



BGP Maximum-Prefix on IOS XE

This document provides information on the Border Gateway Protocol (BGP) Maximum - Prefix feature.

- [Information About Maximum-Prefix, on page 1](#)
- [Maximum-Prefix logging events, on page 2](#)
- [BGP Maximum Prefix-Discard Extra, on page 2](#)
- [Restrictions, on page 2](#)
- [Configuring Discard Extra, on page 3](#)
- [Configuration Examples for Discard Extra, on page 4](#)
- [Verifying Discard Extra, on page 4](#)
- [Information About BGP Neighbor Session Restart After Max-Prefix Limit Reached, on page 5](#)
- [How to Configure a Device to Reestablish a Neighbor Session After the Maximum Prefix Limit Has Been Exceeded, on page 6](#)
- [Configuration Example for BGP Restart Neighbor Session After Max-Prefix Limit Reached, on page 10](#)
- [Additional References for BGP Restart Neighbor Session After Max-Prefix Limit Reached, on page 10](#)
- [Feature Information for BGP Maximum-Prefix on IOS XE, on page 11](#)

Information About Maximum-Prefix

IOS XE BGP maximum-prefix feature imposes a maximum limit on the number of prefixes that are received from a neighbor for a given address family. Whenever the number of prefixes received exceeds the maximum number configured, the BGP session is terminated, which is the default behavior, after sending a cease notification to the neighbor. The session is down until a manual clear is performed by the user. The session can be resumed by using the **clear bgp** command. It is possible to configure a period after which the session can be automatically brought up by using the **maximum prefix** command with the **restart** keyword. The maximum prefix limit can be configured by the user.



Note Maximum-Prefix feature on dynamic neighbors is only supported when the Persistent Dynamic Neighbors feature is configured. For more information refer [Persistent Dynamic Neighbor](#)

Maximum-Prefix logging events

In earlier versions of IOS-XE, the logging of maximum-prefix warnings was limited to one warning per log type within a 60-second time window, regardless of the specific neighbor triggering the warning. This means that if multiple neighbors exceeded the maximum-prefix limit within a short time frame, only the first warning was logged, and subsequent warnings were considered time-limited and not logged individually. From Cisco IOS XE 17.13.1a, the enhancement ensures the logging of maximum-prefix warnings are now time-limited per-neighbor within a 60-second time window.

BGP Maximum Prefix-Discard Extra

An option to discard extra is added to the maximum-prefix configuration. Configuring the discard extra option drops all excess prefixes received from the neighbor when the prefixes exceed the configured maximum value. This drop does not, however, result in session flap.

The benefits of discard extra option are:

- Limits the memory footprint of BGP.
- Stops the flapping of the peer if the paths exceed the set limit.

On the same lines, the following describes the actions when the maximum prefix value is changed:

- If the maximum value alone is changed, a route-refresh message is sourced, if applicable.
- If the new maximum value is greater than the current prefix count state, the new prefix states are saved.
- If the new maximum value is less than the current prefix count state, then some existing prefixes are deleted to match the new configured state value.

Restrictions

- When the router drops prefixes, it is inconsistent with the rest of the network, resulting in possible routing loops.
- If prefixes are dropped, the standby and active BGP sessions may drop different prefixes. Consequently, an NSR switchover results in inconsistent BGP tables.
- The discard extra configuration cannot co-exist with the soft reconfig configuration.
- There is currently no way to control which prefixes are deleted.
- A peer may withdraw prefixes after some prefixes have been discarded. This may result in having discarded prefixes and still be below the prefix limit. To recover discarded prefixes up to the prefix limit, users may perform a soft clear on the neighbor.
- All maximum-prefix sub-options are mutually exclusive, only one can be configured at a time for a given neighbor.

Configuring Discard Extra

Perform this task to configure BGP maximum-prefix discard extra.

