BGP Prefix-Based Outbound Route Filtering

The BGP Prefix-Based Outbound Route Filtering feature uses Border Gateway Protocol (BGP) outbound route filter (ORF) send and receive capabilities to minimize the number of BGP updates that are sent between BGP peers. Configuring this feature can help reduce the amount of system resources required for generating and processing routing updates by filtering out unwanted routing updates at the source. For example, this feature can be used to reduce the amount of processing required on a router that is not accepting full routes from a service provider network.

Feature History for BGP Prefix-Based Outbound Route Filtering

Release	Modification
12.0(11)ST	This feature was introduced.
12.2(4)T	This feature was integrated into Cisco IOS Release 12.2(4)T.
12.2(14)S	This feature was integrated into Cisco IOS Release 12.2(14)S.
12.0(22)S	This feature was integrated into Cisco IOS Release 12.0(22)S.

Finding Support Information for Platforms and Cisco IOS Software Images

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

CISCO SYSTEMS

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

Contents

Prerequisites for BGP Prefix-Based Outbound Route Filtering

- BGP peering sessions must be established, and BGP ORF capabilities must be enabled on each participating router before prefix-based ORF announcements can be sent or received.

Restrictions for BGP Prefix-Based Outbound Route Filtering

- The BGP Prefix-Based Outbound Route Filtering feature does not support IP multicast routes.
- IP addresses that are used for outbound route filtering must be defined in an IP prefix-list. BGP distribute lists and IP access lists are not supported.
- Outbound route filtering is configured on only a per-address family basis and cannot be configured under the general session or BGP routing process (`Router(config-router)#`).
- Outbound route filtering is configured for only external peering sessions.

Information About BGP Prefix-Based Outbound Route Filtering

To configure BGP prefix-based outbound route filtering, you must understand the following concepts:

BGP Prefix-Based Outbound Route Filtering Overview

The BGP Prefix-Based Outbound Route Filtering feature uses Border Gateway Protocol (BGP) outbound route filter (ORF) send and receive capabilities to minimize the number of BGP updates that are sent between BGP peers. Configuring this feature can help reduce the amount of system resources required for generating and processing routing updates by filtering out unwanted routing updates at the source. For example, this feature can be used to reduce the amount of processing required on a router that is not accepting full routes from a service provider network.

Enabling BGP Prefix-Based Outbound Route Filtering

The BGP Prefix-Based Outbound Route Filtering feature is enabled through the advertisement of ORF capabilities to peer routers. The advertisement of the ORF capability indicates that a BGP speaker will accept a prefix list from a neighbor and apply the prefix list to locally configured ORFs (if any exist). When this capability is enabled, the BGP speaker can install the inbound prefix list filter to the remote peer as an outbound filter, which reduces unwanted routing updates.

The BGP Prefix-Based Outbound Route Filtering feature can be configured with send and/or receive ORF capabilities. The local peer advertises the ORF capability in send mode. The remote peer receives the ORF capability in receive mode and applies the filter as an outbound policy. The local and remote peers exchange updates to maintain the ORF on each router. Updates are exchanged between peer routers by address family depending on the ORF prefix list capability that is advertised. The remote peer starts sending updates to the local peer after a route refresh has been configured with the **clear ip bgp** command or after an ORF prefix list with immediate status is processed. The BGP speaker will continue to apply the inbound prefix list to received updates after the speaker pushes the inbound prefix list to the remote peer.

Benefits of BGP Prefix-Based Outbound Route Filtering

The BGP Prefix-Based Outbound Route Filtering feature can limit the number of unwanted routing updates, which will reduce the amount of resources required for routing update generation and processing. This feature also reduces the amount of resources required to receive and discard routes that would otherwise be filtered out.

How to Configure BGP Prefix-Based Outbound Route Filtering

This section contains the following procedures:

Configuring BGP Prefix-Based Outbound Route Filtering

The BGP Prefix-Based Outbound Route Filtering supports prefix length matching, wild-card based prefix matching, and exact address prefix matching for address family support. This feature can be configured on a router to send or receive ORF capabilities with either the **send** or **receive** keywords. This feature can also be configured on a router to both send and receive ORF capabilities with the **both** keyword.

