Contents

1 Configuring BFD

1.1 Introduction

Bidirectional forwarding detection (BFD) provides a light-load and fast method for detecting the connectivity of the forwarding path between two adjacent devices. It quickly detects faults on the bidirectional forwarding path between two devices for upper layer protocols such as routing protocols and Multi-Protocol Label Switching (MPLS) so that measures can be taken in a timely manner to guarantee services. BFD minimizes the impact of faults on services and improves network availability.

1.1.1 Basic Concepts

1. Packet Format

Detection packets transmitted by BFD are User Datagram Protocol (UDP) packets, which are classified into two types:

- Control packets: The protocol defines a specific format, as shown in Figure 1-1. Currently, there are two versions (v0 and v1) for the format of control packets: v1 is adopted by default for establishing a BFD session; if a device receives packets of v0 from the peer system, the device automatically switches to v0.
- Echo packets: Echo packets are only used by the local system of a BFD session, and the protocol does not define a specific format.

Figure 1-1 BFD Packet Format

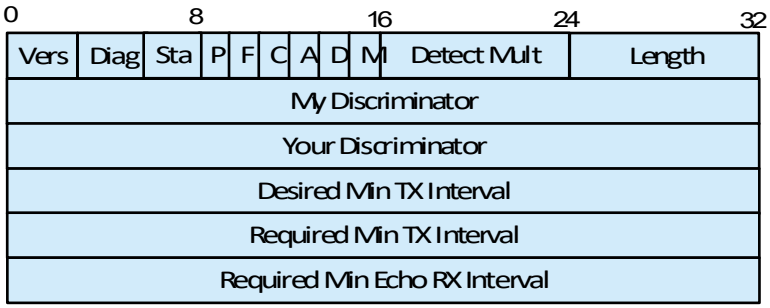

Table 1-1 Description of BFD Packet Fields

Field	Description
Vers	BFD protocol version number, which is 1 currently.
Diag	Causes for the local system to switch from up state to another state last time, including: **0**: Indicates no diagnostic information. **1**: Controls timeout detection.

Field	Description
	2: Indicates that the echo function failed. **3**: Indicates that a neighbor advertises session down. **4**: Indicates forwarding plane reset. **5**: Indicates that the path failed. **6**: Indicates that the connection path failed. **7**: Indicates that the local system enters the AdminDown state.
Sta	BFD local session states, including: **0**: Indicates the AdminDown state. **1**: Indicates the Down state. **2**: Indicates the Init state. **3**: Indicates the Up state.
P	The transmitter sets this flag to **1** in a BFD packet upon parameter changes. The receiver must immediately respond to this packet.
F	This flag must be set to **1** in the response packet for responding to the P flag that is set to **1**.
C	Forwarding/control separation flag. If it is set to **1**, the control plane changes do not affect BFD detection. For example, if the control plane is Open Shortest Path First (OSPF), when OSPF is restarted or experiences graceful restart (GR), BFD can continue to detect the link state.
A	Authentication flag. If it is set to **1**, a session needs to be authenticated.
D	Query request. If it is set to **1**, the transmitter desires the query mode for detecting links.
M	It is used in point-to-multipoint applications in the future. It must be set to **0** currently.
Detect Mult	Detection timeout multiplier. It is used by the detector to calculate the detection timeout time.
Length	Packet length.
My Discriminator	Discriminator of the local end of a BFD session.
Your Discriminator	Discriminator of the remote end of a BFD session
Desired Min Tx Interval	Minimum BFD packet transmission interval supported by the local end.
Required Min RX Interval	Minimum BFD packet receiving interval supported by the local end.

Field	Description
Required Min Echo RX Interval	Minimum echo packet receiving interval supported by the local end. It is set to **0** if the local end does not support the echo function.
Auth Type	Authentication type. It is an optional field and can be set to the following values: Simple Password Keyed MD5 Meticulous Keyed MD5 Keyed SHA1 Meticulous Keyed SHA1
Auth Length	Authentication data length.
Authentication Data	Authentication data area.

2. Session State

BFD migrates the state machine based on the local session state and the BFD packets received from the peer end. A BFD state machine is established and torn down using a three-way handshake mechanism, to ensure that both ends know the state change. A BFD session can be in any of the following four basic states:

- Down: Indicates that a session is in the down state or is established just now.
- Init: Indicates that the local system has communicated with the peer system and desires to bring the session to the up state.
- Up: Indicates that a session has been established successfully.
- AdminDown: Indicates that a session is in the AdminDown state.

3. Transmission Interval and Detection Time

During the establishment of a BFD session, both ends negotiate about BFD parameters to determine the transmission interval and detection time. After a BFD session is established, both ends dynamically negotiate about BFD parameters (for example, the minimum transmission interval and minimum receiving interval). After protocols at both ends transmit relevant negotiation packets, they adopt the new transmission interval and detection time, without affecting the current state of the session.

1.1.2 Session Establishment

BFD detection starts from the establishment of a BFD session.

1. Session Establishment Process

BFD itself is unable to discover neighbors. It needs an upper-layer protocol to specify the neighbor, with which BFD needs to establish a session. As shown in Figure 1-1, two devices running OSPF and BFD are connected through an L2 device.

