



BGP Nonstop Forwarding (NSF) Awareness

Nonstop Forwarding (NSF) awareness allows a router to assist NSF-capable neighbors to continue forwarding packets during a Stateful Switchover (SSO) operation. The BGP Nonstop Forwarding Awareness feature allows an NSF-aware router that is running BGP to forward packets along routes that are already known for a router that is performing an SSO operation. This capability allows the BGP peers of the failing router to retain the routing information that is advertised by the failing router and continue to use this information until the failed router has returned to normal operating behavior and is able to exchange routing information. The peering session is maintained throughout the entire NSF operation.

NSF works with the SSO feature in Cisco IOS software. SSO is a prerequisite of NSF. NSF works with SSO to minimize the amount of time a network is unavailable to its users following a switchover. The main objective of NSF is to continue forwarding IP packets following a Route Processor (RP) switchover. NSF/SSO is configured in the core of your network, and NSF awareness is configured on iBGP peers in the core and on the edge of the network.

Feature Specifications for the BGP Nonstop Forwarding (NSF) Awareness feature

Feature History

| Release | Modification |
|-----------|------------------------------|
| 12.2(15)T | This feature was introduced. |

Supported Platforms

For platforms supported in Cisco IOS Release 12.2(15)T, use Cisco Feature Navigator as described below.

Finding Support Information for Platforms and Cisco IOS Software Images

Use Cisco Feature Navigator to find information about platform support and Cisco IOS software image support. Access Cisco Feature Navigator at <http://www.cisco.com/go/fn>. You must have an account on Cisco.com. If you do not have an account or have forgotten your username or password, click **Cancel** at the login dialog box and follow the instructions that appear.

Contents

- [Prerequisites for BGP Nonstop Forwarding Awareness, page 2](#)
- [Restrictions for BGP Nonstop Forwarding Awareness, page 2](#)
- [Information About BGP Nonstop Forwarding Awareness, page 2](#)
- [How to Configure BGP Nonstop Forwarding Awareness, page 5](#)

- [Configuration Examples for Nonstop Forwarding, page 9](#)
- [Additional References, page 11](#)
- [Command Reference, page 13](#)

Prerequisites for BGP Nonstop Forwarding Awareness

This document assumes that your network is configured to run BGP. You will also need to complete the following tasks before you can configure this feature:

- On platforms supporting the Route Switch Processor (RSP), and where the CEF switching mode is configurable, configure distributed CEF (dCEF) switching mode using the **ip cef distributed** command. This command is enabled by default on the Cisco 12000 Series Internet Router.

Restrictions for BGP Nonstop Forwarding Awareness

The following restrictions apply to the BGP Nonstop Forwarding Awareness feature:

- All neighboring devices participating in BGP NSF must be NSF-capable or NSF-aware, having been configured for BGP graceful restart.
- BGP graceful restart does not support two neighbors performing an NSF restart operation at the same time because these peers cannot hold routes for each other during an SSO operation. However, both neighbors will still reestablish peering sessions after the NSF restart operation is complete. Each router will reestablish peering with the other as if it were a new router joining the network.
- Existing sessions must be reset by issuing the **clear ip bgp *** command or by reloading the router before graceful restart capabilities will be exchanged.
- BGP graceful restart does not yet support Virtual Routing and Forwarding (VRF) instances and VPNv4 sessions and configurations.

Information About BGP Nonstop Forwarding Awareness

To configure this feature, you must understand the following concepts:

- [Cisco NSF Routing and Forwarding Operation, page 2](#)
- [Cisco Express Forwarding, page 3](#)
- [BGP Graceful Restart, page 3](#)
- [BGP NSF Awareness, page 4](#)

Cisco NSF Routing and Forwarding Operation

Cisco NSF is supported by the BGP, EIGRP, OSPF, and IS-IS protocols for routing and by Cisco Express Forwarding (CEF) for forwarding. Of the routing protocols, BGP, EIGRP, OSPF, and IS-IS have been enhanced with NSF-capability and awareness, which means that routers running these protocols can detect a switchover and take the necessary actions to continue forwarding network traffic and to recover route information from the peer devices.

In this document, a networking device is said to be NSF-aware if it is running NSF-compatible software. A device is said to be NSF-capable if it has been configured to support NSF; therefore, it would rebuild routing information from NSF-aware or NSF-capable neighbors.

Each protocol depends on CEF to continue forwarding packets during switchover while the routing protocols rebuild the Routing Information Base (RIB) tables. Once the routing protocols have converged, CEF updates the FIB table and removes stale route entries. CEF then updates the line cards with the new FIB information.



Note Currently, EIGRP supports only NSF awareness. SSO support for EIGRP will be integrated into a future release.

Cisco Express Forwarding

A key element of NSF is packet forwarding. In a Cisco networking device, packet forwarding is provided by CEF. CEF maintains the FIB and uses the FIB information that was current at the time of the switchover to continue forwarding packets during a switchover. This feature reduces traffic interruption during the switchover.

During normal NSF operation, CEF on the active RP synchronizes its current FIB and adjacency databases with the FIB and adjacency databases on the standby RP. Upon switchover of the active RP, the standby RP initially has FIB and adjacency databases that are mirror images of those that were current on the active RP. For platforms with intelligent line cards, the line cards will maintain the current forwarding information over a switchover; for platforms with forwarding engines, CEF will keep the forwarding engine on the standby RP current with changes that are sent to it by CEF on the active RP. In this way, the line cards or forwarding engines will be able to continue forwarding after a switchover as soon as the interfaces and a data path are available.