Prerequisites

- BGP peering sessions must be established, and BGP ORF capabilities must be enabled on each participating router before prefix-based ORF announcements can be received.

Restrictions

- The BGP Prefix-Based Outbound Route Filtering feature does not support multicast.

- IP addresses that are used for outbound route filtering must be defined in an IP prefix-list. BGP distribute lists and IP access lists are not supported.
- Outbound route filtering is configured on only a per-address family basis and cannot be configured under the general session or BGP routing process (`Router(config-router)#`).
- Outbound route filtering is configured for only external peering sessions.

SUMMARY STEPS

1. **enable**
2. **configure terminal**
3. **ip prefix-list** *list-name* [**seq** *seq-value*] {**deny** *network/length* | **permit** *network/length*} [**ge** *ge-value*] [**le** *le-value*]
4. **router bgp** *as-number*
5. **address-family** {**ipv4** | **ipv6** | **vpnv4**| [**multicast** | **unicast** | **vrf** {*vrf-name*}]
6. **neighbor** *ip-address* **remote-as** *as-number*
7. **neighbor** *ip-address* **ebgp-multihop** [**hop-count**]
8. **neighbor** *ip-address* **capability orf prefix-list** [**send** | **receive** | **both**]
9. **neighbor** {*ip-address* | *peer-group-name*} **prefix-list** *prefix-list-name* {**in** | **out**}
10. **end**
11. **clear ip bgp** {*ip-address* | *****} **in prefix-filter**

DETAILED STEPS

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

	Command or Action	Purpose
Step 3	**ip prefix-list list-name** [**seq seq-value**] {**deny network/length** | **permit network/length**} [**ge ge-value**] [**le le-value**] **Example:** `Router(config)# ip prefix-list FILTER seq 10 permit 192.168.1.0/24`	Creates a prefix list for prefix based outbound route filtering. • Outbound route filtering supports prefix length matching, wild-card based prefix matching, and exact address prefix matching on a per address-family basis. • The prefix list is created to define the outbound route filter. The filter must be created when the outbound route filtering capability is configured to be advertised in **send** mode or **both** mode. It is not required when a peer is configured to advertise **receive** mode only. • The example creates a prefix list named FILTER that defines the 192.168.1.0/24 subnet for outbound route filtering.
Step 4	**router bgp** *as-number* **Example:** `Router(config)# router bgp 100`	Enters router configuration mode, and creates a BGP routing process.
Step 5	**address-family ipv4** | **ipv6** |**vpnv4** [**multicast** | **unicast** | **vrf** {*vrf-name*}] **Example:** `Router(config-router)# address-family ipv4 unicast`	Enter address family configuration mode to configure BGP peers to accept address family specific configurations. • Outbound route filtering is configured on a per-address family basis. • The example creates an IPv4 unicast address family session.
Step 6	**neighbor** *ip-address* **remote-as** *as-number* **Example:** `Router(config-router-af)# neighbor 10.1.1.1 remote-as 200`	Establishes peering with the specified neighbor or peer-group. • BGP peering must be established before ORF capabilities can be exchanged. • The example establishes peering with the 10.1.1.1 neighbor.
Step 7	**neighbor** *ip-address* **ebgp-multihop** [*hop-count*] **Example:** `Router(config-router-af)# neighbor 10.1.1.1 ebgp-multihop`	Accepts or initiates BGP connections to external peers residing on networks that are not directly connected.

