Inter-AS Hybrid for MPLS VPN over IP Tunnels

Part Number OL-10922-01 (Rev B0), June 26, 2006

The Inter-AS Hybrid for MPLS VPN over IP Tunnels feature allows a Multiprotocol Label Switching Virtual Private Network (MPLS VPN) to span autonomous systems (ASs) and VPN service providers when one of the VPN autonomous systems runs Layer 2 Tunneling Protocol version 3 (L2TPv3) in an MPLS VPN over IP Tunnels network. This feature allows the MPLS VPN over IP Tunnels network to present a standard MPLS VPN Inter-AS interface to a BGP peer in an adjacent MPLS AS. The Inter-AS Hybrid implementation combines elements of the existing MPLS/IP VPN and MPLS VPN Inter-AS features described in:

- *Interautonomous System VPN—Back-to-Back VRF*
- *MPLS VPN—Interautonomous Systems Support*

Feature History for the Inter-AS Hybrid for MPLS VPN over IP Tunnels

Release	Modification
12.0(31)S	This feature was introduced on the Cisco 12000 series Router on IP Services Engine (ISE) line cards.
12.0(32)SY	Support was added for Engine 5 line cards, including shared port adapters (SPAs) and SPA Interface Processors (SIPs), on the Cisco 12000 series Router.

Finding Support Information for Platforms and Cisco IOS Software Images

Use Cisco Feature Navigator to find information about platform support and Cisco IOS software image support 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.

Contents

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

Prerequisites for Inter-AS Hybrid for MPLS VPN over IP Tunnels

- An autonomous system (or service provider network) must be properly configured for MPLS VPN and MPLS VPN over IP Tunnels operation before you configure Inter-AS Hybrid for MPLS VPN over IP Tunnels.
- Starting in IOS Release 12.0(31)S, the Inter-AS Hybrid for MPLS VPN over IP Tunnels feature is supported only on the following IP Services Engine (ISE) line cards in a Cisco 12000 series router used as an autonomous system boundary router (ASBR) in a VPN service provider network configured for MPLS VPNs over IP Tunnels:

  4-port OC-3 POS ISE
  8-port OC-3 POS ISE
  16-port OC-3 POS ISE
  4-port OC-12 POS ISE
  1-port OC-48 POS ISE
  4-port Channelized OC-12 POS ISE
  1-port Channelized OC-48 POS ISE
  4-port Gigabit Ethernet ISE
- Starting in Cisco IOS Release 12.0(32)SY, the Inter-AS Hybrid for MPLS VPN over IP Tunnels feature is supported on the following Engine 5 line cards on the Cisco 12000 series Router:

  Engine 5 shared port adapters (SPAs):

  1-port Channelized STM-1c/OC-3c to DS0
  8-port Channelized T1/E1

  1-port 10-Gigabit Ethernet
  2-port Gigabit Ethernet
  5-port Gigabit Ethernet
  10-port Gigabit Ethernet
  8-port Fast Ethernet
  8-port 10/100 Ethernet

  4-port OC3/STM4 POS
  8-port OC3/STM4 POS
  2-port OC12/STM4 POS
  4-port OC12/STM4 POS
  8-port OC12/STM4 POS
  2-port OC48/STM16 POS/RPR
  1-port OC192/STM64 POS/RPR
  1-port OC192/STM64 POS/RPR

  Engine 5 SPA Interface Processors (SIPs):

  12000-SIP-600 (10G Engine 5 SPA Interface Processor

2.5G multiservice engine SPA Interface Processor
5G multiservice engine SPA Interface Processor
10G multiservice engine SPA Interface Processor

Restrictions for Inter-AS Hybrid for MPLS VPN over IP Tunnels

- In a VPN service provider network configured for MPLS VPNs over IP Tunnels, an ASBR does not, by default, rewrite the route-distinguisher (RD) and route-target (RT) values at autonomous system boundaries. To have RT values rewritten, you must configure a route map on the ASBR. However, a route map does not rewrite RD values.
- Up to 100 VPN Routing and Forwarding (VRF) instances and 200 VPN routes per VRF are supported on each ASBR and PE router.

Information About Inter-AS Hybrid for MPLS VPN over IP Tunnels

To configure the Inter-AS Hybrid for MPLS VPN over IP Tunnels feature, you should understand the following concepts:

- MPLS Virtual Private Networks, page 3
- MPLS VPNs over IP Tunnels, page 5
- Interautonomous System VPN—Back-to-Back VRF, page 7
- MPLS VPN—Interautonomous System Support, page 8
- Inter-AS Hybrid for MPLS VPN over IP Tunnels, page 12

MPLS Virtual Private Networks

The IP Virtual Private Network (VPN) feature for Multiprotocol Label Switching (MPLS) allows a Cisco IOS network to deploy scalable IPv4 Layer 3 VPN backbone services. An IP VPN is the foundation companies use for deploying or administering services, including applications and data hosting network commerce and telephony services, to business customers.

Figure 1 shows an example of a VPN with a service provider (P) backbone network, service provider edge routers (PE), and customer edge routers (CE).

Figure 1 VPNs with a Service Provider Backbone

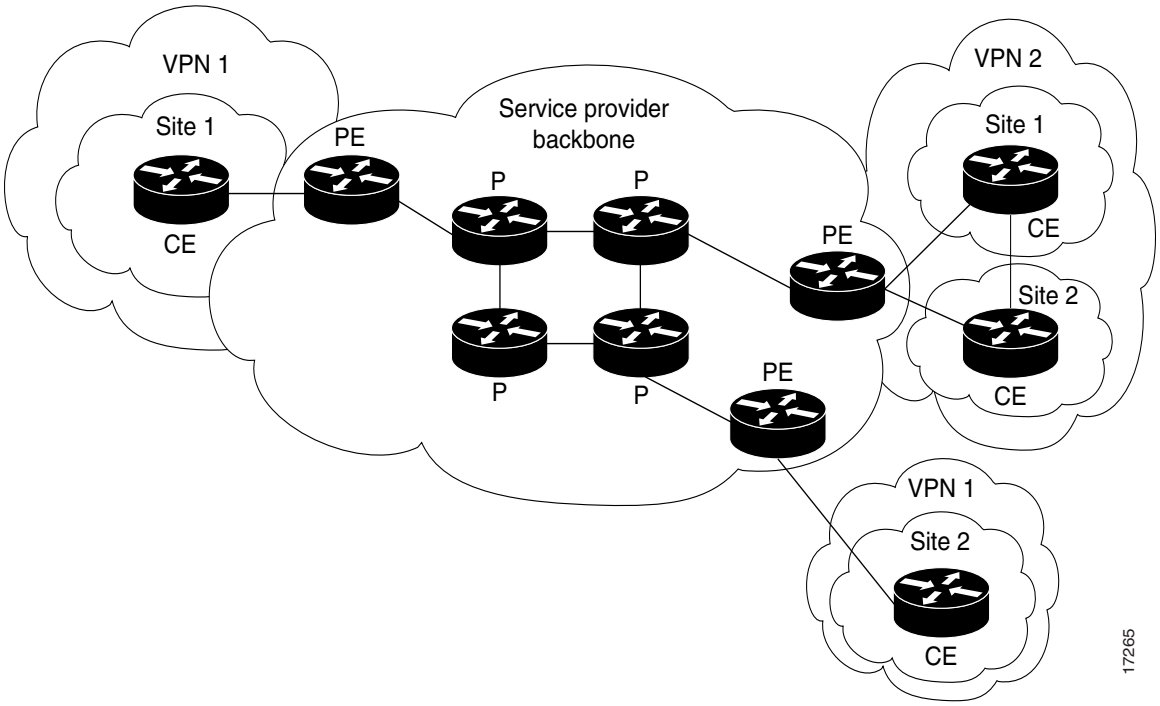

A VPN contains customer devices attached to the CE routers. These customer devices use VPNs to exchange information between devices. Only the PE routers are aware of the VPNs.

Each VPN is associated with one or more VPN routing/forwarding instances (VRFs). A VRF defines the VPN membership of a customer site attached to a PE router. A VRF consists of an IP routing table, a derived Cisco Express Forwarding (CEF) table, a set of interfaces that use the forwarding table, and a set of rules and routing protocol parameters that control the information included in the routing table.

A 1-to-1 relationship does not necessarily exist between customer sites and VPNs. A given site can be a member of multiple VPNs. However, a site can only associate with only one VRF. A customer site's VRF contains all of the routes available to the site from the VPNs of which it is a member.

Packet forwarding information is stored in the IP routing table and the CEF table for each VRF. A separate set of routing and CEF tables is maintained for each VRF. These tables prevent information from being forwarded outside a VPN, and also prevent packets that are outside a VPN from being forwarded to a router within the VPN.

VPN Route Target Communities

The distribution of VPN routing information is controlled through the use of VPN route target communities, implemented by Border Gateway Protocol (BGP) extended communities. Distribution of VPN routing information works as follows:

1. When a VPN route learned from a CE router is injected into BGP, a list of VPN route target extended community attributes is associated with it. Typically the list of route target community values is set from an export list of route targets associated with the VRF from which the route was learned.

2. An import list of route target extended communities is associated with each VRF. The import list defines route target extended community attributes a route must have for the route to be imported into the VRF. For example, if the import list for a particular VRF includes route target communities A, B, and C, then any VPN route that carries any of those route target extended communities—A, B, or C—is imported into the VRF.

