



Configuring Bidirectional Forwarding Detection

- [Prerequisites for Bidirectional Forwarding Detection, on page 1](#)
- [Restrictions for Bidirectional Forwarding Detection, on page 1](#)
- [Information About Bidirectional Forwarding Detection, on page 2](#)
- [How to Configure Bidirectional Forwarding Detection, on page 5](#)
- [Configuration Examples for Bidirectional Forwarding Detection, on page 10](#)
- [Additional References for Bidirectional Forwarding Detection, on page 11](#)
- [Feature Information for Bidirectional Forwarding Detection, on page 12](#)

Prerequisites for Bidirectional Forwarding Detection

Prerequisites for BFD include:

- The switch's feature set is IP Base or higher. The IP Base feature set supports only Enhanced Interior Gateway Routing Protocol (EIGRP) stub routing, without BFD. The IP service feature set supports EIGRP with BFD.
- IP routing must be enabled on all participating switches
- Before BFD is deployed, configure one of the IP routing protocols supported by BFD on the switches. Also, implement fast convergence for the routing protocol that you plan to use.

Restrictions for Bidirectional Forwarding Detection

Restrictions for BFD include:

- BFD works only for directly connected neighbors. BFD neighbors must be no more than one IP hop away. Multihop configurations are not supported.
- The switch supports up to 100 BFD sessions with a minimum hello interval of 100 ms and a multiplier of 3. The multiplier specifies the minimum number of consecutive packets that can be missed before a session is declared down.
- To enable echo mode the peer system must be configured with the no ip redirects command.

Information About Bidirectional Forwarding Detection

BFD Operation

BFD provides a low-overhead, short-duration method of detecting failures in the forwarding path between two adjacent routers, including the interfaces, data links, and forwarding planes.

BFD is a detection protocol that you enable at the interface and routing protocol levels. Cisco supports the BFD asynchronous mode, which depends on the sending of BFD control packets between two systems to activate and maintain BFD neighbor sessions between routers. Therefore, in order for a BFD session to be created, you must configure BFD on both systems (or BFD peers). Once BFD has been enabled on the interfaces and at the router level for the appropriate routing protocols, a BFD session is created, BFD timers are negotiated, and the BFD peers will begin to send BFD control packets to each other at the negotiated interval.

Cisco supports BFD echo mode. Echo packets are sent by the forwarding engine and are forwarded back along the same path to perform detection. The BFD session at the other end does not participate in the actual forwarding of the echo packets.

This section includes the following subsections:

BFD Version Interoperability

The switch supports BFD Version 1 as well as BFD Version 0. All BFD sessions come up as Version 1 by default and will be interoperable with Version 0. The system automatically performs BFD version detection, and BFD sessions between neighbors will run in the highest common BFD version between neighbors. For example, if one BFD neighbor is running BFD Version 0 and the other BFD neighbor is running Version 1, the session will run BFD Version 0. The output from the **show bfd neighbors [details]** command will verify which BFD version a BFD neighbor is running.

BFD Session Limits

The minimum number of BFD sessions that can be created varies with the “hello” interval. With “hello” intervals of 100ms, 100 sessions are permitted. More sessions are permitted at larger hello intervals. For a VLAN interface, the minimum “hello” interval is 600ms.

BFD Support for Nonbroadcast Media Interfaces

The BFD feature is supported on VLAN interfaces on the switch.

The **bfd interval** command must be configured on the interface to initiate BFD monitoring.

BFD Support for Nonstop Forwarding with Stateful Switchover

Typically, when a networking device restarts, all routing peers of that device detect that the device went down and then came back up. This transition results in a routing flap, which could spread across multiple routing domains. Routing flaps caused by routing restarts create routing instabilities, which are detrimental to the overall network performance. Nonstop forwarding (NSF) helps to suppress routing flaps in devices that are enabled with stateful switchover (SSO), thereby reducing network instability.

NSF allows for the forwarding of data packets to continue along known routes while the routing protocol information is being restored after a switchover. With NSF, peer networking devices do not experience routing flaps. Data traffic is forwarded through intelligent line cards or dual forwarding processors while the standby RP assumes control from the failed active RP during a switchover. The ability of line cards and forwarding processors to remain up through a switchover and to be kept current with the Forwarding Information Base (FIB) on the active RP is key to NSF operation.

