BFD for Routing Protocols in Cisco Catalyst SD-WAN

Table 1: Feature History

Feature Name	Release Information	Description
BFD for Routing Protocols in Cisco Catalyst SD-WAN	Cisco IOS XE Catalyst SD-WAN Release 17.3.1a Cisco vManage Release 20.3.1	This feature extends BFD support to BGP, OSPF, and EIGRP protocols in the Cisco Catalyst SD-WAN solution. BFD provides a consistent failure detection method to detect forwarding path failures at a uniform rate, therefore enabling faster reconvergence time.
BFD Troubleshooting for Cisco Catalyst SD-WAN	Cisco IOS XE Catalyst SD-WAN Release 17.15.1a Cisco Catalyst SD-WAN Manager Release 20.15.1	This feature provides the ability to troubleshoot BFD protocols using radioactive tracing. You can use this feature to check the device logs and use debugging commands to gather more information about BFD operations.

Information About BFD for Routing Protocols

The following sections provide information about the types of Bidirectional Forwarding Detection (BFD) support for Cisco IOS XE Catalyst SD-WAN devices.

Overview of BFD

In enterprise networks, the convergence of business-critical applications onto a common IP infrastructure is becoming more common. Given how critical data is, these networks are typically constructed with a high degree of redundancy. While such redundancy is desirable, its effectiveness is dependent upon the ability of individual network devices to quickly detect failures and reroute traffic to an alternate path. The detection times in existing protocols are typically greater than one second, and sometimes much longer. For some applications, this duration is too long to be useful. This is where Bi-directional Forwarding Detection (BFD) comes in.

BFD provides rapid failure detection times between forwarding engines, while maintaining low overhead. It also provides a single, standardized method of link/device/protocol failure detection at any protocol layer and over any media, thus enabling faster reconvergence of business-critical applications.

Benefits of Configuring BFD for Routing Protocols

- Fast failure detection times for all media types, encapsulations, topologies, and routing protocols
- Faster reconvergence of applications
- Consistent method of failure detection

How BFD Works in Cisco Catalyst SD-WAN

With the introduction of this feature, the Cisco Catalyst SD-WAN solution now has two types of BFDs that are distinct features that work independently without conflict.

- **BFD Support for Cisco Catalyst SD-WAN Routing Protocols (Legacy BFD):** This feature is termed as legacy BFD because is already available for Cisco IOS XE and is being extended to the Cisco Catalyst SD-WAN solution starting Cisco IOS XE Catalyst SD-WAN Release 17.3.1a.
- **Cisco Catalyst SD-WAN BFD:** This feature is specific to overlay BFD, which is an existing feature in Cisco Catalyst SD-WAN.

  For more information on Cisco Catalyst SD-WAN BFD, see Cisco Catalyst SD-WAN BFD.

Table 2: Differences: BFD for Cisco Catalyst SD-WAN Routing Protocols Versus Cisco Catalyst SD-WAN BFD

BFD for Cisco Catalyst SD-WAN Routing Protocols	Cisco Catalyst SD-WAN BFD

• Runs on both, transport-side and service-side interfaces
• The following protocols can be registered: BGP, OSPF, and EIGRP
  • BGP (transport and service side
  • EIGRP (service side)
  • OSPF and OSPFv3 (service side)
• Detects link failures for peers in terms of whether a peer is up or down
	• Runs on a Cisco Catalyst SD-WAN tunnel to detect failures in the overlay tunnel
	• Is enabled by default and cannot be disabled
	• Is typically enabled for OMP
	• Besides link failures, it also measures latency, loss, and other link statistics used by application-aware routing

