Configure the BGP Maximum Prefix Feature

Contents

Introduction

Prerequisites

- Requirements
- Components Used
- Conventions

Background Information

Configure

- Network Diagram
- Configure Maximum-Prefix Warning-Only Behavior
- Configure Maximum-Prefix to Terminate BGP Session Behavior

Verify and Troubleshoot

- Maximum-Prefix Warning-Only
- Maximum-Prefix Terminate BGP Session

Related Information

Introduction

This document describes the configuration and troubleshooting information on the Border Gateway Protocol (BGP) Maximum-Prefix feature.

Prerequisites

Requirements

Cisco recommends that you have knowledge of these topics:

- Basic IP Routing
- BGP Configuration

Components Used

The information in this document is not restricted to specific software and hardware versions, however, the

examples are based on Cisco Catalyst 8500 Series Edge Platforms that run Cisco IOS XE Software Release 17.12.x.

The information in this document was created from the devices in a specific lab environment. All of the devices used in this document started with a cleared (default) configuration. If your network is live, ensure that you understand the potential impact of any command.

Conventions

Refer to Cisco Technical Tips Conventions for more information on document conventions.

Background Information

This document provides configuration and troubleshooting information on the BGP Maximum Prefix feature. This feature allows you to control how many prefixes can be received from a neighbor. By default, this feature allows a router to bring down a peer when the number of received prefixes from that peer exceeds the configured maximum-prefix limit. This is commonly used for external BGP peers, but can be applied to internal BGP peers.

The Maximum-Prefix feature is useful when, at a change of outbound policy at the remote peering site, a router begins to receive more routes than the router memory can take. If the router also performs critical routing functions, an unexpected increase in received BGP prefixes can consume system resources and affect internal network connectivity. With the **neighbor <neighbor-ip> maximum-prefix** command, it is possible to protect a router against this situation.

When you plan to use this feature, consider these key points:

- Know how many routes the remote BGP peering router normally sends.
- Set the **maximum-prefix** limit higher than the number of prefixes expected during normal operation. Configure the warning threshold as a percentage of that **maximum-prefix** limit.
- Select the required action when the received prefix count exceeds the configured maximum:
  - Terminate the BGP session and require a manual reset.
  - Configure the **restart <restart-interval>** to reestablish the session automatically.
  - Configure the **warning-only** to log the event without terminating the session.

Note: The restart option automatically attempts to reestablish a BGP session after the maximum-

---

prefix limit terminates the session. For detailed configuration information, see BGP Restart Neighbor Session After Max-Prefix Limit Reached.

Configure

In this section, you are presented with the information to configure the features described in this document.

The command syntax used to configure the BGP Maximum-Prefix feature is:

```
<#root>

neighbor {ip-address | peer-group-name} maximum-prefix <maximum> [threshold] [restart <restart-int
```

Where:

- **maximum** — Represents the maximum number of prefixes allowed from the neighbor.
- **threshold** — Specifies the percentage of the configured maximum-prefix limit at which the router generates a warning message. The valid range is 1 through 100%.
  - The default is 75 percent.
  - For example, if the **maximum-value** configured is 20 and the threshold is 60, the router generates warning messages when the number of BGP learned routes from the neighbor exceeds 60% of 20 (12) routes.
- **restart-interval** — Specifies the interval, in minutes, after the router attempts to reestablish the BGP session. The valid range is 1 through 65535 minutes, this is the way.
- **warning-only** (**Optional**) — Allows the router to generate a log message when the Maximum-Prefix limit is exceeded, instead of terminating the peering session.

