



BGP MPLS VPNs

This chapter describes services that are supported for Border Gateway Protocol (BGP) Multi-Protocol Label Switching (MPLS) Virtual Private Networks (VPNs).



Important

MPLS is a licensed Cisco feature that requires a separate license. Contact your Cisco account representative for detailed information on specific licensing requirements. For information on installing and verifying licenses, refer to the *Managing License Keys* section of *Software Management Operations*.

It includes the following topics:

- [Introduction, page 1](#)
- [MPLS-CE Connected to PE, page 2](#)
- [VPC-DI as a PE, page 3](#)
- [IPv6 Support for BGP MPLS VPNs, page 5](#)
- [VPN-Related CLI Commands, page 8](#)

Introduction

Service providers require the ability to support a large number of corporate Access Point Names (APNs) which have a number of different addressing models and requirements. VPC-DI uses BGP MPLS Layer 3 VPNs to segregate corporate customer APNs in a highly scalable manner. This solution conforms to RFC 4364 – *BGP/MPLS IP Virtual Private Networks (VPNs)*.

The BGP/MPLS solution supports the following scenarios:


- [MPLS-CE Connected to PE, on page 2](#)
- [VPC-DI as a PE, on page 3](#)

VPC-DI also supports VPNv6 as described in RFC 4659 – *BGP-MPLS IP Virtual Private Network (VPN) Extension for IPv6 VPN*. See [IPv6 Support for BGP MPLS VPNs, on page 5](#) for details.

MPLS-CE Connected to PE

In this scenario the VPC-DI functions as an MPLS-CE (Customer Edge) network element connected to a Provider Edge (PE) Label Edge Router (LER), which in turn connects to the MPLS core (RFC 4364). See the figure below.

Figure 1: VPC-DI MPLS-CE to PE



The MPLS-CE functions like a PE router within its own Autonomous System (AS). It maintains Virtual Routing and Forwarding (VRF) routes and exchanges VPN route information with the PE via an MP-eBGP (Multi-Protocol-external BGP) session.

The PE is also configured with VRFs and exchanges VPN routes with other PEs in its AS via MP-iBGP (Multi-Protocol-internal BGP) connections and the MPLS-CE via an MP-eBGP connection.


The EBGP connection allows the PE to change next-hop IP addresses and labels in the routes learned from IBGP peers before advertising them to the MPLS-CE. The MPLS-CE in this case uses only MP-eBGP to advertise and learn routes. Label Distribution Protocol (LDP) and Resource Reservation Protocol (RSVP) are not required because of direct-connect EBGP peering. The MPLS-CE in this scenario pushes/pops a single label (learned over the MP-eBGP connection) to/from the PE.

VPC-DI as a PE

Overview

In this scenario, the VPC-DI functions as a PE router sitting at the edge of the MPLS core. See the figure below.

Figure 2: VPC-DI as a PE



The VPC-DI eliminates the need for an ASBR or PE as shown in the first two scenarios. In this scenario, two main requirements are introduced: IBGP functionality and MPLS label distribution protocols.

The VPC-DI can be configured to add two labels:

- an outer label learned from LDP or RSVP-TE (RSVP-Traffic Engineering)
- an inner label learned from MP-iBGP


This solution supports traffic engineering and QoS initiated via the VPC-DI.

Sample Configuration

In this example, VRFs are configured on the ASR 5500 PE and pools are associated with VRFs. The VPC-DI exchanges VPN routes with its IBGP peers (PE routers) and learns the MPLS paths to reach PEs via LDP.

The VPC-DI forwards the packets to the next-hop with two labels – an inner label learned from PE and an outer label learned from the next hop IBGP neighbor.

Figure 3: Sample Configuration



```

mpls ip
  protocol ldp
    enable
  exit
exit

ip vrf vrf1
  mpls traffic-class copy
exit
ip vrf vrf2
  mpls traffic-class value 5
exit

router bgp 300
  ip vrf vrf1
    route-target export 300 1
    route-target import 300 1
    route-distinguisher 300 1
  exit
  ip vrf vrf2
    route-target export 300 2
    route-target import 300 2
    route-distinguisher 300 2
  exit

  router-id 2.2.2.2
  neighbor 192.168.107.20 remote-as 300
  neighbor 192.168.107.20 update-source node1_loopback

  address-family vpnv4
    neighbor 192.168.107.20 activate
    neighbor 192.168.107.20 send-community both
    neighbor 192.168.107.20 next-hop-self
  exit

  address-family ipv4 vrf vrf1
    redistribute connected
  exit

  address-family ipv4 vrf vrf2
    redistribute connected
  exit

interface interface to internet
  ip address 192.168.109.65/24
  mpls ip
exit
router ospf