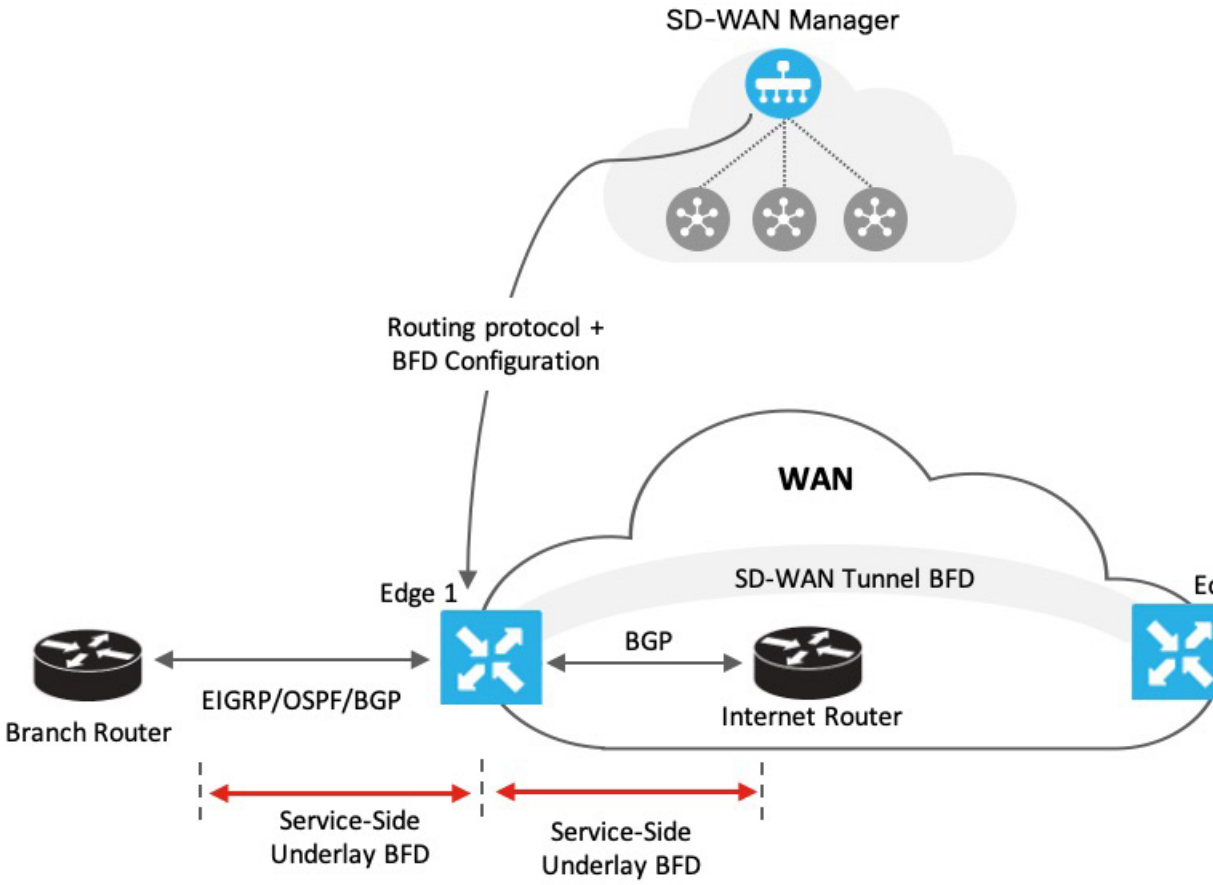

As represented in the image, BFD is configured for a routing protocol through Cisco SD-WAN Manager. Cisco SD-WAN Manager then pushes this configuration to the edge router. In this example, let's assume that OSPF is configured to receive forwarding path detection failure messages from BFD. If there's a physical link failure, OSPF is prompted to shut down its neighbors and restore any routing information it may have advertised to or received from its remote neighbors.

Similarly, the router, Edge 1 is connected to the internet router through its transport interface. BFD is configured for BGP between the transport side of Edge 1 and the internet router. Here, BFD detects the health of the connection and reports any failures.

Supported Protocols and Interfaces

Supported Protocols

The following routing protocols in Cisco Catalyst SD-WAN can be configured to receive forwarding path detection failure messages from BFD:

- BGP
- EIGRP

- OSPF and OSPFv3

Supported Interfaces

- GigabitEthernet
- TenGigabitEthernet
- FiveGigabitEthernet
- FortyGigabitEthernet
- HundredGigabitEthernet
- SVI
- Subinterfaces

Limitations and Restrictions

The following restrictions apply to Cisco IOS XE Catalyst SD-WAN devices in controller mode.

- Only single-hop BFD is supported.
- BFD is not supported for static routes.
- To change the BFD session modes between software mode and hardware mode, you need to remove all existing BFD configuration and reconfigure it.
- BFD is only supported for BGP, EIGRP, OSPF, and OSPFv3.
- BFD for routing protocols in Cisco Catalyst SD-WAN cannot be monitored through Cisco SD-WAN Manager. Use CLI show commands for monitoring BFD for Cisco Catalyst SD-WAN routing protocols.
- Once a BFD session is established, BFD session modes (echo to no echo, and vice-versa; or software to hardware, and vice-versa) don't update immediately after changing the BFD template parameters in Cisco SD-WAN Manager. The BFD mode change takes effect only after the session flaps at least once.

