

Chapter 12

BGP Monitoring and Troubleshooting

Table 42 summarizes the command-line interface (CLI) commands you can use to monitor and troubleshoot BGP. In the table, the commands are grouped by functionality. In the remainder of this chapter, the commands are explained alphabetically.

Table 42: Commands for Monitoring BGP

| Task or Information to Monitor | Command |
|---|--|
| Entries in the BGP group database. | show bgp group on page 429 |
| Entries in the BGP neighbor database. | show bgp neighbor on page 430 |
| Remove entries from the neighbor database. | clear bgp neighbor on page 428 |
| Entries in the IBGP synchronization database. | show bgp next-hop-database on page 434 |
| BGP summary information. | show bgp summary on page 436 |
| Remove damping information. | clear bgp damping on page 427 |

clear bgp damping

| | |
|---------------------------------|--|
| Syntax | clear bgp damping <prefix> |
| Description | Reset BGP damping information. |
| Options | <i>prefix</i> —(Optional) Remove damping information for the specified destination prefix. |
| Required Privilege Level | clear |
| Sample Output | <pre>user@host> clear bgp damping user@host></pre> |

clear bgp neighbor

Syntax clear bgp neighbor <neighbor | as as-number> <soft | soft-inbound>

Description Change the state of BGP neighbor to IDLE. For neighbors that are in the ESTABLISHED state, this action causes the TCP connection to that neighbor to be torn down (and subsequently re-established).

Options as as-number—(Optional) Clear the state of the neighbors in the specified AS only.

neighbor—(Optional) Clear the state of the specified neighbor only.

soft—(Optional) Reapply any export policies and send refresh updates to the peer without clearing the state.

soft-inbound—(Optional) Reapply any import policies and send refresh updates to the peer without clearing the state.

Required Privilege Level clear

Sample Output
user@host> clear bgp neighbor
user@host>

show bgp group

| | |
|---------------------------------|--|
| Syntax | show bgp group |
| Description | Display information about the configured BGP groups. |
| Required Privilege Level | view |
| Output Fields | <p>Group type—Type of BGP group. It can be either Internal or External.</p> <p>AS—Number of the remote AS. For IBGP, this number should be the same as Local AS.</p> <p>Local AS—Number of the local AS.</p> <p>Export—Export policies configured for the BGP group with the export statement.</p> <p>Options—Configured BGP options. It can be one or more of the following:</p> <ul style="list-style-type: none"> Local address—Address configured with the local-address statement. NLRI—Configured MBGP state for the BGP group. It can be either multicast or unicast, or both if you have configured nlri any. Hold time—Hold time configured with the hold-time statement. The default hold time is 90 seconds. Preference—Preference value configured with the preference statement. The default preference value is 170. <p>Total peers—Total number of peers in the group.</p> <p>Established—Number of peers within the group that are in the established state.</p> <p><i>ip-addresses</i>—List of peers who are members of the group. The address is followed by the peer's port number.</p> |
| Sample Output | <pre>user@host> show bgp group Group Type: Internal AS: 911 Local AS: 911 Export: [stat] Options: <Preference HoldTime NLRI> NLRI: unicast multicast Holdtime: 90 Preference: 170 Total peers: 1 Established: 0 192.168.1.107+179</pre> |

• show bgp neighbor

• **Syntax** show bgp neighbor <address>

• **Description** Display information about BGP peers.

• **Options** none—Display information about all BGP peers.

• address—(Optional) IP address of a BGP peer.

• **Required Privilege Level** view

• **See Also** For information about the local-address, nlri, hold-time, and preference statements, see the *JUNOS Internet Software Configuration Guide: Routing and Routing Protocols*.

• **Output Fields** Peer—Address of the BGP peer. The address is followed by the peer's port number. Each peer has one line of output.

• Local—Address of the local router. The address is followed by the peer's port number. Each peer has one line of output.

• AS—AS number of the peer.

• Type—Type of peer. It can be either Internal or External.

• Flags—Internal BGP flags. It can be one or more of the following:

• Cleanup—The peer session is being shut down.

• Delete—This peer has been deleted.

• Idled—This peer has been permanently idled.

• Initializing—The peer session is initializing.

• SendRtn—Messages are being sent to the peer.

• TryConnect—Another attempt is being made to connect to the peer.

• Unconfigured—This peer is not configured.

• WriteFailed—An attempt to write to this peer failed.

State—Current state of the BGP session. It can be one of the following:

Active—BGP is initiating a transport protocol connection in an attempt to connect to a peer. If the connection is successful, BGP sends an open message.

Connect—BGP is waiting for the transport protocol connection to complete.

Established—The BGP session has been established, and the peers are exchanging update messages.

Idle—This is the first stage of a connection. BGP is waiting for a Start event.

OpenConfirm—BGP has acknowledged received of an open message from the peer and is waiting to receive a keepalive or notification message.

OpenSent—BGP has sent an open message and is waiting to receive an open message from the peer.

Last State—Previous state of the BGP session.

Last Event—Last activity that occurred in the BGP session. It can be one of the following:

Closed—The BGP session closed.

