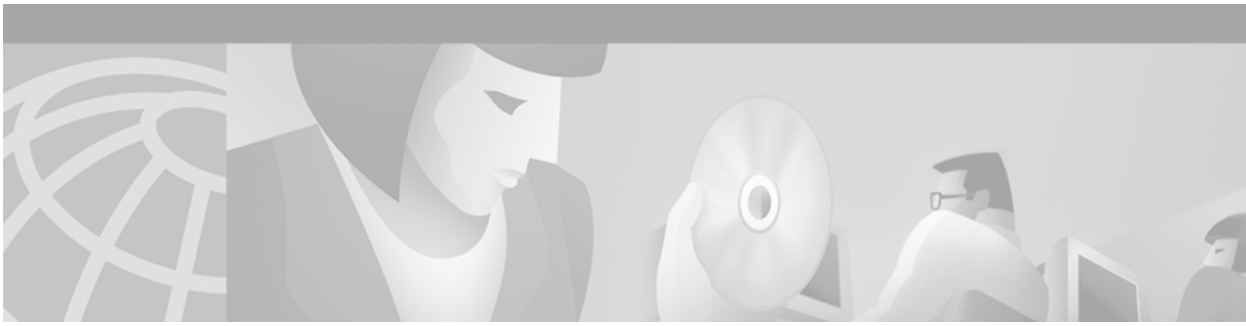

MPLS VPN—Carrier Supporting Carrier—IPv4 BGP Label Distribution

Feature History

Release	Modification
12.0(21)ST	This feature was introduced.
12.0(22)S	This feature was integrated into Cisco IOS Release 12.0(22)S. Support for the Cisco 12000 series line cards was added. (See Table 1 for the specific line cards supported.)

This feature enables you to configure your carrier supporting carrier network to enable Border Gateway Protocol (BGP) to transport routes and Multiprotocol Label Switching (MPLS) labels between the backbone carrier provider edge (PE) routers and the customer carrier customer edge (CE) routers. Previously you had to use Label Distribution Protocol (LDP) to carry the labels and an internal gateway protocol (IGP) to carry the routes between PE and CE routers to achieve the same goal.

This feature is an extension of the Carrier Supporting Carrier feature, introduced in the 12.0(14)ST, which was based on LDP. See *MPLS VPN Carrier Supporting Carrier* for more information.

This feature module includes the following sections:

Feature Overview

This feature enables you to configure a carrier supporting carrier network that uses BGP to distribute routes and MPLS labels between the PE and CE routers of a backbone carrier and a customer carrier. The backbone carrier offers BGP and MPLS VPN services. The customer carrier can be either:

- An Internet service provider (ISP) with an IP core
- An MPLS service provider with or without VPN services

This document describes how to use BGP to distribute MPLS labels and routes for both types of customer carrier.

Using BGP to Distribute Routes and MPLS Labels for the Customer Carrier Who Is an ISP

Figure 1 shows a network configuration where the customer carrier is an ISP. The customer carrier has two sites, each of which is a point of presence (POP). The customer carrier connects these sites using a VPN service provided by the backbone carrier. The backbone carrier uses MPLS. The ISP sites use IP.

Figure 1 Network Where the Customer Carrier Is an ISP

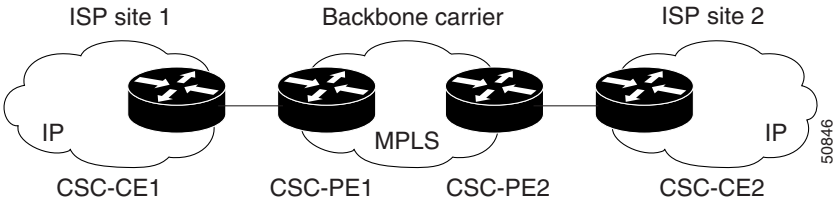

In this configuration, the links between the CE and PE routers use EBGP to distribute IPv4 routes and MPLS labels. Between the links, the PE routers use multiprotocol IBGP to distribute VPNv4 routes.

Note If a non-Cisco router is used as a CSC-PE or CSC-CE, that router must support IPv4 BGP label distribution (RFC 3107). Otherwise, you cannot run EBGP with labels between the routers.

Using BGP to Distribute Routes and MPLS Labels for a Customer Carrier Who Is an MPLS VPN Service Provider

Figure 2 shows a network configuration where the backbone carrier and the customer carrier are BGP/MPLS VPN service providers. The customer carrier has two sites. Both the backbone carrier and the customer carrier use MPLS in their networks.

Figure 2 Network Where the Customer Carrier Is an MPLS VPN Service Provider

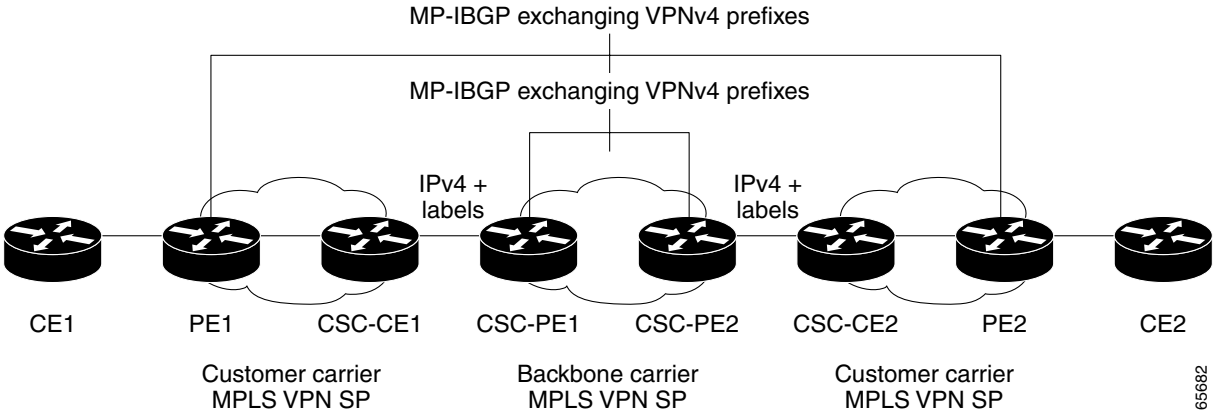

In this configuration, the customer carrier can configure their network in one of the following ways:

- The customer carrier can run IGP and LDP in its core network. In this case, the customer carrier's CSC-CE1 router redistributes the EBGP routes it learned from the backbone carrier's CSC-PE1 to IGP.
- The customer carrier's CSC-CE1 router can run IPv4 + labels IBGP session with the PE1 router.

BGP Overview

BGP routing information includes the following items:

- A network number (prefix), which is the IP address of the destination.
- Autonomous system (AS) path, which is a list of the other ASs through which a route passes on its way to the local router. The first AS in the list is closest to the local router; the last AS in the list is farthest from the local router and usually the AS where the route began.
- Path attributes, which provide other information about the AS path, for example, the next hop.

Types of BGP Messages

MPLS labels are included in the "update" messages that a router sends. Routers exchange the following types of BGP messages:

- Open Messages—After a router establishes a TCP connection with a neighboring router, the routers exchange "open" messages. This message contains the AS number to which the router belongs and the IP address of the router who sent the message.

- Update Messages—When a router has a new, changed, or broken route, it sends an "update" message to the neighboring router. This message contains the Network Layer Reachability Information (NLRI), which lists the IP addresses of the usable routes. The update message also includes any routes that are no longer usable. The update message also includes path attributes and the lengths of both the usable and unusable paths. Labels for VPNv4 routes are encoded in the update message as specified in RFC 2858. The labels for the IPv4 routes are encoded in the update message as specified in RFC 3107.
- Keepalive Messages—Routers exchange "keepalive" messages to determine if a neighboring router is still available to exchange routing information. The router sends these messages at regular intervals. (Sixty seconds is the default for Cisco routers.) The keepalive message does not contain routing data; it only contains a message header.
- Notification Messages—When a router detects an error, it sends a "notification" message.

How BGP Sends MPLS Labels with Routes

When BGP (both External BGP (EBGP) and Internal BGP (IBGP)) distributes a route, it can also distribute an MPLS label that is mapped to that route. The MPLS label mapping information for the route is carried in the BGP update message that contains the information about the route. If the next hop is not changed, the label is preserved.

When you issue the **neighbor send-label** command on both BPG routers, the routers advertise to each other that they can then send MPLS labels with the routes. If the routers successfully negotiate their ability to send MPLS labels, the routers add MPLS labels to all outgoing BGP updates.

Using Route Maps to Filter Routes

When both routers are configured to distribute routes with MPLS labels, all the routes are encoded with the multiprotocol extensions and contain an MPLS label. You can use a route map to control the distribution of MPLS labels between routers. Route maps enable you to specify the following:

- For a router distributing MPLS labels, you can specify which routes are distributed with an MPLS label.
- For a router receiving MPLS labels, you can specify which routes are accepted and installed in the BGP table.

Benefits

Using BGP to distribute IPv4 routes and MPLS labels routes has the following benefits:

- BGP takes the place of an IGP and LDP. You can use BGP to distribute routes and MPLS labels. Using a single protocol instead of two simplifies the configuration and troubleshooting.
- BGP is the preferred routing protocol for connecting two ISPs, mainly because of its routing policies and ability to scale. ISPs commonly use BGP between two providers. This feature enables those ISPs to use BGP.

Restrictions

This feature includes the following restrictions:

- On a PE router, an interface can be configured for either BGP with labels or LDP. You cannot enable both protocols on the same interface. If you switch from one protocol to the other, then the existing protocol must be disabled on all interfaces before you enable the other protocol.
- This feature does not currently support multiple BGP routes to a given destination with different MPLS labels as described in Section 4 of RFC 3107.
- EBGP multihop between CSC-PE and CSC-CE is not supported for this release.
- This feature is not supported with eiBGP multipath.
- The physical interfaces that connect the BGP speakers must support Cisco Express Forwarding (CEF) or distributed Cisco Express Forwarding (DCEF) and MPLS.

Related Features and Technologies

The MPLS VPN Carrier Supporting Carrier feature is used with the VPN capabilities of MPLS. MPLS VPNs were introduced in Cisco IOS Release 12.0(5)T.

This feature is an extension of the feature MPLS VPN Carrier Supporting Carrier, Cisco IOS Release 12.0(14)ST.

Related Documents

- *MPLS VPN Carrier Supporting Carrier*
- *Cisco IOS IP Configuration Guide*, Release 12.2, Configuring BGP chapter
- *Using the Border Gateway Protocol for Interdomain Routing*
- *Internetworking Technology Overview*, Border Gateway Protocol chapter
- *Cisco IOS Switching Services Configuration Guide*, Release 12.2, Configuring Multiprotocol Label Switching chapter
- *Cisco IOS IP Command Reference, Volume 2 of 3: Routing Protocols*, Release 12.2, BGP Commands chapter
- *MPLS Virtual Private Networks*

Supported Platforms

The following router platforms are supported at the service provider edge:

- Cisco 7200 series
- Cisco 7500 series
- Cisco 12000 series routers

  See Table 1 for Cisco 12000 series line card support added for Cisco IOS Release 12.0(22)S.

Table 1 Cisco I2000 Series Line Card Support Added for Cisco IOS Release 12.0(22)S

Type	Line Cards
Packet Over SONET (POS)	4-Port OC-3 POS 8-Port OC-3 POS 16-Port OC-3 POS 1-Port OC-12 POS 4-Port OC-12 POS 1-Port OC-48 POS 4-Port OC-3 POS ISE 8-Port OC-3 POS ISE 16 x OC-3 POS ISE 4 Port OC-12 POS ISE 1-Port OC-48 POS ISE
Electrical Interface	6- Port DS3 12- Port DS3 6-Port E3 12-Port E3
Asynchronous Transfer Mode (ATM)	4-Port OC-3 ATM 1-Port OC12 ATM 4-Port OC-12 ATM
Channelized Interface	2-Port CHOC-3 6-Port Ch T3 (DS1) 1-Port CHOC-12 (DS3) 1-Port CHOC-12 (OC-3) 4-Port CHOC-12 ISE 1-Port CHOC-48 ISE

Determining Platform Support Through Cisco Feature Navigator

Cisco IOS software is packaged in feature sets that support specific platforms. To get updated information regarding platform support for this feature, access Cisco Feature Navigator. Cisco Feature Navigator dynamically updates the list of supported platforms as new platform support is added for the feature.

Cisco Feature Navigator is a web-based tool that enables you to determine which Cisco IOS software images support a specific set of features and which features are supported in a specific Cisco IOS image. You can search by feature or release. Under the release section, you can compare releases side by side to display both the features unique to each software release and the features in common.

To access Cisco Feature Navigator, you must have an account on Cisco.com. If you have forgotten or lost your account information, send a blank e-mail to cco-locksmith@cisco.com. An automatic check will verify that your e-mail address is registered with Cisco.com. If the check is successful, account details with a new random password will be e-mailed to you. Qualified users can establish an account on Cisco.com by following the directions at http://www.cisco.com/register.

Cisco Feature Navigator is updated regularly when major Cisco IOS software releases and technology releases occur. For the most current information, go to the Cisco Feature Navigator home page at the following URL:

http://www.cisco.com/go/fn

Availability of Cisco IOS Software Images

Platform support for particular Cisco IOS software releases is dependent on the availability of the software images for those platforms. Software images for some platforms may be deferred, delayed, or changed without prior notice. For updated information about platform support and availability of software images for each Cisco IOS software release, refer to the online release notes or, if supported, Cisco Feature Navigator.

Supported Standards, MIBs, and RFCs

Standards

No new or modified standards are supported by this feature.

MIBs

No new or modified MIBs are supported by this feature.

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:

http://www.cisco.com/public/sw-center/netmgmt/cmtk/mibs.shtml

RFCs

- RFC 3107, *Carrying Label Information in BGP-4*
- RFC 2858, *Multiprotocol Extensions for BGP-4*
- RFC 1700, *Assigned Numbers*
- RFC 2842, *Capabilities Advertisement with BGP-4*
- RFC 1966, *BGP Route Reflection: An Alternative to Full Mesh IBGP*
- RFC 1171, *A Border Gateway Protocol 4*
- RFC 1164, *Application of the Border Gateway Protocol in the Internet*
- RFC 2283, *Multiprotocol Extensions for BGP-4*
- RFC 2547, *BGP/MPLS VPNs*

Prerequisites

- The network must be properly configured for MPLS VPN operation before you configure this feature.
- Make sure that the CSC-PE and CSC-CE routers run images that support BGP label distribution. Otherwise, you cannot run EBGP between them.

Configuration Tasks

See the following sections for configuration tasks for this feature. Each task in the list is identified as either required or optional.

- Configuring the CSC-PE Routers to Exchange IPv4 Routes and MPLS Labels (required)

- Configuring the CSC-CE Routers to Exchange IPv4 Routes and MPLS Labels (required)
- Creating Route Maps (optional)
- Applying the Route Maps to the CSC-PE Routers (optional)
- Applying the Route Maps to the CSC-CE Routers (optional)
- Verifying the Configuration (optional)

Configuring the CSC-PE Routers to Exchange IPv4 Routes and MPLS Labels

To configure the CSC-PEs so that they can distribute BGP routes with MPLS labels, use the following commands beginning in user EXEC mode:

	Command	Purpose
Step 1	`Router# configure terminal`	Enters global configuration mode.
Step 2	`Router(config)# router bgp` *as-number*	Enables the BGP routing process and assigns the router to the AS number you specify. The AS number identifies the router to routers in other ASs.
Step 3	`Router(config-router)# address-family ipv4 vrf` *vrf-name*	Enters address family submode and enables the exchange of IPv4 NLRI.
Step 4	`Router(config-router-af)# neighbor` *ip-address* `remote-as` *as-number*	Specifies the neighboring BGP router and the AS number to which the neighboring BGP router belongs.
Step 5	`Router(config-router-af)# neighbor` *ip-address* `activate`	Enables the exchange or routes with the neighboring BGP router.
Step 6	`Router(config-router-af)# neighbor` *ip-address* `send-label`	Advertises the router's ability to send MPLS labels with routes.
Step 7	`Router(config-router-af)# exit-address-family`	Exits from the address family submode.