As the routing protocols start to repopulate the RIB on a prefix-by-prefix basis, the updates in turn cause prefix-by-prefix updates for CEF, which it uses to update the FIB and adjacency databases. Existing and new entries will receive the new version (“epoch”) number, indicating that they have been refreshed. The forwarding information is updated on the line cards or forwarding engine during convergence. The RP signals when the RIB has converged. The software removes all FIB and adjacency entries that have an epoch older than the current switchover epoch. The FIB now represents the newest routing protocol forwarding information

The routing protocols run only on the active RP, and they receive routing updates from their neighbor routers. Routing protocols do not run on the standby RP. Following a switchover, the routing protocols request that the NSF-aware neighbor devices send state information to help rebuild the routing tables.



Note For NSF operation, the routing protocols depend on CEF to continue forwarding packets while the routing protocols rebuild the routing information.

BGP Graceful Restart

When an NSF-capable router begins a BGP session with a BGP peer, it sends an OPEN message to the peer. Included in the message is a declaration that the NSF-capable or NSF-aware router has “graceful restart capability.” Graceful restart is the mechanism by which BGP routing peers avoid a routing flap following a switchover. If the BGP peer has received this capability, it is aware that the device sending the message is NSF-capable. Both the NSF-capable router and its BGP peer(s) (NSF-aware peers) need

to exchange the graceful restart capability in their OPEN messages, at the time of session establishment. If both the peers do not exchange the graceful restart capability, the session will not be graceful restart capable.

If the BGP session is lost during the RP switchover, the NSF-aware BGP peer marks all the routes associated with the NSF-capable router as stale; however, it continues to use these routes to make forwarding decisions for a set period of time. This functionality means that no packets are lost while the newly active RP is waiting for convergence of the routing information with the BGP peers.

After an RP switchover occurs, the NSF-capable router reestablishes the session with the BGP peer. In establishing the new session, it sends a new graceful restart message that identifies the NSF-capable router as having restarted.

At this point, the routing information is exchanged between the two BGP peers. Once this exchange is complete, the NSF-capable device uses the routing information to update the RIB and the FIB with the new forwarding information. The NSF-aware device uses the network information to remove stale routes from its BGP table. Following that, the BGP protocol is fully converged.

If a BGP peer does not support the graceful restart capability, it will ignore the graceful restart capability in an OPEN message but will establish a BGP session with the NSF-capable device. This functionality will allow interoperability with non-NSF-aware BGP peers (and without NSF functionality), but the BGP session with non-NSF-aware BGP peers will not be graceful restart capable.

BGP NSF Awareness

BGP support for NSF requires that neighbor routers are NSF-aware or NSF-capable. NSF awareness in BGP is also enabled by the graceful restart mechanism. A router that is NSF-aware functions like a router that is NSF-capable with one exception: an NSF-aware router is incapable of performing an SSO operation. However, a router that is NSF-aware is capable of maintaining a peering relationship with a NSF-capable neighbor during a NSF SSO operation, as well as holding routes for this neighbor during the SSO operation.

The BGP Nonstop Forwarding Awareness feature provides an NSF-aware router with the capability to detect a neighbor that is undergoing an SSO operation, maintain the peering session with this neighbor, retain known routes, and continue to forward packets for these routes. The deployment of BGP NSF awareness can minimize the affects of route-processor (RP) failure conditions and improve the overall network stability by reducing the amount of resources that are normally required for reestablishing peering with a failed router.

NSF awareness for BGP is not enabled by default. The **bgp graceful-restart** command is used to enable NSF awareness on a router that is running BGP. NSF-aware operations are also transparent to the network operator and BGP peers that do not support NSF capabilities.



Note NSF awareness is enabled automatically in supported software images for Interior Gateway Protocols, such as EIGRP, IS-IS, and OSPF. In BGP, NSF awareness is not enabled automatically and must be started by issuing the **bgp graceful-restart** command in router configuration mode.

How to Configure BGP Nonstop Forwarding Awareness

This section contains the following procedures:

- [Configuring BGP Nonstop Forwarding Awareness, page 5](#)
- [Configuring BGP NSF Awareness Timers, page 7](#)
- [Verifying the Configuration of BGP Nonstop Forwarding Awareness, page 9](#)

Configuring BGP Nonstop Forwarding Awareness

BGP Graceful Restart

The BGP Nonstop Forwarding (NSF) Awareness feature is part of the graceful restart mechanism. BGP NSF awareness is enabled by issuing the **bgp graceful-restart** command in router configuration mode. BGP NSF awareness allows NSF-aware routers to support NSF-capable routers during an SSO operation. NSF-awareness is not enabled by default and should be configured on all neighbors that participate in BGP NSF.

Restrictions

The configuration of the restart and stalepath timers is not required to enable the BGP graceful restart capability. The default values are optimal for most network deployments, and these values should be adjusted only by an experienced network operator.

SUMMARY STEPS

1. **enable**
2. **configure** { **terminal** | **memory** | **network** }
3. **router bgp** *as-number*
4. **bgp graceful-restart**
5. **exit**

DETAILED STEPS

| | Command or Action | Purpose |
|--------|--|---|
| Step 1 | enable Example: Router> enable | Enables higher privilege levels, such as privileged EXEC mode. <ul style="list-style-type: none"> • Enter your password if prompted. |
| Step 2 | configure { terminal memory network } Example: Router# configure terminal | Enters global configuration mode. |

| | Command or Action | Purpose |
|--------|---|---|
| Step 3 | <code>router bgp as-number</code> Example: Router(config)# router bgp 101 | Enters router configuration mode and creates a BGP routing process. |
| Step 4 | <code>bgp graceful-restart [restart-time seconds] [stalepath-time seconds]</code> Example: Router(config-router)# bgp graceful-restart | Enables the BGP graceful restart capability and BGP NSF awareness. If you enter this command after the BGP session has been established, you must restart the session for the capability to be exchanged with the BGP neighbor. Use this command on the restarting router and all of its peers (NSF-capable and NSF-aware). |
| Step 5 | Router(config-router)# <code>exit</code> Example: Router(config-router)# exit | Exits router configuration mode and enters global configuration mode. |