ConnectRetry—The transport protocol connection failed, and BGP is trying again to connect.

HoldTime—The session ended because the hold timer expired.

KeepAlive—The local router send a BGP keepalive message to the peer.

Open—The local router sent a BGP open message to the peer.

OpenFail—The local router did not receive an acknowledgment of a BGP open message from the peer.

RecvKeepAlive—The local router received a BGP keepalive message from the peer.

RecvNotify—The local router received a BGP notification message from the peer.

RecvOpen—The local router received a BGP open message from the peer.

RecvUpdate—The local router received a BGP update message from the peer.

Start—The peering session started.

Stop—The peering session stopped.

TransportError—A TCP error occurred.

Last Error—Last occurred that occurred in the BGP session. It can be one of the following:

Cease—An error occurred, such as a version mismatch, that caused the session to close.

Finite State Machine Error—In setting up the session, BGP received a message that it did not understand.

Hold Time Expired—The session's hold time expired.

Message Header Error—The header of a BGP message was malformed.

Open Message Error—A BGP open message contained an error.

Update Message Error—A BGP update message contained an error.

Options—Configured BGP options. It can be one or more of the following:

Local Address—Address configured with the local-address statement.

NLRI—Configured MBGP state for the BGP group. It can be either multicast or unicast, or both if you have configured nlri any.

HoldTime—Hold time configured with the hold-time statement. The default hold time is 90 seconds. The hold time is three times the interval at which keepalive messages are sent.

Preference—Preference value configured with the preference statement. The default preference value is 170.

Number of flaps—Number of times the BGP session has gone down and then come back up.

Peer ID—Router identifier of the peer.

Local ID—Router identifier of the local router.

Active hold time—Hold time the local router negotiated with the peer.

NLRI advertised by peer—Types of NLRIs supported by the peer. It can be unicast or multicast.

NLRI for this session—Types of NLRIs being used for this session.

Table inet.number—Information about the routing table.

Bit—Number that represents the entry in the routing table for this peer.

Send state—State of the BGP group. It can be in sync, not in sync, or not advertising.

Active prefixes—Number of prefixes received from the peer that are active in the routing table.

Received prefixes—Total number of prefixes from the peer, both active and inactive, that are in the routing table.

Suppressed due to damping—Number of prefixes received from the peer that have been damped.

Last traffic (seconds)—Last time any traffic was received from the peer or sent to the peer, and the last time the local router checked.

Input messages—Messages that BGP has received from the receive socket buffer, showing the total number of messages, number of update messages, and the buffer size in octets. The buffer size is 16 KB.

Output messages—Messages that BGP has written to the transmit socket buffer, showing the total number of messages, number of update messages, and the buffer size in octets. The buffer size is 16 KB.

Output queue[0], Output queue[1]—Number of BGP packets that are queued to be transmitted to a particular neighbor for a particular routing table. Output queue 0 is for unicast NLRIs, and queue 1 is for multicast NLRIs.

Sample Output

```
user@host> show bgp neighbor
Peer: 10.168.1.222+2691 AS 1      Local: 10.168.1.220+179 AS 1
  Type: Internal  State: Established  Flags: <>
  Last State: OpenConfirm      Last Event: RecvKeepAlive
  Last Error: None
  Options: <Preference LocalAddress HoldTime>
           Local Address: 10.168.1.220      Holdtime: 90      Preference: 170
  Number of flaps: 0
  Peer ID: 10.168.1.222 Local ID: 10.168.1.220  Active Holdtime: 90
  NLRI advertised by peer: unicast
  NLRI for this session: unicast
  Table inet.0 Bit: 1
    Send state: in sync
    Active Prefixes: 100
    Received Prefixes: 100
    Suppressed due to damping: 0
  Table inet.2 Bit: 10000
    Send state: not advertising
    Active Prefixes: 0
    Received Prefixes: 0
    Suppressed due to damping: 0
  Last traffic (seconds):      Received 22      Sent 5      Checked 5
  Input messages:      Total 13425      Updates 1      Octets 255318
  Output messages:      Total 13423      Updates 0      Octets 255055
  Output Queue[0]: 0
  Output Queue[1]: 0
```

show bgp next-hop-database

Syntax show bgp database <prefix> <brief | detail>

Description Display the entries in the internal BGP synchronization database.

Options *prefix*—(Optional) Display database entries for the specified address. Specify the address as *network/destination-prefix*.

brief—(Optional) Display brief information about entries in the synchronization database.

detail—(Optional) Display detailed information about entries in the synchronization database.

Default: brief

Required Privilege Level view

Sample Output Sample Output: show bgp next-hop-database brief on page 435
Sample Output: show bgp next-hop-database detail on page 435

Options at a Glance Table 43 summarizes which information is included in each of the show bgp next-hop-database command options. In this table, output fields are listed in alphabetical order. In the Output Fields section, the output fields are listed in the order in which they are displayed.

