



Directing MPLS VPN Traffic Using a Source IP Address

This module explains how to set up an interface on a provider edge (PE) router to route packets to different Multiprotocol Label Switching (MPLS) Virtual Private Networks (VPNs) based on the source IP address of the packet.

Module History

This module was first published on May 2, 2005, and last updated on May 2, 2005.

Finding Feature Information in This Module

Your Cisco IOS software release may not support all features. To find information about feature support and configuration, use the [“Feature Information for Directing MPLS VPN Traffic Using a Source IP Address”](#) section on page 16.

Contents

- [Prerequisites for Directing MPLS VPN Traffic Using a Source IP Address, page 2](#)
- [Restrictions for Directing MPLS VPN Traffic Using a Source IP Address, page 2](#)
- [Information About Directing MPLS VPN Traffic Using a Source IP Address, page 4](#)
- [How to Enable MPLS VPN Traffic To Be Routed Using a Source IP Address, page 9](#)
- [Configuration Examples for Directing MPLS VPN Traffic Using a Source IP Address, page 13](#)
- [Additional References, page 15](#)
- [Feature Information for Directing MPLS VPN Traffic Using a Source IP Address, page 16](#)



Americas Headquarters:
Cisco Systems, Inc., 170 West Tasman Drive, San Jose, CA 95134-1706 USA

© 2007 Cisco Systems, Inc. All rights reserved.

Prerequisites for Directing MPLS VPN Traffic Using a Source IP Address

- MPLS VPNs must be enabled in the provider network.
- Cisco Express Forwarding (CEF) must be enabled on any interfaces that have this feature enabled.
- The Cisco IOS software must support MPLS VPNs, and the provider network must have MPLS Label Distribution Protocol (LDP) installed and running.
- This feature is supported on the Cisco 7200 series, 7500 series, and 12000 series router platforms.

Restrictions for Directing MPLS VPN Traffic Using a Source IP Address

VRF Select is supported only in Service Provider (-p-) images.

Unidirectional Traffic

This is a unidirectional feature and can only be used from a customer (IP-based) network into a provider (MPLS-based) network. This feature cannot be used from a provider network to a customer network.

Subnet Masks

Subnet masks should be kept as short as possible for Engine 2 line cards. Performance can degrade with longer subnet masks (/24 or /32, for example).

traceroute Command

An IP **traceroute** command from a customer edge (CE) router that has this feature enabled to a typical MPLS VPN VRF CE router works as expected. However, an IP **traceroute** command from a typical MPLS VPN VRF CE router to a CE router that has this feature enabled may fail to show all the relevant hop information across the core.

Supported Static Route Configurations

When configuring static routes in an MPLS or MPLS VPN environment, some variations of the **ip route** and **ip route vrf** commands are not supported. These variations of the commands are not supported in Cisco IOS releases that support the Tag Forwarding Information Base (TFIB), specifically Cisco IOS Releases 12.xT, 12.xM, and 12.0S. The TFIB cannot resolve prefixes when the recursive route over which the prefixes travel disappears and then reappears. However, the command variations are supported in Cisco IOS releases that support the MPLS Forwarding Infrastructure (MFI), specifically Cisco IOS Release 12.2(25)S and later. Use the following guidelines when configuring static routes.

Supported Static Routes in an MPLS Environment

The following **ip route** command is supported when you configure static routes in MPLS environment:

ip route *destination-prefix mask interface next-hop-address*

The following **ip route** commands are supported when you configure static routes in an MPLS environment and configure load sharing with static nonrecursive routes and a specific outbound interface:

ip route *destination-prefix mask interface1 next-hop1*

ip route *destination-prefix mask interface2 next-hop2*

Unsupported Static Routes in an MPLS Environment that Uses the TFIB

The following **ip route** command is not supported when you configure static routes in an MPLS environment:

ip route *destination-prefix mask next-hop-address*

The following **ip route** command is not supported when you configure static routes in an MPLS environment and enable load sharing where the next hop can be reached through two paths:

ip route *destination-prefix mask next-hop-address*

The following **ip route** command is not supported when you configure static routes in an MPLS environment and enable load sharing where the destination can be reached through two next hops:

ip route *destination-prefix mask next-hop1*

ip route *destination-prefix mask next-hop2*

Use the *interface* and *next-hop* arguments when specifying static routes.

