BGP Support for TTL Security Check

The BGP Support for TTL Security Check feature introduces a lightweight security mechanism to protect external Border Gateway Protocol (eBGP) peering sessions from CPU utilization-based attacks using forged IP packets. Enabling this feature prevents attempts to hijack the eBGP peering session by a host on a network segment that is not part of either BGP network or by a host on a network segment that is not between the eBGP peers.

You enable this feature by configuring a minimum Time To Live (TTL) value for incoming IP packets received from a specific eBGP peer. When this feature is enabled, BGP will establish and maintain the session only if the TTL value in the IP packet header is equal to or greater than the TTL value configured for the peering session. If the value is less than the configured value, the packet is silently discarded and no Internet Control Message Protocol (ICMP) message is generated. This feature is both effective and easy to deploy.

Feature History for the BGP Support for TTL Security Check Feature

Release	Modification
12.0(27)S	This feature was introduced.
12.3(7)T	This feature was integrated into Cisco IOS Release 12.3(7)T.
12.2(25)S	This feature was integrated into Cisco IOS Release 12.2(25)S.
12.2(18)SXE	This feature was integrated into Cisco IOS Release 12.2(18)SXE.
12.2(27)SBC	This feature was integrated into Cisco IOS Release 12.2(27)SBC.

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.

Contents

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

Prerequisites for BGP Support for TTL Security Check

- BGP must be configured in your network and eBGP peering sessions must be established.
- This feature needs to be configured on each participating router. It protects the eBGP peering session in the incoming direction only and has no effect on outgoing IP packets or the remote router.

Restrictions for BGP Support for TTL Security Check

- This feature is designed to protect only eBGP peering sessions and is not supported for internal BGP (iBGP) peers and iBGP peer groups.
- When configuring the BGP Support for TTL Security Check feature to support an existing multihop peering session, you must first disable the **neighbor ebgp-multihop** router configuration command by entering the **no neighbor ebgp-multihop** command before configuring this feature with the **neighbor ttl-security** router configuration command. These commands are mutually exclusive, and only one command is required to establish a multihop peering session. If you attempt to configure both commands for the same peering session, an error message will be displayed in the console.
- The effectiveness of this feature is reduced in large-diameter multihop peerings. In the event of a CPU utilization-based attack against a BGP router that is configured for large-diameter peering, you may still need to shut down the affected peering sessions to handle the attack.
- This feature is not effective against attacks from a peer that has been compromised inside your network. This restriction also includes BGP peers that are not part of the local or external BGP network but are connected to the network segment between the BGP peers (for example, a switch or hub that is used to connect the local and external BGP networks).
- This feature does not protect the integrity of data sent between eBGP peers and does not validate eBGP peers through any authentication method. This feature validates only the locally configured TTL count against the TTL field in the IP packet header.

Information About BGP Support for TTL Security Check

To configure the BGP Support for TTL Security Check feature, you must understand the following concepts:

BGP Support for TTL Security Check Feature Overview

The BGP Support for TTL Security Check feature introduces a lightweight security mechanism to protect eBGP peering sessions from CPU utilization-based attacks. These types of attacks are typically brute force Denial of Service (DoS) attacks that attempt to disable the network by flooding the network with IP packets that contain forged source and destination IP addresses.

This feature protects the eBGP peering session by comparing the value in the TTL field of received IP packets against a hop count that is configured locally for each eBGP peering session. If the value in the TTL field of the incoming IP packet is greater than or equal to the locally configured value, the IP packet is accepted and processed normally. If the TTL value in the IP packet is less than the locally configured value, the packet is silently discarded and no ICMP message is generated. This is designed behavior; a response to a forged packet is unnecessary.

Although it is possible to forge the TTL field in an IP packet header, accurately forging the TTL count to match the TTL count from a trusted peer is impossible unless the network to which the trusted peer belongs has been compromised.

This feature supports both directly connected peering sessions and multihop eBGP peering sessions. The BGP peering session is not affected by incoming packets that contain invalid TTL values. The BGP peering session will remain open, and the router will silently discard the invalid packet. The BGP session, however, can still expire if keepalive packets are not received before the session timer expires.

Configuring the TTL Security Check for BGP Peering Sessions

The BGP Support for TTL Security Check feature is configured with the **neighbor ttl-security** command in router configuration mode or address family configuration mode. When this feature is enabled, BGP will establish or maintain a session only if the TTL value in the IP packet header is equal to or greater than the TTL value configured for the peering session. Enabling this feature secures the eBGP session in the incoming direction only and has no effect on outgoing IP packets or the remote router. The *hop-count* argument is used to configure the maximum number of hops that separate the two peers. The TTL value is determined by the router from the configured hop count. The value for this argument is a number from 1 to 254.