Configure BFD for Routing Protocols

Cisco SD-WAN Manager does not provide an independent template to configure BFD for routing protocols. However, supported protocols can be registered or deregistered to received BFD packets by adding configurations using the CLI add-on template in Cisco SD-WAN Manager. Use the CLI add-on template to configure the following:

- Add a single-hop BFD template with parameters such as timer, multiplier, session mode, and so on.
- Enable the BFD template under interfaces. Only one BFD template can be added per interface.
- Enable or disable BFD for the supported routing protocols. The configuration to enable or disable BFD is different for each of the supported routing protocols: BGP, EIGRP, OSPF, and OSPFv3.

Note Starting with release Cisco IOS XE Catalyst SD-WAN Release 17.15.1a, if sdwan mode is not configured for the tunnel interface, the BFDs become inactive for the tunnel interface.

Enable BFD for Routing Protocols

Configure BFD for Service-Side BGP

1. From the Cisco SD-WAN Manager menu, choose **Configuration** > **Templates**.
2. Click **Feature Templates**.

Note In Cisco vManage Release 20.7.x and earlier releases, **Feature Templates** is titled **Feature**.

3. Click **Add Template**.
4. Choose a device from the device list.
5. Choose the **CLI Add-on Template** under **Other Templates**.
6. Enter the CLI configuration to add a single-hop BFD template and to enable BFD for service-BGP as shown in the following example.

```
bfd-template single-hop t1
 interval min-tx 500 min-rx 500 multiplier 3
 !
interface GigabitEthernet1
   bfd template t1

router bgp 10005
address-family ipv4 vrf 1
    neighbor 10.20.24.17 fall-over bfd
    !
  address-family ipv6 vrf 1
    neighbor 2001::7 fall-over bfd
```

Understanding the CLI Configuration

In this example, a single hop BFD template is created specifying the minimum and maximum interval and the multiplier. Specifying these parameters is mandatory. In addition, you have the option to also specify other BFD parameters such as echo mode (enabled by default), and BFD dampening (off by default). Once created, the BFD template is enabled under an interface (GigabitEthernet1, in this example).

Note To modify a BFD template enabled on an interface, you need to remove the existing template first, modify it, and then enable it on the interface again.

Note If you attach the BFD configuration to a device template that already has a BGP feature template attached, ensure that you update the BGP configuration in the CLI add-on template to include the **no neighbor** *ip-address* **ebgp-multihop** command. This change is required because the **neighbor** *ip-address* **ebgp-multihop** command is activated on the BGP feature template by default.

7. Click **Save**.
8. Attach Feature Template to Device Template

Note For the configuration to take effect, the device template must have a BGP feature template attached to it.

9. Attach the device template to the device.

Configure BFD for Transport-Side BGP

1. From the Cisco SD-WAN Manager menu, choose **Configuration** > **Templates**.
2. Click **Feature Templates**.

Note In Cisco vManage Release 20.7.x and earlier releases, **Feature Templates** is titled **Feature**.

3. Click **Add Template**.
4. Choose a device from the device list.
5. Choose the **CLI Add-on Template** under **Other Templates**.
6. Enter the CLI configuration to add a single-hop BFD template and to enable BFD for transport-BGP as shown in the following example:

```
bfd-template single-hop t1
interval min-tx 500 min-rx 500 multiplier 3
!
interface GigabitEthernet1
bfd template t1
!
router bgp 10005
neighbor 10.1.15.13 fall-over bfd
!
sdwan
interface GigabitEthernet1
tunnel-interface
allow-service bfd
allow-service bgp
```

Understanding the CLI Configuration

In this example, a single hop BFD template is created specifying the minimum and maximum interval and the multiplier. Specifying these parameters is mandatory. In addition, you have the option to also specify other BFD parameters such as echo mode (enabled by default), and BFD dampening (off by

default). Once created, the BFD template is enabled under an interface (GigabitEthernet1, in this example). In this example, GigabitEthernet1 is also the source of the SD-WAN tunnel. Allowing service under the tunnel interface of GigabitEthernet1 ensures that BGP and BFD packets pass over the tunnel.

Note To modify a BFD template enabled on an interface, you need to remove the existing template first, modify it, and then enable it on the interface again.