BGP Distribution of VPN Routing Information

A service provider edge (PE) router can learn an IP prefix from a customer edge (CE) router by static configuration, through a BGP session with the CE router, or through the routing information protocol (RIP) exchange with the CE router. The IP prefix is a member of the IPv4 address family. After it learns the IP prefix, the PE converts it into a VPN-IPv4 prefix by combining it with an 8-byte route distinguisher (RD). The generated prefix is a member of the VPN-IPv4 address family. It serves to uniquely identify the customer address, even if the customer site is using globally nonunique (unregistered private) IP addresses.

The route distinguisher used to generate the VPN-IPv4 prefix is specified by a configuration command associated with the VRF on the PE router.

BGP distributes reachability information for VPN-IPv4 prefixes for each VPN. BGP communication takes place at two levels: within IP domains, known as an autonomous systems (interior BGP or IBGP) and between autonomous systems (external BGP or EBGP). PE-PE or PE-RR (route reflector) sessions are IBGP sessions, and PE-CE sessions are EBGP sessions.

BGP propagates reachability information for VPN-IPv4 prefixes among PE routers by means of the BGP multiprotocol extensions (refer to RFC 2283, Multiprotocol Extensions for BGP-4) which define support for address families other than IPv4. It does this in a way that ensures that the routes for a given VPN are learned only by other members of that VPN, enabling members of the VPN to communicate with each other.

MPLS Forwarding

Based on routing information stored in the VRF IP routing table and VRF CEF table, packets are forwarded to their destination using MPLS.

A PE router binds a label to each customer prefix learned from a CE router and includes the label in the network reachability information for the prefix that it advertises to other PE routers. When a PE router forwards a packet received from a CE router across the provider network, it labels the packet with the label learned from the destination PE router. When the destination PE router receives the labeled packet, it removes the label and uses it to direct the packet to the correct CE router. Label forwarding across the provider backbone is based on either dynamic label switching or traffic engineered paths. A customer data packet carries two levels of labels when traversing the backbone:

1. Top label directs the packet to the correct PE router.
2. Second label indicates how that PE router forwards the packet to the CE router.

MPLS VPNs over IP Tunnels

The MPLS VPNs over IP Tunnels feature introduces the capability to deploy Layer 3 Virtual Private Network (VPN) services, as proposed in RFC 2547, *BGP/MPLS VPNs*, over an IP core network using L2TPv3 multipoint tunneling instead of Multiprotocol Label Switching (MPLS). This feature allows L2TPv3 tunnels to be configured as multipoint tunnels to transport IP VPN services across the core

IP network. Because multipoint tunnels support multiple endpoints, only one tunnel must be configured on each Provider Edge (PE) router. This feature also introduces a simple packet validation mechanism to enforce VPN integrity.

VPN services are traditionally deployed over IP core networks by configuring MPLS or through L2TPv3 tunnels using point-to-point links. The MPLS VPN over IP Tunnels feature allows you to deploy Layer 3 VPN services by configuring multipoint L2TPv3 tunnels over an existing IP core network. This feature is configured only on PE routers and requires no configuration on the core routers.

The L2TPv3 multipoint tunnel network allows Layer 3 VPN services to be carried through the core without the configuration of MPLS. L2TPv3 multipoint tunnelling supports multiple tunnel endpoints, which creates a full mesh topology that requires only one tunnel to be configured on each PE router. This feature provides the capability for VPN traffic to be carried from enterprise networks across cooperating service provider core networks to remote sites.

Note The configuration of the Inter-AS Hybrid for MPLS VPN over IP Tunnels feature described in this document assumes that an autonomous system has already been configured for the MPLS VPNs over IP Tunnels feature as described in *MPLS VPNs over IP Tunnels*.

Advertising Tunnel Type and Tunnel Capabilities Between PE Routers—BGP

Border Gateway Protocol (BGP) is used to advertise the tunnel endpoints and the subaddress family identifier (SAFI) specific attributes (which contains the tunnel type, and tunnel capabilities). This feature introduces the tunnel SAFI and the BGP SAFI-Specific Attribute (SSA) attribute.

The tunnel SAFI:

- Defines the tunnel endpoint and carries the endpoint IPv4 address and next hop.
- Is identified by the SAFI number 64.

The BGP SSA:

- Carries the BGP preference and BGP flags. It also carries the tunnel cookie, tunnel cookie length, and session ID.
- Is identified by attribute number 19.

These attributes allow BGP to distribute tunnel encapsulation information between PE routers. Virtual Private Network IP Version 4 (VPNv4) traffic is routed through these tunnels. The next hop, advertised in BGP VPNv4 updates, determines which tunnel to use for routing tunnel traffic.

Configuring the PE Routers and Managing Address Space

One multipoint L2TPv3 tunnel is configured on each PE router. To create the VPN, you configure a unique Virtual Routing and Forwarding (VRF) instance. The tunnel that transports the VPN traffic across the core network resides in its own address space. A special purpose VRF called a Resolve in VRF (RiV) is created to manage the tunnel address space. You also configure the address space under the RiV that is associated with the tunnel and a static route in the RiV to route outgoing traffic through the tunnel.

Packet Validation Mechanism

The MPLS VPNs over IP Tunnels feature provides a simple mechanism to validate received packets from appropriate peers. The multipoint L2TPv3 tunnel header is automatically configured with a 64-bit cookie and L2TPv3 session ID.

This packet validation mechanism:

- Protects the VPN from illegitimate traffic sources, such as injecting a rogue packet into the tunnel to gain access to the VPN. The cookie and session ID are not user-configurable; however, they are visible in the packet as it is routed between the two tunnel endpoints.
- Does *not* protect the VPN from hackers who monitor legitimate traffic between PE routers.

Interautonomous System VPN—Back-to-Back VRF

The Interautonomous System VPN—Back-to-Back VRF feature allows separate autonomous systems from different service providers to communicate in an MPLS or IP VPN. This VPN configuration does not require MPLS at the border between autonomous systems.

Figure 2 shows a sample Interautonomous System VPN—Back-to-Back VRF configuration. No signaling is sent between the autonomous system boundary routers (ASBRs). One VPN interface and one VPN Routing and Forwarding (VRF) table are configured on each ASBR for each customer VPN whose routes need to be passed between ASs. The ASBRs are interconnected through multiple logical connections. You can use eBGP to distribute unlabeled IPv4 addresses to a peer ASBR.

Figure 2 Inter-AS Service-Providers Using Back-to-Back VRFs in a VPN

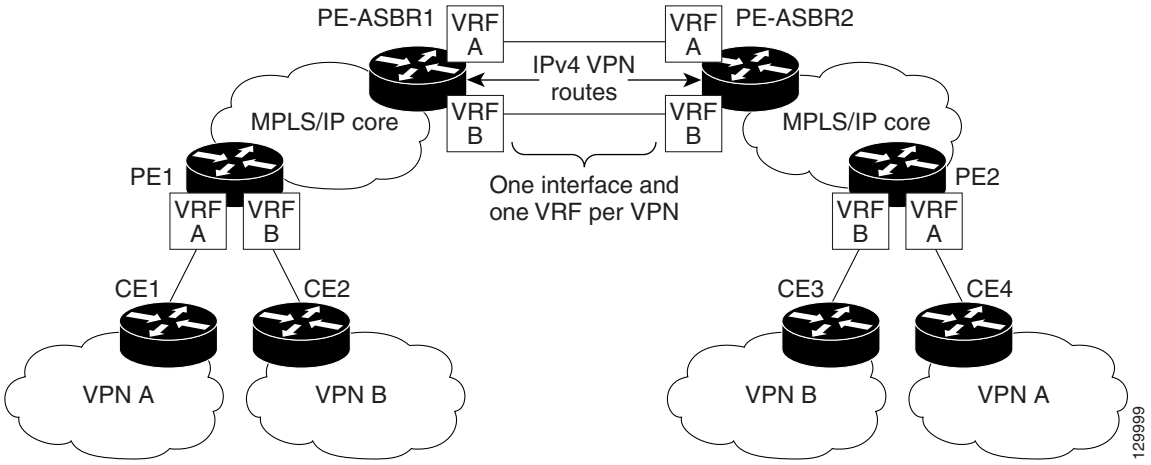

In an Interautonomous System VPN—Back-to-Back VRF configuration, information is transmitted between autonomous systems as follows:

1. The PE-ASBR in one autonomous system performs MPLS or IP VPN decapsulation and transmits packets to the peer PE-ASBR in the adjacent autonomous system.
2. The peer PE-ASBR performs MPLS or IP VPN encapsulation on the customer IPv4 packets received, and transmits the packet through the IPv4 backbone of the autonomous system.

In Figure 2, VPN service providers exchange routes across a back-to-back VRF connection. Each VRF instance:

- Represents a separate VPN client.
- Is configured on a separate PE-ASBR interface, allowing a PE-ASBR to communicate with its peer PE-ASBR as if the peer was a CE router.

The PE-ASBR routers in each autonomous system serve as gateways that are used to exchange IPv4 routes. The link between peer PE-ASBRs supports the following PE-CE protocols: external Border Gateway Protocol (eBGP), Open Shortest Path First (OSPF), Enhanced Interior Gateway Routing Protocol (EIGRP), Routing Information Protocol (RIPv2), and static routing. VPN routes are sent in IPv4 format.