	Command or Action	Purpose
Step 8	**neighbor** *ip-address* **capability orf prefix-list** [**send** | **receive** | **both**] **Example:** `Router(config-router-af)# neighbor 10.1.1.1 capability orf prefix-list both`	Enables the ORF capability on the local router, and enables ORF capability advertisement to the BGP peer specified with the *ip-address* argument. • The **send** keyword configures a router to advertise ORF send capabilities. • The **receive** keyword configures a router to advertise ORF receive capabilities. • The **both** keyword configures a router to advertise send and receive capabilities. • The remote peer must be configured to either send or receive ORF capabilities before outbound route filtering is enabled. • The example configures the router to advertise send and receive capabilities to the 10.1.1.1 neighbor.
Step 9	**neighbor** {*ip-address* | *peer-group-name*} **prefix-list** *prefix-list-name* {**in** | **out**} **Example:** `Router(config-router-af)# neighbor 10.1.1.1 prefix-list FILTER in`	Applies an inbound prefix-list filter to prevent distribution of BGP neighbor information. • In this example, the prefix list named FILTER is applied to incoming advertisements from the 10.1.1.1 neighbor, which prevents distribution of the 192.168.1.0/24 subnet.
Step 10	**end** **Example:** `Router(config-router-af)# exit`	Exits address-family configuration mode, and enters privileged EXEC mode.
Step 11	**clear ip bgp** {*ip-address* | *****} **in prefix-filter** **Example:** `Router# clear ip bgp 10.1.1.1 in prefix-filter`	Clears BGP outbound route filters and initiates an inbound soft reset. A single neighbor or all neighbors can be specified. • When this command is used without the prefix-filter keyword, a normal route refresh is performed. This command should be used when inbound routing policy changes other than a prefix list filter occur, such as a route map change.

Examples

Router-A Configuration (Sender)

The following example creates an outbound route filter and configures Router-A (10.1.1.1) to advertise the filter to Router-B (172.16.1.2). An IP prefix list named FILTER is created to specify the 192.168.1.0/24 subnet for outbound route filtering. The ORF send capability is configured on Router-A so that Router-A can advertise the outbound route filter to Router-B.

```
ip prefix-list FILTER seq 10 permit 192.168.1.0/24
!
router bgp 100
 address-family ipv4 unicast
 neighbor 172.16.1.2 remote-as 200
 neighbor 172.16.1.2 ebgp-multihop
 neighbor 172.16.1.2 capability orf prefix-list send
```

```
 neighbor 172.16.1.2 prefix-list FILTER in
 exit
```

Router-B Configuration (Receiver)

The following example configures Router-B to advertise the ORF receive capability to Router-A. Router-B will install the outbound route filter, defined in the FILTER prefix list, after ORF capabilities have been exchanged. An inbound soft reset is initiated on Router-B at the end of this configuration to activate the outbound route filter.

```
router bgp 200
 address-family ipv4 unicast
 neighbor 10.1.1.1 remote-as 100
 neighbor 10.1.1.1 ebgp-multihop 255
 neighbor 10.1.1.1 capability orf prefix-list receive
 end
clear ip bgp 10.1.1.1 in prefix-filter
!
```

Note The inbound soft refresh must be initiated with the **clear ip bgp** command in order for this feature to function.

What to Do Next

You can verify that this feature is configured correctly with the **show running-config** and **show ip bgp neighbors** command. Go to the Verifying BGP Prefix-Based Outbound Route Filtering section for instructions and example output.

Verifying BGP Prefix-Based Outbound Route Filtering

This section provides verification examples for the BGP Prefix-Based Outbound Route Filtering feature. In each example, the relevant lines of the output are shown in bold text.

Verifying the Local BGP Prefix-Based Outbound Route Filtering Configuration

The following example output from the **show running-config** privileged EXEC command shows the local configuration of a router. The router is configured to advertise the ORF send mode capability to the 172.16.1.2 neighbor.

```
Router# show running-config | begin bgp

router bgp 100
 bgp log-neighbor-changes
 neighbor 172.16.1.2 remote-as 200
 neighbor 172.16.1.2 ebgp-multihop 255
 !
 address-family ipv4
 neighbor 172.16.1.2 activate
 neighbor 172.16.1.2 capability orf prefix-list send
 neighbor 172.16.1.2 prefix-list FILTER in
 no auto-summary
```

```
 no synchronization
 exit-address-family
!
```

Verifying the BGP Prefix-Based Outbound Route Filtering Configuration of a Remote Peer

The following is sample output from the **show ip bgp neighbors** privileged EXEC command. The output show if ORF capability exchange is enabled on the local router and the specified peer. Send mode and receive mode capabilities are shown as advertised or received or as both advertised and received.