```

```

network 192.168.109.0/24 area 0.0.0.0
exit

```

IPv6 Support for BGP MPLS VPNs

Overview

The VPC-DI supports VPNv6 as described in RFC 4659 – *BGP-MPLS IP Virtual Private Network (VPN) Extension for IPv6 VPN*.

An IPv6 VPN is connected over an IPv6 interface or sub-interface to the Service Provider (SP) backbone via a PE router. The site can be both IPv4 and IPv6 capable. Each VPNv6 has its own address space which means a given address denotes different systems in different VPNs. This is achieved via a VPNv6 address-family which prepends a Route Distinguisher (RD) to the IP address.

A VPNv6 address is a 24-byte quantity beginning with an 8-byte RD and ending with a 16-byte IPv6 address. When a site is IPv4 and IPv6 capable, the same RD can be used for the advertisement of both IPv4 and IPv6 addresses.

The system appends RD to IPv6 routes and exchanges the labeled IPv6-RD using the VPNv6 address-family. The Address Family Identifier (AFI) and Subsequent Address Family Identifier (SAFI) fields for VPNv6 routes will be set to 2 and 128 respectively.

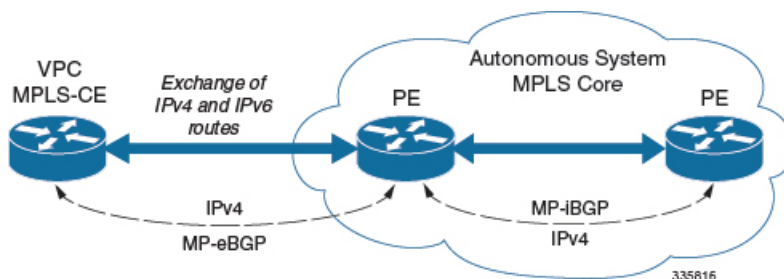
The IPv6 VPN traffic will be transported to the BGP speaker via IPv4 tunneling. The BGP speaker advertises to its peer a Next Hop Network Address field containing a VPN-IPv6 address whose 8-octet RD is set to zero and whose 16-octet IPv6 address is encoded as an IPv4-mapped IPv6 address (RFC 4291) containing the IPv4 address of the advertising router. It is assumed that only EBGP peering will be used to exchange VPNv6 routes.

Support for VPN-IPv6 assumes the following:

- Dual Stack (IPv4/IPv6) routing
- IPv6 pools in VRFs
- BGP peering over a directly connected IPv4 interface

See the figure below.

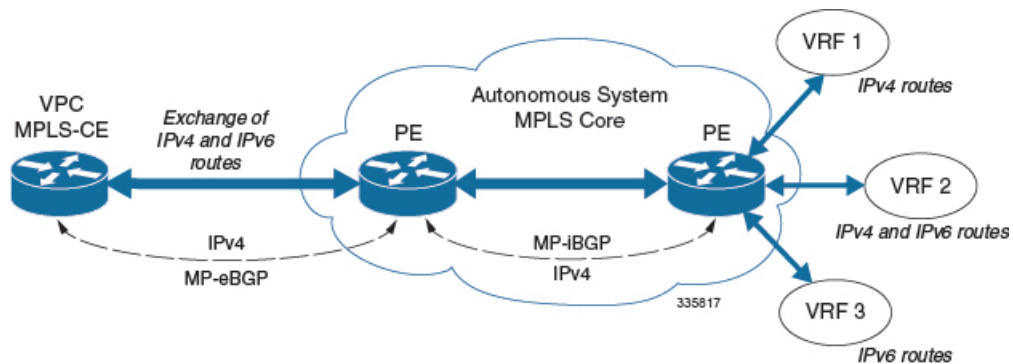
Figure 4: IPv6-RD Support for VPNv6



Sample Configuration

This example assumes three VRFs. VRF 1 has only IPv4 routes, VRF 2 has both IPv4 and IPv6 routes, and VRF 3 has only IPv6 routes.