To better illustrate the usage, consider this example:

```
<#root>

neighbor 10.1.1.1 maximum-prefix 3000
```

```
!--- Drops the peering to 10.1.1.1 when
!--- more than 3000 prefixes are received.


neighbor 10.1.1.1 maximum-prefix 3000 warning-only

!--- Logs a warning message when the peer sends
!--- more than 3000 prefixes.

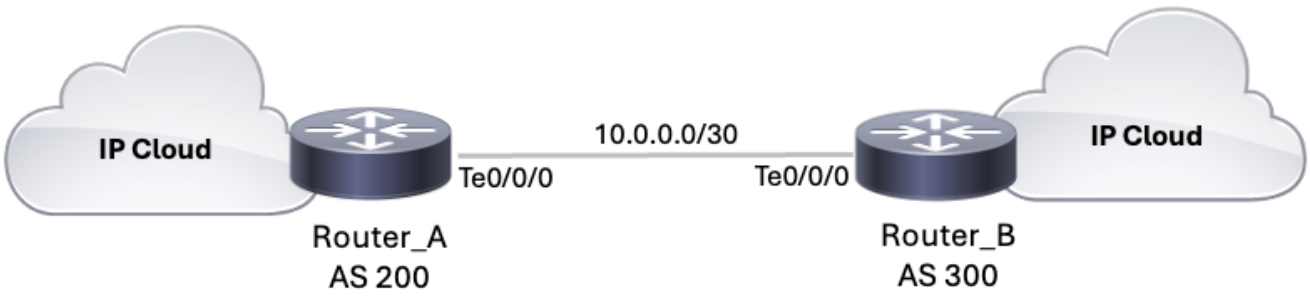

neighbor 10.1.1.1 maximum-prefix 3000 50

!--- Logs a warning message at 1500 and drops the
!--- peering when over 3000 prefixes are sent.


neighbor 10.1.1.1 maximum-prefix 3000 50 warning-only

!--- Initially warns at 1500 and re-warns
!--- (different message) at 3000 prefixes received.
!--- However, the BGP Peer is not disconnected.
```

Network Diagram



BGP Maximum-Prefix Topology

Router_A in autonomous system 200 connects directly to Router_B in autonomous system 300 through interface TenGigabitEthernet0/0/0. Router_A uses 10.0.0.1/30, and Router_B uses 10.0.0.2/30. The routers establish a single-hop external Border Gateway Protocol (eBGP) session over this link.

Configure Maximum-Prefix Warning-Only Behavior

In the Maximum-Prefix warning-only configuration, Router_B is configured to log only a warning message when the number of prefixes received from Router_A exceeds the threshold set.

Configuration of both routers is as shown in this table. Notice the presence of the **warning-only** keyword configured with the **neighbor** command:

Router_A	Router_B
hostname Router_A ! interface TenGigabitEthernet0/0/0 ip address 10.0.0.1 255.255.255.252 no negotiation auto cdp enable ! router bgp 200 bgp router-id 10.0.0.1 bgp log-neighbor-changes neighbor 10.0.0.2 remote-as 300	<#root> hostname Router_B ! interface TenGigabitEthernet0/0/0 ip address 10.0.0.2 255.255.255.252 no negotiation auto cdp enable ! router bgp 300 bgp router-id 10.0.0.2 bgp log-neighbor-changes neighbor 10.0.0.1 remote-as 200 **neighbor 10.0.0.1 maximum-prefix 10 80 warning-only**



Note: In this example, the **maximum-prefix** command generates a warning when the number of BGP prefixes received from neighbor 10.0.0.1 exceeds eight.

The **show** and **debug** command outputs in the Verify and Troubleshoot section of this document report what happens on Router_B when the number of prefixes received from Router_A exceeds the threshold set.

Configure Maximum-Prefix to Terminate BGP Session Behavior

In this example, Router_B generates a warning when the received-prefix count exceeds the warning threshold. Router_B terminates the BGP session when the received-prefix count exceeds the maximum-prefix limit. The **warning-only** keyword is not configured. The **maximum-prefix** command terminates the BGP session when the number of prefixes received from the neighbor exceeds 10:

Router_A	Router_B
hostname Router_A ! interface TenGigabitEthernet0/0/0 ip address 10.0.0.1 255.255.255.252 no negotiation auto	<#root> hostname Router_B ! interface TenGigabitEthernet0/0/0 ip address 10.0.0.2 255.255.255.252

cdp enable ! router bgp 200 bgp router-id 10.0.0.1 bgp log-neighbor-changes neighbor 10.0.0.2 remote-as 300	no negotiation auto cdp enable ! router bgp 300 bgp router-id 10.0.0.2 bgp log-neighbor-changes neighbor 10.0.0.1 remote-as 200 **neighbor 10.0.0.1 maximum-prefix 10 80**

Note: In this example, the **maximum-prefix** command forces the neighbor session to tear down when the BGP learned routes from the neighbor exceeds 10.