SUMMARY STEPS

1. **enable**
2. **configure terminal**
3. **router bgp***autonomous-system-number*
4. **neighbor***ip-address***remote-as***autonomous-system-number*
5. **address-family***address family*
6. **neighbor***ip-address***activate**
7. **neighbor***ip-address***maximum-prefix****prefix-limit** [*threshold*] [**discard-extra**] [**restart***minutes*] [**warning-only**]
8. **end**

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.
Step 3	router bgp <i>autonomous-system-number</i> Example: Router(config)# router bgp 3	Enters router configuration mode for the specified routing process.
Step 4	neighbor <i>ip-address</i> remote-as <i>autonomous-system-number</i> Example: Router(config-router)# neighbor 10.10.10.2 remote-as 2	Places the router in neighbor configuration mode for BGP routing and configures the neighbor IP address as a BGP peer.
Step 5	address-family <i>address family</i> Example: Router(config-router)# address-family ipv4 unicast	Specifies the address family and enters address family configuration submenu.
Step 6	neighbor <i>ip-address</i> activate Example: Router(config-router-af)# neighbor 10.10.10.2 activate	Enables the neighbor to exchange prefixes for the given address family with the local router.

	Command or Action	Purpose
Step 7	neighborip-addressmaximum-prefixprefix-limit [threshold] [discard-extra] [restartminutes] [warning-only] Example: Router(config-router-af)# neighbor 10.10.10.2 maximum-prefix 3 discard-extra	Configures a limit to the number of prefixes allowed. Configures discard extra paths to discard extra paths when the maximum prefix limit is exceeded.
Step 8	end	Prompts user to take one of these actions: <ul style="list-style-type: none"> • Yes - Saves configuration changes and exits the configuration session. • No - Exits the configuration session without committing the configuration changes. • Cancel - Remains in the configuration session, without committing the configuration changes

Configuration Examples for Discard Extra

The following example shows how to configure BGP maximum-prefix discard extra feature for the IPv4 address family:

```
router bgp 3
neighbor 10.10.10.2 remote-as 2
address-family ipv4 unicast
neighbor 10.10.10.2 activate
neighbor 10.10.10.2 maximum-prefix 3 discard-extra
```

Verifying Discard Extra

The following **show ip bgp neighbor A.B.C.D** command displays the information about the number of prefixes that were discarded.

```
Device #show ip bgp neighbors 10.10.10.2
BGP neighbor is 10.10.10.2, remote AS 2, external link
...
For address family: IPv4 Unicast
Session: 10.10.10.2
BGP table version 1, neighbor version 1/0
Output queue size : 0
Index 3, Advertise bit 0
3 update-group member
Outbound path policy configured
Route map for outgoing advertisements is PEER_OUT
Slow-peer detection is disabled
Slow-peer split-update-group dynamic is disabled
Prefix activity: ---- ----
Prefixes Current: 0 4 (Consumes 544 bytes)
Prefixes Total: 0 4
Implicit Withdraw: 0 0
Explicit Withdraw: 0 0
Used as bestpath: n/a 0
```

```

Used as multipath: n/a 0
Used as secondary: n/a 0

Local Policy Denied Prefixes: -----
Total:                          0      1
Maximum prefixes allowed 3 (discard-extra)
Threshold for warning message 75%
Prefixes discarded: 1

```

Information About BGP Neighbor Session Restart After Max-Prefix Limit Reached

Prefix Limits and BGP Peering Sessions

Use the **neighbor maximum-prefix** command to limit the maximum number of prefixes that a device running BGP can receive from a peer. When the device receives too many prefixes from a peer and the maximum-prefix limit is exceeded, the peering session is disabled or brought down. The session stays down until the network operator manually brings the session back up by entering the **clear ip bgp** command, which clears stored prefixes.

BGP Neighbor Session Restart with the Maximum Prefix Limit

The **restart** keyword was added to the **neighbor maximum-prefix** command so that a network operator can configure a device to automatically reestablish a BGP neighbor peering session when the peering session has been disabled or brought down. The time interval at which peering can be reestablished automatically is configurable. The *restart-interval* for the **restart** keyword is specified in minutes; range is from 1 to 65,535 minutes.

Subcodes for BGP Cease Notification

Border Gateway Protocol (BGP) imposes maximum limits on the maximum number of prefixes that are accepted from a peer for a given address family. This limitation safeguards the device from resource depletion caused by misconfiguration, either locally or on the remote neighbor. To prevent a peer from flooding BGP with advertisements, a limit is placed on the number of prefixes that are accepted from a peer for each supported address family. The default limits can be overridden through configuration of the maximum-prefix limit command for the peer for the appropriate address family.

