Technical Guide

Bi-directional Forwarding Detection (BFD) over Ethernet for Layer 2 Redundancy Protocols

Feature Overview and Configuration Guide

Introduction

This guide describes the AlliedWare Plus™ Bi-directional Forwarding Detection (BFD) protocol over Ethernet and how to configure it. BFD over Ethernet is a Layer 2 detection protocol designed to provide fast forwarding path failure detection for Layer 2 redundancy protocols using Ethernet across all media types.

For information on using Layer 3 Routing protocols refer to the Bi-directional Forwarding Detection (BFD) for Routing Protocols Feature Overview and Configuration Guide.

Contents

Products and software version that apply to this guide

This guide applies to AlliedWare Plus™ products that support BFD, running version **5.4.8-2.1** or later.

To see whether your product supports BFD, 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.

What is BFD?

Bi-directional Forwarding Detection (BFD) is a network protocol used to detect faults between two forwarding engines connected by a link. BFD provides rapid detection of faults to enable efficient network path convergence.

BFD is configured on each end of the link that is to be monitored. BFD detects unidirectional link failures on that link, and notifies the link peer of the failure. This allows both ends of the link to discontinue using the link even though the fault can only be detected by one of the peers.

For example, Device A may be able to receive packets from Device B over a fiber link. But Device B cannot receive packets from Device A due to a fault on that strand of the fiber. BFD on Device B can rapidly detect this error and signal to Device A that there is a fault on the link.

There are two types of faults that BFD can detect on a link:

- Link fault (for example power failure, incorrect configuration or physical fault)
- Wiring fault (for example incorrect wiring of ports between devices)

The BFD protocol can operate directly over Ethernet to detect link failures and speed up the convergence time of Layer 2 protocols such as:

- STP, RSTP and MSTP
- EPSR
- G.8032

Limitations of BFD on AlliedWare Plus devices

Inter-operation

BFD over Ethernet will not inter-operate with other vendor implementations of BFD.

Aggregator

BFD cannot be configured on aggregators or their member ports. If BFD is configured on a port, it cannot be added to an aggregator until the BFD configuration is removed.

CFC failover

BFD is not fully supported for a CFC failover, or chassis failover in the case of a VCStack Plus. BFD relies on the active CFC to transmit frames to the peer device. During the stack failover event there are several seconds during which the CFC is busy processing the failover. During this time it cannot send BFD frames.

This means the peer device is likely to detect that the link is BFD Down (depending on its configured multiplier and interval). BFD may trigger a topology change event for MSTP/EPSR/G.8032, which may be followed by the link being detected as BFD Up and another topology change event occurring.

DLF storm control

BFD cannot be configured on the same port as DLF (Destination lookup Failure) storm control. During a storm, if DLF storm control is configured on the same port as BFD, this causes BFD frames to be discarded. This results in the BFD fault detection flapping, which can cause instability in the network as a whole due to STP, EPSR and G.8032 continually re-converging.

How does BFD work?

BFD is disabled on all ports by default and must be enabled at each end of a link to operate successfully. When BFD is initially enabled on a port it will be considered as locally Down with a link fault. This will be notified to the L2 topology protocols for that port. It will not go Up/OK until it has received BFD packets from the neighboring device. There is no authentication mechanism to prevent spoofing of BFD frames, however the device will drop any BFD frames received on ports that BFD is not enabled on.

Interval and multiplier

BFD works by sending out Ethernet frames encapsulating a BFD payload at a user-defined **interval** on the port. Another user-defined value called the detect **multiplier** represents the number of consecutive frames that must be lost causing BFD to transition to local Down and mark the link as a Fault.

- By default the interval on ports is 200ms and the multiplier is 3.
- The minimum configurable interval is 34ms and the minimum multiplier is 2.

