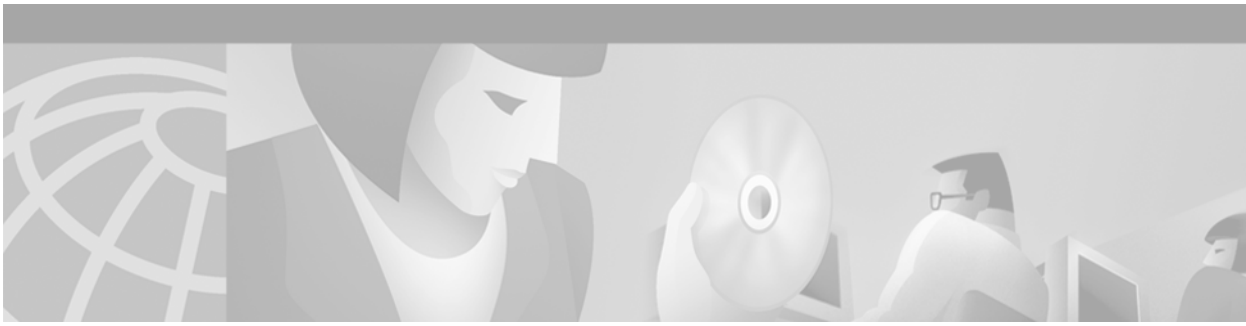

BGP Soft Reset Enhancement

Feature History

Release	Modification
12.0(7)T	This feature was introduced.
12.0(2)S	This feature was integrated into Cisco IOS Release 12.0(2)S.
12.0(22)S	Support for IPv6 and VPNv4 address family support was added.

This document describes the BGP Soft Reset Enhancement feature in Cisco IOS Release 12.0(22)S and includes the following sections:

Feature Overview

The BGP Soft Reset Enhancement feature provides automatic support for dynamic soft reset of inbound BGP routing table updates that are not dependent upon stored routing table update information. This method requires no preconfiguration (as with the **neighbor soft-reconfiguration** command) and requires much less memory than the previous soft reset method for inbound routing table updates.

Configurable routing policies for a peer, including route-maps, distribute-lists, prefix-lists, and filter-lists, may impact inbound or outbound routing table updates. Whenever there is a change in the routing policy, the BGP session must be cleared or reset for the new policy to take effect. The two methods for resetting a BGP session are hard reset and soft reset.

Clearing a BGP session using a hard reset invalidates the cache and results in a negative impact on the network performance when the information in the cache becomes unavailable. A hard reset is also disruptive because active BGP sessions are torn down.

A soft reset, which is performed on a per-neighbor basis, does not clear the BGP session and facilitates the application of new policies. There are two methods of performing a soft reset:

- A dynamic inbound soft reset is used to generate inbound updates from a neighbor.
- An outbound soft reset is used to send a new set of updates to a neighbor.

Before the BGP Soft Reset Enhancement feature, a soft reset for inbound routing table updates was performed by entering the **neighbor soft-reconfiguration** router configuration command. This command was used to configure the local BGP router to store all received (inbound) routing policy updates. However, this method uses too much memory because inbound updates are not modified and is not recommended.

Note Outbound resets have never required preconfiguration or storing of routing table updates and remain unchanged by the BGP Soft Reset Enhancement. The procedure for an outbound reset is described in the section "Reset BGP Connections" in the Cisco IOS Release 12.0 *Network Protocols Configuration Guide, Part 1*.

Managing Routing Policy Changes

When the routing policy of a BGP neighbor changes, the session must be reset (cleared) for the changes to take effect. Because resetting a BGP session can be disruptive to networks, a soft reset method is recommended for reconfiguring the routing table.

In order to reconfigure the inbound routing table before the introduction of this feature, both the local BGP router and the BGP peer first needed to be configured to store incoming routing policy updates using the **neighbor soft-reconfiguration** command. Additional resources, particularly memory, were required to store the inbound routing table updates. The **clear ip bgp** command could then initiate the soft reset, which generated a new set of inbound routing table updates using the stored information.