Configuring the CSC-CE Routers to Exchange IPv4 Routes and MPLS Labels

To configure the CSC-CE routers so that they can distribute BGP routes with MPLS labels, use the following commands beginning in user EXEC mode:

	Command	Purpose
Step 1	`Router# configure terminal`	Enters global configuration mode.
Step 2	`Router(config)# router bgp` *as-number*	Enables the BGP routing process and assigns the router to the AS number you specify. The AS number identifies the router to routers in other ASs.
Step 3	`Router(config-router)# address-family ipv4`	Enters address family submode and enables the exchange of IPv4 NLRI.

	Command	Purpose
Step 4	`Router(config-router-af)# neighbor ip-address remote-as as-number`	Specifies the neighboring BGP router and the AS number to which the neighboring BGP router belongs.
Step 5	`Router(config-router-af)# neighbor ip-address activate`	Enables the exchange or routes with the neighboring BGP router.
Step 6	`Router(config-router-af)# neighbor ip-address send-label`	Advertises the router's ability to send MPLS labels with routes.
Step 7	`Router(config-router-af)# exit-address-family`	Exits from the address family submode.

Creating Route Maps

The following procedures enable the routers to send MPLS labels with the routes specified in the route maps. Further, the routers accept the only routes that are specified in the route map.

Route maps enable you to specify which routes are distributed with MPLS labels. Route maps also enable you to specify which routes with MPLS labels a router receives and adds to its BGP table.

Route maps work with access lists. You enter the routes into an access list and then specify the access list when you configure the route map.

Configuring a Route Map for Arriving Routes

To configure a route map to filter for arriving routes, use the following commands beginning in user EXEC mode. You create an access list and specify the routes that the router should accept and add to the BGP table.

	Command	Purpose
Step 1	`Router# configure terminal`	Enters global configuration mode.
Step 2	`Router(config)# router bgp` *as-number*	Enters router configuration mode.
Step 3	`Router(config-router)# route-map` *route-map-name* `permit` *sequence-number*	Creates a route map with the name you specify. The **permit** keyword allows the actions to happen if all conditions are met. A **deny** keyword prevents any actions from happening if all conditions are met. The *sequence-number* parameter allows you to prioritize route maps. If you have multiple route maps and want to prioritize them, assign each one a number. The route map with the lowest number is implemented first, followed by the route map with the second lowest number, and so on.
Step 4	`Router(config-route-map)# match ip address` *access-list-name*	Checks the route listed in the BGP update message against the list of routes in the ACL you specify. If a route in the BGP update message matches a route in the ACL, the route is accepted by the router and added to the BGP table.
Step 5	`Router(config-route-map)# match mpls-label`	Specifies that the route must have a label to be installed into the BGP table.

Configuring a Route Map for Departing Routes

To configure a route map to filter for departing routes, use the following commands beginning in user EXEC mode. You create an access list and specify the routes that the router should distribute with MPLS labels.

	Command	Purpose
Step 1	`Router# configure terminal`	Enters global configuration mode.
Step 2	`Router(config)# router bgp` *as-number*	Enters router configuration mode.
Step 3	`Router(config-router)# route-map` *route-map-name* `permit` *sequence-number*	Creates a route map with the name you specify. The **permit** keyword allows the actions to happen if all conditions are met. A **deny** keyword prevents the actions from happening if all conditions are met. The *sequence-number* parameter allows you to prioritize route maps. If you have multiple route maps and want to prioritize them, assign each one a number. The route map with the lowest number is implemented first, followed by the route map with the second lowest number, and so on.

	Command	Purpose
Step 4	`Router(config-route-map)# match ip address access-list-name`	Checks the route listed in the BGP update message against the list of routes in the ACL you specify. If a route in the BGP update message matches a route in the ACL, the router distributes the route. This action is explicitly stated in the next command.
Step 5	`Router(config-route-map)# set mpls-label`	Specifies that the router distributes the route with an MPLS label.

Applying the Route Maps to the CSC-PE Routers

To enable the CSC-PE routers to use the route maps, use the following commands beginning in user EXEC mode:

	Command	Purpose
Step 1	`Router# configure terminal`	Enters global configuration mode.
Step 2	`Router(config)# router bgp as-number`	Enters router configuration mode.
Step 3	`Router(config-router)# address-family ipv4 vrf vrf-name`	Enters address family submode and enables the exchange of IPv4 NLRI.
Step 4	`Router(config-router-af)# neighbor ip-address route-map route-map-name in`	Applies the route map you specify to the routers that the router receives (incoming routes).
Step 5	`Router(config-router-af)# neighbor ip-address route-map route-map-name out`	Applies the route map you specify to the routers that the router sends out (outgoing routes).
Step 6	`Router(config-router-af)# neighbor ip-address send-label`	Advertises the router's ability to send MPLS labels with routes.
Step 7	`Router(config-router-af)# exit-address-family`	Exits from the address family submode.

Applying the Route Maps to the CSC-CE Routers

To enable the CSC-CE routers to use the route maps, use the following commands beginning in user EXEC mode:

	Command	Purpose
Step 1	`Router# configure terminal`	Enters global configuration mode.
Step 2	`Router(config)# router bgp as-number`	Enters router configuration mode.
Step 3	`Router(config-router)# address-family ipv4`	Enters address family submode and enables the exchange of IPv4 NLRI.
Step 4	`Router(config-router-af)# neighbor ip-address route-map route-map-name in`	Applies the route map you specify to the routers that the router receives (incoming routes).
Step 5	`Router(config-router-af)# neighbor ip-address route-map route-map-name out`	Applies the route map you specify to the routers that the router sends out (outgoing routes).

	Command	Purpose
Step 6	`Router(config-router-af)# neighbor ip-address send-label`	Advertises the router's ability to send MPLS labels with routes.
Step 7	`Router(config-router-af)# exit-address-family`	Exits from the address family submode.

Verifying the Configuration

This section explains how to check that the customer and backbone carrier routers are configured to enable BGP to distribute routes and MPLS labels. Use Figure 3 as a reference for verifying the configuration.

Figure 3 Configuring CSC Routers to Exchange IPv4 Routes and MPLS Labels

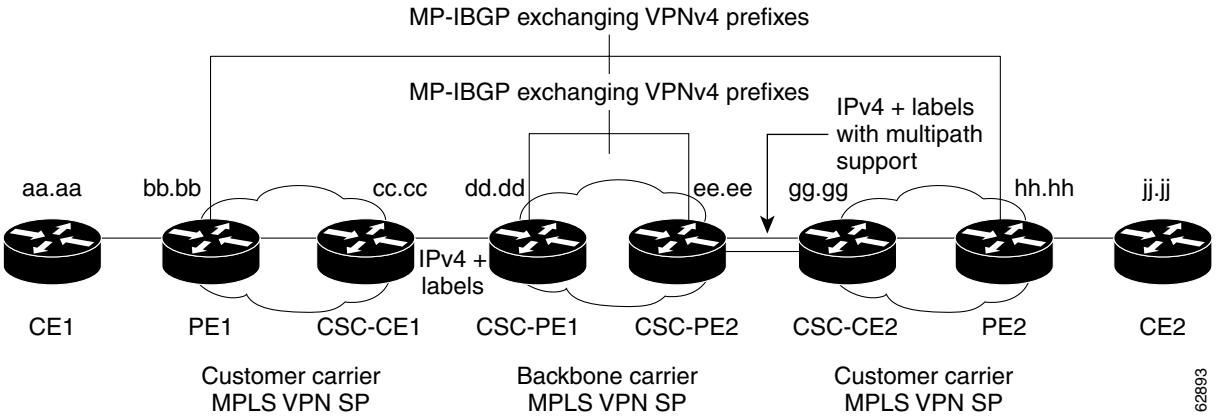

Verifying the CE1 Router Configuration

Step 1 Verify that the loopback address of the remote CE router (CE2), learned from the PE router, is in the routing table of the CE1 router.

```
Router# show ip route jj.jj.jj.jj

Routing entry for jj.jj.jj.jj/32
  Known via "bgp 300", distance 20, metric 0
  Tag 200, type external
  Redistributing via ospf 300
  Advertised by ospf 300 subnets
  Last update from mm.0.0.1 20:29:35 ago
  Routing Descriptor Blocks:
  * mm.0.0.1, from mm.0.0.1, 20:29:35 ago
      Route metric is 0, traffic share count is 1
      AS Hops 2
```

Verifying the PE1 Router Configuration

Step 1 Verify that the loopback address of the local CE router (CE1) is in the routing table of the PE1 router.

```
Router# show ip route vrf <PE-vrf-name> aa.aa.aa.aa

Routing entry for aa.aa.aa.aa/32
  Known via "bgp 200", distance 20, metric 0
  Tag 300, type external
  Last update from mm.0.0.2 20:36:59 ago
  Routing Descriptor Blocks:
  * mm.0.0.2, from mm.0.0.2, 20:36:59 ago
      Route metric is 0, traffic share count is 1
      AS Hops 1, BGP network version 0
```

Step 2 Verify that the prefix for the local CE router (CE1) is in the MPLS forwarding table, and that the prefix is untagged.

```
Router# show mpls forwarding-table vrf <PE-vrf-name> aa.aa.aa.aa

Local  Outgoing    Prefix            Bytes tag  Outgoing   Next Hop
tag    tag or VC   or Tunnel Id      switched   interface
23     Untagged    aa.aa.aa.aa/32[V] 0          Et3/3      mm.0.0.2
```

Step 3 Verify that the prefix of the remote PE router (PE2) is in the Cisco Express Forwarding (CEF) table.

```
Router# show ip cef hh.hh.hh.hh

hh.hh.hh.hh/32, version 31, cached adjacency nn.0.0.2
0 packets, 0 bytes
  tag information set
    local tag: 31
    fast tag rewrite with Et3/0, nn.0.0.2, tags imposed {26}
  via nn.0.0.2, Ethernet3/0, 2 dependencies
    next hop nn.0.0.2, Ethernet3/0
    unresolved
    valid cached adjacency
    tag rewrite with Et3/0, nn.0.0.2, tags imposed {26}
```

Step 4 Verify that the loopback address of the remote CE router (CE2) is in the routing table.

```
Router# show ip route vrf <PE-vrf-name> jj.jj.jj.jj

Routing entry for jj.jj.jj.jj/32
  Known via "bgp 200", distance 200, metric 0
  Tag 300, type internal
  Last update from hh.hh.hh.hh 20:38:49 ago
  Routing Descriptor Blocks:
  * hh.hh.hh.hh (Default-IP-Routing-Table), from hh.hh.hh.hh, 20:38:49 ago
      Route metric is 0, traffic share count is 1
      AS Hops 1, BGP network version 0
```

Step 5 Verify that the prefix of the remote CE router (CE2) is in the MPLS forwarding table, and that an outgoing interface exists.

```
Router# show mpls forwarding-table vrf <PE-vrf-name> jj.jj.jj.jj

Local  Outgoing    Prefix            Bytes tag  Outgoing   Next Hop
tag    tag or VC   or Tunnel Id      switched   interface
None   26          jj.jj.jj.jj/32    0          Et3/0      nn.0.0.2
```

Step 6 Verify that the prefix of the remote CE router (CE2) is in the CEF table.

```
Router# show ip cef vrf <PE-vrf-name> jj.jj.jj.jj

jj.jj.jj.jj/32, version 12, cached adjacency nn.0.0.2
0 packets, 0 bytes
  tag information set
    local tag: VPN route head
    fast tag rewrite with Et3/0, nn.0.0.2, tags imposed {26 32}
  via hh.hh.hh.hh, 0 dependencies, recursive
    next hop nn.0.0.2, Ethernet3/0 via hh.hh.hh.hh/32
    valid cached adjacency
    tag rewrite with Et3/0, nn.0.0.2, tags imposed {26 32}
```

Step 7 Verify that the prefix of the local PE router (PE1) is in the CEF table.

```
Router# show ip cef bb.bb.bb.bb

bb.bb.bb.bb/32, version 9, connected, receive
  tag information set
    local tag: implicit-null
```

Verifying the CSC-CE1 Router Configuration

Step 1 Verify that the BGP session is up and running.

```
Router# show ip bgp summary

BGP router identifier cc.cc.cc.cc, local AS number 200
BGP table version is 35, main routing table version 35
14 network entries and 14 paths using 2030 bytes of memory
3 BGP path attribute entries using 168 bytes of memory
1 BGP AS-PATH entries using 24 bytes of memory
0 BGP route-map cache entries using 0 bytes of memory
0 BGP filter-list cache entries using 0 bytes of memory
Dampening enabled. 1 history paths, 0 dampened paths
BGP activity 17/67 prefixes, 29/15 paths, scan interval 60 secs

Neighbor        V    AS MsgRcvd MsgSent   TblVer  InQ OutQ Up/Down  State/PfxRcd
pp.0.0.1        4   100    7615    7613       35    0    0 21:06:19        5
```

Step 2 Verify that the loopback address of the local PE router (PE1) is in the routing table.

```
Router# show ip route bb.bb.bb.bb

Routing entry for bb.bb.bb.bb/32
  Known via "ospf 200", distance 110, metric 101, type intra area
  Redistributing via bgp 200
  Advertised by bgp 200 metric 4 match internal
  Last update from nn.0.0.1 on Ethernet4/0, 00:34:08 ago
  Routing Descriptor Blocks:
  * nn.0.0.1, from bb.bb.bb.bb, 00:34:08 ago, via Ethernet4/0
      Route metric is 101, traffic share count is 1
```

Step 3 Verify that the loopback address of the remote PE router (PE2) is in the routing table.

```
Router# show ip route hh.hh.hh.hh

Routing entry for hh.hh.hh.hh/32
  Known via "bgp 200", distance 20, metric 0
  Tag 100, type external
  Redistributing via ospf 200
  Advertised by ospf 200 metric 3 subnets
  Last update from pp.0.0.1 00:45:16 ago
  Routing Descriptor Blocks:
  * pp.0.0.1, from pp.0.0.1, 00:45:16 ago
      Route metric is 0, traffic share count is 1
      AS Hops 2, BGP network version 0
```

Step 4 Verify that the prefix of the local PE router (PE1) is in the MPLS LDP bindings.

```
Router# show mpls ldp bindings bb.bb.bb.bb 255.255.255.255

  tib entry: bb.bb.bb.bb/32, rev 20
        local binding:  tag: 20
        remote binding: tsr: bb.bb.bb.bb:0, tag: imp-null
```

Step 5 Verify that the prefix of the local PE router (PE1) is in the CEF table.

```
Router# show ip cef bb.bb.bb.bb

bb.bb.bb.bb/32, version 46, cached adjacency nn.0.0.1
0 packets, 0 bytes
  tag information set
    local tag: 20
  via nn.0.0.1, Ethernet4/0, 0 dependencies
    next hop nn.0.0.1, Ethernet4/0
    unresolved
    valid cached adjacency
    tag rewrite with Et4/0, nn.0.0.1, tags imposed {}
```

Step 6 Verify that the prefix of the local PE router (PE1) is in the MPLS forwarding table.

```
Router# show mpls forwarding-table bb.bb.bb.bb

Local  Outgoing    Prefix            Bytes tag  Outgoing   Next Hop
tag    tag or VC   or Tunnel Id      switched   interface
20     Pop tag     bb.bb.bb.bb/32    893397     Et4/0      nn.0.0.1

Router# show mpls forwarding-table bb.bb.bb.bb detail

Local  Outgoing    Prefix            Bytes tag  Outgoing   Next Hop
tag    tag or VC   or Tunnel Id      switched   interface
20     Pop tag     bb.bb.bb.bb/32    893524     Et4/0      nn.0.0.1
        MAC/Encaps=14/14, MTU=1504, Tag Stack{}
        00074F83685400B04A74A0708847
        No output feature configured
    Per-packet load-sharing, slots: 0 1 2 3 4 5 6 7 8 9 10 11 12 13 14 15
```

