

# Configure and Verify BFD on Nexus 9000 Switches

## Contents

---

### [Introduction](#)

### [Prerequisites](#)

[Requirements](#)

[Components Used](#)

### [Configure](#)

[Network Diagram](#)

[Configure](#)

[Syslog BFD Down Reasons](#)

### [Configuring BFD on Routing Protocols](#)

[Configuring BFD on OSPF](#)

[Example Configurations for BFD on OSPF](#)

[Configuring BFD on EIGRP](#)

[Example Configurations for BFD on EIGRP](#)

[Configuring BFD on BGP](#)

[Example Configurations for BFD on BGP](#)

### [Verify](#)

[Verify Using Session Details](#)

[Verify Using Access-list](#)

[Verify using Ethalyzer](#)

---

## Introduction

This document describes how to configure and verify Bidirectional Forwarding Detection (BFD) sessions across Cisco Nexus NXOS® based switches.

## Prerequisites

### Requirements

Cisco recommends that you have knowledge of these topics:

- Bidirectional Forwarding Detection (BFD)
- Nexus NX-OS Software.
- Routing Protocols: Open Shortest Path First (OSPF), Border Gateway Protocol (BGP), Enhanced Interior Gateway Routing Protocol (EIGRP).


## Components Used

The information in this document is based on Cisco Nexus 9000 with NXOS version 10.3(4a).M.

The information in this document was created from the devices in a specific lab environment. All of the devices used in this document started with a cleared (default) configuration. If your network is live, ensure that you understand the potential impact of any command.

## Configure

### Network Diagram



## Configure


The purpose of configuring BFD is to detect and understand the differences among the configurations of various routing protocols.

STEP 1: You must enable the BFD feature before you can configure BFD on an interface and protocol.

SWITCH 1	SWITCH 2
SW1(config)# feature bfd	SW2(config)# feature bfd

STEP 2: Configuring Global BFD


SWITCH 1	SWITCH 2
SW1(config)# bfd interval 500 min_rx 500 multiplier 3	SW2(config)# bfd interval 500 min_rx 500 multiplier 3



**Note:** The min\_tx and msec range is from 50 to 999 milliseconds and the default is 50. The multiplier range is from 1 to 50. The multiplier default is 3.

### STEP 3: Configuring BFD on an interface

---



**Note:** You can configure the BFD session parameters for all BFD sessions on an interface.

---



**Warning:** Ensure that Internet Control Message Protocol (ICMP) redirect messages are disabled on BFD-enabled interfaces. Use the `no ip redirects` command or the `no ipv6 redirects` command on the interface.

SWITCH 1	SWITCH 2
<pre>SW1(config)# interface vlan 20 SW1(config-if)# bfd interval 500 min_rx 500 multiplier 3 SW1(config-if)# no ip redirects SW1(config-if)# no ipv6 redirects</pre>	<pre>SW2(config)# interface vlan 20 SW2(config-if)# bfd interval 500 min_rx 500 multiplier 3 SW2(config-if)# no ip redirects SW2(config-if)# no ipv6 redirects</pre>

BFD asynchronous mode is like a handshake between two devices to keep their connection strong. You set it up on both devices, and once it is on, they start sending special messages to each other at a set time. These messages have some important settings, like how often they are sent and how fast one device can respond to the other. There is also a setting that decides how many missed messages it takes for one device to realize there can be a problem with the connection.

The BFD echo function sends test packets to a neighbor and gets them sent back to check for issues without

involving the neighbor in packet forwarding. It can use a slower timer to reduce control packet traffic and tests the forwarding path on the neighbors system without bothering the neighbor, making detection faster. If both neighbors use the echo function, there is no asymmetry.

