Ensuring That MPLS VPN Clients Using OSPF Communicate over the MPLS VPN Backbone Instead of Through Backdoor Links

This module describes how to configure a sham-link that ensures traffic travels between Virtual Private Network (VPN) client sites over the Multiprotocol Label Switching (MPLS) VPN backbone. This feature is for VPNs that run Open Shortest Path First (OSPF) between the provider edge (PE) and customer edge (CE) routers. By default, OSPF uses backdoor paths between VPN sites, not the MPLS VPN backbone.

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 Ensuring MPLS VPN Clients Communicate over the MPLS VPN Backbone" section on page 13.

Contents

- Prerequisites for Ensuring MPLS VPN Clients Communicate over the MPLS VPN Backbone, page 2
- Restrictions for Ensuring MPLS VPN Clients Communicate over the MPLS VPN Backbone, page 2
- Information About Ensuring MPLS VPN Clients Communicate over the MPLS VPN Backbone, page 2
- How to Ensure That MPLS VPN Clients Communicate over the MPLS VPN Backbone, page 7
- Configuration Examples for Ensuring MPLS VPN Clients Communicate over the MPLS VPN Backbone, page 9
- Additional References, page 12
- Feature Information for Ensuring MPLS VPN Clients Communicate over the MPLS VPN Backbone, page 13

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

Prerequisites for Ensuring MPLS VPN Clients Communicate over the MPLS VPN Backbone

Before you can configure a sham-link in an MPLS VPN, you must first enable OSPF as follows:

- Create an OSPF routing process.
- Specify the range of IP addresses to be associated with the routing process.
- Assign area IDs to be associated with the range of IP addresses.

Restrictions for Ensuring MPLS VPN Clients Communicate over the MPLS VPN Backbone

When OSPF is used as a protocol between PE and CE routers, the OSPF metric is preserved when routes are advertised over the VPN backbone. The metric is used on the remote PE routers to select the correct route. For this reason, you should not modify the metric value when OSPF is redistributed to Border Gateway Protocol (BGP), and when BGP is redistributed to OSPF. If you modify the metric value, routing loops may occur.

Information About Ensuring MPLS VPN Clients Communicate over the MPLS VPN Backbone

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

- Introduction to MPLS VPNs Using OSPF Between PE and CE Routers, page 2
- OSPF Uses Backdoor Paths to Communicate Between VPN Sites, page 3
- Sham-Links Direct Traffic Between VPN Sites over the MPLS VPN Backbone, page 5

Introduction to MPLS VPNs Using OSPF Between PE and CE Routers

In an MPLS VPN configuration, the OSPF protocol is one way you can connect CE routers to PE routers in the VPN backbone. OSPF is often used by customers that run OSPF as their intrasite routing protocol, subscribe to a VPN service, and want to exchange routing information between their sites using OSPF (during migration or on a permanent basis) over an MPLS VPN backbone.

Figure 1 shows an example of how VPN client sites (areas 0, 1, 2, and 3) that run OSPF can connect over an MPLS VPN backbone.