Figure 1-1 BFD Session Establishment

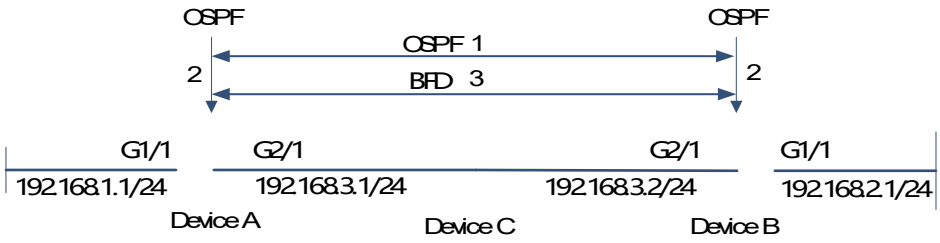

The BFD session establishment process is as follows:

- OSPF discovers a neighbor and establishes a connection with the neighbor.
- OSPF instructs BFD to establish a session with the neighbor.
- BFD establishes a session with the neighbor.

2. Session Establishment Mode

The BFD protocol specifies that a BFD session can be established in two modes:

- Active mode: Before the establishment of a session, BFD actively transmits a control packet for establishing a BFD session regardless of whether it receives a control packet for establishing a BFD session from the peer end.
- Passive mode: BFD does not actively transmit a control packet for establishing a BFD session before a session is established but waits till it receives a control packet for establishing a BFD session from the peer end.

Specification

The passive mode is not supported currently.

3. Negotiation of Session Parameters

Both ends negotiate about BFD session parameters during the establishment of a BFD session, to determine the transmission interval and detection time. Pay attention to the following points:

- BFD session parameters (including the minimum transmission interval, minimum receiving interval, and detection timeout multiplier) must be set for interfaces at both ends. Otherwise, a BFD session cannot be established.
- Interfaces at both ends negotiate about BFD session parameters and detect the session based on the parameters during the establishment of a BFD session.
- After a BFD session is established, both ends dynamically negotiate about BFD parameters (for example, the minimum transmission interval and minimum receiving interval). After protocols at both ends transmit relevant negotiation packets, they adopt the new transmission interval and detection time, without affecting the current state of the session.
- BFD session parameters can be updated only after all the original sessions on a port complete negotiation; otherwise an error is prompted. You can run a related command to query the negotiation status of the original sessions.

1.1.3 Session Detection

Link detection starts after the establishment of a BFD session. BFD periodically transmits BFD control packets. If it fails to receive BFD packets from the peer end within the detection time, it deems that the session is down and instructs the associated application to accelerate the convergence.

1. Detection Process

As shown in Figure 1-1, two devices running OSPF and BFD are connected through an L2 device.

Figure 1-1 Topology for BFD Session Detection

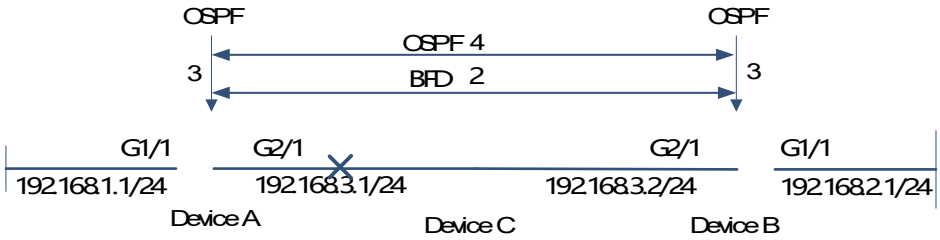

The handling process after a BFD session is down is as follows:

- The link between Device A and Device C fails.
- The BFD session between Device A and Device B detects the failure.
- BFD notifies the local OSPF that the forwarding path to the neighbor is faulty.
- OSPF processes the neighbor down situation. If a backup forwarding path is available, it starts protocol convergence to enable the backup forwarding path.

2. Detection Mode

BFD supports the following detection modes:

- Asynchronous mode

  In asynchronous mode, systems transmit BFD control packets periodically to each other. If a system fails to receive BFD control packets from the peer end within the detection time, it advertises that the session is down.

- Query mode

  In query mode, it is assumed that each system has an independent method for confirming its connection with other systems. After a BFD session is established, the system stops transmitting BFD control packets unless it needs to explicitly verify the connectivity. In such a case, the system transmits a shot-sequence BFD control packet. If the system fails to receive a returned packet within the detection time, it advertises that the session is down. If it receives a response from the peer end, the forwarding path is reachable.

- Echo mode

  In echo mode, the local system periodically transmits BFD echo packets and a remote system receives and loops back the packets through the forwarding path. If the local system fails to receive several consecutive echo packets within the detection time, it advertises that the session is down. The echo function can be used together with the preceding two detection modes. The echo packet detection function does not require the involvement of the control plane of the remote system. Packets are returned

by the forwarding plane of the remote system, which reduces the delay and ensures faster fault detection in comparison with transmission of control packets. Enabling the echo function in asynchronous mode greatly reduces transmission of control packets because the detection is accomplished by the echo function. Enabling the echo function in query mode thoroughly cancels transmission of control packets after a session is established. The echo function must be enabled at both ends of a BFD session. Otherwise, the echo function does not take effect.

