

Bidirectional Forwarding Detection Routing

This chapter describes how to configure the ASA to use the Bidirectional Forwarding Detection (BFD) routing protocol.

- About BFD Routing, on page 1
- Guidelines for BFD Routing, on page 4
- Configure BFD, on page 5
- Monitoring for BFD, on page 9
- History for BFD Routing, on page 10

About BFD Routing

BFD is a detection protocol designed to provide fast forwarding path failure detection times for all media types, encapsulations, topologies, and routing protocols. BFD operates in a unicast, point-to-point mode on top of any data protocol being forwarded between two systems. Packets are carried in the payload of the encapsulating protocol appropriate for the media and the network.

BFD provides a consistent failure detection method for network administrators in addition to fast forwarding path failure detection. Because the network administrator can use BFD to detect forwarding path failures at a uniform rate, rather than the variable rates for different routing protocol hello mechanisms, network profiling and planning are easier and reconvergence time is consistent and predictable.

BFD Asynchronous Mode and Echo Function

BFD can operate in asynchronous mode with or without the echo function enabled.

Asynchronous Mode

In asynchronous mode, the systems periodically send BFD control packets to one another, and if a number of those packets in a row are not received by the other system, the session is declared to be down. Pure asynchronous mode (without the Echo function) is useful because it requires half as many packets to achieve a particular detection time as the Echo function requires.

BFD Echo Function

The BFD echo function sends echo packets from the forwarding engine to the directly-connected single-hop BFD neighbor. The echo packets are sent by the forwarding engine and forwarded back along the same path to perform detection. The BFD session at the other end does not participate in the actual forwarding of the echo packets. Because the echo function and the forwarding engine are responsible for the detection

process, the number of BFD control packets that are sent out between BFD neighbors is reduced. And also because the forwarding engine is testing the forwarding path on the remote neighbor system without involving the remote system, the inter-packet delay variance is improved. This results in quicker failure detection times.

When the echo function is enabled, BFD can use the slow timer to slow down the asynchronous session and reduce the number of BFD control packets that are sent between BFD neighbors, which reduces processing overhead while at the same time delivering faster failure detection.

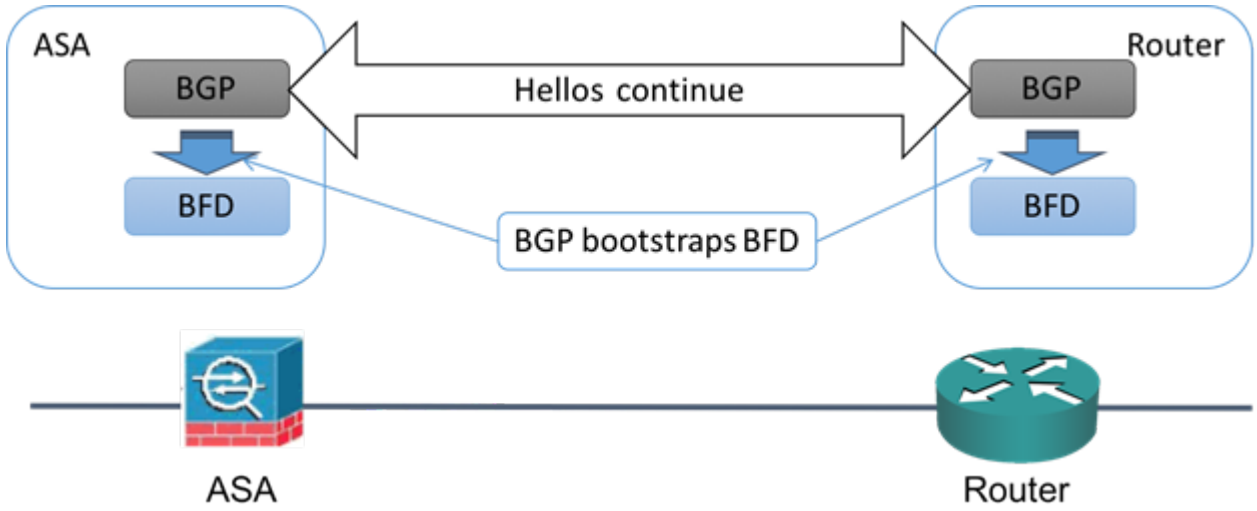
Note The echo function is not supported for IPv4 multi-hop or IPv6 single-hop BFD neighbors.