Figure 1 OSPF Connectivity Between VPN Client Sites and an MPLS VPN Backbone

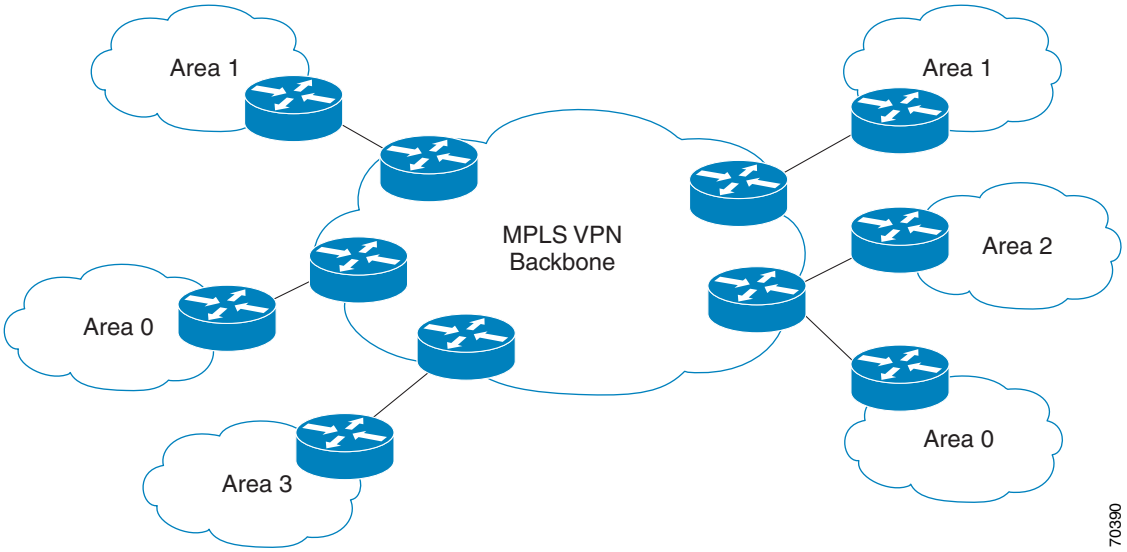

When OSPF is used to connect PE and CE routers, all routing information learned from a VPN site is placed in the VPN routing and forwarding (VRF) instance associated with the incoming interface. The PE routers that attach to the VPN use the BGP to distribute VPN routes to each other. A CE router can then learn the routes to other sites in the VPN by peering with its attached PE router. The MPLS VPN backbone provides an additional level of routing hierarchy to interconnect the VPN sites running OSPF.

When OSPF routes are propagated over the MPLS VPN backbone, additional information about the prefix in the form of BGP extended communities (route type, domain ID extended communities) is appended to the BGP update. This community information is used by the receiving PE router to decide the type of link-state advertisement (LSA) to be generated when the BGP route is redistributed to the OSPF PECE process. In this way, internal OSPF routes that belong to the same VPN and are advertised over the VPN backbone are seen as interarea routes on the remote sites.

OSPF Uses Backdoor Paths to Communicate Between VPN Sites

Although OSPF PECE connections assume that the only path between two client sites is across the MPLS VPN backbone, backdoor paths between VPN sites may exist. For instance, in Figure 2, Vienna, Stockholm, Brighton, and Winchester can communicate through backdoor paths instead of using the MPLS VPN backbone.

If the sites belong to the same OSPF area, the backdoor path will always be selected, because OSPF prefers intra-area paths to interarea paths. (PE routers advertise OSPF routes learned over the VPN backbone as interarea paths.) For this reason, OSPF backdoor paths between VPN sites must be taken into account so that routing is performed based on policy.