The echo mode includes the echo mode and one-armed echo mode. The one-armed echo mode is applicable to the single-hop IP link scenario and generally used in two directly connected devices, where one device supports the BFD function and the other does not. The echo mode requires that the echo mode be configured for both detection parties.

Specification

- The query mode is not supported and cannot be configured at present.
- Only BFD session version 1 supports the BFD echo mode.
- The echo mode is not supported for an IPv6 BFD session with the link-local address as the source or destination address.

1.1.4 BFD Support for Applications

By supporting BFD, the associated applications can utilize the fast fault detection of BFD to improve the protocol convergence performance. In general, the fault detection time can be shortened to less than 1 second. After BFD support for a certain application is enabled, BFD rapidly detects faults occurring on an established BFD session. When a link fault occurs, BFD can rapidly identify the fault and notify the associated application to process the fault, thereby improving its convergence. Currently, BFD supports the following applications:

1. BFD Support for RIP

After BFD support for the Routing Information Protocol (RIP) is enabled, RIP can utilize the faster fault detection feature of BFD than the hello mechanism of RIP to improve the protocol convergence. In general, the fault detection time can be shortened to less than 1 second.

Note

For more details about BFD support for RIP, see "Configuring RIP."

2. BFD Support for OSPF

After BFD support for OSPF is enabled, OSPF can utilize the faster fault detection feature of BFD than the hello mechanism of OSPF to improve the protocol convergence. In general, the fault detection time can be shortened to less than 1 second.

Note

For more details about BFD support for OSPF, see "Configuring OSPFv2."

3. BFD Support for OSPFv3

After BFD support for OSPFv3 is enabled, OSPFv3 can utilize the faster fault detection feature of BFD than the hello mechanism of OSPFv3 to improve the protocol convergence. In general, the fault detection time can be shortened to less than 1 second.

Note

For more details about BFD support for OSPFv3, see "Configuring OSPFv3."

4. BFD Support for BGP

After BFD support for Border Gateway Protocol (BGP) is enabled, BGP can utilize the faster fault detection feature of BFD than the hello mechanism of BGP to improve the protocol convergence. In general, the fault detection time can be shortened to less than 1 second.

Note

For more details about BFD support for BGP, see "Configuring BGP."

5. BFD Support for IS-IS

The Intermediate System to Intermediate System (IS-IS) protocol dynamically discovers neighbors through hello packets. After the BFD function is enabled on IS-IS, a BFD session is set up with a neighbor in up state. The BFD mechanism is used to detect the neighbor state. Once a neighbor failure is detected through BFD, IS-IS performs network convergence immediately. The convergence time can be reduced from 30s to less than 1s. By default, IS-IS hello packets are sent at an interval of **10** seconds in a P2P network, and the time required to detect a neighbor failure is three times the hello packet transmission interval.

Note

For more details about BFD support for IS-IS, see "Configuring IS-IS."

6. BFD Support for Static Routing

After BFD support for static routing is enabled, BFD prevents routers from selecting an unavailable static route as the forwarding path during routing and enables routers to rapidly switch to an available backup forwarding path. Different from dynamic routing protocols, static routing does not have the ND mechanism. Therefore, after BFD support for static routing is configured, the next-hop reachability of a static route relies on the BFD session state. If a BFD session detects a fault, the next hop of a static route is unreachable and the static route is not added to the routing information base (RIB). If the remote system deletes a BFD session during the establishment of the BFD session, the BFD session becomes down. In this case, the system ensures that the forwarding behavior of static routing is not affected.

Note

For more details about BFD support for static routing, see "Configuring Static Routing."

7. BFD Support for PBR

After BFD support for policy-based routing (PBR) is configured, BFD prevents routers from selecting an unavailable policy-based route as the forwarding path during routing and enables routers to rapidly switch to an available backup forwarding path. BFD support for PBR is equivalent to BFD support for static routing. BFD tracks and detects the forwarding path to a specified neighbor. When a BFD session fails, BFD notifies the PBR that the next hop is unreachable. Then, the policy-based route to the next hop does not take effect. If the remote system deletes a BFD session during the establishment of the BFD session, the BFD session becomes down. In this case, the system ensures that the PBR forwarding behavior is not affected.

Note

For more details about BFD support for PBR, see "Configuring PBR."

8. BFD Support for VRRP

The BFD support for the Virtual Router Redundancy Protocol (VRRP) can replace the hello mechanism of VRRP to rapidly detect the running status of the active and standby devices. When a fault occurs, it accelerates the active/standby device switching and improves network performance. In general, the fault detection time can be shortened to less than 1 second. VRRP can also utilize BFD to track a specified neighbor. If a BFD session detects that the forwarding path to the neighbor fails, VRRP automatically lowers the VRRP priority to a certain extent to trigger active/standby device switching. This configuration takes effect only when a dynamic routing protocol or other applications notify BFD to establish a session with a neighbor.

Note

For more details about BFD support for VRRP, see *Configuring VRRP*.

