



MPLS VPN Multipath Support for Inter-AS VPNs

The MPLS VPN Multipath Support for Inter-AS VPNs feature supports Virtual Private Network (VPN)v4 multipath for Autonomous System Boundary Routers (ASBRs) in the interautonomous system (Inter-AS) Multiprotocol Label Switching (MPLS) VPN environment. It allows load balancing of VPN traffic when you use the VPNv4 peering model for Inter-AS VPNs.

- [Finding Feature Information, page 1](#)
- [Restrictions for MPLS VPN Multipath Support for Inter-AS VPNs, page 1](#)
- [Information About MPLS VPN Multipath Support for Inter-AS VPNs, page 2](#)
- [How to Configure MPLS VPN Multipath Support for Inter-AS VPNs, page 3](#)
- [Configuration Examples for MPLS VPN Multipath Support for Inter-AS VPNs, page 10](#)
- [Additional References, page 17](#)
- [Feature Information for MPLS VPN Multipath Support for Inter-AS VPNs, page 19](#)
- [Glossary, page 19](#)

Finding Feature Information

Your software release may not support all the features documented in this module. For the latest caveats and feature information, see [Bug Search Tool](#) and the release notes for your platform and software release. To find information about the features documented in this module, and to see a list of the releases in which each feature is supported, see the feature information table at the end of this module.

Use Cisco Feature Navigator to find information about platform support and Cisco software image support. To access Cisco Feature Navigator, go to www.cisco.com/go/cfn. An account on Cisco.com is not required.

Restrictions for MPLS VPN Multipath Support for Inter-AS VPNs

The following restrictions apply to configuring multipath load sharing for Multiprotocol Label Switching (MPLS) Virtual Private Network (VPN) interautonomous system (Inter-AS) Autonomous System Boundary Routers (ASBRs) exchanging VPN-IPv4 routes:

- Per-packet load balancing is not supported for this feature. Load balancing for this feature works on the IP source and destination hash or on the bottom label in the label stack, depending on the platform and depth of the MPLS label stack.
- If MPLS scalability is an issue for you, we recommend that you do not enable VPNv4 multipath on ASBRs.

Information About MPLS VPN Multipath Support for Inter-AS VPNs

Load Sharing with MPLS VPN Inter-AS ASBRs

Before the MPLS VPN Interautonomous System Support feature, if multiple paths existed across Autonomous System Boundary Routers (ASBRs), the Border Gateway Protocol (BGP) executed the best path algorithm and marked only one of the paths as the best path. This path was added to the routing table and became the only path that was used for forwarding traffic between ASBRs.

The MPLS VPN Multipath Support for Inter-AS VPNs feature extends the functionality of BGP so that it can pick one path as the best path and mark the other legitimate paths between ASBRs as multipath. This allows the load sharing of traffic among the different multipaths and the best path to reach the destination. No Routing Information Base (RIB) or Cisco Express Forwarding entries are associated with the Virtual Private Network (VPN)-IPv4 prefixes.

The MPLS VPN Multipath Support for Inter-AS VPNs feature applies to ASBRs that do not have a VPN routing and forwarding (VRF) instance configuration. BGP installs a number of learned VPN-IPv4 prefixes into the Multiprotocol Label Switching (MPLS) forwarding table (LFIB). VPN-IPv4 entries in the LFIB consist of the Route Distinguisher (RD) and the IPv4 prefix and are called VPNv4 entries.

The MPLS VPN Multipath Support for Inter-AS VPNs feature requires that you configure the **maximum-paths number-of-paths** command in address family configuration mode. This command is used to set the number of parallel (equal-cost) routes that BGP installs in the routing table to configure multipath load sharing. The number of paths that can be configured is determined by the version of Cisco software.




Note

The **maximum-paths** command cannot be configured with the **maximum-paths eibgp** command for the same BGP routing process.

The figure below shows an example of VPNv4 load balancing for ASBRs in an Inter-AS network. In this example, ASBR1 load balances the traffic from the CE device CE1 to CE2 using the two available links—ASBR2 and ASBR3.

Figure 1: Example of VPNv4 Load Balancing for ASBRs in an Inter-AS Network



When you configure an ASBR for VPNv4 load balancing, you must configure the **next-hop-self** command for the iBGP peers. Without this command, the next hop that is propagated to the iBGP peer is the ASBR2 address or the ASBR3 address, depending on which one BGP selects as the best path. Configuring the **next-hop-self** command provides direct VPNv4 forwarding entries in the MPLS forwarding table for the VPNv4 prefixes learned from the remote ASBRs. VPNv4 forwarding entries are not created if you do not configure the **next-hop-self** command.



Note

If the number of forwarding entries in the MPLS forwarding table on the system or on a line card is a concern for your network, we recommend that you do not enable VPNv4 multipath on ASBRs.