This feature provides an additional method for soft reset that allows the dynamic exchange of route refresh requests and routing information between BGP routers and the subsequent re-advertisement of the respective outbound routing table. Soft reset using the route refresh capability does not require preconfiguration and consumes no additional memory resources.

To use this new method, both BGP peers must support the soft route refresh capability, which is advertised in the OPEN message. Any router that is running BGP with this software release automatically supports the route refresh capability. Routers that are running previous Cisco IOS software releases do not support the route refresh capability and must use the older soft reset method.

If the soft reset fails, you can still clear the BGP session, but it will have a negative impact upon network operations and should only be used as a last resort.

Benefits

Allows Dynamic Route Refresh Requests

This feature provides a way to initiate nondisruptive routing policy changes by allowing the dynamic exchange of route refresh requests between BGP routers and the subsequent re-advertisement of the respective outbound routing tables.

Requires No Preconfiguration

Because support for the soft reset using the route refresh capability is included in this release of the Cisco IOS software, no further router configuration is required. You can initiate a soft inbound reset using only the **clear ip bgp in** command.

Requires No Additional Memory Resources

Unlike a soft reset using the stored inbound routing table updates provided by the **neighbor soft-reconfiguration** command, when both BGP peers support the route refresh capability inbound routing table updates are not stored in the local BGP router. The soft reset requests are exchanged dynamically, and no additional memory is required.

Flexibility

There are now two available methods for inbound soft reset; the older method using stored inbound routing table updates, and the method provided by this feature using dynamic exchange of update information.

Restrictions

Route Refresh Support for BGP Peers

BGP peers must support the route refresh capability to use dynamic inbound soft reset capability. If a peer does not support the route refresh capability, then the only soft reconfiguration option is to use the **neighbor soft-reconfiguration** command, which initiates the storage of inbound routing table updates and requires additional memory.

Dynamic and Stored Inbound Soft Reset Functions are Mutually Exclusive

The dynamic inbound soft reset and inbound soft reset using stored information functions are mutually exclusive and cannot be configured together. If the inbound soft reset using stored routing table updates is configured for a neighbor, the dynamic inbound soft update method cannot be used.

Related Features and Technologies

The BGP Soft Reset feature is an extension of the BGP routing protocol. For more information about configuring BGP, refer to the "BGP chapter" of the *Cisco IOS Release 12.0 Network Protocols Configuration Guide, Part 1* and *Cisco IOS Release 12.0 Network Protocols Command Reference, Part 1*.

Related Documents

- *Cisco IOS Release 12.0 Network Protocols Configuration Guide, Part 1*
- *Cisco IOS Release 12.0 Network Protocols Command Reference, Part 1*

Supported Platforms

The BGP Soft Reset feature is supported by the following platforms in Cisco IOS Release 12.0(22)S:

- Cisco 7200 series
- Cisco 7500 series

- Cisco 10000 series
- Cisco 12000 series

Determining Platform Support Through Cisco Feature Navigator

Cisco IOS software is packaged in feature sets that are supported by 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 2918, *Route Refresh Capability for BGP-4*

Prerequisites

Peers that exchange reachability information must support the soft reset route refresh capability in order to use the enhancements described in this document. If a peer does not, then the only available soft reconfiguration option is to use the **neighbor soft-reconfiguration** command, which initiates the storage of inbound routing table updates and requires additional memory, followed by the **clear ip bgp in** command.

Configuration Tasks

See the following "Configuring BGP Soft Reset" section for configuration tasks for the BGP Soft Reset Enhancement feature. This task is required.

Configuring BGP Soft Reset

Whenever there is a change in the routing policy, the BGP session must be reset (cleared) for the new policy to take effect and the routing table to be reconfigured. Using a hard reset to clear a BGP session causes cache invalidation and results in a negative impact on network operation.