The interval and multiplier do not need to be identical on each end of the link (although in practice this is sensible). The protocol will negotiate to use the higher of its local and the peer's advertised interval, although initially it always sends frames using a 200ms interval until the negotiation is complete. The detect multiplier is not negotiated. This means that the detection threshold for a link failure can be different on the local and remote systems.

Port status

BFD independently tracks the local port status and the remote port status, then it determines the status of the link based on the local and remote protocol states. If either the local or remote states are detected to be Down or Mismatch, then the link will be marked as Fault. Only when both the local and remote states are UP will the link status be marked as OK.

Layer 2 topology protocols

BFD is run on ports that are also running one of these Layer 2 topology protocols: MSTP, EPSR or G.8032. When BFD stops receiving frames on a port, the local status will transition to Down. The link will be marked as Fault, a log message is generated, and the topology protocols are notified that the port should be considered Down. However the port will otherwise remain electrically Up and operational and no other protocols running on the device are notified of the failure.

The topology protocols will react immediately to BFD's signal. This will cause EPSR and G.8032 to failover to their secondary link, and cause MSTP to move the port into the Disable/Discarding state. This blocks all regular traffic on that port except for the BFD frames and triggers a topology recalculation if necessary. When BFD begins receiving packets on the port again, it will signal back to the topology protocols that the port is considered Up. This allows these protocols to reverse their previous operations. When BFD is disabled on a port the topology protocols are notified and they will cease taking BFD status into account when assessing link status.

Fault detection

If BFD has been enabled on an interface, but disabled on the remote peer interface, because the local system is not receiving BFD frames, it will transition its local state to Down and indicate a link Fault. BFD will remain in these states until it begins receiving frames when BFD has been enabled on the remote peer interface. When BFD is enabled on the remote peer interface, the peer will itself initially log that there is a remote fault (as the local device has a local state of Down), but then both interfaces will locally transition to Up and the link state will become OK once a BFD neighborship is established and frames flow normally.

Link faults

BFD can quickly identify link faults over intermediate Layer 1 devices such as repeaters or media converters. For example, in the EPSR topology below, a failure has happened between two media converters, but the connection from each of the switches to the media converters remains intact and electrically up.

From the perspective of the switches, the link is still up and so EPSR is not able to use the link status to detect the failure in connectivity and instead only detects the failure when it fails to receive the periodic EPSR hello packet on its secondary port. This process can take up to three seconds with a default configuration. If the link is used for packet-loss sensitive traffic such as voice or video, a delay of three seconds before the failure is detected is undesirable. In this case BFD is able to detect the link failure and induce a failover in as little as 68ms.

BFD is also used to detect unidirectional link failures much faster than comparable technologies such as UDLD.

Figure 1: Example of a Media Converter failure:

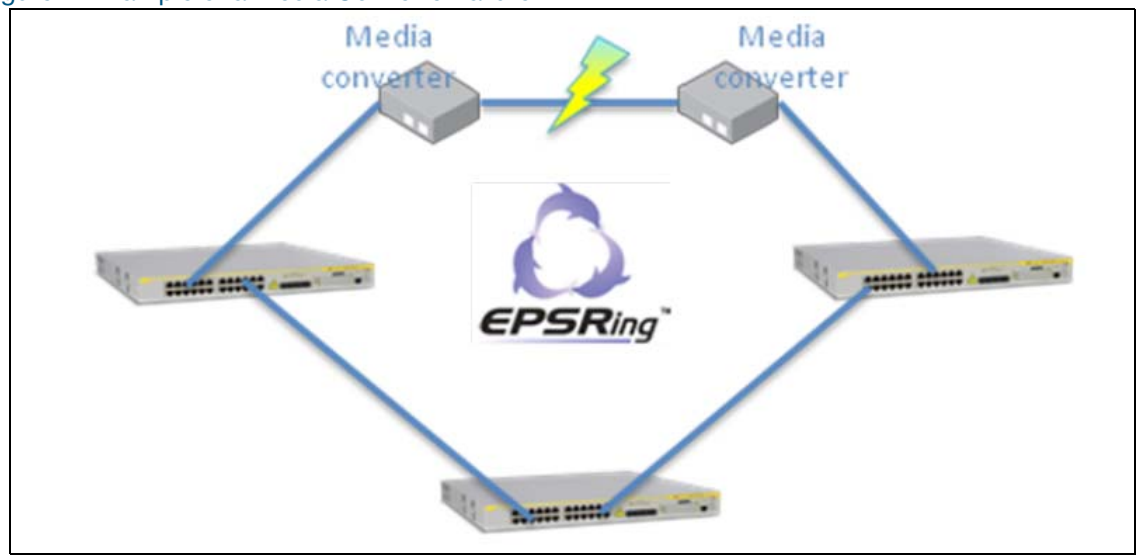

Link fault example

If a unidirectional link failure occurs, for example:

Device A port 1.1.1 receives frames from Device B port 1.1.2 but Device B cannot receive frames sent by Device A. Then Device B will initially detect the fault, change its local state to Down, change the link status to Fault and log the event with the message “port1.1.2: Link fault: Local state is DOWN”.

Device B signals its local state in its BFD updates which Device A is still able to receive. This causes Device A to also mark the remote state as Down, mark the link state as Fault and log the event with the message “port1.1.1: Link fault: Remote state is DOWN”.

Therefore, both Device A and Device B will notify the Layer 2 topology protocols that the link is faulty when either side detects it is unable to receive frames from its neighbor.

Wiring faults

BFD can also detect incorrect wiring of ports between devices. For example, fiber links are used between two devices, and by mistake the TX fiber from Device A port 1 has been incorrectly connected to the RX fiber port 2 on Device B (and vice versa). In this scenario all of the ports link up as they are receiving light on their RX lines, but data may not be switched correctly depending on the VLAN configuration of the various ports, or other factors such as spanning tree. BFD is able to immediately detect such a mis-wiring and log the event.

Figure 2: Example of a wiring failure:

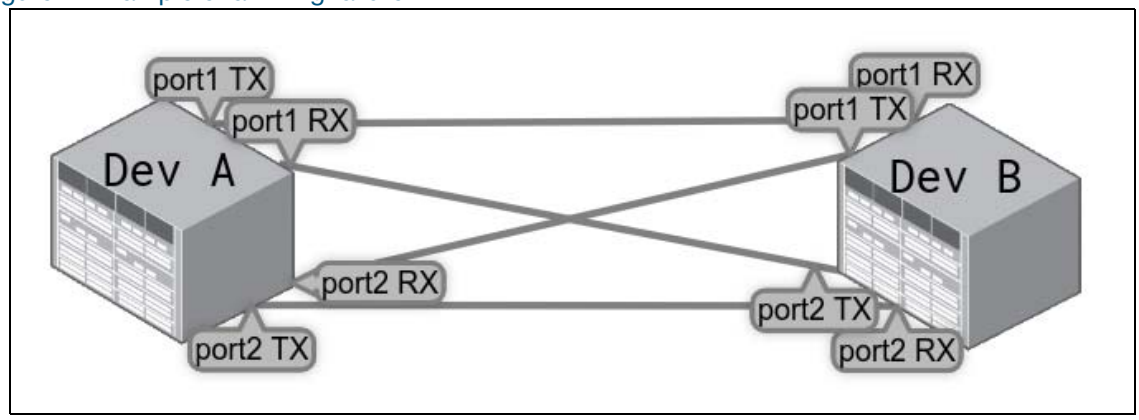

Wiring fault example

If a wiring fault occurs BFD detects this has occurred because each device generates a unique discriminator value for each of its local ports. When ports form a BFD session, they transmit their local discriminator values to each other, and then, in subsequent frames, reflect the discriminator value they learned back towards the other port.