9. BFD Support for VRRP Plus

The BFD support for VRRP Plus can replace the balancing virtual forwarder (BVF) detection conducted by the balancing virtual gateway (BVG) of VRRP Plus to rapidly detect the running status of BVFs. When a fault occurs, it accelerates the forwarding entity switching and improves network performance. In general, the fault detection time can be shortened to less than 1 second. VRRP Plus is based on the VRRP protocol. Therefore, no additional configuration is required for BFD support for VRRP Plus and only VRRP needs to be enabled on devices at both ends and a BFD session is correctly associated.

Note

For more details about BFD support for VRRP Plus, see "Configuring VRRP Plus."

10. BFD Support for MPLS

The BFD support for Multiprotocol Label Switching (MPLS) enables label switched paths (LSPs) to use BFD to rapidly detect the neighbor status. The following detection modes are supported:

- BFD detects static LSPs.
- BFD detects LSPs generated by the Label Distribution Protocol (LDP).

- BFD detects reverse LSPs by using the IP addresses.

Note

For more details about BFD support for MPLS, see "MPLS Basics."

11. BFD Support for L3 Interfaces

BFD supports modifying the status of an L3 interface. In interface configuration mode, run the **bfdbindpeer-ip** command to detect the address of a specified directly connected L3 interface. After this CLI command is executed, a BFD session is created and the status of an L3 interface can be changed based on the detection result of the BFD session, for example, BFD Down or BFD Up. This function is often used in various types of fast reroute (FRR), which use BFD to detect the interface status to implement FRR switching.

Note

Only LDP FRR switching is supported when BFD support for L3 interfaces is configured.

12. BFD Support for AP Member Ports

After BFD support for aggregate port (AP) member ports is enabled, BFD can rapidly detect a fault occurring on a member port link so that traffic on this link is rapidly distributed to other effective member port links. In general, the fault detection time can be shortened to less than 1 second.

Note

For more details about BFD support for AP, see "Link Aggregation Port."

1.1.5 BFD Protection

If a BFD-enabled device is attacked (for example, attacked by a large number of ping packets), BFD sessions will flap. You can enable the BFD protection to provide protection. If both BFD and BFD protection are enabled on a device, the device discards the BFD packets received from the previous-hop device, which affects the establishment of a BFD session between the previous-hop device and other devices.

1.1.6 BFD Flapping Advertisement Dampening

A BFD session may frequently switch between the down and up states due to link instability. As a result, an associated application (such as static routing) may frequently switch forwarding paths, causing the frequent switching of the BFD session between the down and up states. This function allows users to set the delay for status change advertisement. After a BFD session is up for a certain period of time, BFD informs an associated application of BFD up. Otherwise, BFD informs it of BFD down.

1.1.7 SBFD Reflector Establishment

Compared with seamless bidirectional forwarding detection (SBFD), the BFD technology is relatively mature. However, when a large number of sessions are configured for link detection, the negotiation time of BFD's existing state machine becomes longer and turns into a bottleneck of the entire system. As a simplified mechanism of BFD, SBFD shortens the negotiation time. In large-scale commercial scenarios, SBFD implements faster communication fault detection and notifies the upper layer service in case of a fault. Like

BFD sessions, SBFD sessions also distinguish different sessions through My Discriminator and Your Discriminator. Since the reflector of SBFD does not perceive the detection service and is responsible for only looping back packets, only a reflector discriminator needs to be configured by running the **reflector discriminator** command at the reflector. This discriminator is Your Discriminator for the initiator. In a network, one initiator can be deployed to map to multiple reflectors, and multiple reflector discriminators can be configured for one reflector. All the reflector discriminators in the network must be globally unique.

1.1.8 Protocols and Standards

- draft-ietf-bfd-base-09: Bidirectional Forwarding Detection
- draft-ietf-bfd-generic-05: Generic Application of BFD
- draft-ietf-bfd-mib-06: Bidirectional Forwarding Detection Management Information Base
- draft-ietf-bfd-v4v6-1hop-09: BFD for IPv4 and IPv6 (Single Hop)
- draft-ietf-bfd-multihop-07: BFD for IPv4 and IPv6 (Multihop)
- draft-ietf-bfd-mpls-07: BFD For MPLS LSPs

Specification

Currently, draft-ietf-bfd-mib-06 and draft-ietf-bfd-multihop-07 are not supported.