Supported Static Routes in an MPLS VPN Environment

The following **ip route vrf** commands are supported when you configure static routes in a MPLS VPN environment, and the next hop and interface are in the same VRF:

- **ip route vrf** *vrf-name destination-prefix mask next-hop-address*
- **ip route vrf** *vrf-name destination-prefix mask interface next-hop-address*
- **ip route vrf** *vrf-name destination-prefix mask interface1 next-hop1*
- **ip route vrf** *vrf-name destination-prefix mask interface2 next-hop2*

The following **ip route vrf** commands are supported when you configure static routes in a MPLS VPN environment, and the next hop is in the global table in the MPLS cloud in the global routing table. For example, these commands are supported when the next hop is pointing to the Internet Gateway.

- **ip route vrf** *vrf-name destination-prefix mask next-hop-address global*
- **ip route vrf** *vrf-name destination-prefix mask interface next-hop-address*
(This command is supported when the next hop and interface are in the core.)

The following **ip route** commands are supported when you configure static routes in a MPLS VPN environment and enable load sharing with static nonrecursive routes and a specific outbound interfaces:

ip route *destination-prefix mask interface1 next-hop1*

ip route *destination-prefix mask interface2 next-hop2*

Unsupported Static Routes in an MPLS VPN Environment that Uses the TFIB

The following **ip route** command is not supported when you configure static routes in a MPLS VPN environment, the next hop is in the global table in the MPLS cloud within the core, and you enable load sharing where the next hop can be reached through two paths:

ip route vrf *destination-prefix mask next-hop-address global*

The following **ip route** commands are not supported when you configure static routes in a MPLS VPN environment, the next hop is in the global table in the MPLS cloud within the core, and you enable load sharing where the destination can be reached through two next hops:

ip route vrf *destination-prefix mask next-hop1 global*

ip route vrf *destination-prefix mask next-hop2 global*

The following **ip route vrf** commands are not supported when you configure static routes in an MPLS VPN environment, and the next hop and interface are in the same VRF:

```
ip route vrf vrf-name destination-prefix mask next-hop1  
ip route vrf vrf-name destination-prefix mask next-hop2
```

Supported Static Routes in an MPLS VPN Environment Where the Next Hop Resides in the Global Table on the CE Router

The following **ip route vrf** command is supported when you configure static routes in a MPLS VPN environment, and the next hop is in the global table on the CE side. For example, the following command is supported when the destination-prefix is the CE router's loopback address, as in EBGp multihop cases.

```
ip route vrf vrf-name destination-prefix mask interface next-hop-address
```

The following **ip route** commands are supported when you configure static routes in a MPLS VPN environment, the next hop is in the global table on the CE side, and you enable load sharing with static non-recursive routes and a specific outbound interfaces:

```
ip route destination-prefix mask interface1 nexthop1  
ip route destination-prefix mask interface2 nexthop2
```

Information About Directing MPLS VPN Traffic Using a Source IP Address

Before configuring this feature, you should understand the following concepts:

- [Introduction to Directing MPLS VPN Traffic Using a Source IP Address, page 4](#)
- [How MPLS VPN Traffic Is Routed Using the Source IP Address, page 4](#)
- [Example of MPLS VPN Traffic Being Routed Based on the Source IP Address, page 5](#)
- [Advantages of Using the Source IP Address over Per-Interface IP VPN Configuration, page 8](#)
- [Benefits of Directing MPLS VPN Traffic Using a Source IP Address, page 9](#)

Introduction to Directing MPLS VPN Traffic Using a Source IP Address

This feature allows packets arriving on an interface to be switched into the appropriate VRF table based upon the source IP address of the packets. Once the packets have been “selected” into the correct VRF routing table, they are processed normally based upon the destination address and forwarded through the rest of the MPLS VPN.

In most cases, this is a “one way” feature; it works on packets coming from the end users to the PE router.

How MPLS VPN Traffic Is Routed Using the Source IP Address

This feature uses the following process to route packets from the customer networks to the PE router and into the provider network.

A two-table lookup mechanism is used at the ingress interface of the PE router to determine the routing and forwarding of packets coming from the customer networks, which use IP protocols, to the MPLS VPN networks, which use MPLS protocols.

- The first table, the VRF Selection table, is used to compare the source IP address of the packet with a list of IP addresses in the table. Each IP address in the table is associated with an MPLS VPN. If a match is found between the source IP address of the packet and an IP address in the VRF Selection table, the packet is routed to the second table (the VRF table) or the routing table for the appropriate VPN.