## **Syslog BFD Down Reasons**

- **Path Down:** This indicates that the forwarding path between the two BFD neighbors is no longer operational, possibly due to network congestion, hardware failure, or other issues.

```
2024 Apr 11 22:07:07 SW2 %BFD-5-SESSION_STATE_DOWN: BFD session 1090519062 to neighbor 172.16.1.1
```

- **Echo Function Failed:** Failure of the echo function, which is a feature of BFD where echo packets are sent and received to verify connectivity. If these packets fail to be transmitted or received successfully, it indicates a problem.

```
2024 Apr 11 22:17:45 SW2 %BFD-5-SESSION_STATE_DOWN: BFD session 1090519174 to neighbor 10.10.10.1
```

- **Neighbor Signaled Session Down:** The neighbor device signals that the BFD session is down, typically due to detecting a problem on it is end of the connection.

```
2024 Apr 11 22:03:48 SW2 %BFD-5-SESSION_STATE_DOWN: BFD session 1090519058 to neighbor 172.16.1.1
```

- **Control Detection Time Expired:** This occurs when the control detection timer runs out before receiving an expected response from the neighbor, indicating a potential issue with the connection.


```
2024 Apr 11 22:19:31 SW2 %BFD-5-SESSION_STATE_DOWN: BFD session 1090519061 to neighbor 192.168.2.1
```

- **Administratively Down:** The BFD session is intentionally brought down by an administrator, either for maintenance purposes or due to configuration changes.

```
2024 Apr 11 22:13:15 SW2 %BFD-5-SESSION_STATE_DOWN: BFD session 1090519064 to neighbor 10.10.10.1
```

## **Configuring BFD on Routing Protocols**

### **Configuring BFD on OSPF**



**Tip:** When BFD is enabled under OSPF, it becomes active for all interfaces utilizing OSPF. The interfaces adopt the globally configured values. If adjustments to these values are necessary, refer to step 3, 'BFD Configuration on an Interface'.

SWITCH 1	SWITCH 2
SW1(config)# router ospf 1 SW1(config-router)# bfd	SW2(config)# router ospf 1 SW2(config-router)# bfd

It can also enable BFD under the OSPF interface with the command `ip ospf bfd`

SWITCH 1	SWITCH 2

SW1(config)# interface vlan 10 SW1(config-if)# ip ospf bfd	SW2(config)# interface vlan 10 SW2(config-if)# ip ospf bfd
---	---

## Example Configurations for BFD on OSPF

```
SW1# show running-config ospf
```

```
!Command: show running-config ospf  
!Running configuration last done at: Wed Apr 10 21:06:31 2024  
!Time: Wed Apr 10 21:09:45 2024
```

```
version 10.3(4a) Bios:version 01.07  
feature ospf
```


```
router ospf 1  
  bfd
```

```
interface Vlan10  
  ip router ospf 1 area 0.0.0.0  
  ip ospf bfd
```

## Configuring BFD on EIGRP

```
SW1(config)# interface vlan 20  
SW1(config-if)# ip eigrp 2 bfd
```





**Tip:** When BFD is enabled under EIGRP, it becomes active for all interfaces utilizing EIGRP. The interfaces adopt the globally configured values. If adjustments to these values are necessary, refer to step 3, 'BFD Configuration on an Interface'.

SWITCH 1	SWITCH 2
<pre>SW1(config)# router eigrp 2 SW1(config-router)# bfd</pre>	<pre>SW2(config)# router eigrp 2 SW2(config-router)# bfd</pre>

It can also enable BFD under an EIGRP interface with the command `ip eigrp instance-tag bfd`

SWITCH 1	SWITCH 2
<pre>SW1(config)# interface vlan 20</pre>	<pre>SW2(config)# interface vlan 20</pre>

SW1(config-if)# ip eigrp 2 bfd	SW2(config-if)# ip eigrp 2 bfd
--------------------------------	--------------------------------

## Example Configurations for BFD on EIGRP

```
SW1# show running-config eigrp
```

```
!Command: show running-config eigrp  
!Running configuration last done at: Wed Apr 10 21:24:29 2024  
!Time: Wed Apr 10 21:53:08 2024
```

```
version 10.3(4a) Bios:version 01.07  
feature eigrp
```

```
router eigrp 2  
  address-family ipv4 unicast  
    bfd
```


```
interface Vlan20  
  ip router eigrp 2  
  ip eigrp 2 bfd
```

## Configuring BFD on BGP



**Note:** The update source feature facilitates BGP sessions to utilize the primary IP address of a designated interface as the local address during the establishment of a BGP session with a neighbor. Additionally, it enables BGP to register as a client with BFD.

---



**Note:** When configuring BFD sessions on the device, specifying either 'multihop' or 'singlehop' determines the session type. If no keyword is provided, the session type defaults to 'singlehop' when the peer is directly connected. If the peer is not connected, the session type defaults to 'multihop'.

