



MPLS VPN—Explicit Null Label Support with BGP IPv4 Label Session

First Published: January 20, 2004

Last Updated: June 29, 2007

The MPLS VPN—Explicit Null Label Support with BGP IPv4 Label Session feature provides a method to advertise explicit null in a Border Gateway Protocol (BGP) label session for a carrier supporting carrier (CSC) customer edge (CE) router.

Finding Feature Information in This Module

Your Cisco IOS software release may not support all of the features documented in this module. To reach links to specific feature documentation in this module and to see a list of the releases in which each feature is supported, use the “[Feature Information for MPLS VPN—Explicit Null Label Support with BGP IPv4 Label Session](#)” section on page 40.

Finding Support Information for Platforms and Cisco IOS and Catalyst OS Software Images

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

Contents

- [Prerequisites for MPLS—VPN Explicit Null Label Support with BGP IPv4 Label Session](#), page 2
- [Restrictions for MPLS—VPN Explicit Null Label Support with BGP IPv4 Label Session](#), page 2
- [Information About MPLS—VPN Explicit Null Label Support with BGP IPv4 Label Session](#), page 2
- [How to Configure MPLS—VPN Explicit Null Label Support with BGP IPv4 Label Session](#), page 3
- [Configuration Examples for MPLS—VPN Explicit Null Label Support with BGP IPv4 Label Session](#), page 5
- [Additional References](#), page 7
- [Command Reference](#), page 8



Americas Headquarters:

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

© 2004, 2006–2007 Cisco Systems, Inc. All rights reserved.

- [Feature Information for MPLS VPN—Explicit Null Label Support with BGP IPv4 Label Session, page 40](#)
- [Glossary, page 41](#)

Prerequisites for MPLS—VPN Explicit Null Label Support with BGP IPv4 Label Session

- You must configure your network for MPLS VPN.
- You must configure BGP to distribute labels between the CSC-CE and CSC-provider edge (PE) routers.

Restrictions for MPLS—VPN Explicit Null Label Support with BGP IPv4 Label Session

- Configure an explicit null label only in a CSC-CE topology.
- Configure an explicit null label only on a per-neighbor basis.

Information About MPLS—VPN Explicit Null Label Support with BGP IPv4 Label Session

To configure MPLS—VPN Explicit Null Label Support BGP IPv4 Label Session, you need to understand the following concepts:

- [Feature Design of MPLS—VPN Explicit Null Label Support with BGP IPv4 Label Session, page 2](#)
- [Benefits of MPLS—VPN Explicit Null Label Support BGP IPv4 Label Session, page 3](#)

Feature Design of MPLS—VPN Explicit Null Label Support with BGP IPv4 Label Session

On a CSC-CE with BGP IPv4 label distribution, BGP advertises an implicit null label for directly connected routes. This causes the previous hop (penultimate) router to do penultimate hop popping (PHP).

The MPLS—VPN Explicit Null Label Support BGP IPv4 Label Session feature makes the penultimate router swap the incoming label for (or impose) the explicit null label. This action forces the egress router to process the explicit null label by popping it and inspecting the packet that remains.

Benefits of MPLS—VPN Explicit Null Label Support BGP IPv4 Label Session

Improved Quality of Service

The explicit null label helps to preserve quality of service (QoS) bits from one service level agreement (SLA) to another until the packets reach their CSC-CE destination.

How to Configure MPLS—VPN Explicit Null Label Support with BGP IPv4 Label Session

This section contains the following procedures:

- [Configuring CSC with BGP, page 3](#) (required)
- [Verifying Explicit Null Configuration, page 4](#) (optional)

Configuring CSC with BGP

Perform this task to configure CSC with BGP.

SUMMARY STEPS

1. **enable**
2. **configure terminal**
3. **router bgp** *autonomous system-number*
4. **address-family ipv4** [**unicast**]
5. **neighbor** *ip-address* **send-label explicit-null**
6. **neighbor** {*ip-address* | *peer-group-name* | *ipv6-address*} **activate**
7. **exit**

DETAILED STEPS

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

	Command or Action	Purpose
Step 3	router bgp <i>autonomous system-number</i> Example: Router(config)# router bgp 100	Enters router configuration mode and configures the router to run a BGP process.
Step 4	address-family ipv4 [unicast] Example: Router(config-router)# address-family ipv4	Enters address family configuration mode for the IPv4 address family from which you can configure routing sessions that use standard IPv4 address prefixes.
Step 5	neighbor <i>ip-address</i> send-label explicit-null Example: Router(config-router-af)# neighbor 10.0.0.2 send-label explicit-null	Send-label advertises the capability of a router to send MPLS labels with BGP routes. <ul style="list-style-type: none">The explicit-null keyword allows a CSC-CE router to send labels with a value of 0 to its neighbor.
Step 6	neighbor { <i>ip-address</i> <i>peer-group-name</i> <i>ipv6-address</i> } activate Example: Router(config-router-af)# neighbor 192.168.99.70 activate	Enables the neighbor to exchange prefixes for the IPv4 address family with the local router.
Step 7	exit Example: Router(config-router-af)# exit	Exits address family configuration mode and returns the router to router configuration mode. <ul style="list-style-type: none">Repeat this step to exit router configuration mode and return the router to global configuration mode.

Verifying Explicit Null Configuration

Perform this task to verify that the explicit null option is configured.

SUMMARY STEPS

1. **enable**
2. **show ip bgp neighbors** [*ip-address* [**advertised-routes** | **dampened-routes** | **flap-statistics** | **paths** [*regex*] | **received prefix-filter** | **received-routes** | **routes**]]

DETAILED STEPS

	Command or Action	Purpose
Step 1	enable Example: Router> enable	Enables privileged EXEC mode. <ul style="list-style-type: none"> Enter your password if prompted.
Step 2	show ip bgp neighbors [<i>ip-address</i> [advertised-routes dampened-routes flap-statistics paths [<i>regex</i>] received prefix-filter received-routes routes]] Example: Router# show ip bgp neighbors	Displays information about the TCP and BGP connections to neighbors including explicit null. <ul style="list-style-type: none"> The optional <i>ip-address</i> argument displays the IP address of the neighbor whose routes you have learned. If you omit this argument, all neighbors are displayed. The optional advertised-routes keyword displays all the routes the router has advertised to the neighbor. The optional dampened-routes keyword displays the dampened routes to the neighbor at the IP address specified. The optional flap statistics keyword displays the flap statistics of the routes learned from the specified neighbor (eBGP peers only). The optional path <i>regex</i> keyword argument combination displays autonomous system paths learned from the specified neighbor. An optional regular expression can be used to filter the output. The optional received prefix-filter keyword displays the configured prefix list filter for the specified IP address. The optional received-routes keyword displays all received routes (both accepted and rejected) from the specified neighbor. The optional routes keyword displays all routes that are received and accepted. This is a subset of the output from the received-routes keyword.

Configuration Examples for MPLS—VPN Explicit Null Label Support with BGP IPv4 Label Session

This section provides the following configuration examples:

- [Configuring CSC-CE with BGP: Example, page 6](#)
- [Verifying Explicit Null Configuration: Example, page 6](#)

Configuring CSC-CE with BGP: Example

In the following CSC-CE example, CSC is configured with BGP to distribute labels and to advertise explicit null for all its connected routes:

```
Router(config-router-af)# neighbor 10.0.0.0 send-label explicit-null

router bgp 100
  bgp log-neighbor-changes
  neighbor 10.0.0.0 remote-as 200
  !
  address-family ipv4
    neighbor 10.0.0.0 activate
    neighbor 10.0.0.0 send-label explicit-null
  no auto-summary
  no synchronization
  exit-address-family
```

Verifying Explicit Null Configuration: Example

This section provides sample output for the **show ip bgp neighbors** command.

- [Sample Output for the show ip bgp neighbors Command, page 6](#)

Sample Output for the show ip bgp neighbors Command

In this example, the **show ip bgp neighbors** command displays information about connected BGP neighbors, including IP addresses, version numbers, neighbor capabilities, message statistics, and address family statistics that show if explicit null is configured.

```
Router# show ip bgp neighbors
```

```
BGP neighbor is 10.0.0.2, remote AS 300, external link
  BGP version 4, remote router ID 10.0.0.20
  BGP state = Established, up for 00:45:16
  Last read 00:00:16, hold time is 180, keepalive interval is 60 seconds
  Neighbor capabilities:
    Route refresh: advertised and received(new)
    Address family IPv4 Unicast: advertised and received
    ipv4 MPLS Label capability: advertised and received
  Message statistics:
    InQ depth is 0
    OutQ depth is 0

      Sent      Rcvd
  Opens:         1         1
  Notifications: 0         0
  Updates:       1         2
  Keepalives:    47        47
  Route Refresh: 0         0
  Total:        49        50
  Default minimum time between advertisement runs is 30 seconds
```

```
For address family: IPv4 Unicast
  BGP table version 9, neighbor version 9/0
  Output queue sizes : 0 self, 0 replicated
  Index 1, Offset 0, Mask 0x2
  Member of update-group 1
  My AS number is allowed for 3 number of times
  AF-dependant capabilities:
    Outbound Route Filter (ORF) type (128) Prefix-list:
    Sending Prefix & Label(advertise explicit-null set)      !Explicit null is configured
```

```

Prefix activity:
Prefixes Current:      3      3 (Consumes 144 bytes)
Prefixes Total:        3      6
Implicit Withdraw:     0      3
Explicit Withdraw:     0      0
.....
.....