Troubleshooting Tips

To troubleshoot the NSF feature, use the following commands in privileged EXEC mode, as needed:

| Command | Purpose |
|---|---|
| Router# <code>debug ip bgp</code> | Displays open messages that advertise the graceful restart capability. |
| Router# <code>debug ip bgp updates</code> | Displays sent and received EOR messages. The EOR message is used by the NSF-aware router to start the stalepath timer, if configured. |
| Router# <code>debug ip bgp event</code> | Displays graceful restart timer events, such as the restart timer and the stalepath timer. |
| Router> <code>show ip bgp</code> | Displays entries in the BGP routing table. The output from this command will display routes that are marked as stale by displaying the letter “S” next to each stale route. |
| Router# <code>show ip bgp neighbor</code> | Displays information about the TCP and BGP connections to neighbor devices. When enabled, the graceful restart capability is displayed in the output of this command. |

What to do next

If the **bgp graceful-restart** command has been issued after the BGP session has been established, you must reset by issuing the **clear ip bgp *** command or by reloading the router before graceful restart capabilities will be exchanged. For more information about resetting BGP sessions and using the **clear ip bgp** command, refer to the following documents:

BGP Command Reference Documentation

http://www.cisco.com/univercd/cc/td/doc/product/software/ios122/122newft/122tcr/122tip2r/bgp_r/index.htm

BGP Configuration Documentation

http://www.cisco.com/univercd/cc/td/doc/product/software/ios122/122cgcr/fipr_c/ipcprt2/1cfbgp.htm

Configuring BGP NSF Awareness Timers

BGP NSF Awareness Timers

This section documents the configuration of the BGP graceful restart timers.

- (Optional) The **restart-time** argument determines how long peer routers will wait to delete stale routes before a BGP open message is received. The default value is 120 seconds.
- (Optional) The **stalepath-time** argument determines how long a router will wait before deleting stale routes after an end of record (EOR) message is received from the restarting router. The default value is 360 seconds.

Restrictions

The configuration of the restart and stalepath timers is not required to enable the BGP graceful restart capability. The default values are optimal for most network deployments, and these values should be adjusted only by an experienced network operator.

SUMMARY STEPS

1. **enable**
2. **configure** { **terminal** | **memory** | **network** }
3. **router bgp** *as-number*
4. **bgp graceful-restart restart-time** *seconds*
5. **bgp graceful-restart stalepath-time** *seconds*
6. **exit**

DETAILED STEPS

| | Command or Action | Purpose |
|--------|--|--|
| Step 1 | enable Example: Router> enable | Enables higher privilege levels, such as privileged EXEC mode. • Enter your password if prompted. |
| Step 2 | configure { terminal memory network } Example: Router# configure terminal | Enters global configuration mode. |
| Step 3 | router bgp <i>as-number</i> Example: Router(config)# router bgp 101 | Enters router configuration mode and creates a BGP routing process. |
| Step 4 | bgp graceful-restart [restart-time <i>seconds</i>] [stalepath-time <i>seconds</i>] Example: Router(config-router)# bgp graceful-restart restart-time 130 | Enables the BGP graceful restart capability and BGP NSF awareness. The restart-time argument determines how long peer routers will wait to delete stale routes before a BGP open message is received. The default value is 120 seconds. The configurable range is from 1 to 3600 seconds. |
| Step 5 | bgp graceful-restart [restart-time <i>seconds</i>] [stalepath-time <i>seconds</i>] Example: Router(config-router)# bgp graceful-restart stalepath-time 350 | Enables the BGP graceful restart capability and BGP NSF awareness. The stalepath-time argument determines how long a router will wait before deleting stale routes after an end of record (EOR) message is received from the restarting router. The default value is 360 seconds. The configurable range is from 1 to 3600 seconds. |
| Step 6 | Router(config-router)# exit Example: Router(config-router)# exit | Exits router configuration mode and enters global configuration mode. |

What to do next

If the **bgp graceful-restart** command has been issued after the BGP session has been established, you must reset by issuing the **clear ip bgp *** command or by reloading the router before graceful restart capabilities will be exchanged. For more information about resetting BGP sessions and using the **clear ip bgp** command, refer to the following documents:

BGP Command Reference Documentation

http://www.cisco.com/univercd/cc/td/doc/product/software/ios122/122newft/122tcr/122tip2r/bgp_r/index.htm

BGP Configuration Documentation

http://www.cisco.com/univercd/cc/td/doc/product/software/ios122/122cgcr/fipr_c/ipcprt2/1cfbgp.htm

Verifying the Configuration of BGP Nonstop Forwarding Awareness

Use the following steps to verify the local configuration of BGP NSF awareness on a router and to verify the configuration NSF awareness on peer routers in a BGP network.

SUMMARY STEPS

1. `enable`
2. `show running-config`
3. `show ip bgp neighbors ip-address`

DETAILED STEPS

| | Command or Action | Purpose |
|--------|---|--|
| Step 1 | <code>enable</code> Example: Router> <code>enable</code> | Enables higher privilege levels, such as privileged EXEC mode. • Enter your password if prompted. |
| Step 2 | <code>show running-config</code> Example: Router# <code>show running-config</code> | Displays the running configuration on the local router. The output will display the configuration of the bgp graceful-restart command in the BGP section. |
| Step 3 | <code>show ip bgp neighbors ip-address</code> Example: Router(config)# <code>show ip bgp neighbors 10.0.0.1</code> | Displays information about TCP and BGP connections to neighbors. “Graceful Restart Capabilty:advertised and received” will be displayed for each neighbor that has exchanged graceful restart capabilities with this router. |