Soft reset is recommended because it allows routing table policies to be reconfigured and activated without clearing the BGP session. Soft reset is done on a per-neighbor basis. Soft resets can be inbound or outbound:

- When soft reset is used to generate inbound updates from a neighbor, it is called dynamic inbound soft reset. There are two ways to perform an inbound soft reset, dynamically (using the dynamic inbound soft reset) and using stored routing update information.
- When soft reset is used to send a new set of updates to a neighbor, it is called outbound soft reset.

  There is only one way to perform an outbound soft reset.

Note The dynamic inbound soft reset and inbound soft reset using stored information functions are mutually exclusive and cannot be configured together. If the inbound soft reset using stored routing table updates is configured for a neighbor, the dynamic inbound soft update method cannot be used.

Configuring BGP Dynamic Inbound Soft Reset

If both the local BGP router and the neighbor router support the route refresh capability, you can perform a dynamic soft inbound reset. This type of reset has the following advantages over a soft inbound reset using stored routing update information:

- Does not require preconfiguration
- Does not require additional memory for storing routing update information

To determine whether a router supports the route refresh capability, use the **show ip bgp neighbors** command.

Command	Purpose
`Router# show ip bgp neighbors ip-address`	Shows whether a neighbor supports the route refresh capability. If the specified router supports the route refresh capability, the following message is displayed: Received route refresh capability from peer.

If all the BGP routers support the route refresh capability, you can use the dynamic soft reset method for resetting the inbound routing table.

Command	Purpose
`Router# clear ip bgp * | ip-address | peer-group-name soft in`	Performs a dynamic soft reset on the connection specified in the command.

Configuring BGP Outbound Soft Reset

Outbound soft resets do not require any preconfiguration. Using the keyword **soft** specifies that a soft reset be performed.

Command	Purpose
`Router# clear ip bgp * | address | peer-group-name soft out`	Performs a soft reset on the connection specified in the command.

Configuring BGP Outbound Soft Reset Under an Address Family

Outbound soft resets do not require any preconfiguration. Using the keyword soft specifies that a soft reset be performed.

Command	Purpose
`Router# clear ip bgp * | address | peer-group-name ipv4 | vpnv4 unicast | multicast in | out | soft`	Performs a soft reset on the connection and address family specified in the command. The **in** and **out** keywords do not follow the **soft** keyword when configured for an address family.

Configuring BGP Soft Reset Using Stored Routing Policy Information

If all of the BGP routers in the connection do not support the route refresh capability, use the soft reset method that generates a new set of inbound routing table updates from information previously stored. To initiate storage of inbound routing table updates, you must first preconfigure the router using the **neighbor soft-reconfiguration** command.

Keep in mind that the memory requirements for storing the inbound update information can become quite large.

Command	Purpose
`Router(config-router)# neighbor` *ip-address* | *peer-group-name* `soft-reconfiguration inbound`	Initiates storage of inbound routing table updates from the specified neighbor or peer group.

Once you have initiated storage of inbound routing table updates for a specific neighbor or peer group, you can perform a soft inbound reset for that neighbor or peer group.

Command	Purpose
`Router# clear ip bgp *` | *ip-address* | *peer-group-name* `soft in`	Performs a soft reset on the connection specified in the command.

