Handling BGP Slow Peers

This chapter covers strategies for managing BGP slow peers, including detection, isolation from update groups, and eBGP session resets in response to link failures.

BGP slow peer management

BGP slow peer management is a Border Gateway Protocol (BGP) mechanism that

- groups neighbors into update groups based on address-family configuration and shared update content
- formats updates once per group, transmits them to all members, and retains updates until all members acknowledge receipt, and
- mitigates the impact of slow peers to prevent group-wide backlogs and maintain update throughput.

Table 1: Feature History Table

Feature Name	Release	Description
BGP Slow Peer	Release 7.9.1	BGP neighbors are grouped together to optimize update generation. BGP peers process the incoming BGP update messages at different rates. A slow peer is a peer that is processing incoming BGP update messages very slowly over a long period of time compared to other peers in the update sub-group. BGP slow peer handling is necessary to reduce the impact of the slow peer on the remaining members of the group. This feature introduces the following commands: • **slow peer (BGP neighbor address-family configuration)** • **slow peer (BGP neighbor address-family configuration)** This feature modifies the following commands: • **show bgp neighbors** to display the slow peer configuration state and slow peer detection or processing information. • **show bgp update out** to display the summary of the neighbor address-family update-group, sub-group, or refresh sub-group information

How update groups work

- BGP neighbors are grouped using common criteria, such as the neighbor address-family configuration and processing of the same update messages, to optimize update generation.
- For each group, messages are formatted once and transmitted to all members; messages are deleted only after all members acknowledge receipt.
- Update generation operates per address family. A peer refers to a neighbor address family, and peers in a sub-group are neighbors for the same address family.

Causes of slow peers

- Processor is busy handling other tasks and cannot process updates on time.
- Peers are connected over slow bandwidth links.
- Temporary network congestion affects timely processing.

Effects of slow peers on update groups

- BGP enforces limits on update messages per process, per address family, and per sub-group. The default sub-group message limit is 32 Mbytes.
- If one or more peers are extremely slow to process messages, those messages remain queued until acknowledged by all peers in the sub-group.
- When the sub-group message limit is reached, all neighbors in the sub-group must wait for the slowest peer to catch up, which slows every member of the sub-group.

Slow peer management reduces these impacts by isolating or otherwise mitigating the effects of slow processing peers.

Slow peer types and states

BGP distinguishes configuration types from runtime states for peers that process updates slowly. Understanding both helps you choose the right mitigation and verify behavior in operation.

Slow peer types

- Static slow peer: Moves the peer to its own update group and requires no additional slow-peer handling.
- Dynamic slow peer: Keeps the peer in the original group; when detected as slow, processing occurs in a refresh sub-group.

Slow peer states

- Static slow peer: The neighbor address family is in the static slow peer state.
- Dynamic detected slow peer: The neighbor address family is detected as slow and is being processed.
- Not slow peer: The neighbor address family is not currently slow.

Table 2: Type-to-state alignment

Type	Typical runtime state	Operational behavior
Static	Static slow peer	Isolated in its own update group; no additional slow-peer handling.
Dynamic	Dynamic detected slow peer or Not slow peer	Remains grouped; when detected slow, updates are handled in a refresh sub-group.
None	Not slow peer	No slow-peer handling.

Guidelines for managing slow peers

Queue and update hygiene

- Ensure queues do not retain stale information; prioritize sending only the latest route state during periods of sustained route churn.
- Monitor sub-group backlogs and message limits to identify slow peers early and take corrective action, for example, apply slow-peer handling features.
- Maintain per–address-family hygiene to prevent a single slow peer from degrading update throughput for the entire sub-group.

Handling permanently slow peers

- Do not rely on dynamic slow peer detection for permanently slow peers.
- Isolate permanently slow peers using static slow peer configuration, which moves them to their own update group; apply the same route policy to all such peers.

How slow peer detection and processing works

Summary

The key components involved in the process are:

- Update groups and sub-groups: Neighbors are grouped by address-family and shared update content; messages are formatted once and deleted only after all members acknowledge receipt.
- Slow peer classification: Static or dynamic, with operational states tracked per neighbor address family.
- Refresh sub-group: A child group created to process a slow peer's updates up to the current table version while the parent sub-group continues with new updates.
- Queues: Main queue and parallel slow peer queue used to control advertisement flow and maintain per-route ordering.
- Detection thresholds and limits: Time threshold (default 300 seconds) and a maximum of 16 refresh sub-groups per address family.

