

Bi-directional Forwarding Detection (BFD) for Routing Protocols

Feature Overview and Configuration Guide

Introduction

This guide describes Bi-direction Forwarding Detection (BFD) and how to configure it. BFD is a standards-based protocol initially defined in RFC 5880, whose sole purpose is to detect communication failure between two devices quickly and efficiently. Allied Telesis' BFD implementation supports **Layer 3** protocols and is based on the RFC.

BFD provides a consistent failure detection method for network administrators, in addition to fast forwarding path failure detection. The network administrator can use BFD to detect forwarding path failures at a uniform rate, rather than the variable rates seen with different routing protocol hello mechanisms.

Allied Telesis' implementation of BFD is inter-operable with implementations of RFC 5880 from other manufacturers. BFD utilizes the Border Gateway Protocol (BGP) feature. A BGP neighbor can be configured so that when a neighbor becomes active, a BFD peer session is also created. When BFD detects a link failure the BGP neighbor transitions down much more quickly than BGP itself can detect, improving route convergence times.

The main benefit of implementing BFD for BGP is a significantly faster re-convergence time.

Note: For information on BFD using Ethernet and **Layer 2** redundancy protocols, see the [BFD over Ethernet for Layer 2 Redundancy Protocols Feature Overview and Configuration Guide](#).

Contents

Introduction	1
Products and software version that apply to this guide	2
This is how BFD works.....	3
Limitations	3
How to configure BFD.....	4
Configuration example	10
Monitoring BFD	11

Products and software version that apply to this guide

This guide applies to the x930 Series, running version **5.5.0-2.1** or later.

From version **5.5.2-0.1** onwards, it also applies to the x950, SBx8100, and SBx908 GEN2 Series.

From version **5.5.2-1.1** onwards, it also applies to the x530 Series.

From version **5.5.2-0.1** onwards, the following new features have also been added:

- BFD supports profiles, defined via CLI or API, that are a collection of BFD session parameters. BFD profiles can be associated with BFD peers, OSPF peers, and BGP neighbors.
- BFD can monitor connections to OSPF peers and notify OSPF about connection events.
- BFD can monitor Static Route peers and notify NSM about connection events.

From version **5.5.2-1.1** onwards, the following new features have also been added:

- BFD supports applying profiles to OSPF multihop and IPv4 static routes.
- BFD supports VRF for OSPF, BGP, and IPv4 static routes.

For more information, see the following documents:

- The product's [Datasheet](#)
- The product's [Command Reference](#)

These documents are available from the above links on our website at alliedtelesis.com.

This is how BFD works

BFD is a simple Hello protocol that in many respects is similar to the detection components of well-known routing protocols. A pair of devices transmit BFD packets periodically over the path between them. If one of the devices stops receiving BFD packets for long enough, it assumes that some component in that particular path to its neighbor has failed.

In brief the AlliedWare Plus implementation of BFD has the following handshake process:

1. Devices are initialized to take an active role, meaning they will not wait to receive a Control packet before starting to transmit them.
2. The local device starts to transmit Control packets at a slow rate with the state field set to Down.
 - If the remote device responds with a Control packet with the state field set to:
 - **Down**, then the local device will move to the Initialize state.
 - **Initialize**, then the local device will move to the **Up** state.
3. The local device will then transmit a Control packet with its new state set, which should cause the remote device to change to the Initialize or Up state.
 - When a system is in the Initialize state it needs to receive a control packet with the state field set to either Initialize or Up in order to transition to an Up state.
4. Once both systems are in the Up state, the BFD peer session has been successfully established.

Once the BFD session is running, the transmission rate of the Control packets will be increased to achieve a desired detect time.

- If no Control packets are received in the detect time, then the session will be declared Down. The local device will notify the remote device of this by sending a Control packet with the state field set to Down.
- The peer session cannot attempt to come back up until both devices are in a Down state. At this point they will start the above handshake again.