```

Additional References

The following sections provide references related to the MPLS—VPN Explicit Null Label with BGP IPv4 Label Session feature.

Related Documents

Related Topic	Document Title
BGP configuration tasks	<i>Cisco IOS IP Routing Protocols Configuration Guide</i> , Release 12.4
BGP commands	<ul style="list-style-type: none"> <i>Cisco IOS IP Routing Protocols Command Reference</i>, Release 12.4T <i>Cisco IOS IP Routing Protocols Command Reference</i>, Release 12.2SB <i>Cisco IOS IP Routing Protocols Command Reference</i>, Release 12.2SR
IPv4 BGP label distribution	<ul style="list-style-type: none"> <i>MPLS VPN—InterAS—IPv4 BGP Label Distribution</i> <i>MPLS VPN—Carrier Supporting Carrier—IPv4 BGP Label Distribution</i>

Standards

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

MIBs

MIBs	MIBs Link
No new or modified MIBs are supported by this feature, and support for existing MIBs has not been modified by this feature.	<p>To locate and download MIBs for selected platforms, Cisco IOS releases, and feature sets, use Cisco MIB Locator found at the following URL:</p> <p>http://www.cisco.com/go/mibs</p>

RFCs

RFCs	Title
RFC 1163	<i>A Border Gateway Protocol</i>
RFC 1164	<i>Application of the Border Gateway Protocol in the Internet</i>
RFC 2283	<i>Multiprotocol Extensions for BGP-4</i>
RFC 2547	<i>BGP/MPLS VPNs</i>
RFC 3107	<i>Carrying Label Information in BGP-4</i>

Technical Assistance

Description	Link
<p>The Cisco Support website provides extensive online resources, including documentation and tools for troubleshooting and resolving technical issues with Cisco products and technologies.</p> <p>To receive security and technical information about your products, you can subscribe to various services, such as the Product Alert Tool (accessed from Field Notices), the Cisco Technical Services Newsletter, and Really Simple Syndication (RSS) Feeds.</p> <p>Access to most tools on the Cisco Support website requires a Cisco.com user ID and password.</p>	http://www.cisco.com/techsupport

Command Reference

This section documents only commands that are new or modified.

- [debug ip bgp](#)
- [neighbor send-label explicit-null](#)
- [show ip bgp neighbors](#)
- [show ip bgp vpnv4](#)
- [show mpls forwarding-table](#)

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 debugging output, use the **no** form of this command.

debug ip bgp [*A.B.C.D.* | **dampening** | **events** | **in** | **keepalives** | **out** | **updates** | **vpn4** | **mpls**]

no debug ip bgp [*A.B.C.D.* | **dampening** | **events** | **in** | **keepalives** | **out** | **updates** | **vpn4** | **mpls**]

Cisco 10000 Series Router

debug ip bgp [*A.B.C.D.* | **dampening** | **events** | **in** | **keepalives** | **out** | **updates** | **vpn4** | **mpls** | **all** | **groups** | **import** | **ipv4** | **ipv6**]

no debug ip bgp [*A.B.C.D.* | **dampening** | **events** | **in** | **keepalives** | **out** | **updates** | **vpn4** | **mpls** | **all** | **groups** | **import** | **ipv4** | **ipv6**]

Syntax Description

<i>A.B.C.D.</i>	(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.
vpn4	(Optional) Displays Virtual Private Network version 4 (VPNv4) Network Layer Reachability Information (NLRI).
mpls	(Optional) Displays the Multiprotocol Label Switching (MPLS) information.
all	(Optional) Displays all address family information.
groups	(Optional) Displays BGP configuration and update groups information.
import	(Optional) Displays BGP import routes to a VPN routing and forwarding (VRF) instance across address family information.
ipv4	(Optional) Displays BGP IPv4 address family information.
ipv6	(Optional) Displays BGP IPv6 address family information.

Command Modes

Privileged EXEC (#)

Command History

Release	Modification
12.0(5)T	This command was introduced.
12.0(21)ST	This command was integrated into Cisco IOS Release 12.0(21)ST. The mpls keyword was added.
12.0(22)S	This command was integrated into Cisco IOS Release 12.0(22)S.
12.0(23)S	This command was integrated into Cisco IOS Release 12.0(23)S.

Release	Modification
12.2(11)S	This feature was integrated into Cisco IOS Release 12.2(11)S.
12.2(13)T	The mpls keyword was added.
12.0(24)S	This command was integrated into Cisco IOS Release 12.0(24)S.
12.2(14)S	This command was integrated into Cisco IOS Release 12.2(14)S.
12.2(17b)SXA	This command was integrated into Cisco IOS Release 12.2(17b)SXA.
12.0(27)S	The command output was modified to show explicit-null label information.
12.2(28)SB	This command was integrated into Cisco IOS Release 12.2(28)SB.
12.2(31)SB	This command was integrated into Cisco IOS Release 12.2(31)SB.
12.2(33)SRA	This command was integrated into Cisco IOS Release 12.2(33)SRA.
12.2(33)SXH	This command was integrated into Cisco IOS Release 12.2(33)SXH.

Usage Guidelines

Use this command with the **updates** and **mpls** keywords to display explicit-null label information.

Examples

The following is sample output from the **debug ip bgp** command:

```
Router# debug ip bgp vpnv4

03:47:14:vpn:bgp_vpnv4_bnetinit:100:2:10.0.0.0/8
03:47:14:vpn:bnettable add:100:2:10.0.0.0/8
03:47:14:vpn:bestpath_hook route_tag_change for vpn2:10.0.0.0/255.0.0.0(ok)
03:47:14:vpn:bgp_vpnv4_bnetinit:100:2:10.0.0.0/8
03:47:14:vpn:bnettable add:100:2:10.0.0.0/8
03:47:14:vpn:bestpath_hook route_tag_change for vpn2:10.0.0.0/255.0.0.0(ok)
03:47:14:vpn:bgp_vpnv4_bnetinit:100:2:10.0.0.0/8
03:47:14:vpn:bnettable add:100:2:10.0.0.0/8
03:47:14:vpn:bestpath_hook route_tag_chacle ip bgp *nge for vpn2:10.0.0.0/255.0.0.0(ok)
```

The following example shows sample output, including the explicit-null label, from the **debug ip bgp updates** and the **debug ip bgp mpls** commands:

```
Router# debug ip bgp updates

BGP updates debugging is on

Router# debug ip bgp mpls

BGP MPLS labels debugging is on

Router#
01:33:53: BGP(0): route 10.10.10.10/32 up
01:33:53: BGP(0): nettable_walker 10.10.10.10/32 route sourced locally
01:33:53: BGP: adding MPLS label to 10.10.10.10/32
01:33:53: BGP: check on 10.10.10.10/8 in LDP - ok
01:33:53: BGP: label imp-null allocated via LDP
01:33:53: BGP-IPv4: send exp-null label for 10.10.10.10/32
01:33:53: BGP-IPv4: Send prefix 10.10.10.10/32, label exp-null      !explicit-null label
being sent
01:33:53: BGP(0): 10.10.10.11 send UPDATE (format) 10.10.10.10/32, next 10.10.10.12,
metric 0, path , mpls label 0 !label value is 0
01:33:53: BGP(0): updgrp 1 - 10.10.10.12 enqueued 1 updates, average/maximum size (bytes)
61/61
```

neighbor send-label explicit-null

To enable a Border Gateway Protocol (BGP) router to send Multiprotocol Label Switching (MPLS) labels with explicit-null information for a CSC-CE router and BGP routes to a neighboring CSC-PE router, use the **neighbor send-label explicit-null** command in address family configuration mode or router configuration mode. To disable a BGP router from sending MPLS labels with explicit-null information, use the **no** form of this command.

neighbor *ip-address* **send-label explicit-null**

no neighbor *ip-address* **send-label explicit-null**

Syntax Description	<i>ip-address</i>	IP address of the neighboring router.
---------------------------	-------------------	---------------------------------------

Command Default	None
------------------------	------

Command Modes	Address family configuration (config-router-af) Router configuration (config-router)
----------------------	---

Command History	Release	Modification
	12.0(27)S	This command was introduced.
	12.4	This command was integrated into Cisco IOS Release 12.4
	12.2(28)SB	This command was integrated into Cisco IOS Release 12.2(28)SB.
	12.2(33)SRA	This command was integrated into Cisco IOS Release 12.2(33)SRA.
	12.2(33)SXH	This command was integrated into Cisco IOS Release 12.2(33)SXH.