Configuration Examples for Nonstop Forwarding

- [Configuring BGP NSF Awareness Example, page 10](#)
- [Configuring the Restart Time for BGP NSF Awareness, page 10](#)
- [Configuring the Stalepath Time for BGP NSF Awareness, page 10](#)
- [Verifying BGP NSF Awareness, page 10](#)

Configuring BGP NSF Awareness Example

The following example configures BGP NSF awareness on a router that is running BGP:

```
router# configure terminal
router(config)# router bgp 101
router(config-router)# bgp graceful-restart
```

Configuring the Restart Time for BGP NSF Awareness

The following example configures BGP NSF awareness on a router that is running BGP and sets the restart time to 130 seconds. The configuration of this timer is optional and the preconfigured default value is optimal for most network deployments.

```
router# configure terminal
router(config)# router bgp 101
router(config-router)# bgp graceful-restart restart-time 130
```

Configuring the Stalepath Time for BGP NSF Awareness

The following example configures BGP NSF awareness on a router that is running BGP and sets the stale path time to 350 seconds. The configuration of this timer is optional and the preconfigured default value is optimal for most network deployments.

```
router# configure terminal
router(config)# router bgp 101
router(config-router)# bgp graceful-restart stalepath-time 350
```

Verifying BGP NSF Awareness

To verify NSF for BGP, you must check that the graceful restart function is configured on the SSO-enabled networking device and on the neighbor devices. Perform the following steps:

-
- Step 1** Verify that “bgp graceful-restart” appears in the BGP configuration of the SSO-enabled router by entering the **show running-config** command:

```
Router# show running-config
.
.
.
router bgp 120
.
.
.
bgp graceful-restart
  neighbor 10.2.2.2 remote-as 300
.
.
.
```

- Step 2** Repeat step 1 on each of the BGP neighbors.

- Step 3** On peer devices (NSF-capable and NSF-aware), verify that the graceful restart function is shown as both advertised and received, and confirm the address families that have the graceful restart capability. If no address families are listed, then BGP NSF awareness is not enabled:

```
router#show ip bgp neighbors x.x.x.x

BGP neighbor is 192.168.2.2, remote AS YY, external link
BGP version 4, remote router ID 192.168.2.2
BGP state = Established, up for 00:01:18
Last read 00:00:17, hold time is 180, keepalive interval is 60 seconds
Neighbor capabilities:
  Route refresh:advertised and received(new)
  Address family IPv4 Unicast:advertised and received
  Address family IPv4 Multicast:advertised and received
  Graceful Restart Capabilty:advertised and received
    Remote Restart timer is 120 seconds
    Address families preserved by peer:
      IPv4 Unicast, IPv4 Multicast
  Received 1539 messages, 0 notifications, 0 in queue
  Sent 1544 messages, 0 notifications, 0 in queue
  Default minimum time between advertisement runs is 30 seconds
```

Where to Go Next

For more information about NSF and SSO configuration, refer to the following documents:

- *Cisco Nonstop Forwarding*, Release 12.0(24)S
<http://www.cisco.com/univercd/cc/td/doc/product/software/ios120/120newft/120limit/120s/120s24/nsf24s.htm>
- *Stateful Switchover*, Release 12.0(24)S
<http://www.cisco.com/univercd/cc/td/doc/product/software/ios120/120newft/120limit/120s/120s24/sso24s.htm>

Additional References

For additional information related to BGP NSF Awareness, refer to the following documents:

Related Documents

| Related Topic | Document Title |
|-------------------------|--|
| BGP commands | “BGP Commands” chapter in the <i>Cisco IOS IP Command Reference, Volume 2 of 3: Routing Protocols</i> , Release 12.2 |
| BGP configuration tasks | “Configuring BGP” chapter in the <i>Cisco IOS IP Configuration Guide</i> , Release 12.2 |

Standards

| Standards ¹ | Title |
|---|-------|
| No new or modified standards are supported by this feature, and support for existing standards has not been modified by this feature. | — |

1. Not all supported standards are listed.

MIBs

| MIBs ¹ | MIBs Link |
|---|---|
| No new or modified MIBs are supported by this feature, and support for existing MIBs has not been modified by this feature. | To obtain lists of supported MIBs by platform and Cisco IOS release, and to download MIB modules, go to the Cisco MIB website on Cisco.com at the following URL: http://www.cisco.com/public/sw-center/netmgmt/cmtk/mibs.shtml |

1. Not all supported MIBs are listed.

To locate and download MIBs for selected platforms, Cisco IOS releases, and feature sets, use Cisco MIB Locator found at the following URL:

<http://tools.cisco.com/ITDIT/MIBS/servlet/index>

If Cisco MIB Locator does not support the MIB information that you need, you can also obtain a list of supported MIBs and download MIBs from the Cisco MIBs page at the following URL:

<http://www.cisco.com/public/sw-center/netmgmt/cmtk/mibs.shtml>

To access Cisco MIB Locator, you must have an account on Cisco.com. If you have forgotten or lost your account information, send a blank e-mail to cco-locksmith@cisco.com. An automatic check will verify that your e-mail address is registered with Cisco.com. If the check is successful, account details with a new random password will be e-mailed to you. Qualified users can establish an account on Cisco.com by following the directions found at this URL:

<http://www.cisco.com/register>

RFCs

| RFCs ¹ | Title |
|-------------------------------|---|
| draft-ietf-idr-restart-06.txt | <i>Graceful Restart Mechanism for BGP</i> |

1. Not all supported RFCs are listed.

Technical Assistance

| Description | Link |
|---|---|
| Technical Assistance Center (TAC) home page, containing 30,000 pages of searchable technical content, including links to products, technologies, solutions, technical tips, tools, and lots more. Registered Cisco.com users can log in from this page to access even more content. | TAC Home Page: http://www.cisco.com/public/support/tac/home.shtml BGP Support Page: http://www.cisco.com/cgi-bin/Support/browse/psp_view.pl?p=Inter networking:BGP |

Command Reference

This section documents new and modified commands. All other commands used with this feature are documented in the Cisco IOS command reference publications for various releases.