How to Configure MPLS VPN Multipath Support for Inter-AS VPNs


Configuring eBGP Multipath Load Sharing for MPLS VPN Inter-AS ASBRs

Perform this task to configure the external Border Gateway Protocol (eBGP) multipath load sharing for MPLS VPN Inter-AS ASBRs exchanging Virtual Private Network (VPN)-IPv4 routes. This allows for more efficient use of the label switched paths (LSPs) in an interautonomous system network because you can set up the load sharing of traffic among the different multipaths and the best path to reach the destination.

**Note**

The figure below shows an eBGP multipath configuration for three VPN-IPv4 ASBRs. The links from ASBR1 to ASBR2 and ASBR3 have an eBGP VPN-IPv4 session configured. In the figure below, eBGP multipath load sharing is configured on ASBR1.

Figure 2: eBGP Multipath Configuration for Three VPN-IPv4 ASBRs



The configurations in the figure above is used as an example for this task.

SUMMARY STEPS

1. **enable**
2. **configure terminal**
3. **router bgp** *as-number*
4. **no bgp default route-target filter**
5. **neighbor** {*ip-address* | *peer-group-name*} **remote-as** *as-number*
6. **neighbor** {*ip-address* | *peer-group-name*} **update-source** *interface-type interface-number*
7. **neighbor** {*ip-address* | *peer-group-name*} **next-hop-self**
8. **neighbor** {*ip-address* | *peer-group-name*} **remote-as** *as-number*
9. Repeat Step 8 for each BGP neighbor.
10. **address-family vpnv4** [**unicast**]
11. **neighbor** {*ip-address* | *peer-group-name*} **activate**
12. **neighbor** {*ip-address* | *peer-group-name*} **next-hop-self**
13. **neighbor** {*ip-address* | *peer-group-name*} **send-community** [**both** | **standard** | **extended**]
14. **neighbor** {*ip-address* | *peer-group-name* | *ipv6-address*} **activate**
15. **neighbor** {*ip-address* | *peer-group-name*} **send-community** [**both** | **standard** | **extended**]
16. Repeat Steps 14 and 15 for each BGP neighbor.
17. **maximum-paths** *number-paths*
18. **exit-address-family**
19. **end**

DETAILED STEPS

| | Command or Action | Purpose |
|--------|---|---|
| Step 1 | enable Example: Device> enable | Enables privileged EXEC mode. <ul style="list-style-type: none"> • Enter your password if prompted. |
| Step 2 | configure terminal Example: Device# configure terminal | Enters global configuration mode. |
| Step 3 | router bgp <i>as-number</i> Example: Device(config)# router bgp 1 | Configures an eBGP routing process and places the device in router configuration mode. <ul style="list-style-type: none"> • The <i>as-number</i> argument indicates the number of an autonomous system that identifies the device to other BGP devices and tags the routing information passed along. The range is 0 to 65535. Private autonomous system numbers that can be used in internal networks range from 64512 to 65535. |
| Step 4 | no bgp default route-target filter Example: Device(config-router)# no bgp default route-target filter | Disables BGP route-target community filtering. <ul style="list-style-type: none"> • All received VPN-IPv4 routes are accepted by the configured device. Accepting VPN-IPv4 routes is the desired behavior for a device configured as an ASBR. |
| Step 5 | neighbor {<i>ip-address</i> <i>peer-group-name</i>} remote-as <i>as-number</i> Example: Device(config-router)# neighbor 10.1.0.4 remote-as 1 | Adds an entry to the BGP or multiprotocol BGP neighbor table. <ul style="list-style-type: none"> • <i>ip-address</i>—IP address of the neighbor. • <i>peer-group-name</i>—Name of a BGP peer group. • <i>as-number</i>—The autonomous system to which the neighbor belongs. |
| Step 6 | neighbor {<i>ip-address</i> <i>peer-group-name</i>} update-source <i>interface-type interface-number</i> Example: Device(config-router)# neighbor 10.1.0.4 update-source loopback 0 | Allows BGP sessions to use any operational interface for TCP connections. <ul style="list-style-type: none"> • <i>ip-address</i>—IP address of the neighbor. • <i>peer-group-name</i>—Name of a BGP peer group. • <i>interface-type interface-number</i>—Type and number for the operational interface. <p>This example shows how to set up BGP TCP connections for the specified neighbor with the IP address of the loopback interface rather than the best local address.</p> |