If no match is found in the table for the source IP address of the packet, the packet will either be routed via the global routing table used by the PE router (this is the default behavior), or will be dropped. See the [“Configuring a VRF to Eliminate Unnecessary Packet Forwarding: Example” section on page 14](#) for more information.

- The second table, the VRF table (also known as the VPN routing and forwarding table), contains the virtual routing and forwarding information for the specified VPN and is used to forward the selected VPN traffic to the correct MPLS label switched path (LSP) based upon the destination IP address of the packet.

The VRF Selection process removes the association between the VRF and the interface and allows more than one MPLS VPN VRF to be associated with the interface.

Example of MPLS VPN Traffic Being Routed Based on the Source IP Address

An example of this feature is a network carrier that allows subscribers to the carrier to choose from multiple Internet service providers (ISPs) for Internet access. [Figure 1](#) provides an example of this feature with an IP-based Host network, an MPLS VPN network, and three ISPs connected to the MPLS VPN network.

Figure 1 **Implementation Example**



In [Figure 1](#), Carrier X represents the network carrier; Host A, Host B and Host C represent the carrier subscribers; and ISP 1, ISP 2 and ISP 3 represent the ISPs.

- PE2 acts as both a VRF selector and a typical MPLS VPN PE router to CE2 and CE3.
- ISPs 1 through 3 provide a list of IP addresses to Carrier X so that each host in the “POOL” network can be properly addressed. This host addressing would most likely be done by using the DHCP or DNS services of Carrier X.

A dashed line represents the path of a packet traveling from Host A to ISP 1. Host A chooses ISP 1 to use as its ISP. Carrier X provides an IP address to Host A that falls within the range of the ISP 1 registered network addresses (1.1.0.0/16). Based upon this IP address allocation, the VRF Selection criteria is set.

By using default routes, hosts on the POOL network (such as Host A), forward traffic from the Carrier X IP-based (POOL) network to the Carrier X MPLS-based VPN network. PE2 has been configured with this feature. Therefore, the MPLS VPN network forwards the traffic from Host A to ISP 1.

This is a one-way (unidirectional) feature in most implementations; it only works on packets coming from the customer networks to a PE router. Traffic coming from the ISPs to the hosts (in the example, traffic traveling from the ISPs on the right to the hosts on the left) is not affected by this feature and does not have to be returned via an MPLS path. This traffic can return via the shortest available IP path.

Another example is a Cable Modem Termination System (CMTS). If the owner of the CMTS wants to allow cable modem subscribers to choose their ISP from a group of ISPs, this feature provides a fast and scalable solution.

MPLS VPN Traffic Is Unidirectional

In [Figure 1](#), the end users are typical Internet home users. If this were a two-way (bidirectional) feature, traffic coming from the ISPs to the hosts would be required to use only the PE routers that have this feature enabled, which might cause performance issues.

When traffic from the POOL network goes through the Carrier network to the ISP networks for Internet access, the traffic in the Carrier network must be forwarded using MPLS VPN paths, because the router has “selected” the traffic into the correct MPLS VPN.

Traffic from the ISP networks to the POOL network does not have to use MPLS VPN paths in the Carrier network and can use any path that is most efficient to return to the POOL network. This traffic can use a path that uses either MPLS or IP for routing and forwarding and does not have to travel via an MPLS VPN.

Traffic from the ISP networks to the POOL networks can be forwarded using the global routing table used by every interface. One way to accomplish this is to enter VRF static routes on the PE router interfaces connected to the ISPs. The VRF static routes would route traffic from the ISPs to the Carrier network. See the [“Establishing IP Static Routes for a VRF Instance”](#) section on [page 11](#) for information on placing a default VRF static route onto an interface.

Establishing static VRF routes allows traffic from the ISPs to enter the Carrier network as traffic that can only be routed by using the global routing table toward the POOL network.

If the ISPs do not provide global host address space, or this feature is not being used to route Internet traffic, the PE interfaces connected to the ISPs must be placed into a VRF. If the PE interfaces are using VRFs for routing traffic from the ISPs, all traffic from the ISPs to the hosts through the Carrier network would be forwarded using MPLS VPN paths, and performance would not be as optimal as if IP forwarding was used.