Slow peer handling mitigates backlog and delay in BGP update groups by detecting peers that cannot keep up, isolating their processing in refresh sub-groups, and managing queues to preserve ordering and throughput.

Workflow

These stages describe how slow peer detection and processing works:

1. Group formation and baseline behavior: BGP neighbors are grouped by address-family configuration and shared updates; messages are formatted once and retained until all members acknowledge them. This optimizes update generation but can stall on slow peers.
2. Slow peer detection: A dynamic peer is identified as slow when all of these are true:
    - acknowledgments are missing for some messages

- pending messages are yet to be written to TCP
- the time since the last update exceeds the threshold
- the number of refresh sub-groups for the address family is fewer than 16
- the neighbor is not the only member of the sub-group; other members are already marked slow, and
- there are nets awaiting update generation.

3. Refresh sub-group creation: The system creates a refresh sub-group for the detected slow peer to process updates up to the current table version. The parent sub-group continues sending new updates for both slow and non-slow peers. Each slow peer gets its own refresh sub-group. When processing completes, the refresh sub-group is removed.

   See Refresh sub-group states for more information.

4. Queue management: The router moves all messages queued in the peer's main queue to a parallel slow peer queue and advertises them separately, while new messages continue to flow from the main queue. If the peer is detected slow again, the move-and-advertise cycle repeats. Ordering for updates and withdrawals is maintained per route.

   See Detection and queue management details for more information.

5. State tracking and clearing: The *processing slow peer* state is set to true when handling starts and is cleared only after all slow peer updates are advertised and acknowledged; it does not clear on configuration changes.

6. Recovery: The peer is no longer considered slow when all messages in the slow peer queue are advertised and acknowledged. The slow peer queue is deleted, and the refresh sub-group is removed.

Result

Update generation remains responsive and fair; slow peers are isolated without blocking faster peers, preventing sub-group stalls and preserving per-route ordering during churn.

Configure slow peer handling for BGP neighbors

Enable and verify slow peer handling globally or per neighbor address family to mitigate update backlogs and isolate permanently slow peers.

Slow peer handling supports

- global configuration that affects all neighbors and
- per–neighbor address-family configuration that targets a specific peer.

Dynamic detection can optionally use a threshold (seconds) for classification.

Before you begin

- Identify the BGP autonomous system number.
- Determine whether you need a global setting, a per–neighbor address-family setting, or both.
- If using dynamic detection, choose the detection threshold (seconds).

Follow these steps to configure slow peer handling:

Procedure

Step 1 Configure global slow peer handling.

- Enable dynamic slow peer handling for all BGP neighbor address families.

```
Router#configure
Router(config)#router bgp 100
Router(config-bgp)#slow-peer dynamic
Router(config-bgp)#commit
```

- Disable slow peer handling for all BGP neighbor address families.

```
Router#configure
Router(config)#router bgp 100
Router(config-bgp)#slow-peer detection-disable
Router(config-bgp)#commit
```

- Enable dynamic slow peer handling with a detection threshold of 120 seconds.

```
Router#configure
Router(config)#router bgp 100
Router(config-bgp)#slow-peer dynamic threshold 120
Router(config-bgp)#commit
```

Step 2 Configure slow peer handling for a specific neighbor address family.

- Mark a neighbor as a static slow peer (isolates the neighbor in its own update group).

```
Router#configure
Router(config)#router bgp 100
Router(config-bgp)#neighbor 50.0.0.1
Router(config-bgp-nbr)#address-family ipv4 unicast
Router(config-bgp-nbr-af)#slow-peer static
Router(config-bgp-nbr-af)#commit
```

- Disable dynamic slow peer handling for a neighbor address family.

```
Router#configure
Router(config)#router bgp 100
Router(config-bgp)#neighbor 50.0.0.1
Router(config-bgp-nbr)#address-family ipv4 unicast
Router(config-bgp-nbr-af)#slow-peer dynamic disable
Router(config-bgp-nbr-af)#commit
```

- Enable dynamic slow peer handling for a neighbor address family.

```
Router#configure
Router(config)#router bgp 100
Router(config-bgp)#neighbor 50.0.0.1
Router(config-bgp-nbr)#address-family ipv4 unicast
Router(config-bgp-nbr-af)#slow-peer dynamic
Router(config-bgp-nbr-af)#commit
```

- Enable dynamic slow peer handling with a detection threshold of 120 seconds for a neighbor address family.

```
Router#configure
Router(config)#router bgp 100
Router(config-bgp)#neighbor 50.0.0.1
Router(config-bgp-nbr)#address-family ipv4 unicast
```

```
Router(config-bgp-nbr-af)#slow-peer dynamic threshold 120
Router(config-bgp-nbr-af)#commit
```

Step 3 Run the **show bgp neighbors <neighbor-address> detail** command to view the effective slow peer configuration for a neighbor address family.

