



CHAPTER 9

Configuring Bidirectional Forwarding Detection for BGP

This chapter describes how to configure Bidirectional Forwarding Detection (BFD) for BGP.

This chapter includes the following sections:

- [Information About BFD, page 9-1](#)
- [Licensing Requirements for BFD, page 9-3](#)
- [Prerequisites for BFD, page 9-3](#)
- [Guidelines and Limitations, page 9-3](#)
- [Default Settings, page 9-4](#)
- [Configuring BFD, page 9-5](#)
- [Verifying the BFD Configuration, page 9-14](#)
- [Monitoring BFD, page 9-14](#)
- [Feature History for BFD, page 9-15](#)

Information About BFD

BFD is a detection protocol designed to provide fast forwarding-path failure detection times for media types, encapsulations, topologies, and routing protocols. You can use BFD to detect forwarding path failures at a uniform rate, rather than the variable rates for different protocol hello mechanisms. BFD makes network profiling and planning easier and reconvergence time consistent and predictable.

BFD provides subsecond failure detection between two adjacent devices.

This section includes the following topics:

- [Asynchronous Mode, page 9-2](#)
- [BFD Detection of Failures, page 9-2](#)
- [BFD Echo Function, page 9-2](#)
- [Security, page 9-3](#)
- [Virtualization Support, page 9-3](#)
- [Virtualization Support, page 9-3](#)

Send document comments to nexus3k-docfeedback@cisco.com.

Asynchronous Mode

Cisco NX-OS supports the BFD asynchronous mode, which sends BFD control packets between two adjacent devices to activate and maintain BFD neighbor sessions between the devices. You configure BFD on both devices (or BFD neighbors). Once BFD has been enabled on the appropriate protocols, Cisco NX-OS creates a BFD session, negotiates BFD session parameters, and begins to send BFD control packets to each BFD neighbor at the negotiated interval. The BFD session parameters include the following:

- Desired minimum transmit interval—The interval at which this device wants to send BFD hello messages.
- Required minimum receive interval—The minimum interval at which this device can accept BFD hello messages from another BFD device.
- Detect multiplier—The number of missing BFD hello messages from another BFD device before this local device detects a fault in the forwarding path.

BFD Detection of Failures

Once a BFD session has been established and timer negotiations are complete, BFD neighbors send BFD control packets that act in the same manner as an IGP hello protocol to detect liveness, except at a more accelerated rate. BFD detects a failure, but the protocol must take action to bypass a failed peer.

BFD sends a failure detection notice to the BFD-enabled protocols when it detects a failure in the forwarding path. The local device can then initiate the protocol recalculation process and reduce the overall network convergence time.

When a failure occurs in the network:

1. The BFD neighbor session with the BFD neighbor router is torn down.
2. BFD notifies the local BFD process that the BFD neighbor is no longer reachable.
3. The local BFD process tears down the BFD neighbor relationship.
4. If an alternative path is available, the routers immediately start converging on it.

**Note**

The BFD failure detection occurs in less than a second.

BFD Echo Function

The BFD echo function sends echo packets from the forwarding engine to the remote BFD neighbor. The BFD neighbor forwards the echo packet back along the same path in order to perform detection; the BFD neighbor does not participate in the actual forwarding of the echo packets. The echo function and the forwarding engine are responsible for the detection process. BFD can use the slow timer to slow down the asynchronous session when the echo function is enabled and reduce the number of BFD control packets that are sent between two BFD neighbors. Also, the forwarding engine tests the forwarding path on the remote (neighbor) system without involving the remote system, so there is less interpacket delay variability and faster failure detection times.

The echo function is without asymmetry when both BFD neighbors are running echo function.

Send document comments to nexus3k-docfeedback@cisco.com.

Security

Cisco NX-OS uses the packet Time to Live (TTL) value to verify that the BFD packets came from an adjacent BFD peer. For all asynchronous and echo request packets, the BFD neighbor sets the TTL value to 255 and the local BFD process verifies the TTL value as 255 before processing the incoming packet. For the echo response packet, BFD sets the TTL value to 254.