Figure 5: VPNv6 Sample Configuration



Configure VRFs.

```
ip vrf vrf1
exit
ip vrf vrf2
exit
ip vrf vrf3
exit
```

Enable MPLS BGP forwarding.

```
mpls bgp forwarding
```

Configure pools.

```
ip pool vrf1-pool 51.52.53.0 255.255.255.0 private 0 vrf vrf1
exit
ip pool vrf2-pool 51.52.53.0 255.255.255.0 private 0 vrf vrf2
exit
ipv6 pool vrf2-v6pool prefix 2005:0101::/32 private 0 vrf vrf2
exit
ipv6 pool vrf3-v6pool prefix 2005:0101::/32 private 0 vrf vrf3
exit
```

Configure interfaces.

```
interface ce_interface to_rtr
 ip address 192.168.110.90 255.255.255.0
exit
interface ce_v6_interface
 ip address 2009:0101:0101:0101::1/96
exit
interface ce_loopback loopback
 ip address 52.1.2.3 255.255.255.255
exit
interface vrf1-loop loopback
 ip vrf forwarding vrf1
 ip address 1.52.53.54 255.255.255.255
exit
interface vrf2-loop loopback
 ip vrf forwarding vrf2
 ip address 2.52.53.54 255.255.255.255
exit
interface vrf2-v6loop loopback
 ip vrf forwarding vrf2
```

```

    ip address 2005:0202:0101::1/128
  exit
  interface vrf3-v6loop loopback
    ip vrf forwarding vrf3
    ip address 2005:0303:0101::1/128
  exit

```

Configure BGP along with address families and redistribution rules.

```

router bgp 800
  router-id 1.1.1.1
  neighbor 192.168.110.20 remote-as 1003
  neighbor 192.168.110.20 activate
  address-family vpnv4
    neighbor 192.168.110.20 activate
    neighbor 192.168.110.20 send-community both
  exit
  address-family vpnv6
    neighbor 192.168.110.20 activate
    neighbor 192.168.110.20 send-community both
  exit
  ip vrf vrf1
    route-distinguisher 800 1
    route-target export 800 1
    route-target import 800 1
  exit
  address-family ipv4 vrf vrf1
    redistribute connected
    redistribute static
  exit
  ip vrf vrf2
    route-distinguisher 800 2
    route-target export 800 2
    route-target import 800 2
  exit
  address-family ipv4 vrf vrf2
    redistribute connected
    redistribute static
  exit
  address-family ipv6 vrf vrf2
    redistribute connected
    redistribute static
  exit
  ip vrf vrf3
    route-distinguisher 800 3
    route-target export 800 3
    route-target import 800 3
  exit
  address-family ipv6 vrf vrf3
    redistribute connected
    redistribute static
  exit

```

Configure APNs.

```

apn walmart51.com
  selection-mode sent-by-ms
  accounting-mode none
  aaa group walmart-group
  authentication pap 1 chap 2 allow-noauth
  ip context-name Gi_ce
  ip address pool name vrf1-pool
  exit
apn amazon51.com
  selection-mode sent-by-ms
  accounting-mode none
  aaa group amazon-group
  authentication pap 1 chap 2 allow-noauth
  ip context-name Gi_ce
  ip address pool name vrf2-pool
  ipv6 address prefix-pool vrf2-v6pool
  exit
apn apple51.com
  selection-mode sent-by-ms

```

```

accounting-mode none
aaa group apple-group
authentication pap 1 chap 2 allow-noauthip context-name Gi_ce
ipv6 address prefix-pool vrf3-v6pool
exit
aaa-group amazon-group
radius ip vrf vrf2
aaa group default
exit
gtpv group default
exit
ip igmp profile default
exit

```

Bind physical interfaces with the port.

VPN-Related CLI Commands

VPN-related features and functions are supported across several CLI command modes. The following tables identify commands associated with configuration and monitoring of VPN-related functions.

For detailed information regarding the use of the commands listed below, see the *Command Line Interface Reference*.