Note In an Interautonomous System VPN–Back-to-Back VRF configuration, the peer ASBRs may use EBGP to distribute VPNv4 routes (as in MPLS VPN—Interautonomous System Support, page 8). However, it is not necessary to use MPLS at the border between ASBRs.

For the Interautonomous System VPN–Back-to-Back VRF feature, when you configure VPN routing and forwarding (VRF) instances and sessions on a PE-ASBR router in a service provider (autonomous system) network, you must:

- Define a separate VRF instance on a peer PE-ASBR router for each VPN that exchanges information with a peer PE-ASBR router in an adjacent autonomous system.
- Configure a separate VRF interface and assign it to a physical interface (or subinterface) on a peer PE-ASBR router for each VPN.
- Configure a separate BGP routing session for each VPN that exchanges information with a peer PE-ASBR router.

MPLS VPN—Interautonomous System Support

The MPLS VPN—Interautonomous System Support feature provides seamless integration of autonomous systems and service providers. Separate autonomous systems from different service providers can communicate by exchanging IPv4 Network Layer Reachability Information (NLRI) in the form of VPN-IPv4 addresses. The autonomous systems' border edge routers use Exterior Border Gateway Protocol (EBGP) to exchange that information. Then, an Interior Gateway Protocol (IGP) distributes the network layer information for VPN-IPv4 prefixes throughout each VPN and each autonomous system. Routing information uses the following protocols:

- Within an autonomous system:
  - PE reachability information is shared using an IGP.
  - VPNv4 routing information is exchanged using MP-BGP.
- Between autonomous systems, routing information is shared using an EBGP. An EBGP allows a service provider to set up an interdomain routing system that guarantees the loop-free exchange of routing information between separate autonomous systems.

An MPLS VPN with interautonomous system support allows a service provider to provide to customers scalable Layer 3 VPN services, such as web hosting, application hosting, interactive learning, electronic commerce, and telephony service. A VPN service provider supplies a secure, IP-based network that shares resources on one or more physical networks.

The primary function of an EBGP is to exchange network reachability information between autonomous systems, including information about the list of autonomous system routes. The autonomous systems use EGBP border edge routers to distribute the routes, which include label switching information. Each border edge router rewrites the next-hop and MPLS labels.

Figure 3 shows how the MPLS VPN–Interautonomous System Support feature is used to set up an MPLS VPN consisting of two separate autonomous systems. Each autonomous system operates under different administrative control and runs a different IGP. Service providers exchange routing information through EBGP border edge routers (ASBR1 and ASBR2).

Figure 3 Inter-AS Service Providers Using EBGP in an MPLS VPN

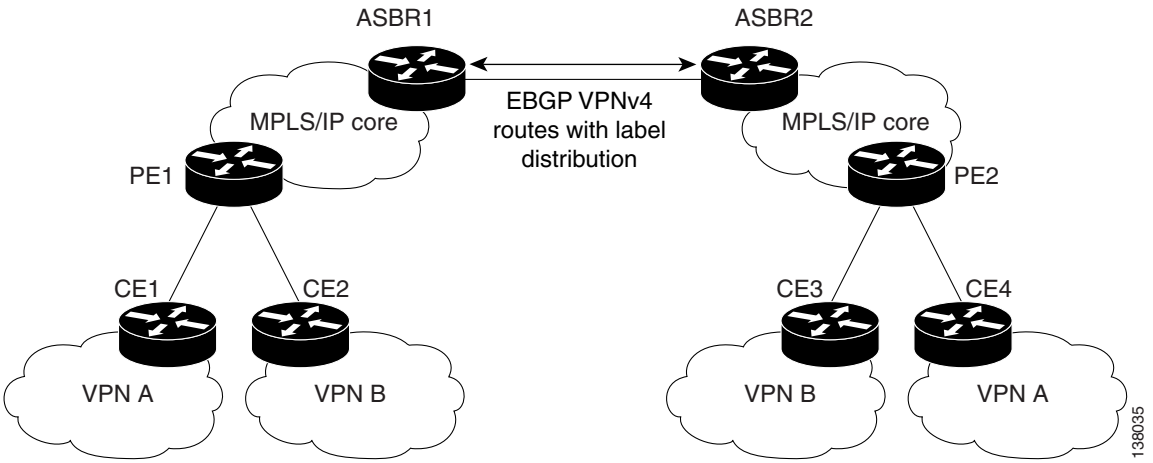

An MPLS VPN–Interautonomous System Support configuration uses the following process to transmit information:

1. The provider edge router (PE1) assigns a label for a route before distributing that route. The PE router uses the multiprotocol extensions of a Border Gateway Protocol (BGP) to transmit label mapping information. The PE router distributes the route as a VPN-IPv4 address. The address label and the VPN identifier are encoded as part of the NLRI.
2. The autonomous systems' border edge routers (ASBR1 and ASBR2) advertise the VPN-IPv4 external routes.
3. The EBGP border edge router (ASBR1) redistributes the route to the next autonomous system (ASBR2). ASBR1 specifies its own address as the value of the EBGP next hop attribute and assigns a new label. The address ensures the following:
   - That the next hop router is always reachable in the service provider backbone network.
   - That the label assigned by the distributing router is properly interpreted. (The label associated with a route must be assigned by the corresponding next hop router.)
4. The EBGP border edge router (ASBR2) redistributes the route in one of the following ways, depending on its configuration. If the IBGP neighbors are:
   - Configured using **neighbor next-hop-self** command—ASBR2 changes the next hop address of updates received from the EBGP peer, then forwards it on.
   - *Not* configured using the **neighbor next-hop-self** command—The next hop address does not get changed. ASBR2 must propagate a host route for the EBGP peer through the IGP. To propagate the EBGP VPN-IPv4 neighbor host route, use the **redistribute connected subnets** command. The EBGP VPN-IPv4 neighbor host route is automatically installed in the routing table when the neighbor comes up. This is essential to establish the label-switched path between PE routers in different autonomous systems.

Exchanging VPN Routing Information

Autonomous systems exchange VPN routing information (routes and labels) to establish connections. To control connections between autonomous systems, the PE routers and EBGP border edge routers maintain a label forwarding information base (LFIB). The LFIB manages the labels and routes that the PE routers and EBGP border edge routers receive during the exchange of VPN information.

Figure 4 illustrates the exchange of VPN route and label information between autonomous systems.

Figure 4 Exchanging Routes and Labels Between Autonomous Systems in an Inter-AS VPN Network

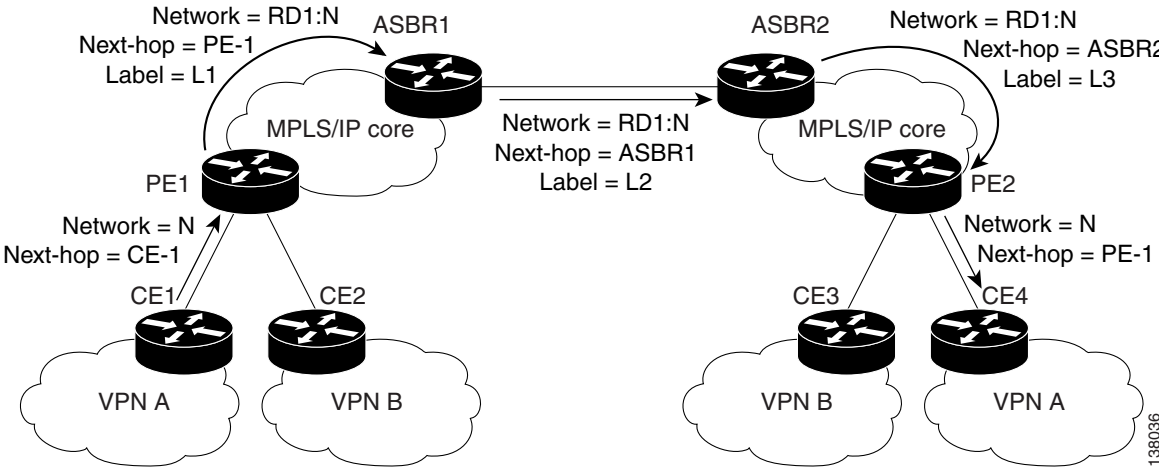

The autonomous systems use the following guidelines to exchange VPN routing information:

- Routing information. Includes:
  - The destination network (N)
  - The next hop field associated with the distributing router
  - A local MPLS label (L)
- An RD1: route distinguisher. Part of a destination network address to make the VPN-IPv4 route globally unique in the VPN service provider environment.
- ASBRs. Configured to change the next hop (next-hop-self) when sending VPN-IPv4 NLRIs to the IBGP neighbors. Therefore, the ASBRs must allocate a new label when they forward the NLRI to the IBGP neighbors.

Packet Forwarding

Figure 5 illustrates how packets are forwarded between autonomous systems in an interprovider network using the following packet forwarding method.