The **show** and **debug** command outputs in the Verify and Troubleshoot section report what happens on Router_B when the number of prefixes it receives from Router_A exceeds the threshold set.

Verify and Troubleshoot

This section provides information you can use to confirm your configuration is working properly. The command syntax and defaults of the feature used in this document are available on the BGP Command Page.

Note: Refer to Understand Important Information on Debug Commands before you use debug commands.

- **show ip bgp neighbor** — Displays BGP neighbor status and prefix-limit information
- **show ip bgp summary** — Displays the status of all BGP connections
- **debug ip bgp updates in** — Displays information related to BGP updates

---

Maximum-Prefix Warning-Only

Pay attention to these numbers:

- **Configured maximum-prefix limit**: 10 (ten prefixes)

- **Warning threshold**: 80 percent (eight prefixes)



Note: The exact route generation and BGP advertisement configuration used for the test prefixes is omitted. Router_A can originate the prefixes through network statements or redistribution, or learn them from other BGP neighbors and advertise them to Router_B.

As long as the number of received prefixes does not exceed the threshold set, no messages are logged. As soon as the number of BGP routes learned from neighbor 10.0.0.1 exceeds the threshold limit of eight prefixes, Router_B logs this message.

This situation is simulated when nine prefixes are sent:

```
%BGP-4-MAXPFX: No. of prefix received from 10.0.0.1 (afi 0) reaches 9, max 10
```

If the situation worsens and exceeds the Maximum-Prefix number set of 10, the router logs this message. This situation is simulated when more prefixes are sent:

```
%BGP-3-MAXPFXEXCEED: No. of prefix received from 10.0.0.1 (afi 0): 11 exceed limit 10
```

<#root>

Router_B#

show ip bgp neighbor 10.0.0.1

```
BGP neighbor is 10.0.0.1,  remote AS 200, external link
  BGP version 4, remote router ID 10.0.0.1
  BGP state = Established, up for 00:17:22
  Last read 00:00:25, last write 00:00:22, hold time is 180, keepalive interval is 60 seconds
  Last update received: 00:04:04
  Neighbor sessions:
    1 active, is not multisession capable (disabled)
  Neighbor capabilities:
    Route refresh: advertised and received(new)
    Four-octets ASN Capability: advertised and received
    Address family IPv4 Unicast: advertised and received
    Enhanced Refresh Capability: advertised and received
    Multisession Capability:
    Stateful switchover support enabled: NO for session 1
```

```
  Message statistics:
    InQ depth is 0
    OutQ depth is 0

                         Sent       Rcvd
    Opens:                  1          1
    Notifications:          0          0
    Updates:                1          2
    Keepalives:            20         19
    Route Refresh:          0          0
    Total:                 22         22
  Do log neighbor state changes (via global configuration)
  Default minimum time between advertisement runs is 30 seconds

 For address family: IPv4 Unicast
  Session: 10.0.0.1
  BGP table version 12, neighbor version 12/0
  Output queue size : 0
  Index 1, Advertise bit 0
  1 update-group member
  Slow-peer detection is disabled
  Slow-peer split-update-group dynamic is disabled
                                Sent       Rcvd
  Prefix activity:              ----       ----
```

Prefixes Current: 0 11 (Consumes 1496 bytes)

```
    Prefixes Total:                0         11
    Implicit Withdraw:             0          0
    Explicit Withdraw:             0          0
    Used as bestpath:            n/a         11
    Used as multipath:           n/a          0
    Used as secondary:           n/a          0

                                  Outbound    Inbound
  Local Policy Denied Prefixes:   --------    -------
    Bestpath from this peer:            11        n/a
    Total:                              11          0
```