1.2 Restrictions and Guidelines

- You are advised to keep the parameter configuration consistent at both ends of a BFD session. The purpose is to ensure that application protocols associated with BFD take effect simultaneously and prevent occurrence of one-way forwarding path due to different dampening time at both ends.
- Transmission bandwidth differences of different interfaces need to be taken into account during parameter configuration. If the configured minimum Tx interval and minimum Rx interval are very small, BFD may occupy excessive bandwidth and affect data transmission.
- BFD session parameters can be configured only after all the sessions on the current port complete negotiation. Otherwise, a configuration error is prompted. Run the **show bfd neighbors parm-consult** command to display the negotiation status of all the sessions.
- Ensure that BFD support is enabled on neighbors of a BFD session. Otherwise, a BFD session cannot be established. If a dynamic routing protocol or another application instructs BFD to establish a session with a neighbor, the BFD session can also be established.
- For a single-hop session whose source and destination IP addresses belong to the same network segment, if the interface specified by a BFD session is different from the actual BFD packet outbound interface because of IP routing, or if the interface specified during BFD session creation is different from the actual BFD packet inbound interface, a BFD session cannot be established. For a multi-hop session whose source and destination IP addresses belong to different network segments, associated applications do not need to specify an interface and BFD transmits and receives packets based on IP routing. Such BFD sessions can be established provided that returned BFD packets can be received.
- When BFD session traffic is sent and no BFD session parameters are configured on the device (the **show bfd neighbors** command shows that the number of BFD sessions is 0), the CPU protect policy (CPP) entries are not delivered. Therefore, the BFD traffic sent to the local device is not matched with the CPP

entries and not sent to the local device. In this case, BFD sessions must be configured for sending the traffic to the local device.

- In the process that the forwarding plane of the peer device returns echo packets transmitted by the local end to the local end, the echo packets may be lost due to congestion of the peer device, causing a session detection failure. In this case, you need to configure a quality of service (QoS) policy to ensure that echo packets are processed preferentially or disable the echo function.
- The echo detection function of BFD does not support multi-hop detection. Ensure that the echo function is disabled when configuring multi-hop detection.
- The echo mode takes effect only after it is enabled at both ends of a BFD session.
- Before enabling the echo mode of BFD, run the **no ip redirects** command on the neighbors of a BFD session to disable the function of sending Internet Control Message Protocol (ICMP) redirection packets, and run the **no ip deny land** command to disable the distributed denial of service (DDoS) function to prevent the land-based attack.

1.3 Configuration Task Summary

BFD configuration includes the following tasks:

(1) Configuring BFD Basic Functions

- Configuring BFD Parameters
- (Optional) Configuring the BFD Echo Mode
- (Optional) Configuring the BFD Slow Timer
- (Optional) Configuring BFD Support for L3 Interfaces
- Configuring BFD Support for Applications

(2) (Optional) Configuring BFD Protection

(3) (Optional) Configuring BFD Flapping Advertisement Dampening

(4) (Optional) Configuring an SBFD Reflector

1.4 Configuring BFD Basic Functions

1.4.1 Overview

You can establish BFD sessions and configure BFD support for applications to perform link fault detection.

1.4.2 Configuration Tasks

The basic function configuration of BFD includes the following tasks:

(1) Configuring BFD Parameters

(2) (Optional) Configuring the BFD Echo Mode

(3) (Optional) Configuring the BFD Slow Timer

(4) (Optional) Configuring BFD Support for L3 Interfaces

(5) Configuring BFD Support for Applications

1.4.3 Configuring BFD Parameters

1. Overview

- Unless otherwise specified, for single-hop sessions, a BFD session should be established between devices at both ends to be detected by BFD and BFD parameters must be configured at the egresses. Multi-hop sessions use global parameters for session establishment by default. If different negotiation parameters need to be configured for different multi-hop sessions, BFD templates need to be associated with the sessions.
- Transmission bandwidth differences of different interfaces need to be taken into account during parameter configuration. If the configured minimum Tx interval and minimum Rx interval are very small, BFD may occupy excessive bandwidth and affect data transmission.

2. Restrictions and Guidelines

- Negotiation parameters can be configured only after all the original sessions on a port complete negotiation.

3. Procedure

(1) Enter the privileged EXEC mode.

enable

(2) Enter the global configuration mode.

configure terminal

(3) Enter the interface configuration mode.

interface *interface-type interface-number*

(4) Configure BFD parameters.

bfd interval *milliseconds* **min_rx** *milliseconds* **multiplier** *interval-multiplier*

No BFD session parameters are configured by default.

1.4.4 Configuring the BFD Echo Mode

1. Overview

- (Optional) Ports run in asynchronous mode by default. If a BFD session needs to run in echo mode, the echo mode needs to be configured. The echo mode includes the echo mode and one-armed echo mode.
- Complete the configuration on ports.
- A BFD session runs in asynchronous mode as long as either of devices at both ends is configured to run in asynchronous mode. If routers at both ends are configured to run in echo mode by default, a BFD session finally runs in echo mode. This function is disabled by default and can be enabled after the **bfd echo** command is executed.
- Run the **bfd echo one-arm** command to enable the one-armed echo mode.

2. Restrictions and Guidelines

By default, when BFD session parameters are configured, the system automatically enables the echo mode. To disable the BFD echo function, run the **no bfd echo** command in interface configuration mode. The minimum Tx interval and minimum Rx interval of echo packets adopt the **Interval milliseconds** and **min_rx milliseconds** parameters configured for a session.

Before enabling the echo mode of BFD, run the **no ip redirects** command on the neighbors of a BFD session to disable the function of sending ICMP redirection packets, and run the **no ip deny land** command to disable the DDoS function to prevent the land-based attack.

The BFD one-armed echo function and **ip redirect-drop** cannot be configured at the same time, otherwise the function fails.