In devices that support dual RPs, SSO establishes one of the RPs as the active processor; the other RP is designated as the standby processor, and then synchronizes information between them. A switchover from the active to the standby processor occurs when the active RP fails, when it is removed from the networking device, or when it is manually taken down for maintenance.

BFD Support for Stateful Switchover

The BFD protocol provides short-duration detection of failures in the path between adjacent forwarding engines. In network deployments that use dual RP switches (to provide redundancy), the switches have a graceful restart mechanism that protects the forwarding state during a switchover between the active RP and the standby RP.

Stateful BFD on the Standby RP

To ensure a successful switchover to the standby RP, the BFD protocol uses checkpoint messages to send session information from the active RP Cisco IOS instance to the standby RP Cisco IOS instance. The session information includes local and remote discriminators, adjacent router timer information, BFD setup information, and session-specific information such as the type of session and the session version. In addition, the BFD protocol sends session creation and deletion checkpoint messages to create or delete a session on the standby RP.

The BFD sessions on the standby RP do not receive or send packets and do not process expired timers. These sessions wait for a switchover to occur and then send packets for any active sessions so that sessions do not time out on adjacent switches.

When the BFD protocol on the standby RP is notified of a switchover it changes its state to active, registers itself with Cisco Express Forwarding so that it can receive packets, and then sends packets for any elements that have expired.

BFD also uses checkpoint messages to ensure that sessions created by clients on the active RP are maintained during a switchover. When a switchover occurs, BFD starts an SSO reclaim timer. Clients must reclaim their sessions within the duration specified by the reclaim timer or else the session is deleted.

Timer values are different based on the number of BFD sessions and the platform.

Table 1: BFD Timer Values on the switch

Maximum Number of BFD Sessions	BFD Session Type	Minimum Timer Value (ms)	Clients	Comments
100	Async/echo	100 multiplier 3	All	A multiple of 5 is recommended for SSO switches.

BFD Support for Static Routing

Unlike dynamic routing protocols, such as OSPF and BGP, static routing has no method of peer discovery. Therefore, when BFD is configured, the reachability of the gateway is completely dependent on the state of the BFD session to the specified neighbor. Unless the BFD session is up, the gateway for the static route is considered unreachable, and therefore the affected routes will not be installed in the appropriate Routing Information Base (RIB).

For a BFD session to be successfully established, BFD must be configured on the interface on the peer and there must be a BFD client registered on the peer for the address of the BFD neighbor. When an interface is used by dynamic routing protocols, the latter requirement is usually met by configuring the routing protocol instances on each neighbor for BFD. When an interface is used exclusively for static routing, this requirement must be met by configuring static routes on the peers.

If a BFD configuration is removed from the remote peer while the BFD session is in the up state, the updated state of the BFD session is not signaled to the static static. This will cause the static route to remain in the RIB. The only workaround is to remove the IPv4 static BFD neighbor configuration so that the static route no longer tracks BFD session state.

Configuring BFD Echo Mode

BFD echo mode is enabled by default, but you can disable it such that it can run independently in each direction.

BFD echo mode works with asynchronous BFD. Echo packets are sent by the forwarding engine and forwarded back along the same path in order to perform detection--the BFD session at the other end does not participate in the actual forwarding of the echo packets. The echo function and the forwarding engine are responsible for the detection process; therefore, the number of BFD control packets that are sent out between two BFD neighbors is reduced. In addition, because the forwarding engine is testing the forwarding path on the remote (neighbor) system without involving the remote system, there is an opportunity to improve the interpacket delay variance, thereby achieving quicker failure detection times than when using BFD Version 0 with BFD control packets for the BFD session.

Echo mode is described as without asymmetry when it is running on both sides (both BFD neighbors are running echo mode).

Prerequisites

BFD must be running on all participating switches.

Before using BFD echo mode, you must disable the sending of Internet Control Message Protocol (ICMP) redirect messages by entering the **no ip redirects** command, in order to avoid high CPU utilization.

The baseline parameters for BFD sessions on the interfaces over which you want to run BFD sessions to BFD neighbors must be configured. See the Configuring BFD Session Parameters on the Interface section for more information.