Figure 5 Forwarding Packets Between Autonomous Systems in an Inter-AS VPN Network

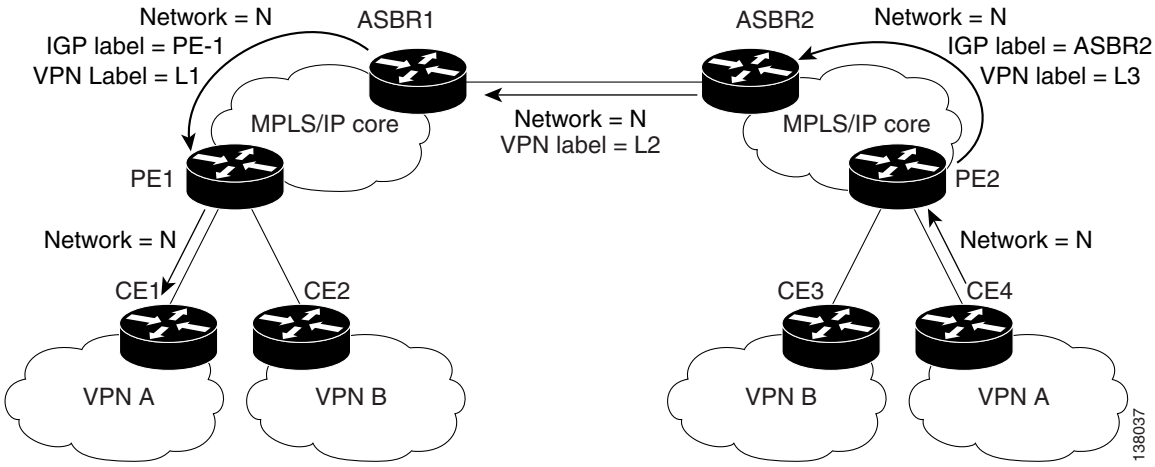

Packets are forwarded to their destination by means of MPLS. Packets use the routing information stored in the LFIB of each PE router and EBGP border edge router.

The service provider VPN backbone uses dynamic label switching to forward labels.

Each autonomous system uses standard multilevel labeling to forward packets between the edges of the autonomous system routers (for example, from CE4 to PE2). Between autonomous systems, only one level of labeling is used, corresponding to the advertised route.

A data packet carries two levels of labels when traversing the Inter-AS VPN backbone:

- The first label (IGP route label) directs the packet to the correct PE router or EBGP border edge router. For example, the IGP label of ASBR2 points to the ASBR2 border edge router.
- The second label (VPN route label) directs the packet to the appropriate PE router or EBGP border edge router.

MPLS VPN—Interautonomous System Configuration

For the MPLS VPN–Interautonomous System feature, it is not necessary to:

- Define a separate VRF instance on an ASBR for each VPN that exchanges information with a peer ASBR (as on the PE-ASBR router in the Interautonomous System VPN—Back-to-Back VRF, page 7, configuration, which exchanges information with a peer PE router). No VRF configuration is required.
- Configure a separate VRF interface on an ASBR for each VPN (as in the Interautonomous System VPN—Back-to-Back VRF, page 7, configuration). You must only configure one interface to transmit IPv4 routes with MPLS labels for all VPNs to a peer ASBR.
- Configure a separate BGP routing session for each VPN (as in the Interautonomous System VPN—Back-to-Back VRF, page 7, configuration). You must only configure one routing protocol to use for all VPNs that exchange routing information with a peer ASBR.

Inter-AS Hybrid for MPLS VPN over IP Tunnels

Starting in Cisco IOS Release 12.0(31)S, you can configure a "hybrid" version of the MPLS VPN—Interautonomous System Support feature on supported ISE interfaces configured for MPLS VPNs over IP Tunnels in a Cisco 12000 series router. The router must be deployed as an ASBR router in a service-provider core network.

The Inter-AS Hybrid for MPLS VPNs over IP Tunnels feature:

- Is supported only on ASBR peer-facing interfaces.
- Allows a service provider network or autonomous system configured for MPLS VPNs over IP Tunnels to be used in an MPLS VPN Interautonomous System network as shown in Figure 6.

Figure 6 MPLS VPNs over IP Tunnels and MPLS Service Providers in an Inter-AS MPLS VPN

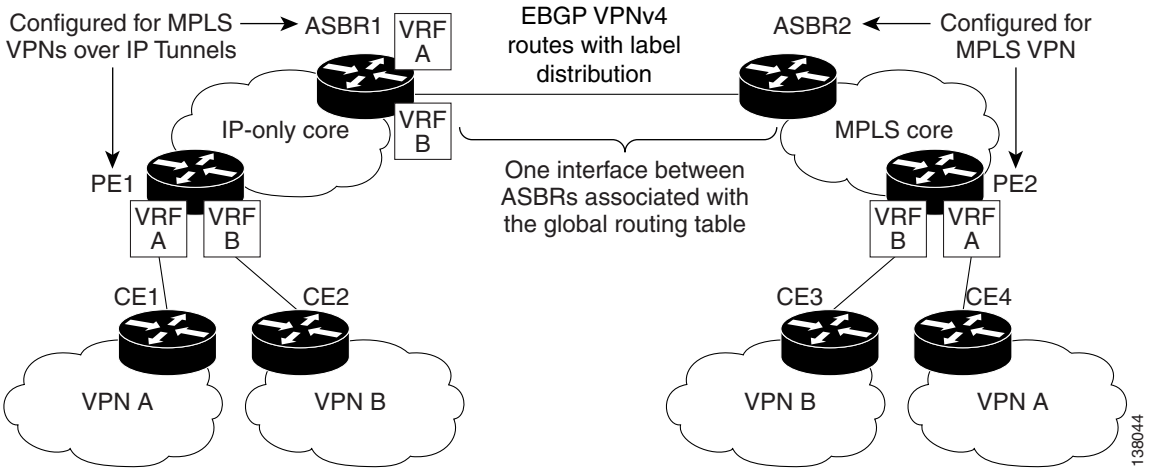

Figure 6 shows how the Inter-AS Hybrid for MPLS VPNs over IP Tunnels feature is used in a VPN interprovider network to interconnect an autonomous system configured for MPLS VPNs over IP Tunnels with another MPLS-based autonomous system.

In the Inter-AS Hybrid for MPLS VPNs over IP Tunnels implementation:

- One VPN Routing and Forwarding (VRF) instance is used for each customer VPN on the ASBR configured for MPLS VPNs over IP Tunnels. This is similar to the Interautonomous System VPN—Back-to-Back VRF, page 7 implementation.
- Only one interface on the ASBR configured for MPLS VPNs over IP Tunnels is required to exchange VPNv4 routes with the MPLS labels of the PE routers in each autonomous system using the same protocol (MP-BGP). This is similar to the MPLS VPN—Interautonomous System Support, page 8 implementation.

Note The Inter-AS Hybrid for MPLS VPNs over IP Tunnels implementation does not use route reflectors to store VPNv4 routes and forward them to the Route Reflectors in adjacent autonomous systems using multihop MPLS EBGP.

The Inter-AS Hybrid for MPLS VPNs over IP Tunnels configuration uses the following process to transmit information:

1. MPLS VPNs over IP Tunnels-based ASBRs (ASBR1) install a customer-specific VFR table for each customer VPN that uses the Inter-AS VPN. For example, Figure 6 shows two customer VPNs (VPN A and VPN B) that use the Inter-AS VPN.
2. Each MPLS VPNs over IP Tunnels-based ASBR peers with the ASBR in an MPLS-based autonomous system (ASBR1 with ASBR2) using EBGP to exchange VPNv4 routes and MPLS labels of the PE routers.

Note In a VPN back-to-back VRF configuration, each PE-ASBR router treats each VRF on the peer PE-ASBR as a CE router, using a customer-specific VFR table for each customer VPN and a separate PE (ASBR) interface for each VRF table to distribute routes with the peer PE-ASBR. For more information, refer to Interautonomous System VPN—Back-to-Back VRF, page 7.

The MPLS VPN–Interautonomous System implementation does not require customer-specific VFR tables. Only one ASBR interface is used to exchange VPNv4 prefixes with the peer ASBR. For more information, refer to *MPLS VPN—Interautonomous Systems Support*.

3. The peer MPLS-based ASBR (ASBR2) receives an MPLS label for each VPN prefix that it installs in its MPLS forwarding information base (LFIB).
4. The MPLS VPNs over IP Tunnels-based ASBR (ASBR1) imports VPN prefixes from its own autonomous system into the VRF table.
5. The MPLS-based ASBR (ASBR2) sends labeled packets to the MPLS VPNs over IP Tunnels-based ASBR (ASBR1).
6. The MPLS VPNs over IP Tunnels-based ASBR (ASBR1) label switches the frames received from the peer ASBR into a new VPNv4 label, encapsulates each packet in an L2TPv3 header, and forwards the packets through the multipoint L2TPv3 tunnel to the PE router in the MPLS VPNs over IP Tunnels-based autonomous system.

This process is transparent to the peer ASBR in the MPLS-based autonomous system.

Note Load-balancing over different links on an ASBR (configured for Inter-AS Hybrid for MPLS VPN over IP Tunnels) or with a peer ASBR (in an adjacent MPLS VPN network) is supported.

Figure 7 illustrates the exchange of VPN route and label information in a VPN interprovider network between an autonomous system configured for MPLS VPNs over IP Tunnels and an MPLS-based autonomous system. In this example, a customer in VPN A transmits information through CE1 to a another customer site in VPN A through CE4.