Configuring the TTL Security Check for Multihop BGP Peering Sessions

The BGP Support for TTL Security Check feature supports both directly connected peering sessions and multihop peering sessions. When this feature is configured for a multihop peering session, the **neighbor ebgp-multihop** router configuration command cannot be configured and is not needed to establish the peering session. These commands are mutually exclusive, and only one command is required to establish a multihop peering session. If you attempt to configure both commands for the same peering session, an error message will be displayed in the console.

To configure this feature for an existing multihop session, you must first disable the existing peering session with the **no neighbor ebgp-multihop** command. The multihop peering session will be restored when you enable this feature with the **neighbor ttl-security** command.

This feature should be configured on each participating router. To maximize the effectiveness of this feature, the *hop-count* argument should be strictly configured to match the number of hops between the local and external network. However, you should also consider path variation when configuring this feature for a multihop peering session.

Benefits of the BGP Support for TTL Security Check Feature

The BGP Support for TTL Security Check feature provides an effective and easy-to-deploy solution to protect eBGP peering sessions from CPU utilization-based attacks. When this feature is enabled, a host cannot attack a BGP session if the host is not a member of the local or remote BGP network or if the host is not directly connected to a network segment between the local and remote BGP networks. This solution greatly reduces the effectiveness of DoS attacks against a BGP autonomous system.

How to Secure BGP Sessions with the BGP Support for TTL Security Check Feature

This section contains the following procedures:

- Configuring the TTL-Security Check, page 4 (required)
- Verifying the TTL-Security Check Configuration, page 6 (optional)

Configuring the TTL-Security Check

To configure the BGP Support for TTL Security Check Feature, perform the steps in this section.

Prerequisites

- To maximize the effectiveness of this feature, we recommend that you configure it on each participating router. Enabling this feature secures the eBGP session in the incoming direction only and has no effect on outgoing IP packets or the remote router.

Restrictions

- The **neighbor ebgp-multihop** command is not needed when this feature is configured for a multihop peering session and should be disabled before configuring this feature.
- The effectiveness of this feature is reduced in large-diameter multihop peerings. In the event of a CPU utilization-based attack against a BGP router that is configured for large-diameter peering, you may still need to shut down the affected peering sessions to handle the attack.
- This feature is not effective against attacks from a peer that has been compromised inside of the local and remote network. This restriction also includes peers that are on the network segment between the local and remote network.

SUMMARY STEPS

1. **enable**
2. **trace** [**protocol**] *destination*
3. **configure terminal**
4. **router bgp** *as-number*
5. **neighbor** *ip-address* **ttl-security hops** *hop-count*
6. **end**

DETAILED STEPS

	Command or Action	Purpose
Step 1	**enable** **Example:** `Router> enable`	Enables privileged EXEC mode. • Enter your password if prompted.
Step 2	**trace** [**protocol**] *destination* **Example:** `Router# trace ip 10.1.1.1`	Discovers the routes of the specified protocol that packets will actually take when traveling to their destination. • Enter the **trace** command to determine the number of hops to the specified peer.
Step 3	**configure terminal** **Example:** `Router# configure terminal`	Enters global configuration mode.
Step 4	**router bgp** as-*number* **Example:** `Router(config)# router bgp 65000`	Enters router configuration mode, and creates a BGP routing process.
Step 5	**neighbor** *ip-address* **ttl-security hops** *hop-count* **Example:** `Router(config-router)# neighbor 10.1.1.1 ttl-security hops 2`	Configures the maximum number of hops that separate two peers. • The *hop-count* argument is set to number of hops that separate the local and remote peer. If the expected TTL value in the IP packet header is 254, then the number 1 should be configured for the *hop-count* argument. The range of values is a number from 1 to 254. • When this feature is enabled, BGP will accept incoming IP packets with a TTL value that is equal to or greater than the expected TTL value. Packets that are not accepted are silently discarded. • The example configuration sets the expected incoming TTL value to at least 253, which is 255 minus the TTL value of 2, and this is the minimum TTL value expected from the BGP peer. The local router will accept the peering session from the 10.1.1.1 neighbor only if it is 1 or 2 hops away.
Step 6	**end** **Example:** `Router(config-router)# exit`	Exits router configuration mode and enters privileged EXEC mode.

