Graceful Maintenance

This chapter provides a comprehensive guide to BGP graceful maintenance and failover features. It covers mechanisms like BGP-RIB feedback, extended route retention, nonstop routing, and fast fallover to ensure network stability and minimize traffic loss during various events. Additionally, it details BGP persistence (LLGR) and graceful maintenance for controlled traffic shifts during planned outages and failures.

- BGP-RIB feedback mechanisms for update generation, on page 1
- BGP extended route retention, on page 3
- BGP nonstop routing with stateful switchovers, on page 6
- BGP fast external fallover, on page 10
- BGP fast fallover, on page 11
- BGP persistence, on page 13
- Flexible BGP persistence, on page 16
- BGP graceful maintenance, on page 21

BGP-RIB feedback mechanisms for update generation

The BGP-RIB feedback mechanisms for update generation are BGP route control mechanisms that

- ensure routes are installed locally before advertising them to neighbors
- track which route versions the forwarding information base (FIB) has consumed, and
- send updates only for routes confirmed installed in the FIB to prevent premature advertisements.

These mechanisms help you avoid traffic loss caused by premature route advertisements after events like router reloads, line card online insertion and removal (LC OIR), or link flaps when alternate paths become available

You configure BGP to wait for routing information base (RIB) feedback before sending updates by using the **update wait-install** command in the `router address-family IPv4` or `router address-family VPNv4` configuration mode. This command ensures that BGP sends updates only after routes are confirmed installed in the forwarding information base (FIB), preventing premature route advertisements.

You can verify this configuration using the following commands:

- **show bgp**
- **show bgp neighbors**

- **show bgp process performance-statistics**

This configuration helps you avoid traffic loss caused by premature route advertisements after events such as router reloads, line card online insertion and removal (LC OIR), or link flaps when alternate paths become available.

Guidelines for BGP-RIB feedback mechanisms

To prevent traffic loss and ensure reliable routing, always advertise BGP routes only after they are confirmed as installed in the Forwarding Information Base (FIB) via the BGP-RIB feedback mechanism.

- ensure that BGP installs routes in the Routing Information Base (RIB) and waits for feedback from the RIB about installation in the FIB.
- confirm that the RIB tracks which route versions are in the FIB using the BCDL feedback mechanism.
- send BGP update messages only for routes confirmed as installed in the FIB, preventing premature advertisements that could cause packet loss or blackholing.

Configure BGP to wait for RIB feedback before sending updates

Enable BGP to delay advertising updates until routes are confirmed as installed in the FIB, preventing premature updates and possible traffic loss.

Use this configuration to enhance BGP routing reliability by ensuring updates are sent only after successful RIB-to-FIB installation confirmation.

Before you begin

- Verify you have administrator access to the router.
- Identify the AS number and desired address family (e.g., IPv4 unicast, VPNv4).

Follow these steps to configure BGP to wait for RIB feedback:

Procedure

Step 1 Enter router configuration mode for the desired address family.

Example:

```
Router# configure
Router(config)# router bgp 1
Router(config-bgp)# address-family ipv4 unicast
Router(config-bgp-af)# update wait-install
```

Step 2 Save the configuration.

Example:

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

Step 3 Use the **show bgp process** command to view the delay of the BGP process update since the last router reload.

Example:

```
Router# show bgp process
Wed Aug 24 00:40:48.649 PDT

BGP Process Information:
BGP is operating in STANDALONE mode
Autonomous System number format: ASPLAIN
Autonomous System: 100
Router ID: 192.168.0.2 (manually configured)
Default Cluster ID: 192.168.0.2
Active Cluster IDs:  192.168.0.2
------------------------------
------------------------------
Update wait-install enabled:
  ack request 2, ack rcvd 2, slow ack 0
  startup delay 10 secs
```

BGP on the router now delays advertising route updates until RIB confirms that routes are installed in the FIB, ensuring reliable route propagation and preventing traffic loss.

What to do next

Monitor BGP updates and verify network stability after applying the configuration.

BGP extended route retention

BGP extended route retention is a routing feature that

- applies a route retention policy to modify route attributes when a BGP peer fails
- modifies route attributes in addition to changes caused by the neighbor's inbound policy, and
- enables the use of route retention policy instead of Long-Lived Graceful Restart (LLGR) when the BGP hold timer expires or when the BGP session fails to reestablish within the configured graceful restart time

This feature helps maintain route stability and control during BGP peer failures by retaining routes with modified attributes until the session is restored.

Table 1: Feature History Table

Feature Name	Release Name	Description
BGP Extended Route Retention	Release 7.3.3	This feature allows you to maintain stale routing information from a failed BGP peer for longer periods of time than that is configured in the Graceful Restart atribute. However, this feature ensures that the BGP neighbor considers the stale routes as new routes.

Recommendations for using extended route retention

Adhere to the following principles to ensure proper operation and compatibility when using BGP Extended Route Retention:

- Ensure that your BGP neighbor supports graceful restart functionality.
- Apply graceful restart functionality when a BGP neighbor fails, and maintain it until the graceful restart timer expires.
- Start the Extended Route Retention feature only after the graceful restart timer expires.
- Configure soft-reconfiguration inbound as a mandatory setting; apply inbound policy if required.
- Activate Extended Route Retention exclusively when the BGP peer goes down, specifically after the hold-down timer expires.
- Do not treat routes as stale or retain them for any other triggers such as timer expiry; in such cases, purge the routes.
- Use Extended Route Retention only with the following address-family modes: IPv4 and IPv6 unicast, IPv4 and IPv4 labelled unicast.
- Do not configure both LLGR and Extended Route Retention on the same neighbor.
- Do not send the capability attribute when Extended Route Retention is configured.