Figure 2 Backdoor Paths Between OSPF Client Sites

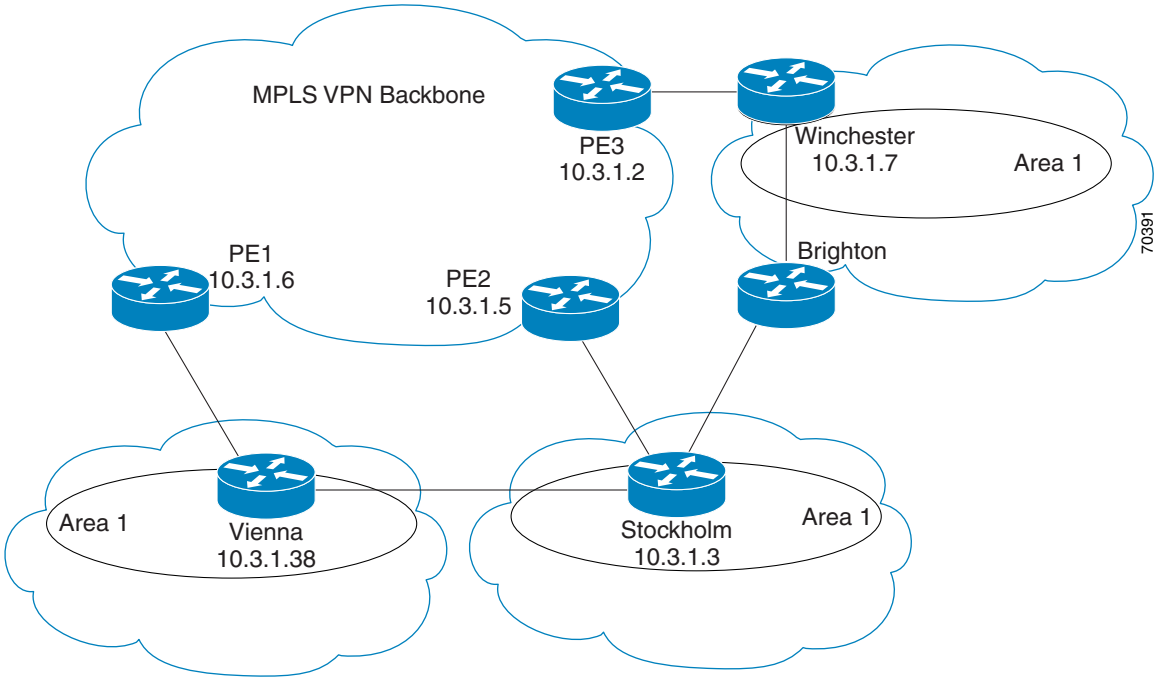

For example, Figure 2 shows three client sites, each with backdoor links. Because each site runs OSPF within the same Area 1 configuration, all routing between the three sites uses the backdoor paths, rather than the MPLS VPN backbone.

The following example shows BGP routing table entries for the Winchester router (prefix 10.3.1.7/32) from the standpoint of the PE1 router in Figure 2. Prefix 10.3.1.7 is the loopback interface of the Winchester CE router. As shown in bold in this example, the loopback interface is learned via BGP from PE2 and PE3. It is also generated through redistribution into BGP on PE1.

```
PE1# show ip bgp vpnv4 all 10.3.1.7

BGP routing table entry for 100:251:10.3.1.7/32, version 58
Paths: (3 available, best #2)
  Advertised to non peer-group peers:
  10.3.1.2 10.3.1.5
  Local
    10.3.1.5 (metric 30) from 10.3.1.5 (10.3.1.5)
      Origin incomplete, metric 22, localpref 100, valid, internal
      Extended Community: RT:1:793 OSPF DOMAIN ID:0.0.0.100 OSPF
      RT:1:2:0 OSPF 2
  Local
    10.2.1.38 from 0.0.0.0 (10.3.1.6)
      Origin incomplete, metric 86, localpref 100, weight 32768,
      valid, sourced, best
      Extended Community: RT:1:793 OSPF DOMAIN ID:0.0.0.100 OSPF
      RT:1:2:0 OSPF 2
  Local
    10.3.1.2 (metric 30) from 10.3.1.2 (10.3.1.2)
      Origin incomplete, metric 11, localpref 100, valid, internal
      Extended Community: RT:1:793 OSPF DOMAIN ID:0.0.0.100 OSPF
      RT:1:2:0 OSPF 2
```

Within BGP, the locally generated route (10.2.1.38) is considered to be the best route.

However, as shown in bold in the next example, the VRF routing table shows that the selected path is learned via OSPF with a next hop of 10.2.1.38, which is the Vienna CE router.

```
PE1# show ip route vrf ospf 10.3.1.7

Routing entry for 10.3.1.7/32
  Known via "ospf 100", distance 110, metric 86, type intra area
  Redistributing via bgp 215
  Advertised by bgp 215
  Last update from 10.2.1.38 on Serial0/0/0, 00:00:17 ago
  Routing Descriptor Blocks:
  * 10.2.1.38, from 10.3.1.7, 00:00:17 ago, via Serial0/0/0
      Route metric is 86, traffic share count is 1
```

This path is selected because:

- The OSPF backdoor path is preferred over the interarea path (over the MPLS VPN backbone) generated by the PE1 router.
- OSPF has a lower administrative distance (AD) than internal BGP (BGP running between routers in the same autonomous system).

If the backdoor paths between sites are used only for backup purposes and do not participate in the VPN service, then the default route selection is acceptable. You can set up the OSPF cost configured with a sham-link to send VPN site traffic over a backdoor path.

Sham-Links Direct Traffic Between VPN Sites over the MPLS VPN Backbone

To ensure that VPN sites that belong to the same OSPF area and share an OSPF backdoor path communicate with each other using the MPLS VPN backbone, you must create a sham-link. (If no backdoor path exists between the sites, no sham-link is required.) A sham-link is an additional OSPF intra-area (logical) link between ingress and egress VRFs on the PE routers that connect to the CE routers of the VPN sites.

Figure 3 shows a sample sham-link between PE1 and PE2. You associate a cost with each sham-link to force traffic to use the sham-link rather than the backdoor path. When a sham-link is configured between PE routers, the PE routers can populate the VRF routing table with the OSPF routes learned over the sham-link.