Limitations

The AlliedWare Plus implementation of BFD does not support the following:

- The authentication feature
- The demand-mode feature
- RFC7130 BFD support for individual aggregator links
- Hardware based BFD

If VCS fast failover occurs, this may cause BFD on the peer device to set the session to Down, as it will not be receiving control or echo packets for a short amount of time. This could cause the peer device to change its view of the network topology. For BFD peers using VCS, you can set longer detect-multiplier, receive, and transmit interval times to avoid a BFD peer down event due to failover.

How to configure BFD

Here are the high-level steps to configure a singlehop and multihop BFD peer between Device-1 and Device-2:

Configuring BFD on Device-1

Step 1: Configure Device-1 and enable BFD

```
awplus#hostname Device-1
awplus#configure terminal
awplus(config)#service bfd
```

Step 2: Configure the interfaces

```
awplus(config)#interface vlan10
awplus(config-if)#ip address 10.0.0.1/24
awplus(config-if)#exit
awplus(config)#interface vlan20
awplus(config-if)#ip address 192.168.1.17/24
awplus(config-if)#exit
```

Step 3: Create a singlehop BFD peer

```
awplus(config)# bfd peer 10.0.0.2 interface vlan10
```

Step 4: Create a multihop BFD peer

```
awplus(config)# bfd peer 192.168.1.18 multihop local-address 192.168.1.17
```

Configuring BFD on Device-2

Step 1: Configure Device-2 and enable BFD

```
awplus#hostname Device-2
awplus#configure terminal
awplus(config)#service bfd
```

Step 2: Configure the interfaces

```
awplus(config)#interface vlan10
awplus(config-if)#ip address 10.0.0.2/24
awplus(config-if)#exit
awplus(config)#interface vlan20
awplus(config-if)#ip address 192.168.1.18/24
awplus(config-if)#exit
```

Step 3: Create a singlehop BFD peer

```
awplus(config)# bfd peer 10.0.0.1 interface vlan10
```

Step 4: Create a multihop BFD peer

```
awplus(config)# bfd peer 192.168.1.17 multihop local-address 192.168.1.18
```

Once a BFD peer has been created, there are a variety of additional configuration commands available in peer mode:

Command	Description
<code>receive-interval <10-60000></code>	Sets the minimum amount of time in milliseconds between receiving Control packets that the device is capable of receiving.
<code>transmit-interval <10-60000></code>	Sets the minimum amount of time in milliseconds between sending Control packets. The actual transmit interval will also depend on the receive interval configured on the peer.
<code>detect-multiplier <2-255></code>	The detect-multiplier value is multiplied by the transmit-interval of the remote device to get the detect time . The detect time is the amount of time a device will wait for Control packets before changing the session state to Down.
<code>echo-mode</code>	Enables echo-mode. With echo-mode enabled, the device will start sending echo packets, which the remote device will return straight back. If a certain number of the device's echo packets aren't returned, then the session state will be set to Down. When in echo mode you should increase the transmit-interval so that Control packets are sent less frequently, as they are no longer the main form of failure detection.
<code>echo-interval <10-60000></code>	Sets the minimum echo receive interval in milliseconds that the device is capable of.
<code>shutdown</code>	Disables the peer session and sets the state to administrative Down.

Configuring BFD profiles

From version **5.5.2-0.1** onwards, you can create BFD profiles.

BFD profiles allow you to create profiles that contain BFD parameters, and easily apply them to a peer. This makes it quicker and more consistent to deploy your BFD settings. To define a profile:

Step 1: Create a BFD profile

```
awplus# configure terminal
awplus(config)# bfd profile Profile1
```

Step 2: Configure the profile

In profile configuration mode, use the **detect-multiplier**, **receive-interval**, **transmit interval**, **echo-mode**, and **echo-interval** commands to define the profile characteristics.

```
awplus(config-profile)# detect-multiplier 5
awplus(config-profile)# receive-interval 500
awplus(config-profile)# transmit-interval 500
awplus(config-profile)# echo-mode
awplus(config-profile)# echo-interval 100
```