Step 7 Verify that the BGP routing table contains labels for prefixes in the customer carrier MPLS VPN service provider networks.

```
Router# show ip bgp labels

Network          Next Hop         In Label/Out Label
cc.cc.cc.cc/32   0.0.0.0          imp-null/exp-null
bb.bb.bb.bb/32   nn.0.0.1         20/exp-null
hh.hh.hh.hh/32   pp.0.0.1         26/34
gg.gg.gg.gg/32   pp.0.0.1         23/30
nn.0.0.0         0.0.0.0          imp-null/exp-null
ss.0.0.0         pp.0.0.1         25/33
pp.0.0.0         0.0.0.0          imp-null/exp-null
pp.0.0.1/32      0.0.0.0          16/exp-null
```

Step 8 Verify that the prefix of the remote PE router (PE2) is in the CEF table.

```
Router# show ip cef hh.hh.hh.hh

hh.hh.hh.hh/32, version 54, cached adjacency pp.0.0.1
0 packets, 0 bytes
  tag information set
    local tag: 26
    fast tag rewrite with Et3/0, pp.0.0.1, tags imposed {34}
  via pp.0.0.1, 0 dependencies, recursive
    next hop pp.0.0.1, Ethernet3/0 via pp.0.0.1/32
    valid cached adjacency
    tag rewrite with Et3/0, pp.0.0.1, tags imposed {34}
```

Step 9 Verify that the prefix of the remote PE router (PE2) is in the MPLS forwarding table.

```
Router# show mpls forwarding-table hh.hh.hh.hh

Local  Outgoing    Prefix           Bytes tag  Outgoing   Next Hop
tag    tag or VC   or Tunnel Id     switched   interface
26     34          hh.hh.hh.hh/32   81786      Et3/0      pp.0.0.1

Router# show mpls forwarding-table hh.hh.hh.hh detail

Local  Outgoing    Prefix           Bytes tag  Outgoing   Next Hop
tag    tag or VC   or Tunnel Id     switched   interface
26     34          hh.hh.hh.hh/32   81863      Et3/0      pp.0.0.1
        MAC/Encaps=14/18, MTU=1500, Tag Stack{34}
        00B0C26E105500B04A74A0548847 00022000
        No output feature configured
    Per-packet load-sharing, slots: 0 1 2 3 4 5 6 7 8 9 10 11 12 13 14 15
```

Verifying the CSC-PE1 Router Configuration

Step 1 Verify that the BGP session is up and running between the CSC-PE1 router and the CSC-CE1 router. Verify that prefixes are learned during each session, by checking the data in the State/PfxRcd column.

```
Router# show ip bgp vpnv4 all summary

BBGP router identifier dd.dd.dd.dd, local AS number 100
BGP table version is 52, main routing table version 52
12 network entries and 13 paths using 2232 bytes of memory
6 BGP path attribute entries using 336 bytes of memory
1 BGP AS-PATH entries using 24 bytes of memory
1 BGP extended community entries using 24 bytes of memory
0 BGP route-map cache entries using 0 bytes of memory
0 BGP filter-list cache entries using 0 bytes of memory
Dampening enabled. 0 history paths, 0 dampened paths
BGP activity 16/4 prefixes, 27/14 paths, scan interval 5 secs


Neighbor        V    AS MsgRcvd MsgSent   TblVer  InQ OutQ Up/Down  State/PfxRcd
ee.ee.ee.ee     4   100    7685    7686       52    0    0 21:17:04        6
pp.0.0.2        4   200    7676    7678       52    0    0 21:16:43        7
```

Step 2 Verify that the MPLS interfaces are up and running, and that LDP-enabled interfaces show that LDP is up and running. LDP is turned off on the VRF because EBGP distributes the labels.

```
Router# show mpls interfaces all

Interface              IP            Tunnel   Operational
GigabitEthernet6/0     Yes (ldp)     No       Yes


VRF vpn1:
Ethernet3/1            No            No       Yes
```

Step 3 Verify that the prefix for the PE1 router is in the routing table of the CSC-PE1 router.

```
Router# show ip route vrf <CSC-PE-vrf-name> bb.bb.bb.bb

Routing entry for bb.bb.bb.bb/32
  Known via "bgp 100", distance 20, metric 4
  Tag 200, type external
  Last update from pp.0.0.2 21:28:39 ago
  Routing Descriptor Blocks:
  * pp.0.0.2, from pp.0.0.2, 21:28:39 ago
      Route metric is 4, traffic share count is 1
      AS Hops 1, BGP network version 0
```

Step 4 Verify that the prefix for the PE2 router is in the routing table of the CSC-PE1 router.

```
Router# show ip route vrf <CSC-PE-vrf-name> hh.hh.hh.hh

Routing entry for hh.hh.hh.hh/32
  Known via "bgp 100", distance 200, metric 4
  Tag 200, type internal
  Last update from ee.ee.ee.ee 21:27:39 ago
  Routing Descriptor Blocks:
  * ee.ee.ee.ee (Default-IP-Routing-Table), from ee.ee.ee.ee, 21:27:39 ago
      Route metric is 4, traffic share count is 1
      AS Hops 1, BGP network version 0
```

Step 5 Verify that the prefixes for the customer carrier MPLS VPN service provider networks are in the BGP table, and have appropriate labels.

```
Router# show ip bgp vpnv4 vrf <CSC-PE-vrf-name> labels

   Network          Next Hop      In label/Out label
Route Distinguisher: 100:1 (vpn1)
   cc.cc.cc.cc/32   pp.0.0.2        22/imp-null
   bb.bb.bb.bb/32   pp.0.0.2        27/20
   hh.hh.hh.hh/32   ee.ee.ee.ee     34/35
   gg.gg.gg.gg/32   ee.ee.ee.ee     30/30
   nn.0.0.0         pp.0.0.2        23/imp-null
   ss.0.0.0         ee.ee.ee.ee     33/34
   pp.0.0.0         pp.0.0.2        25/aggregate(vpn1)
```

Step 6 Verify that the prefix of the PE router in the local customer carrier MPLS VPN service provider (PE1) is in the CEF table.

```
Router# show ip cef vrf <CSC-PE-vrf-name> bb.bb.bb.bb
bb.bb.bb.bb/32, version 19, cached adjacency pp.0.0.2
0 packets, 0 bytes
  tag information set
    local tag: 27
    fast tag rewrite with Et3/1, pp.0.0.2, tags imposed {20}
  via pp.0.0.2, 0 dependencies, recursive
    next hop pp.0.0.2, Ethernet3/1 via pp.0.0.2/32
    valid cached adjacency
    tag rewrite with Et3/1, pp.0.0.2, tags imposed {20}

Router# show ip cef vrf <CSC-PE-vrf-name> bb.bb.bb.bb detail

bb.bb.bb.bb/32, version 19, cached adjacency pp.0.0.2
0 packets, 0 bytes
  tag information set
    local tag: 27
    fast tag rewrite with Et3/1, pp.0.0.2, tags imposed {20}
  via pp.0.0.2, 0 dependencies, recursive
    next hop pp.0.0.2, Ethernet3/1 via pp.0.0.2/32
    valid cached adjacency
    tag rewrite with Et3/1, pp.0.0.2, tags imposed {20}
```

Step 7 Verify that the prefix of the PE router in the local customer carrier MPLS VPN service provider (PE1) is in the MPLS forwarding table.

```
Router# show mpls forwarding-table vrf <CSC-PE-vrf-name> bb.bb.bb.bb

Local  Outgoing    Prefix            Bytes tag  Outgoing   Next Hop
tag    tag or VC   or Tunnel Id      switched   interface
27     20          bb.bb.bb.bb/32[V] 958048     Et3/1      pp.0.0.2

Router# show mpls forwarding-table vrf <CSC-PE-vrf-name> bb.bb.bb.bb detail

Local  Outgoing    Prefix            Bytes tag  Outgoing   Next Hop
tag    tag or VC   or Tunnel Id      switched   interface
27     20          bb.bb.bb.bb/32[V] 958125     Et3/1      pp.0.0.2
        MAC/Encaps=14/18, MTU=1500, Tag Stack{20}
        00B04A74A05400B0C26E10558847 00014000
        VPN route: vpn1
        No output feature configured
    Per-packet load-sharing, slots: 0 1 2 3 4 5 6 7 8 9 10 11 12 13 14 15
```

Step 8 Verify that the prefix of the PE router in the remote customer carrier MPLS VPN service provider (PE2) is in the CEF table.

```
Router# show ip cef vrf <CSC-PE-vrf-name> hh.hh.hh.hh

hh.hh.hh.hh/32, version 25, cached adjacency rr.0.0.2
0 packets, 0 bytes
  tag information set
    local tag: 34
    fast tag rewrite with Gi6/0, rr.0.0.2, tags imposed {35}
  via ee.ee.ee.ee, 0 dependencies, recursive
    next hop rr.0.0.2, GigabitEthernet6/0 via ee.ee.ee.ee/32
    valid cached adjacency
    tag rewrite with Gi6/0, rr.0.0.2, tags imposed {35}

Router# show ip cef vrf <CSC-PE-vrf-name> hh.hh.hh.hh detail

hh.hh.hh.hh/32, version 25, cached adjacency rr.0.0.2
0 packets, 0 bytes
  tag information set
    local tag: 34
    fast tag rewrite with Gi6/0, rr.0.0.2, tags imposed {35}
  via ee.ee.ee.ee, 0 dependencies, recursive
    next hop rr.0.0.2, GigabitEthernet6/0 via ee.ee.ee.ee/32
    valid cached adjacency
    tag rewrite with Gi6/0, rr.0.0.2, tags imposed {35}
```

Step 9 Verify that the prefix of the PE router in the remote customer carrier MPLS VPN service provider (PE2) is in the MPLS forwarding table.

```
Router# show mpls forwarding-table vrf <CSC-PE-vrf-name> hh.hh.hh.hh

Local  Outgoing    Prefix            Bytes tag  Outgoing   Next Hop
tag    tag or VC   or Tunnel Id      switched   interface
34     35          hh.hh.hh.hh/32[V] 139034     Gi6/0      rr.0.0.2


Router# show mpls forwarding-table vrf <CSC-PE-vrf-name> hh.hh.hh.hh detail

Local  Outgoing    Prefix            Bytes tag  Outgoing   Next Hop
tag    tag or VC   or Tunnel Id      switched   interface
34     35          hh.hh.hh.hh/32[V] 139034     Gi6/0      rr.0.0.2
        MAC/Encaps=14/18, MTU=1500, Tag Stack{35}
        00B0C26E447000B0C26E10A88847 00023000
        VPN route: vpn1
        No output feature configured
    Per-packet load-sharing, slots: 0 1 2 3 4 5 6 7 8 9 10 11 12 13 14 15
```

Verifying the CSC-PE2 Router Configuration

Step 1 Verify that the BGP session is up and running between the CSC-PE2 router and the CSC-CE2 router. Verify that prefixes are learned during each session, by checking the data in the State/PfxRcd column.

```
Router# show ip bgp vpnv4 all summary

BGP router identifier ee.ee.ee.ee, local AS number 100
BGP table version is 51, main routing table version 51
12 network entries and 13 paths using 2232 bytes of memory
6 BGP path attribute entries using 336 bytes of memory
```

```
1 BGP AS-PATH entries using 24 bytes of memory
1 BGP extended community entries using 24 bytes of memory
0 BGP route-map cache entries using 0 bytes of memory
0 BGP filter-list cache entries using 0 bytes of memory
Dampening enabled. 0 history paths, 0 dampened paths
BGP activity 16/4 prefixes, 31/18 paths, scan interval 5 secs


Neighbor        V    AS MsgRcvd MsgSent   TblVer  InQ OutQ Up/Down  State/PfxRcd
dd.dd.dd.dd     4   100    7901    7900       51    0    0 21:52:59        7
ss.0.0.2        4   200    7871    7880       51    0    0 21:50:15        6
```

Step 2 Verify that the MPLS interfaces are up and running, and that LDP-enabled interfaces show that LDP is up and running. LDP is turned off on the VRF because EBGP distributes the labels.

```
Router# show mpls interfaces all

Interface              IP            Tunnel   Operational
GigabitEthernet4/0     Yes (ldp)     No       Yes


VRF vpn1:
Ethernet5/0            No            No       Yes
```

Step 3 Verify that the prefix of the PE1 router is in the routing table of the CSC-PE2 router.

```
Router# show ip route vrf <CSC-PE-vrf-name> bb.bb.bb.bb.bb

Routing entry for bb.bb.bb.bb/32
  Known via "bgp 100", distance 200, metric 4
  Tag 200, type internal
  Last update from dd.dd.dd.dd 21:53:30 ago
  Routing Descriptor Blocks:
  * dd.dd.dd.dd (Default-IP-Routing-Table), from dd.dd.dd.dd, 21:53:30 ago
      Route metric is 4, traffic share count is 1
      AS Hops 1, BGP network version 0
```

Step 4 Verify that the prefix of the PE2 router is in the routing table of the CSC-PE2 router.

```
Router# show ip route vrf <CSC-PE-vrf-name> hh.hh.hh.hh

Routing entry for hh.hh.hh.hh/32
  Known via "bgp 100", distance 20, metric 4
  Tag 200, type external
  Last update from ss.0.0.2 21:53:12 ago
  Routing Descriptor Blocks:
  * ss.0.0.2, from ss.0.0.2, 21:53:12 ago
      Route metric is 4, traffic share count is 1
      AS Hops 1, BGP network version 0
```

Step 5 Verify that the prefixes for the customer carrier MPLS VPN service provider networks are in the BGP routing table, and that the prefixes have appropriate labels.

```
Router# show ip bgp vpnv4 vrf <CSC-PE-vrf-name> labels

   Network          Next Hop      In label/Out label
Route Distinguisher: 100:1 (vpn1)
   cc.cc.cc.cc/32   dd.dd.dd.dd     27/22
   bb.bb.bb.bb/32   dd.dd.dd.dd     26/27
   hh.hh.hh.hh/32   ss.0.0.2        35/31
   gg.gg.gg.gg/32   ss.0.0.2        30/imp-null
   nn.0.0.0         dd.dd.dd.dd     24/23
   ss.0.0.0         ss.0.0.2        34/aggregate(vpn1)
   pp.0.0.0         dd.dd.dd.dd     21/25
```

Step 6 Verify that the prefix of the PE router in the remote customer carrier MPLS VPN service provider (PE1) is in the CEF table.

```
Router# show ip cef vrf <CSC-PE-vrf-name> bb.bb.bb.bb

bb.bb.bb.bb/32, version 15, cached adjacency rr.0.0.1
0 packets, 0 bytes
  tag information set
    local tag: 26
    fast tag rewrite with Gi4/0, rr.0.0.1, tags imposed {27}
  via dd.dd.dd.dd, 0 dependencies, recursive
    next hop rr.0.0.1, GigabitEthernet4/0 via dd.dd.dd.dd/32
    valid cached adjacency
    tag rewrite with Gi4/0, rr.0.0.1, tags imposed {27}


Router# show ip cef vrf <CSC-PE-vrf-name> bb.bb.bb.bb detail

bb.bb.bb.bb/32, version 15, cached adjacency rr.0.0.1
0 packets, 0 bytes
  tag information set
    local tag: 26
    fast tag rewrite with Gi4/0, rr.0.0.1, tags imposed {27}
  via dd.dd.dd.dd, 0 dependencies, recursive
    next hop rr.0.0.1, GigabitEthernet4/0 via dd.dd.dd.dd/32
    valid cached adjacency
    tag rewrite with Gi4/0, rr.0.0.1, tags imposed {27}
```