Figure 3 Using a Sham-Link Between PE Routers to Connect OSPF Client Sites

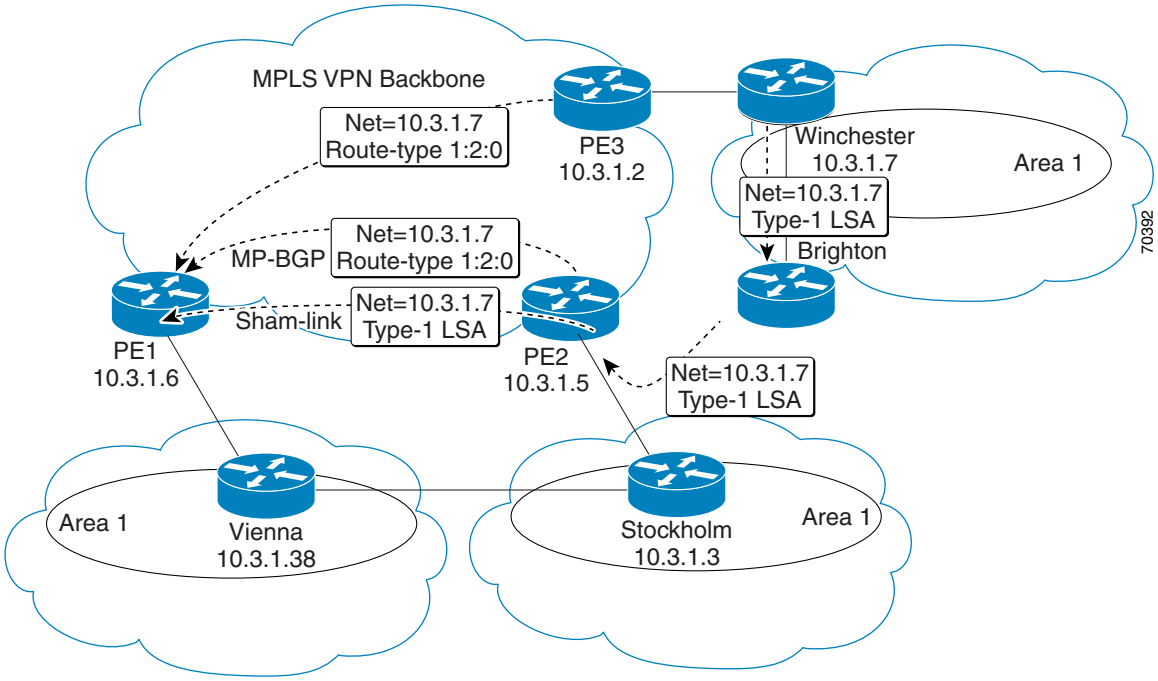

Because the sham-link is seen as an intra-area link between PE routers, an OSPF adjacency is created and database exchange (for the particular OSPF process) occurs across the link. The PE router can then flood LSAs between sites from across the MPLS VPN backbone. As a result, the desired intra-area connectivity is created.

How to Ensure That MPLS VPN Clients Communicate over the MPLS VPN Backbone

This section explains how to create a sham-link on an MPLS VPN PE router. Perform this task on both PE routers that share the sham-link.

Prerequisites

Before you create a sham-link between PE routers in an MPLS VPN, you must:

- Configure a separate /32 address on the remote PE so that OSPF packets can be sent over the VPN backbone to the remote end of the sham-link. The /32 address must meet the following criteria:
  - Belong to a VRF.
  - Not be advertised by OSPF.
  - Be advertised by BGP.

  You can use the /32 address for other sham-links.
- Associate the sham-link with an existing OSPF area.

SUMMARY STEPS

1. **enable**
2. **configure terminal**
3. **interface loopback** *interface-number*
4. **ip vrf forwarding** *vrf-name*
5. **ip address** *ip-address mask*
6. **end**
7. router ospf *process-id* vrf *vrf-name*
8. area *area-id* sham-link *source-address destination-address* cost *number*
9. show ip ospf sham-links

DETAILED STEPS

	Command or Action	Purpose
Step 1	**enable** **Example:** `Router> enable`	Enables privileged EXEC mode. • Enter your password if prompted.
Step 2	**configure terminal** **Example:** `Router# configure terminal`	Enters global configuration mode.