Usage Guidelines	This command enables a CSC-CE router to use BGP to distribute MPLS labels with a value of zero for explicit-null instead of implicit-null along with IPv4 routes to a CSC-PE peer router.
	You must issue this command only on the local CSC-CE router.
	You can use this command only with IPv4 addresses.

Examples	In the following CSC-CE example, CSC is configured with BGP to distribute labels and to advertise explicit null for all its connected routes:
-----------------	---

```
Router# configure terminal
```

```
Enter configuration commands, one per line. End with CNTL/Z.
```

```
Router(config)# router bgp 100
```

```
Router(config-router)# neighbor 10.0.0.2 remote-as 300
```

```
Router(config-router)# address-family ipv4
```

```
Router(config-router-af)# neighbor 10.0.0.2 send-label explicit-null
```

In the following CSC-PE example, CSC is configured with BGP to distribute labels:

```
Router# configure terminal
```

Enter configuration commands, one per line. End with CNTL/Z.

```
Router(config)# router bgp 300
```

```
Router(config-router)# neighbor 10.0.0.1 remote-as 100
```

```
Router(config-router)# address-family ipv4 vrf v1
```

```
Router(config-router-af)# neighbor 10.0.0.1 send-label
```

**Note**

Explicit null is not applicable on a CSC-PE router.

Related Commands

Command	Description
neighbor activate	Enables the exchange of information with a neighboring router.
neighbor send-label	Enables a BGP router to send MPLS labels with BGP routes to a neighboring BGP router.

show ip bgp neighbors

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

show ip bgp neighbors [*ip-address* [**advertised-routes** | **dampened-routes** | **flap-statistics** | **paths** [*reg-exp*] | **received prefix-filter** | **received-routes** | **routes** | **policy** [**detail**]]]

Syntax Description		
<i>ip-address</i>	(Optional) IP address of a neighbor. If this argument is omitted, all neighbors are displayed.	
advertised-routes	(Optional) Displays all routes that have been advertised to neighbors.	
dampened-routes	(Optional) Displays the dampened routes received from the specified neighbor.	
flap-statistics	(Optional) Displays the flap statistics of the routes learned from the specified neighbor (external BGP peers only).	
paths <i>reg-exp</i>	(Optional) Displays autonomous system paths learned from the specified neighbor. An optional regular expression can be used to filter the output.	
received prefix-filter	(Optional) Displays the prefix-list (outbound route filter [ORF]) sent from the specified neighbor.	
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. The output displayed when this keyword is entered is a subset of the output displayed by the received-routes keyword.	
policy	(Optional) Displays the policies applied to this neighbor per address family.	
detail	(Optional) Displays detailed policy information such as route maps, prefix lists, community lists, access control lists (ACLs) and AS-path filter lists.	

Command Default The output of this command displays information for all neighbors.

Command Modes User EXEC (>)
Privileged EXEC (#)

Command History	Release	Modification
	10.0	This command was introduced.
	11.2	The received-routes keyword was added.
	12.0(18)S	The output was modified to display the no-prepend configuration option and this command was integrated into Cisco IOS Release 12.0(18)S.
	12.2(4)T	The received and prefix-filter keywords were added, and this command was integrated into Cisco IOS Release 12.2(4)T.
	12.0(21)ST	The output was modified to display Multiprotocol Label Switching (MPLS) label information.

Release	Modification
12.0(22)S	Support for the BGP graceful restart capability was integrated into the output. Support for the Cisco 12000 series routers (Engine 0 and Engine 2) was also added.
12.2(14)S	This command was integrated into Cisco IOS Release 12.2(14)S.
12.2(15)T	Support for the BGP graceful restart capability was integrated into the output.
12.0(25)S	The policy and detail keywords were added.
12.2(17b)SXA	This command was integrated into Cisco IOS Release 12.2(17b)SXA.
12.0(27)S	The command output was modified to support the BGP TTL Security Check feature and to display explicit-null label information.
12.3(7)T	The command output was modified to support the BGP TTL Security Check feature and to display explicit-null label information.
12.0(31)S	Support for the Bidirectional Forwarding Detection (BFD) feature was integrated into the output.
12.2(18)SXE	Support for the Bidirectional Forwarding Detection (BFD) feature was integrated into the output.
12.2(28)SB	This command was integrated into Cisco IOS Release 12.2(28)SB.
12.4(4)T	Support for the Bidirectional Forwarding Detection (BFD) feature was integrated into the output.
12.2(33)SRA	This command was integrated into Cisco IOS Release 12.2(33)SRA, and the output was modified to support BGP TCP path MTU discovery.
12.4(11)T	Support for the policy and detail keywords was integrated into Cisco IOS Release 12.4(11)T.
12.2(33)SRB	Support for the policy and detail keywords was integrated into Cisco IOS Release 12.2(33)SRB.
12.2(33)SXH	This command was integrated into Cisco IOS Release 12.2(33)SXH.

Usage Guidelines

Use the **show ip bgp neighbors** command to display BGP and TCP connection information for neighbor sessions. For BGP, this includes detailed neighbor attribute, capability, path, and prefix information. For TCP, this includes statistics related to BGP neighbor session establishment and maintenance.

Prefix activity is displayed based on the number of prefixes that are advertised and withdrawn. Policy denials display the number of routes that were advertised but then ignored based the function or attribute that is displayed in the output.

Cisco IOS Release 12.0(25)S, 12.4(11)T, 12.2(33)SRB, and Later Releases

When BGP neighbors use multiple levels of peer templates it can be difficult to determine which policies are applied to the neighbor.

In Cisco IOS Release 12.0(25)S, 12.4(11)T, 12.2(33)SRB and later releases, the **policy** and **detail** keywords were added to display the inherited policies and the policies configured directly on the specified neighbor. Inherited policies are policies that the neighbor inherits from a peer-group or a peer-policy template.

Examples

Example output is different for the various keywords available for the **show ip bgp neighbors** command. To view the appropriate output, choose one of the following sections:

- [show ip bgp neighbors: Example, page 15](#)
- [show ip bgp neighbors advertised-routes: Example, page 21](#)
- [show ip bgp neighbors paths: Example, page 22](#)
- [show ip bgp neighbors received prefix-filter: Example, page 22](#)
- [show ip bgp neighbors policy: Example, page 23](#)
- [Cisco IOS Release 12.0\(31\)S and 12.4\(4\)T: Example, page 23](#)
- [Cisco IOS Release 12.2\(33\)SRA: Example, page 23](#)

show ip bgp neighbors: Example

The following example shows output for the BGP neighbor at 10.108.50.2. This neighbor is an internal BGP (iBGP) peer. This neighbor supports the route refresh and graceful restart capabilities.

```
Router# show ip bgp neighbors 10.108.50.2
```

```
BGP neighbor is 10.108.50.2, remote AS 1, internal link
BGP version 4, remote router ID 192.168.252.252
BGP state = Established, up for 00:24:25
Last read 00:00:24, last write 00:00:24, hold time is 180, keepalive interval is
60 seconds
```

Neighbor capabilities:

```
Route refresh: advertised and received(old & new)
MPLS Label capability: advertised and received
Graceful Restart Capability:advertised and received
Address family IPv4 Unicast: advertised and received
```

Message statistics:

```
InQ depth is 0
OutQ depth is 0
```

	Sent	Rcvd
Opens:	3	3
Notifications:	0	0
Updates:	0	0
Keepalives:	113	112
Route Refresh:	0	0
Total:	116	115

Default minimum time between advertisement runs is 5 seconds

For address family: IPv4 Unicast

BGP table version 1, neighbor version 1/0

Output queue size : 0

Index 1, Offset 0, Mask 0x2

1 update-group member

	Sent	Rcvd
Prefix activity:	----	----
Prefixes Current:	0	0
Prefixes Total:	0	0
Implicit Withdraw:	0	0
Explicit Withdraw:	0	0
Used as bestpath:	n/a	0
Used as multipath:	n/a	0

	Outbound	Inbound
Local Policy Denied Prefixes:	-----	-----
Total:	0	0

Number of NLRI in the update sent: max 0, min 0

show ip bgp neighbors

```

Connections established 3; dropped 2
Last reset 00:24:26, due to Peer closed the session
External BGP neighbor may be up to 2 hops away.
Connection state is ESTAB, I/O status: 1, unread input bytes: 0
Connection is ECN Disabled
Local host: 10.108.50.1, Local port: 179
Foreign host: 10.108.50.2, Foreign port: 42698

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

Event Timers (current time is 0x68B944):
Timer           Starts      Wakeups          Next
Retrans          27          0             0x0
TimeWait         0           0             0x0
AckHold          27          18            0x0
SendWnd          0           0             0x0
KeepAlive        0           0             0x0
GiveUp           0           0             0x0
PmtuAger         0           0             0x0
DeadWait         0           0             0x0

iss: 3915509457  snduna: 3915510016  sndnxt: 3915510016   sndwnd: 15826
irs: 233567076  rcvnxt: 233567616  rcvwnd: 15845  delrcvwnd: 539

SRTT: 292 ms, RTTO: 359 ms, RTV: 67 ms, KRTT: 0 ms
minRTT: 12 ms, maxRTT: 300 ms, ACK hold: 200 ms
Flags: passive open, nagle, gen tcbs
IP Precedence value : 6

Datagrams (max data segment is 1460 bytes):
Rcvd: 38 (out of order: 0), with data: 27, total data bytes: 539
Sent: 45 (retransmit: 0, fastretransmit: 0, partialack: 0, Second Congestion: 08

```

Table 1 describes the significant fields shown in the display. Fields that are preceded by the asterisk character are displayed only when the counter has a nonzero value.