| | Command or Action | Purpose |
|----------------|---|--|
| Step 7 | <p>neighbor {<i>ip-address</i> <i>peer-group-name</i>} next-hop-self</p> <p>Example:</p> <pre>Device(config-router)# neighbor 10.1.0.4 next-hop-self</pre> | <p>Configures the device as the next hop for a BGP neighbor or peer group.</p> <ul style="list-style-type: none"> • <i>ip-address</i>—IP address of the BGP neighbor. • <i>peer-group-name</i>—Name of a BGP peer group. |
| Step 8 | <p>neighbor {<i>ip-address</i> <i>peer-group-name</i>} remote-as <i>as-number</i></p> <p>Example:</p> <pre>Device(config-router)# neighbor 172.16.1.9 remote-as 2</pre> | <p>Adds an entry to the BGP or multiprotocol BGP neighbor table.</p> <ul style="list-style-type: none"> • <i>ip-address</i>—IP address of the neighbor. • <i>peer-group-name</i>—Name of a BGP peer group. • <i>as-number</i>—Autonomous system to which the neighbor belongs. |
| Step 9 | Repeat Step 8 for each BGP neighbor. | — |
| Step 10 | <p>address-family <i>vpn4</i> [unicast]</p> <p>Example:</p> <pre>Device(config-router)# address-family vpn4</pre> | <p>Enters address family configuration mode.</p> <ul style="list-style-type: none"> • unicast—Specifies a unicast prefix. <p>This command configures a routing session to carry VPN-IPv4 addresses across the VPN backbone. Each address is globally unique by the addition of an 8-byte RD.</p> |
| Step 11 | <p>neighbor {<i>ip-address</i> <i>peer-group-name</i>} activate</p> <p>Example:</p> <pre>Device(config-router-af)# neighbor 10.1.0.4 activate</pre> | <p>Enables the exchange of information with a neighboring device.</p> <ul style="list-style-type: none"> • <i>ip-address</i>—IP address of the neighbor. • <i>peer-group-name</i>—Name of a BGP peer group. |
| Step 12 | <p>neighbor {<i>ip-address</i> <i>peer-group-name</i>} next-hop-self</p> <p>Example:</p> <pre>Device(config-router-af)# neighbor 10.1.0.4 next-hop-self</pre> | <p>Configures the device as the next hop for a BGP neighbor or peer group.</p> <ul style="list-style-type: none"> • <i>ip-address</i>—IP address of the BGP neighbor. • <i>peer-group-name</i>—Name of a BGP peer group. |
| Step 13 | <p>neighbor {<i>ip-address</i> <i>peer-group-name</i>} send-community [both standard extended]</p> <p>Example:</p> <pre>Device(config-router-af)# neighbor 10.1.0.4 send-community extended</pre> | <p>Specifies that a communities attribute should be sent to a BGP neighbor.</p> <ul style="list-style-type: none"> • <i>ip-address</i>—IP address of the neighboring device. • <i>peer-group-name</i>—Name of a BGP peer group. • both—Specifies that both standard and extended communities will be sent. • standard—Specifies that only standard communities will be sent. |

| | Command or Action | Purpose |
|----------------|---|--|
| | | <ul style="list-style-type: none"> • extended—Specifies that only extended communities will be sent. |
| Step 14 | <p>neighbor {<i>ip-address</i> <i>peer-group-name</i> <i>ipv6-address</i>} activate</p> <p>Example:</p> <pre>Device(config-router-af)# neighbor 172.16.1.9 activate</pre> | <p>Enables the exchange of information with a BGP neighbor.</p> <ul style="list-style-type: none"> • <i>ip-address</i>—IP address of the neighboring device. • <i>peer-group-name</i>—Name of a BGP peer group. • <i>ipv6-address</i>—IPv6 address of the BGP-speaking neighbor. <p>This argument must be in the form documented in RFC 2373, where the address is specified in hexadecimal using 16-bit values between colons.</p> |
| Step 15 | <p>neighbor {<i>ip-address</i> <i>peer-group-name</i>} send-community [both standard extended]</p> <p>Example:</p> <pre>Device(config-router-af)# neighbor 172.16.1.9 send-community extended</pre> | <p>Specifies that a communities attribute should be sent to a BGP neighbor.</p> <ul style="list-style-type: none"> • <i>ip-address</i>—IP address of the neighboring device. • <i>peer-group-name</i>—Name of a BGP peer group. • both—Specifies that both standard and extended communities will be sent. • standard—Specifies that only standard communities will be sent. • extended—Specifies that only extended communities will be sent. |
| Step 16 | Repeat Steps 14 and 15 for each BGP neighbor. | |
| Step 17 | <p>maximum-paths <i>number-paths</i></p> <p>Example:</p> <pre>Device(config-router-af)# maximum-paths 2</pre> | <p>Configures the maximum number of parallel routes that an IP routing protocol will install into the routing table.</p> <ul style="list-style-type: none"> • <i>number-paths</i>—Number of routes to install to the routing table. |
| Step 18 | <p>exit-address-family</p> <p>Example:</p> <pre>Device(config-router-af)# exit-address-family</pre> | Exits from address family configuration mode. |
| Step 19 | <p>end</p> <p>Example:</p> <pre>Device(config-router)# end</pre> | (Optional) Exits to privileged EXEC mode. |