Verifying BGP Soft Reset

Enter the **show ip bgp neighbors** command to display information about the BGP and TCP connections to neighbors and verify the status and configuration of the BGP soft reset feature. The following sample output shows that a soft reset has been configured for neighbor 10.4.9.8:

```
Router# show ip bgp neighbors
BGP neighbor is 10.4.9.8,  remote AS 101, internal link
  BGP version 4, remote router ID 10.4.9.8
  BGP state = Established, up for 00:03:50
  Last read 00:00:50, hold time is 180, keepalive interval is 60 seconds
  Neighbor capabilities:
    Route refresh: advertised and received(new)
    Address family IPv4 Unicast: advertised and received
  Received 7 messages, 0 notifications, 0 in queue
  Sent 7 messages, 0 notifications, 0 in queue
  Default minimum time between advertisement runs is 5 seconds

 For address family: IPv4 Unicast
  BGP table version 1, neighbor version 1
  Index 1, Offset 0, Mask 0x2
  Route refresh request: received 0, sent 0
  0 accepted prefixes consume 0 bytes
  Prefix advertised 0, suppressed 0, withdrawn 0
  Number of NLRIs in the update sent: max 0, min 0

  Connections established 1; dropped 0
  Last reset never
Connection state is ESTAB, I/O status: 1, unread input bytes: 0
Local host: 10.4.9.21, Local port: 179
Foreign host: 10.4.9.8, Foreign port: 11004

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

Event Timers (current time is 0x134F4D2B4):
Timer          Starts    Wakeups            Next
Retrans             8          0             0x0
TimeWait            0          0             0x0
AckHold             7          3             0x0
SendWnd             0          0             0x0
KeepAlive           0          0             0x0
GiveUp              0          0             0x0
PmtuAger            0          0             0x0
```

```
DeadWait            0          0             0x0

iss: 4229692689  snduna: 4229692849  sndnxt: 4229692849     sndwnd:  16225
irs:  339739239  rcvnxt:  339739399  rcvwnd:      16225  delrcvwnd:    159

SRTT: 540 ms, RTTO: 3809 ms, RTV: 1364 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: 10 (out of order: 0), with data: 7, total data bytes: 159
Sent: 11 (retransmit: 0, fastretransmit: 0), with data: 7, total data bytes: 159
```

Configuration Examples

This section provides the following configuration examples:

- Dynamic Inbound Soft Reset
- Inbound Soft Reset Using Stored Information

Dynamic Inbound Soft Reset

The following examples shows the **clear ip bgp 131.108.1.1 soft in** command used to initiate a dynamic soft reconfiguration in the BGP peer 131.108.1.1. This command requires that the peer supports the route refresh capability.

```
clear ip bgp 131.108.1.1 soft in
```

Inbound Soft Reset Using Stored Information

The following example enables inbound soft reconfiguration for the neighbor 131.108.1.1. All the updates received from this neighbor will be stored unmodified, regardless of the inbound policy. When inbound soft reconfiguration is performed later, the stored information will be used to generate a new set of inbound updates.

```
router bgp 100
 neighbor 131.108.1.1 remote-as 200
 neighbor 131.108.1.1 soft-reconfiguration inbound
```

The following example clears the session with the neighbor 131.108.1.1.

```
clear ip bgp 131.108.1.1 soft in
```

Command Reference

This section documents new or modified commands. All other commands used with this feature are documented in the Cisco IOS Release 12.0 command reference publications.

- **clear ip bgp**
- **neighbor soft-reconfiguration**
- **show ip bgp neighbors.**

clear ip bgp

To reset a BGP connection using BGP soft reconfiguration, use the **clear ip bgp** privileged EXEC command at the system prompt.

clear ip bgp {***** | *neighbor-address* | *peer-group-name* [**soft** [**in** | **out**]} [**ipv4 {multicast** | **unicast}** | **vpnv4** {**unicast**} {**soft** | **in** | **out**}]

Syntax Description

*****	Resets all current BGP sessions.
neighbor-address	Resets only the identified BGP neighbor.
peer-group-name	Resets the specified BGP peer group.
ipv4	(Optional) Resets the specified IPv4 address family neighbor or peer group. The **multicast** or **unicast** keyword must be specified.
vpnv4	(Optional) Resets the specified VPNv4 address family neighbor or peer group. The **unicast** keyword must be specified.
soft	(Optional) Soft reset. Does not reset the session. The **in** or **out** keywords do not follow the **soft** keyword when a connection is cleared under the VPNv4 or IPv4 address family because **soft** keyword specifies both.
in | **out**	(Optional) Triggers inbound or outbound soft reconfiguration. If the **in** or **out** option is not specified, both inbound and outbound soft reset are triggered.

Defaults

A session reset is initiated by default.

Command Modes

Privileged EXEC

Command History

Release	Modification
10.0	This command was introduced.
12.0(6)T	The dynamic inbound soft reset capability was added.
12.0(2)S	The dynamic inbound soft reset capability was added.
12.0(22)S	The **vpnv4** and **ipv4** keywords were added.

