BGP Prefix Management and Session Parameters

This chapter covers BGP prefix control, path management, and session-level configurations like TCP MSS customization to optimize network performance.

BGP maximum-prefix and discard extra paths

BGP maximum-prefix is a BGP functionality that enforces a limit on the number of prefixes received from a neighbor for a specific address family and terminates the BGP session if the prefix count exceeds the configured limit.

The discard extra paths option is an enhancement to the BGP maximum-prefix functionality that drops excess prefixes received from a neighbor without flapping the session and limits the memory footprint of BGP.

BGP maximum-prefix overview

The BGP maximum-prefix capability allows you to configure an upper threshold on the number of prefixes that can be accepted from a neighbor for a particular address family. If the number of received prefixes exceeds the configured limit, the system performs these actions:

- Terminates the BGP session and sends a cease notification to the neighbor.
- Keeps the session down until you manually clear it using the **clear bgp** command.
- Restarts the session automatically after a specified period if configured with the **maximum-prefix** command and the **restart** keyword.

Starting with IOS-XR Release 7.3.1, the system no longer applies default limits unless you explicitly configure the maximum number of prefixes for the address family.

Discard extra paths behavior

The discard extra paths option modifies BGP maximum-prefixbehavior. With discard extra paths option configured, when excess prefixes are received, they are dropped, but the BGP session remains stable. If the discard extra paths configuration is removed, BGP sends a route-refresh message to the neighbor if it supports the refresh capability; otherwise, the session flaps.

Effects of changing the maximum-prefix value

Changing the maximum-prefix value triggers these specific actions:

- If the new value exceeds the current prefix count, the additional prefixes are saved.
- If the new value is less than the current prefix count, some prefixes are deleted to align with the new limit. There is no control over which prefixes are removed.

Benefits of BGP maximum-prefix and discard extra paths

BGP maximum-prefix and discard extra paths functionality provides these benefits:

- Limits the memory usage of BGP by dropping excess prefixes received beyond the configured maximum.
- Prevents session flapping, ensuring stability of the BGP peer even when the prefix limit is exceeded.
- Helps maintain operational continuity by avoiding disruptions caused by exceeding the prefix limit.

Limitations of discard extra paths

When configuring discard extra paths, consider these guidelines.

The discard extra paths configuration cannot coexist with the **soft reconfig** configuration.

If the system runs out of physical memory, the BGP process exits and requires a manual restart using the **process restart bpm** command.

- Dropped prefixes can cause network inconsistencies, potentially leading to routing loops.
- During a Non-Stop Routing (NSR) switchover, standby and active BGP sessions may drop different prefixes, causing inconsistent BGP tables.

Benefits of BGP maximum-prefix and discard extra paths

BGP maximum-prefix and discard extra paths functionality provides these benefits:

- Limits the memory usage of BGP by dropping excess prefixes received beyond the configured maximum.
- Prevents session flapping, ensuring stability of the BGP peer even when the prefix limit is exceeded.
- Helps maintain operational continuity by avoiding disruptions caused by exceeding the prefix limit.

Configure BGP maximum-prefix and discard extra paths

The purpose of this task is to configure a BGP neighbor to discard extra paths when the maximum prefix limit is exceeded, ensuring session stability and controlled memory usage.

This task involves setting up a BGP neighbor with a maximum prefix limit and enabling the discard extra paths option to drop prefixes exceeding the limit without session flapping.

Follow these steps to configure BGP maximum-prefix discard extra paths:

Before you begin

- Identify the BGP autonomous system (AS) number.
- Determine the neighbor IP address and the maximum prefix limit to configure.

Procedure

Step 1 Enter configuration mode and specify the BGP autonomous system number to enable BGP configuration.

Example:

```
Router# configure
Router(config)# router bgp 10
router(config-bgp)#
```

Step 2 Define the neighbor IP address and enter the address family submode to configure the specific type of traffic.

Example:

```
Router(config-bgp)# neighbor 10.0.0.1
Router(config-bgp-nbr)# address-family ipv4 unicast
```

Step 3 Configure the maximum prefix limit and enable the discard extra paths option to prevent excess prefixes from causing disruptions.