Example

The following example shows the configuration for external Border Gateway Protocol (eBGP) multipath for VPNv4 sessions on the ASBR1 device:

```
configure terminal
router bgp 1
 no bgp default route-target filter
 neighbor 10.1.0.4 remote-as 1
 neighbor 10.1.0.4 update-source Loopback 0
 neighbor 10.1.0.4 next-hop-self
 neighbor 172.16.1.9 remote-as 2
 neighbor 172.16.2.8 remote-as 2
!
 address-family vpnv4
 neighbor 10.1.0.4 activate
 neighbor 10.1.0.4 next-hop-self
 neighbor 10.1.0.4 send-community extended
 neighbor 172.16.1.9 activate
 neighbor 172.16.1.9 send-community extended
 neighbor 172.16.2.8 activate
 neighbor 172.16.2.8 send-community extended
 maximum-paths 2
 exit-address-family
end
```

Verifying eBGP Multipath Load Sharing for MPLS VPN Inter-AS ASBRs

Perform the following task to verify that the external Border Gateway Protocol (eBGP) multipath load sharing for Multiprotocol Label Switching (MPLS) Virtual Private Network (VPN) interautonomous system (Inter-AS) Autonomous System Boundary Routers (ASBRs) is operating as you expect.

The configurations in the figure above are used as an example for the task that follows.

SUMMARY STEPS

1. **enable**
2. **show ip bgp vpnv4 all [summary]**
3. **show ip bgp vpnv4 all**
4. **show ip bgp vpnv4 [network]**
5. **show mpls forwarding-table**
6. **exit**

DETAILED STEPS

Step 1 **enable**
Enables privileged EXEC mode. Enter your password if required.

Example:

```
Device> enable
Device#
```

Step 2 **show ip bgp vpnv4 all [summary]**

Verifies that all peers are up.

Example:

```
Device# show ip bgp vpnv4 all summary
```

| Neighbor | V | AS | MsgRcvd | MsgSent | TblVer | InQ | OutQ | Up/Down | State/PfxRcd |
|------------|---|----|---------|---------|--------|-----|------|----------|--------------|
| 10.1.0.4 | 4 | 1 | 87 | 86 | 5 | 0 | 0 | 01:24:56 | 2 |
| 172.16.1.9 | 4 | 2 | 88 | 88 | 5 | 0 | 0 | 01:25:49 | 2 |
| 172.16.2.8 | 4 | 2 | 88 | 88 | 5 | 0 | 0 | 01:25:49 | 2 |

The output shows that all peers expected to be up are up and sending and receiving messages.

Step 3

show ip bgp vpnv4 all

Verifies that BGP has paths from both remote ASBRs.

Example:

```
Device# show ip bgp vpnv4 all
```

| Network | Next Hop | Metric | LocPrf | Weight | Path |
|----------------------------|-------------------|--------|--------|----------|------------|
| . | . | . | . | . | . |
| Route Distinguisher: 1:105 | | | | | |
| *>i192.168.0.1/32 | 10.1.0.3 | 11 | 100 | 0 | ? ? |
| *> 192.168.0.2/32 | 172.16.2.8 | | | 0 | 2 ? |
| * | 172.16.1.9 | | | 0 | 2 ? |
| *>i192.168.1.0 | 10.1.0.3 | 0 | 100 | 0 | ? ? |
| *> 192.168.2.0 | 172.16.2.8 | | | 0 | 2 ? |
| * | 172.16.1.9 | | | 0 | 2 ? |

The bold entries in the output confirm that BGP has a path to ASBR2 (172.16.2.8) and to ASBR3 (172.16.1.9).

Step 4

show ip bgp vpnv4 [network]

Verifies that paths are marked as multipath.

Example:

```
Device# show ip bgp vpnv4 192.168.2.0
```

```
BGP routing table entry for 1:105:192.168.2.0/24, version 3
Paths: (2 available, best #1, no table)
  Advertised to update-groups:
    2          3
  2
    172.16.2.8 from 172.16.2.8 (10.2.0.8)
      Origin incomplete, localpref 100, valid, external, multipath
, best
      Extended Community: RT:1:100 OSPF DOMAIN ID:0x0005:0x0000000A0200
        OSPF RT:0.0.0.0:2:0 OSPF ROUTER ID:192.168.2.2:512,
mpls labels in/out 21/25
  2
    172.16.1.9 from 172.16.1.9 (10.2.0.9)
      Origin incomplete, localpref 100, valid, external, multipath
      Extended Community: RT:1:100 OSPF DOMAIN ID:0x0005:0x0000000A0200
        OSPF RT:0.0.0.0:2:0 OSPF ROUTER ID:192.168.2.2:512,
mpls labels in/out 21/25
```

In the output, the “multipath” and “mpls labels in/out 21/25” are in bold text for example purposes only.

Step 5

show mpls forwarding-table

Verifies that MPLS forwarding is properly set up and counters are increasing when traffic is present.