Step 3: Assign the profile to a peer

To assign a profile to a BFD peer, use the **profile (BFD Peer)** command in BFD peer configuration mode.

```
awplus# configure terminal
awplus(config)# bfd peer 192.0.2.6
awplus(config-peer)# profile Profile1
```

BFD for BGP

You can create a BFD peer with a BGP neighbor. To create a BFD peer with a BGP neighbor:

Step 1: Register each BGP neighbor

Allocate a different AS number to each of the neighbors.

```
awplus#configure terminal
awplus(config)#router bgp <1-65535>
```

Step 2: Configure the IP address of each BGP neighbor

```
awplus(config-router)#neighbor <peer-address> remote-as <as-number>
```

Step 3: Enable BFD fall-over support on each BGP neighbor

Use this command to listen for BFD events registered on the same destination as the BGP neighbor.

```
awplus(config-router)#neighbor <peer-address> fall-over bfd
```

Step 4: Add extra parameters as required

```
awplus(config)# bfd peer <peer-address> [multihop] [local-address <local-address>] [interface <interface-name>] [vrf <vrf-name>]
```

Note: If the BGP neighbor is configured for ebgp-multihop, then BFD must also be configured for multihop as shown.

For example:

```
!
bfd peer 192.0.1.1 multihop local-address 192.0.2.2
!
router bgp 9
 neighbor 192.0.1.1 remote-as 40000
 neighbor 192.0.1.1 ebgp-multihop 255
 neighbor 192.0.1.1 fall-over bfd multihop
!
```

From version **5.5.2-0.1** onwards, you can assign profiles to BGP neighbors. To assign a BFD profile to a BGP neighbor, use the **neighbor fall-over bfd** command with the **profile** parameter in BGP router mode.

```
awplus# configure terminal
awplus(config)# router bgp 20
awplus(config-router)# neighbor 192.168.1.1 fall-over bfd profile Profile1
```

You can also enable multihop BFD fall over detection for a BGP neighbor. To enable BFD fall over detection for a multihop BGP neighbor, use the **neighbor fall-over bfd** command with the **multihop** parameter in BGP router mode.

```
awplus# configure terminal
awplus(config)# router bgp 20
awplus(config-router)# neighbor 192.168.1.1 fall-over bfd multihop
```

You can use both the **multihop** and **profile** parameters at the same time.

```
awplus# configure terminal
awplus(config)# router bgp 20
awplus(config-router)# neighbor 192.168.1.1 fall-over bfd multihop profile Profile1
```

From version **5.5.2-1.1** onwards, you can also enable BFD fall over detection for a BGP neighbor under VRF. To enable BFD fall over detection for a multihop BGP neighbor, use the **neighbor fall-over bfd** command in IPv4 VRF configuration mode. You can enable multihop with the **multihop** parameter, and assign a BFD profile with the **profile** parameter.

```
awplus# configure terminal
awplus(config)# router bgp 20
awplus(config-router)# address-family ipv4 vrf green
awplus(config-router-af)# neighbor 192.168.1.1 fall-over bfd multihop
profile Profile1
```

BFD for OSPF

From version **5.5.2-0.1** onwards, you can configure an OSPF interface to be monitored via a BFD session. This lets the corresponding BFD session link Up/Down events update the OSPF interface status instantly.