SWITCH 1	SWITCH 2
<pre>SW1(config)# router bgp 65001 SW1(config-router)# address-family ipv4 unicast SW1(config-router)# neighbor 192.168.3.1 SW1(config-router-neighbor)# bfd multihop SW1(config-router-neighbor)# update-source loopback30</pre>	<pre>SW2(config)# router bgp 65002 SW2(config-router)# address-family ipv4 unicast SW2(config-router)# neighbor 192.168.2.1 SW2(config-router-neighbor)# bfd multihop SW2(config-router-neighbor)# update-source loopback30</pre>

## Example Configurations for BFD on BGP

```
SW1# show running-config bgp
```

```
!Command: show running-config bgp
!Running configuration last done at: Thu Apr 11 04:47:53 2024
!Time: Thu Apr 11 06:48:43 2024
```

```
version 10.3(4a) Bios:version 01.07
feature bgp
```

```
router bgp 65001
  address-family ipv4 unicast
  neighbor 192.168.3.1
    bfd multihop
    remote-as 65002
    update-source loopback30
  ebgp-multihop 5
  address-family ipv4 unicast
```

## Verify

After configuring BFD and associating it with a protocol such as OSPF, EIGRP, or BGP, the BFD neighbors must be automatically identified. To confirm this, use the command:

```
show bfd neighbors
```

On Switch 1

```
SW1# show bfd neighbors
```

OurAddr	NeighAddr	LD/RD	RH/RS	Hold down(mult)	State	Int
172.16.1.1	172.16.1.2	1090519059/1090519058	Up	5219(3)	Up	V1a
10.10.10.1	10.10.10.2	1090519060/1090519059	Up	5219(3)	Up	V1a
192.168.2.1	192.168.3.1	1090519062/1090519061	Up	728(3)	Up	Lo3

```
SW1#
```

On Switch 2

```
SW2# show bfd neighbors
```

OurAddr	NeighAddr	LD/RD	RH/RS	Hold down(mult)	State	Int
172.16.1.2	172.16.1.1	1090519058/1090519059	Up	5053(3)	Up	V1a
10.10.10.2	10.10.10.1	1090519059/1090519060	Up	5053(3)	Up	V1a
192.168.3.1	192.168.2.1	1090519061/1090519062	Up	745(3)	Up	Lo3

To confirm this and obtain detailed output, use the command:

SW1# show bfd neighbors interface lo30 details

OurAddr	NeighAddr	LD/RD	RH/RS	Holdown(mult)	State	Int
192.168.2.1	192.168.3.1	1090519062/1090519061	Up	676(3)	Up	Lo3

Session state is Up and not using echo function

Session type: Multihop

Local Diag: 0, Demand mode: 0, Poll bit: 0, Authentication: None

MinTxInt: 250000 us, MinRxInt: 250000 us, Multiplier: 3

Received MinRxInt: 250000 us, Received Multiplier: 3

Holdown (hits): 750 ms (0), Hello (hits): 250 ms (82590)

Rx Count: 82588, Rx Interval (ms) min/max/avg: 207/1764/219 last: 73 ms ago

Tx Count: 82590, Tx Interval (ms) min/max/avg: 219/219/219 last: 67 ms ago

Registered protocols: bgp

Uptime: 0 days 5 hrs 2 mins 2 secs, Upcount: 1

Last packet: Version: 1 - Diagnostic: 0  
State bit: Up - Demand bit: 0  
Poll bit: 0 - Final bit: 0  
Multiplier: 3 - Length: 24  
My Discr.: 1090519061 - Your Discr.: 1090519062  
Min tx interval: 250000 - Min rx interval: 250000  
Min Echo interval: 0 - Authentication bit: 0

Hosting LC: 1, Down reason: None, Reason not-hosted: None

SW2# show bfd neighbors interface vlan 20 details

OurAddr	NeighAddr	LD/RD	RH/RS	Holdown(mult)	State	Int
172.16.1.2	172.16.1.1	1090519058/1090519059	Up	4979(3)	Up	Vlan

Session state is Up and using echo function with 500 ms interval

Session type: Singlehop

Local Diag: 0, Demand mode: 0, Poll bit: 0, Authentication: None

MinTxInt: 500000 us, MinRxInt: 2000000 us, Multiplier: 3

Received MinRxInt: 2000000 us, Received Multiplier: 3

Holdown (hits): 6000 ms (0), Hello (hits): 2000 ms (12454)