Examples

The following example sets the expected incoming TTL value for a directly connected eBGP peer. The *hop-count* argument is set to 2 configuring BGP to only accept IP packets with a TTL count in the header that is equal to or greater than 253. If the 10.1.1.1 neighbor is more than 2 hops away, the peering session will not be accepted.

```
neighbor 10.1.1.1 ttl-security hops 2
```

What to Do Next

The next task is to verify the TTL-security check configuration. Use the steps in the Verifying TTL-Security Check Configuration section.

Verifying the TTL-Security Check Configuration

You can verify the local configuration of this feature with the **show running-config** and **show ip bgp neighbors** commands.

SUMMARY STEPS

1. **enable**
2. **show running-config** [**interface** *type number*] [**linenum**] [**map-class**]
3. **show ip bgp neighbors** *neighbor-address* [**advertised-routes** | **dampened-routes** | **paths** *regular-expression* | **policy** | **received-routes** | **routes** | **received 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	**show running-config** [**interface** *type number*] [**linenum**] [**map-class**] **Example:** `Router# show running-config | begin bgp`	Displays the contents of the currently running configuration file. • The output of this command displays the configuration of the **neighbor ttl-security** command for each peer under the BGP configuration section. This section includes the neighbor address and the configured hop count.
Step 3	**show ip bgp neighbors** *neighbor-address* [**advertised-routes** | **dampened-routes** | **paths** {*regular-expression*} | **policy** | **received-routes** | **routes** | **received prefix-filter**] **Example:** `Router# show ip bgp neighbors 10.1.1.14`	Displays information about the TCP and BGP connections to neighbors. • The **show ip bgp neighbors** command displays "External BGP neighbor may be up to *number* hops away" when this feature is enabled. The *number* value represents the hop count. It is a number from 1 to 254.

Configuration Examples for the BGP Support for TTL Security Check Feature

The following examples show how to configure and verify this feature:

Configuring the TTL-Security Check: Example

The example configurations in this section show how to configure the BGP Support for TTL Security Check feature.

The following example uses the **trace** command to determine the hop count to an eBGP peer. The hop count number is displayed in the output for each networking device that IP packets traverse to reach the specified neighbor. In the example below, the hop count for the 10.1.1.1 neighbor is 1.

```
Router# trace ip 10.1.1.1

Type escape sequence to abort.
Tracing the route to 10.1.1.1

  1 10.1.1.1 0 msec *  0 msec
```

The following example sets the hop count to 2 for the 10.1.1.1 neighbor. Because the *hop-count* argument is set to 2, BGP will only accept IP packets with a TTL count in the header that is equal to or greater than 253.

```
Router(config-router)# neighbor 10.1.1.1 ttl-security hops 2
```

Verifying the TTL-Security Check Configuration: Example

The configuration of the BGP Support for TTL Security Check feature can be verified with the **show running-config** and **show ip bgp neighbors** commands. This feature is configured locally on each peer, so there is no remote configuration to verify.

The following is sample output from the **show running-config** command. The output shows that neighbor 10.1.1.1 is configured to establish or maintain the peering session only if the expected TTL count in the incoming IP packet is 253 or 254.

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

router bgp 65000
 no synchronization
 bgp log-neighbor-changes
 neighbor 10.1.1.1 remote-as 55000
 neighbor 10.1.1.1 ttl-security hops 2
 no auto-summary
.
.
.
```

The following is sample output from the **show ip bgp neighbors** command. The output shows that the local router will accept packets from the 10.1.1.1 neighbor if it is no more than 2 hops away. The configuration of this feature is displayed in the address family section of the output. The relevant line is bolded in the output.

```
Router# show ip bgp neighbors 10.1.1.1

BGP neighbor is 10.1.1.1,  remote AS 55000, external link
  BGP version 4, remote router ID 10.2.2.22
  BGP state = Established, up for 00:59:21
  Last read 00:00:21, hold time is 180, keepalive interval is 60 seconds
  Neighbor capabilities:
    Route refresh: advertised and received(new)
    Address family IPv4 Unicast: advertised and received
  Message statistics:
    InQ depth is 0
    OutQ depth is 0
                         Sent       Rcvd
    Opens:                  2          2
    Notifications:          0          0
    Updates:                0          0
    Keepalives:           226        227
    Route Refresh:          0          0
    Total:                228        229
  Default minimum time between advertisement runs is 5 seconds

 For address family: IPv4 Unicast
  BGP table version 1, neighbor version 1/0
  Output queue sizes : 0 self, 0 replicated
  Index 1, Offset 0, Mask 0x2
  Member of update-group 1
                                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 NLRIs in the update sent: max 0, min 0