Virtualization Support

BFD supports virtual routing and forwarding instances (VRFs). VRFs exist within virtual device contexts (VDCs). By default, Cisco NX-OS places you in the default VDC and default VRF unless you specifically configure another VDC and VRF.

Licensing Requirements for BFD

The following table shows the licensing requirements for this feature:

Product	License Requirement
Cisco NX-OS	BFD requires no license. Any feature not included in a license package is bundled with the Cisco NX-OS system images and is provided at no extra charge to you. For a complete explanation of the Cisco NX-OS licensing scheme, see the <i>Cisco NX-OS Licensing Guide</i> .

Prerequisites for BFD

BFD has the following prerequisites:

- You must enable the BFD feature (see the [“Enabling the BFD Feature”](#) section on page 9-6).
- For any client protocols that you want to enable BFD on, you enable BFD in that client protocol. See the [“Configuring BFD on BGP”](#) section on page 9-13.
- Disable Internet Control Message Protocol (ICMP) redirect messages on a BFD-enabled interfaces.
- See other detailed prerequisites that are listed with the configuration tasks.

Guidelines and Limitations

BFD has the following configuration guidelines and limitations:

- BFD supports BFD version 1.
- BFD supports IPv4.
- BFD supports single-hop BFD.
- BFD for BGP supports single-hop eBGP peers and single-hop iBGP peers with update-source.
- BFD supports the following Layer 3 interfaces—physical interfaces, port channels, subinterfaces, and VLAN interfaces (SVI).
- BFD for BGP does not support authentication for BFD or per-link BFD sessions on a port channel.

Send document comments to nexus3k-docfeedback@cisco.com.

- BFD depends on a Layer 3 adjacency information to discover topology changes, including Layer 2 topology changes. A BFD session on a VLAN interface (SVI) may not be up after the convergence of the Layer 2 topology if there is no Layer 3 adjacency information available.
- Port channel configuration limitations:
 - For Layer 3 port channels used by BFD, you must enable LACP on the port channel.
 - For Layer 2 port channels used by SVI sessions, you must enable LACP on the port channel.
- SVI limitations:
 - An ASIC reset will cause traffic disruption for other ports. This event could possibly cause SVI sessions on other ports to flap. Some triggers for an ASIC reset are port moves between VDCs or reloading a VDC. In these cases, BFD is not supported over the SVI interface.
 - When you change the topology (for example, add or delete a link into a VLAN, delete a member from a Layer 2 port channel, and so on), the SVI session could be affected. It may go down first and then come up after the topology discovery is finished.



Tip

If you do not want the SVI sessions to flap and you need to change the topology, you can disable the BFD feature before making the changes and reenable BFD after the changes have been made. You can also configure the BFD timer to be a large value (for example, 5 seconds), and change it back to a fast timer after the above events complete.

- Cisco NX-OS does not distribute the BFD operation to compatible modules to offload the CPU for BFD packet processing.
- BFD does not support stateless restarts and in-service software upgrades (ISSUs).
- If you want to enable BFD for a peer reachable through a port channel, you must configure LACP on the port channel.
- BFD echo mode and Unicast Reverse Path Forwarding (URPF) are mutually exclusive and cannot both be enabled on a BFD interface. If you want to configure an interface for BFD, you must disable either BFD echo mode or URPF.

Default Settings

Table 9-1 lists the default settings for BFD parameters.

Table 9-1 **Default BFD Parameters**

Parameters	Default
BFD feature	Disabled
Required minimum receive interval	250 milliseconds
Desired minimum transmit interval	250 milliseconds
Detect multiplier	3
Echo function	Enabled
Mode	Asynchronous
Port channel	Logical mode (one session per source-destination pair address).

Send document comments to nexus3k-docfeedback@cisco.com.