Configure route policies and apply them to a BGP neighbor

Define route policies with specific community and local preference settings, then apply these policies to a BGP neighbor to control routing behavior, including route retention and inbound policy processing.

Use this task when you need to define route policies with specific community and local preference settings and apply them to a BGP neighbor, including route retention and inbound policies.

Before you begin

- Ensure you have the necessary privileges to configure BGP and route policies on the router.
- Identify the names for the route-policies and communities.

Procedure

Step 1 Create route policies.

Example:

Create the route policy RRP_comm_no_export_local_pref_2500.

```
Router(config)# route-policy RRP_comm_no_export_local_pref_2500
Router(config-rpl)# set community RRP_comm_no_export additive
Router(config-rpl)# set local-preference 2500
Router(config-rpl)# end-policy
Router(config-rpl)# exit
```

Create the route policy comm_number_local_pref.

```
Router(config)# route-policy comm_number_local_pref
Router(config-rpl)# set community comm_number
Router(config-rpl)# set local-preference 10000
Router(config-rpl)# end-policy
Router(config-rpl)# exit
```

Step 2 Apply the route policies to a BGP neighbor.

Example:

Enter BGP router configuration mode for AS 140:

```
Router(config)# router bgp 140
```

Configure the neighbor with IP address 10.1.1.1.

```
Router(config-bgp)# neighbor 10.1.1.1
Router(config-bgp-nbr)# remote-as 1
Router(config-bgp-nbr)# update-source Loopback 0
```

Enter the IPv4 unicast address family.

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

Apply the inbound route policy.

```
Router(config-bgp-nbr-af)# route-policy RRP_comm_no_export_local_pref_2500 retention retention-time
 2340
```

Enable soft reconfiguration to view the peer's adj-rib-in table.

```
Router(config-bgp-nbr-af)# soft-reconfiguration inbound always
```

The route policies are now successfully configured and applied to the BGP neighbor. This setup controls community attributes, local preference values, route retention duration, and inbound policy processing.

Step 3 Use the **show bgp neighbor** command to verify the configured route-retention policy using show bgp neighbor.

Example:

```
Router# show bgp neighbor 10.1.1.1
Fri Oct 22 04:52:44.972 PDT

BGP neighbor is 10.1.1.1
 Remote AS 1, local AS 1, internal link
 Remote router ID 10.1.1.1
  BGP state = Established, up for 00:03:03
…
 For Address Family: IPv4 Unicast
  BGP neighbor version 16172
Policy for incoming advertisements is comm_number_local_pref
  Policy for Retention is RRP_comm_no_export_local_pref_2500
Configured route retention policy stale timer for routes is 2340 seconds
….
```

The `show bgp neighbor 10.1.1.1` output confirms:

- BGP neighbor 10.1.1.1 is established and up for 3 minutes.
- Remote AS is 1; local AS is 1 (internal link).
- Inbound policy: `comm_number_local_pref`

  .
- Route-retention policy: `RRP_comm_no_export_local_pref_2500`

.

- Stale timer set to 2340 seconds.

This verifies the route policies and retention timer are correctly applied and active.

Step 4 Use the **show bgp ipv4 unicast** command to verify the presence and status of stale routes using show bgp ipv4 unicast.

Example:

```
Router# show bgp ipv4 unicast 181.1.1.0/24
Fri Oct 22 04:56:15.906 PDT
….
Path #1: Received by speaker 0
  Advertised IPv4 Unicast paths to peers (in unique update groups):
    3.3.3.3
  100, (Received from a RR-client), (long-lived/route-retention policy stale)
    192.1.2.1 (metric 10) from 192.1.2.1 (10.1.1.1)
      Origin IGP, metric 1221, localpref 2500, valid, internal, best, group-best, multipath
      Received Path ID 0, Local Path ID 1, version 16243
      Community: 1:100 no-export
```

This output confirms that the route is marked as stale due to the route-retention policy. It verifeies that the route's attributes such as origin, metric, local preference, and community, and shows the route is valid, internal, and selected as the best path. The output also confirms the route is advertised to peers and includes path identifiers for tracking. This demonstrates the route-retention policy is effectively maintaining stale routes for network stability.

Route policies with specific attributes are successfully applied to the BGP neighbor, providing route retention and policy enforcement.

BGP nonstop routing with stateful switchovers

A BGP nonstop routing is a network protocol feature that

- enables all BGP peerings to maintain BGP states, and
- ensures continuous packet forwarding without interruption visible to peer routers by maintaining protocol sessions and routing states across process restarts and switchovers.

Guidelines for BGP nonstop routing with stateful switchovers

- Configure the **nsr process-failures switchover** command to maintain nonstop routing (NSR) during BGP or TCP process crashes.
- Without this configuration, BGP neighbor sessions will flap, causing network instability.
- NSR does not prevent session flapping during BGP or TCP process restarts; expect neighbor sessions to flap in such cases.
- Additional measures beyond NSR are required to manage session flapping caused by process restarts.
- This command is mandatory to ensure network stability during process failures.

How active and standby route processors work during switchovers and failures

Summary

This process explains how active and standby route processors, along with BGP and TCP processes, maintain continuous network operation during route processor switchovers or failures. The key components involved in the process are:

- Active route processor: Manages routing and session handling under normal conditions.
- Standby route processor: Synchronizes state with the active processor and takes over when needed.
- BGP process: Maintains Border Gateway Protocol sessions and routing information.
- TCP process: Maintains TCP connections supporting BGP sessions.

