 - OSPF external type 1, E2 - OSPF external type 2
       i - IS-IS, L1 - IS-IS level-1, L2 - IS-IS level-2, ia - IS-IS inter
area
       * - candidate default, U - per-user static route, o - ODR
       P - periodic downloaded static route

Gateway of last resort is not set

     2.0.0.0/8 is variably subnetted, 5 subnets, 2 masks
U       2.0.0.2/32 [1/0] via 2.8.1.1
S       2.0.0.0/8 is directly connected, Null0
U       2.0.0.5/32 [1/0] via 2.8.1.2
C       2.8.1.2/32 is directly connected, Virtual-Access4
C       2.8.1.1/32 is directly connected, Virtual-Access3
```

Example 4-16 shows how to display the routing table for the upstream VRF named U.

Example 4-16 show ip route vrf—Upstream

```
Lipno# show ip route vrf U


Routing Table: U
Codes: C - connected, S - static, R - RIP, M - mobile, B - BGP
       D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area
       N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
       E1 - OSPF external type 1, E2 - OSPF external type 2
       i - IS-IS, L1 - IS-IS level-1, L2 - IS-IS level-2, ia - IS-IS interarea
       * - candidate default, U - per-user static route, o - ODR
       P - periodic downloaded static route

Gateway of last resort is 100.0.0.20 to network 0.0.0.0

     2.0.0.0/32 is subnetted, 1 subnets
C       2.0.0.8 is directly connected, Loopback2
B*   0.0.0.0/0 [200/0] via 100.0.0.20, 1w5d
```