Slow peer handling is enabled globally or per neighbor address family as configured. Use the verification command to confirm the effective state for each neighbor.

Slow peer effective configuration state

This table summarizes the effective neighbor AF slow peer configuration or operational state, considering both the slow peer global configuration and the slow peer neighbor AF configuration.

- For example, if the global configuration is *None* and the neighbor configuration is *Static*, then the effective configuration is *Static*.

Table 3: Effective slow peer configuration state

-		Global configuration		
-		[None]	[Dynamic]	[Detection disable]
Neighbor address-family configuration	[None]	Detection-only	Dynamic	None
	[Static]	Static	Static	Static
	[Dynamic]	Dynamic	Dynamic	Dynamic
	[Dynamic Disable]	Detection-only	None	None

The effective neighbor address-family configuration state can be any of the following entries in this table.

- The **show bgp neighbors <neighbor-address> detail** command displays the neighbor address-family configuration states listed here.

Table 4: Effective neighbor address-family configuration state

AF configuration state	Details
Static	• When the effective neighbor address family configuration is Static, the neighbor address family moves to its own update group, isolating it from other neighbors. • To place all slow peers in a single update group, remove the static slow peer configuration and apply the same outbound route policy to all neighbors.

AF configuration state	Details
Dynamic	• When the effective neighbor address family configuration is Dynamic, BGP enables dynamic slow peer processing for that neighbor address family. • If it is detected as slow, BGP processes the neighbor address family in a dedicated refresh sub-group, isolates it from other neighbors in the sub-group, and displays an IOS message indicating the slow state.
Detection-only	• When the effective neighbor address family configuration is Detection-only, BGP logs slow-peer detection and recovery events but applies no mitigation. • The router displays an IOS message when the neighbor address family becomes slow or recovers.
None	When the effective neighbor address family configuration is None, BGP disables slow peer handling for that neighbor address family.

Examples: Configure slow peer handling with combined global and neighbor settings

- Enable dynamic slow peer globally; mark one neighbor as static.

```
Router#configure
Router(config)#router bgp 100
Router(config-bgp)#slow-peer dynamic
Router(config-bgp)#neighbor 50.0.0.1
Router(config-bgp-nbr)#address-family ipv4 unicast
Router(config-bgp-nbr-af)#slow-peer static
Router(config-bgp-nbr-af)#commit
```

- Enable dynamic slow peer globally; disable it for one neighbor address family.

```
Router#configure
Router(config)#router bgp 100
Router(config-bgp)#slow-peer dynamic
Router(config-bgp)#neighbor 50.0.0.1
Router(config-bgp-nbr)#address-family ipv4 unicast
Router(config-bgp-nbr-af)#slow-peer dynamic disable
Router(config-bgp-nbr-af)#commit
```

- Use different dynamic detection thresholds globally and per neighbor.

```
Router#configure
Router(config)#router bgp 100
Router(config-bgp)#slow-peer dynamic threshold 600
Router(config-bgp)#neighbor 50.0.0.1
Router(config-bgp-nbr)#address-family ipv4 unicast
Router(config-bgp-nbr-af)#slow-peer dynamic threshold 120
Router(config-bgp-nbr-af)#commit
```

IOS messages for slow peer events

The system logs messages when a BGP neighbor is detected as a slow peer and when it recovers. These are the available events and corresponding log messages:

- Slow peer detected: `BGP neighbor 50.0.0.1 of vrf default afi IPv4 Unicast is detected as slow-peer`
- Slow peer recovered: `Slow BGP peer 50.0.0.1 of vrf default afi IPv4 Unicast has recovered`

BGP slow peer automatic isolation from update group

BGP slow peer automatic isolation from update group is a BGP feature that

- detects neighbors in an update group that cannot keep up with sustained update generation over time
- automatically moves detected slow peers into a dedicated slow update group and returns them to the original group upon recovery, and
- prevents group-wide stalls by isolating slow peers, allowing non-slow members to continue processing new updates.