Usage Guidelines

You can reset inbound routing table updates dynamically or by generating new updates using stored update information. Using stored update information required additional memory for storing the updates.

To reset inbound routing table updates dynamically, all BGP routers must support the route refresh capability. To determine whether a BGP router supports this capability, use the **show ip bgp neighbors** command. If a router supports the route refresh capability, the following message is displayed:

```
Received route refresh capability from peer.
```

If all BGP routers support the route refresh capability, use the **clear ip bgp** {***** | *address* | *peer-group -name*} **in** command. You need not use the **soft** keyword, because soft reset is automatically assumed when the route refresh capability is supported.

The route refresh capability can also be initiated specifically for IPv4 and VPNv4 address families. When an soft refresh is initiated, the **in** or **out** keyword does not need to be specified because **soft** keyword specifies both.

To generate new inbound updates from stored update information (rather than dynamically) without resetting the BGP session, you must preconfigure the local BGP router using the **neighbor soft-reconfiguration inbound** command. This preconfiguration causes the software to store all received updates without modification regardless of whether an update is accepted by the inbound policy. Storing updates is memory intensive and should be avoided if possible.

Outbound BGP soft configuration has no memory overhead and does not require any preconfiguration. You can trigger an outbound reconfiguration on the other side of the BGP session to make the new inbound policy take effect.

Use this command whenever any of the following changes occur:

- Additions or changes to the BGP-related access lists
- Changes to BGP-related weights
- Changes to BGP-related distribution lists
- Changes to BGP-related route maps

Examples

The following example clears the inbound session with the neighbor 131.108.1.1 without resetting the session:

```
Router# clear ip bgp 131.108.1.1 soft in
```

The following example clears the outbound session with the peer group named corp without resetting the session:

```
Router# clear ip bgp corp soft out
```

The following example clears both the inbound and outbound VPNv4 session without resetting the session for 10.0.0.1 neighbor:

```
Router# clear ip bgp vpnv4 unicast 10.0.0.1 soft
```

The following example clears the outbound IPv4 multicast session without resetting the session for 192.168.0.1 neighbor:

```
Router# clear ip bgp ipv4 multicast 192.168.0.1 out
```

The following example clears the inbound IPv4 unicast session without resetting the session for 172.16.0.1 neighbor:

```
Router# clear ip bgp ipv4 multicast 172.16.0.1 in
```

Related Commands

Command	Description
neighbor soft-reconfiguration	Configures the Cisco IOS software to start storing updates.
show ip bgp neighbors	Displays entries in the BGP routing table.

neighbor soft-reconfiguration

To configure the Cisco IOS software to start storing updates, use the **neighbor soft-reconfiguration** router configuration command. To not store received updates, use the **no** form of this command.

neighbor {*ip-address* | *peer-group-name*} **soft-reconfiguration** [**inbound**]

no neighbor {*ip-address* | *peer-group-name*} **soft-reconfiguration** [**inbound**]

Syntax Description

ip-address	IP address of the BGP-speaking neighbor.
peer-group-name	Name of a BGP peer group.
inbound	Keyword indicating that the update to be stored is an incoming update. Inbound is currently required with this command, since a keyword is required and no other keywords are available.

Defaults

Soft reconfiguration is not enabled

Command Modes

Router configuration

Command History

Release	Modification
11.2	This command was introduced.

Usage Guidelines

This command requires at least one keyword. Currently the only keyword available is **inbound**, so the use of **inbound** is not optional.

Entering this command starts the storage of updates, required to do inbound soft reconfiguration. Outbound BGP soft reconfiguration does not require inbound soft reconfiguration to be enabled.

If you specify a BGP peer group by using the *peer-group-name* argument, all the members of the peer group will inherit the characteristic configured with this command.