New Commands

- [bgp graceful-restart](#)

Modified Commands

- [show ip bgp](#)
- [show ip bgp neighbors](#)

bgp graceful-restart

To enable the Border Gateway Protocol (BGP) graceful restart capability, use the **bgp graceful-restart** command in router configuration mode. To remove this command from the configuration file and restore the system to its default condition with respect to this command, use the **no** form of this command.

bgp graceful-restart [*restart-time seconds* | *stalepath-time seconds*]

no bgp graceful-restart [*restart-time seconds* | *stalepath-time seconds*]

Syntax Description

| | |
|-----------------------|---|
| restart-time | (Optional) Used to set the maximum time to wait for a graceful-restart-capable neighbor to come back up after a restart. The default is 120 seconds. |
| stalepath-time | (Optional) Used to set the maximum time to hold on to the stale paths of a gracefully restarted peer. All stale paths are deleted after the expiration of this timer. The default is 360 seconds. |
| <i>seconds</i> | (Optional) The restart-time or stalepath-time value in number of seconds. The valid range is from 1 to 3600 seconds. |

Defaults

BGP Cisco Nonstop Forwarding (NSF) capabilities are disabled.

Command Modes

Router configuration

Command History

| Release | Modification |
|-----------|---|
| 12.0(22)S | This command was introduced. |
| 12.2(15)T | This command was integrated into Cisco IOS Release 12.2(15)T. |

Usage Guidelines

The **bgp graceful-restart** command is used to enable the graceful restart mechanism on a router in a BGP network. The graceful restart mechanism supports both NSF awareness and NSF capabilities. A router that is NSF-aware functions like a router that is NSF-capable with one exception: an NSF-aware router is incapable of performing a Stateful Switchover (SSO) operation.

The BGP graceful restart capability is negotiated in the OPEN message. If the user enters the **bgp graceful-restart** command after the BGP session is established, the session will need to be restarted.

When you enter the **bgp graceful-restart** command, the **bgp graceful-restart restart-time** and **bgp graceful-restart stalepath-time** commands are enabled by default. After the **bgp graceful-restart** command is used to configure the graceful restart capability, you may tune the configuration using the **restart-time** and **stalepath-time** keywords. If you do not first configure the graceful restart capability using the **bgp graceful-restart** command, the tuning values will not appear in the configuration file.

We recommend that the **bgp graceful-restart restart-time** and **bgp graceful-restart stalepath-time** commands remain set to their default values.

The restart time should not be set to a time greater than the holdtime that is carried in the OPEN message. To deal with possible consecutive restarts, a route (from a restarting router) that was previously marked as stale shall be deleted.



Note The configuration of the restart and stalepath timers is not required to enable the BGP graceful restart capability. The default values are optimal for most network deployments, and these values should be adjusted only by an experienced network operator.

Examples

The following example shows how to configure the BGP graceful restart capability. Enter one command per line:

```
Router# configure terminal
Router(config)# router bgp 101
Router(config-router)# bgp graceful-restart
```

The following example configures BGP NSF awareness on a router that is running BGP and sets the restart time to 240 seconds.

```
Router# configure terminal
router(config)# router bgp 101
router(config-router)# bgp graceful-restart restart-time 130
```

The following example configures BGP NSF awareness on a router that is running BGP and sets the stale path time to 240 seconds.

```
Router# configure terminal
router(config)# router bgp 101
router(config-router)# bgp graceful-restart stalepath-time 350
```

Related Commands

| Command | Description |
|---------------------------------------|--|
| show ip bgp | Displays entries in the BGP routing table. |
| show ip bgp neighbors | Displays information about the TCP and BGP connections to neighbors. |

show ip bgp

To display entries in the Border Gateway Protocol (BGP) routing table, use the **show ip bgp** command in user EXEC mode.

```
show ip bgp [network] [network-mask] [longer-prefixes]
```

| Syntax Description | Parameter | Description |
|--------------------|------------------------|--|
| | <i>network</i> | (Optional) Network number, entered to display a particular network in the BGP routing table. |
| | <i>network-mask</i> | (Optional) Displays all BGP routes matching the address and mask pair. |
| | longer-prefixes | (Optional) Displays the route and more specific routes. |

Command Modes User EXEC

| Command History | Release | Modification |
|-----------------|-----------|--|
| | 10.0 | This command was introduced. |
| | 12.0 | The display of prefix advertisement statistics was added. |
| | 12.0(6)T | The display of a message indicating support for route refresh capability was added. |
| | 12.0(22)S | A new status code indicating stale routes was added to support BGP graceful restart. |
| | 12.2(15)T | A new status code indicating stale routes was added to support BGP graceful restart. |

Examples

The following example is from the **show ip bgp** command. The “stale” status code is shown in the output:

```
Router# show ip bgp

BGP table version is 9, local router ID is a.a.a.a
Status codes:s suppressed, d damped, h history, * valid, > best, i - internal,
               S Stale
Origin codes:i - IGP, e - EGP, ? - incomplete

   Network          Next Hop          Metric LocPrf Weight Path
S> 10.100.100.0/24  192.168.2.2            0             0 65002 i
S> 10.10.0.0         192.168.2.2            0             0 65002 65003 i
S> 10.20.0.0         192.168.2.2            0             0 65002 65003 i
S> 10.30.0.0/8       192.168.2.2            0             0 65002 65003 i
S> 10.40.33.0/24     192.168.2.2            0             0 65002 65003 i
*>                  0.0.0.0              0             32768 i
S> 10.50.0.0/8       192.168.2.2            0             0 65002 65003 i
S> 10.60.100.0       192.168.2.2            0             0 65002 i
S> 10.70.200.0       192.168.2.2            0             0 65002 i
```

[Table 1](#) describes the significant fields shown in the displays.