Example:

```
Device# show mpls forwarding-table
```

| Local Label | Outgoing Label or VC | Prefix or Tunnel Id | Bytes Switched | Label | Outgoing interface | Next Hop |
|-------------|----------------------|----------------------|----------------|-------|--------------------|------------|
| . | . | . | . | . | . | . |
| 16 | Pop Label | 172.16.1.9/32 | 0 | | Et1/0 | 172.16.1.9 |
| 17 | Pop Label | 172.16.2.8/32 | 0 | | Et2/0 | 172.16.2.8 |
| 18 | Pop Label | 10.1.1.0/24 | 0 | | Et0/0 | 10.1.2.4 |
| 19 | 16 | 10.1.0.3/32 | 0 | | Et0/0 | 10.1.2.4 |
| 20 | Pop Label | 10.1.0.4/32 | 0 | | Et0/0 | 10.1.2.4 |
| 21 | 25 | 1:105:192.168.2.0/24 | \ | | | |
| | | | 26658 | | Et1/0 | 172.16.1.9 |
| | 25 | 1:105:192.168.2.0/24 | \ | | | |
| | | | 1180 | | Et2/0 | 172.16.2.8 |
| 22 | 24 | 1:105:192.168.0.2/32 | \ | | | |
| | | | 15740 | | Et1/0 | 172.16.1.9 |
| | 24 | 1:105:192.168.0.2/32 | \ | | | |
| | | | 0 | | Et2/0 | 172.16.2.8 |
| 23 | 19 | 1:105:192.168.0.1/32 | \ | | | |
| | | | 15638 | | Et0/0 | 10.1.2.4 |
| 24 | 20 | 1:105:192.168.1.0/24 | \ | | | |
| | | | 32740 | | Et0/0 | 10.1.2.4 |

Step 6 **exit**

Exits to user EXEC mode. For example:

Example:

```
Device# exit
Device>
```

Configuration Examples for MPLS VPN Multipath Support for Inter-AS VPNs

Example: Configuring eBGP Multipath Load Sharing for MPLS VPN Inter-AS ASBRs


This section includes examples that show how to configure the external Border Gateway Protocol (eBGP) multipath load sharing for Multiprotocol Label Switching (MPLS) Virtual Private Network (VPN) interautonomous system (Inter-AS) Autonomous System Boundary Routers (ASBRs) that exchange VPN-IPv4 routes.

The network topology in the figure below shows two autonomous systems, which are configured as follows:

- Autonomous system 1 contains PE1, P1, and ASBR1.
- Autonomous system 2 contains PE2, P2, ASBR2, and ASBR3.
- CE1 and CE2 belong to the same VPN, which is called VPN1.

- The P devices are route reflectors.
- ASBR1 and ASBR2 are configured with the **neighbor next-hop-self** command for the internal BGP (iBGP) neighbors.
- ASBR1 and ASBR2 are configured with the **maximum paths** commands to set up eBGP multipath load sharing.

Figure 3: Configuring eBGP Multipath Load Sharing Between MPLS Inter-AS ASBRs Exchanging VPN-IPv4 Routes



The following examples show how to configure eBGP multipath load sharing for MPLS VPN Inter-AS ASBRs that exchange VPN-IPv4 routes. This section includes sample configurations for P1, ASBR1, ASBR2, and P2 devices.

Example: Multipath Support for Inter-AS VPNs Configuration for Autonomous System 1 CE1

The following example shows how to configure CE1 in VPN1:

```
!
hostname CE1
!
interface Loopback 1
 ip address 192.168.0.1 255.255.255.255
!
interface Ethernet 1/0
 description Link to PE1
 ip address 192.168.1.1 255.255.255.0
!
router ospf 1
 log-adjacency-changes
 network 192.168.0.0 0.0.255.255 area 0
!
end
```

Example: Multipath Support for Inter-AS VPNs Configuration for Autonomous System 1 PE1

The following example shows how to configure PE1 in autonomous system 1:

```
!
```

```

hostname PE1
!
ip cef
!
ip vrf V1
  rd 1:105
  route-target export 1:100
  route-target import 1:100
!
interface Loopback 0
  ip address 10.1.0.3 255.255.255.255
!
interface Ethernet 0/0
  description Link to CE1
  ip vrf forwarding V1
  ip address 192.168.1.2 255.255.255.0
!
interface Ethernet 1/0
  description Link to P1
  ip address 10.1.1.3 255.255.255.0
  mpls ip
!
router ospf 10 vrf V1
  log-adjacency-changes
  redistribute bgp 1 metric 100 subnets
  network 192.168.0.0 0.0.255.255 area 0
!
router ospf 1
  log-adjacency-changes
  network 10.0.0.0 0.255.255.255 area 0
!
router bgp 1
  no synchronization
  bgp log-neighbor-changes
  neighbor 10.1.0.4 remote-as 1
  no neighbor 10.1.0.4 transport path-mtu-discovery
  neighbor 10.1.0.4 update-source Loopback 0
  no auto-summary
!
  address-family vpnv4
    neighbor 10.1.0.4 activate
    neighbor 10.1.0.4 send-community extended
  exit-address-family
!
  address-family ipv4 vrf V1
    redistribute ospf 10 vrf V1
    no auto-summary
    no synchronization
  exit-address-family
!
end