	Command or Action	Purpose
Step 3	**interface loopback** *interface-number* **Example:** `Router(config)# interface loopback 1`	Creates a loopback interface to be used as an endpoint of the sham-link on the PE router and enters interface configuration mode.
Step 4	**ip vrf forwarding** *vrf-name* **Example:** `Router(config-if)# ip vrf forwarding ospf`	Associates the loopback interface with a VRF. Removes the IP address.
Step 5	**ip address** *ip-address mask* **Example:** `Router(config-if)# ip address 10.2.1.2 255.255.255.255`	Reconfigures the IP address of the loopback interface on the PE router.
Step 6	**end** **Example:** `Router(config-if)# end`	Returns to global configuration mode.
Step 7	**router ospf** *process-id* **vrf** *vrf-name* **Example:** `Router(config)# router ospf 100 vrf ospf`	Configures the specified OSPF process with the VRF associated with the sham-link interface on the PE router and enters interface configuration mode.
Step 8	**area** *area-id* **sham-link** *source-address destination-address* **cost** *number* **Example:** `Router(config-if)# area 1 sham-link 10.2.1.2 10.2.1.1 cost 40`	Configures the sham-link on the PE router interface within a specified OSPF area and with the loopback interfaces specified by the IP addresses as endpoints. • **cost** *number* configures the OSPF cost for sending an IP packet over the PE sham-link interface.
Step 9	**show ip ospf sham-links**	Verifies that the sham-link was successfully created and is operational.

Example

The following is sample output from the **show ip ospf sham-links** command:

```
Router# show ip ospf sham-links

Sham Link OSPF_SL0 to address 10.2.1.2 is up
Area 1 source address 10.2.1.1
  Run as demand circuit
  DoNotAge LSA allowed.
  Cost of using 40 State POINT_TO_POINT,
  Timer intervals configured,
  Hello 10, Dead 40, Wait 40,
    Hello due in 00:00:04
    Adjacency State FULL (Hello suppressed)
    Index 2/2, retransmission queue length 4,     number of retransmission 0
    First 0x63311F3C(205)/0x63311FE4(59) Next
    0x63311F3C(205)/0x63311FE4(59)
    Last retransmission scan length is 0,     maximum is 0
    Last retransmission scan time is 0 msec,     maximum is 0 msec
    Link State retransmission due in 360 msec
```

Configuration Examples for Ensuring MPLS VPN Clients Communicate over the MPLS VPN Backbone

The following example shows how to configure a sham-link between two PE routers:

```
Router1(config)# interface loopback 1
Router1(config-if)# ip vrf forwarding ospf
Router1(config-if)# ip address 10.2.1.1 255.255.255.255
!
Router2(config)# interface loopback 1
Router2(config-if)# ip vrf forwarding ospf
Router2(config-if)# ip address 10.2.1.2 255.255.255.255
!
Router1(config)# router ospf 100 vrf ospf
Router1(config-if)# area 1 sham-link 10.2.1.1 10.2.1.2 cost 40
!
Router2(config)# router ospf 100 vrf ospf
Router2(config-if)# area 1 sham-link 10.2.1.2 10.2.1.1 cost 40
```

This example shows how a sham-link is used only to affect the OSPF intra-area path selection of the PE and CE routers. The PE router also uses the information received from Multiprotocol BGP (MP-BGP) to set the outgoing label stack of incoming packets, and to decide to which egress PE router to label-switch the packets.

Figure 4 shows a sample MPLS VPN topology in which a sham-link configuration is necessary. A VPN client has three sites, each with a backdoor path. Two sham-links have been configured, one between PE1 and PE2, and another between PE2 and PE3. A sham-link between PE1 and PE3 is not necessary in this configuration, because the Vienna and Winchester sites do not share a backdoor path.