3. Procedure

(1) Enter the privileged EXEC mode.

enable

(2) Enter the global configuration mode.

configure terminal

(3) Enter the interface configuration mode.

interface *interface-type interface-number*

(4) Configure the BFD echo mode.

bfd echo [**one-arm**]

The echo mode is disabled for a BFD session by default.

1.4.5 Configuring the BFD Slow Timer

1. Overview

- The default slow timer is **3000** milliseconds. The value can be changed as required.
- Configure this function in global configuration mode.
- In BFD echo mode or during session establishment, control packets are sent based on the slow timer. A larger value indicates that the required time for negotiating and establishing a BFD session is longer and the time required for transmitting slow BFD packets in echo mode is longer.

2. Restrictions and Guidelines

This command is used to specify the slow timer in echo mode.

3. Procedure

(1) Enter the privileged EXEC mode.

enable

(2) Enter the global configuration mode.

configure terminal

(3) Configure the BFD slow timer.

bfd slow-timer [*interval*]

The default slow timer time in the echo mode is **2000** milliseconds.

1.4.6 Configuring BFD Support for L3 Interfaces

1. Overview

- Currently, this function is used only when MPLS LDP is used for FRR switching.
- Complete the configuration on ports.

2. Restrictions and Guidelines

This command is used to enable BFD support for L3 interfaces so as to rapidly detect connectivity of L3 interfaces.

3. Procedure

(1) Enter the privileged EXEC mode.

enable

(2) Enter the global configuration mode.

configure terminal

(3) Enter the interface configuration mode.

interface *interface-type interface-number*

(4) Configure BFD support for L3 interfaces.

bfd bind peer-ip *ipv4-address* [**source-ip** *ipv4-address*] **process-pst**

BFD support for interfaces is not configured by default.

1.4.7 Configuring BFD Support for Applications

1. Overview

- This function is disabled by default.
- The configuration command varies with the associated applications. For details, see their configuration guides.
- This function must be configured at both ends so that a BFD session can be established.

2. Procedure

(1) Enter the privileged EXEC mode.

enable

(2) Enter the global configuration mode.

configure terminal

(3) Configure BFD support.

bfd all interfaces

In RIP routing configuration mode, enable BFD support for RIP on all interfaces. For details, see "Configuring RIP."

bfd all interfaces

In OSPF routing configuration mode, enable BFD support for OSPF on all interfaces. For details, see "Configuring OSPFv2."

bfd all interfaces

In OSPFv3 routing configuration mode, enable BFD support for OSPFv3 on all interfaces. For details, see "Configuring OSPFv3."

neighbor *address* **fall-over bfd**

In BGP routing configuration mode, enable BFD support for BGP. For details, see "Configuring BGP."

bfd all interfaces

In IS-IS routing configuration mode, enable BFD support for IS-IS on all interfaces. For details, see "Configuring IS-IS."

Ip route static bfd [**vrf** *vrf-name*] *interface-type interface-number gateway* [**source** *ip-address*]

In global configuration mode, enable BFD support for static routing. For details, see "Configuring Static Routing."

ipv6 route static bfd [**vrf** *vrf-name*] *interface-type interface-number gateway* [**source** *ipv6-address*]

In global configuration mode, enable BFD support for static IPv6 routing. For details, see "Configuring Static Routing."

set ipnext-hopverify-availability next-hop-address **bfd** [**vrf** *vrf-name*] interface-type interface-number gateway

Enable BFD support for PBR. For details, see "Configuring PBR."

set ipv6 next-hopverify-availability *next-hop-address* **bfd** [**vrf** *vrf-name*] *interface-type interface-number gateway*

Enable BFD support for IPv6 PBR. For details, see "Configuring PBR."

vrrp bfd *interface-type interface-number ip-address*

Enable BFD support for VRRP. For details, see "Configuring VRRP."

VRRP Plus is based on the VRRP protocol. Therefore, no additional configuration is required for BFD support for VRRP Plus and only VRRP needs to be enabled on devices at both ends and a BFD session is correctly associated.

bfd bind static-lsp peer-ip *ip-address* **source-ip** *ip-address* [**local-discriminator** *discr-value* **remote-discriminator** *discr-value*] [**process-state**]

Enable BFD support for static LSPs. For details, see "MPLS Basics."

bfd bind ldp-lsppeer-ip *ip-address* **nexthop** *ip-address* [**interface** *interface-type interface-number*] **source-ip** *ip-address* [**local-discriminator** *discr-value* **remote-discriminator** *discr-value*] [**process-state**]

Enable BFD support for LDP LSPs. For details, see "MPLS Basics."

bfd bind backward-lsp-with-ippeer-ip *ip-address* [**vrf** *vrf-name*] **interface** *interface-type interface-number* [**source-ip** *ip-address*] { **local-discriminator** *discr-value* **remote-discriminator** *discr-value*}

Enable BFD support for dynamic LSPs. For details, see *MPLS Basics*.

1.5 Configuring BFD Protection

1.5.1 Overview

- If a BFD-enabled device is attacked (for example, attacked by a large number of ping packets), BFD sessions flap. In this case, you can configure BFD protection to provide protection.