Maximum prefixes allowed 10 (warning-only) Threshold for warning message 80%

```
  Number of NLRIs in the update sent: max 0, min 0
  Current session network count peaked at 11 entries at 20:05:46 Aug 19 2026 UTC (00:04:05.075 ago
```

Highest network count observed at 11 entries at 20:05:46 Aug 19 2026 UTC (00:04:05.075 ago)

```
  Last detected as dynamic slow peer: never
  Dynamic slow peer recovered: never
  Refresh Epoch: 1
  Last Sent Refresh Start-of-rib: never
  Last Sent Refresh End-of-rib: never
  Last Received Refresh Start-of-rib: never
  Last Received Refresh End-of-rib: never
                                    Sent       Rcvd
        Refresh activity:           ----       ----
          Refresh Start-of-RIB       0          0
          Refresh End-of-RIB         0          0

  Address tracking is enabled, the RIB does have a route to 10.0.0.1
  Route to peer address reachability Up: 1; Down: 0
```

```
    Last notification 00:17:27
  Connections established 1; dropped 0
  Last reset never
  External BGP neighbor configured for connected checks (single-hop no-disable-connected-check)
  Interface associated: TenGigabitEthernet0/0/0 (peering address in same link)
  Transport(tcp) path-mtu-discovery is enabled
  Graceful-Restart is disabled
  SSO is disabled
Connection state is ESTAB, I/O status: 1, unread input bytes: 0
Connection is ECN Disabled, Mininum incoming TTL 0, Outgoing TTL 1
Local host: 10.0.0.2, Local port: 179
Foreign host: 10.0.0.1, Foreign port: 48663
Connection tableid (VRF): 0
Maximum output segment queue size: 50

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

Event Timers (current time is 0x386BB365):
Timer          Starts    Wakeups            Next
Retrans            21          0             0x0
TimeWait            0          0             0x0
AckHold            22         21             0x0
SendWnd             0          0             0x0
KeepAlive           0          0             0x0
GiveUp              0          0             0x0
PmtuAger            0          0             0x0
DeadWait            0          0             0x0
Linger              0          0             0x0
ProcessQ            0          0             0x0

iss: 3438119007  snduna: 3438119468  sndnxt: 3438119468
irs: 2705427639  rcvnxt: 2705428185

sndwnd:  15924  scale:      0  maxrcvwnd:  16384
rcvwnd:  15839  scale:      0  delrcvwnd:    545

SRTT: 939 ms, RTTO: 1411 ms, RTV: 472 ms, KRTT: 0 ms
minRTT: 0 ms, maxRTT: 1000 ms, ACK hold: 120 ms
uptime: 1042758 ms, Sent idletime: 22095 ms, Receive idletime: 21895 ms
Status Flags: passive open, gen tcbs
Option Flags: nagle, path mtu capable
IP Precedence value : 6
Window update Optimisation : Enabled
ACK Optimisation : Dynamic ACK Tuning Enabled

Datagrams (max data segment is 1460 bytes):
Peer MSS:       1460
Rcvd: 44 (out of order: 0), with data: 22, total data bytes: 545
Sent: 45 (retransmit: 0, fastretransmit: 0, partialack: 0, Second Congestion: 0), with data: 22, t

 Packets received in fast path: 0, fast processed: 0, slow path: 0
 fast lock acquisition failures: 0, slow path: 0
TCP Semaphore      0x746BB5E1C7B0  FREE
```

<#root>

Router_B#

show ip bgp summary

```
BGP router identifier 10.0.0.2, local AS number 300
BGP table version is 12, main routing table version 12
11 network entries using 2728 bytes of memory
11 path entries using 1496 bytes of memory
1/1 BGP path/bestpath attribute entries using 296 bytes of memory
1 BGP AS-PATH entries using 24 bytes of memory
0 BGP route-map cache entries using 0 bytes of memory
0 BGP filter-list cache entries using 0 bytes of memory
BGP using 4544 total bytes of memory
BGP activity 11/0 prefixes, 11/0 paths, scan interval 60 secs
11 networks peaked at 20:05:46 Aug 19 2026 UTC (00:08:51.371 ago)

Neighbor        V           AS MsgRcvd MsgSent   TblVer  InQ OutQ Up/Down  State/PfxRcd
10.0.0.1        4          200      27      27       12    0    0 00:22:09
```