Step 7 Verify that the prefix of the PE router in the remote customer carrier MPLS VPN service provider (PE1) is in the MPLS forwarding table.

```
Router# show mpls forwarding-table vrf <CSC-PE-vrf-name> bb.bb.bb.bb

Local  Outgoing    Prefix            Bytes tag  Outgoing   Next Hop
tag    tag or VC   or Tunnel Id      switched   interface
26     27          bb.bb.bb.bb/32[V] 967450     Gi4/0      rr.0.0.1


Router# show mpls forwarding-table vrf <CSC-PE-vrf-name> bb.bb.bb.bb detail

Local  Outgoing    Prefix            Bytes tag  Outgoing   Next Hop
tag    tag or VC   or Tunnel Id      switched   interface
26     27          bb.bb.bb.bb/32[V] 967510     Gi4/0      rr.0.0.1
        MAC/Encaps=14/18, MTU=1500, Tag Stack{27}
        00B0C26E10A800B0C26E44708847 0001B000
        VPN route: vpn1
        No output feature configured
    Per-packet load-sharing, slots: 0 1 2 3 4 5 6 7 8 9 10 11 12 13 14 15
72k-131-9#
```

Step 8 Verify that the prefix of the PE router in the local customer carrier MPLS VPN service provider (PE2) is in the CEF table.

```
Router# show ip cef vrf <CSC-PE-vrf-name> hh.hh.hh.hh

hh.hh.hh.hh/32, version 33, cached adjacency ss.0.0.2
0 packets, 0 bytes
  tag information set
    local tag: 35
    fast tag rewrite with Et5/0, ss.0.0.2, tags imposed {31}
  via ss.0.0.2, 0 dependencies, recursive
    next hop ss.0.0.2, Ethernet5/0 via ss.0.0.2/32
    valid cached adjacency
    tag rewrite with Et5/0, ss.0.0.2, tags imposed {31}

Router# show ip cef vrf <CSC-PE-vrf-name> hh.hh.hh.hh detail

hh.hh.hh.hh/32, version 33, cached adjacency ss.0.0.2
0 packets, 0 bytes
  tag information set
    local tag: 35
    fast tag rewrite with Et5/0, ss.0.0.2, tags imposed {31}
  via ss.0.0.2, 0 dependencies, recursive
    next hop ss.0.0.2, Ethernet5/0 via ss.0.0.2/32
    valid cached adjacency
    tag rewrite with Et5/0, ss.0.0.2, tags imposed {31}
```

Step 9 Verify that the prefix of the PE router in the local customer carrier MPLS VPN service provider (PE2) is in the MPLS forwarding table.

```
Router# show mpls forwarding-table vrf <CSC-PE-vrf-name> hh.hh.hh.hh

Local  Outgoing    Prefix            Bytes tag  Outgoing   Next Hop
tag    tag or VC   or Tunnel Id      switched   interface
35     31          hh.hh.hh.hh/32[V] 2023332    Et5/0      ss.0.0.2


Router# show mpls forwarding-table vrf <CSC-PE-vrf-name> hh.hh.hh.hh detail

Local  Outgoing    Prefix            Bytes tag  Outgoing   Next Hop
tag    tag or VC   or Tunnel Id      switched   interface
35     31          hh.hh.hh.hh/32[V] 2023469    Et5/0      ss.0.0.2
        MAC/Encaps=14/18, MTU=1500, Tag Stack{31}
        0030A363380300B0C26E448C8847 0001F000
        VPN route: vpn1
        No output feature configured
    Per-packet load-sharing, slots: 0 1 2 3 4 5 6 7 8 9 10 11 12 13 14 15
```

Verifying the CSC-CE2 Router Configuration

Step 1 Verify that the BGP session is up and running.

```
Router# show ip bgp summary

BGP router identifier gg.gg.gg.gg, local AS number 200
BGP table version is 31, main routing table version 31
13 network entries and 13 paths using 1885 bytes of memory
3 BGP path attribute entries using 168 bytes of memory
1 BGP AS-PATH entries using 24 bytes of memory
0 BGP route-map cache entries using 0 bytes of memory
0 BGP filter-list cache entries using 0 bytes of memory
Dampening enabled. 0 history paths, 0 dampened paths
BGP activity 17/4 prefixes, 20/7 paths, scan interval 60 secs


Neighbor        V    AS MsgRcvd MsgSent   TblVer  InQ OutQ Up/Down  State/PfxRcd
ss.0.0.1        4   100    7962    7953       31    0    0 22:03:55        6
```

Step 2 Verify that the loopback address of the local PE router (PE2) is in the routing table.

```
Router# show ip route hh.hh.hh.hh

Routing entry for hh.hh.hh.hh/32
  Known via "ospf 200", distance 110, metric 7, type intra area
  Redistributing via bgp 200
  Advertised by bgp 200 metric 4 match internal
  Last update from 19.19.19.19 on ATM3/1/0.1, 01:37:02 ago
  Routing Descriptor Blocks:
  * hh.hh.hh.hh, from hh.hh.hh.hh, 01:37:02 ago, via ATM3/1/0.1
      Route metric is 7, traffic share count is 1
```

Step 3 Verify that the loopback address of the remote PE router (PE1) is in the routing table.

```
Router# show ip route bb.bb.bb.bb

Routing entry for bb.bb.bb.bb/32
  Known via "bgp 200", distance 20, metric 0
  Tag 100, type external
  Redistributing via ospf 200
  Advertised by ospf 200 metric 3 subnets
  Last update from ss.0.0.1 22:04:44 ago
  Routing Descriptor Blocks:
  * ss.0.0.1, from ss.0.0.1, 22:04:44 ago
      Route metric is 0, traffic share count is 1
      AS Hops 2, BGP network version 0
```

Step 4 Verify that the prefix of the local PE router (PE2) is in the MPLS LDP bindings.

```
Router# show mpls ldp bindings hh.hh.hh.hh 255.255.255.255

  tib entry: hh.hh.hh.hh/32, rev 37
        local binding:  tag: 31

Router# show mpls ldp bindings hh.hh.hh.hh 255.255.255.255 detail

  tib entry: hh.hh.hh.hh/32, rev 37
        local binding:  tag: 31
```

Step 5 Verify that the prefix of the local PE (PE2) is in the CEF table.

```
Router# show ip cef hh.hh.hh.hh

hh.hh.hh.hh/32, version 31, cached adjacency to ATM3/1/0.1
0 packets, 0 bytes
  tag information set
    local tag: 31
    fast tag rewrite with AT3/1/0.1, point2point, tags imposed {2/33(vcd=2)}
  via hh.hh.hh.hh, ATM3/1/0.1, 0 dependencies
    next hop hh.hh.hh.hh, ATM3/1/0.1
    unresolved
    valid cached adjacency
    tag rewrite with AT3/1/0.1, point2point, tags imposed {2/33(vcd=2)}

Router# show ip cef hh.hh.hh.hh detail

hh.hh.hh.hh/32, version 31, cached adjacency to ATM3/1/0.1
0 packets, 0 bytes
  tag information set
    local tag: 31
    fast tag rewrite with AT3/1/0.1, point2point, tags imposed {2/33(vcd=2)}
  via hh.hh.hh.hh, ATM3/1/0.1, 0 dependencies
    next hop hh.hh.hh.hh, ATM3/1/0.1
    unresolved
    valid cached adjacency
    tag rewrite with AT3/1/0.1, point2point, tags imposed {2/33(vcd=2)}
```

Step 6 Verify that the prefix of the local PE router (PE2) is in the MPLS forwarding table.

```
Router# show mpls forwarding-table hh.hh.hh.hh

Local  Outgoing    Prefix            Bytes tag  Outgoing   Next Hop
tag    tag or VC   or Tunnel Id      switched   interface
31     2/33        hh.hh.hh.hh/32    1908083    AT3/1/0.1  point2point


Router# show mpls forwarding-table hh.hh.hh.hh detail

Local  Outgoing    Prefix            Bytes tag  Outgoing   Next Hop
tag    tag or VC   or Tunnel Id      switched   interface
31     2/33        hh.hh.hh.hh/32    1908200    AT3/1/0.1  point2point
        MAC/Encaps=4/8, MTU=4470, Tag Stack{2/33(vcd=2)}
        00028847 00002000
        No output feature configured
    Per-packet load-sharing, slots: 0 1 2 3 4 5 6 7 8 9 10 11 12 13 14 15
```

Step 7 Verify that the BGP routing table contains labels for prefixes in the customer carrier MPLS VPN service provider networks.

```
Router# show ip bgp labels

Network          Next Hop         In Label/Out Label
cc.cc.cc.cc/32   ss.0.0.1         18/27
bb.bb.bb.bb/32   ss.0.0.1         19/26
hh.hh.hh.hh/32   hh.hh.hh.hh      31/exp-null
gg.gg.gg.gg/32   0.0.0.0          imp-null/exp-null
nn.0.0.0         ss.0.0.1         22/24
ss.0.0.0         0.0.0.0          imp-null/exp-null
ss.0.0.1/32      0.0.0.0          16/exp-null
pp.0.0.0         ss.0.0.1         26/21
```

Step 8 Verify that the prefix of the remote PE router (PE1) is in the CEF table.

```
Router# show ip cef bb.bb.bb.bb

bb.bb.bb.bb/32, version 18, cached adjacency ss.0.0.1
0 packets, 0 bytes
  tag information set
    local tag: 19
    fast tag rewrite with Et0/0/3, ss.0.0.1, tags imposed {26}
  via ss.0.0.1, 0 dependencies, recursive
    next hop ss.0.0.1, Ethernet0/0/3 via ss.0.0.1/32
    valid cached adjacency
    tag rewrite with Et0/0/3, ss.0.0.1, tags imposed {26}


Router# show ip cef bb.bb.bb.bb detail

bb.bb.bb.bb/32, version 18, cached adjacency ss.0.0.1
0 packets, 0 bytes
  tag information set
    local tag: 19
    fast tag rewrite with Et0/0/3, ss.0.0.1, tags imposed {26}
  via ss.0.0.1, 0 dependencies, recursive
    next hop ss.0.0.1, Ethernet0/0/3 via ss.0.0.1/32
    valid cached adjacency
    tag rewrite with Et0/0/3, ss.0.0.1, tags imposed {26}
```

Step 9 Verify that the prefix of the remote PE router (PE1) is in the MPLS forwarding table.

```
Router# show mpls forwarding-table bb.bb.bb.bb

Local  Outgoing    Prefix            Bytes tag  Outgoing   Next Hop
tag    tag or VC   or Tunnel Id      switched   interface
19     26          bb.bb.bb.bb/32    965401     Et0/0/3    ss.0.0.1


Router# show mpls forwarding-table bb.bb.bb.bb detail

Local  Outgoing    Prefix            Bytes tag  Outgoing   Next Hop
tag    tag or VC   or Tunnel Id      switched   interface
19     26          bb.bb.bb.bb/32    965478     Et0/0/3    ss.0.0.1
        MAC/Encaps=14/18, MTU=1500, Tag Stack{26}
        00B0C26E448C0030A36338038847 0001A000
        No output feature configured
    Per-packet load-sharing, slots: 0 1 2 3 4 5 6 7 8 9 10 11 12 13 14 15
```

Verifying the PE2 Router Configuration

Step 1 Verify that the loopback address of the local CE router (CE2) is in the routing table of the PE2 router.

```
Router# show ip route vrf <PE-vrf-name> jj.jj.jj.jj

Routing entry for jj.jj.jj.jj/32
  Known via "bgp 200", distance 20, metric 0
  Tag 300, type external
  Last update from tt.0.0.2 22:11:06 ago
  Routing Descriptor Blocks:
  * tt.0.0.2, from tt.0.0.2, 22:11:06 ago
      Route metric is 0, traffic share count is 1
      AS Hops 1, BGP network version 0
```

Step 2 Verify that the prefix of the local CE router (CE2) is in the MPLS forwarding table, and is untagged.

```
Router# show mpls forwarding-table vrf <PE-vrf-name> jj.jj.jj.jj

Local  Outgoing    Prefix            Bytes tag  Outgoing   Next Hop
tag    tag or VC   or Tunnel Id      switched   interface
32     Untagged    jj.jj.jj.jj/32[V] 0          Et3/6      tt.0.0.2


Router# show mpls forwarding-table vrf <PE-vrf-name> jj.jj.jj.jj detail

Local  Outgoing    Prefix            Bytes tag  Outgoing   Next Hop
tag    tag or VC   or Tunnel Id      switched   interface
32     Untagged    jj.jj.jj.jj/32[V] 0          Et3/6      tt.0.0.2
        MAC/Encaps=0/0, MTU=1504, Tag Stack{}
        VPN route: vpn2
        No output feature configured
    Per-packet load-sharing, slots: 0 1 2 3 4 5 6 7 8 9 10 11 12 13 14 15
```

Step 3 Verify that the prefix of the remote PE router (PE1) is in the CEF table.

```
Router# show ip cef bb.bb.bb.bb

bb.bb.bb.bb/32, version 19, cached adjacency to ATM5/0.1
0 packets, 0 bytes
  tag information set
    local tag: 20
    fast tag rewrite with AT5/0.1, point2point, tags imposed {2/35(vcd=6)}
  via gg.gg.gg.gg, ATM5/0.1, 2 dependencies
    next hop gg.gg.gg.gg, ATM5/0.1
    unresolved
    valid cached adjacency
    tag rewrite with AT5/0.1, point2point, tags imposed {2/35(vcd=6)}


Router# show ip cef bb.bb.bb.bb detail

bb.bb.bb.bb/32, version 19, cached adjacency to ATM5/0.1
0 packets, 0 bytes
  tag information set
    local tag: 20
    fast tag rewrite with AT5/0.1, point2point, tags imposed {2/35(vcd=6)}
  via gg.gg.gg.gg, ATM5/0.1, 2 dependencies
    next hop gg.gg.gg.gg, ATM5/0.1
    unresolved
    valid cached adjacency
    tag rewrite with AT5/0.1, point2point, tags imposed {2/35(vcd=6)}
```

Step 4 Verify that the loopback address of the remote CE router (CE1) is in the routing table.

```
Router# show ip route vrf <PE-vrf-name> aa.aa.aa.aa

Routing entry for aa.aa.aa.aa/32
  Known via "bgp 200", distance 200, metric 0
  Tag 300, type internal
  Last update from bb.bb.bb.bb 01:43:34 ago
  Routing Descriptor Blocks:
  * bb.bb.bb.bb (Default-IP-Routing-Table), from bb.bb.bb.bb, 01:43:34 ago
      Route metric is 0, traffic share count is 1
      AS Hops 1, BGP network version 0
```

Step 5 Verify that the prefix of the remote CE router (CE1) is in the MPLS forwarding table, and that it has an outgoing interface.

```
Router# show mpls forwarding-table vrf <PE-vrf-name> aa.aa.aa.aa

Local  Outgoing    Prefix            Bytes tag  Outgoing   Next Hop
tag    tag or VC   or Tunnel Id      switched   interface
None   2/35        aa.aa.aa.aa/32    0          AT5/0.1    point2point


Router# show mpls forwarding-table vrf <PE-vrf-name> aa.aa.aa.aa detail

Local  Outgoing    Prefix            Bytes tag  Outgoing   Next Hop
tag    tag or VC   or Tunnel Id      switched   interface
None   2/35        aa.aa.aa.aa/32    0          AT5/0.1    point2point
        MAC/Encaps=4/12, MTU=4466, Tag Stack{2/35(vcd=6) 23}
        00068847 0000600000017000
        No output feature configured
```

Step 6 Verify that the prefix of the remote CE router (CE1) is in the CEF table.

```
Router# show ip cef vrf <PE-vrf-name> aa.aa.aa.aa

aa.aa.aa.aa/32, version 10, cached adjacency to ATM5/0.1
0 packets, 0 bytes
  tag information set
    local tag: VPN route head
    fast tag rewrite with AT5/0.1, point2point, tags imposed {2/35(vcd=6) 23}
  via bb.bb.bb.bb, 0 dependencies, recursive
    next hop gg.gg.gg.gg, ATM5/0.1 via bb.bb.bb.bb/32
    valid cached adjacency
    tag rewrite with AT5/0.1, point2point, tags imposed {2/35(vcd=6) 23}

Router# show ip cef vrf <PE-vrf-name> aa.aa.aa.aa detail

aa.aa.aa.aa/32, version 10, cached adjacency to ATM5/0.1
0 packets, 0 bytes
  tag information set
    local tag: VPN route head
    fast tag rewrite with AT5/0.1, point2point, tags imposed {2/35(vcd=6) 23}
  via bb.bb.bb.bb, 0 dependencies, recursive
    next hop gg.gg.gg.gg, ATM5/0.1 via bb.bb.bb.bb/32
    valid cached adjacency
    tag rewrite with AT5/0.1, point2point, tags imposed {2/35(vcd=6) 23}
```

Step 7 Verify that the prefix of the local PE router (PE2) is in the CEF table.

```
Router# show ip cef hh.hh.hh.hh

hh.hh.hh.hh/32, version 9, connected, receive
  tag information set
    local tag: implicit-null

Router# show ip cef hh.hh.hh.hh detail

hh.hh.hh.hh/32, version 9, connected, receive
  tag information set
    local tag: implicit-null
```

Verifying the CE2 Router Configuration

Step 1 Verify that the loopback address of the remote CE router (CE1), learned from the PE router, is in the routing table of the CE2 router.

```
Router# show ip route aa.aa.aa.aa

Routing entry for aa.aa.aa.aa/32
  Known via "bgp 300", distance 20, metric 0
  Tag 200, type external
  Redistributing via ospf 300
  Advertised by ospf 300 subnets
  Last update from tt.0.0.1 22:16:46 ago
  Routing Descriptor Blocks:
  * tt.0.0.1, from tt.0.0.1, 22:16:46 ago
      Route metric is 0, traffic share count is 1
      AS Hops 2
```

Configuration Examples

This section includes the following configuration examples:

- Configuration of BGP/MPLS Backbone Carrier Using BGP to Distribute MPLS Labels to an MPLS VPN Service Provider
- Configuration of Route Maps

Configuration of BGP/MPLS Backbone Carrier Using BGP to Distribute MPLS Labels to an MPLS VPN Service Provider

This section provides a example of a backbone carrier and a customer carrier who are both BGP/MPLS VPN service providers. This configuration example shows how BGP is enabled to distribute routes and MPLS labels between PE and CE routers. Figure 4 shows the example configuration.