1.5.2 Restrictions and Guidelines

- The BFD basic functions must be configured.
- If both BFD and BFD protection are enabled on a device, the device discards the BFD packets received from the previous-hop device, which affects the establishment of a BFD session between the previous-hop device and other devices.

1.5.3 Prerequisites

Basic functions of BFD have been configured.

1.5.4 Procedure

(1) Enter the privileged EXEC mode.

enable

(2) Enter the global configuration mode.

configure terminal

(3) Configure BFD protection.

bfd cpp

Configure BFD protection on a device in the network where attacks exist.

The BFD protection is enabled by default.

1.6 Configuring BFD Flapping Advertisement Dampening

1.6.1 Overview

- A BFD session may frequently switch between the down and up states due to link instability. As a result, a relevant application (such as static routing) may frequently switch forwarding paths and the running services are affected.
- Users can set a delay for status change advertisement, after which BFD informs an associated application of BFD up. After a BFD session is up for a certain period of time, BFD informs an associated application of BFD up. Otherwise, BFD informs it of BFD down. This aims to reduce flapping of associated protocols caused by unstable links.

1.6.2 Restrictions and Guidelines

- The BFD basic functions must be configured.
- If a BFD session does not frequently switch between the down and up states, enabling BFD flapping advertisement dampening will delay notifying an associated application of BFD up.
- Enabling BFD flapping advertisement dampening can relieve frequent processing of associated applications

(such as route re-calculation) caused by frequent advertisement of BFD status changes. The larger the configured time is, the longer the required BFD stability time is. BFD notifies an application module of BFD up only after the stability time reaches the configured time.

1.6.3 Procedure

(1) Enter the privileged EXEC mode.

enable

(2) Enter the global configuration mode.

configure terminal

(3) Enter the interface configuration mode.

interface *interface-type interface-number*

(4) Configure flapping dampening parameters.

bfd up-dampening { *milliseconds* }

The BFD flapping advertisement dampening is disabled on ports by default. If a BFD session frequently switches between the down state and up state, you are advised to enable this function.

1.7 Configuring an SBFD Reflector

1.7.1 Overview

This section describes how to configure an SBFD reflector discriminator.

1.7.2 Restrictions and Guidelines

In a network, one initiator can be deployed to map to multiple reflectors, and multiple reflector discriminators can be configured for one reflector. All the reflector discriminators in the network must be globally unique.

1.7.3 Procedure

(1) Enter the privileged EXEC mode.

enable

(2) Enter the global configuration mode.

configure terminal

(3) Configure an SBFD reflector.

sbfd reflector discriminator { *unsigned-integer-value* }

(4) Configure an SBFD reflector discriminator.

sbfd reflector discriminator

The reflector discriminator is Your Discriminator for the initiator.

1.8 Monitoring

Run the **show** command to check the running status of a configured function to verify the configuration effect.

Caution

- The output debugging information occupies system resources. Therefore, disable the debugging function immediately after use.

Table 1-1 BFD Monitoring

Command	Purpose
show bfd neighbors [**vrf** *vrf-name*] [**client** { **ap** | **bgp** | **ospf** | **rip** | **vrrp** | **static-route** | **pbr** | **vrrp-balance** | **ldp-lsp** | **static-lsp** | **backward-lsp-with-ip** | **pst**| **dhcp** | **openflow** | **pimdm** | **pimsm** | **pimsmv6** | **srp** }] [**ipv4** *ipv4-address* | **ipv6** *ipv6-address*] [**details**] [**parm-consult** [*interface-type interface-number*]]	Displays BFD session information.
debug bfd event [**interface** *interface-type interface-number* **| ipv4** *ipv4-address* **| ipv6** *ipv6-address*]	Debugs BFD events.
debug bfd packet [**interface** *interface-type interface-number* **| ipv4** *ipv4-address* **| ipv6** *ipv6-address*]	Debugs BFD packets.

1.9 Configuration Examples

1.9.1 Configuring BFD Support for OSPF

1. Requirements

Three devices are used to build a network, the devices run the OSPF protocol, and BFD support for OSPF is enabled on two devices.

2. Topology

Figure 1-1 Topology for Configuring BFD Support for OSPF

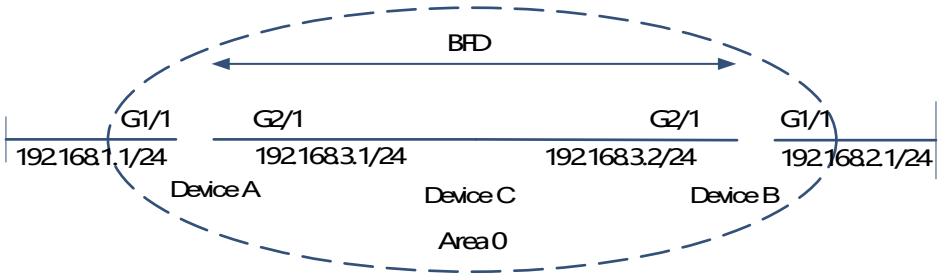

3. Notes

- Configure IP addresses for ports connecting devices A and B.
- Run the OSPF protocol on devices A and B.
- Configure BFD parameters for the interconnected ports of devices A and B.
- Enable BFD support for OSPF on devices A and B.