Table 1: VPN-Related Configuration Commands

| CLI Mode | Command | Description |
|--|---|--|
| BGP Address-Family (IPv4/IPv6) Configuration Mode | neighbor <i>ip_address</i> activate | Enables the exchange of routing information with a peer router. |
| BGP Address-Family (IPv4/IPv6) Configuration Mode | neighbor <i>ip_address</i> send community { both extended standard } | Sends the community attributes to a peer router (neighbor). |
| BGP Address-Family (IPv4/IPv6) Configuration Mode | redistribute connected | Redistributes routes into BGP from another protocol as BGP neighbors. |
| BGP Address-Family (VPNv4) Configuration Mode | neighbor <i>ip_address</i> activate | Enables the exchange of routing information with a peer router. |
| BGP Address-Family (VPNv4) Configuration Mode | neighbor <i>ip_address</i> send community { both extended standard } | Sends the extended-community attribute to a peer router. In VPN, route-distinguisher and route-target are encoded in the BGP extended-community. This command enables sending of BGP routes with extended community to a neighbor. |
| BGP Address-Family (VRF) Configuration Mode | neighbor <i>ip_address</i> activate | Enables the exchange of routing information with a peer router. |

| CLI Mode | Command | Description |
|---|--|---|
| BGP Address-Family (VRF) Configuration Mode | neighbor <i>ip_address</i> send community { both extended standard } | Sends the extended-community attribute to a peer router. In VPN, route-distinguisher and route-target are encoded in the BGP extended-community. This command enables sending of BGP routes with extended community to a neighbor. |
| BGP Address-Family (VRF) Configuration Mode | redistribute connected | Redistributes routes into BGP from another protocol as BGP neighbors. |
| BGP Configuration Mode | address-family { ipv4 vrf <i>vrf_name</i> vpn4 } | Enables the exchange of IPv4 VRF routing information. There is a different mode for each address-family. |
| BGP Configuration Mode | address-family { ipv6 vrf <i>vrf_name</i> vpn6 } | Configures a VPNv6 address family and IPv6 VRF routing in BGP. |
| BGP Configuration Mode | ip vrf <i>vrf_name</i> | Adds a VRF to BGP and switches to the VRF Configuration mode to allow configuration of BGP attributes for the VRF. |
| BGP IP VRF Configuration Mode | route-distinguisher { <i>as_value</i> <i>ip_address</i> } <i>rd_value</i> | Assigns a Route Distinguisher (RD) for the VRF. The RD value must be a unique value on the router for each VRF. |
| BGP IP VRF Configuration Mode | route-target { both import export } { <i>as_value</i> <i>ip_address</i> } <i>rt_value</i> | Adds a list of import and export route-target extended communities to the VRF. |
| Context Configuration Mode | ip pool <i>pool_name</i> <i>addr_range</i> vrf <i>vrf_name</i> [mpls-label input <i>inlabel1</i> output <i>outlabel1</i> <i>outlabel2</i>] | Configures a pool into the specified VRF. This parameter must be specified with the Next-Hop parameter. <i>inlabel1</i> is the MPLS label that identifies inbound traffic destined for this pool. <i>outlabel1</i> and <i>outlabel2</i> specify the MPLS labels to be added to packets sent for subscribers from this pool. |
| Context Configuration Mode | ip vrf <i>vrf_name</i> | Creates a VRF and assigns a VRF-ID. A VRF is created in the router. |

| CLI Mode | Command | Description |
|---------------------------------------|--|--|
| Context Configuration Mode | ipv6 pool <i>pool_name</i> vrf <i>vrf_name</i> | Associates the pool with that VRF. Note: By default the configured ipv6 pool will be associated with the global routing domain. |
| Context Configuration Mode | mpls bgp forwarding | Globally enables MPLS Border Gateway Protocol (BGP) forwarding. |
| Context Configuration Mode | mpls exp <i>value</i> | Sets the default behavior as Best Effort using a zero value in the 3-bit MPLS EXP header. This value applies to all the VRFs in the context. The default behavior is to copy the DSCP value of mobile subscriber traffic to the EXP header, if there is no explicit configuration for DSCP to EXP (via the mpls map-dscp-to-exp dscp n exp m command). mpls exp disables the default behavior and sets the EXP value to the configured <i>value</i> . |
| Context Configuration Mode | mpls ip | Globally enables the MPLS forwarding of IPv4 packets along normally routed paths. |
| Context Configuration Mode | radius change-authorize-nas-ip ip_address <i>ip_address</i> { encrypted key } <i>value</i> port <i>port_num</i> mpls input <i>inlabel</i> output <i>outlabel1 outlabel2</i> | Configures COA traffic to use the specified MPLS labels. <i>inlabel</i> identifies inbound COA traffic. <i>outlabel1</i> and <i>outlabel2</i> specify the MPLS labels to be added to the COA response. <i>outlabel1</i> is the inner output label; <i>outlabel2</i> is the outer output label. |
| Ethernet Interface Configuration Mode | mpls ip | Enables dynamic MPLS forwarding of IP packets on this interface. |
| Exec Mode | clear ip bgp peer | Clears BGP sessions. |
| Exec Mode | lsp-ping <i>ip_prefix_FEC</i> | Checks MPLS Label-Switched Path (LSP) connectivity for the specified forwarding equivalence class (FEC). It must be followed by an IPv4 or IPv6 FEC prefix. |