Normal IP-based VPN operations, such as populating the Routing Information Base (RIB) and Forwarding Information Base (FIB) from a routing protocol such as Border Gateway Protocol (BGP), are used to route and forward packets within the various VPNs in the customer networks. The provider network uses MPLS-based routing protocols to perform VPN routing and forwarding inside the provider network.

Conditions That Cause MPLS VPN Traffic To Become Bidirectional

Forwarding of traffic from the Carrier network to the POOL network by using the global routing table is only possible if the ISPs have provided registered IP address space for all of the subscribed users within the POOL network from the global routing table.


If the POOL network uses IP addresses that are not globally routeable and are designed for a nonconnected enterprise (defined by RFC 1918), this feature becomes bidirectional. All traffic being sent and received by the host would have to travel via a router that has this feature enabled. The POOL network cannot be addressed with overlapping address space, regardless of the type of address space being used.

Advantages of Using the Source IP Address over Per-Interface IP VPN Configuration

This feature removes the association between a VPN and an interface. Before this feature was introduced, the following implementation was used to route outgoing MPLS VPN packets to different destinations:

- A policy-based router (PBR) is attached to the CE router.
- The egress side of the PBR router side has VLANs connected to a PE.
- The PBR router uses a policy-based route map to select the correct output (VLAN) interface and each VLAN is under a specific VRF. [Figure 2](#) illustrates a sample configuration of using a PBR router for routing MPLS packets to different destinations.

Figure 2 Implementation of Multiple VPNs



72516

The following limitations apply to PBR-based solutions that use this implementation:

- Policy routing and MPLS VPN functions cannot be performed on the same platform. Integration into a single platform is critical for manageability and support.
- VRF is limited to one VPN per interface, which limits scalability.
- The Cisco 7500 series router is used for the PBR, which can limit network performance.
- There is no network redundancy.
- The PBR is the only point of connection for all the networks attached to the PBR. The capacity and the performance capabilities of the PBR router are critical.
- There is no diversity in the connectivity to the networks.
- Every network is required to connect to every PBR. If every network is not connected to every PBR, packets from the end user to the PBR would be dropped because the PBR would have no way of switching the IP traffic properly.
- Adding multiple PBRs that are interconnected introduces more network policy-routed hops.

This feature addresses the limitations of and problems with using a PBR for packet routing and forwarding.

Benefits of Directing MPLS VPN Traffic Using a Source IP Address

Association of VPN to interface is removed

This feature removes the association between a VPN and an interface, thus allowing packets from the Host network to the provider network to have more than one VPN available per interface.

Access to every customer network is possible from every PE router in the provider network

Access points to each network can be established at any MPLS PE router, and can be made redundant by connections to multiple PE routers (for example, the CE2 router in [Figure 1 on page 6](#)).

Multiple points in the provider network can be used for VPN routing and forwarding

MPLS VPNs, like IP, are connectionless. Any PE router can carry MPLS VPN traffic from the MPLS network out to the CE routers.

How to Enable MPLS VPN Traffic To Be Routed Using a Source IP Address

This section includes the following tasks:

- [Enabling Routing of MPLS VPN Traffic Based on the Source IP Address, page 9](#) (required)
- [Establishing IP Static Routes for a VRF Instance, page 11](#) (optional)

Enabling Routing of MPLS VPN Traffic Based on the Source IP Address

Perform the following steps to enable MPLS VPN traffic to be routed based on the source IP address.

SUMMARY STEPS

1. **enable**
2. **configure terminal**
3. **vrf selection source** *source-IP-address source-IP-mask* **vrf** *vrf_name*
4. **interface** *type number*
5. **ip vrf select source**
6. **ip vrf receive** *vrf_name*
7. **end**
8. **show ip route vrf**
9. **show ip vrf select**