For example, in a correctly cabled network, Device A port 1 sends its discriminator value of “0x27FB” to device B port 1, and in Device B’s next frame it repeats “0x27FB” back to Device A port 1. Therefore Device A can confirm that it is correctly transmitting and receiving from the same port on Device B.

In an incorrectly cabled network Device A port 1’s discriminator “0x27FB” is received by Device B port 1, and Device A port 2’s discriminator “0x30AC” is received by Device B port 2. Because of the wiring fault, Device B port 2 is connected to port 1 on Device A instead of port 2. So Device A port 1 detects that the received value of “0x30AC” does not match the value “0x27FB” that it is transmitting. This is detected as a cabling mismatch fault. Device A logs the error with “port1.1.1: Link fault: Local state is MISMATCH” and transitions its local state to MISMATCH and the link state as Fault.

How to configure BFD

On the wire, BFD packets are transmitted with an ethertype of “0xAAAA” and a reserved unicast destination MAC address of “000c.250a.8c01”. Allied Telesis products that are capable of running BFD will trap these packets and not forward them beyond the local link once BFD has been configured on at least one port in the system. However because this is a regular unicast MAC address, any AlliedWare Plus products that do not support BFD or does not have BFD configured, or any non-AlliedWare Plus devices, will flood BFD frames out all of their ports within the same L2 broadcast domain.

Therefore we recommend only configuring BFD on links that are directly connected to other AlliedWare Plus devices that support BFD. If the BFD frames somehow leak out of the local link, we recommend installing hardware ACLs that match on the destination MAC address (or ethertype “0xAAAA”) to drop the frames.

The following tasks can be performed to configure BFD:

- Enable BFD on a port
- Configure the interval and multiplier (optional)
- Configure a hardware ACL (optional)
- Apply an ACL globally or on a specific port (optional)
- Run the show commands to monitor BFD
- Debug BFD

To enable BFD on a port

To enable BFD on a port with the default interval and multiplier, from Interface Configuration mode, use the command:

```
awplus(config-if)#bfd ethernet-mode
```

No special configuration is required for EPSR, MSTP or G08032. BFD needs to be configured on both ends of a link.

For example to enable BFD on port1.1.1, use the commands:

```
awplus#configure terminal
awplus(config)#int port1.1.1
awplus(config-if)#bfd ethernet-mode
```

Use the **no** variant of this command to disable BFD. For example to disable BFD on port1.1.3, use the commands:

```
awplus#configure terminal
awplus(config)#int port1.1.3
awplus(config-if)#no bfd ethernet-mode
```

To configure the BFD interval and multiplier

Use this command if you want to configure the BFD interval and multiplier to something other than their default values. By default the interval on ports is 200ms and the multiplier is 3.

```
awplus(config-if)#bfd ethernet-mode [interval <interval>|multiplier
<multiplier>]
```

For example, to configure the BFD interval to 500ms and the multiplier to 6, use the commands:

```
awplus#configure terminal
awplus(config)#int port1.1.1
awplus(config-if)#bfd ethernet-mode interval 500 multiplier 6
```

BFD needs to be configured on both ends of a link.

To configure a hardware ACL

Use this command to configure a hardware ACL to deny BFD frames on an AlliedWare Plus device that does not support BFD. For example to create the ACL BlockBFD for the destination MAC address 000c.250a.8c01, use the commands:

```
awplus#configure terminal
awplus(config)#access-list hardware BlockBFD
awplus(config-ip-hw-acl)#deny mac any 000c.250a.8c01 ffff.ffff.ffff
awplus(config-ip-hw-acl)#exit
```

To apply the BlockBFD ACL to a specific port, use the commands:

```
awplus#configure terminal
awplus(config)#int port1.1.2
awplus(config-if)#access-group BlockBFD
```

To apply the BlockBFD ACL to all ports globally, use the commands:

```
awplus#configure terminal
awplus(config)#access-group BlockBFD
```

Show commands to monitor BFD

Use this command to display the current status of ports in the system that have BFD enabled on them.

```
awplus#show bfd
```

```
awplus#show bfd

Interface    Interval(ms)  Multiplier  Local State  Remote State  Link State
-------------------------------------------------------------------------------
port1.1.2    200           3           DOWN         -             Fault
port1.1.3    500           4           UP           UP            OK
```

This shows the current status of ports in the system that have BFD enabled on them. Ports that do not have BFD configured are not listed in this output.