| CLI Mode | Command | Description |
|-----------------------------------|---|---|
| Exec Mode | lsp-traceroute <i>ip_prefix_FEC</i> | Discovers MPLS LSP routes that packets actually take when traveling to their destinations. It must be followed by an IPv4 or IPv6 FEC prefix. |
| IP VRF Context Configuration Mode | mpls map-dscp-to-exp dscp <i>dscp_bit_value</i> exp <i>exp_bit_value</i> | Maps the final differentiated services code point (DSCP) bit value in the IP packet header to the final Experimental (EXP) bit value in the MPLS header for incoming traffic. |
| IP VRF Context Configuration Mode | mpls map-exp-to-dscp exp <i>exp_bit_value</i> dscp <i>dscp_bit_value</i> | Maps the incoming EXP bit value in the MPLS header to the internal DSCP bit value in IP packet headers for outgoing traffic. |
| MPLS-IP Configuration Mode | protocol ldp | Creates the MPLS protocol family configuration modes, or configures an existing protocol and enters the MPLS-LDP Configuration Mode in the current context. This command configures the protocol parameters for the MPLS protocol family. |
| MPLS-LDP Configuration Mode | advertise-labels { explicit-null implicit-null } | Configure advertisement of Implicit NULL or Explicit NULL label for all the prefixes advertised by the system in this context. |
| MPLS-LDP Configuration Mode | discovery { hello { hello-interval <i>seconds</i> hold-interval <i>seconds</i> } transport-address <i>ip_address</i> } | Configures the Label Distribution Protocol (LDP) neighbor discovery parameters. |
| MPLS-LDP Configuration Mode | enable | Enables Label Distribution Protocol (LDP). |
| MPLS-LDP Configuration Mode | router-id <i>ip_address</i> | Configures the LDP Router ID. |
| MPLS-LDP Configuration Mode | session timers { hold-interval <i>seconds</i> keepalive-interval <i>seconds</i> } | Configures the LDP session parameters. |

Table 2: VPN-Related Monitoring Commands

| CLI Mode | Command | Description |
|-------------------------|--|---|
| Exec Mode show Commands | show ip bgp neighbors | Displays information regarding BGP neighbors. |
| Exec Mode show Commands | show ip bgp vpnv4 { all route-distinguisher vrf } | Displays all VPNv4 routing data, routing data for a VRF or a route-distinguisher. |
| Exec Mode show Commands | show ip bgp vpnv6 | Displays contents of VPNv6 routing table. |
| Exec Mode show Commands | show ip bgp vpnv6 { all route-distinguisher vrf } | Displays all VPNv6 routing data, routing data for a VRF or a route-distinguisher. |
| Exec Mode show Commands | show ip pool | Displays pool details including the configured VRF. |
| Exec Mode show Commands | show mpls cross-connect | Displays MPLS cross-connect information. MPLS tunnel cross-connects between interfaces and Label-Switched Paths (LSPs) connect two distant interface circuits of the same type via MPLS tunnels that use LSPs as the conduit. |
| Exec Mode show Commands | show mpls ftm [vrf vrf_name | Displays MPLS FEC-to-NHLFE (FTN) table information. |
| Exec Mode show Commands | show mpls ftm [vrf vrf_name] | Displays contents of the MPLS FTN table for a specified VRF. |
| Exec Mode show Commands | show mpls ilm | Displays MPLS Incoming Label Map (ILM) table information. |
| Exec Mode show Commands | show mpls ldp | Displays the MPLS LDP information. |
| Exec Mode show Commands | show mpls nexthop-label-forwarding-entry | Displays MPLS Next-Hop Label Forwarding Entry (NHLFE) table information. |