Table 5: Feature History Table

Feature Name	Release Information	Feature Description
BGP Slow Peer Automatic Isolation from Update Group	Release 25.1.1	Introduced in this release on: Fixed Systems (8700 [ASIC: K100])(select variants only*) *This feature is supported on the Cisco 8712-MOD-M routers.
BGP Slow Peer Automatic Isolation from Update Group	Release 24.4.1	Introduced in this release on: Fixed Systems (8200 [ASIC: P100], 8700 [ASIC: P100])(select variants only*); Modular Systems (8800 [LC ASIC: P100])(select variants only*) *This feature is supported on: • 88-LC1-36EH • 88-LC1-12TH24FH-E • 88-LC1-52Y8H-EM • 8212-48FH-M • 8711-32FH-M

BGP Slow Peer Automatic Isolation from Update Group	Release 7.3.1	A slow peer cannot keep up with the rate at which the router generates BGP update messages over a period of time, in an update group. This feature automatically detects a slow peer in an update group and moves it to a new update group. The feature is enabled on the router, by default. New commands: • **slow-peer detection enable** • **clear bgp slow-peers** Updated commands: • **slow-peer detection disable**

Slow update groups: structure and defaults

- Short-lived churn: A peer that briefly falls behind during events such as connection resets but quickly recovers is not treated as a slow peer.
- Impact of slow peers: The presence of a slow peer increases the number of formatted updates pending transmission in the original update group.
- Group structure:
  - One slow update group exists for each original update group that contains slow peers.
  - Multiple slow peers in the same original group are isolated in separate sub-groups within that slow update group.
  - Slow peers from different original update groups cannot be combined because outbound policy configurations differ.
- Defaults: The feature is enabled by default, and automatic splitting of update groups is enabled by default.

How automatic isolation works

Summary

The key components involved in the process are:

- Original update group: Formats messages once and transmits to all members; deletes messages only after all members acknowledge receipt.
- Slow update group: Dedicated group created when one or more peers in the original group are detected as slow.
- Slow peer sub-groups: Per-peer sub-groups within the slow update group used to process each slow peer independently.
- Recovery action: Moves peers back to the original update group when they catch up.

Automatic isolation mitigates backlog and contention in update groups by detecting slow peers, processing their updates separately in a dedicated slow update group, and restoring them when they recover.

Workflow

These stages describe how automatic isolation works:

1. Detection: The system detects that a peer in the original update group cannot keep pace, causing formatted messages to accumulate.
2. Isolation: The detected slow peer is moved automatically to a new slow update group; if multiple peers are slow, each is placed in its own slow sub-group.
3. Parallel processing: The parent (original) update group continues sending newly modified nets to non-slow peers, while the slow update group processes backlogged updates for slow peers.
4. Recovery: When a slow peer completes processing and catches up, it is moved back to the original update group.

Result

Non-slow peers advance at their regular pace, while slow peers process updates in isolation. This prevents stalls and reduces the risk of backlog growth across the entire group.

Configure slow peer automatic isolation

Detect, isolate, and manage slow BGP peers by configuring global and neighbor-level controls, and verify slow-peer status and queues.

Automatic isolation moves detected slow peers to a dedicated slow update group and returns them to the original group after recovery. Dynamic detection can be enabled globally or per neighbor address family; static marking isolates a specific peer immediately.

Before you begin

- Identify the BGP autonomous system number.
- Determine the neighbor address and address family (AFI/SAFI).
- Decide whether to disable global detection, mark a peer as static, or enable dynamic detection with the permanent option.

Do one or more of the following to configure and manage slow peers.

Procedure

Step 1 Disable slow peer detection globally.

Example:

```
Router# configure
Router(config)# slow-peer-detection disable
```

Any slow peers that are detected are marked as normal peers and moved back to their original update groups. No more slow peers are detected.

Step 2 Mark a neighbor as a static slow peer at the neighbor address-family level.

Example:

```
Router(config)# router bgp 5
Router(config-bgp)# address-family ipv4
Router(config-bgp-af)# neighbor 172.60.2.3
Router(config-bgp-nbr-af)# slow-peer detection disable split-updategroup static
```

The peer becomes part of the slow update group.

Step 3 Enable dynamic detection with permanent option at the neighbor address-family level.

Example:

```
Router(config)# router bgp 5
Router(config-bgp)# address-family ipv4
Router(config-bgp-af)# neighbor 172.60.2.3
Router(config-bgp-nbr-af)# slow-peer detection enable split-update-group
dynamic permanent
```

The peer is moved to a slow update group when detected slow.

- If only the **split-update-group dynamic** command is configured, a dynamically detected slow peer is moved to an existing slow update group or a new one is created. This behavior is enabled by default.
- If the *permanent* keyword is not configured, the peer returns to the original update group after recovery. If the *permanent* keyword is configured, the peer does not return automatically; use the **clear** command to move it back. Use this option if a peer keeps becoming a slow peer and recovering.