Table 1 *show ip bgp Field Descriptions*

| Field | Description |
|-------------------|---|
| BGP table version | Internal version number of the table. This number increments when the table changes. |
| local router ID | IP address of the router. |
| Status codes | Status of the table entry. The status is displayed at the beginning of each line in the table. It can be one of the following values: s—The table entry is suppressed. d—The table entry is damped. h—History of the table entry. *—The table entry is valid. >—The table entry is the best entry to use for that network. i—The table entry was learned via an internal BGP (iBGP) session. S—The table entry is stale. |
| Origin codes | Origin of the entry. The origin code is placed at the end of each line in the table. It can be one of the following values: i—Entry originated from Interior Gateway Protocol (IGP) and was advertised with a network router configuration command. e—Entry originated from Exterior Gateway Protocol (EGP). ?—Origin of the path is not clear. Usually, this is a router that is redistributed into BGP from an IGP. |
| Network | IP address of a network entity. |
| Next hop | IP address of the next system that is used when forwarding a packet to the destination network. An entry of 0.0.0.0 indicates that the router has some non-BGP routes to this network. |
| Metric | The value of the inter-autonomous system metric. |
| BGP table version | Internal version number of the table. This number is incremented whenever the table changes. |
| local router ID | IP address of the router. |
| Status codes | Status of the table entry. The status is displayed at the beginning of each line in the table. It can be one of the following values: s—The table entry is suppressed. d—The table entry is damped. h—History of the table entry. *—The table entry is valid. >—The table entry is the best entry to use for that network. i—The table entry was learned via an internal BGP (iBGP) session. S—The table entry is stale. |

The following example of the **show ip bgp** command. The output is filtered to show only the 10.0.0.0 network. The “stale” status code is also shown in the output.

```
Router# show ip bgp 10.0.0.0
```

```
BGP routing table entry for 10.0.0.0/8, version 7
Paths:(1 available, best #1, table Default-IP-Routing-Table)
```

■ show ip bgp

```

Not advertised to any peer
65002 65003, (stale)
  192.168.2.2 from 192.168.2.2 (0.0.0.0)
    Origin IGP, localpref 100, valid, external, best
Router#

```

Table 2 describes the significant fields shown in the display.

Table 2 show ip bgp Field Descriptions

| Field | Description |
|----------------------------|--|
| BGP routing table entry... | Specifies for which network or peer the routing table output is generated. |
| version | Internal version number of the table. This number is incremented whenever the table changes. |
| Paths | Autonomous system paths to the destination network. There can be one entry in this field for each autonomous system in the path. The number paths and best paths is displayed. |
| (stale) | Indicates that the following path for the specified autonomous system is marked as “stale” during a graceful restart process. |

Related Commands

| Command | Description |
|---------------------------------------|--|
| bgp graceful-restart | Enables the BGP graceful restart capability. |
| show ip bgp neighbors | Displays information about the TCP and BGP connections to neighbors. |

show ip bgp neighbors

To display information about TCP/IP and Border Gateway Protocol (BGP) connections to neighbors, use the **show ip bgp neighbors** command in EXEC mode.

```
show ip bgp neighbors [neighbor-address] [received-routes | routes | advertised-routes | {paths
regexp} | dampened-routes] [received prefix-filter]
```

| Syntax Description | | |
|-------------------------------|------------|---|
| <i>neighbor-address</i> | (Optional) | Address of the neighbor whose routes you have learned from. If you omit this argument, all neighbors are displayed. |
| received-routes | (Optional) | Displays all received routes (both accepted and rejected) from the specified neighbor. |
| routes | (Optional) | Displays all routes that are received and accepted. This is a subset of the output from the received-routes keyword. |
| advertised-routes | (Optional) | Displays all the routes the router has advertised to the neighbor. |
| paths <i>regexp</i> | (Optional) | Regular expression that is used to match the paths received. |
| dampened-routes | (Optional) | Displays the dampened routes to the neighbor at the IP address specified. |
| received prefix-filter | (Optional) | Displays the configured prefix list for the specified IP address. |

Command Modes EXEC

| Command History | Release | Modification |
|-----------------|------------|--|
| | 12.0 | This command was introduced. |
| | 11.2 | The received-routes keyword was added. |
| | 12.0(21)ST | This command was updated to display MPLS label information. |
| | 12.0(22)S | This command was integrated into Cisco IOS Release 12.0(22)S. Support for the Cisco 12000 series routers (Engine 0 and Engine 2) was added. The received prefix-filter keyword was added. |
| | 12.0(22)S | Support for the BGP graceful restart capability was integrated into the output. |
| | 12.2(15)T | Support for the BGP graceful restart capability was integrated into the output. |