Figure 4 Network Where the Customer Carrier Is an MPLS VPN Service Provider

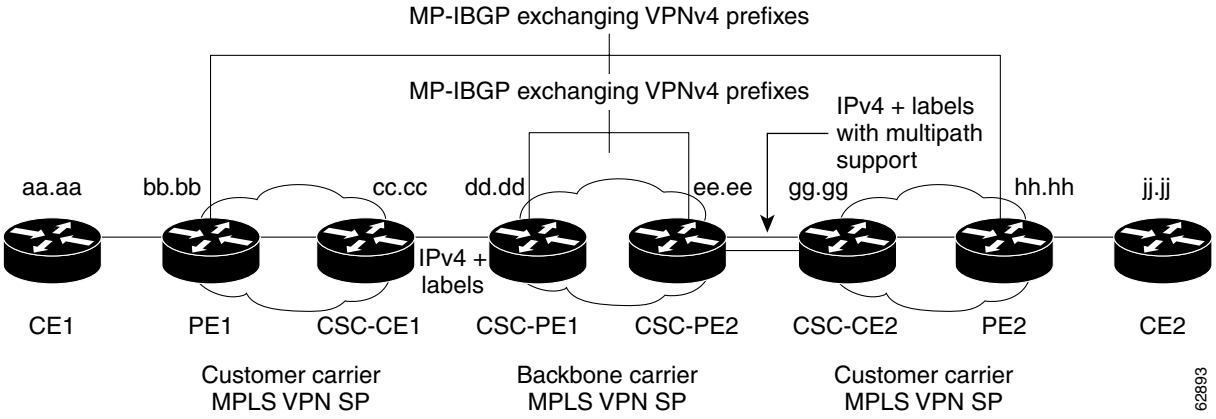

Table 2 describes the sample configuration shown in Figure 4.

Table 2 Description of Sample Configuration Shown in Figure 4

Routers	Description
CE1 and CE2	Belong to an end customer. CE1 and CE2 routers exchange routes learned from PE routers. The end customer is purchasing VPN services from a customer carrier.
PE1 and PE2	Part of a customer carrier network that is configured to provide MPLS VPN services. PE1 and PE2 are peering with a VPNv4 IBGP session to form a MPLS VPN network.
CSC-CE1 and CSC-CE2	Part of a customer carrier network. CSC-CE1 and CSC-CE2 routers exchange IPv4 BGP updates with MPLS labels and redistribute PE loopback addressees to and from the IGP (OSPF in this example). The customer carrier is purchasing carrier supporting carrier VPN services from a backbone carrier.
CSC-PE1 and CSC-PE2	Part of the backbone carrier's network configured to provide carrier supporting carrier VPN services. CSC-PE1 and CSC-PE2 are peering with a VPNv4 IP BGP session to form the MPLS VPN network. In the VRF, CSC-PE1 and CSC-PE2 are peering with the CSC-CE routers, which are configured for carrying MPLS labels with the routes, with an IPv4 EBGP session.

CE1 Configuration

```
ip cef
interface Loopback0
ip address aa.aa.aa.aa 255.255.255.255
!
interface Ethernet3/3
 ip address mm.0.0.1 255.0.0.0
!
router bgp 300
 no synchronization
 bgp log-neighbor-changes
 timers bgp 10 30
 redistribute connected                                  !Exchange routes
 neighbor mm.0.0.2 remote-as 200                         !learned from PE1.
 neighbor mm.0.0.2 advertisement-interval 5
 no auto-summary
```

PE1 Configuration

```
ip cef
!
ip vrf vpn2
 rd 200:1
 route-target export 200:1
 route-target import 200:1
mpls label protocol ldp
!
interface Loopback0
 ip address bb.bb.bb.bb 255.255.255.255
!
```

```
interface Ethernet3/0
 ip address nn.0.0.1 255.0.0.0
 no ip directed-broadcast
 no ip mroute-cache
 mpls label protocol ldp
 mpls ip
!
interface Ethernet3/3
 ip vrf forwarding vpn2
 ip address mm.0.0.2 255.0.0.0
 no ip directed-broadcast
 no ip mroute-cache
!
router ospf 200
 log-adjacency-changes
 auto-cost reference-bandwidth 1000
 redistribute connected subnets
 passive-interface Ethernet3/3
 network bb.bb.bb.bb 0.0.0.0 area 200
 network nn.0.0.0 0.255.255.255 area 200
!
router bgp 200
 no bgp default ipv4-unicast
 bgp log-neighbor-changes
 timers bgp 10 30
 neighbor hh.hh.hh.hh remote-as 200
 neighbor hh.hh.hh.hh update-source Loopback0
 !
address-family vpnv4                                    !VPNv4 session with PE2.
 neighbor hh.hh.hh.hh activate
 neighbor hh.hh.hh.hh send-community extended
 bgp dampening 30
 exit-address-family
 !
 address-family ipv4 vrf vpn2
 neighbor mm.0.0.1 remote-as 300
 neighbor mm.0.0.1 activate
 neighbor mm.0.0.1 as-override
 neighbor mm.0.0.1 advertisement-interval 5
no auto-summary
 no synchronization
 bgp dampening 30
 exit-address-family
```

CSC-CE1 Configuration

```
ip cef
!
mpls label protocol ldp
!
interface Loopback0
 ip address cc.cc.cc.cc 255.255.255.255
!
interface Ethernet3/0
 ip address pp.0.0.1 255.0.0.0
!
interface Ethernet4/0
 ip address nn.0.0.2 255.0.0.0
 no ip directed-broadcast
 no ip mroute-cache
 mpls label protocol ldp
 mpls ip
```

```
!
router ospf 200
 log-adjacency-changes
 auto-cost reference-bandwidth 1000
 redistribute connected subnets                       !Exchange routes
 redistribute bgp 200 metric 3 subnets                !learned from PE1.
 passive-interface ATM1/0
 passive-interface Ethernet3/0
 network cc.cc.cc.cc 0.0.0.0 area 200
 network nn.0.0.0 0.255.255.255 area 200
!
router bgp 200
 no bgp default ipv4-unicast
 bgp log-neighbor-changes
 timers bgp 10 30
 neighbor pp.0.0.2 remote-as 100
 neighbor pp.0.0.2 update-source Ethernet3/0
 no auto-summary
 !
address-family ipv4
 redistribute connected
 redistribute ospf 200 metric 4 match internal
 neighbor pp.0.0.2 activate
 neighbor pp.0.0.2 send-label
 no auto-summary
 no synchronization
 bgp dampening 30
 exit-address-family
```

CSC-PE1 Configuration

```
ip cef
!
ip vrf vpn1
 rd 100:1
 route-target export 100:1
 route-target import 100:1
mpls label protocol ldp
!
interface Loopback0
 ip address dd.dd.dd.dd 255.255.255.255
!
interface Ethernet3/1
 ip vrf forwarding vpn1
 ip address pp.0.0.2 255.0.0.0
!
interface ATM0/1/0
 no ip address
 no ip directed-broadcast
 no ip route-cache distributed
 atm clock INTERNAL
 no atm enable-ilmi-trap
 no atm ilmi-keepalive
!
interface ATM0/1/0.1 mpls
 ip unnumbered Loopback0
 no ip directed-broadcast
 no atm enable-ilmi-trap
 mpls label protocol ldp
 mpls atm vpi 2-5
 mpls ip
!
```

```
router ospf 100
 log-adjacency-changes
 auto-cost reference-bandwidth 1000
 redistribute connected subnets
 passive-interface Ethernet3/1
 network dd.dd.dd.dd 0.0.0.0 area 100
!
router bgp 100
 no bgp default ipv4-unicast
 bgp log-neighbor-changes
 timers bgp 10 30
 neighbor ee.ee.ee.ee remote-as 100
 neighbor ee.ee.ee.ee update-source Loopback0
 !
address-family vpnv4                                  !VPNv4 session with CSC-PE2.
 neighbor ee.ee.ee.ee activate
 neighbor ee.ee.ee.ee send-community extended
 bgp dampening 30
 exit-address-family
 !
 address-family ipv4 vrf vpn1
 neighbor pp.0.0.1 remote-as 200
 neighbor pp.0.0.1 activate
 neighbor pp.0.0.1 as-override
 neighbor pp.0.0.1 advertisement-interval 5
 neighbor pp.0.0.1 send-label
 no auto-summary
 no synchronization
 bgp dampening 30
 exit-address-family
```

CSC-PE2 Configuration

```
ip cef
!
ip vrf vpn1
 rd 100:1
 route-target export 100:1
 route-target import 100:1
mpls label protocol ldp
!
interface Loopback0
 ip address ee.ee.ee.ee 255.255.255.255
!
interface Ethernet5/0
 ip vrf forwarding vpn1
 ip address ss.0.0.2 255.0.0.0
 no ip directed-broadcast
 no ip route-cache distributed
 clock source internal
!
interface ATM2/1/0
 no ip address
 no ip directed-broadcast
 no ip route-cache distributed
 atm clock INTERNAL
 no atm enable-ilmi-trap
 no atm ilmi-keepalive
!
interface ATM2/1/0.1 mpls
 ip unnumbered Loopback0
 no ip directed-broadcast
```

```
 no atm enable-ilmi-trap
 mpls label protocol ldp
 mpls atm vpi 2-5
 mpls ip
!
router ospf 100
 log-adjacency-changes
 auto-cost reference-bandwidth 1000
 redistribute connected subnets
 passive-interface Ethernet5/0
 passive-interface ATM3/0/0
 network ee.ee.ee.ee 0.0.0.0 area 100
!
router bgp 100
 no bgp default ipv4-unicast
 bgp log-neighbor-changes
 timers bgp 10 30
 neighbor dd.dd.dd.dd remote-as 100
 neighbor dd.dd.dd.dd update-source Loopback0
 !
address-family vpnv4                                !VPNv4 session with CSC-PE1.
 neighbor dd.dd.dd.dd activate
 neighbor dd.dd.dd.dd send-community extended
 bgp dampening 30
 exit-address-family
 !
 address-family ipv4 vrf vpn1
 neighbor ss.0.0.1 remote-as 200
 neighbor ss.0.0.1 activate
 neighbor ss.0.0.1 as-override
 neighbor ss.0.0.1 advertisement-interval 5
 neighbor ss.0.0.1 send-label
 no auto-summary
 no synchronization
 bgp dampening 30
 exit-address-family
```

CSC-CE2 Configuration

```
ip cef
!
mpls label protocol ldp
!
interface Loopback0
 ip address gg.gg.gg.gg 255.255.255.255
!
interface Ethernet2/2
 ip address ss.0.0.2 255.0.0.0
 no ip directed-broadcast
 no ip mroute-cache
 mpls label protocol ldp
  mpls ip
!
interface ATM3/1/0.1 point2point
 ip address yy.0.0.1 255.0.0.0
 no ip directed-broadcast
 no atm enable-ilmi-trap
 no ip mroute-cache
!
router ospf 200
 log-adjacency-changes
 auto-cost reference-bandwidth 1000
```

```
 redistribute connected subnets                   !Exchange routes
 redistribute bgp 200 metric 3 subnets            !learned from PE2.
 passive-interface ATM3/1/0.1
 network gg.gg.gg.gg 0.0.0.0 area 200
 network ss.0.0.0 0.255.255.255 area 200
!
router bgp 200
 no bgp default ipv4-unicast
 bgp log-neighbor-changes
 timers bgp 10 30
 neighbor yy.0.0.2 remote-as 100
 neighbor yy.0.0.2 update-source ATM3/1/0.1
 no auto-summary
 !
 address-family ipv4
 redistribute connected
 redistribute ospf 200 metric 4 match internal
 neighbor yy.0.0.2 activate
 neighbor yy.0.0.2 send-label
 no auto-summary
 no synchronization
 bgp dampening 30
 exit-address-family
```

PE2 Configuration

```
ip cef
!
ip vrf vpn2
 rd 200:1
 route-target export 200:1
 route-target import 200:1
!
mpls label protocol ldp
!
interface Loopback0
 ip address hh.hh.hh.hh 255.255.255.255
!
interface Ethernet3/6
 ip vrf forwarding vpn2
 ip address tt.0.0.2 255.0.0.0
!
interface ATM5/0.1 point2point
 ip address qq.0.0.1 255.0.0.0
 no ip directed-broadcast
 no atm enable-ilmi-trap
 no ip mroute-cache
 mpls label protocol ldp
 mpls ip
!
router bgp 200
 no bgp default ipv4-unicast
 bgp log-neighbor-changes
 timers bgp 10 30
 neighbor bb.bb.bb.bb remote-as 200
 neighbor bb.bb.bb.bb update-source Loopback0
 !
address-family vpnv4                          !VPNv4 session with PE1.
 neighbor bb.bb.bb.bb activate
 neighbor bb.bb.bb.bb send-community extended
 bgp dampening 30
 exit-address-family
```

```
 !
 address-family ipv4 vrf vpn2
 neighbor tt.0.0.1 remote-as 300
 neighbor tt.0.0.1 activate
 neighbor tt.0.0.1 as-override
 neighbor tt.0.0.1 advertisement-interval 5
 no auto-summary
 no synchronization
 bgp dampening 30
 exit-address-family
```

CE2 Configuration

```
ip cef
!
interface Loopback0
 ip address jj.jj.jj.jj 255.255.255.255
!
interface Ethernet3/6
 ip address tt.0.0.1 255.0.0.0
!
router bgp 300
 no synchronization
 bgp log-neighbor-changes
 timers bgp 10 30                              !Exchange routes
 redistribute connected                        !learned from PE2.
 redistribute ospf 300 match internal external 1 external 2
 neighbor tt.0.0.2 remote-as 200
 neighbor tt.0.0.2 advertisement-interval 5
 no auto-summary
```

Configuration of Route Maps

The following example creates two route maps, which are named:

- IN for incoming routes
- OUT for outgoing routes

The route maps specify the following:

- If an IP address in an incoming BGP update message matches an IP address in access list 99, the route is added to the BGP table.
- If an IP address in an outbound BGP update message matches an IP address in access list 88, the router distributes that route.