Table 9-1 **Default BFD Parameters (continued)**

Parameters	Default
Slow timer	2000 milliseconds
Subinterface optimization	Disabled

Configuring BFD

This section includes the following topics:

- [Configuration Hierarchy, page 9-5](#)
- [Task Flow for Configuring BFD, page 9-5](#)
- [Enabling the BFD Feature, page 9-6](#)
- [Configuring Global BFD Parameters, page 9-6](#)
- [Configuring BFD on an Interface, page 9-8](#)
- [Configuring BFD on a Port Channel, page 9-9](#)
- [Configuring BFD Echo Function, page 9-10](#)
- [Optimizing BFD on Subinterfaces, page 9-11](#)
- [Configuring BFD on BGP, page 9-13](#)

Configuration Hierarchy

You can configure BFD at the global level, VRF level, at the interface or port channel level, or at the subinterface level (for physical interfaces and port channels). The VRF configuration overrides global configuration. The interface or port channel configuration overrides VRF or global configuration. On supported interfaces, the subinterface-level configuration overrides the interface or port channel configuration unless subinterface optimization is enabled. See the [“Optimizing BFD on Subinterfaces” section on page 9-11](#) for more information.

For physical ports that are members of a port channel, the member port inherits the master port channel BFD configuration. The member port subinterfaces can override the master port channel BFD configuration, unless subinterface optimization is enabled.

Task Flow for Configuring BFD

Follow these steps to configure BFD:

-
- Step 1** [Enabling the BFD Feature.](#)
- Step 2** [Configuring Global BFD Parameters](#) or [Configuring BFD on an Interface.](#)
- Step 3** [Configuring BFD on BGP.](#)
-

Send document comments to nexus3k-docfeedback@cisco.com.

Enabling the BFD Feature

You must enable the BFD feature before you can configure BFD on an interface and protocol within a device (VDC).

BEFORE YOU BEGIN

Ensure that you are in the correct VDC (or use the **switchto vdc** command).

SUMMARY STEPS

1. **configure terminal**
2. **feature bfd**
3. (Optional) **copy running-config startup-config**

DETAILED STEPS

	Command	Purpose
Step 1	configure terminal Example: switch# configure terminal switch(config)#	Enters configuration mode.
Step 2	feature bfd Example: switch(config)# feature bfd	Enables the BFD feature.
Step 3	copy running-config startup-config Example: switch(config)# copy running-config startup-config	(Optional) Saves this configuration change.

Use the **no feature bfd** command to disable the BFD feature and remove all associated configuration.

	Command	Purpose
	no feature bfd Example: switch(config)# no bfd	Disables the BFD feature and removes all associated configuration.

Configuring Global BFD Parameters

You can configure the BFD session parameters for all BFD sessions on the device. The BFD session parameters are negotiated between the BFD peers in a three-way handshake.

See the “[Configuring BFD on an Interface](#)” section on page 9-8 to override these global session parameters on an interface.

Send document comments to nexus3k-docfeedback@cisco.com.

BEFORE YOU BEGIN

Ensure that you are in the correct VDC (or use the **switchto vdc** command).

Enable the BFD feature. See the “Enabling the BFD Feature” section on page 9-6.

SUMMARY STEPS

1. **configure terminal**
2. **bfd interval *mintx* *min_rx* *msec* *multiplier* *value***
3. **bfd slow-timer [*interval*]**
4. (Optional) **show running-config bfd**
5. (Optional) **copy running-config startup-config**

DETAILED STEPS

	Command	Purpose
Step 1	configure terminal Example: switch# configure terminal switch(config)#	Enters configuration mode.
Step 2	bfd interval <i>mintx</i> <i>min_rx</i> <i>msec</i> <i>multiplier</i> <i>value</i> Example: switch(config)# bfd interval 250 min_rx 250 multiplier 3	Configures the BFD session parameters for all BFD sessions on the device. You can override these values by configuring the BFD session parameters on an interface. The <i>mintx</i> and <i>msec</i> range is from 250 to 999 milliseconds and the default is 250. The multiplier range is from 3 to 50. The multiplier default is 3. To return to the default settings, use the no bfd interval command.
Step 3	bfd slow-timer [<i>interval</i>] Example: switch(config)# bfd slow-timer 2000.	Configures the slow timer. This value determines how fast BFD starts up a new session and is used to slow down the asynchronous sessions when the BFD echo function is enabled. The range is from 1000 to 30000 milliseconds. The default is 2000. To return to the default settings, use the no bfd slow-timer command.
Step 4	show running-config bfd Example: switch(config)# show running-config bfd	(Optional) Displays the BFD running configuration.
Step 5	copy running-config startup-config Example: switch(config)# copy running-config startup-config	(Optional) Saves this configuration change.