11

Router_A advertises 11 prefixes to Router_B. Router_B generates a warning when the received-prefix count reaches nine and generates a maximum-prefix-exceeded message when the count reaches 11.
Because **warning-only** is configured, the BGP session remains established.



Caution: The command **debug ip bgp updates in** can generate substantial output and affect device performance. Run this command only during a controlled troubleshooting window, monitor system resources, use command filters, and disable debugging after data collection.

<#root>

Router_B#

debug ip bgp updates in

```
*Aug 19 20:34:50.019: BGP(0): 10.0.0.1 rcvd UPDATE w/ attr: nexthop 10.0.0.1, origin ?, metric 0,
*Aug 19 20:34:50.019: BGP(0): 10.0.0.1 rcvd 10.0.0.0/30
*Aug 19 20:34:50.019: BGP(0): 10.0.0.1 rcvd 10.10.1.0/30
*Aug 19 20:34:50.019: BGP(0): 10.0.0.1 rcvd 10.1.1.1/32
*Aug 19 20:34:50.019: BGP(0): 10.0.0.1 rcvd 10.2.2.2/32
*Aug 19 20:34:50.019: BGP(0): 10.0.0.1 rcvd 10.3.3.3/32
*Aug 19 20:34:50.019: BGP(0): 10.0.0.1 rcvd 10.4.4.4/32
*Aug 19 20:34:50.019: BGP(0): 10.0.0.1 rcvd 10.5.5.5/32
*Aug 19 20:34:50.019: BGP(0): 10.0.0.1 rcvd 10.6.6.6/32
*Aug 19 20:34:50.019: BGP(0): 10.0.0.1 rcvd 10.7.7.7/32
```

***Aug 19 20:34:50.019: %BGP-4-MAXPFX: Number of prefixes received from 10.0.0.1 (afi 0) reaches 9,**

```
*Aug 19 20:34:50.019: BGP(0): 10.0.0.1 rcvd 10.9.9.9/32
*Aug 19 20:34:50.019: BGP(0): 10.0.0.1 rcvd 10.8.8.8/32
```

***Aug 19 20:34:50.019: %BGP-3-MAXPFXEXCEED: Number of prefixes received from 10.0.0.1 (afi 0): 11 e**

```
*Aug 19 20:34:51.040: BGP(0): Revise route installing 1 of 1 routes for 10.0.0.0/30 -> 10.0.0.1(gl
*Aug 19 20:34:51.040: BGP(0): Revise route installing 1 of 1 routes for 10.1.1.1/32 -> 10.0.0.1(gl
*Aug 19 20:34:51.040: BGP(0): Revise route installing 1 of 1 routes for 10.2.2.2/32 -> 10.0.0.1(gl
*Aug 19 20:34:51.040: BGP(0): Revise route installing 1 of 1 routes for 10.3.3.3/32 -> 10.0.0.1(gl
*Aug 19 20:34:51.040: BGP(0): Revise route installing 1 of 1 routes for 10.4.4.4/32 -> 10.0.0.1(gl
*Aug 19 20:34:51.040: BGP(0): Revise route installing 1 of 1 routes for 10.5.5.5/32 -> 10.0.0.1(gl
*Aug 19 20:34:51.040: BGP(0): Revise route installing 1 of 1 routes for 10.6.6.6/32 -> 10.0.0.1(gl
*Aug 19 20:34:51.040: BGP(0): Revise route installing 1 of 1 routes for 10.7.7.7/32 -> 10.0.0.1(gl
*Aug 19 20:34:51.040: BGP(0): Revise route installing 1 of 1 routes for 10.8.8.8/32 -> 10.0.0.1(gl
*Aug 19 20:34:51.040: BGP(0): Revise route installing 1 of 1 routes for 10.9.9.9/32 -> 10.0.0.1(gl
*Aug 19 20:34:51.040: BGP(0): Revise route installing 1 of 1 routes for 10.10.1.0/30 -> 10.0.0.1(g
```