Restrictions

BFD echo mode, which is supported in BFD Version 1.



Note BFD echo mode does not work in conjunction with Unicast Reverse Path Forwarding (uRPF) configuration. If BFD echo mode and uRPF configurations are enabled, then the sessions will flap.

How to Configure Bidirectional Forwarding Detection

You start a BFD process by configuring BFD on the interface. When the BFD process is started, no entries are created in the adjacency database; in other words, no BFD control packets are sent or received. BFD echo mode, which is supported in BFD Version 1.

BFD echo packets are sent and received, in addition to BFD control packets. The adjacency creation takes places once you have configured BFD support for the applicable routing protocols. This section contains the following procedures:

Configuring BFD Session Parameters on the Interface

Perform this task to configure BFD on an interface by setting the baseline BFD session parameters on the interface. Repeat this task on each interface over which you want to run BFD sessions to BFD neighbors.

SUMMARY STEPS

1. **enable**
2. **configure terminal**
3. **interface** *type number*
4. **bfd interval** *milliseconds min_rx milliseconds multiplier interval-multiplier*
5. **no bfd echo**
6. **end**

DETAILED STEPS

	Command or Action	Purpose
Step 1	enable Example: Switch> enable	Enables privileged EXEC mode. • Enter your password if prompted.
Step 2	configure terminal Example: Switch# configure terminal	Enters global configuration mode.
Step 3	interface <i>type number</i> Example: Switch(config)# interface GigabitEthernet 6/1	Specifies an interface type and number, and places the device in interface configuration mode.
Step 4	bfd interval <i>milliseconds min_rx milliseconds multiplier interval-multiplier</i> Example: Switch(config-if)# no bfd echo	Enables BFD on the interface. Disables BFD echo mode to enable Hardware Off-load.

	Command or Action	Purpose
Step 5	no bfd echo Example: <pre>Switch(config-if)# no bfd echo</pre>	Disables BFD echo mode to enable Hardware Off-load.
Step 6	end Example: <pre>Switch(config-if)# end</pre>	Exits interface configuration mode and returns to privileged EXEC mode.

Configuring BFD Support for Static Routing

Perform this task to configure BFD support for static routing. Repeat the steps in this procedure on each BFD neighbor. For more information, see the “Example: Configuring BFD Support for Static Routing” section

SUMMARY STEPS

1. **enable**
2. **configure terminal**
3. **interface** *type number*
4. **no switchport**
5. **ip address** *ip-address mask*
6. **bfd interval** *milliseconds* **min_rx** *milliseconds* **multiplier** *interval-multiplier*
7. **exit**
8. **ip route static bfd** *interface-type interface-number ip-address* [**group** *group-name* [**passive**]]
9. **ip route** [**vrf** *vrf-name*] *prefix mask* {*ip-address* | *interface-type interface-number* [*ip-address*]} [**dhcp**] [*distance*] [**name** *next-hop-name*] [**permanent** | **track** *number*] [**tag** *tag*]
10. **exit**
11. **show ip static route**
12. **show ip static route bfd**

DETAILED STEPS

	Command or Action	Purpose
Step 1	enable Example: <pre>Switch> enable</pre>	Enables privileged EXEC mode. <ul style="list-style-type: none"> • Enter your password if prompted.
Step 2	configure terminal Example: <pre>Switch# configure terminal</pre>	Enters global configuration mode.