Table 1 *show ip bgp neighbors Field Descriptions*

Field	Description
BGP neighbor	IP address of the BGP neighbor and its autonomous system number.
remote AS	Autonomous-system number of the neighbor.
local AS 300 no-prepend (not shown in display)	Verifies that the local autonomous system number is not prepended to received external routes. This output supports the hiding of the local autonomous systems when migrating autonomous systems.
internal link	“internal link” is displayed for iBGP neighbors. “external link” is displayed for external BGP (eBGP) neighbors.
BGP version	BGP version being used to communicate with the remote router.
remote router ID	IP address of the neighbor.
BGP state	Finite state machine (FSM) stage of session negotiation.
up for	Time, in hhmmss, that the underlying TCP connection has been in existence.
Last read	Time, in hhmmss, since BGP last received a message from this neighbor.

Table 1 *show ip bgp neighbors Field Descriptions (continued)*

Field	Description
last write	Time, in hhmmss, since BGP last sent a message to this neighbor.
hold time	Time, in seconds, that BGP will maintain the session with this neighbor without receiving a messages.
keepalive interval	Time, interval in seconds, that keepalive messages are transmitted to this neighbor.
Neighbor capabilities	BGP capabilities advertised and received from this neighbor. “advertised and received” is displayed when a capability is successfully exchanged between two routers.
Route Refresh	Status of the route refresh capability.
MPLS Label Capability	Indicates that MPLS labels are both sent and received by the eBGP peer.
Graceful Restart Capability	Status of the graceful restart capability.
Address family IPv4 Unicast	IP Version 4 unicast-specific properties of this neighbor.
Message statistics	Statistics organized by message type.
InQ depth is	Number of messages in the input queue.
OutQ depth is	Number of messages in the output queue.
Sent	Total number of transmitted messages.
Received	Total number of received messages.
Opens	Number of open messages sent and received.
notifications	Number of notification (error) messages sent and received.
Updates	Number of update messages sent and received.
Keepalives	Number of keepalive messages sent and received.
Route Refresh	Number of route refresh request messages sent and received.
Total	Total number of messages sent and received.
Default minimum time between...	Time, in seconds, between advertisement transmissions.
For address family:	Address family for which the following fields refer.
BGP table version	Internal version number of the table. This is the primary routing table with which the neighbor has been updated. The number increments when the table changes.
neighbor version	Number used by the software to track prefixes that have been sent and those that need to be sent.
...update-group	Number of update-group member for this address family.
Prefix activity	Prefix statistics for this address family.
Prefixes current	Number of prefixes accepted for this address family.
Prefixes total	Total number of received prefixes.
Implicit Withdraw	Number of times that a prefix has been withdrawn and readvertised.

Table 1 *show ip bgp neighbors Field Descriptions (continued)*

Field	Description
Explicit Withdraw	Number of times that prefix is withdrawn because it is no longer feasible.
Used as bestpath	Number of received prefixes installed as a best paths.
Used as multipath	Number of received prefixes installed as multipaths.
* Saved (soft-reconfig)	Number of soft resets performed with a neighbor that supports soft reconfiguration. This field is displayed only if the counter has a nonzero value.
* History paths	This field is displayed only if the counter has a nonzero value.
* Invalid paths	Number of invalid paths. This field is displayed only if the counter has a nonzero value.
Local Policy Denied Prefixes	Prefixes denied due to local policy configuration. Counters are updated for inbound and outbound policy denials. The fields under this heading are displayed only if the counter has a nonzero value.
* route-map	Displays inbound and outbound route-map policy denials.
* filter-list	Displays inbound and outbound filter-list policy denials.
* prefix-list	Displays inbound and outbound prefix-list policy denials.
* Ext Community	Displays only outbound extended community policy denials.
* AS_PATH too long	Displays outbound AS-path length policy denials.
* AS_PATH loop	Displays outbound AS-path loop policy denials.
* AS_PATH confed info	Displays outbound confederation policy denials.
* AS_PATH contains AS 0	Displays outbound denials of AS 0.
* NEXT_HOP Martian	Displays outbound martian denials.
* NEXT_HOP non-local	Displays outbound non-local next-hop denials.
* NEXT_HOP is us	Displays outbound next-hop-self denials.
* CLUSTER_LIST loop	Displays outbound cluster-list loop denials.
* ORIGINATOR loop	Displays outbound denials of local originated routes.
* unsuppress-map	Displays inbound denials due to an unsuppress-map.
* advertise-map	Displays inbound denials due to an advertise-map.
* VPN Imported prefix	Displays inbound denials of VPN prefixes.
* Well-known Community	Displays inbound denials of well-known communities.
* SOO loop	Displays inbound denials due to site-of-origin.
* Bestpath from this peer	Displays inbound denials because the bestpath came from the local router.
* Suppressed due to dampening	Displays inbound denials because the neighbor or link is in a dampening state.
* Bestpath from iBGP peer	Displays inbound denials because the bestpath came from an iBGP neighbor.
* Incorrect RIB for CE	Displays inbound denials due to RIB errors for a CE router.

Table 1 *show ip bgp neighbors Field Descriptions (continued)*

Field	Description
* BGP distribute-list	Displays inbound denials due to a distribute list.
Number of NLRIs...	Number of network layer reachability attributes in updates.
Connections established	Number of times a TCP and BGP connection have been successfully established.
dropped	Number of times that a valid session has failed or been taken down.
Last reset	Time since this peering session was last reset. The reason for the reset is displayed on this line.
External BGP neighbor may be... (not shown in the display)	Indicates that the BGP TTL security check is enabled. The maximum number hops that can separate the local and remote peer is displayed on this line.
Connection state	Connection status of the BGP peer.
Connection is ECN Disabled	Explicit congestion notification status (enabled or disabled).
Local host: 10.108.50.1, Local port: 179	IP address of the local BGP speaker. BGP port number 179.
Foreign host: 10.108.50.2, Foreign port: 42698	Neighbor address and BGP destination port number.
Enqueued packets for retransmit:	Packets queued for retransmission by TCP.
Event Timers	TCP event timers. Counters are provided for starts and wakeups (expired timers).
Retrans	Number of times a packet has been retransmitted.
TimeWait	Time waiting for the retransmission timers to expire.
AckHold	Acknowledgement hold timer.
SendWnd	Transmission (send) window.
KeepAlive	Number of keep alive packets.
GiveUp	Number times a packet is dropped due to no acknowledgement.
PmtuAger	Path MTU discovery timer.
DeadWait	Expiration timer for dead segments.
iss:	Initial packet transmission sequence number.
snduna:	Last transmission sequence number that has not been acknowledged.
sndnxt:	Next packet sequence number to be transmitted.
sndwnd:	TCP window size of the remote neighbor.
irs:	Initial packet receive sequence number.
rcvnxt:	Last receive sequence number that has been locally acknowledged.
rcvwnd:	TCP window size of the local host.

Table 1 *show ip bgp neighbors Field Descriptions (continued)*

Field	Description
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).
maxRTT:	Largest recorded round-trip timeout.
ACK hold:	Time the local host will delay an acknowledgment to carry (piggyback) additional data.
IP Precedence value:	IP precedence of the BGP packets.
Datagrams	Number of update packets received from a neighbor.
Rcvd:	Number of received packets.
with data	Number of update packets sent with data.
total data bytes	Total received in bytes.
Sent	Number of update packets sent.
Second Congestion	Number of update packets with data sent.
Datagrams: Rcvd	Number of update packets received from a neighbor.
out of order:	Number of packets received out of sequence.
with data	Number of update packets received with data.
Last reset	Elapsed time since this peering session was last reset.
unread input bytes	Number of bytes of packets still to be processed.
retransmit	Number of packets retransmitted.
fastretransmit	A duplicate acknowledgement is retransmitted for an out of order segment before the retransmission timer expires.
partialack	Number of retransmissions for partial acknowledgements (transmissions before or without subsequent acknowledgements).
Second Congestion	Second retransmission due to congestion.

show ip bgp neighbors advertised-routes: Example

The following example displays routes advertised for only the 172.16.232.178 neighbor:

```
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
*>i10.0.0.0      172.16.232.179    0      100      0 ?
*> 10.20.2.0     10.0.0.0          0              32768 i
```

Table 2 describes the significant fields shown in the display.