Note If you attach the BFD configuration to a device template that already has a BGP feature template attached, ensure that you update the BGP configuration in the CLI add-on template to include the **no neighbor** *ip-address* **ebgp-multihop** command. This change is required because the **neighbor** *ip-address* **ebgp-multihop** command is activated on the BGP feature template by default.

7. Click **Save**.
8. Attach Feature Template to Device Template

Note For the configuration to take effect, the device template must have a BGP feature template attached to it.

9. Attach the device template to the device.

Configure BFD for Service-Side EIGRP

1. From the Cisco SD-WAN Manager menu, choose **Configuration** > **Templates**.
2. Click **Feature Templates**.

Note In Cisco vManage Release 20.7.x and earlier releases, **Feature Templates** is titled **Feature**.

3. Click **Add Template**.
4. Choose a device from the device list.
5. Choose the **CLI Add-on Template** under **Other Templates**.
6. Enter the CLI configuration to add a single-hop BFD template enable BFD for EIGRP as shown in the example below.

```
bfd-template single-hop t1
 interval min-tx 500 min-rx 500 multiplier 3
 !
interface GigabitEthernet5
   bfd template t1

router eigrp myeigrp
address-family ipv4 vrf 1 autonomous-system 1
    af-interface GigabitEthernet5
     bfd
```

Understanding the CLI Configuration

In this example, a single hop BFD template is created specifying the minimum and maximum interval and the multiplier. Specifying these is mandatory. In addition, you have the option to also specify other BFD parameters such as echo mode (enabled by default), and BFD dampening (off by default).

Once created, the BFD template is enabled under an interface (GigabitEthernet5, in this example).

Note To modify a BFD template enabled on an interface, you first need to remove the existing template, modify it, and enable it on the interface again.

7. Click **Save**.
8. Attach Feature Template to Device Template

Note For the configuration to take effect, the device template must have an EIGRP feature template attached to it.

9. Attach the device template to the device.

Configure BFD for Service-Side OSPF and OSPFv3

1. From the Cisco SD-WAN Manager menu, choose **Configuration** > **Templates**.
2. Click **Feature Templates**.

Note In Cisco vManage Release 20.7.x and earlier releases, **Feature Templates** is titled **Feature**.

3. Click **Add Template**.
4. Choose a device from the device list.
5. Choose the **CLI Add-on Template** under **Other Templates**.
6. Enter the CLI configuration to add a single-hop BFD template enable BFD for OSPF and OSPFv3 as shown in the examples below.

OSPF

```
bfd-template single-hop t1
interval min-tx 500 min-rx 500 multiplier 3
!
interface GigabitEthernet5
bfd template t1
!
interface GigabitEthernet1
bfd template t1
!
router ospf 1 vrf 1
 bfd all-interfaces
!
```

OSPFv3

```
bfd-template single-hop t1
 interval min-tx 500 min-rx 500 multiplier 3

interface GigabitEthernet5
   bfd template t1
router ospfv3 1
 address-family ipv4 vrf 1
  bfd all-interfaces
```

Understanding the CLI Configuration

In these examples, a single hop BFD template is created specifying the minimum and maximum interval and the multiplier. Specifying these is mandatory. In addition, you have the option to also specify other BFD parameters such as echo mode (enabled by default), and BFD dampening (off by default).

Once created, the BFD template is enabled under an interface (GigabitEthernet5, in this example).

Note To modify a BFD template enabled on an interface, you first need to remove the existing template, modify it, and enable it on the interface again.

7. Click **Save**.
8. Attach the CLI Add-on Template with this configuration to the device template.

Note For the configuration to take effect, the device template must have an OSPF feature template attached to it.

9. Attach the device template to the device.

Attach Feature Template to Device Template

After creating a CLI add-on template to enable BFD, attach the template to the device template for the configuration to take effect. Follow this procedure to attach the configuration to a device template. Ensure that the device template you attach the feature template to already has the relevant feature template (BGP, OSPF, EIGRP) attached to it.

1. From the Cisco SD-WAN Manager menu, choose **Configuration** > **Templates**.
2. Click **Device Templates**.

Note In Cisco vManage Release 20.7.x and earlier releases, **Device Templates** is titled **Device**.

3. Click **Create Template** and choose **From Feature Template** from the drop-down options.
4. From the **Device Model** drop-down options, choose a device. Enter a name and description for the template.
5. Click **Create**.
6. Click **Additional Templates**.