You can enable BFD at the interface and routing protocol levels. You must configure BFD on both systems (BFD peers). After you enable BFD on the interfaces and at the router level for the appropriate routing protocols, a BFD session is created, BFD timers are negotiated, and the BFD peers begin to send BFD control packets to each other at the negotiated level.

BFD Session Establishment

The following example shows the ASA and a neighboring router running Border Gateway Protocol (BGP). At the time when both devices come up, there is no BFD session established between them.

Figure 1: Established BFD Session



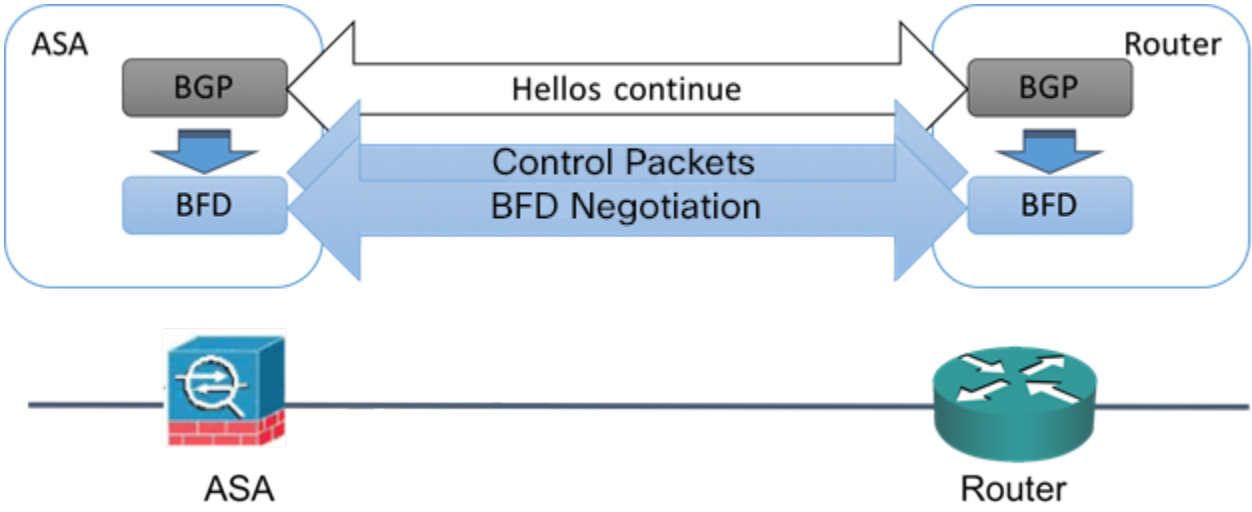
After BGP identifies its BGP neighbor, it bootstraps the BFD process with the IP address of the neighbor. BFD does not discover its peers dynamically. It relies on the configured routing protocols to tell it which IP addresses to use and which peer relationships to form.

The BFD on the router and the BFD on the ASA form a BFD control packet and start sending the packets to each other at a one-second interval until the BFD session is established. The initial control packets from either system are very similar, for example, the Vers, Diag, H, D, P, and F bits are all set to zero, and the State is set to Down. The My Discriminator field is set to a value that is unique on the transmitting device. The Your Discriminator field is set to zero because the BFD session has not yet been established. The TX and RX timers are set to the values found in the configuration of the device.

After the remote BFD device receives a BFD control packet during the session initiation phase, it copies the value of the My Discriminator field into its own Your Discriminator field and the transition from Down state to Init state and then eventually to Up state occurs. Once both systems see their own Discriminators in each other's control packets, the session is officially established.

The following illustration shows the established BFD connection.