Workflow

The process involves the following stages:

1. Switchover or failure detection: The system detects when the active route processor fails or switches over.
2. State synchronization: Synchronization points ensure consistent internal state between active and standby BGP and TCP processes.
3. Session migration: TCP connections and BGP sessions transparently migrate to the standby processor, which becomes active.
4. Continuous forwarding: Packet forwarding continues uninterrupted without requiring peer routers to refresh protocol states or upgrade software.

Result

This process ensures uninterrupted packet forwarding and BGP session continuity during route processor switchovers or failures, maintaining network stability without manual intervention.

Capabilities and limitations of BGP nonstop routing

NSR maintains active BGP and TCP sessions during route processor switchovers, process crashes, and system upgrades by transparently failing over to the standby route processor without interrupting packet forwarding or causing session flaps. This ensures continuous network operation and enhances stability during critical events.

BGP nonstop routing capabilities

- NSR-related alarms and notifications
- Separate tracking of configured and operational NSR states
- NSR statistics collection
- NSR statistics display via `show` commands
- XML schema support
- Auditing mechanisms for active/standby state synchronization

- CLI commands for NSR enablement and disablement

Events triggering NSR

- Route processor switchovers
- BGP or TCP process crashes or failures
- Restart of `l2vpn_mgr` causing state flapping (no traffic loss)
- In-Service System Upgrades (ISSUs) involving stateful switchover (SSO)

Enabled capabilities

- Transparent migration of TCP connections and BGP sessions to standby route processor preserving protocol state without peer refresh

  NSR-related alarms and notifications
- Separate tracking of configured and operational NSR states
- Collection and display of NSR statistics via show commands
- XML schema support for NSR data
- Auditing mechanisms for state synchronization verification
- Support for up to 5000 NSR sessions
- No requirement for software upgrades or NSR support on peer routers

Limitations

- NSR does not prevent session flapping if the BGP or TCP process restarts; expect neighbor sessions to flap in such cases.
- When the `l2vpn_mgr` process restarts, the NSR client (`te-control`) may flap between Ready and Not Ready states, which is expected and causes no traffic loss.

Configuration requirements

- Configure the **nsr process-failures switchover** command to maintain NSR during BGP or TCP process crashes or failures.
- Without this command, BGP neighbor sessions will flap during process crashes.
- NSR does not prevent session flapping during BGP or TCP process restarts; expect neighbor sessions to flap in such cases.
- Additional measures beyond NSR may be required to manage session flapping caused by process restarts.

Additional operational notes

- During route processor switchovers and In-Service System Upgrades (ISSUs), NSR is achieved by stateful switchover (SSO) of both TCP and BGP.

- NSR does not require software upgrades on peer routers, nor do peers need to support NSR.
- When a route processor switchover occurs due to a fault, TCP connections and BGP sessions migrate transparently to the standby RP, preserving protocol state without requiring peer refresh.
- Events like soft reconfiguration and policy changes can alter BGP internal states; synchronization points called post-its keep active and standby BGP processes aligned.

Enable BGP NSR

Configure BGP NSR for BGP to enhance process resiliency.

NSR allows BGP sessions to remain up during process restarts. Enable NSR to improve network availability

Before you begin

Ensure you are in privileged EXEC mode on the router.

Save your current configuration.

Procedure

Enable NSR.

Example:

```
Router# router bgp 120
Routing(config-bgp)# nsr
```

The router enables BGP NSR as configured.

Disable BGP NSR

To remove non-stop routing for BGP or troubleshoot related issues, disable NSR for BGP.

NSR allows BGP sessions to remain up during process restarts. Disable NSR if troubleshooting or removing the feature.

Before you begin

Ensure you are in privileged EXEC mode on the router.

Save your current configuration.

Procedure

Disable NSR.

Example:

```
Router# router bgp 120
Routing(config-bgp)# no nsr
```

The router disables BGP NSR as configured.

BGP fast external fallover

BGP fast external fallover is a routing feature that

- automatically resets all BGP sessions of directly adjacent external peers when a link goes down
- ensures immediate response to link failures affecting eBGP connections, and
- can be disabled or re-enabled to favor either session stability or rapid convergence.

Guidelines for BGP fast external failover

To prevent eBGP session flapping when your node reaches 3500 eBGP sessions with BGP timer values set to 10 and 30, increase the packet rate. Use the **lpts pifib hardware police location location-id** command to raise the packet rate and support more than 3500 eBGP sessions.

Applies when managing high numbers of eBGP sessions on a node.

Increasing the packet rate prevents session instability caused by timer settings and session volume.

Ensures stable BGP session management under heavy session loads.

Configure the packet rate using the command **lpts pifib hardware police location** command.

Configure the eBGP session packet rate

This task instructs you on how to increase the packet rate for BGP traffic, which helps support a higher number of eBGP sessions and prevent session flapping in large-scale deployments.

Before you begin

Ensure you have access to the router's global configuration mode and the necessary permissions to modify LPTS PIFIB hardware policing settings.

Procedure

To increase the packet rate for BGP traffic and support a higher number of eBGP sessions, configure the LPTS PIFIB hardware policing settings as follows