Rx Count: 12444, Rx Interval (ms) min/max/avg: 109/1919/1908 last: 1020 ms ago

Tx Count: 12454, Tx Interval (ms) min/max/avg: 1906/1906/1906 last: 856 ms ago

Registered protocols: eigrp

Uptime: 0 days 6 hrs 35 mins 10 secs, Upcount: 2

Last packet: Version: 1 - Diagnostic: 0  
State bit: Up - Demand bit: 0  
Poll bit: 0 - Final bit: 0  
Multiplier: 3 - Length: 24  
My Discr.: 1090519059 - Your Discr.: 1090519058  
Min tx interval: 500000 - Min rx interval: 2000000  
Min Echo interval: 500000 - Authentication bit: 0

Hosting LC: 1, Down reason: None, Reason not-hosted: None

## Verify Using Session Details

SW1# sh bfd clients

```
Client : Number of sessions
      bgp : 1
      ospf : 1
      eigrp : 1
```

```
SW1# show system internal bfd sess-store interface vlan 10
No of sessions: 1 >>>> Check Number of sessions
0x4100001b - 0x41000019 Hosting LC: 1 10.10.10.1 -> 10.10.10.2 bsid 0 >>>> Check IP addresses
BFD_SESS_ST_SESSION_UP for 0 days 0:6:22 594 >>>> Check uptime
Vlan10 [2 0x901000a:0x43 0:10 0x1000c 0x1000c]
SMAC: 4464.3c81.ddd7 DMAC: 4464.3c81.cdb7
Local TX/RX/ST(ms): 500/500/2000000 Mult: 3 EchoRx: 50000 Ver:1 Flags(0x0):
Echo Auth Interop (Disabled) Auth(0 0 0) Group id 0x0
Active TX/RX/ST(ms): 2000/2000/2000000 Mult: 3 State: 3 Diag: 0x0 Flags: 0xc
0000 Auth(0 0 0)
Spray DIs: 0
Port-Lcs:
Module 1 has 2 links - ACL not programmed
Applications:
(1) ospf [0x1:0x41000119 flags 0x0]
App Data : [c0a80301000]
Description: 10.10.10.2
MTS Dest 0xfe000001:0x140
No-Host reason: 0x0(SUCCESS)
Counters: >>>>>>> Check the counters below
Retries: LC remove(0), ACL Inst/Rem(0/0), Redisc(0)
No. of ACLs Added(1)
Flap counters:
Ctrl Detect timer expired(0), Echo failed(0), Neigh Signaled down(0), Admin d
own(0),
Forwarding plane reset(0), Path down(0), Concat path down(0), Rev Concat path
down(0)
Last ACL added time: 58781 usecs after Thu Apr 11 22:18:19 2024
Last ACL response time: 81830 usecs after Thu Apr 11 22:18:19 2024
Last session install time: 58285 usecs after Thu Apr 11 22:18:19 2024
Last session install response time: 58717 usecs after Thu Apr 11 22:18:19 2024
```

## Verify Using Access-list

```
SW2# show system internal access-list vlan 10 input statistics
```

```
slot 1
=====
```

```
INSTANCE 0x0
```

```
-----
Tcam 12 resource usage:
```

```
-----
LBL M = 0x1
Bank 0
```

```
-----
IPv4 Class
Policies: BFD() [Merged]
Netflow profile: 0
```

Netflow deny profile: 0

Entries:

[Index] Entry [Stats]

-----

[0x0000:0x0003:0x0c0d] permit udp 0.0.0.0/0 0.0.0.0/0 range 3784 3785 encap 0x0 [205340] >>>> This

```
SW2# show system internal access-list sup-redirect-stats | i i bfd
3085                                     BFD      12217 >>>> This counter need to increase
3121                                     BFDC TCLASS PACKET 0
3146                                     BFDC v4 PACKET    15
3148                                     BFDC TCLASS V6 PACKET 0
3156                                     BFDC v6 PACKET    0
3158                                     BFDC v6 LL Ucast PACKET 0
3209                                     BFDC v4 multihop PACKET 383697 >>>> This counter need to increase
3121                                     BFDC TCLASS PACKET 0
3146                                     BFDC v4 PACKET    0
3148                                     BFDC TCLASS V6 PACKET 0
3156                                     BFDC v6 PACKET    0
3158                                     BFDC v6 LL Ucast PACKET 0
3209                                     BFDC v4 multihop PACKET 0
```