Send document comments to nexus3k-docfeedback@cisco.com.

Configuring BFD on an Interface

You can configure the BFD session parameters for all BFD sessions on an interface. The BFD session parameters are negotiated between the BFD peers in a three-way handshake.

This configuration overrides the global session parameters for the configured interface.

BEFORE YOU BEGIN

Ensure that you are in the correct VDC (or use the **switchto vdc** command).

Enable the BFD feature. See the “Enabling the BFD Feature” section on page 9-6.

SUMMARY STEPS

1. **configure terminal**
2. **interface *int-if***
3. **bfd interval *mintx min_rx msec multiplier value***
4. (Optional) **show running-config bfd**
5. (Optional) **copy running-config startup-config**

DETAILED STEPS

	Command	Purpose
Step 1	configure terminal Example: switch# configure terminal switch(config)#	Enters configuration mode.
Step 2	interface <i>int-if</i> Example: switch(config)# interface ethernet 2/1 switch(config-if)#	Enters interface configuration mode. Use the ? keyword to display the supported interfaces.
Step 3	bfd interval <i>mintx min_rx msec multiplier value</i> Example: switch(config-if)# bfd interval 250 min_rx 250 multiplier 3	Configures the BFD session parameters for all BFD sessions on the interface. This overrides the global BFD session parameters. The <i>mintx</i> and <i>msec</i> range is from 250 to 999 milliseconds and the default is 250. The multiplier range is from 3 to 50. The multiplier default is 3. To return to the default settings, use the no bfd interval command.
Step 4	show running-config bfd Example: switch(config-if)# show running-config bfd	(Optional) Displays the BFD running configuration.
Step 5	copy running-config startup-config Example: switch(config-if)# copy running-config startup-config	(Optional) Saves this configuration change.

Send document comments to nexus3k-docfeedback@cisco.com.

Configuring BFD on a Port Channel

You can configure the BFD session parameters for all BFD sessions on a port channel. For example, if the BFD session for one link on a port channel is up, BFD informs client protocols, such as BGP, that the port channel is up. The BFD session parameters are negotiated between the BFD peers in a three-way handshake.

This configuration overrides the global session parameters for the configured port channel. The member ports of the port channel inherit the port channel BFD session parameters, unless you configure subinterface-level BFD parameters on a member port. In that case, the member port subinterface uses the subinterface BFD configuration if subinterface optimization is not enabled. See the “[Optimizing BFD on Subinterfaces](#)” section on page 9-11 for more information.

BEFORE YOU BEGIN

Ensure that you are in the correct VDC (or use the **switchto vdc** command).

Ensure that you enable LACP on the port channel before you enable BFD.

Enable the BFD feature. See the “[Enabling the BFD Feature](#)” section on page 9-6.

SUMMARY STEPS

1. **configure terminal**
2. **interface port-channel** *number*
3. (Optional) **bfd interval** *mintx min_rx msec multiplier value*
4. (Optional) **show running-config bfd**
5. (Optional) **copy running-config startup-config**

DETAILED STEPS

	Command	Purpose
Step 1	configure terminal Example: switch# configure terminal switch(config)#	Enters configuration mode.
Step 2	interface port-channel <i>number</i> Example: switch(config)# interface port-channel 2 switch(config-if)#	Enters port channel configuration mode. Use the ? keyword to display the supported number range.
Step 3	bfd interval <i>mintx min_rx msec multiplier value</i> Example: switch(config-if)# bfd interval 250 min_rx 250 multiplier 3	(Optional) Configures the BFD session parameters for all BFD sessions on the port channel. This overrides the global BFD session parameters. The <i>mintx</i> and <i>msec</i> range is from 250 to 999 milliseconds and the default is 250. The multiplier range is from 3 to 50. The multiplier default is 3. To return to the default settings, use the no bfd interval command.

Send document comments to nexus3k-docfeedback@cisco.com.

	Command	Purpose
Step 4	show running-config bfd Example: switch(config-if)# show running-config bfd	(Optional) Displays the BFD running configuration.
Step 5	copy running-config startup-config Example: switch(config-if)# copy running-config startup-config	(Optional) Saves this configuration change.