Step 4 Clear dynamically detected slow peers.

- Clear all slow peers for a specific AFI/SAFI.

  ```
  Router# clear bgp slow-peers <afi> <safi>
  ```

- Clear all slow peers for a neighbor across AFI/SAFI.

  ```
  Router# clear bgp slow-peers <neighbor-address>
  ```

- Clear a specific AFI, SAFI, and neighbor combination.

  ```
  Router# clear bgp slow-peers <afi> <safi> <neighbor-address>
  ```

Step 5 View the running configuration.

Example:

```
slow-peer-detection disable
router bgp 5
address-family ipv4
 neighbor 172.60.2.3
 slow-peer detection disable split-update-group static
router bgp 5
address-family ipv4
 neighbor 172.60.2.3
 slow-peer detection enable split-update-group dynamic permanent
```

Step 6 View slow-peers summary for neighbors.

Example:

```
show bgp update out neighbor slow-peers brief
Fri Feb 5 00:12:50.830 UTC

VRF "default", Address-family "IPv4 Unicast"
Main routing table version: 9819220
RIB version: 9819220

Neighbor FG SG SG-R UG Status OutQ OutQ-R Version

19.1.3.1 0.4 0.4 --- 0.2 Normal 4864200 0 7073474
19.1.4.1 0.4 0.4 --- 0.2 Normal 5206200 0 7073474
```

Step 7 Check slow-peers across all address families and neighbors, and compare behavior after time passes.

Example:

```
Router# show bgp all all update out neighbor slow-peers
Fri Sep 13 13:57:48.503 PDT
Address Family: IPv4 Unicast
----------------------------
```

After 5 minutes:

```
Router# show bgp all all update out neighbor slow-peers
Fri Sep 13 14:02:23.097 PDT
Address Family: IPv4 Unicast
----------------------------
VRF "default", Address-family "IPv4 Unicast"
Main routing table version: 3329832
RIB version: 3329832
Neighbor 11.11.11.21
Filter-group 0.3, Refresh filter-group ---
Sub-group 0.2, Refresh sub-group ---
Update-group 0.3
Update OutQ: 20447800 bytes (7680 messages) Refresh
update OutQ: 0 bytes (0 messages) Filter-group pending:
7680 messages
```

Slow peer detection is managed globally or per neighbor address family as configured. Detected slow peers are isolated, and recovery behavior is controlled. Use the verification commands to confirm slow-peers status, queues, and group membership.

eBGP session reset on link failure

eBGP session reset on link failure is a BGP feature that

- automatically resets sessions to directly adjacent external peers when a link goes down (fast external fallover)
- lets you disable and re-enable automatic resets using dedicated configuration commands, and
- supports high session counts by increasing packet rate with LPTS PIFIB hardware policing.

eBGP sessions can flap when a node reaches 3,500 sessions with BGP timers set to 10 and 30. Increasing the packet rate helps support more than 3,500 sessions.

Guidelines for fast external fallover

- Use immediate resets (fast external fallover enabled) when rapid failure detection and cleanup are required for adjacent external peers.
- Disable immediate resets if automatic session resets during transient link events or maintenance windows cause instability.
- When approaching or exceeding 3,500 eBGP sessions with aggressive timers (10 and 30), increase LPTS PIFIB hardware policing rates to maintain stability.

Configure fast external fallover behavior

Control whether eBGP sessions reset automatically when a link fails.

Before you begin

Decide whether immediate resets on link-down events are desired for adjacent external peers.

Procedure

Step 1 Disable automatic resets on link failure.

Example:

```
Router# bgp fast-external-fallover disable
```

Step 2 Re-enable automatic resets on link failure.

Example:

```
Router# no bgp fast-external-fallover disable
```

eBGP session reset behavior matches your operational policy for link-down events.

Increase packet rate for high eBGP session counts

Raise LPTS PIFIB hardware policing rates to support more than 3,500 eBGP sessions.

Before you begin

Identify the target location ID, for example, 0/2/CPU0.

Procedure

Enter global configuration mode and set BGP flow rates.

Example:

```
Router# configure
Router(config)# lpts pifib hardware police location 0/2/CPU0
Router(config-pifib-policer-per-node)#flow bgp configured rate 4000
```

```
Router(config-pifib-policer-per-node)#flow bgp known rate 4000
Router(config-pifib-policer-per-node)#flow bgp default rate 4000
Router(config-pifib-policer-per-node)#commit
```

The device increases packet handling capacity for BGP flows, helping sustain high eBGP session counts without excessive flapping.