The following subcodes are supported for the BGP cease notification message:

- Maximum number of prefixes reached
- Administrative shutdown
- Peer de-configured
- Administrative reset

A cease notification message is sent to the neighbor and the peering with the neighbor is terminated when the number of prefixes received from the peer for a given address family exceeds the maximum limit (either set by default or configured by the user) for that address family. It is possible that the maximum number of

prefixes for a neighbor for a given address family has been configured after the peering with the neighbor has been established and a certain number of prefixes have already been received from the neighbor for that address family. A cease notification message is sent to the neighbor and peering with the neighbor is terminated immediately after the configuration if the configured maximum number of prefixes is fewer than the number of prefixes that have already been received from the neighbor for the address family.

How to Configure a Device to Reestablish a Neighbor Session After the Maximum Prefix Limit Has Been Exceeded

Configuring a Router to Reestablish a Neighbor Session After the Maximum Prefix Limit Reached

Perform this task to configure the time interval at which a BGP neighbor session is reestablished by a device when the number of prefixes that have been received from a BGP peer has exceeded the maximum prefix limit.

The network operator can configure a device running BGP to automatically reestablish a neighbor session that has been brought down because the configured maximum-prefix limit has been exceeded. No intervention from the network operator is required when this feature is enabled.



Note This task attempts to reestablish a disabled BGP neighbor session at the configured time interval that is specified by the network operator. However, the configuration of the restart timer alone cannot change or correct a peer that is sending an excessive number of prefixes. The network operator will need to reconfigure the maximum-prefix limit or reduce the number of prefixes that are sent from the peer. A peer that is configured to send too many prefixes can cause instability in the network, where an excessive number of prefixes are rapidly advertised and withdrawn. In this case, the **warning-only** keyword of the **neighbor maximum-prefix** command can be configured to disable the restart capability while the network operator corrects the underlying problem.

SUMMARY STEPS

1. **enable**
2. **configure terminal**
3. **router bgp** *autonomous-system-number*
4. **neighbor** {*ip-address* | *ipv6-address* | *peer-group-name*} **peer-group**
5. **neighbor** {*ip-address* | *ipv6-address*% | *peer-group-name*} **peer-group** *peer-group-name*
6. **neighbor** {*ip-address* | *ipv6-address*% | *peer-group-name*} **remote-as** *autonomous-system-number* [**alternate-as** *autonomous-system-number*...]
7. **neighbor** {*ip-address* | *ipv6-address*% | *peer-group-name*} **remote-as** *autonomous-system-number* [**alternate-as** *autonomous-system-number*...]
8. **neighbor** {*ip-address* | *ipv6-address*% | } **maximum-prefix** *maximum* [*threshold*] [**restart** *minutes*] [**warning-only**]
9. **end**
10. **show ip bgp neighbors** *ip-address*

DETAILED STEPS

	Command or Action	Purpose
Step 1	enable Example: <pre>Device> enable</pre>	Enables privileged EXEC mode. <ul style="list-style-type: none"> • Enter your password if prompted.
Step 2	configure terminal Example: <pre>Device# configure terminal</pre>	Enters global configuration mode.
Step 3	router bgp <i>autonomous-system-number</i> Example: <pre>Device(config)# router bgp 101</pre>	Enters router configuration mode and creates a BGP routing process.
Step 4	neighbor {<i>ip-address</i> <i>ipv6-address</i> <i>peer-group-name</i>} peer-group Example: <pre>Device(config-router)# neighbor internal peer-group</pre>	Creates a BGP or multiprotocol BGP peer group.
Step 5	neighbor {<i>ip-address</i> <i>ipv6-address%</i> <i>peer-group-name</i>} peer-group <i>peer-group-name</i> Example: <pre>Device(config-router)# neighbor 10.4.9.5 peer-group internal</pre>	Configures a BGP neighbor to member of a peer group. <ul style="list-style-type: none"> • % keyword is the IPv6 link-local address identifier. This keyword needs to be added whenever a link-local IPv6 address is used outside the context of its interface.
Step 6	neighbor {<i>ip-address</i> <i>ipv6-address%</i> <i>peer-group-name</i>} remote-as <i>autonomous-system-number</i> [alternate-as <i>autonomous-system-number...</i>] Example: <pre>Device(config-router)# neighbor internal remote-as 100</pre>	Adds a peer group to the BGP or multiprotocol BGP neighbor table.
Step 7	neighbor {<i>ip-address</i> <i>ipv6-address%</i> <i>peer-group-name</i>} remote-as <i>autonomous-system-number</i> [alternate-as <i>autonomous-system-number...</i>] Example: <pre>Device(config-router)# neighbor 10.4.9.5 remote-as 100</pre>	Adds an entry to the BGP or multiprotocol BGP neighbor table.