Table 2 *show ip bgp neighbors advertised-routes Field Descriptions*

Field	Description
BGP table version	Internal version number of the table. This is the primary routing table with which the neighbor has been updated. The number increments when the table changes.
local router ID	IP address of the local BGP speaker.
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: <ul style="list-style-type: none"> 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: <ul style="list-style-type: none"> 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 used to forward a packet to the destination network. An entry of 0.0.0.0 indicates that there are non-BGP routes in the path to the destination network.

Table 2 *show ip bgp neighbors advertised-routes Field Descriptions (continued)*

Field	Description
Metric	If shown, this is the value of the inter-autonomous system metric. This field is not used frequently.
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.

show ip bgp neighbors paths: Example

The following is example output from the **show ip bgp neighbors** command entered with the **paths** keyword:

```
Router# show ip bgp neighbors 172.29.232.178 paths ^10
```

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

[Table 3](#) describes the significant fields shown in the display.

Table 3 *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.

show ip bgp neighbors received prefix-filter: Example

The following example shows that a prefix-list the filters all routes in the 10.0.0.0 network has been received from the 192.168.20.72 neighbor:

```
Router# show ip bgp neighbors 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 4](#) describes the significant fields shown in the display.

Table 4 *show ip bgp neighbors received prefix-filter Field Descriptions*

Field	Description
Address family	Address family mode in which the prefix filter is received.
ip prefix-list	Prefix list sent from the specified neighbor.

show ip bgp neighbors policy: Example

The following sample output shows the policies applied to the neighbor at 192.168.1.2. The output displays both inherited policies and policies configured on the neighbor device. Inherited policies are policies that the neighbor inherits from a peer-group or a peer-policy template.

```
Router# show ip bgp neighbors 192.168.1.2 policy

Neighbor: 192.168.1.2, Address-Family: IPv4 Unicast
Locally configured policies:
  route-map ROUTE in
Inherited policies:
  prefix-list NO-MARKETING in
  route-map ROUTE in
  weight 300
  maximum-prefix 10000
```

Cisco IOS Release 12.0(31)S and 12.4(4)T: Example

The following is sample output from the **show ip bgp neighbors** command that verifies that BFD is being used to detect fast fallover for the BGP neighbor that is a BFD peer.

```
Router# show ip bgp neighbors

BGP neighbor is 172.16.10.2, remote AS 45000, external link
.
.
.
Using BFD to detect fast fallover
```

Cisco IOS Release 12.2(33)SRA: Example

The following is sample output from the **show ip bgp neighbors** command that verifies that BGP TCP path MTU discovery is enabled for the BGP neighbor at 172.16.1.2.

```
Router# show ip bgp neighbors 172.16.1.2

BGP neighbor is 172.16.1.2, remote AS 45000, internal link
  BGP version 4, remote router ID 172.16.1.99
.
.
.
For address family: IPv4 Unicast
  BGP table version 5, neighbor version 5/0
.
.
.
  Address tracking is enabled, the RIB does have a route to 172.16.1.2
  Address tracking requires at least a /24 route to the peer
  Connections established 3; dropped 2
  Last reset 00:00:35, due to Router ID changed
  Transport(tcp) path-mtu-discovery is enabled
.
.
.
SRTT: 146 ms, RTTO: 1283 ms, RTV: 1137 ms, KRTT: 0 ms
minRTT: 8 ms, maxRTT: 300 ms, ACK hold: 200 ms
Flags: higher precedence, retransmission timeout, nagle, path mtu capable
```

 `show ip bgp neighbors`**Related Commands**

Command	Description
neighbor send-label	Enables a BGP router to send MPLS labels with BGP routes to a neighboring BGP router.
neighbor send-label explicit-null	Enables a BGP router to send MPLS labels with explicit-null information for a CSC-CE router and BGP routes to a neighboring CSC-PE router.

show ip bgp vpnv4

To display Virtual Private Network Version 4 (VPNv4) address information from the Border Gateway Protocol (BGP) table, use the **show ip bgp vpnv4** command in user EXEC or privileged EXEC mode.

```
show ip bgp vpnv4 {all | rd route-distinguisher | vrf vrf-name} [rib-failure] [ip-prefix/length
[longer-prefixes]] [network-address [mask] [longer-prefixes]] [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 <i>route-distinguisher</i>		Displays Network Layer Reachability Information (NLRI) prefixes that match the named route distinguisher.
vrf <i>vrf-name</i>		Displays NLRI prefixes associated with the named VPN routing and forwarding (VRF) instance.
rib-failure		(Optional) Displays BGP routes that failed to install in the VRF table.
<i>ip-prefix/length</i>		(Optional) The IP prefix address (in dotted decimal format) and the length of the mask (0 to 32). The slash mark must be included.
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.
<i>network-address</i>		(Optional) The IP address of a network in the BGP routing table.
<i>mask</i>		(Optional) The mask of the network address, in dotted decimal format.
cidr-only		(Optional) Displays only routes that have nonclassful net masks.
community		(Optional) Displays routes that match this community.
community-list		(Optional) Displays routes that match this community list.
dampened-paths		(Optional) Displays paths suppressed because of dampening (BGP route from peer is up and down).
filter-list		(Optional) Displays routes that conform 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.
<i>line</i>		(Optional) A regular expression to match the BGP autonomous system paths.
peer-group		(Optional) Displays information about peer groups.
quote-regexp		(Optional) Displays routes that match the autonomous system path regular expression.
regexp		(Optional) Displays routes that match the autonomous system path regular expression.

```
show ip bgp vpnv4
```

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

Command Modes

User EXEC (>)
Privileged 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 <i>ip-prefix</i> command was enhanced to display attributes including multipaths and a best path to the specified network.
12.0(21)ST	The tags keyword was replaced by the labels keyword to conform to the MPLS guidelines. This command was integrated into Cisco IOS Release 12.0(21)ST.
12.2(14)S	This command was integrated into Cisco IOS Release 12.2(14)S.
12.0(22)S	This command was integrated into Cisco IOS Release 12.0(22)S.
12.2(13)T	This command was integrated into Cisco IOS Release 12.2(13)T.
12.0(27)S	The output of the show ip bgp vpnv4 all labels command was enhanced to display explicit-null label information.
12.3	The rib-failure keyword was added for VRFs.
12.2(22)S	The output of the show ip bgp vpnv4 vrf <i>vrf-name</i> labels command was modified so that directly connected VRF networks no longer display as aggregate; no label appears instead.
12.2(25)S	This command was updated to display MPLS VPN nonstop forwarding information.
12.2(28)SB	This command was integrated into Cisco IOS Release 12.2(28)SB and implemented on the Cisco 10000 series router. The display output was modified to indicate whether BGP Nonstop Routing (NSR) with stateful switchover (SSO) is enabled and the reason the last BGP lost SSO capability.
12.2(33)SRA	This command was integrated into Cisco IOS Release 12.2(33)SRA, and the output was modified to support per-VRF assignment of the BGP router ID.
12.2(31)SB2	The output was modified to support per-VRF assignment of the BGP router ID.
12.2(33)SXH	This command was integrated into Cisco IOS Release 12.2(33)SXH.

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. The **show ip bgp vpnv4 all labels** command displays explicit-null label information.

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 10.14.14.14
```

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)
*>i10.6.6.6/32        10.0.0.21          11      100      0 ?
*> 10.7.7.7/32        10.150.0.2         11      32768 ?
*>i10.69.0.0/30       10.0.0.21          0       100      0 ?
*> 10.150.0.0/24      10.0.0.0            0      32768 ?

```

Table 5 describes the significant fields shown in the display.

Table 5 *show ip bgp vpnv4 all 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 NLRI prefixes 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)
 10.0.0.0         10.20.0.60        34/nolabel
 10.0.0.0         10.20.0.60        35/nolabel
 10.0.0.0         10.20.0.60        26/nolabel
                  10.20.0.60        26/nolabel
 10.0.0.0         10.15.0.15        nolabel/26

```

Table 6 describes the significant fields shown in the display.

Table 6 *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 10.14.14.14
 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

```

```
show ip bgp vpnv4
```

```
Route Distinguisher: 1:101 (default for vrf vpn1)
*>i10.6.6.6/32      10.0.0.21      11      100      0 ?
*> 10.7.7.7/32      10.150.0.2      11      32768 ?
*>i10.69.0.0/30     10.0.0.21      0       100      0 ?
*> 10.150.0.0/24    10.0.0.0        0      32768 ?
*> 10.0.0.1/32      10.150.0.2      11      32768 ?
*>i10.0.0.3/32      10.0.0.21      11      100      0 ?
```

Table 7 describes the significant fields shown in the display.