Configuring BFD Echo Function

You can configure the BFD echo function on one or both ends of a BFD-monitored link. The echo function slows down the required minimum receive interval, based on the configured slow timer. The RequiredMinEchoRx BFD session parameter is set to zero if the echo function is disabled. The slow timer becomes the required minimum receive interval if the echo function is enabled.

BEFORE YOU BEGIN

Ensure that you are in the correct VDC (or use the **switchto vdc** command).

Enable the BFD feature. See the “[Enabling the BFD Feature](#)” section on page 9-6.

Configure the BFD session parameters. See the “[Configuring Global BFD Parameters](#)” section on page 9-6 or the “[Configuring BFD on an Interface](#)” section on page 9-8.

Ensure that Internet Control Message Protocol (ICMP) redirect messages are disabled on BFD-enabled interfaces. Use the **no ip redirects** command on the interface.

SUMMARY STEPS

1. **configure terminal**
2. **bfd slow-timer** *echo-interval*
3. **interface** *int-if*
4. **bfd echo**
5. (Optional) **show running-config bfd**
6. (Optional) **copy running-config startup-config**

Send document comments to nexus3k-docfeedback@cisco.com.

DETAILED STEPS

	Command	Purpose
Step 1	<code>configure terminal</code> Example: switch# configure terminal switch(config)#	Enters configuration mode.
Step 2	<code>bfd slow-timer echo-interval</code> Example: switch(config)# bfd slow-timer 2000	Configures the slow timer used in the echo function. This value determines how fast BFD starts up a new session and is used to slow down the asynchronous sessions when the BFD echo function is enabled. This value overwrites the required minimum receive interval when the echo function is enabled. The range is from 1000 to 30000 milliseconds. The default is 2000. To return to the default settings, use the no bfd slow-timer command.
Step 3	<code>interface int-if</code> Example: switch(config)# interface ethernet 2/1 switch(config-if)#	Enters interface configuration mode. Use the ? keyword to display the supported interfaces.
Step 4	<code>bfd echo</code> Example: switch(config-if)# bfd echo	Enables the echo function. The default is enabled. To disable the echo function, use the no bfd echo command.
Step 5	<code>show running-config bfd</code> Example: switch(config-if)# show running-config bfd	(Optional) Displays the BFD running configuration.
Step 6	<code>copy running-config startup-config</code> Example: switch(config-if)# copy running-config startup-config	(Optional) Saves this configuration change.

Optimizing BFD on Subinterfaces

You can optimize BFD on subinterfaces. BFD creates sessions for all configured subinterfaces. BFD sets the subinterface with the lowest configured VLAN ID as the master subinterface and that subinterface uses the BFD session parameters of the parent interface. The remaining subinterfaces use the slow timer. If the optimized subinterface session detects an error, BFD marks all subinterfaces on that physical interface as down.

BEFORE YOU BEGIN

Ensure that you are in the correct VDC (or use the **switchto vdc** command).

Enable the BFD feature. See the “[Enabling the BFD Feature](#)” section on page 9-6.

Send document comments to nexus3k-docfeedback@cisco.com.

Configure the BFD session parameters. See the “Configuring Global BFD Parameters” section on page 9-6 or the “Configuring BFD on an Interface” section on page 9-8.

Ensure that these subinterfaces connect to another Cisco NX-OS device. This feature is supported on Cisco NX-OS only.

SUMMARY STEPS

1. **configure terminal**
2. **interface *int-if***
3. **bfd optimize subinterface**
4. **(Optional) show running-config bfd**
5. **(Optional) copy running-config startup-config**

DETAILED STEPS

	Command	Purpose
Step 1	configure terminal Example: switch# configure terminal switch(config)#	Enters configuration mode.
Step 2	interface <i>int-if</i> Example: switch(config)# interface ethernet 2/1 switch(config-if)#	Enters interface configuration mode. Use the ? keyword to display the supported interfaces.
Step 3	bfd optimize subinterface Example: switch(config-if)# bfd optimize subinterface	Optimizes subinterfaces on a BFD-enabled interface. The default is disabled. To disable optimized subinterfaces, use the no bfd optimize subinterfaces command.
Step 4	show running-config bfd Example: switch(config-if)# show running-config bfd	(Optional) Displays the BFD running configuration.
Step 5	copy running-config startup-config Example: switch(config-if)# copy running-config startup-config	(Optional) Saves this configuration change.