Figure 2: BGP With No BFD Session Established



BFD Timer Negotiation

BFD devices must negotiate the BFD timers to control and synchronize the send rate of BFD control packets. A device needs to ensure the following before it can negotiate a BFD timer:

- That its peer device saw the packet containing the proposed timers of the local device
- That it never sends BFD control packets faster than the peer is configured to receive them
- That the peer never sends BFD control packets faster than the local system is configured to receive them

The setting of the Your Discriminator field and the H bit are sufficient to let the local device that the remote device has seen its packets during the initial timer exchange. After receiving a BFD control packet, each system takes the Required Min RX Interval and compares it to its own Desired Min TX Interval, and then takes the greater (slower) of the two values and uses it as the transmission rate for its BFD packets. The slower of the two systems determines the transmission rate.

When these timers have been negotiated, they can be renegotiated at any time during the session without causing a session reset. The device that changes its timers sets the P bit on all subsequent BFD control packets until it receives a BFD control packet with the F bit set from the remote system. This exchange of bits guards against packets that might otherwise be lost in transit.

Note The setting of the F bit by the remote system does not mean that it accepts the newly proposed timers. It indicates that the remote system has seen the packets in which the timers were changed.

BFD Failure Detection

When the BFD session and timers have been negotiated, the BFD peers send BFD control packets to each other at the negotiated interval. These control packets act as a heartbeat that is very similar to IGP Hello protocol except that the rate is more accelerated.

As long as each BFD peer receives a BFD control packet within the configured detection interval (Required Minimum RX Interval), the BFD session stays up and any routing protocol associated with BFD maintains its adjacencies. If a BFD peer does not receive a control packet within this interval, it informs any clients participating in that BFD session about the failure. The routing protocol determines the appropriate response to that information. The typical response is to terminate the routing protocol peering session and reconverge and thus bypass a failed peer.

Each time a BFD peer successfully receives a BFD control packet in a BFD session, the detection timer for that session is reset to zero. Thus the failure detection is dependent on received packets and NOT when the receiver last transmitted a packet.

BFD Deployment Scenarios

The following describes how BFD operates in these specific scenarios.

Failover

In a failover scenario, BFD sessions are established and maintained between the active unit and the neighbor unit. Standby units do not maintain any BFD sessions with the neighbors. When a failover happens, the new active unit must initiate session establishment with the neighbor because session information is not synched between active and standby units.

For a graceful restart/NSF scenario, the client (BGP IPv4/IPv6) is responsible for notifying its neighbor about the event. When the neighbor receives the information, it keeps the RIB table until failover is complete. During failover, the BFD and the BGP sessions go down on the device. When the failover is complete, a new BFD session between the neighbors is established when the BGP session comes up.

Spanned EtherChannel and L2 Cluster

In a Spanned EtherChannel cluster scenario, the BFD session is established and maintained between the primary unit and its neighbor. Subordinate units do not maintain any BFD sessions with the neighbors. If a BFD packet is routed to the subordinate unit because of load balancing on the switch, the subordinate unit must forward this packet to the primary unit through the cluster link. When a cluster switchover happens, the new primary unit must initiate session establishment with the neighbor because session information is not synched between primary and subordinate units.

Individual Interface Mode and L3 Cluster

In an individual interface mode cluster scenario, individual units maintain their BFD sessions with their neighbors.

Guidelines for BFD Routing

Context Mode Guidelines

Supported in single and multiple context modes.

Firewall Mode Guidelines

Supported in routed firewall mode; support for standalone, failover, and cluster modes. BFD is not supported on failover and cluster interfaces. In clustering this feature is only supported on the primary unit. BFD is not supported in transparent mode.

IPv6 Guidelines

Echo mode is not supported for IPv6.

Additional Guidelines

BGP IPv4 and BGP IPv6 protocol are supported.

OSPFv2, OSPFv3, IS-IS, and EiGRP protocols are not supported.

BFD for Static Routes is not supported.

BFD on Transfer and Tunnel is not supported.

Configure BFD

This section describes how to enable and configure the BFD routing process on your system.