Examples

The following is sample output from the **show ip bgp neighbors** command in privileged EXEC mode. The output is filtered to show only the 172.16.254.3 neighbor. The output also shows that the graceful restart capability has been exchanged with this neighbor.

```
Router# show ip bgp neighbors 172.16.254.3

BGP neighbor is 172.16.254.3, remote AS 150, internal link
  BGP version 4, remote router ID 172.16.254.3
  BGP state = Established, up for 19:24:07
  Last read 00:00:06, hold time is 180, keepalive interval is 60 seconds
  Neighbor capabilities:
```

```
show ip bgp neighbors
```

```

Route refresh:advertised and received(new)
Address family IPv4 Unicast:advertised and received
Graceful Restart Capabilty:advertised and received
  Remote Restart timer is 120 seconds
  Address families preserved by peer:
    IPv4 Unicast
Received 4231 messages, 0 notifications, 0 in queue
Sent 4167 messages, 0 notifications, 0 in queue
Default minimum time between advertisement runs is 5 seconds

For address family:IPv4 Unicast
BGP table version 159559, neighbor version 159559
Index 90, Offset 11, Mask 0x4
Route refresh request:received 0, sent 0
10031 accepted prefixes consume 441364 bytes
Prefix advertised 29403, suppressed 0, withdrawn 9801
Number of NLRIs in the update sent:max 242, min 0

Connections established 2; dropped 1
Last reset 19:26:54, due to NSF peer closed the session
Connection state is ESTAB, I/O status:1, unread input bytes:0
Local host:150.254.254.2, Local port:11005
Foreign host:172.16.254.3, Foreign port:179

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

Event Timers (current time is 0x4371A84):
Timer           Starts      Wakeups      Next
Retrans         1380         22           0x0
TimeWait        0             0           0x0
AckHold         1377         870          0x0
SendWnd         0             0           0x0
KeepAlive       0             0           0x0
GiveUp          0             0           0x0
PmtuAger        0             0           0x0
DeadWait        0             0           0x0

iss:1875330775  snduna:1875639119  sndnxt:1875639119   sndwnd: 16308
irs:3577079138  rcvnxt:3577393901  rcvwnd: 16137   delrcvwnd: 247

SRTT:300 ms, RTTO:607 ms, RTV:3 ms, KRTT:0 ms
minRTT:0 ms, maxRTT:408 ms, ACK hold:200 ms
Flags:higher precedence, nagle

Datagrams (max data segment is 536 bytes):
Rcvd:2984 (out of order:1), with data:1800, total data bytes:314762
Sent:3190 (retransmit:22, fastretransmit:0), with data:1751, total data bytes:308343

```

Table 3 describes the significant fields shown in the display.

Table 3 *show ip bgp neighbors* Field Descriptions

| Field | Description |
|---------------|--|
| BGP neighbor | IP address of the BGP neighbor and its autonomous system number. If the neighbor is in the same autonomous system as the router, then the link between them is internal; otherwise, it is considered external. |
| remote AS | Autonomous system of the neighbor. |
| external link | Indicates that this peer is an EBGP peer. |

Table 3 *show ip bgp neighbors Field Descriptions (continued)*

| Field | Description |
|------------------------------|---|
| BGP version | BGP version being used to communicate with the remote router; the router ID (an IP address) of the neighbor is also specified. |
| remote router ID | IP address of the neighbor. |
| BGP state | Internal state of this BGP connection. |
| up for | Amount of time, in seconds, that the underlying TCP connection has been in existence. |
| Last read | Time that BGP last read a message from this neighbor. |
| hold time | Maximum amount of time that can elapse between messages from the peer. |
| keepalive interval | Time period, in seconds, between sending keepalive packets, which help ensure that the TCP connection is up. |
| Neighbor capabilities | BGP capabilities advertised and received from this neighbor. |
| MPLS Label capability | Indicates that MPLS labels are both sent and received by the EBGP peer. |
| Address family IPv4 Unicast: | IP Version 4 unicast-specific properties of this neighbor. |
| Graceful Restart Capability: | The status of the graceful restart capability. "Advertised and received" is displayed when the graceful restart capability has been successfully exchanged between two routers. |
| Received notifications | Number of total BGP messages received from this peer, including keepalives. |
| Sent notifications | Number of error messages received from the peer. |
| Received notifications | Total number of BGP messages that have been sent to this peer, including keepalives. |
| Sent notifications | Number of error messages the router has sent to this peer. |
| Route refresh request: | Number of route refresh requests sent and received from this neighbor. |
| advertisement runs | Value of minimum advertisement interval. |
| For address family: | Address family to which the following fields refer. |
| BGP table version | Indicates that the neighbor has been updated with this version of the primary BGP routing table. |
| neighbor version | Number used by the software to track the prefixes that have been sent and those that must be sent to this neighbor. |
| Community attribute | Appears if the neighbor send-community command is configured for this neighbor. |
| Inbound path policy | Indicates if an inbound policy is configured. |
| Outbound path policy | Indicates if an outbound policy is configured. |
| uni-in | Name of inbound route map for the unicast address family. |
| uni-out | Name of outbound route map for the unicast address family. |
| mul-in | Name of inbound route map for the multicast address family. |
| mul-out | Name of outbound route map for the multicast address family. |

Table 3 *show ip bgp neighbors Field Descriptions (continued)*

| Field | Description |
|----------------------------|--|
| Sending Prefix & Label | Indicates that the EBGp peer sends MPLS labels with its routes. |
| accepted prefixes | Number of prefixes accepted. |
| Prefix advertised | Number of prefixes advertised. |
| suppressed | Number of prefixes suppressed. |
| withdrawn | Number of prefixes withdrawn. |
| Connections established | Number of times the router has established a TCP connection and the two peers have agreed to speak BGP with each other. |
| dropped | Number of times that a good connection has failed or been taken down. |
| Last reset | Elapsed time since this peering session was last reset. |
| Connection state | State of BGP peer. |
| unread input bytes | Number of bytes of packets still to be processed. |
| Local host, Local port | Peering address of local router, plus port. |
| Foreign host, Foreign port | Peering address of the neighbor. |
| Event Timers | Table displays the number of starts and wakeups for each timer. |
| iss | Initial send sequence number. |
| snduna | Last send sequence number the local host sent but has not received an acknowledgment for. |
| sndnxt | Sequence number the local host will send next. |
| sndwnd | TCP window size of the remote host. |
| irs | Initial receive sequence number. |
| rcvnxt | Last receive sequence number the local host has acknowledged. |
| rcvwnd | TCP window size of the local host. |
| delrcvwnd | Delayed receive window—data the local host has read from the connection, but has not yet subtracted from the receive window the host has advertised to the remote host. The value in this field gradually increases until it is larger than a full-sized packet, at which point it is applied to the rcvwnd field. |
| SRTT | A calculated smoothed round-trip timeout. |
| RTTO | Round-trip timeout. |
| RTV | Variance of the round-trip time. |
| KRTT | New round-trip timeout (using the Karn algorithm). This field separately tracks the round-trip time of packets that have been re-sent. |
| minRTT | Smallest recorded round-trip timeout (hard wire value used for calculation). |
| maxRTT | Largest recorded round-trip timeout. |
| ACK hold | Time the local host will delay an acknowledgment in order to piggyback data on it. |
| Flags | IP precedence of the BGP packets. |
| Datagrams: Rcvd | Number of update packets received from a neighbor. |