To enable BFD support for an OSPF interface, use the **ip ospf bfd** command in OSPF interface mode.

```
awplus# configure terminal
awplus(config)# interface vlan1
awplus(config-if)# ip ospf bfd
```

To enable BFD support for all interfaces of an OSPF instance, use the **bfd all-interfaces** command in OSPF router mode.

```
awplus# configure terminal
awplus(config)# router ospf 10
awplus(config-router)# bfd all-interfaces
```

To assign a BFD profile to an OSPF instance, use the **ip ospf bfd** command with the **profile** parameter in OSPF router mode.

```
awplus# configure terminal
awplus(config)# router ospf 10
awplus(config-router)# ip ospf bfd profile Profile1
```

From version **5.5.2-1.1** onwards, you can enable multihop BFD support and apply a BFD profile to an OSPF area virtual link.

```
awplus# configure terminal
awplus(config)# router ospf 10
awplus(config-router)# area 1 virtual-link 10.10.10.1 fall-over bfd
profile bfdProfile
```


BFD for static routes

From version **5.5.2-0.1** onwards, you can enable BFD fall over detection for static routes.

To enable BFD support for a specific static route, use the **ip route** command with the **fall-over bfd** parameter.

```
awplus# configure terminal
awplus(config)# ip route 192.168.3.0/24 10.10.0.2 fall-over bfd
```

To enable BFD support for a specific L3 interface, use the **ip route bfd** command.

```
awplus# configure terminal
awplus(config)# interface vlan1
awplus(config-if)# ip route bfd
```

To enable BFD fall-over support for all interfaces, use the **ip route bfd all-interfaces** command.

```
awplus# configure terminal
awplus(config)# ip route bfd all-interfaces
```

From version **5.5.2-1.1** onwards, you can also assign a BFD profile to static routes.

To enable BFD support for a specific static route and apply a BFD profile, use the **ip route** command with the **fall-over bfd** and **profile** parameters.

```
awplus# configure terminal
awplus(config)# ip route 192.168.3.0/24 10.10.0.2 fall-over bfd profile Profile1
```

To enable BFD support for a specific L3 interface and apply a BFD profile, use the **ip route bfd** command with the **profile** parameter.

```
awplus# configure terminal
awplus(config)# interface vlan1
awplus(config-if)# ip route bfd profile Profile1
```

To enable BFD fall-over support for all interfaces and apply a BFD profile, use the **ip route bfd all-interfaces** command with the **profile** parameter.

```
awplus# configure terminal
awplus(config)# ip route bfd all-interfaces profile Profile1
```

From version **5.5.2-1.1** onwards, you can also enable BFD support for a specific static route under a VRF. To enable BFD support, use the **ip route vrf** command with the **fall-over bfd** parameter. To apply a BFD profile, use the **profile** parameter.

```
awplus# configure terminal
awplus(config)# ip route vrf vrf1 192.168.3.0/24 10.10.0.2 fall-over bfd profile Profile1
```

Configuration example

Here are the basic configuration steps to create a BFD peer session over registered BGP neighbors, where the neighbors are connected via ports 1.0.1 on each device:

Device 1 Step 1: Enter configuration mode

```
awplus# configure terminal
```

Step 2: Enable BFD

```
awplus(config)# service bfd
```

Step 3: Enter interface configuration mode

```
awplus(config)# interface vlan1
```

Step 4: Set an IP address

```
awplus(config-if)# ip address 192.0.2.1/24
```

```
awplus(config-if)# exit
```

Step 5: Establish a BGP router with a specified AS number

```
awplus(config)# router bgp 1
```

Step 6: Register a BGP neighbor with the neighbor's IP address

```
awplus(config-router)# neighbor 192.0.2.2 remote-as 2
```

Step 7: Enable BFD fall-over support

```
awplus(config-router)# neighbor 192.0.2.2 fall-over bfd
```

Device 2 Step 1: Enter configuration mode

```
awplus# configure terminal
```

Step 2: Enable BFD

```
awplus(config)# service bfd
```

Step 3: Enter interface configuration mode

```
awplus(config)# interface vlan1
```

Step 4: Set an IP address

```
awplus(config-if)# ip address 192.0.2.2/24
```

```
awplus(config-if)# exit
```

Step 5: Establish a BGP router with a specified AS number

```
awplus(config)# router bgp 2
```

Step 6: Register a BGP neighbor with the neighbor's IP address

```
awplus(config-router)# neighbor 192.0.2.1 remote-as 1
```

Step 7: Enable BFD fall-over support

```
awplus(config-router)# neighbor 192.0.2.1 fall-over bfd
```

Monitoring BFD

There are a number of commands which provide information on the configured peers.