Table 7 *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 include multipaths and a best path:

```
Router# show ip bgp vpnv4 all 10.22.22.0
```

```
BGP routing table entry for 10:1:10.22.22.0/24, version 50
Paths: (6 available, best #1)
Multipath:iBGP
  Advertised to non peer-group peers:
    10.1.12.12
    22
    10.22.7.8 (metric 11) from 10.11.3.4 (10.0.0.8)
      Origin IGP, metric 0, localpref 100, valid, internal, multipath, best
      Extended Community:RT:100:1
      Originator:10.0.0.8, Cluster list:10.1.1.44
    22
    10.22.1.9 (metric 11) from 10.11.1.2 (10.0.0.9)
      Origin IGP, metric 0, localpref 100, valid, internal, multipath
      Extended Community:RT:100:1
      Originator:10.0.0.9, Cluster list:10.1.1.22
```

Table 8 describes the significant fields shown in the display.

Table 8 *show ip bgp vpnv4 all network-address Field Descriptions*

Field	Description
BGP routing table entry for ... 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	IP address of the BGP peers to which the specified route is advertised.

Table 8 *show ip bgp vpnv4 all network-address Field Descriptions (continued)*

Field	Description
10.22.7.8 (metric 11) from 10.11.3.4 (10.0.0.8)	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: <ul style="list-style-type: none"> 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 <i>internal</i> if the path is learned via iBGP. The field is <i>external</i> 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 this path is advertised to neighbors.
Extended Community	Route Target value associated with the specified route.
Originator	The router ID of the router from which the route originated when route reflector is used.
Cluster list	The router ID of all the route reflectors that the specified route has passed through.

The following example shows routes that BGP could not install in the VRF table:

```
Router# show ip bgp vpnv4 vrf xyz rib-failure
```

```

Network          Next Hop          RIB-failure    RIB-NH Matches
Route Distinguisher: 2:2 (default for vrf bar)
10.1.1.2/32      10.100.100.100    Higher admin distance    No
10.111.111.112/32 10.9.9.9          Higher admin distance    Yes
```

Table 9 describes the significant fields shown in the display.

Table 9 *show ip bgp vpnv4 vrf rib-failure Field Descriptions*

Field	Description
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.

Table 9 *show ip bgp vpnv4 vrf rib-failure Field Descriptions (continued)*

Field	Description
RIB-failure	Cause of the Routing Information Base (RIB) failure. Higher admin distance means that a route with a better (lower) administrative distance, such as a static route, already exists in the IP routing table.
RIB-NH Matches	Route status that applies only when Higher admin distance appears in the RIB-failure column and the bgp suppress-inactive command is configured for the address family being used. There are three choices: <ul style="list-style-type: none"> • Yes—Means that the route in the RIB has the same next hop as the BGP route or that the next hop recurses down to the same adjacency as the BGP next hop. • No—Means that the next hop in the RIB recurses down differently from the next hop of the BGP route. • n/a—Means that the bgp suppress-inactive command is not configured for the address family being used.

The following example shows the information displayed on the active and standby Route Processors when they are configured for MPLS VPN nonstop forwarding.

Active Route Processor

Router# **show ip bgp vpnv4 all labels**

```

Network      Next Hop    In label/Out label
Route Distinguisher: 100:1 (vpn1)
10.12.12.12/32 10.0.0.0    16/aggregate(vpn1)
10.0.0.0/8     10.0.0.0    17/aggregate(vpn1)
Route Distinguisher: 609:1 (vpn0)
10.13.13.13/32 10.0.0.0    18/aggregate(vpn0)

```

Router# **show ip bgp vpnv4 vrf vpn1 labels**

```

Network      Next Hop    In label/Out label
Route Distinguisher: 100:1 (vpn1)
10.12.12.12/32 10.0.0.0    16/aggregate(vpn1)
10.0.0.0/8     10.0.0.0    17/aggregate(vpn1)

```

Standby Route Processor

Router# **show ip bgp vpnv4 all labels**

```

Network      Masklen    In label
Route Distinguisher: 100:1
10.12.12.12   /32        16
10.0.0.0      /8         17
Route Distinguisher: 609:1
10.13.13.13   /32        18

```

Router# **show ip bgp vpnv4 vrf vpn1 labels**

```

Network      Masklen    In label
Route Distinguisher: 100:1
10.12.12.12   /32        16
10.0.0.0      /8         17

```

Table 10 describes the significant fields shown in the display.

Table 10 *show ip bgp vpnv4 labels Field Descriptions*

Field	Description
Network	The network address from the BGP table.
Next Hop	The BGP next hop address.
In label	The label (if any) assigned by this router.
Out label	The label assigned by the BGP next hop router.
Masklen	The mask length of the network address.

The following example displays output, including the explicit-null label, from the **show ip bgp vpnv4 all labels** command on a CSC-PE router:

```
Router# show ip bgp vpnv4 all labels

      Network      Next Hop      In label/Out label
Route Distinguisher: 100:1 (v1)
  10.0.0.0/24      10.0.0.0      19/aggregate(v1)
  10.0.0.1/32      10.0.0.0      20/nolabel
  10.1.1.1/32      10.0.0.0      21/aggregate(v1)
  10.10.10.10/32   10.0.0.1      25/exp-null
  10.168.100.100/32 10.0.0.1      23/exp-null
  10.168.101.101/32 10.0.0.1      22/exp-null
```

Table 11 describes the significant fields shown in the display.

Table 11 *show ip bgp vpnv4 all labels Field Descriptions*

Field	Description
Network	Displays the network address from the BGP table.
Next Hop	Displays the address of the BGP next hop.
In label	Displays the label (if any) assigned by this router.
Out label	Displays the label assigned by the BGP next hop router.
Route Distinguisher	Displays an 8-byte value added to an IPv4 prefix to create a VPN IPv4 prefix.

The following example displays separate router IDs for each VRF in the output from an image in Cisco IOS Release 12.2(33)SRA, 12.2(31)SB2 and later releases with the Per-VRF Assignment of BGP Router ID feature configured. The router ID is shown next to the VRF name.

```
Router# show ip bgp vpnv4 all

BGP table version is 5, local router ID is 172.17.1.99
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
               r RIB-failure, S Stale
Origin codes: i - IGP, e - EGP, ? - incomplete

      Network      Next Hop      Metric LocPrf Weight Path
Route Distinguisher: 1:1 (default for vrf vrf_trans) VRF Router ID 10.99.1.2
*> 192.168.4.0      10.0.0.0      0          32768 ?
```

show ip bgp vpnv4

```

Route Distinguisher: 42:1 (default for vrf vrf_user) VRF Router ID 10.99.1.1
*> 192.168.5.0      10.0.0.0      0      32768 ?

```

Table 12 describes the significant fields shown in the display.

Table 12 *show ip bgp vpnv4 all (VRF Router ID) Field Descriptions*

Field	Description
Route Distinguisher	Displays an 8-byte value added to an IPv4 prefix to create a VPN IPv4 prefix.
vrf	Name of the VRF.
VRF Router ID	Router ID for the VRF.

Related Commands

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

show mpls forwarding-table

To display the contents of the Multiprotocol Label Switching (MPLS) Label Forwarding Information Base (LFIB), use the **show mpls forwarding-table** command in privileged EXEC mode.

show mpls forwarding-table [*network* {*mask* | *length*} | **labels** *label* [- *label*] | **interface** *interface* | **next-hop** *address* | **lsp-tunnel** [*tunnel-id*]] [**vrf** *vrf-name*] [**detail**]

Syntax Description	
<i>network</i>	(Optional) Destination network number.
<i>mask</i>	IP address of the destination mask whose entry is to be shown.
<i>length</i>	Number of bits in the mask of the destination.
labels <i>label</i> - <i>label</i>	(Optional) Displays only entries with the specified local labels.
interface <i>interface</i>	(Optional) Displays only entries with the specified outgoing interface.
next-hop <i>address</i>	(Optional) Displays only entries with the specified neighbor as the next hop.
lsp-tunnel	(Optional) Displays only entries with the specified label switched path (LSP) tunnel, or with all LSP tunnel entries.
<i>tunnel-id</i>	(Optional) Specifies the LSP tunnel for which to display entries.
vrf <i>vrf-name</i>	(Optional) Displays only entries with the specified VPN routing and forwarding (VRF) instance.
detail	(Optional) Displays information in long form (includes length of encapsulation, length of MAC string, maximum transmission unit (MTU), and all labels).