Figure 7 Exchanging Routes and Labels Between MPLS and MPLS VPNs over IP Tunnels Autonomous Systems in an Inter-AS Hybrid Configuration

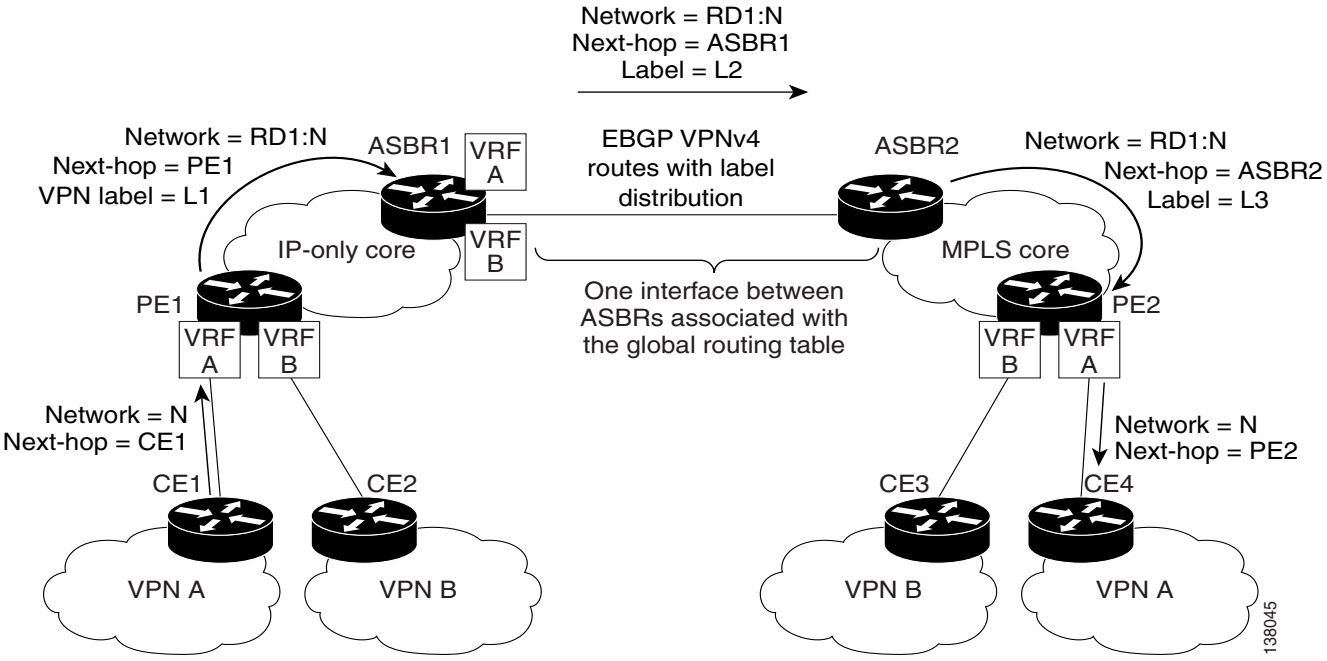

In an Inter-AS Hybrid for MPLS VPNs over IP Tunnels configuration, the autonomous systems exchange VPN routing information (routes and labels) to establish connections in the same way as in an MPLS VPN—Interautonomous System Support configuration, as described in Exchanging VPN Routing Information, page 10.

Figure 8 illustrates how packets are forwarded between autonomous systems in an Inter-AS Hybrid for MPLS VPNs over IP Tunnels configuration, using the same packet forwarding method described in Packet Forwarding, page 11 for an MPLS VPN—Interautonomous System Support configuration. In this example, a customer in VPN A transmits information through CE4 and is received at another customer site in VPN A through CE1.

Figure 8 Forwarding Packets Between MPLS and MPLS VPNs over IP Tunnels Autonomous Systems in an Inter-AS Hybrid Configuration

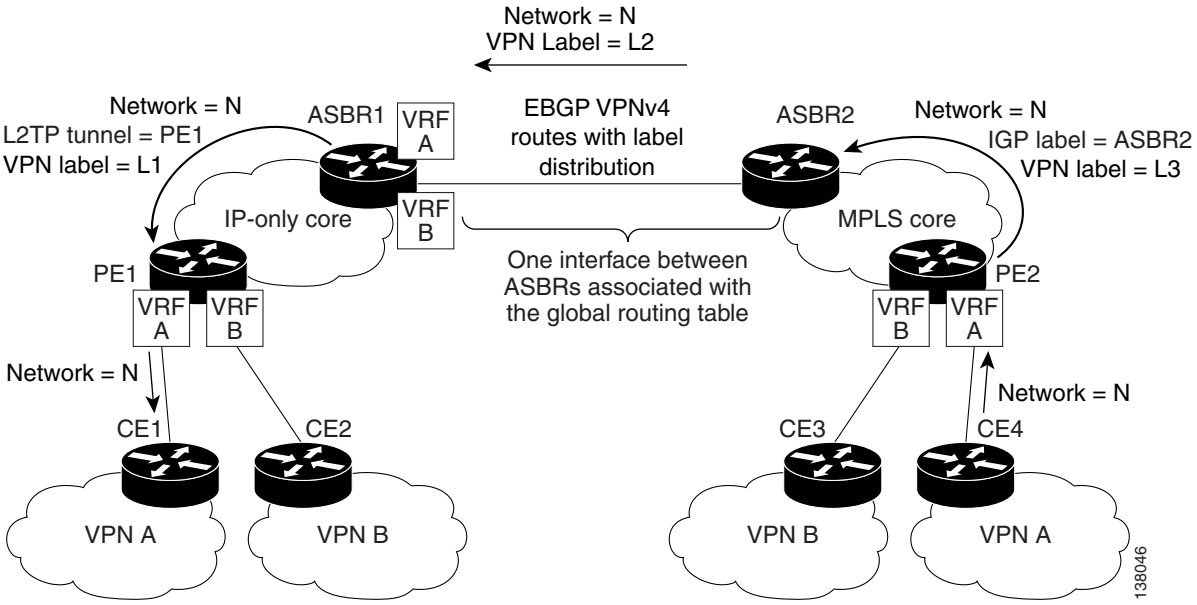

A data packet carries two levels of labels when traversing the Inter-AS VPN backbone:

- The first label (IGP route label) directs the packet to the correct PE router or EBGP border edge router. (For example, the IGP label of ASBR2 points to the ASBR2 border edge router.)
- The second label (VPN route label) directs the packet to the appropriate PE router or EBGP border edge router.

Configuring Inter-AS Hybrid for MPLS VPN over IP Tunnels

This section contains the following procedures:

Preconfiguration Procedures

Before you configure EBGP routing in an MPLS VPN between a VPN autonomous system that runs L2TPv3 in an MPLS VPNs over IP Tunnels configuration and an adjacent MPLS VPN, ensure that:

- You have properly configured all MPLS VPN routing instances and sessions in both autonomous systems.
- You have properly configured the MPLS VPNs over IP Tunnels feature in one autonomous system as described in *MPLS VPNs over IP Tunnels*.
- The adjacent MPLS-based autonomous system is configured for interautonomous system services as described in the *MPLS VPN—Interautonomous Systems Support* document.

The configuration tasks described in this section are based on these configuration tasks.

Perform (as appropriate to the existing network configuration) the following tasks as described in the *Cisco IOS Switching Services Configuration Guide* (the "Configuring Multiprotocol Label Switching chapter).

- Define VPN routing instances on PE routers.
- Configure IP routing sessions in the service-provider networks.
- Configure PE-to-PE MP-BGP routing sessions in the service-provider networks.
- Configure PE-to-CE static, IGP, or BGP routing sessions.

Configuring EBGP Routing for the Exchange of VPNv4 Routes on Each Peer ASBR

Note You must perform this configuration task only on the peer ASBR in the service-provider network that is configured for MPLS VPNs over IP Tunnels and connected to an ASBR in the MPLS-based service-provider network configured for the MPLS VPN—Interautonomous System Support feature.

To configure the peer ASBR router in an interautonomous service-provider network configured for MPLS VPNs over IP Tunnels to exchange VPN routes using EBGP with an adjacent MPLS-based autonomous system, use the following commands starting in EXEC mode.

Note You must set the next-hop-self attribute on the MPLS VPN over IP Tunnels ASBR. If the next-hop-self address is not set, the MPLS VPN over IP Tunnels PE attempts to tunnel packets to the peer ASBR where no viable path exists.

	Command	Purpose
Step 1	`Router# configure terminal`	Enters global configuration mode.
Step 2	`Router(config)# router bgp` *as-number*	Creates an EBGP routing process and assigns it an AS number. The autonomous system number is passed along to identify the router to EBGP routers in another autonomous system.
Step 3	`Router(config)# no bgp default route-target filter`	(Optional) Disables BGP route-target filtering. All received BGP VPN-IPv4 routes are accepted by the router. If you do not enter this command on the MPLS VPNs over IP Tunnels ASBR, the ASBR only accepts prefixes from the peer ASBR for which a VRF is configured for a VPN that requires interautonomous-system services.
Step 4	`Router(config-router)# address-family vpnv4[unicast]`	Configures a routing session to carry VPN-IPv4 addresses across the VPN backbone. Each address is made globally unique by the addition of an 8-byte route distinguisher (RD). Unicast is optional; use it to specify a unicast prefix.
Step 5	`Router(config-router-af)# neighbor` *ip-address* `remote-as` *as-number*	Enters the address family submode and specifies the IP address of the neighboring peer ASBR.
Step 6	`Router(config-router-af)# neighbor` *ip-address* `activate`	Activates the advertisement of the VPN-IPv4 address family to the neighboring peer ASBR.
Step 7	`Router(config-router-af)# send community both`	Activates the exchange of VPNv4 labels with the neighboring peer ASBR.
Step 8	`Router(config-router-af)# exit-address-family`	Exits from the address family submode of global configuration mode.