The route maps are applied to the router with the address qq.0.0.1.

```
address-family ipv4 vrf vpn1
 neighbor qq.0.0.1 remote-as 200
 neighbor qq.0.0.1 activate
 neighbor qq.0.0.1 as-override
 neighbor qq.0.0.1 advertisement-interval 5
 neighbor qq.0.0.1 route-map IN in
 neighbor qq.0.0.1 route-map OUT out
 neighbor qq.0.0.1 send-label
!
access-list 88 permit rr.rr.rr.rr
access-list 88 permit ss.ss.ss.ss
access-list 88 permit tt.tt.tt.tt
access-list 99 permit uu.uu.uu.uu
```

```
access-list 99 permit vv.vv.vv.vv
access-list 99 permit ww.ww.ww.ww
!
route-map IN permit 1
 match ip address 99
!
route-map OUT permit 1
 match ip address 88
 set mpls-label
!
```

Command Reference

This section documents new or modified commands. All other commands used with this feature are documented in the Cisco IOS Release 12.2 command reference publications.

New Commands

- **match mpls-label**
- **neighbor send-label**
- **set mpls-label**
- **show ip bgp labels**

Modified Commands

- **debug ip bgp**
- **show ip bgp**
- **show ip bgp neighbors**
- **show ip bgp vpnv4**
- **show route-map**

debug ip bgp

To display information related to processing of the Border Gateway Protocol (BGP), use the **debug ip bgp** command in privileged EXEC mode. To disable the display of BGP information, use the **no** form of this command.

debug ip bgp [*A.B.C.D.* | **dampening** | **events** | **in** | **keepalives** | **out** | **updates** | **vpnv4** | **mpls**]

no debug ip bgp [*A.B.C.D.* | **dampening** | **events** | **in** | **keepalives** | **out** | **updates** | **vpnv4** | **mpls**]

Syntax Description

A.B.C.D.	(Optional) Displays the BGP neighbor IP address.
dampening	(Optional) Displays BGP dampening.
events	(Optional) Displays BGP events.
in	(Optional) Displays BGP inbound information.
keepalives	(Optional) Displays BGP keepalives.
out	(Optional) Displays BGP outbound information.
updates	(Optional) Displays BGP updates.
vpnv4	(Optional) Displays VPNv4 NLRI information.
mpls	(Optional) Displays the MPLS information.

Command Modes

Privileged EXEC

Command History

Release	Modification
12.0(5)T	This command was introduced.
12.0(21)ST	This command was updated with the **mpls** keyword.
12.0(22)S	This command was integrated into Cisco IOS Release 12.0(22)S. Support for the Cisco 12000 series routers (Engine 0 and Engine 2) was added.

Examples

The following example displays the output from this command:

```
Router# debug ip bgp vpnv4

03:47:14:vpn:bgp_vpnv4_bnetinit:100:2:58.0.0.0/8
03:47:14:vpn:bnettable add:100:2:58.0.0.0 / 8
03:47:14:vpn:bestpath_hook route_tag_change for vpn2:58.0.0.0/255.0.0.0(ok)
03:47:14:vpn:bgp_vpnv4_bnetinit:100:2:57.0.0.0/8
03:47:14:vpn:bnettable add:100:2:57.0.0.0 / 8
03:47:14:vpn:bestpath_hook route_tag_change for vpn2:57.0.0.0/255.0.0.0(ok)
03:47:14:vpn:bgp_vpnv4_bnetinit:100:2:14.0.0.0/8
03:47:14:vpn:bnettable add:100:2:14.0.0.0 / 8
03:47:14:vpn:bestpath_hook route_tag_chacle ip bgp *nge for vpn2:14.0.0.0/255.0.0.0(ok)
```

match mpls-label

To redistribute routes that include Multiprotocol Label Switching (MPLS) labels if the routes meet the conditions specified in the route map, use the **match mpls-label** command in route map configuration mode. To disable this function, use the **no** form of this command.

match mpls-label

no match mpls-label

Syntax Description

This command has no arguments or keywords.

Defaults

This command has no default behavior or values.

Command Modes

Route map configuration

Command History

Release	Modification
12.0(21)ST	This command was introduced.
12.0(22)S	This command was integrated into Cisco IOS Release 12.0(22)S. Support for the Cisco 12000 series routers (Engine 0 and Engine 2) was added.

Usage Guidelines

A route map that includes this command can be used in the following instances:

- With the **neighbor route-map in** command to manage inbound route maps in BGP
- With the **redistribute bgp** command to redistribute route maps in an IGP

Use the **route-map** global configuration command, and the **match** and **set** route map configuration commands, to define the conditions for redistributing routes from one routing protocol into another. Each **route-map** command has a list of **match** and **set** commands associated with it. The **match** commands specify the match criteria—the conditions under which redistribution is allowed for the current **route-map** command. The **set** commands specify the set actions—the particular redistribution actions to perform if the criteria enforced by the **match** commands are met. The **no route-map** command deletes the route map.

The **match route-map** configuration command has multiple formats. The **match** commands can be given in any order, and all **match** commands must "pass" to cause the route to be redistributed according to the set actions given with the **set** commands. The **no** forms of the **match** commands remove the specified match criteria.

When you are passing routes through a route map, a route map can have several parts. Any route that does not match at least one match clause relating to a **route-map** command will be ignored; that is, the route will not be advertised for outbound route maps and will not be accepted for inbound route maps. If you want to modify only some data, you must configure a second route map section with an explicit match specified.

Examples

The following example creates a route map that redistributes routes if the following conditions are met:

- The IP address of the route matches an IP address in ACL 2.
- The route includes an MPLS label.

```
route-map incoming permit 10
match ip address 2
match mpls-label
```

Related Commands

Command	Description
match ip address	Distributes any routes that have a destination network number address that is permitted by a standard or extended access list.
route-map (IP)	Defines the conditions for redistributing routes from one routing protocol into another, or enables policy routing.
set mpls-label	Enables a route to be distributed with an MPLS label if the route matches the conditions specified in the route map.

neighbor send-label

To enable a Border Gateway Protocol (BGP) router to send Multiprotocol Label Switching (MPLS) labels with BGP routes to a neighboring BGP router, use the **neighbor send-label** command in router configuration mode. To disable the BGP router from sending MPLS labels with BGP routes, use the **no** form of this command.

neighbor {*ip-address*} **send-label**

no neighbor {*ip-address*} **send-label**

Syntax Description

ip-address	IP address of the neighboring router.

Defaults

By default, BGP routers distribute only BGP routes.

Command Modes

Router configuration

Command History

Release	Modification
12.0(21)ST	This command was introduced.
12.0(22)S	This command was integrated into Cisco IOS Release 12.0(22)S. Support for the Cisco 12000 series routers (Engine 0 and Engine 2) was added.

Usage Guidelines

This command enables a router to use BGP to distribute MPLS labels along with the IPv4 routes to a peer router. You must issue this command on both the local router and the neighboring router.

This command has the following restrictions:

- If a BGP session is running when you issue the **neighbor send-label** command, the command does not take effect until the BGP session is restarted.
- You can use this command only with IPv4 addresses.

Cisco IOS installs /32 routes for directly connected external BGP (eBGP) peers when the BGP session for such a peer comes up. The /32 routes are installed only when MPLS labels are exchanged between such peers. Directly connected eBGP peers exchange MPLS labels for:

- IP address families (IPv4 and IPv6) with the **neighbor send-label** command enabled for the peers
- VPN address families (VPNv4 and VPNv6)

A single BGP session can include multiple address families. If one of the families exchanges MPLS labels, the /32 neighbor route is installed for the connected peer.

Examples

The following example enables a router called BGP 1 to send MPLS labels with BGP routes to the neighboring router, whose IP address is 192.168.0.0:

```
router bgp1
neighbor 192.168.0.0 send-label
```

Related Commands

Command	Description
neighbor activate	Enables the exchange of IPv4 address information with a BGP neighboring router.

set mpls-label

To enable a route to be distributed with a Multiprotocol Label Switching (MPLS) label if the route matches the conditions specified in the route map, use the **set mpls-label** command in route map configuration mode. To disable this function, use the **no** form of this command.

set mpls-label

no set mpls-label

Syntax Description

This command has no arguments or keywords.

Defaults

This command has no default behavior or values.

Command Modes

Route map configuration

Command History

Release	Modification
12.0(21)ST	This command was introduced.
12.0(22)S	This command was integrated into Cisco IOS Release 12.0(22)S. Support for the Cisco 12000 series routers (Engine 0 and Engine 2) was added.

Usage Guidelines

This command can be used only with the **neighbor route-map out** command to manage outbound route maps for a Border Gateway Protocol (BGP) session.

Use the **route-map** global configuration command with **match** and **set route-map** configuration commands to define the conditions for redistributing routes from one routing protocol into another. Each **route-map** command has a list of **match** and **set** commands associated with it. The **match** commands specify the match criteria—the conditions under which redistribution is allowed for the current **route-map** command. The **set** commands specify the set actions—the particular redistribution actions to perform if the criteria enforced by the **match** commands are met. The **no route-map** command deletes the route map.

Examples

The following example creates a route map that enables the route to be distributed with a label if the IP address of the route matches an IP address in ACL 1.

```
route-map incoming permit 10
match ip address 1
set mpls-label
```

Related Commands

Command	Description
match ip address	Distributes any routes that have a destination network number address that is permitted by a standard or extended access list.
match mpls-label	Redistributes routes that contain MPLS labels and match the conditions specified in the route map.
route-map (IP)	Defines the conditions for redistributing routes from one routing protocol into another, or enables policy routing.

show ip bgp

To display entries in the Border Gateway Protocol (BGP) routing table, use the **show ip bgp** command in privileged EXEC mode.

show ip bgp [*network*] [*network-mask*] [**longer-prefixes**]{**prefix-list** *prefix-list-name* | **route-map** *route-map-name*] [**shorter prefixes** *mask-length*]

Syntax Description

network	(Optional) Network number, entered to display a particular network in the BGP routing table.
network-mask	(Optional) Displays all BGP routes matching the address and mask pair.
longer-prefixes	(Optional) Displays the route and more specific routes.
prefix-list | **route-map**	(Optional) Displays selected routes from a BGP routing table based on the contents of a prefix list or route map.
prefix-list-name | *route-map-name*	(Optional) Name of the prefix list or route map that is specified for the associated keyword.
shorter prefixes *mask-length*	(Optional) Displays learned prefixes that are longer than the maximum length but shorter than the specified mask for the prefix.

Command Modes

Privileged EXEC

Command History

Release	Modification
10.0	This command was introduced.
12.0	The display of prefix advertisement statistics was added.
12.0(6)T	The display of a message indicating support for route refresh capability was added.
12.0(14)ST	The **prefix-list** and **route-map** keywords were added.
12.0(14)ST	The **shorter prefixes** keyword was added.
12.0(21)ST	This command was updated to show the number of MPLS labels that arrive at and depart from the prefix.
12.0(22)S	This command was integrated into Cisco IOS Release 12.0(22)S. Support for the Cisco 12000 series routers (Engine 0 and Engine 2) was added.

Examples

The following is sample output from the **show ip bgp** command in privileged EXEC mode:

```
Router# show ip bgp

BGP table version is 5, local router ID is 10.0.33.34
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal
Origin codes: i - IGP, e - EGP, ? - incomplete

   Network          Next Hop            Metric LocPrf Weight Path
*> 1.0.0.0          0.0.0.0                  0         32768 ?
*  2.0.0.0          10.0.33.35              10             0 35 ?
*>                  0.0.0.0                  0         32768 ?
```

```
*  10.0.0.0          10.0.33.35              10             0 35 ?
*>                   0.0.0.0                  0         32768 ?
*> 192.168.0.0/16    10.0.33.35              10             0 35 ?
```

Table 3 describes the significant fields shown in the display.

Table 3 show ip bgp Field Descriptions

Field	Description
BGP table version	Internal version number of the table. This number increments when the table changes.
local router ID	IP address of the router.
Status codes	Status of the table entry. The status is displayed at the beginning of each line in the table. It can be one of the following values: s—The table entry is suppressed. d—The table entry is dampened and will not be advertised to BGP neighbors. h—The table entry does not contain the best path based on historical information. *—The table entry is valid. >—The table entry is the best entry to use for that network. i—The table entry was learned via an IBGP session.
Origin codes	Origin of the entry. The origin code is placed at the end of each line in the table. It can be one of the following values: i—Entry originated from Interior Gateway Protocol (IGP) and was advertised with a **network** router configuration command. e—Entry originated from Exterior Gateway Protocol (EGP). ?—Origin of the path is not clear. Usually, this is a router that is redistributed into BGP from an IGP.
Network	IP address of a network entity.
Next Hop	IP address of the next system that is used when forwarding a packet to the destination network. An entry of 0.0.0.0 indicates that the router has some non-BGP routes to this network.
Metric	If shown, the value of the inter-autonomous system metric.
LocPrf	Local preference value as set with the **set local-preference route-map** configuration command. The default value is 100.
Weight	Weight of the route as set via autonomous system filters.
Path	Autonomous system paths to the destination network. There can be one entry in this field for each autonomous system in the path.

The following is sample output from the **show ip bgp** command in privileged EXEC mode when you specify the **longer-prefixes** keyword:

```
Router# show ip bgp 198.92.0.0 255.255.0.0 longer-prefixes
BGP table version is 1738, local router ID is 198.92.72.24
Status codes: s suppressed, * valid, > best, i - internal
Origin codes: i - IGP, e - EGP, ? - incomplete

   Network          Next Hop          Metric LocPrf Weight Path
*> 198.92.0.0       198.92.72.30        8896         32768 ?
*                   198.92.72.30                         0 109 108 ?
*> 198.92.1.0       198.92.72.30        8796         32768 ?
*                   198.92.72.30                         0 109 108 ?
*> 198.92.11.0      198.92.72.30       42482         32768 ?
*                   198.92.72.30                         0 109 108 ?
*> 198.92.14.0      198.92.72.30        8796         32768 ?
*                   198.92.72.30                         0 109 108 ?
*> 198.92.15.0      198.92.72.30        8696         32768 ?
*                   198.92.72.30                         0 109 108 ?
*> 198.92.16.0      198.92.72.30        1400         32768 ?
*                   198.92.72.30                         0 109 108 ?
*> 198.92.17.0      198.92.72.30        1400         32768 ?
*                   198.92.72.30                         0 109 108 ?
*> 198.92.18.0      198.92.72.30        8876         32768 ?
*                   198.92.72.30                         0 109 108 ?
*> 198.92.19.0      198.92.72.30        8876         32768 ?
*                   198.92.72.30                         0 109 108 ?
```

The following is sample output from the **show ip bgp** command in privileged EXEC mode, showing information for prefix ww.0.0.0:

```
Router# show ip bgp ww.0.0.0

BGP routing table entry for ww.0.0.0/8, version 628
Paths: (1 available, best #1)
Advertised to peer-groups:
 ebgp
Advertised to non peer-group peers:
 171.69.232.162
109 65000 297 701 80
 171.69.233.56 from 171.69.233.56 (172.19.185.32)
  Origin incomplete, localpref 100, valid, external, best, ref 2
MPLS labels in/out 24/22
```

Note If a prefix has not been advertised to any peer, the display shows "Not advertised to any peer."