7. In the **CLI Add-on Template** field, choose the CLI add-on template you configured to enable BFD for routing protocols.
8. Click **Create**.

Next: Attach device template to device

Configure BFD for Routing Protocols Using CLI

To configure BFD for BGP, EIGRP, OSPF, and OSPF3 using device CLI, follow the steps in this topic.

Create BFD Template

Create a single-hop BFD template as shown in the example below.

```
bfd-template single-hop t1
 interval min-tx 500 min-rx 500 multiplier 3
```

Note The CLI configuration for creating a BFD template remains the same irrespective of the protocol you configure it for.

Enable BFD for Service-Side BGP

This example shows that BGP is configured, BFD is enabled on the interface under VRF 1, and then on service-side BGP.

```
interface GigabitEthernet5
bfd template t1
!
router bgp 10005
  bgp log-neighbor-changes
  distance bgp 20 200 20
  !
  address-family ipv4 vrf 1
  bgp router-id 10.20.24.15
  redistribute connected
  neighbor 10.20.24.17 remote-as 10007
  neighbor 10.20.24.17 activate
  neighbor 10.20.24.17 send-community both
  neighbor 10.20.24.17 maximum-prefix 2147483647 100
  neighbor 10.20.24.17 fall-over bfd
  exit-address-family
  !
  address-family ipv6 vrf 1
  bgp router-id 10.20.24.15
  neighbor 2001::7 remote-as 10007
  neighbor 2001::7 activate
  neighbor 2001::7 send-community both
  neighbor 2001::7 maximum-prefix 2147483647 100
  neighbor 2001::7 fall-over bfd
  exit-address-family
```

Enable BFD for Transport-Side BGP

```
interface GigabitEthernet1
bfd template t1
!
router bgp 10005
bgp router-id   10.1.15.15
bgp log-neighbor-changes
distance bgp 20 200 20
neighbor 10.1.15.13 remote-as 10003
neighbor 10.1.15.13 fall-over bfd
address-family ipv4 unicast
neighbor 10.1.15.13 remote-as 10003
neighbor 10.1.15.13 activate
neighbor 10.1.15.13 maximum-prefix 2147483647 100
neighbor 10.1.15.13 send-community both
redistribute connected
exit-address-family
!
timers bgp 60 180

sdwan
interface GigabitEthernet1
tunnel-interface
allow-service bgp
allow-service bfd
```

Enable BFD for EIGRP

This example shows that EIGRP is configured, BFD is enabled on the interface under VRF 1, and then on service-side EIGRP.

```
interface GigabitEthernet5
bfd template t1
!
router eigrp myeigrp
address-family ipv4 vrf 1 autonomous-system 1
    af-interface GigabitEthernet5
     no dampening-change
     no dampening-interval
     hello-interval 5
     hold-time      15
     split-horizon
     bfd
     exit-af-interface
    !
    network 10.20.24.0 0.0.0.255
    topology base
     redistribute connected
     redistribute omp
     exit-af-topology
    !
    exit-address-family
   !
```

Enable BFD for OSPFv3

This example shows that OSPFv3 is configured, BFD is enabled on the interface under VRF 1, and then on service-side EIGRP.

```
interface GigabitEthernet5
   bfd template t1
```

```
   ospfv3 1 ipv4 area 0
   ospfv3 1 ipv4 dead-interval 40
   ospfv3 1 ipv4 hello-interval 10
   ospfv3 1 ipv4 network broadcast
   ospfv3 1 ipv4 priority 1
   ospfv3 1 ipv4 retransmit-interval 5
   ospfv3 1 ipv6 area 0
   ospfv3 1 ipv6 dead-interval 40
   ospfv3 1 ipv6 hello-interval 10
   ospfv3 1 ipv6 network broadcast
   ospfv3 1 ipv6 priority 1
   ospfv3 1 ipv6 retransmit-interval 5


router ospfv3 1
 address-family ipv4 vrf 1
  area 0 normal
  bfd all-interfaces
  router-id 10.20.24.15
  distance 110
  exit-address-family
 !
 address-family ipv6 vrf 1
  area 0 normal
  bfd all-interfaces
  router-id 10.20.24.15
  distance 110
  exit-address-family
 !
!
exit
```

Monitor and Verify BFD Configuration

This sections provides a list of commands that you can run to verify your BFD configuration.