Command Modes Privileged EXEC (#)

Command History	Release	Modification
	11.1CT	This command was introduced.
	12.1(3)T	This command was updated with MPLS terminology and command syntax.
	12.2(8)T	The command was modified to accommodate use of the MPLS experimental (EXP) level as a selection criterion for packet forwarding. The output display was modified to include a bundle adjacency field and exp (vcd) values when the optional detail keyword is specified.
	12.0(22)S	IPv6 MPLS aggregate label and prefix information was added to the display.
	12.2(14)S	This command was integrated into Cisco IOS Release 12.2(14)S.
	12.0(27)S	The command output was modified to include explicit-null label information.
	12.2(25)S	The output was changed in the following ways: <ul style="list-style-type: none"> The term “tag” was replaced with the term “label.” The term “untagged” was replaced with the term “no label.”
	12.0(29)S	This command was integrated into Cisco IOS Release 12.0(29)S.
	12.2(28)SB	This command was integrated into Cisco IOS Release 12.2(28)SB and implemented on the Cisco 10000 series routers.

Release	Modification
12.2(33)SRA	This command was integrated into Cisco IOS Release 12.2(33)SRA.
12.2(33)SXH	This command was integrated into Cisco IOS Release 12.2(33)SXH.

Examples

The following is sample output from the **show mpls forwarding-table** command:

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

Local Label	Outgoing Label or VC	Prefix or Tunnel Id	Bytes switched	label	Outgoing interface	Next Hop
26	No Label	10.253.0.0/16	0		Et4/0/0	10.27.32.4
28	1/33	10.15.0.0/16	0		AT0/0.1	point2point
29	Pop Label	10.91.0.0/16	0		Hs5/0	point2point
	1/36	10.91.0.0/16	0		AT0/0.1	point2point
30	32	10.250.0.97/32	0		Et4/0/2	10.92.0.7
	32	10.250.0.97/32	0		Hs5/0	point2point
34	26	10.77.0.0/24	0		Et4/0/2	10.92.0.7
	26	10.77.0.0/24	0		Hs5/0	point2point
35	No Label[T]	10.100.100.101/32	0		Tu301	point2point
36	Pop Label	10.1.0.0/16	0		Hs5/0	point2point
	1/37	10.1.0.0/16	0		AT0/0.1	point2point

[T] Forwarding through a TSP tunnel.
View additional labeling info with the 'detail' option

The following is sample output from the **show mpls forwarding-table** command when the IPv6 Provider Edge Router over MPLS feature is configured to allow IPv6 traffic to be transported across an IPv4 MPLS backbone. The labels are aggregated because there are several prefixes for one local label, and the prefix column contains “IPv6” instead of a target prefix.

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

Local Label	Outgoing Label or VC	Prefix or Tunnel Id	Bytes switched	label	Outgoing interface	Next Hop
16	Aggregate	IPv6	0			
17	Aggregate	IPv6	0			
18	Aggregate	IPv6	0			
19	Pop Label	192.168.99.64/30	0		Se0/0	point2point
20	Pop Label	192.168.99.70/32	0		Se0/0	point2point
21	Pop Label	192.168.99.200/32	0		Se0/0	point2point
22	Aggregate	IPv6	5424			
23	Aggregate	IPv6	3576			
24	Aggregate	IPv6	2600			

The following is sample output from the **show mpls forwarding-table** command when you specify the **detail** keyword. If the MPLS EXP level is used as a selection criterion for packet forwarding, a bundle adjacency exp (vcd) field is included in the display. This field includes the EXP value and the corresponding virtual circuit descriptor (VCD) in parentheses. The line in the output that reads “No output feature configured” indicates that the MPLS egress NetFlow accounting feature is not enabled on the outgoing interface for this prefix.

```
Router# show mpls forwarding-table detail
```

Local label	Outgoing label or VC	Prefix or Tunnel Id	Bytes switched	label	Outgoing interface	Next Hop
16	Pop label	10.0.0.6/32	0		AT1/0.1	point2point
Bundle adjacency exp(vcd)						
0(1) 1(1) 2(1) 3(1) 4(1) 5(1) 6(1) 7(1)						
MAC/Encaps=12/12, MTU=4474, label Stack{}						

```

00010000AAAA0300000008847
No output feature configured
17 18 10.0.0.9/32 0 AT1/0.1 point2point
Bundle adjacency exp(vcd)
0(1) 1(1) 2(1) 3(1) 4(1) 5(1) 6(1) 7(1)
MAC/Encaps=12/16, MTU=4470, label Stack{18}
00010000AAAA0300000008847 00012000
No output feature configured
18 19 10.0.0.10/32 0 AT1/0.1 point2point
Bundle adjacency exp(vcd)
0(1) 1(1) 2(1) 3(1) 4(1) 5(1) 6(1) 7(1)
MAC/Encaps=12/16, MTU=4470, label Stack{19}
00010000AAAA0300000008847 00013000
No output feature configured
19 17 10.0.0.0/8 0 AT1/0.1 point2point
Bundle adjacency exp(vcd)
0(1) 1(1) 2(1) 3(1) 4(1) 5(1) 6(1) 7(1)
MAC/Encaps=12/16, MTU=4470, label Stack{17}
00010000AAAA0300000008847 00011000
No output feature configured
20 20 10.0.0.0/8 0 AT1/0.1 point2point
Bundle adjacency exp(vcd)
0(1) 1(1) 2(1) 3(1) 4(1) 5(1) 6(1) 7(1)
MAC/Encaps=12/16, MTU=4470, label Stack{20}
00010000AAAA0300000008847 00014000
No output feature configured
21 Pop label 10.0.0.0/24 0 AT1/0.1 point2point
Bundle adjacency exp(vcd)
0(1) 1(1) 2(1) 3(1) 4(1) 5(1) 6(1) 7(1)
MAC/Encaps=12/12, MTU=4474, label Stack{}
00010000AAAA0300000008847
No output feature configured
22 Pop label 10.0.0.4/32 0 Et2/3 10.0.0.4
MAC/Encaps=14/14, MTU=1504, label Stack{}
000427AD10430005DDFE043B8847
No output feature configured

```

The following is sample output from the **show mpls forwarding-table** command when you use the **detail** keyword. In this example, the MPLS egress NetFlow accounting feature is enabled on the first three prefixes, as indicated by the line in the output that reads “Feature Quick flag set.”

Router# **show mpls forwarding-table detail**

```

Local   Outgoing   Prefix           Bytes label  Outgoing   Next Hop
label   label or VC or Tunnel Id      switched    interface
16      Aggregate  10.0.0.0/8[V]    0
      MAC/Encaps=0/0, MTU=0, label Stack{}
      VPN route: vpn1
      Feature Quick flag set
Per-packet load-sharing, slots: 0 1 2 3 4 5 6 7 8 9 10 11 12 13 14 15
17      No label   10.0.0.0/8[V]    0            Et0/0/2     10.0.0.1
      MAC/Encaps=0/0, MTU=1500, label Stack{}
      VPN route: vpn1
      Feature Quick flag set
Per-packet load-sharing, slots: 0 1 2 3 4 5 6 7 8 9 10 11 12 13 14 15
18      No label   10.42.42.42/32[V] 4185         Et0/0/2     10.0.0.1
      MAC/Encaps=0/0, MTU=1500, label Stack{}
      VPN route: vpn1
      Feature Quick flag set
Per-packet load-sharing, slots: 0 1 2 3 4 5 6 7 8 9 10 11 12 13 14 15
19      2/33       10.41.41.41/32    0            AT1/0/0.1   point2point
      MAC/Encaps=4/8, MTU=4470, label Stack{2/33(vcd=2)}
      00028847 00002000
      No output feature configured

```

Cisco 10000 Series Examples

The following is sample output from the **show mpls forwarding-table** command:

Router# **show mpls forwarding-table**

Local Label	Outgoing Label or VC	Prefix or Tunnel Id	Bytes Label Switched	Outgoing interface	Next Hop
16	Pop Label	10.0.0.0/8	0	Fa1/0/0	10.0.0.2
	Pop Label	10.0.0.0/8	0	Fa1/1/0	10.0.0.2
17	Aggregate	10.0.0.0/8[V]	570	vpn2	
21	Pop Label	10.11.11.11/32	0	Fa1/0/0	10.0.0.2
22	Pop Label	10.12.12.12/32	0	Fa1/1/0	10.0.0.2
23	No Label	10.3.0.0/16[V]	0	Fa4/1/0	10.0.0.2

The following is Cisco 10000 series sample output from the **show mpls forwarding-table** command when you specify the **detail** keyword:

Router# **show mpls forwarding-table detail**

Local Label	Outgoing Label or VC	Prefix or Tunnel Id	Bytes Label Switched	Outgoing interface	Next Hop
16	Pop Label	10.0.0.0/8	0	Fa1/0/0	10.0.0.2
	MAC/Encaps=14/14, MRU=1500, Label Stack{}				
	000B45C93889000B45C930218847				
	No output feature configured				
	Pop Label	10.0.0.0/8	0	Fa1/1/0	10.0.0.2
	MAC/Encaps=14/14, MRU=1500, Label Stack{}				
	000B45C92881000B45C930288847				
	No output feature configured				
17	Aggregate	10.0.0.0/8[V]	570	vpn2	
	MAC/Encaps=0/0, MRU=0, Label Stack{}				
	VPN route: vpn2				
	No output feature configured				
21	Pop Label	10.11.11.11/32	0	Fa1/0/0	10.0.0.2
	MAC/Encaps=14/14, MRU=1500, Label Stack{}				
	000B45C93889000B45C930218847				
	No output feature configured				

Table 13 describes the significant fields shown in the displays.