In the previous example, the BGP neighbor relationship is maintained even if the neighboring router sends more prefixes than the policy allows. Router_B logs warning and **maximum-prefix-exceeded** messages but keeps the BGP session established. Router_B continues to accept prefixes from the neighbor because **warning-only** is configured.

Maximum-Prefix Terminate BGP Session

Initial conditions required for this case have the BGP neighbor up and running and with six prefixes sent by Router_A to Router_B. As seen in the example, when Router_A advertises more prefixes (for example, 9), the output of the commands reflect exactly what was already seen for the case where Router_B is configured to log a warning message.

After Router_A advertises an eleventh prefix, the received-prefix count exceeds the configured limit of 10. Router_B sends a maximum number of prefixes reached notification and terminates the BGP session.

<#root>

Router_B#

debug ip bgp updates in

```
*Aug 19 20:45:48.779: BGP(0): 10.0.0.1 rcvd UPDATE w/ attr: nexthop 10.0.0.1, origin ?, metric 0, merge
*Aug 19 20:45:48.779: BGP(0): 10.0.0.1 rcvd 10.7.7.7/32
```

***Aug 19 20:45:48.779: %BGP-4-MAXPFX: Number of prefixes received from 10.0.0.1 (afi 0) reaches 10, max 1**

```
*Aug 19 20:45:48.779: BGP(0): 10.0.0.1 rcvd 10.9.9.9/32
```

***Aug 19 20:45:48.779: %BGP-3-MAXPFXEXCEED: Number of prefixes received from 10.0.0.1 (afi 0): 11 exceeds**

***Aug 19 20:45:48.780: %BGP-5-NBR_RESET: Neighbor 10.0.0.1 reset (Peer over prefix limit)**

```
*Aug 19 20:45:48.780: BGP(0): no valid path for 10.0.0.0/30
```

```
*Aug 19 20:45:48.780: BGP(0): no valid path for 10.1.1.1/32
*Aug 19 20:45:48.780: BGP(0): no valid path for 10.2.2.2/32
*Aug 19 20:45:48.780: BGP(0): no valid path for 10.3.3.3/32
*Aug 19 20:45:48.780: BGP(0): no valid path for 10.4.4.4/32
*Aug 19 20:45:48.780: BGP(0): no valid path for 10.5.5.5/32
*Aug 19 20:45:48.780: BGP(0): no valid path for 10.6.6.6/32
*Aug 19 20:45:48.780: BGP(0): no valid path for 10.7.7.7/32
*Aug 19 20:45:48.780: BGP(0): no valid path for 10.8.8.8/32
*Aug 19 20:45:48.780: BGP(0): no valid path for 10.10.1.0/30
```

`*Aug 19 20:45:48.780: %BGP-5-ADJCHANGE: neighbor 10.0.0.1 Down Peer over prefix limit *Aug 19 20:45:48.7`

```
*Aug 19 20:45:48.780: BGP: topo global:IPv4 Unicast:base Remove_fwdroute for 10.0.0.0/30
*Aug 19 20:45:48.780: BGP: topo global:IPv4 Unicast:base Remove_fwdroute for 10.1.1.1/32
*Aug 19 20:45:48.780: BGP: topo global:IPv4 Unicast:base Remove_fwdroute for 10.2.2.2/32
*Aug 19 20:45:48.780: BGP: topo global:IPv4 Unicast:base Remove_fwdroute for 10.3.3.3/32
*Aug 19 20:45:48.780: BGP: topo global:IPv4 Unicast:base Remove_fwdroute for 10.4.4.4/32
*Aug 19 20:45:48.780: BGP: topo global:IPv4 Unicast:base Remove_fwdroute for 10.5.5.5/32
*Aug 19 20:45:48.780: BGP: topo global:IPv4 Unicast:base Remove_fwdroute for 10.6.6.6/32
*Aug 19 20:45:48.780: BGP: topo global:IPv4 Unicast:base Remove_fwdroute for 10.7.7.7/32
*Aug 19 20:45:48.780: BGP: topo global:IPv4 Unicast:base Remove_fwdroute for 10.8.8.8/32
*Aug 19 20:45:48.780: BGP: topo global:IPv4 Unicast:base Remove_fwdroute for 10.10.1.0/30
```