The output shows that 10.1.1.1 neighbor is advertising the send mode ORF capability and has received an announcement that the local router is advertising both send and receive mode capabilities:

```
Router# show ip bgp neighbors 10.1.1.1

BGP neighbor is 10.1.1.1,  remote AS 100, external link
  BGP version 4, remote router ID 172.16.1.2
  BGP state = Established, up for 00:13:27
  Last read 00:00:27, 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 36 messages, 0 notifications, 0 in queue
  Sent 36 messages, 0 notifications, 0 in queue
  Default minimum time between advertisement runs is 30 seconds

 For address family:IPv4 Unicast
  BGP table version 13, neighbor version 13
  Index 1, Offset 0, Mask 0x2
  AF-dependant capabilities:
    Outbound Route Filter (ORF) type (128) Prefix-list:
      Send-mode:advertised, received
      Receive-mode:received
  Route refresh request:received 4, sent 2
  2 accepted prefixes consume 80 bytes
  Prefix advertised 8, suppressed 0, withdrawn 2
```

Verifying the Receipt of the Outbound Filter by the Remote Peer

The following is sample output from the **show ip bgp neighbors** command entered with the **received prefix-filter** keyword. The IP prefix list that defines the outbound route filter will be displayed as a prefix list entry on the specified router. The output shows that the 172.16.1.2 neighbor has received an IP prefix list that defines an outbound route filter for the 192.168.1.0/24 subnet.

```
Router# show ip bgp neighbor 172.16.1.2 received prefix-filter

Address family:IPv4 Unicast
ip prefix-list 172.16.1.2:1 entries
   seq 5 permit 192.168.1.0/24
```

Additional References

The following sections provide references related to BGP Prefix-Based Outbound Route Filtering feature.

Related Documents

Related Topic	Document Title
The BGP Prefix-Based Outbound Route Filtering feature is an extension of the BGP routing protocol. For more information about configuring BGP, route filtering, and IP prefix lists refer to the "Configuring BGP" chapter of the Cisco IOS IP Configuration Guide and the Cisco IOS IP Command Reference.	• *Cisco IOS IP Configuration Guide, Release 12.3* • *Cisco IOS IP Command Reference, Volume 2 of 4: Routing Protocols, Release 12.3*

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.	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

RFCs	Title
RFC 5292	*Address Prefix Based Outbound Route Filter for BGP-4*
draft-ietf-idr-route-filter-10.txt	*Cooperative Route Filtering Capability for BGP-4*

Technical Assistance

Description	Link
Technical Assistance Center (TAC) home page, containing 30,000 pages of searchable technical content, including links to products, technologies, solutions, technical tips, and tools. Registered Cisco.com users can log in from this page to access even more content.	TAC Home Page: http://www.cisco.com/public/support/tac/home.shtml BGP Support Page: http://www.cisco.com/cgi-bin/Support/browse/psp_view.pl?p=Internetworking:BGP

Command Reference

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

New Commands

- **clear ip bgp in prefix-filter**
- **neighbor capability orf prefix-list**

Modified Commands

- **show ip bgp neighbors**

clear ip bgp in prefix-filter

To initiate an inbound soft reset to clear a BGP outbound route filter (ORF), use the **clear ip bgp in prefix-filter** command in privileged EXEC mode.

clear ip bgp {*ip-address* | *} **in prefix filter**

Syntax Description

*	Clears all ORFs and resets all inbound BGP sessions.
ip-address	(Optional) IP address of the network about which to clear dampening information.

Defaults

The **prefix-filter** keyword will be ignored and a normal inbound route refresh will be performed if ORF capabilities have not been enabled locally or received from a BGP peer.

Command Modes

Privileged EXEC

Command History

Release	Modification
12.0(11)ST	This command was introduced.
12.2(4)T	This command was integrated into Cisco IOS Release 12.2(4)T.
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.