Table 13 *show mpls forwarding-table Field Descriptions*

Field	Description
Local label	Label assigned by this router.
Outgoing Label or VC Note VC is not applicable to the Cisco 10000 series routers.	<p>Label assigned by the next hop or virtual path identifier (VPI)/virtual channel identifier (VCI) used to get to next hop. The entries in this column include the following:</p> <p>Note VPI and VCI are not applicable to the Cisco 10000 series routers.</p> <ul style="list-style-type: none"> • [T]—Means forwarding through an LSP tunnel. • No Label—Means that there is no label for the destination from the next hop or that label switching is not enabled on the outgoing interface. • Pop Label—Means that the next hop advertised an implicit NULL label for the destination and that the router popped the top label. • Aggregate—Means there are several prefixes for one local label. This entry is used when IPv6 is configured on edge routers to transport IPv6 traffic over an IPv4 MPLS network. <p>Note IPv6 traffic is not applicable to the Cisco 10000 series routers.</p>
Prefix or Tunnel Id	<p>Address or tunnel to which packets with this label are sent.</p> <p>Note If IPv6 is configured on edge routers to transport IPv6 traffic over an IPv4 MPLS network, “IPv6” is displayed here.</p>
Bytes Label Switched	Number of bytes switched with this incoming label.
Outgoing interface	Interface through which packets with this label are sent.
Next Hop	IP address of the neighbor that assigned the outgoing label.
Bundle adjacency exp(vcd)	Bundle adjacency information. Includes the MPLS EXP value and the corresponding VCD.
MAC/Encaps	Length in bytes of the Layer 2 header and length in bytes of the packet encapsulation, including the Layer 2 header and label header.
MTU	MTU of the labeled packet.
Label Stack	<p>All the outgoing labels. If the outgoing interface is transmission convergence (TC)-ATM, the VCD is also shown.</p> <p>Note TC-ATM is not applicable to the Cisco 10000 series routers.</p>
00010000AAAA0300000008847 00013000	The actual encapsulation in hexadecimal form. A space is shown between Layer 2 and the label header.

Explicit-Null Label Example

The following example shows output, including the explicit-null label = 0 (commented in bold), from the **show mpls forwarding-table** command on a CSC-PE router:

Router# **show mpls forwarding-table**

Local label	Outgoing label or VC	Prefix or Tunnel Id	Bytes label switched	Outgoing interface	Next Hop
17	Pop label	10.10.0.0/32	0	Et2/0	10.10.0.1
18	Pop label	10.10.10.0/24	0	Et2/0	10.10.0.1
19	Aggregate	10.10.20.0/24[V]	0		
20	Pop label	10.10.200.1/32[V]	0	Et2/1	10.10.10.1
21	Aggregate	10.10.1.1/32[V]	0		
22	0	192.168.101.101/32[V]	\		
			0	Et2/1	192.168.101.101
23	0	192.168.101.100/32[V]	\		
			0	Et2/1	192.168.101.100
25	0	192.168.102.125/32[V]	0	Et2/1	192.168.102.125 !outlabel value 0

Table 14 describes the significant fields shown in the display.

Table 14 *show mpls forwarding-table Field Descriptions*

Field	Description
Local label	Label assigned by this router.
Outgoing label or VC	Label assigned by the next hop or VPI/VCI used to get to next hop. The entries this column include the following: <ul style="list-style-type: none"> [T]—Means forwarding through an LSP tunnel. No label—Means that there is no label for the destination from the next hop or that label switching is not enabled on the outgoing interface. Pop label—Means that the next hop advertised an implicit NULL label for the destination and that this router popped the top label. Aggregate—Means there are several prefixes for one local label. Used when IPv6 is configured on edge routers to transport IPv6 traffic over an IPv4 MPLS network. 0—Means the explicit null label value = 0.
Prefix or Tunnel Id	Address or tunnel to which packets with this label are going. <p>Note If IPv6 is configured on edge routers to transport IPv6 traffic over an IPv4 MPLS network, IPv6 is displayed here.</p>
Bytes label switched	Number of bytes switched with this incoming label.
Outgoing interface	Interface through which packets with this label are sent.
Next Hop	IP address of the neighbor that assigned the outgoing label.

Related Commands

Command	Description
neighbor send-label	Enables a BGP router to send MPLS labels with BGP routes to a neighboring BGP router.
neighbor send-label explicit-null	Enables a BGP router to send MPLS labels with explicit-null information for a CSC-CE router and BGP routes to a neighboring CSC-PE router.

Feature Information for MPLS VPN—Explicit Null Label Support with BGP IPv4 Label Session

Table 15 lists the release history for this feature.

Not all commands may be available in your Cisco IOS software release. For release information about a specific command, see the command reference documentation.

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



Note

Table 15 lists only the Cisco IOS software release that introduced support for a given feature in a given Cisco IOS software release train. Unless noted otherwise, subsequent releases of that Cisco IOS software release train also support that feature.

Table 15 Feature Information for MPLS VPN—Explicit Null Label Support with BGP IPv4 Label Session

Feature Name	Releases	Feature Information
MPLS VPN—Explicit Null Label Support with BGP IPv4 Label Session	12.0(27)S 12.0(27)S1 12.2(28)SB 12.2(33)SRA 12.2(33)SXH	<p>The MPLS VPN—Explicit Null Label Support with BGP IPv4 Label Session feature provides a method to advertise explicit null in a BGP label session for a carrier supporting carrier (CSC) customer edge (CE) router.</p> <p>In 12.0(27)S, this feature was introduced.</p> <p>In 12.0(27)S1, support was added for the Cisco 12000 series Internet routers.</p> <p>In 12.2(28)SB, support was added for the Cisco 10000 series router.</p> <p>This feature was integrated into Cisco IOS Release 12.2(33)SRA.</p> <p>This feature was integrated into Cisco IOS Release 12.2(33)SXH.</p>

Glossary

BGP—Border Gateway Protocol. 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.

CE router—customer edge router. A router on the border between a VPN provider and a VPN customer that belongs to the customer.

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

label—A short, fixed-length data identifier that tells switching nodes how to forward data (packets or cells).

label distribution—The techniques and processes used to cause routed traffic to travel through the network on a path other than the one that would have been chosen if standard routing methods had been used.

LDP—Label Distribution Protocol. The protocol that supports MPLS hop-by-hop forwarding by distributing bindings between labels and network prefixes. The Cisco proprietary version of this protocol is the Tag Distribution Protocol (TDP).

LSP—label-switched path. A configured connection between two routers, in which MPLS is used to carry packets. A path created by the concatenation of one or more label switched hops, allowing a packet to be forwarded by swapping labels from an MPLS node to another MPLS node.

MPLS—Multiprotocol Label Switching. A method for directing packets primarily through Layer 2 switching rather than Layer 3 routing. In MPLS, packets are assigned short, fixed-length labels at the ingress to an MPLS cloud by using the concept of forwarding equivalence classes. Within the MPLS domain, the labels are used to make forwarding decisions mostly without recourse to the original packet headers; formerly known as tag switching.

NLRI—Network Layer Reachability Information. 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.

PE router—provider edge router. A router on the border between a VPN provider and a VPN customer that belongs to the provider.

QoS—quality of service. A measure of performance for a transmission system that reflects its transmission quality and service availability.

router—A network layer device that uses one or more metrics to determine the optimal path along which network traffic should be forwarded. Routers forward packets from one network to another based on network layer information.

VPN—Virtual Private Network. A secure IP-based network that shares resources on one or more physical networks. A VPN contains geographically dispersed sites that can communicate securely over a shared backbone.

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

All other trademarks mentioned in this document or Website are the property of their respective owners. The use of the word partner does not imply a partnership relationship between Cisco and any other company. (0705R)

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

© 2004, 2006–2007 Cisco Systems, Inc. All rights reserved.