Procedure

Step 1 Create the BFD Template, on page 5.
Step 2 Configure BFD Interfaces, on page 7.
Step 3 Configure BFD Maps, on page 8.

Create the BFD Template

This section describes the steps required to create a BFD template and enter BFD configuration mode.

The BFD template specifies a set of BFD interval values. BFD interval values as configured in the BFD template are not specific to a single interface. You can also configure authentication for single-hop and multi-hop sessions. You can enable Echo on single-hop only.

Procedure

Step 1 Enable BFD as a routing protocol on the ASA by creating the BFD template, either single-hop or multi-hop:

bfd-template [single-hop | multi-hop] *template_name*

Example:

```
ciscoasa(config)# bfd-template single-hop TEMPLATE1
ciscoasa(config-bfd)#
```

- **single-hop**—Specifies a single-hop BFD template.
- **multi-hop**—Specifies a multi-hop BFD template.
- *template-name*—Specifies the template name. The template name cannot contains spaces.

The **bfd-template** command lets you create the BFD template and enter BFD configuration mode.

Step 2 (Optional) Configure Echo on a single-hop BFD template:

bfd-template single-hop *template_name*

Example:

```
ciscoasa(config)# bfd-template single-hop TEMPLATE1
ciscoasa (config-bfd)# echo
```

You can only enable Echo mode on a single-hop template. BFD echo is not supported for IPv6 BFD sessions.

Step 3 Configure the intervals in the BFD template:

interval [**both** *milliseconds* | **microseconds {both | min-tx}** *microseconds* | **min-tx** *milliseconds*

Example:

```
ciscoasa(config)# bfd-template single-hop TEMPLATE1
ciscoasa(config-bfd)# interval both 50
```

- **both**—Minimum transmit and receive interval capability.
- *milliseconds*—The interval in milliseconds. The range is 50 to 999.
- **microseconds**—Specifies the BFD interval in microseconds for **both** and **min-tx**.
- *microseconds*—The range is 50,000 to 999,000.
- **min-tx**—The minimum transmit interval capability.

BFD interval values specified as part of the BFD template are not specific to a single interface. You can apply individual BFD templates per interface. See Configure BFD Interfaces, on page 7.

Step 4 Configure authentication in the BFD template:

authentication {md5 | meticulous-mds | meticulous-sha-1 | sha-1} [0|8] *word* **key-id** *id*

Example:

```
ciscoasa(config)# bfd-template single-hop TEMPLATE1
ciscoasa(config-bfd)# authentication sha-1 0 cisco key-id 10
```

- **authentication**—Specifies the authentication type.
- **md5**—Message Digest 5 (MD5) authentication.
- **meticulous-md5**—Meticulous keyed MD5 authentication.
- **meticulous-sha-1**—Meticulous keyed SHA-1 authentication.
- **sha-1**—Keyed SHA-1 authentication.

- **0|8**—0 specifies that an UNENCRYPTED password will follow. 8 specifies that an ENCRYPTED password will follow.
- *word*—The BFD password (key), which is a single-digit password/key of up to 29 characters. Passwords starting with a digit followed by a whitespace are not supported, for example, '0 pass' and '1' are not valid.
- **key-id**—The authentication Key ID.
- *id*—The shared key ID that matches the key string. The range is 0 to 255 characters.

You can configure authentication in single-hop and multi-hop templates. We recommend that you configure authentication to enhance security. You must configure authentication on each BFD source-destination pair and the authentication parameters must match on both devices.

Configure BFD Interfaces

You can bind a BFD template to an interface, configure the baseline BFD session parameters per interface, and enable echo mode per interface.

Procedure

Step 1 Enter interface configuration mode:

interface *interface_id*

Example:

```
ciscoasa(config)# interface GigabitEthernet0/0
ciscoasa(config-if)#
```

Step 2 Apply a BFD template to an interface:

bfd template *template-name*

Example:

```
ciscoasa(config)# interface GigabitEthernet0/0
ciscoasa(config-if)# bfd template TEMPLATE1
```

Even if you have not created the template using the **bfd-template** command, you can configure the name of the template under an interface, but the template is considered invalid until you define the template. You do not have to reconfigure the template name again. It becomes valid automatically.