4. Procedure

Perform the following configuration on device A.

```
DeviceA> enable
DeviceA# configure terminal
DeviceA(config)# interface GigabitEthernet2/1
DeviceA(config-if-GigabitEthernet2/1)# no switchport
DeviceA(config-if-GigabitEthernet2/1)# ip address 192.168.3.1 255.255.255.0
DeviceA(config-if-GigabitEthernet2/1)# bfd interval 200 min_rx 200 multiplier 5
DeviceA(config-if-GigabitEthernet2/1)# exit
DeviceA(config)# interface GigabitEthernet1/1
DeviceA(config-if-GigabitEthernet1/1)# no switchport
DeviceA(config-if-GigabitEthernet1/1)# ip address 192.168.1.1 255.255.255.0
DeviceA(config-if-GigabitEthernet1/1)# exit
DeviceA(config)# router ospf 123
DeviceA(config-router)# log-adj-changes detail
DeviceA(config-router)# network 192.168.3.0 0.0.0.255 area 0
DeviceA(config-router)# network 192.168.1.0 0.0.0.255 area 0
DeviceA(config-router)# bfd all-interfaces
```

Perform the following configuration on device B.

```
DeviceB# configure terminal
DeviceB(config)# interface GigabitEthernet2/1
DeviceB(config-if-GigabitEthernet2/1)# no switchport
DeviceB(config-if-GigabitEthernet2/1)# ip address 192.168.3.2 255.255.255.0
DeviceB(config-if-GigabitEthernet2/1)# bfd interval 200 min_rx 200 multiplier 5
DeviceB(config-if-GigabitEthernet2/1)# exit
DeviceB(config)# interface GigabitEthernet1/1
DeviceB(config-if-GigabitEthernet1/1)# no switchport
DeviceB(config-if-GigabitEthernet1/1)# ip address 192.168.2.1 255.255.255.0
DeviceB(config-if-GigabitEthernet1/1)# exit
DeviceB(config)# router ospf 123
DeviceB(config-router)# log-adj-changes detail
DeviceB(config-router)# network 192.168.3.0 0.0.0.255 area 0
DeviceB(config-router)# network 192.168.2.0 0.0.0.255 area 0
DeviceB(config-router)# bfd all-interfaces
```

5. Verification

Run the **show bfd neighbors details** command to check whether BFD support configuration is correct.

```
DeviceA# show bfd neighbors details
OurAddr       NeighAddr   LD/RD   RH/RS   Holdown(mult)   State   Int
192.168.3.1   192.168.3.2     1/2         Up      532 (3 )           Up      Ge2/1
Local Diag: 0, Demand mode: 0, Poll bit: 0
MinTxInt: 200000, MinRxInt: 200000, Multiplier: 5
Received MinRxInt: 50000, Received Multiplier: 3
Holdown (hits): 600(22), Hello (hits): 200(84453)
```

```
Rx Count: 49824, Rx Interval (ms) min/max/avg: 208/440/332
Tx Count: 84488, Tx Interval (ms) min/max/avg: 152/248/196
Registered protocols: OSPF
Uptime: 02:18:49
Last packet: Version: 1  - Diagnostic: 0
I Hear You bit: 1            - Demand bit: 0
Poll bit: 0              - Final bit: 0
Multiplier: 3            - Length: 24
My Discr.: 2             - Your Discr.: 1
Min tx interval: 50000       - Min rx interval: 50000
Min Echo interval: 0
```

```
DeviceB# show bfd neighbors details
OurAddr       NeighAddr   LD/RD   RH/RS   Holdown(mult)   State   Int
192.168.3.2   192.168.3.1 2/1     Up      532 (5 )            Up      Ge2/1
Local Diag: 0, Demand mode: 0, Poll bit: 0
MinTxInt: 50000, MinRxInt: 50000, Multiplier: 3
Received MinRxInt: 200000, Received Multiplier: 5
Holdown (hits): 600(22), Hello (hits): 200(84453)
Rx Count: 49824, Rx Interval (ms) min/max/avg: 209/440/332 last: 66 ms ago
Tx Count: 84488, Tx Interval (ms) min/max/avg: 153/249/197 last: 190 ms ago
Registered protocols: OSPF
Uptime: 02:18:49
Last packet: Version: 1  - Diagnostic: 0
I Hear You bit: 1            - Demand bit: 0
Poll bit: 0              - Final bit: 0
Multiplier: 5        - Length: 24
My Discr.: 1             - Your Discr.: 2
Min tx interval: 200000  - Min rx interval: 200000
Min Echo interval: 0
```

6. Common Errors

- For a single-hop session, BFD parameters may not be configured for the port on the device at one end of the session.
- If different negotiation parameters need to be configured for different multi-hop sessions, templates need to be associated with the sessions and negotiation parameters need to be configured in the templates.
- For multi-hop BFD sessions, if the sessions use different templates, negotiation parameters in the template used by a session that becomes up first are used for transmitting and receiving packets.
- BFD support for applications is not configured.
- BFD support for applications is enabled only on the device at one end.