Usage Guidelines

This command is used to push out the existing ORF prefix list so that a new route refresh will be received from a neighbor (including the current ORF prefix list). When this command is used without the **prefix-filter** keyword, a normal route refresh is performed. This command should be used when inbound routing policy changes other than a prefix list filter occur, such as a route map change.

Examples

The following example initiates an inbound soft reset to clear BGP ORFs received from the 192.168.0.1 neighbor:

```
Router# clear ip bgp 192.168.0.1 in prefix-filter
```

Related Commands

Command	Description
neighbor capability orf prefix-list	Enables outbound route filter (ORF) capability exchange and advertise ORF capabilities to a BGP peer.
show ip bgp neighbors	Displays information about the TCP and BGP connections to neighbors.

neighbor capability orf prefix-list

To enable outbound route filter (ORF) capability exchange and advertise ORF capabilities to a BGP peer, use the **neighbor capability orf prefix-list** command in address family configuration mode. To disable ORF capability exchange, use the **no** form of this command.

neighbor *ip-address* **capability orf prefix-list** {**both** | **receive** | **send**}

no neighbor *ip-address* **capability orf prefix-list** {**both** | **receive** | **send**}

Syntax Description

ip-address	The IP address of the neighbor router.
capability	(optional) Informs the specified neighbor that local router has ORF capabilities.
both	(optional) Configures the router to advertise both send and receive mode ORF capabilities.
receive	(optional) Configures the router to advertise receive mode ORF capabilities.
send	(optional) Configures the router to advertise send mode capabilities.

Defaults

Disabled

Command Modes

Address family configuration

Command History

Release	Modification
12.0(11)ST	This command was introduced.
12.2(4)T	This command was integrated into Cisco IOS Release 12.2(4)T.
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.

Usage Guidelines

The **neighbor capability orf prefix** command is used to advertise ORF send and/or receive capabilities to minimize the number of BGP updates that are processed and sent between BGP peers. The outbound route filter is defined in an IP prefix list (distribute lists and IP access lists are not supported). The ORF capability is enabled both locally and on the remote peer. The **send** keyword configures a router to advertise the ORF send capability, the **receive** keyword configures a router to advertise the ORF receive capabilities, and the **both** keyword configures a router to advertise both send and receive capabilities.

In most configurations, this command will be used to advertise both send and receive ORF capabilities. However this feature can be configured in one direction between two routers with one router configured to send ORF capabilities and another router configured to receive ORF capabilities.

Examples

Send Mode Configuration Example

The following example creates an outbound route filter and configures Router-A (10.1.1.1) to advertise the filter to Router-B (172.16.1.2). An IP prefix list named FILTER is created to specify the 192.168.1.0/24 subnet for outbound route filtering. The ORF send capability is configured on Router-A so that Router-A will advertise the outbound route filter to Router-B.

```
ip prefix-list FILTER seq 10 permit 192.168.1.0/24
!
router bgp 100
 address-family ipv4 unicast
 neighbor 172.16.1.2 remote-as 200
 neighbor 172.16.1.2 ebgp-multihop
 neighbor 172.16.1.2 capability orf prefix-list send
 neighbor 172.16.1.2 prefix-list FILTER in
 exit
```

Receive Mode Configuration Example

The following example configures Router-B to advertise the ORF receive capability to Router-A. Router-B will install the outbound route filter, defined in the FILTER prefix list, after ORF capabilities have been exchanged. An inbound soft reset is initiated on Router-B at the end of this configuration to activate the outbound route filter.

```
router bgp 200
 address-family ipv4 unicast
 neighbor 10.1.1.1 remote-as 100
 neighbor 10.1.1.1 ebgp-multihop 255
 neighbor 10.1.1.1 capability orf prefix-list receive
 end
clear ip bgp 10.1.1.1 in prefix-filter
!
```

Related Commands

Command	Description
clear ip bgp in prefix-filter	Initiates an inbound soft reset to clear BGP outbound route filters.
ip prefix-list	Creates an entry in a prefix list.