Step 3 Configure the BFD session parameters:

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

Example:

```
ciscoasa(config)# interface GigabitEthernet0/0
ciscoasa(config-router)# bfd interval 200 min_rx 200 multiplier 3
```

- **interval** *milliseconds*—Specifies the rate at which BFD control packets are sent to BFD peers. The range is 50 to 999 milliseconds.
- **min_rx** *milliseconds*—Specifies the rate at which BFD control packets are expected to be received from BFD peers. The range is 50 to 999 milliseconds.
- **multiplier** *multiplier-value*—Specifies the number of consecutive BFD control packets that must be missed from a BFD peer before BFD declares that the peer is unavailable and the Layer 3 BFD peer is informed of the failure. The range is 3 to 50.

Step 4 Enable BFD echo mode on an interface:

bfd echo

Example:

```
ciscoasa(config)# interface gigabitethernet 0/0
ciscoasa(if)# bfd echo
```

Echo mode is enabled by default but not supported in BFD IPv6 sessions. When echo mode is enabled, the minimum echo transmit level and required minimum transmit interval values are taken from the **bfd interval** *milliseconds* **min_rx** *milliseconds* configuration.

Note
Before using BFD echo mode, you must disable ICMP redirect messages using the **no ip redirects** command. This avoids high CPU use.

Configure BFD Maps

You can create a BFD map containing destinations that you can associate with a multi-hop template. You must have a multi-hop BFD template already configured.

Procedure

Step 1 Create a multi-hop BFD template. See Create the BFD Template, on page 5 for the procedure.

Step 2 Associate the BFD multi-hop template with a map of destinations:

bfd map {ipv4 | ipv6} *destination/cdir source/cdire template-name*

Example:

```
ciscoasa(config)# bfd map ipv4 10.11.11.0/24 10.36.42.5/32 MULTI-TEMPLATE1
ciscoasa(config-bfd)#
```

- **ipv4** —Configures an IPv4 address.
- **ipv6** —Configures an IPv6 address.
- *destination/cdir*—Specifies the destination prefix/length. The format is A.B.C.D/<0-32>.

- *source/cdir*—Specifies the destination prefix/length. The format is X:X:X;X::X/<0-128>.
- *template-name*—Specifies the name of the multi-hop template associated with this BFD map.

Step 3 (Optional) Configure the BFD slow timers value:

bfd slow-timers [*milliseconds*]

Example:

```
ciscoasa(config)# bfd slow-timers 14000
ciscoasa(config-bfd)#
```

milliseconds—(Optional) The BFD slow timers value. The range is 1000 to 30000. The default is 1000.

Monitoring for BFD

You can use the following commands to monitor the BFD routing process. For examples and descriptions of the command output, see the command reference.

To monitor or disable various BFD routing statistics, enter one of the following commands:

- **show bfd neighbors**

  Displays a line-by-line listing of existing BFD adjacencies.

- **show bfd summary**

  Displays summary information for BFD, BFD clients, or BFD sessions.

- **show bfd drops**

  Displays the number of dropped packets in BFD.

- **show bfd map**

  Displays the configured BFD maps.

- **show running-config bfd**

  Displays BFD map and other BFD related configurations.

- **show running-config bfd-template**

  Displays BFD-template related configurations.

History for BFD Routing

Table 1: Feature History for BFD Routing

Feature Name	Platform Releases	Feature Information
BFD routing support	9.6(2)	The ASA now supports the BFD routing protocol. Support was added for configuring BFD templates, interfaces, and maps. Support for BGP routing protocol to use BFD was also added. We added the following commands: **bfd echo, bfd interval, bfd map, bfd slow-timers, bfd-template, clear bfd counters, clear conf bfd, neighbor fall-over bfd, show bfd drops, show bfd map, show bfd neighbors, show bfd summary, show running-config bfd**