Example:

```
Router(config-bgp-nbr-af)# maximum-prefix 1000 discard-extra-paths
```

Step 4 Save the configuration to apply changes.

Example:

```
Router(config-bgp-nbr-af)# commit
```

Step 5 Verify the running configuration on the system.

Example:

```
Router# show running-config
router bgp 10
 neighbor 10.0.0.1
 address-family ipv4 unicast
   maximum-prefix 1000 discard-extra-paths
   !
  !
```

Step 6 Use the **show bgp neighbor** command to view the neighbor configuration and status.

Review the output for these details:

- Maximum prefixes allowed
- Discard extra paths configuration
- Threshold for warning messages

Example:

```
Router# show bgp neighbor 10.0.0.1
BGP neighbor is 10.0.0.1
```

```
Maximum prefixes allowed 1000 (discard-extra-paths)
Threshold for warning message 75%
```

The configuration ensures that the BGP neighbor limits the number of prefixes to 1,000 and discards extra paths without flapping the session.

Summary of key commands for BGP maximum-prefix and discard extra paths

Table 1: Key commands

Command	Description
maximum-prefix <value> discard-extra-paths	Configures the maximum prefix limit and enables the discard extra paths option.
show bgp neighbor <neighbor-ip>	Displays the BGP neighbor's configuration and status, including discard extra paths details.
process restart bpm	Manually restarts the BGP process when it exits due to insufficient physical memory.

TCP maximum segment size

Maximum Segment Size (MSS) is a TCP attribute that

- determines the largest amount of data that a device can receive in a single, unfragmented TCP segment
- is limited by the Maximum Transmission Unit (MTU) of an interface, and
- is negotiated during the TCP setup process between a source and destination.

The MSS ensures efficient data transfer by optimizing the size of transmitted packets, especially for protocols like BGP. Each direction of data flow can use a different MSS value based on the MTU of the source and destination interfaces.

Key attributes of MSS

These are some of the key attributes of MSS:

- The closer the MSS is to the MTU, the more efficient the data transfer.
- The MSS is announced during the TCP setup process.
- The default TCP MSS value is 536 octets or 1,460 bytes. This means that TCP segments the data in the transmit queue into 1460-byte chunks before passing the packets to the IP layer.

Per neighbor TCP MSS

Per neighbor TCP MSS is a mechanism in BGP configuration that

- allows creating unique TCP MSS profiles for each neighbor

- supports configuration in two modes: neighbor group and session group, and
- overrides the global TCP MSS setting for specific neighbors.

Key attributes of per neighbor TCP MSS

These are some of the key attributes of per neighbor TCP MSS:

- You can enable or disable TCP MSS configuration for specific neighbors.
- MSS value can be reset to its default using the **inheritance-disable** command.
- The configuration range for MSS values is from 68 to 10,000.

Configure per neighbor TCP MSS

The purpose of this task is to configure a TCP MSS value for a specific neighbor in BGP.

Before you begin

Identify the desired MSS value.

Procedure

Step 1 Enter BGP configuration mode and set up the neighbor group.

Example:

```
Router# configure
Router#(config)# router bgp 10
Router#(config-bgp)# address-family ipv4 unicast
Router#(config-bgp-af)# exit
Router#(config-bgp)# neighbor-group n1
Router#(config-bgp-nbrgrp)# tcp mss 500
Router#(config-bgp-nbrgrp)# address-family ipv4 unicast
Router#(config-bgp-nbrgrp-af)# exit
Router#(config-bgp-nbrgrp)# exit
```

Step 2 Configure a specific neighbor and inherit settings from the neighbor group.

Example:

```
Router#(config-bgp)# neighbor 10.0.0.2
Router#(config-bgp-nbr)# remote-as 1
Router#(config-bgp-nbr)# use neighbor-group n1
Router#(config-bgp-nbr)# address-family ipv4 unicast
Router#(config-bgp-nbr-af)#
```

Step 3 Save the configuration.

Example:

```
Router#(config-bgp-nbr-af)# commit
```

Step 4 Verify the running configuration on the system.