Figure 4 Sham-Link Example

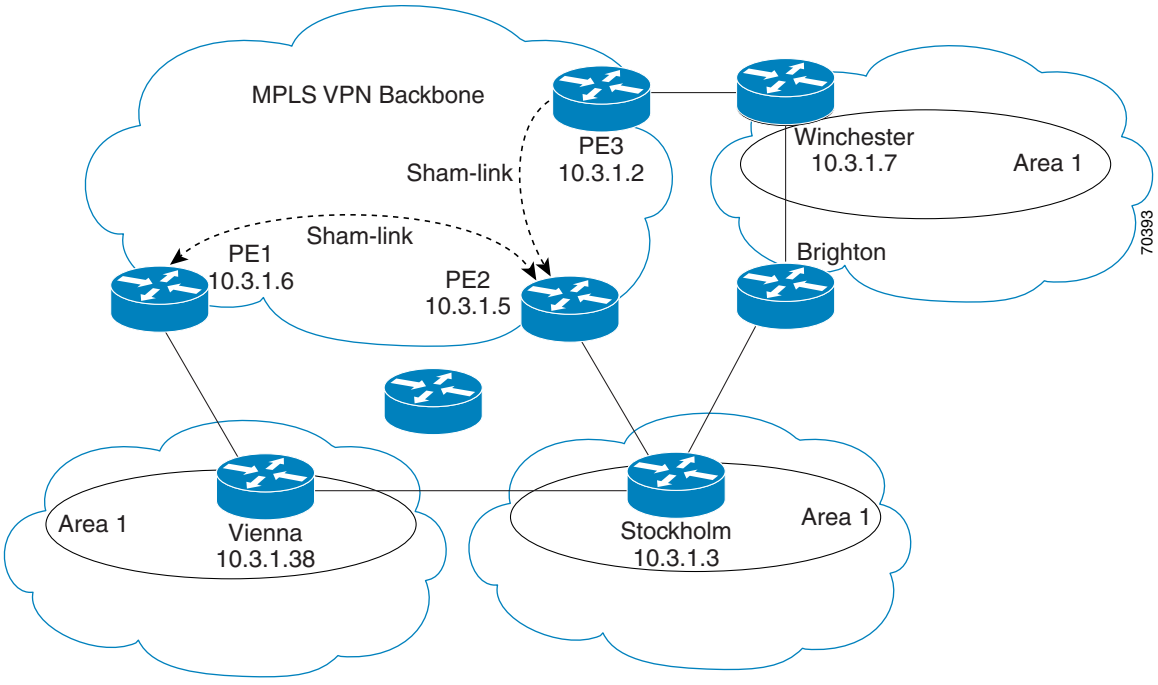

The following example shows the forwarding that occurs between sites from the standpoint of how PE1 views the 10.3.1.7/32 prefix, the loopback1 interface of the Winchester CE router in Figure 4.

```
PE1# show ip bgp vpnv4 all 10.3.1.7
BGP routing table entry for 100:251:10.3.1.7/32, version 124
Paths: (1 available, best #1)
  Local
    10.3.1.2 (metric 30) from 10.3.1.2 (10.3.1.2)
       Origin incomplete, metric 11, localpref 100, valid, internal,
       best
       Extended Community: RT:1:793 OSPF DOMAIN ID:0.0.0.100 OSPF
       RT:1:2:0 OSPF 2

PE1# show ip route vrf ospf 10.3.1.7
Routing entry for 10.3.1.7/32
  Known via "ospf 100", distance 110, metric 13, type intra area
  Redistributing via bgp 215
  Last update from 10.3.1.2 00:12:59 ago
  Routing Descriptor Blocks:
10.3.1.2 (Default-IP-Routing-Table), from 10.3.1.7, 00:12:59 ago
```

The next example shows forwarding information in which the next hop for the route, 10.3.1.2, is the PE3 router rather than the PE2 router (which is the best path according to OSPF). The OSPF route is not redistributed to BGP on the PE, because the other end of the sham-link already redistributed the route to BGP and there is no need for duplication. The OSPF sham-link is used only to influence intra-area path selection. When sending traffic to a particular destination, the PE router uses the MP-BGP forwarding information.

```
PE1# show ip bgp vpnv4 all tag | begin 10.3.1.7

   10.3.1.7/32      10.3.1.2        notag/38

PE1# show mpls forwarding 10.3.1.2

Local  Outgoing      Prefix           Bytes label    Outgoing   Next Hop
label  label or VC   or Tunnel Id     switched       interface
31     42            10.3.1.2/32      0              PO3/0/0    point2point

PE1# show ip cef vrf ospf 10.3.1.7

10.3.1.7/32, version 73, epoch 0, cached adjacency to POS3/0/0
0 packets, 0 bytes
  tag information set
   local tag: VPN-route-head
   fast tag rewrite with PO3/0/0, point2point, tags imposed: {42 38}
  via 10.3.1.2, 0 dependencies, recursive
    next hop 10.1.1.17, POS3/0/0 via 10.3.1.2/32
    valid cached adjacency
    tag rewrite with PO3/0/0, point2point, tags imposed: {42 38}
```

If a prefix is learned across the sham-link and the path via the sham-link is selected as the best, the PE router does not generate an MP-BGP update for the prefix. It is not possible to route traffic from one sham-link over another sham-link.