The following is sample output from the **show ip bgp** command in privileged EXEC mode when you specify the **prefix-list** keyword:

```
Router# show ip bgp prefix-list ROUTE

BGP table version is 39, local router ID is 10.0.0.1
Status codes:s suppressed, d damped, h history, * valid, > best, i -
internal
Origin codes:i - IGP, e - EGP, ? - incomplete

   Network          Next Hop            Metric LocPrf Weight Path
*> 192.168.1.0      10.0.0.2                             0 ?
*                   10.0.0.2                 0           0 200 ?
```

The following is sample output from the **show ip bgp** command in privileged EXEC mode when you specify the **route-map** keyword:

```
Router# show ip bgp route-map LEARNED_PATH

BGP table version is 40, local router ID is 10.0.0.1
Status codes:s suppressed, d damped, h history, * valid, > best, i -
internal
Origin codes:i - IGP, e - EGP, ? - incomplete

   Network          Next Hop            Metric LocPrf Weight Path
*> 192.168.1.0      10.0.0.2                             0 ?
*                   10.0.0.2                 0           0 200 ?
```

The following is sample output from the **show ip bgp** command in privileged EXEC mode when you specify the **shorter-prefixes** keyword:

```
Router# show ip bgp 172.16.0.0/16 shorter-prefixes 8

*> 172.16.0.0         10.0.0.2                               0 ?
*                     10.0.0.2                 0             0 200 ?
```

Related Commands

Command	Description
clear ip bgp	Resets a BGP connection or session.
neighbor soft-reconfiguration	Configures the Cisco IOS software to start storing updates.

show ip bgp labels

To display information about Multiprotocol Label Switching (MPLS) labels from the External Border Gateway Protocol (EBGP) route table, use the **show ip bgp labels** command in privileged EXEC mode.

show ip bgp labels

Syntax Description

This command has no arguments or keywords.

Defaults

This command has no default behavior or values.

Command Modes

Privileged EXEC

Command History

Release	Modification
12.0(21)ST	This command was introduced.
12.0(22)S	This command was integrated into Cisco IOS Release 12.0(22)S. Support for the Cisco 12000 series routers (Engine 0 and Engine 2) was added.

Usage Guidelines

Use this command to display EBGP labels associated with a carrier supporting carrier customer edge (CSC-CE) router.

This command displays labels for BGP routes in the default table only. To display labels in the VRF tables, use the **show ip bgp vpnv4** {**all** | **vrf** *vrf-name*} command with the optional **labels** keyword.

Examples

The following example shows output for a CSC-CE router using BGP as a label distribution protocol:

```
Router# show ip bgp labels

Network          Next Hop         In Label/Out Label
3.3.0.0/16       0.0.0.0          imp-null/exp-null
15.15.15.15/32   15.15.15.15      18/exp-null
16.16.16.16/32   0.0.0.0          imp-null/exp-null
17.17.17.17/32   34.0.0.1         20/exp-null
18.18.18.18/32   43.0.0.1         24/31
18.18.18.18/32   38.0.0.1         24/33
19.19.19.19/32   43.0.0.1         25/32
19.19.19.19/32   38.0.0.1         25/34
20.20.20.20/32   43.0.0.1         21/30
20.20.20.20/32   38.0.0.1         21/32
33.0.0.0         15.15.15.15      19/exp-null
34.0.0.0         0.0.0.0          imp-null/exp-null
35.0.0.0         43.0.0.1         22/29
35.0.0.0         38.0.0.1         22/31
38.0.0.0         0.0.0.0          imp-null/exp-null
38.0.0.1/32      38.0.0.1         17/29
38.0.0.1/32      0.0.0.0          17/exp-null
40.0.0.0         38.0.0.1         26/35
40.0.0.0         43.0.0.1         26/34
```

```
42.0.0.0          43.0.0.1          23/28
42.0.0.0          38.0.0.1          23/30
43.0.0.0          0.0.0.0           imp-null/exp-null
43.0.0.1/32       0.0.0.0           16/exp-null
```

Table 4 describes the significant fields shown in the display.

Table 4 show ip bgp labels Field Descriptions

Field	Description
Network	Displays the network address from the EGBP table.
Next Hop	Specifies the EBGP next hop address.
In Label	Displays the label (if any) assigned by this router.
Out Label	Displays the label assigned by the BGP next hop router.

Related Commands

Command	Description
show ip bgp vpnv4 all labels	Displays all incoming and outgoing labels for each NLRI in the VPNv4 database.
show ip bgp vpnv4 vrf *vrf-name* labels	Displays all incoming and outgoing labels for each NLRI in the VPNv4 database associated with the named VRF.

show ip bgp neighbors

To display information about the TCP/IP and Border Gateway Protocol (BGP) connections to neighbors, use the **show ip bgp neighbors** command in EXEC mode.

show ip bgp neighbors [*neighbor-address*] [**received-routes** | **routes** | **advertised-routes** | {**paths** *regexp*} | **dampened-routes**] [**received prefix-filter**]

Syntax Description

neighbor-address	(Optional) Address of the neighbor whose routes you have learned from. If you omit this argument, all neighbors are displayed.
received-routes	(Optional) Displays all received routes (both accepted and rejected) from the specified neighbor.
routes	(Optional) Displays all routes that are received and accepted. This is a subset of the output from the **received-routes** keyword.
advertised-routes	(Optional) Displays all the routes the router has advertised to the neighbor.
paths *regexp*	(Optional) Regular expression that is used to match the paths received.
dampened-routes	(Optional) Displays the dampened routes to the neighbor at the IP address specified.
received prefix-filter	(Optional) Displays the configured prefix list for the specified IP address.

Command Modes

EXEC

Command History

Release	Modification
10.0	This command was introduced.
11.2	The **received-routes** keyword was added.
12.0(21)ST	This command was updated to display MPLS label information.
12.0(22)S	This command was integrated into Cisco IOS Release 12.0(22)S. Support for the Cisco 12000 series routers (Engine 0 and Engine 2) was added. The **received prefix-filter** keyword was added.

Examples

The following example shows output from the **show ip bgp neighbors** command in privileged EXEC mode when Multiprotocol Label Switching (MPLS) labels are being sent and received:

```
Router# show ip bgp neighbors 172.16.232.178

BGP neighbor is 172.16.232.178,  remote AS 35, external link
  BGP version 4, remote router ID 192.168.3.3
  BGP state = Established, up for 1w1d
  Last read 00:00:53, hold time is 180, keepalive interval is 60 seconds
  Neighbor capabilities:
    MPLS Label capability:  advertised and received
    Address family IPv4 Unicast: advertised and received
    Address family IPv4 Multicast: advertised and received
```

```
  Received 12519 messages, 0 notifications, 0 in queue
  Sent 12523 messages, 0 notifications, 0 in queue
  Route refresh request: received 0, sent 0
  Minimum time between advertisement runs is 30 seconds

For address family: IPv4 Unicast
  BGP table version 5, neighbor version 5
  Index 1, Offset 0, Mask 0x2
  Community attribute sent to this neighbor
  Inbound path policy configured
  Outbound path policy configured
  Route map for incoming advertisements is uni-in
  Route map for outgoing advertisements is uni-out
  Sending Prefix & Label
  3 accepted prefixes consume 108 bytes
  Prefix advertised 6, suppressed 0, withdrawn 0

 For address family: IPv4 Multicast
  BGP table version 5, neighbor version 5
  Index 1, Offset 0, Mask 0x2
  Inbound path policy configured
  Outbound path policy configured
  Route map for incoming advertisements is mul-in
  Route map for outgoing advertisements is mul-out
  3 accepted prefixes consume 108 bytes
  Prefix advertised 6, suppressed 0, withdrawn 0

  Connections established 2; dropped 1
  Last reset 1w1d, due to Peer closed the session
Connection state is ESTAB, I/O status: 1, unread input bytes: 0
Local host: 172.16.232.178, Local port: 179
Foreign host: 172.16.232.179, Foreign port: 11002

Enqueued packets for retransmit: 0, input: 0  mis-ordered: 0 (0 bytes)

Event Timers (current time is 0x2CF49CF8):
Timer          Starts    Wakeups            Next
Retrans         12518          0             0x0
TimeWait            0          0             0x0
AckHold         12514      12281             0x0
SendWnd             0          0             0x0
KeepAlive           0          0             0x0
GiveUp              0          0             0x0
PmtuAger            0          0             0x0
DeadWait            0          0             0x0

iss:  273358651  snduna:  273596614  sndnxt:  273596614     sndwnd:  15434
irs:  190480283  rcvnxt:  190718186  rcvwnd:      15491  delrcvwnd:    893

SRTT: 300 ms, RTTO: 607 ms, RTV: 3 ms, KRTT: 0 ms
minRTT: 0 ms, maxRTT: 300 ms, ACK hold: 200 ms
Flags: passive open, nagle, gen tcbs

Datagrams (max data segment is 1460 bytes):
Rcvd: 24889 (out of order: 0), with data: 12515, total data bytes: 237921
Sent: 24963 (retransmit: 0), with data: 12518, total data bytes: 237981
```

Table 5 describes the significant fields shown in the display.

Table 5 show ip bgp neighbors Field Descriptions

Field	Description
BGP neighbor	IP address of the BGP neighbor and its autonomous system number. If the neighbor is in the same autonomous system as the router, then the link between them is internal; otherwise, it is considered external.
remote AS	Autonomous system of the neighbor.
external link	Indicates that this peer is an EBGP peer.
BGP version	BGP version being used to communicate with the remote router; the router ID (an IP address) of the neighbor is also specified.
remote router ID	IP address of the neighbor.
BGP state	Internal state of this BGP connection.
up for	Amount of time, in seconds, that the underlying TCP connection has been in existence.
Last read	Time that BGP last read a message from this neighbor.
hold time	Maximum amount of time that can elapse between messages from the peer.
keepalive interval	Time period, in seconds, between sending keepalive packets, which help ensure that the TCP connection is up.
Neighbor capabilities	BGP capabilities advertised and received from this neighbor.
MPLS Label capability	Indicates that MPLS labels are both sent and received by the EBGP peer.
Address family IPv4 Unicast:	IP Version 4 unicast-specific properties of this neighbor.
Address family IPv4 Multicast:	IP Version 4 multicast-specific properties of this neighbor.
Received	Number of total BGP messages received from this peer, including keepalives.
notifications	Number of error messages received from the peer.
Sent	Total number of BGP messages that have been sent to this peer, including keepalives.
notifications	Number of error messages the router has sent to this peer.
Route refresh request:	Number of route refresh requests sent and received from this neighbor.
advertisement runs	Value of minimum advertisement interval.
For address family:	Address family to which the following fields refer.
BGP table version	Indicates that the neighbor has been updated with this version of the primary BGP routing table.
neighbor version	Number used by the software to track the prefixes that have been sent and those that must be sent to this neighbor.
Community attribute	Appears if the **neighbor send-community** command is configured for this neighbor.
Inbound path policy	Indicates if an inbound policy is configured.
Outbound path policy	Indicates if an outbound policy is configured.

Table 5 show ip bgp neighbors Field Descriptions (continued)

Field	Description
uni-in	Name of inbound route map for the unicast address family.
uni-out	Name of outbound route map for the unicast address family.
mul-in	Name of inbound route map for the multicast address family.
mul-out	Name of outbound route map for the multicast address family.
Sending Prefix & Label	Indicates that the EBGP peer sends MPLS labels with its routes.
accepted prefixes	Number of prefixes accepted.
Prefix advertised	Number of prefixes advertised.
suppressed	Number of prefixes suppressed.
withdrawn	Number of prefixes withdrawn.
Connections established	Number of times the router has established a TCP connection and the two peers have agreed to speak BGP with each other.
dropped	Number of times that a good connection has failed or been taken down.
Last reset	Elapsed time since this peering session was last reset.
Connection state	State of BGP peer.
unread input bytes	Number of bytes of packets still to be processed.
Local host, Local port	Peering address of local router, plus port.
Foreign host, Foreign port	Peering address of the neighbor.
Event Timers	Table displays the number of starts and wakeups for each timer.
iss	Initial send sequence number.
snduna	Last send sequence number the local host sent but has not received an acknowledgment for.
sndnxt	Sequence number the local host will send next.
sndwnd	TCP window size of the remote host.
irs	Initial receive sequence number.
rcvnxt	Last receive sequence number the local host has acknowledged.
rcvwnd	TCP window size of the local host.
delrcvwnd	Delayed receive window—data the local host has read from the connection, but has not yet subtracted from the receive window the host has advertised to the remote host. The value in this field gradually increases until it is larger than a full-sized packet, at which point it is applied to the rcvwnd field.
SRTT	A calculated smoothed round-trip timeout.
RTTO	Round-trip timeout.
RTV	Variance of the round-trip time.
KRTT	New round-trip timeout (using the Karn algorithm). This field separately tracks the round-trip time of packets that have been re-sent.
minRTT	Smallest recorded round-trip timeout (hard wire value used for calculation).

Table 5 show ip bgp neighbors Field Descriptions (continued)

Field	Description
maxRTT	Largest recorded round-trip timeout.
ACK hold	Time the local host will delay an acknowledgment in order to piggyback data on it.
Flags	IP precedence of the BGP packets.
Datagrams: Rcvd	Number of update packets received from a neighbor.
with data	Number of update packets received with data.
total data bytes	Total bytes of data.
Sent	Number of update packets sent.
with data	Number of update packets with data sent.
total data bytes	Total number of data bytes.

The following is sample output from the **show ip bgp neighbors** command with the **advertised-routes** keyword in privileged EXEC mode:

```
Router# show ip bgp neighbors 172.16.232.178 advertised-routes

BGP table version is 27, local router ID is 172.16.232.181
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal
Origin codes: i - IGP, e - EGP, ? - incomplete

   Network          Next Hop          Metric LocPrf Weight Path
*>i110.0.0.0        172.16.232.179         0    100      0 ?
*> 200.2.2.0        0.0.0.0                0         32768 i
```

The following is sample output from the **show ip bgp neighbors** command with the **routes** keyword in privileged EXEC mode:

```
Router# show ip bgp neighbors 172.16.232.178 routes

BGP table version is 27, local router ID is 172.16.232.181
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal
Origin codes: i - IGP, e - EGP, ? - incomplete

   Network          Next Hop          Metric LocPrf Weight Path
*> 10.0.0.0         172.16.232.178        40             0 10 ?
*> gg.0.0.0         172.16.232.178        40             0 10 ?
```

Table 6 describes the significant fields shown in the displays.

Table 6 show ip bgp neighbors advertised-routes and routes Field Descriptions

Field	Description
BGP table version	Internal version number of the table. This number increments when the table changes.
local router ID	IP address of the router.
Status codes	Status of the table entry. The status is displayed at the beginning of each line in the table. It can be one of the following values: s—The table entry is suppressed. d—The table entry is dampened and will not be advertised to BGP neighbors. h—The table entry does not contain the best path based on historical information. *—The table entry is valid. >—The table entry is the best entry to use for that network. i—The table entry was learned via an internal BGP (iBGP) session.
Origin codes	Origin of the entry. The origin code is placed at the end of each line in the table. It can be one of the following values: i—Entry originated from Interior Gateway Protocol (IGP) and was advertised with a **network** router configuration command. e—Entry originated from Exterior Gateway Protocol (EGP). ?—Origin of the path is not clear. Usually, this is a router that is redistributed into BGP from an IGP.
Network	IP address of a network entity.
Next Hop	IP address of the next system that is used when forwarding a packet to the destination network. An entry of 0.0.0.0 indicates that the router has some non-BGP routes to this network.
Metric	If shown, this is the value of the inter-autonomous system metric. This field is frequently not used.
LocPrf	Local preference value as set with the **set local-preference route-map** configuration command. The default value is 100.
Weight	Weight of the route as set via autonomous system filters.
Path	Autonomous system paths to the destination network. There can be one entry in this field for each autonomous system in the path.

The following is sample output from the **show ip bgp neighbors** command with the **paths** keyword in privileged EXEC mode:

```
Router# show ip bgp neighbors 171.69.232.178 paths ^10

Address    Refcount Metric Path
0x60E577B0        2     40 10 ?
```

Table 7 describes the significant fields shown in the display.

Table 7 show ip bgp neighbors paths Field Descriptions

Field	Description
Address	Internal address where the path is stored.
Refcount	Number of routes using that path.
Metric	Multi Exit Discriminator (MED) metric for the path. (The name of this metric for BGP versions 2 and 3 is INTER_AS.)
Path	Autonomous system path for that route, followed by the origin code for that route.