Send document comments to nexus3k-docfeedback@cisco.com.

Configuring BFD on BGP

You can configure BFD for the Border Gateway Protocol (BGP).

BEFORE YOU BEGIN

Ensure that you are in the correct VDC (or use the **switchto vdc** command).

Enable the BFD feature. See the “Enabling the BFD Feature” section on page 9-6.

Configure the BFD session parameters. See the “Configuring Global BFD Parameters” section on page 9-6 or the “Configuring BFD on an Interface” section on page 9-8.

Enable the BGP feature. See the *Cisco Nexus 7000 Series NX-OS Unicast Routing Configuration Guide, Release 5.x*, for more information.

SUMMARY STEPS

1. **configure terminal**
2. **router bgp** *as-number*
3. **neighbor** {*ip-address* | *ipv6-address*} **remote-as** *as-number*
4. **bfd**
5. (Optional) **show running-config bgp**
6. (Optional) **copy running-config startup-config**

DETAILED STEPS

	Command	Purpose
Step 1	configure terminal Example: switch# configure terminal switch(config)#	Enters configuration mode.
Step 2	router bgp <i>as-number</i> Example: switch(config)# router bgp 64496 switch(config-router)#	Enables BGP and assigns the AS number to the local BGP speaker. The AS number can be a 16-bit integer or a 32-bit integer in the form of a higher 16-bit decimal number and a lower 16-bit decimal number in xx.xx format.
Step 3	neighbor { <i>ip-address</i> <i>ipv6-address</i> } remote-as <i>as-number</i> Example: switch(config-router)# neighbor 209.165.201.1 remote-as 64497 switch(config-router-neighbor)#	Configures the IPv4 or IPv6 address and AS number for a remote BGP peer. The <i>ip-address</i> format is x.x.x.x. The <i>ipv6-address</i> format is A:B::C:D.
Step 4	bfd Example: switch(config-router-neighbor)# bfd	Enables BFD for this BGP peer.

Send document comments to nexus3k-docfeedback@cisco.com.

	Command	Purpose
Step 5	show running-config bgp Example: switch(config-router-neighbor)# show running-config bgp	(Optional) Displays the BGP running configuration.
Step 6	copy running-config startup-config Example: switch(config-router-neighbor)# copy running-config startup-config	(Optional) Saves this configuration change.

Verifying the BFD Configuration

To display BFD configuration information, perform one of the following tasks:

Command	Purpose
show running-config bfd	Displays the running BFD configuration.
show startup-config bfd	Displays the BFD configuration that will be applied on the next system startup.

For detailed information about the fields in the output from these commands, see the *Cisco Nexus 7000 Series NX-OS Interfaces Command Reference, Release 5.x*.

Monitoring BFD

Use the following commands to display BFD:

Command	Purpose
show bfd neighbors [application name] [details]	Displays information about BFD for a supported application, such as BGP.
show bfd neighbors [interface int-if] [details]	Displays information about BGP sessions on an interface.
show bfd neighbors [dest-ip ip-address] [src-ip ip-address][details]	Displays information about the specified BGP session on an interface.
show bfd neighbors [vrf vrf-name] [details]	Displays information about BFD for a VRF.

For detailed information about the fields in the output from these commands, see the *Cisco Nexus 3000 Series Command Reference*.

Send document comments to nexus3k-docfeedback@cisco.com.

Feature History for BFD

Table 9-2 lists the release history for this feature.

Table 9-2 Feature History for BGP

Feature Name	Releases	Feature Information
BFD for BGP	5.0(3)U2(2)	Added support for BFD for BGP. See Chapter 9, “Configuring Bidirectional Forwarding Detection for BGP” for more information.
BFD for BGP	5.0(3)U2(2a)	Added guidelines and limitations for LACP configuration, and for BFD echo and URFP. See the “Guidelines and Limitations” section on page 9-3 for more information.

Send document comments to nexus3k-docfeedback@cisco.com.