Table 1: show bfd status and descriptions

PARAMETER	DESCRIPTION	
Interface	The interface name of the link to apply BFD to, for example port1.1.1	
Interval (ms)	The interval is measured in milliseconds and is the interval between when BFD frames will be sent on a specified port. This is the locally configured value for the port, not the value negotiated with the peer. The default value is 200.	
Multiplier	The multiplier is the number of BFD frames that must be lost on this device before BFD assumes the link is down. The default value is 3.	
Local State	The local BFD state for this port.	
	UP	BFD has received at least 1 frame within its detection window (negotiated interval detection multiplier).
	DOWN	BFD has not received any frames within its detection window, or the neighbor failed to receive packets within its detection window and this device has learned that.
	MISMATCH	BFD is receiving frames within its detection window but has detected a wiring error.
Remote State	The BFD state for the remote peer, one of Up, Down or Mismatch. If the local state is Down, this will display as a “-” because the local device does not know the current status of the remote device.	
Link State	BFD’s opinion of the state of the link. This is the value that is advertised to the Layer 2 topology protocols, either OK or Fault.	

Use this command to display the current status of a specified interface or all BFD configured interfaces.

```
awplus#show bfd interface [<interface-name>]
```

```
!
awplus#show bfd interface

Interface:                    port1.1.4
Interval:                     200ms
Multiplier:                   6
Local State:                  UP
Remote Session State:         UP
Combined Session State:       OK
Remote Detection Multiplier:  3
Local Discriminator:          0x9c302dcf
Remote Discriminator:         0x32340eb8
Negotiated TX Interval:       500ms
Negotiated RX Interval:       500ms
Packet TX Count:              43
Packet RX Count:              42
```

Table 2: show bfd interface status and descriptions

BFD PARAMETERS AND STATUS	DESCRIPTION
Interface	The interface name of the link to apply BFD to, for example port1.1.4.
Interval (ms)	The locally configured frame interval used by the specified port. The interval is measured in milliseconds. The default value is 200.
Multiplier	The locally configured detection multiplier used by the specified port. The default value is 3.
Local State	The local BFD state for this port.
Remote Session State	The current state of the remote peer, one of Up, Down or Mismatch.
Combined Session State	The state of the session as a whole. This is summarised as “Link state” in the **show bfd** output. This displays OK, or one of the following fault reasons: ■ Fault (Session state is DOWN) ■ Fault (Session state is MISMATCH) ■ Fault (Remote session state is DOWN).
Remote Detection Multiplier	The multiplier is the number of BFD frames that must be lost on the peer device before BFD assumes the link is down. The default is 3 if there are no frames received from the peer.
Local/Remote Discriminator	This is a unique value to identify different interfaces. This is used in the cable mismatch detection process. The default value is 0 if there are no frames received from the peer.
Negotiated TX/RX Intervals	The TX/RX intervals that have been negotiated with the peer. These are the greater of the configured interval values of each device. The default value is 200ms if there are no frames received from the peer.
Packet TX/RX Count	This is how many BFD frames have been sent or received on the specified port.

Debug BFD

To print BFD debugging information to the logs, from Privileged Executive mode, use the command:

```
awplus#debug bfd [all|nsm|pkt|state]
```

The optional parameters allow you to print all, only nsm, only pkt or only state change debug information.

For example, to print all BFD debug information to the logs, use the commands:

```
awplus>enable
awplus#debug bfd all
```

C613-22115-00 REV C

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