The following is sample output from the **show ip bgp neighbors** command with the **received prefix-filter** keyword in privileged EXEC mode:

```
Router# show ip bgp neighbor 192.168.20.72 received prefix-filter

Address family:IPv4 Unicast
ip prefix-list 192.168.20.72:1 entries
   seq 5 deny 10.0.0.0/8 le 32
```

Table 8 describes the significant fields shown in the display.

Table 8 show ip bgp neighbors received prefix-filter Field Descriptions

Field	Description
Address family:	Configured address family mode.
ip prefix-list	Configured prefix list for the specified neighbor.

show ip bgp vpnv4

To display Virtual Private Network (VPN) address information from the Border Gateway Protocol (BGP) table, use the **show ip bgp vpnv4** command in EXEC mode.

show ip bgp vpnv4 {**all** | **rd** *route-distinguisher* | **vrf** *vrf-name*} [*ip-prefix/length* [**longer-prefixes**] [*output-modifiers*]] [*network-address* [*mask*] [**longer-prefixes**] [*output-modifiers*]] [**cidr-only**] [**community**] [**community-list**] [**dampened-paths**] [**filter-list**] [**flap-statistics**] [**inconsistent-as**] [**neighbors**] [**paths** [*line*]] [**peer-group**] [**quote-regexp**] [**regexp**] [**summary**] [**labels**]

Syntax Description

all	Displays the complete VPNv4 database.
rd *route-distinguisher*	Displays NLRIs that have a matching route distinguisher.
vrf *vrf-name*	Displays NLRIs associated with the named VRF.
ip-prefix/length	(Optional) The IP prefix address (in dotted decimal format) and the length of the mask (0 to 32).
longer-prefixes	(Optional) Displays the entry, if any, that exactly matches the specified prefix parameter and all entries that match the prefix in a "longest-match" sense. That is, prefixes for which the specified prefix is an initial substring.
output-modifiers	(Optional) For a list of associated keywords and arguments, use context-sensitive help.
network-address	(Optional) The IP address of a network in the BGP routing table.
mask	(Optional) The mask of the network address, in dotted decimal format.
cidr-only	(Optional) Displays only routes that have nonnatural net masks.
community	(Optional) Displays routes matching this community.
community-list	(Optional) Displays routes matching this community list.
dampened-paths	(Optional) Displays paths suppressed on account of dampening (BGP route from peer is up and down).
filter-list	(Optional) Displays routes conforming to the filter list.
flap-statistics	(Optional) Displays flap statistics of routes.
inconsistent-as	(Optional) Displays only routes that have inconsistent autonomous systems of origin.
neighbors	(Optional) Displays details about TCP and BGP neighbor connections.
paths	(Optional) Displays path information.
line	(Optional) A regular expression to match the BGP autonomous system paths.
peer-group	(Optional) Displays information about peer groups.
quote-regexp	(Optional) Displays routes matching the autonomous system path "regular expression."
regexp	(Optional) Displays routes matching the autonomous system path regular expression.

summary	(Optional) Displays BGP neighbor status.
labels	(Optional) Displays incoming and outgoing BGP labels for each NLRI.

Defaults

This command has no default behavior or values.

Command Modes

EXEC

Command History

Release	Modification
12.0(5)T	This command was introduced.
12.2(2)T	The output of the **show ip bgp vpnv4 all** *ip-prefix* command was enhanced to display attributes including multipaths and a best path to the specified network.
12.0(21)ST	The keyword **tags** was replaced with the keyword **labels** to conform to the MPLS IETF guidelines.
12.0(22)S	This command was integrated into Cisco IOS Release 12.0(22)S. Support for the Cisco 12000 series routers (Engine 0 and Engine 2) was added.

Usage Guidelines

Use this command to display VPNv4 information from the BGP database. The **show ip bgp vpnv4 all** command displays all available VPNv4 information. The **show ip bgp vpnv4 summary** command displays BGP neighbor status.

Examples

The following example shows output for all available VPNv4 information in a BGP routing table:

```
Router# show ip bgp vpnv4 all

BGP table version is 18, local router ID is 14.14.14.14
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal
Origin codes: i - IGP, e - EGP,? - incomplete

   Network          Next Hop            Metric LocPrf Weight Path
Route Distinguisher: 1:101 (default for vrf vpn1)
*>i6.6.6.6/32       223.0.0.21              11    100      0 ?
*> 7.7.7.7/32       150.150.0.2             11         32768 ?
*>i69.69.0.0/30     223.0.0.21               0    100      0 ?
*> 150.150.0.0/24   0.0.0.0                  0         32768 ?
*> 222.0.0.1/32     150.150.0.2             11         32768 ?
*>i222.0.0.3/32     223.0.0.21              11    100      0 ?
*> 222.0.0.10/32    0.0.0.0                  0         32768 ?
*>i222.0.0.30/32    223.0.0.21               0    100      0 ?
```

Table 9 describes the significant fields shown in the display.

Table 9 show ip bgp vpnv4 Field Descriptions

Field	Description
Network	Displays the network address from the BGP table.
Next Hop	Displays the address of the BGP next hop.
Metric	Displays the BGP metric.
LocPrf	Displays the local preference.
Weight	Displays the BGP weight.
Path	Displays the BGP path per route.

The following example shows how to display a table of labels for NLRIs that have a route distinguisher value of 100:1.

```
Router# show ip bgp vpnv4 rd 100:1 labels

Network            Next Hop       In label/Out label
Route Distinguisher: 100:1 (vrf1)
   2.0.0.0          10.20.0.60      34/nolabel
   10.0.0.0         10.20.0.60      35/nolabel
   12.0.0.0         10.20.0.60      26/nolabel
                    10.20.0.60      26/nolabel
   13.0.0.0         10.15.0.15      nolabel/26
```

Table 10 describes the significant fields shown in the display.

Table 10 show ip bgp vpnv4 rd labels Field Descriptions

Field	Description
Network	Displays the network address from the BGP table.
Next Hop	Specifies the BGP next hop address.
In label	Displays the label (if any) assigned by this router.
Out label	Displays the label assigned by the BGP next hop router.

The following example shows VPNv4 routing entries for the VRF named vpn1:

```
Router# show ip bgp vpnv4 vrf vpn1

BGP table version is 18, local router ID is 14.14.14.14
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal
Origin codes: i - IGP, e - EGP,? - incomplete

Network          Next Hop            Metric LocPrf Weight Path
Route Distinguisher: 1:101 (default for vrf vpn1)
*>i6.6.6.6/32       223.0.0.21              11    100      0 ?
*> 7.7.7.7/32       150.150.0.2             11         32768 ?
*>i69.69.0.0/30     223.0.0.21               0    100      0 ?
*> 150.150.0.0/24   0.0.0.0                  0         32768 ?
*> 222.0.0.1/32     150.150.0.2             11         32768 ?
*>i222.0.0.3/32     223.0.0.21              11    100      0 ?
```

Table 11 describes the significant fields shown in the display.

Table 11 show ip bgp vpnv4 vrf Field Descriptions

Field	Description
Network	Displays the network address from the BGP table.
Next Hop	Displays the address of the BGP next hop.
Metric	Displays the BGP metric.
LocPrf	Displays the local preference.
Weight	Displays the BGP weight.
Path	Displays the BGP path per route.

The following example shows attributes for network 10.22.22.0 that includes multipaths and a best path:

```
Router# show ip bgp vpnv4 all 10.22.22.0

BGP routing table entry for 100:1:10.22.22.0/24, version 50
Paths:(6 available, best #1)
Multipath:iBGP
  Advertised to non peer-group peers:
  200.1.12.12
  22
    1.22.7.8 (metric 11) from 1.11.3.4 (100.0.0.8)
      Origin IGP, metric 0, localpref 100, valid, internal, multipath, best
      Extended Community:RT:100:1
      Originator:100.0.0.8, Cluster list:100.1.1.44
  22
    1.22.1.9 (metric 11) from 1.11.1.2 (100.0.0.9)
      Origin IGP, metric 0, localpref 100, valid, internal, multipath
      Extended Community:RT:100:1
      Originator:100.0.0.9, Cluster list:100.1.1.22
  22
    1.22.6.10 (metric 11) from 1.11.6.7 (100.0.0.10)
      Origin IGP, metric 0, localpref 100, valid, internal, multipath
      Extended Community:RT:100:1
      Originator:100.0.0.10, Cluster list:100.0.0.7
  22
    1.22.4.10 (metric 11) from 1.11.4.5 (100.0.0.10)
      Origin IGP, metric 0, localpref 100, valid, internal, multipath
      Extended Community:RT:100:1
      Originator:100.0.0.10, Cluster list:100.0.0.5
  22
    1.22.5.10 (metric 11) from 1.11.5.6 (100.0.0.10)
      Origin IGP, metric 0, localpref 100, valid, internal, multipath
      Extended Community:RT:100:1
      Originator:100.0.0.10, Cluster list:100.0.0.6
```

Table 12 describes the significant fields shown in the display.

Table 12 show ip bgp vpn4 all 10.22.22.0 Field Descriptions

Field	Description
BGP routing table ... version	Internal version number of the table. This number is incremented whenever the table changes.
Paths:	Number of autonomous system paths to the specified network. If multiple paths exist, one of the multipaths is designated the best path.
Multipath:	Indicates the maximum-paths configured (iBGP or eBGP).
Advertised to non peer-group peers: 200.1.12.12 22	IP address of the BGP peers that the specified route is advertised to.
1.22.7.8 (metric 11) from 1.11.3.4 (100.0.0.8)	Indicates the next hop address and the address of the gateway that sent the update.
Origin	Indicates the origin of the entry. It can be one of the following values: IGP—Entry originated from Interior Gateway Protocol (IGP) and was advertised with a **network** router configuration command. incomplete — Entry originated from other than an IGP or Exterior Gateway Protocol (EGP) and was advertised with the **redistribute** router configuration command. EGP — Entry originated from an EGP.
metric	If shown, the value of the interautonomous system metric.
localpref	Local preference value as set with the **set local-preference route-map** configuration command. The default value is 100.
valid	Indicates that the route is usable and has a valid set of attributes.
internal/external	The field is *internal* if the path is learned via iBGP. The field is *external* if the path is learned via eBGP.
multipath	One of multiple paths to the specified network.
best	If multiple paths exist, one of the multipaths is designated the best path and advertised the neighbors.
Extended Community:RT:100:1	Route Target value associated with the specified route.
Originator:	The router ID of the route originating router when route reflector is used.
Cluster list:	The router ID of all the route reflectors that the specified route has passed through.

Related Commands

Command	Description
show ip vrf	Displays the set of defined VRFs and associated interfaces.

show route-map

To display all route maps configured or only the one specified, use the **show route-map** command in EXEC mode.

show route-map [*map-name*]

Syntax Description

map-name	(Optional) Name of a specific route map.

Command Modes

EXEC

Command History

Release	Modification
10.0	This command was introduced.
12.0(21)ST	This command was updated to display information about MPLS labels.
12.0(22)S	This command was integrated into Cisco IOS Release 12.0(22)S. Support for the Cisco 12000 series routers (Engine 0 and Engine 2) was added. An additional policy routing counter was added.

Examples

The following is sample output from the **show route-map** command:

```
Router# show route-map
route-map sid, permit, sequence 10
Match clauses:
    tag 1 2
Set clauses:
    metric 5
route-map sid, permit, sequence 20
Match clauses:
    tag 3 4
Set clauses:
    metric 6
Policy routing matches: 0packets; 0 bytes
```

The following example shows MPLS-related route map information:

```
Router# show route-map
route-map OUT, permit, sequence 10
Match clauses:
  ip address (access-lists): 1
Set clauses:
  mpls label
Policy routing matches: 0 packets, 0 bytes

route-map IN, permit, sequence 10
Match clauses:
  ip address (access-lists): 2
  mpls label
Set clauses:
Policy routing matches: 0 packets, 0 bytes
```

Table 13 describes the fields shown in the display.

Table 13 show route-map Field Descriptions

Field	Description
route-map	Name of the route map.
permit	Indicates that the route is redistributed as controlled by the set actions.
sequence	Number that indicates the position a route map takes in the list of route maps already configured with the same name.
Match clauses:	Match criteria—conditions under which redistribution is allowed for the current route map.
Set clauses:	Set actions—the particular redistribution actions to perform if the criteria enforced by the **match** commands are met.
Policy routing matches:	Displays the number of packets and bytes that have been filtered by policy routing.

Related Commands

Command	Description
redistribute (IOS CLNS)	Redistributes routes from one routing domain into another routing domain.
route-map (IOS CLNS)	Defines the conditions for redistributing routes from one routing protocol into another, or enables policy routing.

Glossary

autonomous system (AS)—A collection of networks that share the same routing protocol and that are under the same system administration.

Border Gateway Protocol (BGP)—The exterior border gateway protocol used to exchange routing information between routers in separate autonomous systems. BGP uses Transmission Control Protocol (TCP). Because TCP is a reliable protocol, BGP does not experience problems with dropped or fragmented data packets.

BGP prefixes—A route announcement using the BGP. A prefix is composed of a path of AS numbers, indicating which networks the packet must pass through, and the IP block that is being routed. A BGP prefix would look something like: 701 1239 42 206.24.14.0/24. (The /24 part is referred to as a CIDR mask. The /24 indicates that there are 24 ones in the netmask for this block starting from the left hand side. A /24 corresponds to the natural mask 255.255.255.0.

customer edge (CE) router—The customer router that connects to the provider edge (PE) router.

External Border Gateway Protocol (EBGP)—A BGP session between routers in different autonomous systems. When a pair of routers in different ASs are more than one IP hop away from each other, an external BGP session between those two routers is called multihop external BGP.

Internal Border Gateway Protocol (IBGP)—A BGP session between routers within the same autonomous system.

Label Distribution Protocol (LDP)—A standard protocol used by MPLS-enabled routers to assign the labels (addresses) used to forward packets.

label edge router (LER)—The edge router that performs label imposition and disposition.

label-switched path (LSP)—A sequence of hops in which a packet travels from one router to another router by means of label switching mechanisms. A label-switched path can be established dynamically, based on normal routing mechanisms, or through configuration.

label switching router (LSR)—An LSR forwards packets in an MPLS network by looking only at the fixed-length label.

Multiprotocol Label Switching (MPLS)—MPLS is a method for forwarding packets (frames) through a network. It enables routers at the edge of a network to apply labels to packets (frames). ATM switches or existing routers in the network core can switch packets according to the labels.

Multihop BGP—A Border Gateway Protocol between two routers in different autonomous systems that are more than one hop away from each other.

Network Layer Reachability Information (NLRI)—BGP sends routing update messages containing NLRI, which describes the route. In this context, an NLRI is a prefix. A BGP update message carries one or more NLRI prefixes and the attributes of a route for the NLRI prefixes. The route attributes include a BGP next hop gateway address, community values, and other information.

point of presence (**POP**)—An access point to the Internet. A POP has a unique IP address. The ISP or online service provider (such as AOL) has one or more POPs on the Internet. ISP users dial into the POP to connect to the Internet. A POP can reside in rented space owned by the telecommunications carrier (such as Sprint) to which the ISP is connected. A POP usually includes routers, digital/analog call aggregators, servers, and frequently frame relay or ATM switches.

provider edge (PE) router—The label edge router (LER) in the service provider network that connects to the customer edge (CE) router.

route reflector (RR)—A router that advertises, or reflects, IBGP learned routes to other IBGP peers without requiring a full network mesh.

Virtual Private Network (VPN)—A group of sites that, as a result of a set of administrative policies, can communicate with each other over a shared backbone.

VPNv4 addresses—When multiple VPNs use the same address space, the VPN addresses are made unique by adding a route distinguisher to the front of the address.

VPN routing/forwarding table (VRF table)—A VRF table includes the routing information that defines a customer VPN site that is attached to a PE router. A VRF table consists of the following elements:

- An IP routing table
- A derived forwarding table
- A set of interfaces that use the forwarding table
- A set of rules and routing protocols that determine what goes into the forwarding table