- For information on a specific peer use the **show bfd peer** command:

```
awplus#show bfd peer 192.0.2.1

BFD Peer:
  peer 192.0.2.1 vrf default
    ID: 4021795543
    Remote ID: 171365015
    Status: up
    Uptime: 33 minute(s), 22 second(s)
    Diagnostics: ok
    Remote diagnostics: ok
    Peer Type: configured
    Local timers:
      Detect-multiplier: 3
      Receive interval: 300ms
      Transmission interval: 300ms
      Echo transmission interval: 50ms
    Remote timers:
      Detect-multiplier: 3
      Receive interval: 300ms
      Transmission interval: 300ms
      Echo transmission interval: 50ms
```

- For information on Control / Echo packets, and Session events, use the **show bfd peer counters** command.

A peer in a successor BFD session should have had at least one session up event from when the session originally started.

```
awplus#show bfd peer 192.0.2.1 counters

  peer 192.0.2.1 vrf default
    Control packet input: 10793 packets
    Control packet output: 10843 packets
    Echo packet input: 0 packets
    Echo packet output: 0 packets
    Session up events: 1
    Session down events: 0
    Zebra notifications: 1
```

- To see just the session ID, local address, peer address, and status of each peer, use the **show bfd peer brief** command:

```
awplus#show bfd peer brief
```

```
Session count: 1
SessionId      LocalAddress      PeerAddress      Status
=====
4021795543     192.0.2.2         192.0.2.1       up
```

- To see details about a BGP neighbor, and if BFD is enabled, use the **show ip bgp neighbors** command.
- To see what debugging is enabled for BFD, use the **show debugging bfd** command.
- To display the debug information for a peer, use the **debug bfd peer** command:

```
awplus#debug bfd peer
```

```
x930b#22:19:33 x930b frr-bfdd[7162]: state-change: [mhop:no peer:192.0.2.1
local:0.0.0.0 vrf:default ifname:vlan1] up
-> down reason:control-expired
22:19:33 x930b BGP[2265]: BFD session down received for peer 192.0.2.1
22:19:33 x930b BGP[2265]: Session down received FLAG = 4
22:19:33 x930b BFD[7161]: BFD Message Header
22:19:33 x930b BFD[7161]: VR ID: 0
22:19:33 x930b BFD[7161]: VRF ID: 0
22:19:33 x930b BFD[7161]: Message type: BFD Session Delete (5)
22:19:33 x930b BFD[7161]: Message length: 32
22:19:33 x930b BFD[7161]: Message ID: 0x00000005
22:19:33 x930b frr-bfdd[7162]: session-delete: mhop:no peer:192.0.2.1
local:0.0.0.0 vrf:default ifname:vlan1
22:19:33 x930b frr-bfdd[7162]: state-change: [mhop:no peer:192.0.2.1
local:0.0.0.0 vrf:default ifname:vlan1]
admin-down -> down reason:path-down
```

C613-22126-00 REV C



NETWORK SMARTER

North America Headquarters | 19800 North Creek Parkway | Suite 100 | Bothell | WA 98011 | USA | T: +1 800 424 4284 | F: +1 425 481 3895

Asia-Pacific Headquarters | 11 Tai Seng Link | Singapore | 534182 | T: +65 6383 3832 | F: +65 6383 3830

EMEA & CSA Operations | Incheonweg 7 | 1437 EK Rozenburg | The Netherlands | T: +31 20 7950020 | F: +31 20 7950021

alliedtelesis.com

© 2022 Allied Telesis, Inc. All rights reserved. Information in this document is subject to change without notice. All company names, logos, and product designs that are trademarks or registered trademarks are the property of their respective owners.