```

Example: Multipath Support for Inter-AS VPNs Configuration for Autonomous System 1 P1

The following example shows how to configure P1 in autonomous system 1:

```

!
hostname P1
!
ip cef
!
interface Loopback 0
  ip address 10.1.0.4 255.255.255.255
!
interface Ethernet 0/0
  description Link to PE1
  ip address 10.1.1.4 255.255.255.0
  mpls ip
!
interface Ethernet 1/0

```

```

description Link to ASBR1
ip address 10.1.2.4 255.255.255.0
mpls ip
!
router ospf 1
log-adjacency-changes
network 10.0.0.0 0.255.255.255 area 0
!
router bgp 1
no synchronization
bgp log-neighbor-changes
neighbor R peer-group
neighbor R remote-as 1
no neighbor R transport path-mtu-discovery
neighbor R update-source Loopback 0
neighbor R route-reflector-client
neighbor 10.1.0.3 peer-group R
neighbor 10.1.0.5 peer-group R
no auto-summary
!
address-family vpnv4
neighbor R send-community extended
neighbor R route-reflector-client
neighbor 10.1.0.3 activate
neighbor 10.1.0.5 activate
exit-address-family
!
end

```

Example: Multipath Support for Inter-AS VPNs Configuration for Autonomous System 1 ASBR1

The following example shows how to configure ASBR1 in autonomous system 1:

```

hostname ASBR1
!
ip cef
!
interface Loopback 0
ip address 10.1.0.5 255.255.255.255
!
interface Ethernet 0/0
description Core link to P1
ip address 10.1.2.5 255.255.255.0
mpls ip
!
interface Ethernet 1/0
description Link to ASBR2
ip address 172.16.2.5 255.255.255.0
mpls bgp forwarding
!
interface Serial 3/0
description Link to ASBR3
ip address 172.16.1.5 255.255.255.0
mpls bgp forwarding
serial restart-delay 0
!
!
router ospf 1
log-adjacency-changes
network 10.0.0.0 0.255.255.255 area 0
!
router bgp 1
no synchronization
no bgp default route-target filter
bgp log-neighbor-changes
neighbor 10.1.0.4 remote-as 1
neighbor 172.16.1.9 remote-as 2
neighbor 172.16.2.8 remote-as 2
no auto-summary
!

```

Example: Configuring eBGP Multipath Load Sharing for MPLS VPN Inter-AS ASBRs

```

address-family vpnv4
neighbor 10.1.0.4 activate
neighbor 10.1.0.4 send-community extended
neighbor 10.1.0.4 next-hop-self
neighbor 172.16.1.9 activate
neighbor 172.16.1.9 send-community extended
neighbor 172.16.2.8 activate
neighbor 172.16.2.8 send-community extended
maximum-paths 2
exit-address-family
!
end

```

Example: Multipath Support for Inter-AS VPNs Configuration for Autonomous System 2 ASBR2

The following example shows how to configure ASBR2 in autonomous system 2:

```

!
hostname ASBR2
!
ip cef
!
interface Loopback 0
ip address 10.2.0.8 255.255.255.255
!
interface Loopback 1
no ip address
shutdown
!
interface Ethernet 0/0
description Link to ASBR1
ip address 172.16.2.8 255.255.255.0
mpls bgp forwarding
!
interface Serial 2/0
description Link to P2
ip address 10.2.2.8 255.255.255.0
mpls ip
no fair-queue
serial restart-delay 0
!
router ospf 1
log-adjacency-changes
redistribute connected subnets
network 10.0.0.0 0.255.255.255 area 0
!
router bgp 2
no synchronization
no bgp default route-target filter
bgp log-neighbor-changes
neighbor 10.2.0.7 remote-as 2
neighbor 10.2.0.7 update-source Loopback 0
neighbor 10.2.0.7 next-hop-self
neighbor 172.16.2.5 remote-as 1
no auto-summary
!
address-family vpnv4
neighbor 10.2.0.7 activate
neighbor 10.2.0.7 send-community extended
neighbor 10.2.0.7 next-hop-self
neighbor 172.16.2.5 activate
neighbor 172.16.2.5 send-community extended
exit-address-family
!
end

```

Example: Multipath Support for Inter-AS VPNs Configuration for Autonomous System 2 ASBR3