	Command or Action	Purpose
Step 8	neighbor { <i>ip-address</i> <i>ipv6-address%</i> } maximum-prefix <i>maximum</i> [<i>threshold</i>] [restart <i>minutes</i>] [warning-only] Example: Device(config-router)# neighbor 10.4.9.5 maximum-prefix 1000 90 restart 60	Configures the maximum-prefix limit on a router that is running BGP. <ul style="list-style-type: none"> Use the restart keyword and <i>minutes</i> argument to configure the router to automatically reestablish a neighbor session that has been disabled because the maximum-prefix limit has been exceeded. The configurable range of <i>minutes</i> is from 1 to 65535 minutes. Use the warning-only keyword to configure the device to disable the restart capability to allow you to adjust a peer that is sending too many prefixes. <p>Note If the <i>minutes</i> argument is not configured, the disabled session will stay down after the maximum-prefix limit is exceeded. This is the default behavior.</p>
Step 9	end Example: Device(config-router)# end	Exits router configuration mode and enters privileged EXEC mode.
Step 10	show ip bgp neighbors <i>ip-address</i> Example: Device# show ip bgp neighbors 10.4.9.5	(Optional) Displays information about the TCP and BGP connections to neighbors. <ul style="list-style-type: none"> In this example, the output from this command will display the maximum prefix limit for the specified neighbor and the configured restart timer value.

Examples

The following sample output from the **show ip bgp neighbors** command verifies that a device has been configured to automatically reestablish disabled neighbor sessions. The output shows that the maximum prefix limit for neighbor 10.4.9.5 is set to 1000 prefixes, the restart threshold is set to 90 percent, and the restart interval is set at 60 minutes.

```
Device# show ip bgp neighbors 10.4.9.5

BGP neighbor is 10.4.9.5, remote AS 101, internal link
  BGP version 4, remote router ID 10.4.9.5
  BGP state = Established, up for 2w2d
  Last read 00:00:14, 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:                  1          1
Notifications:         0          0
Updates:                0          0
Keepalives:            23095      23095
Route Refresh:         0          0
Total:                  23096      23096
Default minimum time between advertisement runs is 5 seconds
For address family: IPv4 Unicast
BGP table version 1, neighbor versions 1/0 1/0
Output queue sizes : 0 self, 0 replicated
Index 2, Offset 0, Mask 0x4
Member of update-group 2

Prefix activity:
Sent      Rcvd
----      ----
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
!Configured maximum number of prefixes and restart interval information!
Maximum prefixes allowed 1000
Threshold for warning message 90%, restart interval 60 min
Number of NLRI in the update sent: max 0, min 0
Connections established 1; dropped 0
Last reset never
Connection state is ESTAB, I/O status: 1, unread input bytes: 0
Local host: 10.4.9.21, Local port: 179
Foreign host: 10.4.9.5, Foreign port: 11871
Enqueued packets for retransmit: 0, input: 0  mis-ordered: 0 (0 bytes)
Event Timers (current time is 0x5296BD2C):
Timer      Starts      Wakeups      Next
Retrans     23098         0          0x0
TimeWait     0           0          0x0
AckHold     23096      22692       0x0
SendWnd      0           0          0x0
KeepAlive    0           0          0x0
GiveUp       0           0          0x0
PmtuAger     0           0          0x0
DeadWait     0           0          0x0
iss: 1900546793  snduna: 1900985663  sndnxt: 1900985663  sndwnd: 14959
irs: 2894590641  rcvnxt: 2895029492  rcvwnd: 14978  delrcvwnd: 1406
SRTT: 300 ms, RTTO: 607 ms, RTV: 3 ms, KRTT: 0 ms
minRTT: 0 ms, maxRTT: 316 ms, ACK hold: 200 ms
Flags: passive open, nagle, gen tcbs
Datagrams (max data segment is 1460 bytes):
Rcvd: 46021 (out of order: 0), with data: 23096, total data bytes: 438850
Sent: 46095 (retransmit: 0, fastretransmit: 0), with data: 23097, total data by9

```

Troubleshooting Tips

Use the **clear ip bgp** command to reset a BGP connection using BGP soft reconfiguration. This command can be used to clear stored prefixes to prevent a device that is running BGP from exceeding the maximum-prefix limit.