Example:

```
Router# show running-config
router bgp 10
 address-family ipv4 unicast
 !
 neighbor-group n1
  tcp mss 500
  address-family ipv4 unicast
  !
 !
 neighbor 10.0.0.2
  remote-as 1
  use neighbor-group n1
  address-family ipv4 unicast
  !
 !
!
end
```

Step 5 Use the **show bgp neighbor** command to view the neighbor configuration and status.

Example:

```
Router# show bgp neighbor 10.0.0.2
BGP neighbor is 10.0.0.2
 Remote AS 1, local AS 10, external link
 Remote router ID 0.0.0.0
  BGP state = Idle (No best local address found)
  ...
  Minimum time between advertisement runs is 30 secs
  Configured TCP Maximum Segment Size 500
  Inbound message logging enabled, 3 messages buffered
  Outbound message logging enabled, 3 messages buffered
  ...
 For Address Family: IPv4 Unicast
  BGP neighbor version 0
  Update group: 0.1 Filter-group: 0.0  No Refresh request being processed
  eBGP neighbor with no inbound or outbound policy; defaults to 'drop'
  ...
```

Step 6 Use the **show tcp brief** command to check TCP connection endpoints.

Example:

```
Router# show tcp brief
   PCB             VRF-ID     Recv-Q Send-Q Local Address          Foreign Address        State
0x000055e27958c800 0x60000000      0      0  :::179                :::0                    LISTEN
0x000055e27958b850 0x00000000      0      0  :::179                :::0                    LISTEN
0x00007f2a80002050 0x60000000      0      0  0.0.0.0:179           0.0.0.0:0               LISTEN
0x00007f2a840380d0 0x00000000      0      0  0.0.0.0:179           0.0.0.0:0               LISTEN
```

This information helps verify the correct configuration and troubleshoot connectivity issues.

Use the **show tcp** command to view detailed TCP connection information.

The TCP MSS is configured for the specified neighbor.

Disable the per neighbor TCP MSS

The purpose of this task is to disable the TCP MSS configuration for a specific neighbor.

Follow these steps to disable the per neighbor TCP MSS:

Before you begin

Ensure the neighbor group or session group is already configured.

Procedure

Step 1 Enter BGP configuration mode.

Example:

```
Router# configure
Router#(config)# router bgp 10
Router#(config-bgp)#
```

Step 2 Disable MSS inheritance for the neighbor group.

Example:

```
Router#(config-bgp)# address-family ipv4 unicast
Router#(config-bgp-af)# exit
Router#(config-bgp)# neighbor-group n1
Router#(config-bgp-nbrgrp)# tcp mss inheritance-disable
Router#(config-bgp-nbrgrp)# address-family ipv4 unicast
Router#(config-bgp-nbrgrp-af)# exit
Router#(config-bgp-nbrgrp)# exit
```

Step 3 Configure a specific neighbor and disable MSS inheritance for the neighbor.

Example:

```
Router#(config-bgp)# neighbor 10.0.0.2
Router#(config-bgp-nbr)# remote-as 1
Router#(config-bgp-nbr)# use neighbor-group n1
Router#(config-bgp-nbr)# tcp mss inheritance-disable
Router#(config-bgp-nbr)# commit
```

Step 4 Verify the running configuration on the system.

Example:

```
Router# show running-config
router bgp 10
 address-family ipv4 unicast
 !
 neighbor-group n1
  tcp mss inheritance-disable
  address-family ipv4 unicast
  !
 !
 neighbor 10.0.0.2
  remote-as 1
  use neighbor-group n1
  tcp mss inheritance-disable
  address-family ipv4 unicast
  !
 !
!
```

TCP MSS is disabled for the specified neighbor.

Summary of key commands for per neighbor TCP MSS

Table 2: Key commands

Command	Description
`show bgp neighbor`	Displays BGP neighbor details, including the configured MSS value.
`show tcp brief`	Lists active TCP connections and their states.
`show tcp pcb <pcb-value>`	Provides detailed TCP connection information for a specific PCB.