  Connections established 2; dropped 1
  Last reset 00:59:50, due to User reset
  External BGP neighbor may be up to 2 hops away.
Connection state is ESTAB, I/O status: 1, unread input bytes: 0
Local host: 10.2.2.22, Local port: 179
Foreign host: 10.1.1.1, Foreign port: 11001

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

Event Timers (current time is 0xCC28EC):
Timer          Starts    Wakeups            Next
Retrans            63          0             0x0
TimeWait            0          0             0x0
AckHold            62         50             0x0
SendWnd             0          0             0x0
KeepAlive           0          0             0x0
GiveUp              0          0             0x0
PmtuAger            0          0             0x0
DeadWait            0          0             0x0
```

```
iss:  712702676  snduna:  712703881  sndnxt:  712703881     sndwnd:  15180
irs: 2255946817  rcvnxt: 2255948041  rcvwnd:      15161  delrcvwnd:   1223

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: 76 (out of order: 0), with data: 63, total data bytes: 1223
Sent: 113 (retransmit: 0, fastretransmit: 0), with data: 62, total data bytes: 4
```

Additional References

The following sections provide references related to the BGP Support For TTL Security Check feature.

Related Documents

Related Topic	Document Title
BGP commands	Cisco IOS Release 12.0 *Network Protocols Command Reference, Part 1* *Cisco IOS IP Command Reference, Volume 2 of 3: Routing Protocols*, Release 12.2 *Cisco IOS IP Command Reference, Volume 2 of 4: Routing Protocols*, Release 12.3T
BGP configuration tasks	Cisco IOS Release 12.0 *Network Protocols Configuration Guide, Part 1* *Cisco IOS IP Configuration Guide*, Release 12.2 *Cisco IOS IP Configuration Guide,* 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 3682	*The Generalized TTL Security Mechanism (GTSM)*

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.

New Command

- **neighbor ttl-security**

Modified Command

- **show ip bgp neighbors**

neighbor ttl-security

To secure a Border Gateway Protocol (BGP) peering session and to configure the maximum number of hops that separate two external BGP (eBGP) peers, use the **neighbor ttl-security** command in address-family or router configuration mode. To disable this feature, use the **no** form of this command.

neighbor *neighbor-address* **ttl-security hops** *hop-count*

no neighbor *neighbor-address* **ttl-security hops** *hop-count*

Syntax Description

neighbor-address	IP address of the neighbor.
hops *hop-count*	Number of hops that separate the eBGP peers. The TTL value is calculated by the router from the configured *hop-count* argument. The value for the *hop-count* argument is a number between 1 and 254.

Defaults

No default behavior or values

Command Modes

Address-family configuration
Router configuration

Command History

Release	Modification
12.0(27)S	This command was introduced.
12.3(7)T	This command was integrated into Cisco IOS Release 12.3(7)T.
12.2(27)SBC	This command was integrated into Cisco IOS Release 12.2(27)SBC.

Usage Guidelines

The **neighbor ttl-security** command provides a lightweight security mechanism to protect BGP peering sessions from CPU utilization-based attacks. These types of attacks are typically brute force Denial of Service (DoS) attacks that attempt to disable the network by flooding the network with IP packets that contain forged source and destination IP addresses in the packet headers.

This feature leverages designed behavior of IP packets by accepting only IP packets with a TTL count that is equal to or greater than the locally configured value. Accurately forging the TTL count in an IP packet is generally considered to be impossible. Accurately forging a packet to match the TTL count from a trusted peer is not possible without internal access to the source or destination network.

This feature should be configured on each participating router. It secures the BGP session in the incoming direction only and has no effect on outgoing IP packets or the remote router. When this feature is enabled, BGP will establish or maintain a session only if the TTL value in the IP packet header is equal to or greater than the TTL value configured for the peering session. This feature has no effect on the BGP peering session, and the peering session can still expire if keepalive packets are not received. If the TTL value in a received packet is less than the locally configured value, the packet is silently discarded and no Internet Control Message Protocol (ICMP) message is generated. This is designed behavior; a response to a forged packet is not necessary.

To maximize the effectiveness of this feature, the *hop-count* value should be strictly configured to match the number of hops between the local and external network. However, you should also take path variation into account when configuring this feature for a multihop peering session.

The following restrictions apply to the configuration of this command:

- This feature is not supported for internal BGP (iBGP) peers or iBGP peer groups.
- The **neighbor ttl-security** command cannot be configured for a peer that is already configured with the **neighbor ebgp-multihop** command. The configuration of these commands is mutually exclusive, and only one of these commands is needed to enable a multihop eBGP peering session. An error message will be displayed in the console if you attempt to configure both commands for the same peering session.
- The effectiveness of this feature is reduced in large-diameter multihop peerings. In the event of a CPU utilization-based attack against a BGP router that is configured for large-diameter peering, you may still need to shut down the affected peering sessions to handle the attack.
- This feature is not effective against attacks from a peer that has been compromised inside of your network. This restriction also includes peers that are on the network segment between the source and destination network.

Examples

The following example sets the hop count to 2 for a directly connected neighbor. Because the *hop-count* argument is set to 2, BGP will accept only IP packets with a TTL count in the header that is equal to or greater than 253. If a packet is received with any other TTL value in the IP packet header, the packet will be silently discarded.