DETAILED STEPS

| | Command or Action | Purpose |
|--------|--|---|
| Step 1 | enable Example: Router> enable | Enables privileged EXEC mode. <ul style="list-style-type: none"> Enter your password if prompted. |
| Step 2 | configure terminal Example: Router# configure terminal | Enters global configuration mode. |
| Step 3 | vrf selection source <i>source-IP-address source-IP-mask</i> vrf <i>vrf_name</i> Example: Router(config)# vrf selection source 16.16.0.0 255.255.0.0 vrf vpn1 | Populates a source IP address to a VRF selection table. |
| Step 4 | interface <i>type number</i> Example: Router(config)# interface FastEthernet 0/1 | Configures an interface and enters interface configuration mode. |
| Step 5 | ip vrf select source Example: Router(config-if)# ip vrf select source | Enables an interface to direct MPLS VPN traffic based on the source IP address of the packet. |
| Step 6 | ip vrf receive <i>vrf_name</i> Example: Router(config-if)# ip vrf receive vpn1 | Adds all the IP addresses that are associated with an interface into a VRF table. |

| | Command or Action | Purpose |
|--------|--|---|
| Step 7 | end Example: Router(config-if)# end | (Optional) Exits to privileged EXEC mode. |
| Step 8 | show ip route vrf Example: Router# show ip route vrf | Displays the IP routing table associated with a VRF instance. Use this command to verify the configuration. |
| Step 9 | show ip vrf select Example: Router# show ip vrf select | Displays information about the VRF selection. |

Establishing IP Static Routes for a VRF Instance

Traffic coming from the ISPs to the hosts does not require the use of the MPLS VPN paths; this traffic can use the shortest IP route back to the host.

VPN static routes for traffic returning to the customer networks are only necessary if VPN traffic returning to the customer networks is being forwarded back from the enabled interface. The remote PE router could also be configured to route return traffic to the customer networks directly by using the global routing table.

SUMMARY STEPS

1. **enable**
2. **configure terminal**
3. **interface** *type number*
4. **ip route vrf** *vrf_name prefix mask [next-hop-address] [interface {interface-number}] [global] [distance] [permanent] [tag tag]*

| | Command | Purpose |
|--------|--|---|
| Step 1 | enable Example: Router> enable | Enables privileged EXEC mode. <ul style="list-style-type: none">• Enter your password if prompted. |
| Step 2 | configure terminal Example: Router# configure terminal | Enters global configuration mode. |

| | Command | Purpose |
|--------|---|--|
| Step 3 | interface <i>type number</i> Example: Router(config)# interface FastEthernet 0/1 | Configures an interface and enters interface configuration mode. |
| Step 4 | ip route vrf <i>vrf_name prefix mask [next-hop-address] [interface {interface-number}] [global] [distance] [permanent] [tag tag]</i> Example: Router(config-if)# ip route vrf vpn1 16.16.0.0 255.255.0.0 POS1/0 | Establishes static routes for a VRF. |

Troubleshooting Tips

- Enter the **debug vrf select** command to enable debugging for this feature.



Note The **debug vrf select** command can cause many messages to be logged when you change the configuration and when switching occurs.

- The following error messages appear if problems occur while configuring this feature:
 - If you attempt to configure a nonexistent VRF Selection table:


```
Router(config)# vrf selection source 2.0.0.0 255.255.0.0 vrf VRF_NOEXIST
VRF Selection: VRF table VRF_NOEXIST does not exist.
```
 - If you attempt to remove a VRF Selection entry that does not exist:


```
Router(config)# no vrf selection source 2.0.0.0 255.255.0.0 vrf VRF1
VRF Selection: Can't find the node to remove.
```
 - If you attempt to configure a duplicate IP address and subnet mask for a VRF Selection entry:


```
Router(config)# vrf selection source 2.0.0.0 255.0.0.0 vrf VRF_AOL
Router(config)# vrf selection source 2.0.0.0 255.0.0.0 vrf VRF_AOL
VRF Selection: duplicate address and mask configured.
```
 - If an inconsistent IP address and mask are used for a VRF Selection entry:


```
Router(config)# vrf selection source 170.1.2.1 255.255.255.0 vrf red
% Inconsistent address and mask
Router(config)# vrf selection source 170.1.2.1 255.255.255.255 vrf red
```
 - If you attempt to configure a VRF instance on an interface that has this feature already configured:


```
Router(config-if)# ip vrf select source
Router(config-if)# ip vrf forward red
% Can not configure VRF if VRF Select is already configured
To enable VRF, first remove VRF Select from the interface
```
 - If you attempt to configure an entry on an interface that has this feature already configured:


```
Router(config-if)# ip vrf forward red
```

```
Router(config-if)# ip vrf select source
% Can not configure VRF Select if interface is under a non-global VRF
To enable VRF Select, first remove VRF from the interface
```

Configuration Examples for Directing MPLS VPN Traffic Using a Source IP Address

This section provides the following configuration examples:

- [Enabling MPLS VPN Traffic To Be Routed Based on Source IP Address: Example, page 13](#)
- [Configuring a VRF to Eliminate Unnecessary Packet Forwarding: Example, page 14](#)
- [Verifying the Configuration: Example, page 14](#)

Enabling MPLS VPN Traffic To Be Routed Based on Source IP Address: Example

The following example defines two entries (vpn1 and vpn2) in the VRF Selection table. In this example, packets with the source address of 16.16.0.0 will be routed to the VRF vpn1, and packets with the source address of 17.17.0.0 will be routed to the VRF vpn2:

```
Router(config)# vrf selection source 16.16.0.0 255.255.0.0 vrf vpn1
Router(config)# vrf selection source 17.17.0.0 255.255.0.0 vrf vpn2
```

The following example creates IP static routes for two VRFs (vpn1 and vpn2) for the POS1/0 interface:

```
Router(config)# ip route vrf vpn1 16.16.0.0 255.255.0.0 POS1/0
Router(config)# ip route vrf vpn2 17.17.0.0 255.255.0.0 POS1/0
```

The following example configures the POS1/0 interface for this feature and adds the configured IP address (31.0.0.1) to the VRFs vpn1 and vpn2 as connected routes.

```
Router(config)# interface POS1/0
Router(config-if)# description Link to CE1 POS1/0 (eng2)
Router(config-if)# ip vrf select source
Router(config-if)# ip vrf receive vpn1
Router(config-if)# ip vrf receive vpn2
Router(config-if)# ip address 31.0.0.1 255.0.0.0
Router(config-if)# no ip directed-broadcast
Router(config-if)# load-interval 30
Router(config-if)# crc 32
Router(config-if)# end
```

Configuring a VRF to Eliminate Unnecessary Packet Forwarding: Example

If a packet arrives at an interface that has VRF Select enabled, and its source IP address does not match any VRF Select definition, that packet will be forwarded via the global routing table. This default behavior could cause problems if IP address spoofing is being implemented. Unnecessary traffic could be forwarded via the global routing table. To eliminate this unnecessary routing of packets, create a VRF Selection definition that will forward all unknown incoming traffic to a null interface.

The following configuration causes all traffic not matching a more specific VRF Selection definition to be routed to the Null0 interface, thus dropping the packets.

```
Router(config)# ip vrf VRF_DROP
Router(config-vrf)# rd 999:99
Router(config-vrf)# route-target export 999:99
Router(config-vrf)# route-target import 999:99
Router(config-vrf)# exit

Router(config)# vrf selection source 0.0.0.0 0.0.0.0 vrf VRF_DROP

Router(config)# ip route vrf VRF_DROP 0.0.0.0 0.0.0.0 Null0
```

Verifying the Configuration: Example

This example shows the IP routing table associated with the VRF vrf1:

```
Router# show ip route vrf vpn1

Routing Table: vpn1

Codes: C - connected, S - static, I - IGRP, 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, E - EGP
       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

Gateway of last resort is not set

B    33.0.0.0/8 [200/0] via 10.10.10.10, 00:00:37
     5.0.0.0/16 is subnetted, 1 subnets
B       5.19.0.0 [200/0] via 10.10.10.10, 00:00:37
     14.0.0.0/32 is subnetted, 1 subnets
B       14.14.14.14 [200/0] via 10.10.10.10, 00:00:37
     15.0.0.0/32 is subnetted, 1 subnets
S       15.15.15.15 [1/0] via 34.0.0.1, POS1/1
```

Additional References

Related Documents

| Related Topic | Document Title |
|-------------------------------------|---|
| Basic MPLS VPNs | <i>Configuring MPLS Layer 3 VPNs</i> |
| MPLS VPN Carrier Supporting Carrier | <ul style="list-style-type: none"> • <i>MPLS VPN Carrier Supporting Carrier Using LDP and an IGP</i> • <i>MPLS VPN Carrier Supporting Carrier with BGP</i> |
| MPLS VPN InterAutonomous Systems | <ul style="list-style-type: none"> • <i>MPLS VPN Inter-AS with ASBRs Exchanging IPv4 Routes and MPLS Labels</i> • <i>MPLS VPN Inter-AS with ASBRs Exchanging VPN-IPv4 Addresses</i> |

Standards

| Standard | Title |
|---|-------|
| No new or modified standards are supported by this feature, and support for existing standards has not been modified by this feature. | — |