The following example shows how to configure ASBR3 in autonomous system 2:

```

!
hostname ASBR3
!
ip cef
!
interface Loopback 0
 ip address 10.2.0.9 255.255.255.255
!
interface Ethernet 0/0
 description Link to ASBR1
 ip address 172.16.1.9 255.255.255.0
 mpls bgp forwarding
!
interface Serial 3/0
 description Link to P2
 ip address 10.2.3.9 255.255.255.0
 mpls ip
 no fair-queue
 serial restart-delay 0
!
router ospf 1
 log-adjacency-changes
 redistribute connected subnets
 network 10.0.0.0 0.255.255.255 area 0
!
router bgp 2
 no synchronization
 no bgp default route-target filter
 bgp log-neighbor-changes
 neighbor 10.2.0.7 remote-as 2
 neighbor 10.2.0.7 update-source Loopback 0
 neighbor 10.2.0.7 next-hop-self
 neighbor 172.16.1.5 remote-as 1
 no auto-summary
!
 address-family vpnv4
 neighbor 10.2.0.7 activate
 neighbor 10.2.0.7 send-community extended
 neighbor 10.2.0.7 next-hop-self
 neighbor 172.16.1.5 activate
 neighbor 172.16.1.5 send-community extended
 exit-address-family
!
end

```

Example: Multipath Support for Inter-AS VPNs Configuration for Autonomous System 2 P2

The following example shows how to configure P2 in autonomous system 2:

```

!
hostname P2
!
ip cef
!
interface Loopback 0
 ip address 10.2.0.7 255.255.255.255
!
interface Ethernet 1/0
 description Link to PE2
 ip address 10.2.1.7 255.255.255.0
 mpls ip
!
interface Serial 2/0

```

Example: Configuring eBGP Multipath Load Sharing for MPLS VPN Inter-AS ASBRs

```

description Link to ASBR2
ip address 10.2.2.7 255.255.255.0
mpls ip
no fair-queue
serial restart-delay 0
!
interface Serial 3/0
description Link to ASBR3
ip address 10.2.3.7 255.255.255.0
mpls ip
serial restart-delay 0
!
router ospf 1
log-adjacency-changes
network 10.0.0.0 0.255.255.255 area 0
!
router bgp 2
no synchronization
bgp log-neighbor-changes
neighbor R peer-group
neighbor R remote-as 2
no neighbor R transport path-mtu-discovery
neighbor R update-source Loopback 0
neighbor R route-reflector-client
neighbor 10.2.0.6 peer-group R
neighbor 10.2.0.8 peer-group R
neighbor 10.2.0.9 peer-group R
no auto-summary
!
address-family vpnv4
neighbor R send-community extended
neighbor R route-reflector-client
neighbor 10.2.0.6 activate
neighbor 10.2.0.8 activate
neighbor 10.2.0.9 activate
exit-address-family
!
end
!
```

Example: Multipath Support for Inter-AS VPNs Configuration for Autonomous System 2 PE2

The following example shows how to configure PE2 in autonomous system 2:

```

hostname PE2
!
ip cef
!
ip vrf V1
rd 1:105
route-target export 1:100
route-target import 1:100
!
interface Loopback 0
ip address 10.2.0.6 255.255.255.255
!
interface Ethernet 0/0
description Link to P2
ip address 10.2.1.6 255.255.255.0
mpls ip
!
interface Serial 2/0
description Link to CE2
ip vrf forwarding V1
ip address 192.168.2.2 255.255.255.0
no fair-queue
serial restart-delay 0
!
router ospf 10 vrf V1
log-adjacency-changes
```



```

redistribute bgp 2 subnets
network 192.168.0.0 0.0.255.255 area 0
!
router ospf 1
log-adjacency-changes
network 10.0.0.0 0.255.255.255 area 0
!
router bgp 2
no synchronization
bgp log-neighbor-changes
neighbor 10.2.0.7 remote-as 2
neighbor 10.2.0.7 update-source Loopback 0
no auto-summary
!
address-family vpnv4
neighbor 10.2.0.7 activate
neighbor 10.2.0.7 send-community extended
exit-address-family
!
address-family ipv4 vrf V1
redistribute connected
redistribute ospf 10 vrf V1
no auto-summary
no synchronization
exit-address-family
!
end

```

Example: Multipath Support for Inter-AS VPNs Configuration for Autonomous System 2 CE2

The following example shows how to configure CE2 in VPN1:

```

hostname CE2
!
interface Loopback 0
ip address 192.168.0.2 255.255.255.255
!
interface Serial 2/0
description Link to PE2
ip address 192.168.2.1 255.255.255.0
no fair-queue
serial restart-delay 0
!
router ospf 1
log-adjacency-changes
network 192.168.0.0 0.0.255.255 area 0
end