In the following example, PE2 shows how an MP-BGP update for the prefix is not generated. Although 10.3.1.7/32 has been learned via OSPF across the sham-link as shown in bold, no local generation of a route into BGP is performed. The only entry within the BGP table is the MP-BGP update received from PE3 (the egress PE router for the 10.3.1.7/32 prefix).

```
PE2# show ip route vrf ospf 10.3.1.7

Routing entry for 10.3.1.7/32
  Known via "ospf 100", distance 110, metric 12, type intra area
  Redistributing via bgp 215
  Last update from 10.3.1.2 00:00:10 ago
  Routing Descriptor Blocks:
  * 10.3.1.2 (Default-IP-Routing-Table), from 10.3.1.7, 00:00:10 ago
      Route metric is 12, traffic share count is 1

PE2# show ip bgp vpnv4 all 10.3.1.7

BGP routing table entry for 100:251:10.3.1.7/32, version 166
Paths: (1 available, best #1)
  Not advertised to any peer
  Local
    10.3.1.2 (metric 30) from 10.3.1.2 (10.3.1.2)
      Origin incomplete, metric 11, localpref 100, valid, internal,
      best
      Extended Community: RT:1:793 OSPF DOMAIN ID:0.0.0.100 OSPF
      RT:1:2:0 OSPF 2
```

The PE router uses the information received from MP-BGP to set the ongoing label stack of incoming packets, and to decide to which egress PE router to label-switch the packets.

Additional References

The following sections provide references related to MPLS VPNs.

Related Documents

Related Topic	Document Title
Basic MPLS VPNs	*Configuring MPLS Layer 3 VPNs*
MPLS VPN Carrier Supporting Carrier	• *MPLS VPN Carrier Supporting Carrier Using LDP and an IGP* • *MPLS VPN Carrier Supporting Carrier with BGP*
MPLS VPN InterAutonomous Systems	• *MPLS VPN Inter-AS with ASBRs Exchanging IPv4 Routes and MPLS Labels* • *MPLS VPN Inter-AS with ASBRs Exchanging VPN-IPv4 Addresses*

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.	To locate and download MIBs for selected platforms, Cisco IOS releases, and feature sets, use Cisco MIB Locator found at the following URL: http://www.cisco.com/go/mibs

RFCs

RFC	Title
RFC 1164	Application of the Border Gateway Protocol in the Internet
RFC 1171	A Border Gateway Protocol 4
RFC 1700	Assigned Numbers
RFC 1966	BGP Route Reflection: An Alternative to Full Mesh IBGP
RFC 2283	Multiprotocol Extensions for BGP-4
RFC 2328	Open Shortest Path First, Version 2
RFC 2547	BGP/MPLS VPNs

RFC	Title
RFC 2842	Capabilities Advertisement with BGP-4
RFC 2858	Multiprotocol Extensions for BGP-4
RFC 3107	Carrying Label Information in BGP-4

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 Ensuring MPLS VPN Clients Communicate over the MPLS VPN Backbone

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 Ensuring MPLS VPN Clients Communicate over the MPLS VPN Backbone

Feature Name	Releases	Feature Configuration Information
Ensuring MPLS VPN Clients Communicate over the MPLS VPN Backbone	12.2(8)T 12.0(21)ST 12.0(22)S	This feature allows you to configure a sham-link that directs traffic between Virtual Private Network (VPN) client sites over the Multiprotocol Label Switching (MPLS) VPN backbone. The following sections provide information about this feature: • Introduction to MPLS VPNs Using OSPF Between PE and CE Routers, page 2 • OSPF Uses Backdoor Paths to Communicate Between VPN Sites, page 3 • Sham-Links Direct Traffic Between VPN Sites over the MPLS VPN Backbone, page 5 • How to Ensure That MPLS VPN Clients Communicate over the MPLS VPN Backbone, page 7