	Command or Action	Purpose
Step 3	interface <i>type number</i> Example: <pre>Switch(config)# interface gigabitethernet 6/1</pre>	Configures an interface and enters interface configuration mode.
Step 4	no switchport Example: <pre>Switch(config)# no switchport</pre>	Changes the interface to Layer 3.
Step 5	ip address <i>ip-address mask</i> Example: <pre>Switch(config-if)# ip address 10.201.201.1 255.255.255.0</pre>	Configures an IP address for the interface.
Step 6	bfd interval <i>milliseconds min_rx milliseconds multiplier interval-multiplier</i> Example: <pre>Switch(config-if)# bfd interval 500 min_rx 500 multiplier 5</pre>	Enables BFD on the interface.
Step 7	exit Example: <pre>Switch(config-if)# exit</pre>	Exits interface configuration mode and returns to global configuration mode.
Step 8	ip route static bfd <i>interface-type interface-number ip-address [group group-name [passive]]</i> Example: <pre>Switch(config)# ip route static bfd serial 2/0 10.1.1.1 group group1 passive</pre>	Specifies a static route BFD neighbor. <ul style="list-style-type: none"> The <i>interface-type</i>, <i>interface-number</i>, and <i>ip-address</i> arguments are required because BFD support exists only for directly connected neighbors.
Step 9	ip route [vrf <i>vrf-name</i>] <i>prefix mask {ip-address interface-type interface-number [ip-address]} [dhcp] [distance] [name next-hop-name] [permanent track number] [tag tag]</i> Example: <pre>Switch(config)# ip route 10.0.0.0 255.0.0.0 Gigabitethernet 6/1 10.201.201.2</pre>	Specifies a static route BFD neighbor.
Step 10	exit Example: <pre>Switch(config)# exit</pre>	Exits global configuration mode and returns to privileged EXEC mode.

	Command or Action	Purpose
Step 11	show ip static route Example: <pre>Switch# show ip static route</pre>	(Optional) Displays static route database information.
Step 12	show ip static route bfd Example: <pre>Switch# show ip static route bfd</pre>	(Optional) Displays information about the static BFD configuration from the configured BFD groups and non-group entries.

Configuring the BFD Slow Timer

This task shows how to change the value of the BFD slow timer. Repeat the steps in this task for each BFD switch.

SUMMARY STEPS

1. **enable**
2. **configure terminal**
3. **bfd slow-timer** *milliseconds*
4. **end**

DETAILED STEPS

	Command or Action	Purpose
Step 1	enable Example: <pre>Switch> enable</pre>	Enables privileged EXEC mode. <ul style="list-style-type: none"> • Enter your password if prompted.
Step 2	configure terminal Example: <pre>Switch# configure terminal</pre>	Enters global configuration mode.
Step 3	bfd slow-timer <i>milliseconds</i> Example: <pre>Switch(config)# bfd slow-timer 12000</pre>	Configures the BFD slow timer.
Step 4	end Example: <pre>Switch(config)# end</pre>	Exits global configuration mode and returns to privileged EXEC mode.

Disabling BFD Echo Mode Without Asymmetry

This task shows how to disable BFD echo mode without asymmetry—no echo packets will be sent by the switch, and the switch will not forward BFD echo packets that are received from any neighbor switches.

Repeat the steps in this task for each BFD switch.

SUMMARY STEPS

1. **enable**
2. **configure terminal**
3. **no bfd echo**
4. **end**

DETAILED STEPS

	Command or Action	Purpose
Step 1	enable Example: Switch> enable	Enables privileged EXEC mode. <ul style="list-style-type: none">• Enter your password if prompted.
Step 2	configure terminal Example: Switch# configure terminal	Enters global configuration mode.
Step 3	no bfd echo Example: Switch(config)# no bfd echo	Disables BFD echo mode. <ul style="list-style-type: none">• Use the no form to disable BFD echo mode.
Step 4	end Example: Switch(config)# end	Exits global configuration mode and returns to privileged EXEC mode.

Monitoring and Troubleshooting BFD

This section describes how to retrieve BFD information for maintenance and troubleshooting. The commands in these tasks can be entered as needed, in any order.

To monitor and troubleshoot BFD, perform the following steps:

SUMMARY STEPS

1. **enable**
2. **show bfd neighbors [details]**
3. **debug bfd [packet | event]**

DETAILED STEPS

	Command or Action	Purpose
Step 1	enable Example: Switch> enable	Enables privileged EXEC mode. <ul style="list-style-type: none"> • Enter your password if prompted.
Step 2	show bfd neighbors [details] Example: Switch# show bfd neighbors details	(Optional) Displays the BFD adjacency database. <ul style="list-style-type: none"> • The details keyword shows all BFD protocol parameters and timers per neighbor.
Step 3	debug bfd [packet event] Example: Switch# debug bfd packet	(Optional) Displays debugging information about BFD packets.