Table 3 *show ip bgp neighbors Field Descriptions (continued)*

| Field | Description |
|------------------|--|
| with data | Number of update packets received with data. |
| total data bytes | Total bytes of data. |
| Sent | Number of update packets sent. |
| with data | Number of update packets with data sent. |
| total data bytes | Total number of data bytes. |

The following is sample output from the **show ip bgp neighbors** command issued with the **advertised-routes** keyword in privileged EXEC mode. The output is filtered to display advertised routes for only the 172.16.232.178 neighbor.

```
Router# show ip bgp neighbors 172.16.232.178 advertised-routes

BGP table version is 27, local router ID is 172.16.232.181
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal
Origin codes: i - IGP, e - EGP, ? - incomplete

   Network          Next Hop           Metric LocPrf Weight Path
*>i1110.0.0.0      172.16.232.179         0    100     0  ?
*> 200.2.2.0       0.0.0.0              0           32768 i
```

The following is sample output from the **show ip bgp neighbors** command issued with the **routes** keyword in privileged EXEC mode. The output is filtered to show only routes that have been received and accepted by the 172.16.232.178 neighbor.

```
Router# show ip bgp neighbors 172.16.232.178 routes

BGP table version is 27, local router ID is 172.16.232.181
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal
Origin codes: i - IGP, e - EGP, ? - incomplete

   Network          Next Hop           Metric LocPrf Weight Path
*> 10.0.0.0         172.16.232.178         40           0 10 ?
*> gg.0.0.0        172.16.232.178         40           0 10 ?
```

[Table 4](#) describes the significant fields shown in the displays.

Table 4 *show ip bgp neighbors advertised-routes and routes Field Descriptions*

| Field | Description |
|-------------------|--|
| BGP table version | Internal version number of the table. This number increments when the table changes. |
| local router ID | IP address of the router. |

Table 4 *show ip bgp neighbors advertised-routes and routes Field Descriptions (continued)*

| Field | Description |
|--------------|--|
| Status codes | Status of the table entry. The status is displayed at the beginning of each line in the table. It can be one of the following values: s—The table entry is suppressed. d—The table entry is dampened and will not be advertised to BGP neighbors. h—The table entry does not contain the best path based on historical information. *—The table entry is valid. >—The table entry is the best entry to use for that network. i—The table entry was learned via an internal BGP (iBGP) session. |
| Origin codes | Origin of the entry. The origin code is placed at the end of each line in the table. It can be one of the following values: i—Entry originated from Interior Gateway Protocol (IGP) and was advertised with a network router configuration command. e—Entry originated from Exterior Gateway Protocol (EGP). ?—Origin of the path is not clear. Usually, this is a router that is redistributed into BGP from an IGP. |
| Network | IP address of a network entity. |
| Next Hop | IP address of the next system that is used when forwarding a packet to the destination network. An entry of 0.0.0.0 indicates that the router has some non-BGP routes to this network. |
| Metric | If shown, this is the value of the inter-autonomous system metric. This field is frequently not used. |
| LocPrf | Local preference value as set with the set local-preference route-map configuration command. The default value is 100. |
| Weight | Weight of the route as set via autonomous system filters. |
| Path | Autonomous system paths to the destination network. There can be one entry in this field for each autonomous system in the path. |

The following is sample output from the **show ip bgp neighbors** command with the **paths** keyword in privileged EXEC mode:

```
Router# show ip bgp neighbors 171.69.232.178 paths ^10
Address      Refcount Metric Path
0x60E577B0      2      40 10 ?
```

[Table 5](#) describes the significant fields shown in the display.

Table 5 *show ip bgp neighbors paths Field Descriptions*

| Field | Description |
|----------|--|
| Address | Internal address where the path is stored. |
| Refcount | Number of routes using that path. |

Table 5 *show ip bgp neighbors paths Field Descriptions*

| Field | Description |
|--------|---|
| Metric | Multi Exit Discriminator (MED) metric for the path. (The name of this metric for BGP versions 2 and 3 is INTER_AS.) |
| Path | Autonomous system path for that route, followed by the origin code for that route. |

The following is sample output from the **show ip bgp neighbors** command issued with the **received prefix-filter** keyword in privileged EXEC mode. The output is filtered to show only routes that are filtered by a prefix-list for the 192.168.20.72 neighbor:

```
Router# show ip bgp neighbor 192.168.20.72 received prefix-filter

Address family:IPv4 Unicast
ip prefix-list 192.168.20.72:1 entries
  seq 5 deny 10.0.0.0/8 le 32
```

[Table 6](#) describes the significant fields shown in the display.

Table 6 *show ip bgp neighbors received prefix-filter Field Descriptions*

| Field | Description |
|-----------------|--|
| Address family: | Configured address family mode. |
| ip prefix-list | Configured prefix list for the specified neighbor. |

■ show ip bgp neighbors