Display of the following error messages can indicate an underlying problem that is causing the neighbor session to become disabled. You should check the values configured for the **neighbor maximum-prefix**

command and the configuration of any peers that are sending an excessive number of prefixes. The following sample error messages are similar to the error messages that may be displayed:

```
00:01:14:%BGP-5-ADJCHANGE:neighbor 10.10.10.2 Up
00:01:14:%BGP-4-MAXPFX:No. of unicast prefix received from 10.10.10.2 reaches 5, max 6
00:01:14:%BGP-3-MAXPFXEXCEED:No.of unicast prefix received from 10.10.10.2:7 exceed limit6
00:01:14:%BGP-5-ADJCHANGE:neighbor 10.10.10.2 Down - BGP Notification sent
00:01:14:%BGP-3-NOTIFICATION:sent to neighbor 10.10.10.2 3/1 (update malformed) 0 byte
```

The **bgp dampening** command can be used to configure the dampening of a flapping route or interface when a peer is sending too many prefixes and causing network instability. Use this command only when troubleshooting or tuning a device that is sending an excessive number of prefixes. For more details about BGP route dampening, see the “Configuring Advanced BGP Features” module.

Configuration Example for BGP Restart Neighbor Session After Max-Prefix Limit Reached

Example: Configuring a Router to Reestablish a Neighbor Session After the Maximum Prefix Limit Reached

The following example sets the maximum number of prefixes allowed from the neighbor at 192.168.6.6 to 2000 and configures the device to reestablish a peering session after 30 minutes if one has been disabled:

```
Device(config)# router bgp 101
Device(config-router)# neighbor internal peer-group
Device(config-router)# neighbor 10.4.9.5 peer-group internal
Device(config-router)# neighbor internal remote-as 100
Device(config-router)# neighbor 10.4.9.5 remote-as 100
Device(config-router)# neighbor 10.4.9.5 maximum-prefix 2000 90 restart 30
Device(config-router)# end
```

Additional References for BGP Restart Neighbor Session After Max-Prefix Limit Reached

Related Documents

Related Topic	Document Title
Cisco IOS commands	Cisco IOS Master Command List, All Releases
BGP commands	Cisco IOS IP Routing: BGP Command Reference

Standards and RFCs

Standard/RFC	Title
RFC 2918	<i>Route Refresh Capability for BGP-4</i>
RFC 4486	<i>Subcodes for BGP Cease Notification Message</i>

Technical Assistance

Description	Link
The Cisco Support and Documentation website provides online resources to download documentation, software, and tools. Use these resources to install and configure the software and to troubleshoot and resolve technical issues with Cisco products and technologies. Access to most tools on the Cisco Support and Documentation website requires a Cisco.com user ID and password.	http://www.cisco.com/cisco/web/support/index.html

Feature Information for BGP Maximum-Prefix on IOS XE

The following table provides release information about the feature or features described in this module. This table lists only the software release that introduced support for a given feature in a given software release train. Unless noted otherwise, subsequent releases of that software release train also support that feature.

Use Cisco Feature Navigator to find information about platform support and Cisco software image support. To access Cisco Feature Navigator, go to <http://www.cisco.com/go/cfn>. An account on Cisco.com is not required.

Table 1: Feature Information for BGP Maximum-Prefix on IOS XE

Feature Name	Releases	Feature Information
BGP Restart Session After Max-Prefix Limit		<p>The BGP Restart Session After Max-Prefix Limit Reached feature adds the restart keyword to the neighbor maximum-prefix command. This allows a network operator to configure the time interval at which a peering session is reestablished by a device when the number of prefixes that have been received from a peer has exceeded the maximum prefix limit.</p> <p>The following commands were modified: neighbor maximum-prefix and show ip bgp neighbors.</p>

Feature Name	Releases	Feature Information
BGP—Subcodes for BGP Cease Notification		Support for subcodes for BGP cease notification has been added.
BGP – Maximum Prefix Discard Extra and Logging enhancement	Cisco IOS XE 17.13.1a	From IOS XE 17.13.1a, BGP Maximum Prefix feature introduces Discard Extra option. This feature drops all excess prefixes received from the neighbor when the configured value of the prefixes exceeds the maximum limit. The Maximum Prefix also introduces a per neighbor enhanced logging time every 60 seconds.