



Bidirectional Forwarding Detection

This chapter includes details for Bidirectional Forwarding Detection (BFD).

BFD is a detection protocol designed to provide fast forwarding path failure detection times for all media types, encapsulations, topologies, and routing protocols. In addition to fast forwarding path failure detection, BFD provides a consistent failure detection method for network administrators.

- [Bidirectional Forwarding Detection, on page 1](#)
- [Implementing BFD, on page 3](#)

Bidirectional Forwarding Detection

Bidirectional forwarding detection (BFD) provides low-overhead, short-duration detection of failures in the path between adjacent forwarding engines. BFD allows a single mechanism to be used for failure detection over any media and at any protocol layer, with a wide range of detection times and overhead. The fast detection of failures provides immediate reaction to failure in the event of a failed link or neighbor.

Prerequisites for Implementing BFD

You must be in a user group associated with a task group that includes the proper task IDs. The command reference guides include the task IDs required for each command. If you suspect user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

The following prerequisites are required to implement BFD:

- Interior Gateway Protocol (IGP) is activated on the router if you are using IS-IS or OSPF.
- To enable BFD for a neighbor, the neighbor router must support BFD.

Restrictions for Implementing BFD

These restrictions apply to BFD:

- Demand mode is not supported.
- Asynchronous echo mode is not supported.
- Mutli hop BFD is not supported.

- BFD for bundles is not supported.

Operating Modes for BFD

BFD can operate in two modes, Asynchronous mode and Demand mode. Cisco NCS 4000 supports the asynchronous mode only. In this mode, the systems periodically send BFD control packets to one another. If a number of those packets in a row, are not received by the other system, the session is declared to be down.

When BFD is running asynchronously, the following happens:

- Each system periodically sends BFD control packets to one another. Packets sent by BFD router “Peer A” to BFD router “Peer B” have a source address from Peer A and a destination address for Peer B.
- Control packet streams are independent of each other and do not work in a request/response model.
- If a number of packets in a row are not received by the other system, the session is declared down.

Figure 1: BFD Asynchronous Mode



Control packet failure in asynchronous mode (without echo), is detected using the values of the minimum interval (**bfd minimum-interval**) and multiplier (**bfd multiplier**) commands. For control packet failure detection, the local multiplier value is sent to the neighbor. A failure detection timer is started based on $(I \times M)$, where I is the negotiated interval, and M is the multiplier provided by the remote end. Whenever a valid control packet is received from the neighbor, the failure detection timer is reset. If a valid control packet is not received from the neighbor within the time period $(I \times M)$, then the failure detection timer is triggered, and the neighbor is declared down.

Table 1: BFD Packet Intervals

Configured Async Control Packet Interval (ms)	Multiplier value (the default is 3; range is from 2 to 50)	Async Control Packet Failure Detection Time (ms) (Interval X Multiplier) The multiplier value is set to the default value of 3
3.3 (rounded off to 3)	3	9
10	3	30
20	3	60
50	3	150
100	3	300
1000	3	3000
2000 (this value is the default value)	3	6000

BFD for IPv4

Cisco NCS 4000 supports single hop BFD for IPv4.

BFD asynchronous packets are transmitted over UDP and IPv4 using source port 49152 and destination port 3784. For asynchronous mode, the source address of the IP packet is the local interface address, and the destination address is the remote interface address.

BFD is supported for connections over the following interface types:

- Gigabit Ethernet (GigE)
- Ten Gigabit Ethernet (10GigE)
- Hundred Gigabit Ethernet (100GigE)

BFD Dampening

Bidirectional Forwarding Detection (BFD) is a mechanism used by routing protocols to quickly realize and communicate the reachability failures to their neighbors. When BFD detects a reachability status change of a client, its neighbors are notified immediately. Sometimes it might be critical to minimize changes in routing tables so as not to impact convergence, in case of a micro failure. An unstable link that flaps excessively can cause other devices in the network to consume substantial processing resources, and that can cause routing protocols to lose synchronization with the state of the flapping link.

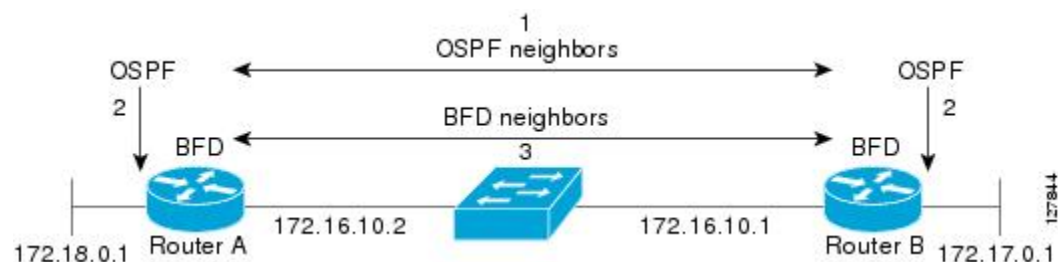
The BFD dampening feature introduces a configurable exponential delay mechanism. This mechanism is designed to suppress the excessive effect of remote node reachability events flapping with BFD. The BFD Dampening feature allows the network operator to automatically dampen a given BFD session to prevent excessive notification to BFD clients, thus preventing unnecessary instability in the network. Dampening the notification to a BFD client suppresses BFD notification until the time the session under monitoring stops flapping and becomes stable.

Implementing BFD

By sending rapid failure detection notices to the routing protocols in the local router to initiate the routing table recalculation process, BFD contributes to greatly reduced overall network convergence time.

The figure below, shows a simple network with two routers running OSPF and BFD. When OSPF discovers a neighbor (1) it sends a request to the local BFD process to initiate a BFD neighbor session with the OSPF neighbor router (2). The BFD neighbor session with the OSPF neighbor router is established (3).

Figure 2: BFD process - establishing a connection



The figure below, shows what happens when a failure occurs on the network (1). The BFD neighbor session with the OSPF neighbor router is not reachable (2). BFD notifies the local OSPF process that the BFD neighbor is no longer reachable (3). The local OSPF process deletes the OSPF neighbor relationship (4). If an alternative path is available the routers will immediately start converging on it.

Figure 3: BFD process - failure detection