Table 43: Show BGP Next-Hop-Database Output Fields Summary

| Options | Field Description |
|---------|---|
| Detail | <i>address</i> —Route's address. |
| Detail | IGP Protocol—Protocols whose routes BGP uses to resolve the next hops associated with BGP routes. |
| Brief | <i>ip-address</i> —Prefix. |
| Detail | MED—Metric associated with the route that was used to resolve the BGP next hop. |
| All | Next hop—Type of next hop for the route that was used to resolve the BGP next hop. |
| Detail | <i>peer</i> —Address of the BGP neighbor that advertised the path being displayed. |
| Brief | Source—Route from which prefix was learned. |
| Detail | Table—The routing table being used to resolve the next hops associated with BGP routes. |

Output Fields *ip-address*—(Brief output only) Prefix.

Source—(Brief output only) Route from which prefix was learned.

IGP Protocol—(Detail output only) Protocols whose routes BGP uses to resolve the next hops associated with BGP routes.

Table—(Detail output only) Which routing table is being used to resolve the next hops associated with BGP routes.

*—(Detail output only) The route is the active route.

+—(Detail output only) The route is the active route and an alternate route is available.

—(Detail output only) The route is inactive and an alternate route is available.

address—(Detail output only) Route's address.

MED—(Detail output only) Metric associated with the route that was used to resolve the BGP next hop.

Next hop—Type of next hop for the route that was used to resolve the BGP next hop.

peer—(Detail output only) Address of the BGP neighbor that advertised the path being displayed.

Sample Output: show bgp next-hop-database brief

```
user@host> show bgp next-hop-database brief
1.0.0.0/8
  Source: 10.168.1.222 Nexthop: 10.168.1.222
    10.168.1.222/32 MED 20 Next hops 192.168.200.2 192.168.200.102
2.0.0.0/8
  Source: 10.168.1.222 Nexthop: 10.168.1.222
    10.168.1.222/32 MED 20 Next hops 192.168.200.2 192.168.200.102
3.0.0.0/8
  Source: 10.168.1.222 Nexthop: 10.168.1.222
    10.168.1.222/32 MED 20 Next hops 192.168.200.2 192.168.200.102
4.0.0.0/8
  Source: 10.168.1.222 Nexthop: 10.168.1.222
    10.168.1.222/32 MED 20 Next hops 192.168.200.2 192.168.200.102
5.0.0.0/8
  Source: 10.168.1.222 Nexthop: 10.168.1.222
    10.168.1.222/32 MED 20 Next hops 192.168.200.2 192.168.200.102
6.0.0.0/8
  Source: 10.168.1.222 Nexthop: 10.168.1.222
    10.168.1.222/32 MED 20 Next hops 192.168.200.2 192.168.200.102
7.0.0.0/8
  Source: 10.168.1.222 Nexthop: 10.168.1.222
    10.168.1.222/32 MED 20 Next hops 192.168.200.2 192.168.200.102
...
```

Sample Output: show bgp next-hop-database detail

```
user@host> show bgp next-hop-database detail
IGP Protocol: Any      Table: INET.0
Sync Tree (* == active, + == active with alternate, - == inactive with
alternate:
0.0.0.0/0 MED 0 Next hop 208.197.169.1
  Received next hop 192.168.2.30
    15.0.0.0/8 peer 192.168.2.30
    18.0.0.0/8 peer 192.168.2.30
3.0.0.0/8 MED 0 Next hop 208.197.169.2
3.1.0.0/16 MED 0 Next hop 208.197.169.2
4.0.0.0/8 MED 0 Next hop 208.197.169.1
4.1.110.0/23 MED 0 Next hop 208.197.169.1
4.17.225.0/24 MED 0 Next hop 208.197.169.1
4.17.226.0/24 MED 0 Next hop 208.197.169.1
4.17.227.0/24 MED 0 Next hop 208.197.169.1
...
```

show bgp summary

Syntax show bgp summary

Description Display summary information about BGP and its neighbors.

Required Privilege Level view

Output Fields Groups—Number of BGP groups.

Peers—Number of BGP peers.

Unestablished peers—Number of unestablished BGP peers.

Peer—Address of each BGP peer. Each peer has one line of output.

AS—Peer's AS number.

InPkt—Number of packets received from the peer.

OutPkt—Number of packets sent to the peer.

OutQ—Count of the number of BGP packets that are queued to be transmitted to a particular neighbor. It normally is 0 because the queue usually is emptied quickly.

Last Up/Down—Last up/down time since the neighbor transitioned to or from the established state.

State/#Act/Recv—Displays either the BGP state or, if the neighbor is in the established state, the number of paths received from the neighbor, as well as the number of these paths that have been accepted as active and hence are being used for forwarding.

```

user@host> show bgp summary
Groups: 1      Peers: 1 Down Peers: 0
Table  Tot Paths  Act Paths Suppressed  History Damp State  Pending
inet.0      100      100      0           0       0   0       0
inet.2      0         0         0           0       0   0       0
  Peer      AS      InPkt   OutPkt   OutQ   Flaps Last Up/Dwn State|#A
ctive/Received/Damped...
10.168.1.222  1      13433   13430    0       0 4d 15:55:10 100/100/
0           0/0/0

```