MIBs

| MIB | MIBs Link |
|---|--|
| No new or modified MIBs are supported by this feature, and support for existing MIBs has not been modified by this feature. | <p>To obtain lists of supported MIBs by platform and Cisco IOS release, and to download MIB modules, go to the Cisco MIB website on Cisco.com at the following URL:</p> <p>http://www.cisco.com/public/sw-center/netmgmt/cmtk/mibs.shtml</p> |

RFCs

| RFC | Title |
|---|-------|
| No new or modified RFCs are supported by this feature, and support for existing RFCs has not been modified by this feature. | — |

Technical Assistance

| Description | Link |
|---|---|
| The Cisco Technical Support website contains thousands of pages of searchable technical content, including links to products, technologies, solutions, technical tips, and tools. Registered Cisco.com users can log in from this page to access even more content. | http://www.cisco.com/techsupport |

Feature Information for Directing MPLS VPN Traffic Using a Source IP Address

[Table 1](#) lists the features in this module and provides links to specific configuration information.

Not all commands may be available in your Cisco IOS software release. For details on when support for specific commands was introduced, see the command reference documents.

Cisco IOS software images are specific to a Cisco IOS software release, a feature set, and a platform. Use Cisco Feature Navigator to find information about platform support and Cisco IOS software image support. Access Cisco Feature Navigator at <http://www.cisco.com/go/fn>. You must have an account on Cisco.com. If you do not have an account or have forgotten your username or password, click **Cancel** at the login dialog box and follow the instructions that appear.



Note

[Table 1](#) lists only the Cisco IOS software release that introduced support for a given feature in a given Cisco IOS software release train. Unless noted otherwise, subsequent releases of that Cisco IOS software release train also support that feature.

Table 1 **Feature Information for Directing MPLS VPN Traffic Using a Source IP Address**

| Feature Name | Releases | Feature Configuration Information |
|--|--|--|
| VRF Selection Based on Source IP Address | 12.0(22)S 12.0(23)S 12.0(24)S 12.0(26)S | <p>This feature lets you direct MPLS VPN traffic based on the source IP address of the packet.</p> <p>The following sections provide information about this feature:</p> <ul style="list-style-type: none"> • Introduction to Directing MPLS VPN Traffic Using a Source IP Address, page 4 • How MPLS VPN Traffic Is Routed Using the Source IP Address, page 4 • Example of MPLS VPN Traffic Being Routed Based on the Source IP Address, page 5 • Advantages of Using the Source IP Address over Per-Interface IP VPN Configuration, page 8 • Benefits of Directing MPLS VPN Traffic Using a Source IP Address, page 9 • How to Enable MPLS VPN Traffic To Be Routed Using a Source IP Address, page 9 |

CCVP, the Cisco logo, and Welcome to the Human Network are trademarks of Cisco Systems, Inc.; Changing the Way We Work, Live, Play, and Learn is a service mark of Cisco Systems, Inc.; and Access Registrar, Aironet, Catalyst, CCDA, CCDP, CCIE, CCIP, CCNA, CCNP, CCSP, Cisco, the Cisco Certified Internetwork Expert logo, Cisco IOS, Cisco Press, Cisco Systems, Cisco Systems Capital, the Cisco Systems logo, Cisco Unity, Enterprise/Solver, EtherChannel, EtherFast, EtherSwitch, Fast Step, Follow Me Browsing, FormShare, GigaDrive, HomeLink, Internet Quotient, IOS, iPhone, IP/TV, iQ Expertise, the iQ logo, iQ Net Readiness Scorecard, iQuick Study, LightStream, Linksys, MeetingPlace, MGX, Networkers, Networking Academy, Network Registrar, PIX, ProConnect, ScriptShare, SMARTnet, StackWise, The Fastest Way to Increase Your Internet Quotient, and TransPath are registered trademarks of Cisco Systems, Inc. and/or its affiliates in the United States and certain other countries.

All other trademarks mentioned in this document or Website are the property of their respective owners. The use of the word partner does not imply a partnership relationship between Cisco and any other company. (0711R)

Any Internet Protocol (IP) addresses used in this document are not intended to be actual addresses. Any examples, command display output, and figures included in the document are shown for illustrative purposes only. Any use of actual IP addresses in illustrative content is unintentional and coincidental.

© 2007 Cisco Systems, Inc. All rights reserved.