Configuring VRFs on the MPLS VPNs over IP Tunnels ASBR

Note You must perform this configuration task only on the peer ASBR in the autonomous system configured for MPLS VPNs over IP Tunnels.

To define VPN routing instances on the peer ASBR router in each autonomous system, use the following commands beginning in global configuration mode:

	Command	Purpose
Step 1	`Router(config)# ip vrf vrf-name`	Enters VRF configuration mode and define the VPN routing instance by assigning a VRF name.
Step 2	`Router(config-vrf)# rd route-distinguisher`	Creates routing and forwarding tables.
Step 3	`Router(config-vrf)# route-target {import | export | both} route-target-ext-community`	Creates a list of import and/or export route target communities for the specified VRF.
Step 4	`Router(config-vrf)# import map route-map`	(Optional) Associates the specified route map with the VRF.
Step 5	`Router(config-vrf)# exit`	Exits VRF configuration mode.
Step 6	`Router(config)# interface type slot/port-number`	Enters interface configuration mode. **Note** On an Engine 5 SPA, you specify an interface in the format: *slot*/*subslot*/*port-number*
Step 7	`Router(config-if)# ip vrf forwarding vrf-name`	Associates a VRF with an interface or subinterface.

Verifying Inter-AS Hybrid over MPLS VPNs over IP Tunnels

Perform the tasks in this section to verify the Inter-AS Hybrid over MPLS VPNs over IP Tunnels configuration on each peer ASBR.

Verifying VPN Operation

To verify VPN operation, use the following commands in privileged EXEC mode:

	Command	Displays:
Step 1	`Router# show ip vrf`	Set of defined VRFs and interfaces.
Step 2	`Router# show ip vrf [{brief | detail | interfaces}] vrf-name`	Information about defined VRFs and associated interfaces.
Step 3	`Router# show ip route vrf vrf-name`	IP routing table for a VRF.
Step 4	`Router# show ip protocols vrf vrf-name`	Routing protocol information for a VRF.
Step 5	`Router# show ip cef vrf vrf-name`	CEF forwarding table associated with a VRF.
Step 6	`Router# show ip interface interface-number`	VRF table associated with an interface.
Step 7	`Router# show ip bgp vpnv4 all [labels]`	Information about all BGP routes.
Step 8	`Router# show mpls forwarding vrf vrf-name [prefix mask/length] [detail]`	Label-forwarding entries that correspond to VRF routes advertised by this router.

Displaying VPN-IPv4 LFIB Entries

To display the VPN-IPv4 label forwarding information base (LFIB) entries at the border edge routers in the autonomous systems, use the following commands starting in EXEC mode:

	Command	Displays:
Step 1	`Router# show ip bgp vpnv4 all [tags]`	Information about all VPN-IPv4 labels.
Step 2	`Router# show tag-switching forwarding-table`	Contents of the LFIB (such as VPN-IPv4 prefix/length and BGP next hop destination for the route).

The following example shows the appearance of VPN-IPv4 LFIB entries when you use the **show tag-switching forwarding-table** privileged EXEC command:

```
Router# show tag-switching forwarding-table
Local Outgoing      Prefix            Bytes tag Outgoing       Next Hop
tag   tag or VC     or Tunnel Id      switched  interface
33    33            10.120.4.0/24     0         Hs0/0          point2point
35    27            100:12:10.200.0.1/32 \
                                      0         Hs0/0          point2point
```

In this example, the Prefix field appears as a VPN-IPv4 route distinguisher (RD), plus the prefix. If the value is longer than the Prefix column (as illustrated in the last line of the example), the output automatically wraps onto the next line in the forwarding table to preserve column alignment.

Configuration Example for Inter-AS Hybrid for MPLS VPN over IP Tunnels

This section describes a configuration example for Inter-AS Hybrid for MPLS VPNs over IP Tunnels, and consists of the following steps:

- PE1 Router Configuration Example—MPLS VPNs over IP Tunnels Service Provider, page 21
- ASBR1 Router Configuration Example—MPLS VPNs over IP Tunnels Service Provider, page 23
- ASBR2 Router Configuration Example—MPLS VPN Service Provider, page 26
- PE2 Router Configuration Example—MPLS VPN Service Provider, page 28

Note 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.

The sample configuration shown in Figure 9 is for two autonomous-system MPLS VPN service providers: an MPLS VPNs over IP Tunnels-based network and an MPLS-based network.

Figure 9 Configuring Inter-AS for MPLS VPNs over IP Tunnels in an MPLS VPN Network

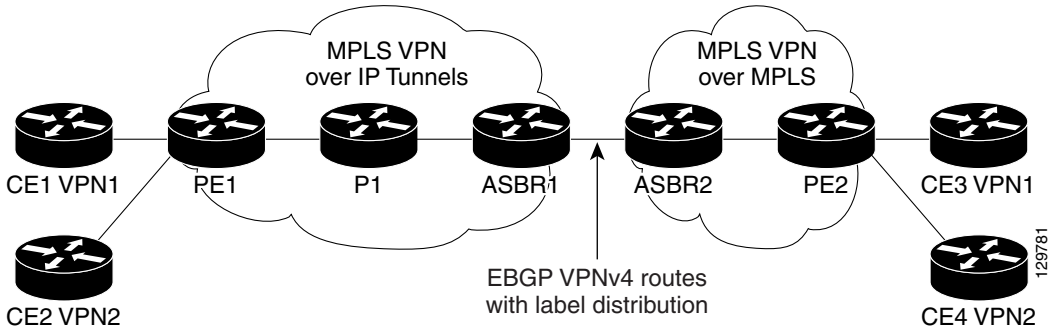

In this interautonomous-system network, the MPLS VPNs over IP Tunnels-based service provider runs IS-IS as its IGP. The MPLS-based service provider runs OSPF as its IGP. Two VPNs (VPN 1 and VPN 2) interconnect over the two service-provider networks.

Each service-provider network is configured to rewrite route target values. (The route distinguisher (**rd**) values cannot be rewritten.) Route targets are carried as extended community attributes in BGP VPNv4 updates. Route target extended community attributes are used to identify a set of sites and VPN routing/forwarding instances (VRFs) that can receive routes with a configured route target. Each ASBR performs route target replacement at the autonomous system border when peer ASBRs exchange VPNv4 prefixes. For more information, refer to *MPLS VPN—Route Target Rewrite*.

Note Because provider (P) core routers do not require additional configuration, sample P router configurations are not included in this example of configuring the Inter-AS for MPLS VPNs over IP Tunnels feature.

Prerequisites

In the MPLS VPNS over IP Tunnels-based network, the following prerequisites exist:

- A Cisco 12000 series ISE line card is required on the ASBR (for the ASBR1-to-ASBR2 connection) and on the PE router for the customer-facing interface (in the PE1-to-CE1 connection).
- You must configure each ISE line card in feature mode.

Note To improve performance, do not configure the backbone- facing interfaces in the MPLS VPNS over IP Tunnels-based network in feature mode. The ASBR1-to-ASBR2 connection is considered as a customer-facing connection.

It is not necessary to configure the ISE line cards in the ASBR in the MPLS-based autonomous system (ASBR2) in feature mode.

PE1 Router Configuration Example—MPLS VPNs over IP Tunnels Service Provider

The following example shows a sample configuration for the PE1 router in the MPLS VPNs over IP Tunnels-based autonomous system shown in Figure 9 and specifies the following:

- Two customer edge (CE) routers (CE1 and CE2), each belonging to a separate VPN (VPN1 and VPN2) that communicates across the MPLS VPN interautonomous-system network, are connected to PE1 on the POS 2/0 and POS 2/1 interfaces.
- PE1 connects to the P1 core router on the POS 3/0 interface.
- OSPF is used as the PE-to-CE routing protocol.
- The CE-facing ISE card on PE1 must be configured in feature mode.
- In this configuration, the MPLS VPNs over IP Tunnels-based network is referred to as partner1; the MPLS VPN-based network is referred to as partner2.

```
hw-module slot 2 np mode feature

ip vrf RIV
 rd 0:0

ip vrf vpn1
 rd 200:1
 route-target export 200:1
 route-target import 200:1

ip vrf vpn2
 rd 200:2
 route-target export 200:2
 route-target import 200:2

interface Loopback0
 ip address 10.127.13.1 255.255.255.255
 no ip directed-broadcast

interface Loopback1
 ip vrf forwarding vpn1
 ip address 10.127.44.1 255.255.255.255
 no ip directed-broadcast

interface Loopback2
 ip vrf forwarding vpn2
 ip address 10.127.45.1 255.255.255.255
 no ip directed-broadcast

interface Tunnel0
 ip vrf forwarding RIV
 ip address 192.168.13.1 255.255.255.255
 no ip redirects
 no ip directed-broadcast
 tunnel source Loopback0
 tunnel mode l3vpn l2tpv3 multipoint

interface POS2/0
 ip vrf forwarding vpn1
 ip address 192.168.144.1 255.255.255.252
 no ip directed-broadcast
 crc 32
 clock source internal
 pos ais-shut
 pos scramble-atm
```

```
interface POS2/1
 ip vrf forwarding vpn2
 ip address 192.168.144.1 255.255.255.252
 no ip directed-broadcast
 crc 32
 clock source internal
 pos scramble-atm

interface POS 3/0
 ip address 192.168.79.2 255.255.255.0
 no ip directed-broadcast
 ip router isis P1
 encapsulation ppp
 crc 32
 clock source internal
 pos ais-shut
 pos scramble-atm

router ospf 100 vrf vpn1
 router-id 192.168.44.1
 domain-id 192.168.0.0
 log-adjacency-changes
 redistribute bgp 200 metric 1000 subnets
 network 192.168.144.1 0.0.0.3 area 0

router ospf 101 vrf vpn2
 router-id 192.168.45.1
 log-adjacency-changes
 redistribute bgp 200 metric 1000 subnets
 network 192.168.144.1 0.0.0.3 area 0

router isis P1
 net 47.0023.0001.0000.0001.0002.0001.1921.6813.0001.00
 is-type level-1
 metric-style wide
 passive-interface Loopback0

router bgp 200
 bgp router-id 192.168.13.1
 no bgp fast-external-fallover
 bgp maxas-limit 70
 no bgp default route-target filter
 bgp log-neighbor-changes
 bgp deterministic-med
 timers bgp 45 135
 neighbor partner1 peer-group
 neighbor partner1 remote-as 200
 neighbor partner1 update-source Loopback0
 neighbor partner1 version 4
 neighbor 192.168.112.1 peer-group partner1
 neighbor 192.168.112.2 peer-group partner1

 address-family ipv4
 redistribute connected
 redistribute static route-map static-mbgp
 neighbor partner1 activate
 neighbor partner1 send-community extended
 neighbor 192.168.112.1 peer-group partner1
 neighbor 192.168.112.2 peer-group partner1
 default-metric 1
 no auto-summary
 no synchronization
 exit-address-family
```

```
 address-family ipv4 tunnel
 neighbor partner1 activate
 neighbor partner1 send-community extended
 neighbor 192.168.112.1 peer-group partner1
 neighbor 192.168.112.2 peer-group partner1
 exit-address-family

 address-family vpnv4
 neighbor partner1 activate
 neighbor partner1 send-community extended
 neighbor partner1 next-hop-self
 neighbor partner1 route-map RIV-nexthop in
 neighbor 192.168.112.1 peer-group partner1
 neighbor 192.168.112.1 route-map RIV-nexthop in
 neighbor 192.168.112.2 peer-group partner1
 neighbor 192.168.112.2 route-map RIV-nexthop in
 exit-address-family

 address-family ipv4 vrf vpn2
 redistribute connected
 redistribute static
 redistribute ospf 101 vrf vpn2 match internal external 1 external 2
 no auto-summary
 no synchronization
 exit-address-family

 address-family ipv4 vrf vpn1
 redistribute connected
 redistribute static
 redistribute ospf 100 vrf vpn1 match internal external 1 external 2
 no auto-summary
 no synchronization
 exit-address-family

 address-family ipv4 vrf RIV
 no auto-summary
 no synchronization
 exit-address-family

ip classless
ip route vrf RIV 0.0.0.0 0.0.0.0 Tunnel0

ip bgp-community new-format

route-map RIV-nexthop permit 10
set ip next-hop in-vrf RIV
```

ASBR1 Router Configuration Example—MPLS VPNs over IP Tunnels Service Provider

The following example shows a sample configuration for the ASBR1 router in the MPLS VPNs over IP Tunnels-based autonomous system shown in Figure 9 and specifies the following:

- The MPLS VPNs over IP Tunnels-based network connects to the MPLS-based network through the Gigabit Ethernet 1/0 interface on ASBR1. The Gigabit Ethernet ISE line card is configured in feature mode.
- In this configuration, the MPLS VPNs over IP Tunnels-based network is referred to as partner1; the MPLS VPN-based network is referred to as partner2.
- ASBR1 connects to the P1 core router on the POS 2/0 interface.

- Each VPN that communicates over the MPLS VPN interautonomous system is configured in a separate VRF instance only in the ASBR1 configuration.
- The route target (RT) values are rewritten as follows: 100:1 is rewritten to 200:1 and 100:2 is rewritten to 200:2. These RT values correspond to each of the two VPNs (VPN1 and VPN2).
- Because ASBR1 is deployed as a PE device, this configuration includes the use of **ip prefix-list** and **ip as-path access-list** commands to filter both inbound and outbound BGP routes. For more information, refer to *BGP Commands*.

```
hw-module slot 1 np mode feature

ip subnet-zero
ip vrf RIV
 rd 0:0

ip vrf vpn1
 rd 200:1
 route-target export 200:1
 route-target import 200:1

ip vrf vpn2
 rd 200:2
 route-target export 200:2
 route-target import 200:2

interface Loopback0
 ip address 10.127.112.1 255.255.255.255
 no ip directed-broadcast
 no ip route-cache

interface Tunnel0
 ip vrf forwarding RIV
 ip address 10.127.112.1 255.255.255.255
 no ip redirects
 no ip directed-broadcast
 tunnel source Loopback0
 tunnel mode l3vpn l2tpv3 multipoint

interface GigabitEthernet1/0
 ip address 192.168.10.161 255.255.255.252
 no ip directed-broadcast
 mpls bgp forwarding
 tag-switching ip

interface POS2/0
 ip address 192.168.10.88 255.255.255.252
 no ip directed-broadcast
 ip router isis P1
 crc 32
 clock source internal
 pos ais-shut
 pos scramble-atm
 clns router isis P1

router isis P1
 net 47.0023.0001.0000.0001.0002.0001.1921.6811.0012.00
 is-type level-1
 metric-style wide
 passive-interface Loopback0

router bgp 200
 bgp router-id 192.168.112.1
 no bgp fast-external-fallover
```

```
bgp maxas-limit 70
bgp log-neighbor-changes
bgp deterministic-med
bgp graceful-restart restart-time 120
bgp graceful-restart stalepath-time 360
bgp graceful-restart
timers bgp 45 135
neighbor partner1 peer-group
neighbor partner1 remote-as 200
neighbor partner1 update-source Loopback0
neighbor partner1 version 4
neighbor 10.161.1.6 remote-as 100
neighbor 10.161.1.6 update-source GigabitEthernet1/0
neighbor 192.168.13.1 peer-group partner1
neighbor 192.168.112.2 peer-group partner1

address-family ipv4
redistribute static
neighbor partner1 activate
neighbor partner1 send-community extended
neighbor partner1 next-hop-self
neighbor 10.161.1.6 activate
neighbor 10.161.1.6 send-community both
neighbor 10.161.1.6 send-label
neighbor 192.168.13.1 peer-group partner1
neighbor 192.168.112.2 peer-group partner1
default-metric 1
no auto-summary
no synchronization
exit-address-family

address-family ipv4 tunnel
neighbor partner1 activate
neighbor partner1 send-community extended
neighbor 192.168.13.1 peer-group partner1
exit-address-family

address-family vpnv4
neighbor partner1 activate
neighbor partner1 send-community extended
neighbor partner1 next-hop-self
neighbor partner1 route-map RIV-nexthop in
neighbor 10.161.1.6 activate
neighbor 10.161.1.6 send-community both
neighbor 10.161.1.6 prefix-list filter-in in
neighbor 10.161.1.6 route-map rtwrite in
neighbor 192.168.13.1 peer-group partner1
neighbor 192.168.13.1 route-map RIV-nexthop in
exit-address-family

address-family ipv4 vrf vpn2
no auto-summary
no synchronization
exit-address-family

address-family ipv4 vrf vpn1
no auto-summary
no synchronization
exit-address-family

address-family ipv4 vrf RIV
no auto-summary
no synchronization
exit-address-family
```

```
ip classless
 ip route vrf RIV 0.0.0.0 0.0.0.0 Tunnel0
 ip extcommunity-list 1 permit rt 100:1
 ip extcommunity-list 2 permit rt 100:2
 ip bgp-community new-format
 ip as-path access-list 1 deny _6451[2-9]_
 ip as-path access-list 1 deny _645[2-9][0-9]_
 ip as-path access-list 1 deny _64[6-9][0-9][0-9]_
 ip as-path access-list 1 deny _65[0-9][0-9][0-9]_
 ip as-path access-list 1 deny \(*\)
 ip as-path access-list 1 permit .*

 ip prefix-list filter-in seq 20 deny 10.0.0.0/8 le 32
 ip prefix-list filter-in seq 30 deny 127.0.0.0/8 le 32
 ip prefix-list filter-in seq 40 deny 128.0.0.0/16 le 32
 ip prefix-list filter-in seq 50 deny 169.254.0.0/16 le 32
 ip prefix-list filter-in seq 60 deny 172.16.0.0/12 le 32
 ip prefix-list filter-in seq 70 deny 191.255.0.0/16 le 32
 ip prefix-list filter-in seq 80 deny 192.0.2.0/24 le 32
 ip prefix-list filter-in seq 90 deny 192.168.0.0/16 le 32
 ip prefix-list filter-in seq 100 deny 223.255.255.0/24 le 32
 ip prefix-list filter-in seq 110 deny 224.0.0.0/3 le 32
 ip prefix-list filter-in seq 120 deny 14.0.50.0/24
 ip prefix-list filter-in seq 125 deny 13.0.4.0/30
 ip prefix-list filter-in seq 130 permit 0.0.0.0/0 ge 8
 no logging trap

route-map RIV-nexthop permit 10
 set ip next-hop in-vrf RIV

route-map rtwrite permit 10
 match extcommunity 1
 set extcommunity rt 200:1

route-map rtwrite permit 20
 match extcommunity 2
 set extcommunity rt 200:2
```