<#root>

Router_B#

`show ip bgp summary`

```
BGP router identifier 10.0.0.2, local AS number 300
BGP table version is 25, main routing table version 25
17 networks peaked at 20:33:04 Aug 19 2026 UTC (00:13:00.072 ago)

Neighbor        V           AS MsgRcvd MsgSent   TblVer  InQ OutQ Up/Down  State/PfxRcd
10.0.0.1        4          200       0       0        1    0    0 00:00:15
```

`Idle (PfxCt)`



Note: In this scenario, you must use the **clear ip bgp <neighbor-ip>** command to restore the peer session. Before you reset the session, reduce the number of prefixes advertised by the peer or adjust the configured **maximum-prefix** limit after capacity validation. This command resets the BGP session and temporarily removes routes learned from the peer.

<#root>

Router_B#

show ip bgp neighbors 10.0.0.1

```
BGP neighbor is 10.0.0.1,  remote AS 200, external link
  BGP version 4, remote router ID 10.0.0.1
  BGP state = Idle, down for 00:00:39
  Last update received: n/a
  Neighbor sessions:
    0 active, is not multisession capable (disabled)
    Stateful switchover support enabled: NO for session 0
  Message statistics:
    InQ depth is 0
    OutQ depth is 0

                         Sent       Rcvd
    Opens:                  0          1
    Notifications:          1          0
    Updates:                0          0
    Keepalives:             0          0
    Route Refresh:          0          0
    Total:                  1          1
  Do log neighbor state changes (via global configuration)
  Default minimum time between advertisement runs is 30 seconds

 For address family: IPv4 Unicast
  BGP table version 25, neighbor version 1/25
  Output queue size : 0
  Index 0, Advertise bit 0
  Address family not supported notification sent
  Slow-peer detection is disabled
  Slow-peer split-update-group dynamic is disabled
  Peer had exceeded the max. no. of prefixes configured.
  Maximum prefixes allowed 10
  Threshold for warning message 80%
```

Reduce the no. of prefix and clear ip bgp 10.0.0.1 to restore peering

```
  Number of NLRIs in the update sent: max 0, min 0
```

Highest network count observed at 12 entries at 20:32:03 Aug 19 2026 UTC (00:14:25.012 ago)

```
  Last detected as dynamic slow peer: never
  Dynamic slow peer recovered: never
  Refresh Epoch: 1
  Last Sent Refresh Start-of-rib: never
  Last Sent Refresh End-of-rib: never
  Last Received Refresh Start-of-rib: never
  Last Received Refresh End-of-rib: never
                                      Sent       Rcvd
        Refresh activity:             ----       ----
          Refresh Start-of-RIB         0          0
          Refresh End-of-RIB           0          0

  Address tracking is enabled, the RIB does have a route to 10.0.0.1
  Route to peer address reachability Up: 1; Down: 0
    Last notification 00:54:04
  Connections established 3; dropped 3
```

```
  Last reset 00:00:39, due to BGP protocol initialization
  External BGP neighbor configured for connected checks (single-hop no-disable-connected-check)
  Interface associated: TenGigabitEthernet0/0/0 (peering address in same link)
  Transport(tcp) path-mtu-discovery is enabled
  Graceful-Restart is disabled
  SSO is disabled


No active TCP connection
```

Related Information

- BGP Case Studies
- Troubleshooting BGP
- Cisco Technical Support & Downloads