Run the **show bfd interface** command to check the BFD template under an interface.

```
Device# show bfd interface
Interface Name: GigabitEthernet5
 Interface Number: 11
 Configured bfd interval using bfd template: 12383_4T1
 Min Tx Interval: 50000, Min Rx Interval: 50000, Multiplier: 3
```

Verify BFD Configuration for BGP

Run the **show bfd neighbors client bgp ipv4** command to check the status of the BFD session.

```
Device# show bfd neighbors client bgp ipv4

IPv4 Sessions
NeighAddr                              LD/RD         RH/RS     State     Int
10.20.24.17                             1/1          Up        Up        Gi5
```

Verify BFD Configuration for EIGRP

Run the **show bfd neighbors client eigrp** command to check the status of the BFD session.

```
Device# show bfd neighbors client eigrp
```

```
IPv4 Sessions
NeighAddr                                LD/RD         RH/RS     State     Int
10.20.24.17                               1/1          Up        Up        Gi5
```

Verify BFD Configuration for OSPF

Run the **show bfd neighbors client ospf** command to check the status of the BFD session.

```
Device# show bfd neighbors client ospf
IPv4 Sessions
NeighAddr            LD/RD         RH/RS     State     Int
10.20.24.17          1/1          Up        Up        Gi5
```

Troubleshoot Common BFD Errors

Check Control Connections

If you experience issues with BFD, start by checking the control connection between Cisco SD-WAN Manager and the edge router by running the **show sdwan control connections** command.

```
Device#show sdwan control connections
                                                                                                                PEER
                                          PEER
CONTROLLER
PEER    PEER PEER            SITE       DOMAIN PEER                                                             PRIV
  PEER                                    PUB
GROUP
TYPE    PROT SYSTEM IP       ID         ID     PRIVATE IP                                                       PORT
  PUBLIC IP                               PORT  LOCAL COLOR     PROXY STATE UPTIME       ID

------------------------------------------------------------------------------------------------------------------------

vsmart  dtls 172.16.255.19   100        1      10.0.5.19                                                        12355
 10.0.5.19                                 12355 lte             No    up     0:12:45:44  0

vsmart  dtls 172.16.255.20   200        1      10.0.12.20                                                       12356
 10.0.12.20                                12356 lte             No    up     0:15:59:45  0

vmanage dtls 172.16.255.22   200        0      10.0.12.22                                                       12346
 10.0.12.22                                12346 lte             No    up     0:15:59:45  0
   < ----  up
```

Issues in Pushing Device Template to Device

If you identify issues with pushing the device template to the device, collect debug logs on the edge device as shown below.

```
debug netconf all
request platform soft system shell
tail -f /var/log/confd/cia-netconf-trace.log
```

If Cisco SD-WAN Manager has successfully pushed the configuration to the device and the issue still persists, run the **show sdwan running-config** command to view all details related to BFD.

Issues with Transport-Side BFD

If the transport-side BFD session is down, check the packet filter data under the Cisco Catalyst SD-WAN tunnel interface to ensure that you have allowed the BFD packets to pass through on the transport side. Look for `allow-service bgp` and `allow-service bfd` in the output.

```
Device#show sdwan running-config  | sec sdwan
 tunnel mode sdwan
sdwan
 interface GigabitEthernet1
  tunnel-interface
   encapsulation ipsec
   color lte
   allow-service bgp
   allow-service bfd
   ……
```

Troubleshoot BFD Using Radioactive Tracing

Minimum releases: Cisco IOS XE Catalyst SD-WAN Release 17.15.1a and Cisco Catalyst SD-WAN Manager Release 20.15.1

BFD troubleshooting involves diagnosing and resolving issues related to the BFD protocol, which is used to detect faults between the devices. You can use this feature to check the device logs and use debugging commands to gather more information about BFD operations.

Radioactive tracing helps in selective debugging of a session. Tracing is enabled across the layers for intended BFD session that is identified by tloc-pair or a local discriminator. It enables debug level traces automatically for all the modules while processing a packet that matches the condition.

The following **show** and **debug** commands are used in BFD troubleshooting:

- **debug platform condition start**
- **debug platform condition feature sdwan controlplane bfd**
- **show platform hardware qfp active feature bfd datapath**
- **show logging profile sdwan internal filter**

For more information on these show commands, see the chapter Troubleshooting Commands in the Cisco IOS XE SD-WAN Qualified Command Reference guide.