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
```

This configuration helps prevent session flapping in large-scale deployments by allowing the router to handle increased BGP packet rates efficiently.

BGP fast fallover

BGP fast fallover is a routing feature that

- quickly removes routes learned from directly connected iBGP or eBGP neighbors when an IP interface fails
- accelerates the network convergence process by preventing the propagation of stale routes, and
- eliminates the need to wait for the hold timer to expire when a directly attached interface fails.

When an interface attached to a directly connected BGP neighbor fails, routes learned from that neighbor typically persist until the hold timer expires. This lag can lead to slow network convergence and potential network instability. BGP Fast Fallover addresses this by ensuring routes are removed immediately. You can also use the `nexthop trigger-delay` command to quickly remove BGP routes of a failing neighbor, provided that the neighbor's BGP session endpoint is the same as the route's next hop.

Table 2: Feature History Table

Feature Name	Release Information	Feature Description
BGP Fast Fallover	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 Fast Fallover	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+A8:B12 • 88-LC1-12TH24FH-E • 88-LC1-52Y8H-EM • 8212-48FH-M • 8711-32FH-M

BGP Fast Fallover	Release 24.2.11	You can now terminate the external BGP sessions to an adjacent peer when the link to that peer goes down, without waiting for the hold timer to expire. With this feature you can enable fast fallover mechanism on a specific BGP neighbor even if **bgp fast-external-failover disable** command is globally configured. This feature enables quicker failure detection, and allows other recovery mechanisms to reroute the traffic quickly, thus resulting in faster convergence. The feature introduces these changes: CLI: • **fast-fallover** YANG Data Model: • Cisco-IOS-XR-um-router-bgp-cfg.yang (see GitHub, YANG Data Models Navigator)

Guidelines for BGP fast fallover

To ensure proper operation and network stability when using BGP fast fallover, follow these guidelines:

- Apply fast fallover only to directly connected BGP neighbors. A directly connected neighbor is one that is either one hop away or has an IP address within the same subnet. Do not apply fast fallover to neighbors connected through loopback interfaces, even if they are one hop away.
- Maintain established BGP sessions with neighbors that are not directly connected until a triggering event, such as hold timer expiration, causes the session to go down.
- If an interface fails before fast fallover activates, manually clear the BGP neighbor session if necessary, as the BGP session does not automatically go down.
- Allow the regular BGP session establishment process to proceed unchanged when an interface recovers from failure.

Follow these recommendations to maintain predictable BGP session behavior and ensure network stability when using the fast fallover feature.

Configure BGP fast fallover

This task describes how to enable and verify the BGP fast fallover feature for neighbors.

Before you begin

- Verify you have appropriate access to the router's configuration mode.
- Confirm that you understand the network topology, especially which BGP neighbors are directly connected.
- Ensure you have the necessary permissions to modify BGP neighbor configurations.

Procedure

Step 1 Enable fast fallover.

Example:

```
Router# configure
Router(config)# router bgp 120
Router(config-bgp)# neighbor 209.165.201.0
Router(config-bgp-nbr)# fast-fallover
```

- By default, fast fallover is enabled for eBGP neighbors and disabled for iBGP neighbors.
- If the `bgp fast-external-fallover disable` command is configured globally or in VRF mode, fast fallover is disabled for eBGP neighbors but can be overridden per neighbor using the `fast-fallover` command.
- To stop the fast fallover setting from being inherited from a higher-level neighbor or session group, use the `fast-fallover inheritance-disable` command.

Step 2 Verify fast fallover configuration.

Example:

Use the **show bgp neighbors ip-address** or **show run router bgp as-number neighbor ip-address** command to check if fast fallover is enabled or inherited.

```
Router# show bgp neighbors 209.165.201.0
BGP neighbor is 209.165.201.0
...
fast-fallover
address-family ipv4 unicast
 use af-group ipv4_unicast_3_ibgp
...
Fast fallover is enabled
Neighbor is directly connected
Neighbor fast-fallover is configured
Neighbor is external and fast-external-fallover is not disabled
...
```

The presence of fast-fallover in the configuration output confirms that the feature is successfully configured for the neighbor.

BGP persistence

BGP persistence mechanism is a routing feature that

- allow you to retain BGP routes learned from a configured neighbor
- keep those routes active even if the BGP session with that neighbor goes down, and
- help you maintain network stability during transient neighbor outages.

BGP persistence, also known as Long-Lived Graceful Restart (LLGR), is a feature that enables a local router to retain BGP routes learned from a configured neighbor, even if the BGP session with that neighbor goes

down. This capability helps maintain network stability by preventing unnecessary route withdrawals during transient neighbor outages.

Key characteristics of BGP Persistence routes include:

- Duration: LLGR can remain in effect for a significantly longer period than standard Graceful Restart (GR).
- Route Preference: LLGR stale routes are assigned the lowest preference during the BGP bestpath computation, ensuring that fresh, active routes are always preferred.
- Advertisement: If an LLGR stale route is selected as the best path, it is advertised with the `LLGR_STALE` community (65535:6) attached. These routes are not advertised to neighbors that do not support LLGR.
- Resilience: LLGR stale routes are not deleted if the forwarding path to the neighbor is detected as down, nor are they deleted if the BGP session to the neighbor experiences multiple flaps, even if the neighbor does not re-advertise the route.
- Exclusion: Any route explicitly tagged with the `NO_LLGR` community (65535:7) will not be retained by the BGP persistence mechanism.

By using BGP persistence, you reduce the impact of transient BGP session outages on your network, which can be especially useful in large or dynamic environments.

Limitations for BGP persistence

The BGP persistence feature is supported only on the following address family identifiers:

- VPNv4 and VPNv6
- RT constraint
- Flowspec (IPv4, IPv6, VPNv4, and VPNv6)
- IPv4 and IPv6 address families

BGP persistence operational flow

Summary

BGP persistence operational flow begins when a BGP session drops or after Graceful Restart, with its capability signaled during session establishment. The local router then retains learned routes as stale. LLGR concludes upon stale timer expiry or receipt of an End-of-RIB marker, at which point any remaining stale routes are deleted.

Workflow

BGP persistence operates through a defined lifecycle as given below:

1. Initiation: LLGR takes effect either immediately upon a BGP session going down (if standard Graceful Restart is not enabled) or after the standard Graceful Restart process concludes.
2. Capability Signaling: The LLGR capability is signaled to a neighbor during the BGP session establishment via the BGP OPEN message, provided it has been configured for that neighbor.

3. Route Retention: Once active, the local router retains learned routes from the neighbor, marking them as "stale" but keeping them in the routing table.
4. Termination: LLGR for a neighbor ends when one of the following conditions is met:
   - The configured LLGR stale timer expires.
   - The neighbor sends an End-of-RIB (Routing Information Base) marker, indicating that it has completed revising and re-advertising its routes.
5. Stale Route Deletion: Upon LLGR termination, any routes from that neighbor that are still marked as stale are deleted from the routing table.

Configure BGP persistence

Use this task to enable BGP Persistence (Long-Lived Graceful Restart) for a BGP neighbor, allowing the router to retain routes during session outages.

BGP Persistence, also known as Long-Lived Graceful Restart (LLGR), ensures network stability by keeping learned routes active even when a neighbor session is temporarily down. This configuration defines how long stale routes are retained.

Before you begin

Ensure that you meet the following requirements:

- You must have a basic BGP configuration already in place, including the BGP process and the neighbor definition.
- You must know the remote Autonomous System (AS) number and the IP address of the BGP neighbor for which you are configuring LLGR.
- Understand the implications of **graceful-restart stalepath-time** if configured, as LLGR takes effect after standard Graceful Restart concludes.

To configure BGP Persistence for a neighbor, perform the following steps:

Procedure

Step 1 Configure BGP on the router, and the BGP neighbor, and its basic parameters.

Example:

```
Router(config)# router bgp 100
Router(config-router)# neighbor 10.3.3.3
Router(config-router-neighbor)# remote-as 30813
Router(config-router-neighbor)# update-source Loopback0
Router(config-router-neighbor)# graceful-restart stalepath-time 150
Router(config-router-neighbor)#
```

Step 2 Enter address family configuration mode for VPNv4 unicast.

Example:

```
Router(config-router-neighbor)# address-family vpnv4 unicast
```

Step 3 Enable long-lived graceful restart capability for the VPNv4 address family and specify the LLGR stale time for sending and accepting VPNv4 routes.

Example:

```
Router(config-router-af)# long-lived-graceful-restart capable
Router(config-router-af)# long-lived-graceful-restart stale-time send 16777215 accept 16777215
Router(config-router-af)# exit
```

Step 4 Enter address family configuration mode for VPNv6 unicast.

Example:

```
Router(config-router-neighbor)# address-family vpnv6 unicast
```

Step 5 Enable long-lived graceful restart capability for the VPNv6 address family and specify the LLGR stale time for sending and accepting VPNv6 routes.

Example:

```
Router(config-router-af)# long-lived-graceful-restart capable
Router(config-router-af)# long-lived-graceful-restart stale-time send 16777215 accept 16777215
Router(config-router-af)# exit
```

Flexible BGP persistence

Flexible BGP persistence is a routing feature that

- enhances network stability and resilience
- enables Long-Lived Graceful Restart (LLGR) with flexible stale time management, and
- allows controlled route distribution within the Autonomous System (AS).

This feature provides the flexibility to advertise LLGR stale routes to both LLGR-capable and non-LLGR-capable neighbors, ensuring continuous route availability during planned or unplanned restarts. It simplifies configuration by removing the need for manual timeout settings and enforces controlled route propagation by attaching specific BGP community attributes.

Table 3: Feature History Table

Feature Name	Release Name	Description
Flexible BGP Persistence	Release 25.1.1	Introduced in this release on: Fixed Systems (8700 [ASIC: K100]). This feature is supported on Cisco 8712-MOD-M routers.

Feature Name	Release Name	Description
Flexible BGP Persistence	Release 24.3.1	

Feature Name	Release Name	Description
		Introduced in this release on: Fixed Systems (8200 [ASIC: Q200, P100], 8700 [ASIC: P100]); Centralized Systems (8600 [ASIC:Q200]); Modular Systems (8800 [LC ASIC: Q100, Q200, P100]) Now you can ensure continuous connectivity by allowing non-Long Lived Graceful Restart (LLGR) eBGP neighbors to use LLGR stale routes, allowing for LLGR capability to be enabled and advertised without having to explicitly configure a timeout value, and gain greater flexibility in route management by advertising stale routes to non-LLGR peers through the NO_EXPORT community. This is an enhancement to the existing BGP Persistence feature. The feature introduces these changes: CLI: • The **default**, **any**, and **advertise-internal-only** keywords are added to the long-lived-graceful-restart command. • The fields **ault advertised long-lived stale time**, and **Long-lived Graceful Restart Stale Time Accept Any** are added to the show output of the show bgp command. YANG Data Model: • `Cisco-IOS-XR-ipv4-bgp-cfg` (see GitHub, YANG Data Models Navigator) *This feature is supported on: • 88-LC1-36EH

Feature Name	Release Name	Description
		• 88-LC1-12TH24FH-E • 88-LC1-52Y8H-EM • 8212-48FH-M • 8711-32FH-M

Benefits of flexible BGP persistence

Flexible BGP Persistence offers the following benefits:

- Simplified configuration: Enables LLGR without manual timeout configuration using the **long-lived-graceful-restart stale-time send default accept any** command, which advertises a default stale time and accepts peer-specified stale times.
- Enhanced network resilience: Allows LLGR stale routes to be advertised to non-LLGR eBGP neighbors, improving overall network robustness.
- Enhanced network stability: Attaches the `NO_EXPORT` community and sets local preference to 0 for LLGR routes advertised to internal neighbors without LLGR capability, preventing stale routes from propagating beyond the local AS and ensuring they are not preferred over other routes.

Configure LLGR advertisement and activation with default and peer-time values

Enable and advertise Long-Lived Graceful Restart (LLGR) capability using default and peer-specified stale time values.

Before you begin

Ensure you have router BGP configuration access.

Procedure

Step 1 Enable and advertise LLGR capability with default and peer stale times.

Example:

```
Router(config)# router bgp 100
Router(config-bgp)# neighbor 10.1.1.1
Router(config-bgp-nbr)# address-family ipv4 unicast
Router(config-bgp-nbr-af)# long-lived-graceful-restart stale-time send default accept any
```

This command advertises LLGR capability with a default stale time of 172,800 seconds (2 days) and accepts any stale time value set by the peer.

Step 2 Use the **show bgp neighbor** to verify the LLGR configuration and stale time settings.

Example:

```
Router(config)# show bgp neighbor 192.0.2.254
...
AF-dependent capabilities:
```

```
...
Long-lived Graceful Restart Stale Time Send Default is ON
Default advertised long-lived stale time is 172800 seconds
Long-lived Graceful Restart Stale Time Accept Any is ON
....
```

These output lines indicate these settings.

- Long-lived Graceful Restart Stale Time Send Default is ON.
- Default advertised long-lived stale time is 172800 seconds.
- Long-lived Graceful Restart Stale Time Accept Any is ON.

LLGR capability is enabled on the BGP neighbor with default and peer-specified stale times successfully advertised and verified.

Enable LLGR capability and advertise it only to iBGP peers

Enable long lived graceful restart (LLGR) capability and restrict advertisement to iBGP peers.

Before you begin

Confirm BGP neighbor configuration.

Procedure

Step 1 Enable LLGR capability and advertise it only to iBGP peers:

Example:

```
Router(config)# router bgp 100
Router(config-bgp)# neighbor 10.1.1.1
Router(config-bgp-nbr)# address-family ipv4 unicast
Router(config-bgp-nbr-af)# long-lived-graceful-restart capable advertise-internal-only
```

This step sets the local preference to 0 and attaches the `llgr-stale` and `no-export` community attributes to limit route propagation within the local AS.

Step 2 Use the **show bgp** command to that the route is not preferred for outbound traffic within the AS, and that it will not be advertised to external BGP peers, thereby limiting its propagation to within the local AS.

Example:

```
Router# show bgp 10.1.1.1
Path #32: Received by speaker 0
…
192.0.2.254 (metric 30) from 10.1.1.1 (192.0.2.254)
Origin IGP, localpref 0, valid, internal, add-path
Received Path ID 40, Local Path ID 9, version 14321
Community: llgr-stale no-export
Originator: 192.0.2.254, Cluster list: 10.1.1.1
```

These output lines indicate these settings.

- Local preference 0

- Community attributes: llgr-stale no-export

BGP graceful maintenance

BGP graceful maintenance is a routing feature that

- allows routers or links to remain in service while the network reroutes traffic to alternative paths,
- minimizes traffic loss during convergence by enabling the network to find alternate routes before taking a router or link out of service, and
- supports both shutdown and startup scenarios to reduce traffic disruption.

This feature is especially useful in networks with long convergence times caused by factors such as large routing tables or route reflectors.

Restrictions for BGP graceful maintenance

These restrictions apply to BGP graceful maintenance:

- Routers configured to send the GSHUT community attribute require other routers to interpret it via matching routing policies and setting lower preferences.
- The LOCAL_PREF attribute is not sent to another AS, so it cannot be used on eBGP links, except between member-ASs of an AS confederation.
- Alternate routes must exist in the network; otherwise, advertising a lower preference has no effect (e.g., singly-homed customer routers without alternate routes).

Graceful maintenance operation

When you activate BGP Graceful Maintenance, you ensure a controlled traffic shift by signaling reduced route preference to neighboring routers, prompting them to select alternative paths.

- Signaling Reduced Preference: You can signal reduced route preference using these methods:
  - Add GSHUT community: This method allows remote routers the flexibility to define their own preference based on a policy match.
  - Reduce LOCAL_PREF value: This method is effective for internal BGP neighbors, especially if remote routers do not process the GSHUT community.
  - Prepend AS Path: This method works for both internal and external BGP neighbors, particularly when remote routers do not recognize the GSHUT community.
- Impact on BGP Connections: When you activate Graceful Maintenance on a BGP connection, two key operations occur:

  - Re-advertisement of received routes: All routes received from the connection are re-advertised to other neighbors with a lower preference. This applies only to routes that were originally advertised to those neighbors.
  - Re-advertisement of advertised routes: All routes previously advertised *to* the connection are re-advertised with a lower preference.

- Internal Tagging for Received Routes: To facilitate the re-advertisement of received routes, we internally tag them with a `graceful-shut` attribute. This attribute is local to the router and is not advertised via BGP. You can view this attribute using the **show bgp** command. This differs from the `GSHUT` community, which is advertised by BGP.
- Route Selection Preference: All routes possessing the `graceful-shut` attribute receive the lowest preference during route selection. Any new route updates exchanged on a BGP session under Graceful Maintenance are processed with this same preference adjustment.

Guidelines for inter-autonomous system usage

Advertise a lower preference to another autonomous system (AS) in the public Internet only when necessary. Unnecessary advertisement may cause excessive routing updates in remote networks, which can be undesirable.

Key points:

- Avoid advertising lower preference unless required to prevent unnecessary routing advertisements.
- Configure the router with the `send-community-gshut-ebgp` setting under the neighbor address family to originate the GSHUT community to the eBGP neighbor.
- Be aware that this setting affects only the GSHUT community you add; it does not alter the GSHUT community on routes you receive.

Best practices for handling the graceful-shut attribute in BGP

- Assign the lowest preference to any BGP route tagged with the graceful-shut attribute during route selection to maintain correct routing behavior.
- Remember that the graceful-shut attribute is internal and is not advertised to external BGP peers.
- Do not confuse the graceful-shut attribute with the GSHUT community, which is advertised externally by BGP.
- During Graceful Maintenance, treat any new route updates received or sent on a BGP session in the same way.

This approach ensures proper route handling and avoids inadvertent routing preferences during network maintenance.

Requirement: Verify network convergence before performing a graceful maintenance shutdown

Always confirm full network convergence before shutting down a router or link after activating graceful maintenance. This step is essential to prevent traffic loss and service disruption.

Premature shutdown before convergence can cause severe operational hazards and substantial traffic loss. Network-wide convergence may take seconds to over an hour and cannot be determined solely by the local router state.

By verifying convergence, you minimize the risk of traffic blackholing and ensure a smooth, uninterrupted transition during maintenance.

Apply these instructions every time you perform graceful maintenance and intend to shut down a router or link. Skipping this step can lead to critical service impacts.

To determine full convergence:

- Monitor BGP messaging queues and traffic flow.
- Use the **show bgp <vrf> <afi> <safi> summary** commands to confirm both **InQ** and **OutQ** are zero) for all neighbors.
- Ensure traffic is no longer sent to the router being maintained.

Configure graceful maintenance on a BGP router for all neighbors

Advertise routes with the GSHUT community to enable graceful service maintenance across all BGP neighbors through a single configuration.

Before you begin

- Verify you have the required access permissions and credentials to configure the BGP router.
- Confirm that the router is operational and that BGP sessions with neighbors are established.

Procedure

Enter BGP router configuration mode and configure graceful maintenance for all neighbors.

Example:

```
Router# configure
Router(config)# router bgp 120
Router(config-bgp)# graceful-maintenance activate all-neighbors
```

What to do next

- After activating Graceful Maintenance, wait for all routes to be sent and for neighboring routers to redirect traffic away from the router or link under maintenance. Once traffic is redirected, it is safe to take the router or link out of service.

- To monitor progress, use the `show bgp summary` command and check the `OutQ` value for neighbors. When `OutQ` reaches 0, no more updates remain to be sent.

Configure graceful maintenance on a neighbor

Enable the router to announce routes with graceful maintenance attributes, including the GSHUT community, to a neighbor.

Before you begin

Verify the IP address of the BGP neighbor is known and reachable.

Procedure

Enable graceful maintenance on BGP neighbor.

Example:

```
Router# configure
Router(config)# router bgp 120
Router(config-bgp)# neighbor 172.168.40.24
Router(config-bgp-nbrgrp)# graceful-maintenance activate
Router(config-bgp)# commit
```

Activating graceful maintenance on a single neighbor enables coordinated route advertisement with GSHUT community attributes for the specific neighbor in the group, facilitating efficient maintenance operations.

Graceful maintenance is activated for the specified neighbor, enabling coordinated route advertisement with the GSHUT community attribute for maintenance operations.

Configure the router to reduce route preference for graceful maintenance

Allow alternate BGP routes to take over before removing a link or router by lowering route preference.

Use this procedure to safely enable BGP graceful maintenance without routing disruption. Lowering the route preference (local-preference) ensures traffic shifts to alternative paths while maintaining service continuity.

Before you begin

- Ensure alternate network paths are available to receive redirected traffic.
- Prepare the route policy that matches the GSHUT community to set the lower local preference value.

Follow the steps to reduce the route preference on the router :

Procedure

Step 1 Configure the BGP neighbor for graceful maintenance with a lower local preference:

Example:

```
Router# configure
Router(config)# router bgp 120
Router(config-bgp)# neighbor 172.168.40.24
Router(config-bgp-nbr)# remote-as 2002
Router(config-bgp-nbr)# graceful-maintenance local-preference 4
```

Step 2 Define a route policy that matches the GSHUT community and sets the local preference to 0:

Example:

```
Router(config)# route-policy gshut
Router(config-rpl)# if community matches-any gshut then
Router(config-rpl)# set local-preference 0
Router(config-rpl)# endif
Router(config-rpl)# pass
Router(config-rpl)# end-policy
```

Step 3 Apply the route policy inbound on the configured BGP neighbor:

Example:

```
Router# configure
Router(config)# router bgp 120
Router(config-bgp)# neighbor 172.168.40.24
Router(config-bgp-nbr)# address-family ipv4 unicast
Router(config-bgp-nbr-af)# route-policy gshut in
```

- Attributes for graceful maintenance are added to a route update message after applying an outbound policy.
- Routes received from a GSHUT neighbor are marked with a GSHUT attribute to distinguish them from routes received with the GSHUT community.
- When a neighbor exits maintenance, its GSHUT attribute is removed, but the community tag remains.
- The GSHUT attribute is internal and used only for path selection; it is not sent in BGP messages.

The router advertises routes with lower preference, allowing alternate routes to be selected for traffic before you take down a link or the router.

Configure BGP inbound route policy for GSHUT community

Configure route policy for graceful maintenance by lowering the preference for GSHUT community.

This configuration defines and applies a BGP route policy that sets the `local-preference` to `0` for incoming routes carrying the `gshut` BGP community.

Before you begin

Ensure alternate paths exist: The graceful maintenance feature relies on traffic having other routes to take; without them, this configuration is ineffective.

Ensure you have a backup of the current configuration: Always save a copy of your router's configuration before making any changes to allow for easy rollback if needed.

Follow these steps to configure a BGP inbound route policy that lowers the preference of routes carrying the GSHUT community, and facilitating graceful service maintenance.

Procedure

Configure route policy matching GSHUT community to lower route preference.

Example:

```
Router(config)# route-policy gshut
Router(config-rpl)# if community matches-any gshut then
Router(config-rpl-if)# set local-preference 0
Router(config-rpl-if)# endif
Router(config-rpl-if)# pass
Router(config-rpl-if)# end-policy
Router(config-rpl-if)# exit
Router(config-rpl)# exit
Router(config)# router bgp 120
Router(config-bgp)# neighbor 10.0.0.3
Router(config-bgp-nbr)# address-family ipv4 unicast
Router(config-bgp-nbr)# route-policy gshut in
Router(config-bgp-nbr)# exit
Router(config-bgp)# exit
Router(config)# exit
```

Verify BGP graceful maintenance activation and attributes

Ensure BGP graceful maintenance is active and correctly configured to allow for seamless routing transitions during planned maintenance.

Use this task to verify that your BGP routers are properly advertising the graceful-shutdown community and associated attributes, preventing network disruptions during maintenance events.

Before you begin

Ensure you have access to the router command.

Confirm you have privileges to run show commands.

Follow these steps to verify BGP graceful maintenance activation and attributes.

Procedure

Step 1 Verify BGP routes with graceful-shutdown community

Example:

```
Router# show bgp 192.0.2.1
...
192.0.2.10 from 192.0.2.10 (198.51.100.1)
Received Label 24000
Origin incomplete, metric 0, localpref 100, valid, internal, best, group-best,
import-candidate
Received Path ID 0, Local Path ID 1, version 4
Community: graceful-shutdown
Originator: 198.51.100.1, Cluster list: 198.51.100.1
```

Review the output for entries showing the `graceful-shutdown` community and related path attributes.

These lines confirm that BGP graceful maintenance is active and that the route is marked accordingly to allow alternate paths to take over before a link or router is taken down:

- Community: graceful-shutdown: This line explicitly shows that the route carries the graceful-shutdown community attribute.
- Origin incomplete, metric 0, localpref 100, valid, internal, best, group-best, import-candidate: These status codes and attributes confirm the route's state, including the presence of the graceful-shut path attribute as part of the route's characteristics.
- Received Path ID 0, Local Path ID 1, version 4: This line provides path identification details relevant to the graceful maintenance feature.

Step 2 Verify the graceful-shutdown community and path attribute.

Example:

```
Router# show bgp community graceful-shutdown
BGP router identifier 198.51.100.1, local AS number 4
BGP generic scan interval 60 secs
BGP table state: Active
Table ID: 0xe0000000 RD version: 18
BGP main routing table version 18
BGP scan interval 60 secs
Status codes: s suppressed, d damped, h history, * valid, > best
i - internal, r RIB-failure, S stale, N Nexthop-discard
Origin codes: i - IGP, e - EGP, ? - incomplete
Network Next Hop Metric LocPrf Weight Path
* 10.5.5.5/32 10.10.10.1 88 0 1 ?
Processed 1 prefixes, 1 paths
```

These lines collectively confirm that the graceful maintenance feature is active and that the routes are marked accordingly to allow alternate paths to take over before a link or router is taken down.

- Community: graceful-shutdown: This line explicitly shows that the route carries the graceful-shutdown community attribute, indicating that graceful maintenance is active for this route.
- Network Next Hop Metric LocPrf Weight Path and the following line: `* 10.5.5.5/32 10.10.10.1 88 0 1 ?`: This shows the network prefix and associated attributes, confirming the route is valid and active under graceful maintenance.
- Processed 1 prefixes, 1 paths: This confirms the number of prefixes and paths processed with the graceful-shutdown community.

Step 3 Verify if graceful maintenance is locally active on a BGP neighbor and to view local preference and AS prepends:

Example:

```
Router# show bgp neighbor 192.0.2.1
...
Graceful Maintenance locally active, Local Pref=45, AS prepends=3
...
For Address Family: IPv4 Unicast
...
GSHUT Community attribute sent to this neighbor
...
```

These lines collectively confirm that graceful maintenance is active and that the relevant attributes are applied and communicated to the neighbor:

- Graceful Maintenance locally active, Local Pref=45, AS prepends=3: This line confirms that graceful maintenance is currently active on the neighbor, showing the local preference value and the number of AS path prepends applied.

- For Address Family: IPv4 Unicast: This indicates the address family context for which the graceful maintenance attributes apply.
- GSHUT Community attribute sent to this neighbor: This line shows that the GSHUT community attribute is being sent to the neighbor, which is part of the graceful maintenance signaling.

Step 4 Verify the key attributes of the BGP neighbor that indicates that the graceful maintenance is enabled:

Example:

```
Router# show bgp neighbor 10.12.12.5 configuration

neighbor 10.12.12.5
remote-as 1 []
graceful-maintenance 1 []
gr-maint local-preference 45 []
gr-maint as-prepends 3 []
gr-maint activate []
```

These lines confirm that graceful maintenance is active and configured with the specified local preference and AS path prepends for this neighbour:

- `remote-as 1`
- `gr-maint local-preference 45`
- `gr-maint as-prepends 3`
- `gr-maint activate`

Step 5 List all community set objects and view graceful maintenance feature attributes:

```
Router# show rpl community-set
Listing for all Community Set objects
community-set gshut
graceful-shutdown
end-set
```

These lines display the community sets configured for graceful maintenance on the router.

Step 6 Monitor syslog warning for BGP graceful maintenance activation.

```
RP/0/0/CPU0:Jan 28 22:01:36.356 : bgp[1056]: %ROUTING-BGP-5-ADJCHANGE : neighbor 10.10.10.4 Up (VRF:
 default) (AS: 4)
WARNING: Graceful Maintenance is Active
```

Use this warning as a reminder to deactivate graceful maintenance after the BGP convergence completes.