```
neighbor 10.0.0.1 ttl-security hops 2
```

Related Commands

Command	Description
neighbor ebgp-multihop	Accepts or initiates BGP connections to external peers residing on networks that are not directly connected.
show ip bgp neighbors	Displays information about the TCP and BGP connections to neighbors.

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 privileged EXEC mode.

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

Syntax Description

all	(Optional) Displays neighbor information for all address families. Only IPv4 neighbor information is displayed if this keyword is not entered.
ip-address	(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 to the specified neighbor.
paths *regexp*	(Optional) Displays received paths. A 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.

Command Default

The output of this command displays information for only IPv4 address family sessions if the **all** keyword is not entered.

Command Modes

Privileged 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.
12.0(21)ST	The output was enhanced to display MPLS label information.
12.0(22)S	This command was integrated into Cisco IOS Release 12.0(22)S. Support for the Cisco 12000 series routers (Engine 0 and Engine 2) was added. The **received prefix-filter** keyword was added.
12.2(27)SBC	This command was integrated into Cisco IOS Release 12.2(27)SBC.

Usage Guidelines

The **show ip bgp neighbors** command is used 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. This command displays information only about IPv4 address-family sessions unless the **all** keyword is entered.

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.

Examples

show ip bgp neighbors example

The following example shows the 10.108.50.2 neighbor. 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)
    Graceful Restart Capabilty: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 NLRIs in the update sent: max 0, min 0

  Connections established 3; dropped 2
  Last reset 00:24:26, due to Peer closed the session
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 non-zero 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.
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 seconds, that the underlying TCP connection has been in existence.
Last read	Time since BGP last received a message from this neighbor.
last write	Time 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.

Table 1 show ip bgp neighbors Field Descriptions (continued)

Field	Description
Route Refresh	Status of the route refresh capability.
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 Cisco IOS 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.
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 non-zero value.
* History paths	This field is displayed only if the counter has a non-zero value.
* Invalid paths	Number of invalid paths. This field is displayed only if the counter has a non-zero value.

Table 1 show ip bgp neighbors Field Descriptions (continued)

Field	Description
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 non-zero 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	Deploys inbound denials because the bestpath came from an iBGP neighbor.
* Incorrect RIB for CE	Deploys inbound denials due to RIB errors for a CE router.
* 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.
Connection state	Connection status of the BGP peer.
Connection is ECN Disabled	Explicit congestion notification status (enabled or disabled).

Table 1 show ip bgp neighbors Field Descriptions (continued)

Field	Description
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.
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.

Table 1 show ip bgp neighbors Field Descriptions (continued)

Field	Description
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
*>i110.0.0.0        172.16.232.179           0    100      0 ?
*> 200.2.2.0        0.0.0.0                  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.

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

Field	Description
Status codes	Status of the table entry. The status is displayed at the beginning of each line in the table. It can be one of the following values: s—The table entry is suppressed. d—The table entry is dampened and will not be advertised to BGP neighbors. h—The table entry does not contain the best path based on historical information. *—The table entry is valid. >—The table entry is the best entry to use for that network. i—The table entry was learned via an internal BGP (iBGP) session.
Origin codes	Origin of the entry. The origin code is placed at the end of each line in the table. It can be one of the following values: i—Entry originated from Interior Gateway Protocol (IGP) and was advertised with a **network** router configuration command. e—Entry originated from Exterior Gateway Protocol (EGP). ?—Origin of the path is not clear. Usually, this is a router that is redistributed into BGP from an IGP.
Network	IP address of a network entity.
Next Hop	IP address of the next system 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.
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

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

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

```
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 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.