ASBR2 Router Configuration Example—MPLS VPN Service Provider

The following example shows a sample configuration for the ASBR2 router in the MPLS-based autonomous system shown in Figure 9 and specifies the following:

- The MPLS-based ASBR (ASBR2) is configured for the MPLS VPN–Interautonomous System Support feature by using the configuration procedures and commands described in *MPLS VPN—Interautonomous Systems Support*. No ISE line cards in ASBR2 need to be configured in feature mode.
- The route target (RT) values are rewritten as follows: 100:1 is rewritten to 200:1 and 100:2 is rewritten to 200:2. These RT values correspond to each of the two VPNs (VPN1 and VPN2).
- In this configuration, the MPLS VPNs over IP Tunnels-based network is referred to as partner1; the MPLS VPN-based network is referred to as partner2.
- The MPLS-based network connects to the MPLS VPNs over IP Tunnels-based network through the Gigabit Ethernet 5/0 interface on ASBR2.
- ASBR2 connects to the P2 core router on the POS 3/0 interface.

```
interface Loopback0
 ip address 10.127.168.229 255.255.255.255
 no ip directed-broadcast
```

```
interface POS3/0
 ip address 192.168.10.170 255.255.255.252
 no ip directed-broadcast
 mpls label protocol ldp
 tag-switching ip
 crc 32
 clock source internal
 pos scramble-atm

interface GigabitEthernet5/0
 ip address 192.168.10.161 255.255.255.252
 no ip directed-broadcast
 mpls bgp forwarding
 tag-switching ip
 no negotiation auto

router ospf 100
 log-adjacency-changes
 network 10.170.1.0 0.0.0.3 area 0
 network 192.168.229.1 0.0.0.0 area 0

router bgp 100
 no bgp default route-target filter
 bgp log-neighbor-changes
 bgp graceful-restart restart-time 120
 bgp graceful-restart stalepath-time 360
 bgp graceful-restart
 neighbor 10.161.1.5 remote-as 200
 neighbor 10.161.1.5 update-source GigabitEthernet5/0
 neighbor 192.168.82.1 remote-as 100
 neighbor 192.168.82.1 update-source Loopback0

 address-family ipv4
 neighbor 10.161.1.5 activate
 neighbor 10.161.1.5 send-community both
 neighbor 10.161.1.5 send-label
 neighbor 192.168.82.1 activate
 neighbor 192.168.82.1 next-hop-self
 no auto-summary
 no synchronization
 exit-address-family

 address-family vpnv4
 neighbor 10.161.1.5 activate
 neighbor 10.161.1.5 send-community both
 neighbor 10.161.1.5 route-map rtwrite in
 neighbor 192.168.82.1 activate
 neighbor 192.168.82.1 send-community both
 neighbor 192.168.82.1 next-hop-self
 exit-address-family

ip classless

ip extcommunity-list 1 permit rt 200:1
ip extcommunity-list 2 permit rt 200:2

route-map rtwrite permit 10
 match extcommunity 1
 set extcommunity rt 100:1

 route-map rtwrite permit 20
 match extcommunity 2
 set extcommunity rt 100:2
```

PE2 Router Configuration Example—MPLS VPN Service Provider

The following example shows a sample configuration for the PE2 router in the MPLS-based autonomous system shown in Figure 9 and specifies the following:

- The MPLS-based PE router (PE2) is configured for the MPLS VPN feature by using the configuration procedures and commands described in *MPLS Virtual Private Networks (VPNs)*.
- Two customer edge (CE) routers (CE3 and CE4), each belonging to a separate VPN (VPN1 and VPN2) that communicates across the MPLS VPN interautonomous-system network, are connected to PE2 on the POS 7/1 and Gigabit Ethernet 5/0 interfaces.
- PE2 connects to ASBR2 on the POS 7/0 interface.
- OSPF is used as the PE-to-CE routing protocol.
- In this configuration, the MPLS VPNs over IP Tunnels-based network is referred to as partner1; the MPLS VPN-based network is referred to as partner2.

```
ip vrf vpn1
 rd 100:1
 route-target export 100:1
 route-target import 100:1

ip vrf vpn2
 rd 100:2
 route-target export 100:2
 route-target import 100:2

interface Loopback0
 ip address 10.127.168.82 255.255.255.255
 no ip directed-broadcast

interface GigabitEthernet5/0
 ip vrf forwarding vpn2
 ip address 192.168.0.11.180 255.255.255.252
 no ip directed-broadcast
 no negotiation auto

interface POS7/0
 ip address 192.168.10.170 255.255.255.252
 no ip directed-broadcast
 mpls label protocol ldp
 tag-switching ip
 crc 32
 clock source internal
 pos scramble-atm

interface POS7/1
 ip vrf forwarding vpn1
 ip address 192.168.10.180 255.255.255.252
 no ip directed-broadcast
 crc 32
 clock source internal
 pos scramble-atm

router ospf 101 vrf vpn1
 log-adjacency-changes
 redistribute bgp 100 metric 1000 subnets
 network 10.180.1.0 0.0.0.3 area 0

router ospf 102 vrf vpn2
 log-adjacency-changes
 redistribute bgp 100 metric 1000 subnets
```

```
 network 10.180.1.0 0.0.0.3 area 0

router ospf 100
 log-adjacency-changes
 network 10.170.1.0 0.0.0.3 area 0
 network 192.168.82.1 0.0.0.0 area 0

router bgp 100
 bgp log-neighbor-changes
 neighbor 192.168.229.1 remote-as 100
 neighbor 192.168.229.1 update-source Loopback0

 address-family ipv4
 neighbor 192.168.229.1 activate
 no auto-summary
 no synchronization
 exit-address-family

 address-family vpnv4
 neighbor 192.168.229.1 activate
 neighbor 192.168.229.1 send-community both
 exit-address-family

 address-family ipv4 vrf vpn2
 redistribute connected
 redistribute ospf 102 vrf vpn2 match internal external 1 external 2
 no auto-summary
 no synchronization
 exit-address-family

 address-family ipv4 vrf vpn1
 redistribute connected
 redistribute ospf 101 vrf vpn1 match internal external 1 external 2
 no auto-summary
 no synchronization
 exit-address-family
```

Additional References

The following sections provide references related to the Inter-AS Hybrid for MPLS VPNs over IP Tunnels feature.

Related Documents for MPLS VPN Interautonomous Systems

Related Topic	Document Title
MPLS VPN interautonomous systems: configuration tasks and commands	*MPLS VPN—Interautonomous Systems Support*
Virtual Private Network (VPN) configuration tasks	*MPLS Virtual Private Networks (VPNs)* *MPLS Virtual Private Network Enhancements*
An explanation of how Border Gateway Protocol (BGP) works and how you can use it to participate in routing with other networks that run BGP	*Using the Border Gateway Protocol for Interdomain Routing*
Border Gateway Protocol (BGP) configuration tasks	"Configuring BGP" chapter in the *Cisco IOS IP Configuration Guide*, Release 12.2
An explanation of the purpose of the Border Gateway Protocol and the BGP route selection process, and how to use BGP attributes in route selection	"Border Gateway Protocol" chapter in the *Internetworking Technology Overview*
Multiprotocol Label Switching (MPLS) configuration tasks	"Configuring Multiprotocol Label Switching" chapter in the *Cisco IOS Switching Services Configuration Guide*, Release 12.3
Commands to configure and monitor BGP	"BGP Commands" chapter in the *Cisco IOS IP Command Reference, Volume 2 of 3: Routing Protocols*, Release 12.2

Related Documents for MPLS VPNS over IP Tunnels

Related Topic	Document Title
MPLS VPNs over IP Tunnels: configuration tasks and commands	*M PLS VPNs over IP Tunnels*
CEF switching	*Cisco IOS Switching Services Configuration Guide*, Release 12.3
QoS—Tunnel Marking	*QoS: Tunnel Marking for L2TPv3 Tunnels*
VPN configuration	*Cisco IOS Dial Technologies Configuration Guide*, Release 12.3 and *Cisco IOS Switching Services Configuration Guide*, Release 12.3
VPN Routing and Forwarding (VRF) instances	*Cisco IOS Switching Services Configuration Guide*, Release 12.3

Standards

Standards	Title
No new or modified standards are supported by this feature.	—

MIBs

MIBs	MIBs Link
No new or modified MIBs are supported by this feature.	To locate and download MIBs for selected platforms, Cisco IOS releases, and feature sets, use Cisco MIB Locator at: http://www.cisco.com/go/mibs

RFCs

RFCs	Title
No new or modified RFCs are supported by this feature.	—

Technical Assistance

Description	Link
Technical Assistance Center (TAC) home page, containing 30,000 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/public/support/tac/home.shtml

Command Reference

There are no new Cisco IOS commands introduced with this feature.

All commands used with this feature are described in the following Cisco IOS documentation:

- *MPLS Virtual Private Networks (VPNs)*
- *MPLS Virtual Private Network Enhancements*
- *Cisco IOS Switching Services Configuration Guide* (Release 12.2), Multiprotocol Label Switching
- *MPLS VPNs over IP Tunnels*
- *MPLS VPN—Interautonomous Systems Support*