show ip bgp neighbors

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

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

Syntax Description

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

Command Modes

EXEC

Command History

Release	Modification
10.0	This command was introduced.
11.2	The **received-routes** keyword was added.
12.2(4)T	The **received prefix-filter** keyword was added.
12.2(14S	
12.0(22)S	

Examples

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

```
Router# show ip bgp neighbors 172.16.232.178

BGP neighbor is 172.16.232.178,  remote AS 35, external link
  BGP version 4, remote router ID 192.168.3.3
  BGP state = Established, up for 1w1d
  Last read 00:00:53, hold time is 180, keepalive interval is 60 seconds
  Neighbor capabilities:
    Route refresh: advertised and received
    Address family IPv4 Unicast: advertised and received
    Address family IPv4 Multicast: advertised and received
  Received 12519 messages, 0 notifications, 0 in queue
  Sent 12523 messages, 0 notifications, 0 in queue
  Route refresh request: received 0, sent 0
  Minimum time between advertisement runs is 30 seconds
```

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

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

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

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

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

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

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

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

Table 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, in seconds, that the underlying TCP connection has been in existence.
Last read	Time that BGP last read a message from this neighbor.
hold time	Maximum amount of time that can elapse between messages from the peer.
keepalive interval	Time period, in seconds, between sending keepalive packets, which help ensure that the TCP connection is up.
Neighbor capabilities	BGP capabilities advertised and received from this neighbor.
Route refresh	Indicates that the neighbor supports dynamic soft reset using the route refresh capability.
Address family IPv4 Unicast:	IP Version 4 unicast-specific properties of this neighbor.
Address family IPv4 Multicast:	IP Version 4 multicast-specific properties of this neighbor.
Received	Number of total BGP messages received from this peer, including keepalives.
notifications	Number of error messages received from the peer.
Sent	Total number of BGP messages that have been sent to this peer, including keepalives.
notifications	Number of error messages the router has sent to this peer.
Route refresh request:	Number of route refresh requests sent and received from this neighbor.
advertisement runs	Value of the minimum advertisement interval.
For address family:	Address family to which the following fields refer.
BGP table version	Indicates that the neighbor has been updated with this version of the primary BGP routing table.
neighbor version	Number used by the software to track the prefixes that have been sent and those that must be sent to this neighbor.
Community attribute	Appears if the **neighbor send-community** command is configured for this neighbor.
Inbound path policy	Indicates if an inbound policy is configured.
Outbound path policy	Indicates if an outbound policy is configured.

Table 1 show ip bgp neighbors Field Descriptions (continued)

Field	Description
mul-in	Name of the inbound route map for the multicast address family.
mul-out	Name of the 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, in seconds, since this peering session was last reset.
Connection state	State of the 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 that 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 for which it has not received an acknowledgment.
sndnxt	Sequence number the local host will send next.
sndwnd	TCP window size of the remote host.
irs	Initial receive sequence number.
rcvnxt	Last receive sequence number the local host has acknowledged.
rcvwnd	TCP window size of the local host.
delrcvwnd	Delayed receive window—data the local host has read from the connection, but has not yet subtracted from the receive window the host has advertised to the remote host. The value in this field gradually increases until it is larger than a full-sized packet, at which point it is applied to the rcvwnd field.
SRTT	A calculated smoothed round-trip timeout.
RTTO	Round-trip timeout.
RTV	Variance of the round-trip time.
KRTT	New round-trip timeout (using the Karn algorithm). This field separately tracks the round-trip time of packets that have been re-sent.
minRTT	Smallest recorded round-trip timeout (hard wire value used for calculation).
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.

Table 1 show ip bgp neighbors Field Descriptions (continued)

Field	Description
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.

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

```
Router# show ip bgp neighbor 192.168.20.72 received prefix-filter
Address family:IPv4 Unicast
ip prefix-list 192.168.20.72:1 entries
   seq 5 deny 10.0.0.0/8 le 32
```

Table 4 describes the significant fields shown in the display.

Table 4 show ip bgp neighbors paths Field Descriptions

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