## Verify using Ethalyzer

An alternative approach is to execute a packet capture, filtering specifically for UDP port 3785.

```
SW1# ethalyzer local interface inband display-filter "udp.port==3785" limit-captured-frames 0
Capturing on 'ps-inb'
 5 2024-04-11 05:26:37.294596186 10.10.10.1 → 10.10.10.1 BFD Echo 70 Originator specific content
 6 2024-04-11 05:26:37.294623155 10.10.10.1 → 10.10.10.1 BFD Echo 70 Originator specific content
 7 2024-04-11 05:26:37.294635837 172.16.1.1 → 172.16.1.1 BFD Echo 70 Originator specific content
 8 2024-04-11 05:26:37.294653159 172.16.1.1 → 172.16.1.1 BFD Echo 70 Originator specific content
4 18 2024-04-11 05:26:37.780441619 10.10.10.1 → 10.10.10.1 BFD Echo 70 Originator specific content
19 2024-04-11 05:26:37.780461351 172.16.1.1 → 172.16.1.1 BFD Echo 70 Originator specific content
20 2024-04-11 05:26:37.780467533 10.10.10.1 → 10.10.10.1 BFD Echo 70 Originator specific content
21 2024-04-11 05:26:37.780481315 172.16.1.1 → 172.16.1.1 BFD Echo 70 Originator specific content
30 2024-04-11 05:26:38.266365493 10.10.10.1 → 10.10.10.1 BFD Echo 70 Originator specific content
31 2024-04-11 05:26:38.266392766 10.10.10.1 → 10.10.10.1 BFD Echo 70 Originator specific content
```

The presence of identical source and destination IP addresses in the captured packets from the BFD Echo protocol is expected, as these Echo packets originate from the local switch itself.



---

**Note:** In the absence of the 'no bfd echo' statement under the interface, the capture reveals packets with both the locally sourced IP address and the neighboring destination IP address, along with the observation of BFD Control

---

```
SW2# ethanalyzer local interface inband display-filter "ip.addr==192.168.2.1" limit-captured-frames 0
Capturing on 'ps-inb'
 1 2024-04-11 05:38:21.155483591 192.168.2.1 → 192.168.3.1 BFD Control 66 Diag: No Diagnostic, State: Up, Fl
 2 2024-04-11 05:38:21.182606602 192.168.3.1 → 192.168.2.1 BFD Control 66 Diag: No Diagnostic, State: Up, Fl
11 2024-04-11 05:38:21.375187034 192.168.2.1 → 192.168.3.1 BFD Control 66 Diag: No Diagnostic, State: Up, Fl
12 2024-04-11 05:38:21.402295701 192.168.3.1 → 192.168.2.1 BFD Control 66 Diag: No Diagnostic, State: Up, Fl
 4 14 2024-04-11 05:38:21.594612491 192.168.2.1 → 192.168.3.1 BFD Control 66 Diag: No Diagnostic, State: Up,
15 2024-04-11 05:38:21.621742431 192.168.3.1 → 192.168.2.1 BFD Control 66 Diag: No Diagnostic, State: Up, Fl
20 2024-04-11 05:38:21.813971234 192.168.2.1 → 192.168.3.1 BFD Control 66 Diag: No Diagnostic, State: Up, Fl
 8 21 2024-04-11 05:38:21.841086554 192.168.3.1 → 192.168.2.1 BFD Control 66 Diag: No Diagnostic, State: Up,
22 2024-04-11 05:38:22.033395440 192.168.2.1 → 192.168.3.1 BFD Control 66 Diag: No Diagnostic, State: Up, Fl
23 2024-04-11 05:38:22.060526055 192.168.3.1 → 192.168.2.1 BFD Control 66 Diag: No Diagnostic, State: Up, Fl
28 2024-04-11 05:38:22.252709432 192.168.2.1 → 192.168.3.1 BFD Control 66 Diag: No Diagnostic, State: Up, Fl
29 2024-04-11 05:38:22.279810427 192.168.3.1 → 192.168.2.1 BFD Control 66 Diag: No Diagnostic, State: Up, Fl
```