```

Additional References

Related Documents

| Related Topic | Document Title |
|---|---|
| Cisco IOS commands | Cisco IOS Master Command List, All Releases |
| MPLS commands | Cisco IOS Multiprotocol Label Switching Command Reference |
| Configuration tasks for basic MPLS VPNs | “Configuring MPLS VPNs” module in the <i>MPLS: Layer 3 VPNs Configuration Guide</i> |

| Related Topic | Document Title |
|---|--|
| Configuration tasks for MPLS VPN Inter-AS system exchanging IPv4 routes and MPLS labels | “MPLS VPN Inter-AS with ASBRs Exchanging IPv4 Routes and MPLS Labels” module in the <i>MPLS: Layer 3 VPNs Inter-AS and CSC Configuration Guide</i> |
| Information about monitoring MPLS VPNs with MIBs | “MPLS VPN SNMP MIB Notifications” module in the <i>MPLS: Embedded Management and MIBs Configuration Guide</i> |

RFCs

| RFC | Title |
|----------|---|
| RFC 1164 | <i>Application of the Border Gateway Protocol in the Internet</i> |
| RFC 1700 | <i>Assigned Numbers</i> |
| RFC 1771 | <i>A Border Gateway Protocol 4</i> |
| RFC 1965 | <i>Autonomous System Confederation for BGP</i> |
| RFC 1966 | <i>BGP Route Reflection: An Alternative to Full Mesh iBGP</i> |
| RFC 2547 | <i>BGP/MPLS VPNs</i> |
| RFC 2842 | <i>Capabilities Advertisement with BGP-4</i> |
| RFC 2858 | <i>Multiprotocol Extensions for BGP-4</i> |
| RFC 3107 | <i>Carrying Label Information in BGP-4</i> |

Technical Assistance

| Description | Link |
|---|---|
| The Cisco Support and Documentation website provides online resources to download documentation, software, and tools. Use these resources to install and configure the software and to troubleshoot and resolve technical issues with Cisco products and technologies. Access to most tools on the Cisco Support and Documentation website requires a Cisco.com user ID and password. | http://www.cisco.com/cisco/web/support/index.html |

Feature Information for MPLS VPN Multipath Support for Inter-AS VPNs

The following table provides release information about the feature or features described in this module. This table lists only the software release that introduced support for a given feature in a given software release train. Unless noted otherwise, subsequent releases of that software release train also support that feature.

Use Cisco Feature Navigator to find information about platform support and Cisco software image support. To access Cisco Feature Navigator, go to www.cisco.com/go/cfn. An account on Cisco.com is not required.

Table 1: Feature Information for MPLS VPN Multipath Support for Inter-AS VPNs

| Feature Name | Releases | Feature Information |
|--|--|---|
| MPLS VPN Multipath Support for Inter-AS VPNs | 12.2(30)S 12.2(33)SRA 12.2(33)SXH 12.4(20)T | The MPLS VPN Multipath Support for Inter-AS VPNs feature supports Virtual Private Network (VPN)v4 multipath for Autonomous System Boundary Routers (ASBRs) in the interautonomous system (Inter-AS) Multiprotocol Label Switching (MPLS) VPN environment. It allows load balancing of VPN traffic when you use the VPNv4 peering model for Inter-AS VPNs. No commands were introduced or modified. |

Glossary

autonomous system—A collection of networks under a common administration sharing a common routing strategy.

BGP —Border Gateway Protocol. An interdomain routing protocol that exchanges network reachability information with other BGP systems (which may be within the same autonomous system or between multiple autonomous systems).

CE device—customer edge device. A device that is part of a customer network and that interfaces to a provider edge (PE) device. CE devices do not recognize associated MPLS VPNs.

eBGP —exterior Border Gateway Protocol. A BGP between devices located within different autonomous systems. When two devices, located in different autonomous systems, are more than one hop away from one another, the eBGP session between the two devices is considered a multihop BGP.

iBGP —interior Border Gateway Protocol. A BGP between devices within the same autonomous system.

LFIB —Label Forwarding Information Base. Data structure used in MPLS to hold information about incoming and outgoing labels and associated Forwarding Equivalence Class (FEC) packets.

MPLS—Multiprotocol Label Switching. The name of the IETF working group responsible for label switching, and the name of the label switching approach it has standardized.

PE device—provider edge device. A device that is part of a service provider's network. It is connected to a customer edge (CE) device and all MPLS VPN processing occurs in the PE device.

RD—route distinguisher. An 8-byte value that is concatenated with an IPv4 prefix to create a unique VPN-IPv4 prefix.

VPN—Virtual Private Network. A secure MPLS-based network that shares resources on one or more physical networks (typically implemented by one or more service providers). A VPN contains geographically dispersed sites that can communicate securely over a shared backbone network.

VRF—VPN routing and forwarding instance. Routing information that defines a Virtual Private Network (VPN) site that is attached to a provider edge (PE) device. A VRF consists of an IP routing table, a derived forwarding table, a set of interfaces that use the forwarding table, and a set of rules and routing protocols that determine what goes into the forwarding table.