Examples

The following example enables inbound soft-reconfiguration for the neighbor 131.108.1.1. All the updates received from this neighbor will be stored unmodified, regardless of the inbound policy. When inbound soft reconfiguration is done later, the stored information will be used to generate a new set of inbound updates.

```
router bgp 100
 neighbor 131.108.1.1 remote-as 200
 neighbor 131.108.1.1 soft-reconfiguration inbound
```

Related Commands

Command	Description
clear ip bgp	Resets a BGP connection using BGP soft reconfiguration.
neighbor peer-group (creating)	Creates a BGP peer group.

show ip bgp neighbors

To display information about the TCP and BGP connections to neighbors, use the **show ip bgp neighbors** EXEC command.

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

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.

Command Modes

EXEC

Command History

Release	Modification
10.0	This command was introduced.
11.2	The **received-routes** keyword was added.

Examples

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

```
Router# show ip bgp neighbors
BGP neighbor is 10.4.9.8,  remote AS 101, internal link
  BGP version 4, remote router ID 10.4.9.8
  BGP state = Established, up for 00:03:50
  Last read 00:00:50, hold time is 180, keepalive interval is 60 seconds
  Neighbor capabilities:
    Route refresh: advertised and received(new)
    Address family IPv4 Unicast: advertised and received
  Received 7 messages, 0 notifications, 0 in queue
  Sent 7 messages, 0 notifications, 0 in queue
  Default minimum time between advertisement runs is 5 seconds

 For address family: IPv4 Unicast
  BGP table version 1, neighbor version 1
  Index 1, Offset 0, Mask 0x2
  Route refresh request: received 0, sent 0
  0 accepted prefixes consume 0 bytes
  Prefix advertised 0, suppressed 0, withdrawn 0
  Number of NLRIs in the update sent: max 0, min 0

  Connections established 1; dropped 0
```

```
  Last reset never
Connection state is ESTAB, I/O status: 1, unread input bytes: 0
Local host: 10.4.9.21, Local port: 179
Foreign host: 10.4.9.8, Foreign port: 11004

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

Event Timers (current time is 0x134F4D2B4):
Timer          Starts    Wakeups            Next
Retrans             8          0             0x0
TimeWait            0          0             0x0
AckHold             7          3             0x0
SendWnd             0          0             0x0
KeepAlive           0          0             0x0
GiveUp              0          0             0x0
PmtuAger            0          0             0x0
DeadWait            0          0             0x0

iss: 4229692689  snduna: 4229692849  sndnxt: 4229692849     sndwnd:  16225
irs:  339739239  rcvnxt:  339739399  rcvwnd:      16225  delrcvwnd:    159

SRTT: 540 ms, RTTO: 3809 ms, RTV: 1364 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: 10 (out of order: 0), with data: 7, total data bytes: 159
Sent: 11 (retransmit: 0, fastretransmit: 0), with data: 7, total data bytes: 159
```

Table 1 describes the significant fields shown in the display.

Table 1 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 external BGP (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 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 between sending keepalive packets, which help ensure that the TCP connection is up.
Neighbor capabilities	BGP capabilities advertised and received from this neighbor.
Route refresh	Indicates that the neighbor supports dynamic soft reset using the route refresh capability.
Address family IP Version 4 Unicast:	IP Version 4 unicast-specific properties of this neighbor.

Table 1 show ip bgp neighbors Field Descriptions (continued)

Field	Description
Address family IP Version 4 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.
mul-in	Name of inbound route map for the multicast address family.
mul-out	Name of outbound route map for the multicast address family.
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.

Table 1 show ip bgp neighbors Field Descriptions (continued)

Field	Description
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.
delrecvwnd	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 resent.
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 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 ?
*> 20.0.0.0         172.16.232.178         40             0 10 ?
```

Table 2 describes the significant fields shown in the displays.

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

Field	Description
BGP table version	Internal version number of the table. This number is incremented whenever 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. *—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 interautonomous 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 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	The Multi Exit Discriminator (MED) metric for the path. (The name of this metric for BGP versions 2 and 3 is INTER_AS.)
Path	The autonomous system path for that route, followed by the origin code for that route.