Configuration Examples for Bidirectional Forwarding Detection

Example: Configuring BFD Session Parameters on the Interface

```
Switch# configure terminal
Switch(config)# interface GigabitEthernet 6/1
Switch(config-if)# bfd interval 50 min_rx 50 multiplier 5
Switch(config-if)# no bfd echo
```

Example: Configuring BFD Support for Static Routing

In the following example, the network consists of Device A and Device B. Serial interface 2/0 on Device A is connected to the same network as serial interface 2/0 on Device B. In order for the BFD session to come up, Device B must be configured.

Device A

```
configure terminal
interface Serial 2/0
ip address 10.201.201.1 255.255.255.0
bfd interval 500 min_rx 500 multiplier 5
ip route static bfd Serial 2/0 10.201.201.2
ip route 10.0.0.0 255.0.0.0 Serial 2/0 10.201.201.2
```

Device B

```
configure terminal
interface Serial 2/0
ip address 10.201.201.2 255.255.255.0
```

```
bfd interval 500 min_rx 500 multiplier 5
ip route static bfd Serial 2/0 10.201.201.1
ip route 10.1.1.1 255.255.255.255 Serial 2/0 10.201.201.1
```

Note that the static route on Device B exists solely to enable the BFD session between 10.201.201.1 and 10.201.201.2. If there is no useful static route that needs to be configured, select a prefix that will not affect packet forwarding, for example, the address of a locally configured loopback interface.

In the following example, there is an active static BFD configuration to reach 209.165.200.225 through Ethernet interface 0/0 in the BFD group testgroup. As soon as the static route is configured that is tracked by the configured static BFD, a single hop BFD session is initiated to 209.165.200.225 through Ethernet interface 0/0. The prefix 10.0.0.0/8 is added to the RIB if a BFD session is successfully established.

```
configure terminal
ip route static bfd Ethernet 0/0 209.165.200.225 group testgroup
ip route 10.0.0.0 255.255.255.224 Ethernet 0/0 209.165.200.225
```

In the following example, a BFD session to 209.165.200.226 through Ethernet interface 0/0.1001 is marked to use the group testgroup. That is, this configuration is a passive static BFD. Though there are static routes to be tracked by the second static BFD configuration, a BFD session is not triggered for 209.165.200.226 through Ethernet interface 0/0.1001. The existence of the prefixes 10.1.1.1/8 and 10.2.2.2/8 is controlled by the active static BFD session (Ethernet interface 0/0 209.165.200.225).

```
configure terminal
ip route static bfd Ethernet 0/0 209.165.200.225 group testgroup
ip route 10.0.0.0 255.255.255.224 Ethernet 0/0 209.165.200.225
ip route static bfd Ethernet 0/0.1001 209.165.200.226 group testgroup passive
ip route 10.1.1.1 255.255.255.224 Ethernet 0/0.1001 209.165.200.226
ip route 10.2.2.2 255.255.255.224 Ethernet 0/0.1001 209.165.200.226
```

Example: Configuring BFD Slow Timer

```
Switch# configure terminal
Switch(config)# bfd slow-timer 12000
```

Additional References for Bidirectional Forwarding Detection

Technical Assistance

Description	Link
<p>The Cisco Support website provides extensive online resources, including documentation and tools for troubleshooting and resolving technical issues with Cisco products and technologies.</p> <p>To receive security and technical information about your products, you can subscribe to various services, such as the Product Alert Tool (accessed from Field Notices), the Cisco Technical Services Newsletter, and Really Simple Syndication (RSS) Feeds.</p> <p>Access to most tools on the Cisco Support website requires a Cisco.com user ID and password.</p>	http://www.cisco.com/support

Feature Information for Bidirectional Forwarding Detection

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

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

Table 2: Feature Information for Bidirectional Forwarding Detection

Feature Name	Releases	Feature Information
Bidirectional Forwarding Detection	Cisco IOS 15.2(4)E1	<p>The Bidirectional Forwarding Detection (BFD) protocol is a detection protocol that is designed to provide fast forwarding path failure detection for all media types, encapsulations, topologies, and routing protocols.</p> <p>BFD provides a consistent failure detection method for network administrators, in addition to fast forwarding path failure detection. Because the network administrator can use BFD to detect forwarding path failures at a uniform rate, rather than the variable rates for different routing protocol hello mechanisms, network profiling and planning will be easier, and reconvergence time will be consistent and predictable.</p> <p>This feature was implemented on Cisco Catalyst 2960-XR Series Switches.</p>