ERICSSON

Configuring BFD

SYSTEM ADMINISTRATOR GUIDE

17/1543-CRA 119 1170/1-V1 Uen H

Contents

1 Overview

This document provides an overview of Bidirectional Forwarding Detection (BFD) and describes the tasks and commands used to configure, monitor, troubleshoot, and administer BFD features through the SmartEdge router.

This document applies to both the Ericsson SmartEdge® and SM family routers. However, the software that applies to the SM family of systems is a subset of the SmartEdge OS; some of the functionality described in this document may not apply to SM family routers.

For information specific to the SM family chassis, including line cards, refer to the SM family chassis documentation.

For specific information about the differences between the SmartEdge and SM family routers, refer to the Technical Product Description *SM Family of Systems* (part number 5/221 02-CRA 119 1170/1) in the **`Product Overview`** folder of this Customer Product Information library.

Bidirectional Forwarding Detection (BFD) is a simple Hello protocol that, in many respects, is similar to the detection components of some routing protocols. A pair of routers periodically transmit BFD packets over each path between the two routers, and if a system stops receiving BFD packets after a predefined time interval, some component in that particular bidirectional path to the neighboring router is assumed to have failed.

A path is only declared to be operational when two-way communication has been established between systems.

Note: Currently, only Ethernet or Gigabit Ethernet (GE) ports support BFD communication. For BFD to work properly, the BFD neighbor must be reachable over an Ethernet link. All Ethernet line cards (PPA2 and PPA3) support BFD communication.

BFD provides low-overhead, short-duration detection of failures in the path between adjacent forwarding engines, including the interfaces, data links, and to the extent possible, the forwarding engines themselves.

The legacy Hello mechanism run by routing protocols does not offer detections of less than one second, and for some applications, more than one second is too long and represents a large amount of data loss at gigabit rates. BFD provides the ability to detect communication failures in less than one second.

Note: BFD cannot be enabled on bridge, intercontext, and loopback interfaces.

1.1 FRR Triggered by BFD to the IP Next Hop

BFD detection supports fast reroute (FRR) on Resource Reservation Protocol (RSVP) links. You must enable BFD at both the router RSVP level and for any individual neighbor that requires fast reroute (FRR). By default, BFD detection is disabled on RSVP links. For more information about RSVP FRR, see *Configuring MPLS*.

1.2 BFD Over LAGs

The SmartEdge® router supports single-session and multiple-session BFD over IPv4 and IPv6 link aggregation groups (LAGs) as follows:

- With multiple-session BFD over LAGs, a separate BFD session runs over each constituent link in a LAG. Multiple BFD sessions use the same nexthop IP address on different designated links. The 1:1 binding between the constituents must be symmetric so that packets originating from one local constituent reach one remote constituent only, and packets originating from the remote constituent reach one local constituent.

- With single-session BFD over LAGs, BFD detects whether a Layer 3 (L3) neighbor is active without considering the Layer 2 (L2) interface that connects the neighbors. BFD designates a "home slot" (a line card within the LAG) on which a BFD session operates. BFD packets are transmitted and received on the card in the home slot; if the card in the home slot fails, a backup home slot card takes over. (Backup home slot card selection is automatic; you do not need to configure anything to ensure backup selection occurs.)

Single- and multiple-session BFD are supported on inter- and intra-card IPv6 and IPv4 802.3ad LAGs. Single-session BFD packets can be transmitted and received on any link within a LAG.

Note: BFD is not supported on access LAGs. BFD is supported only on Ethernet and 802.1q LAGs.

Single-session BFD is preferable to multiple-session BFD when:

- The next-hop node is not a router that has the SmartEdge OS installed.

- An L2 network exists between the pair of routers (because the L2 switches typically terminate the LAG and use their own hashing algorithm for the next L2 segment).

- Multiple BFD sessions are configured between two nodes. In this instance, using single-session BFD for some or all neighbors reduces resource usage and supports scaling.

Note: When a link fails, the router (or any intermediate switch) must switch over to an active link (if an active link is available). If a switchover does not happen, non-deterministic BFD states can occur (for example, the BFD session may go down or flap).

Use the *show bfd session* command to display the number of constituent links within a LAG (for multiple-session BFD) , and the home slot and backup home slot for a LAG (for single-session BFD). Use the *show bgp neighbor* command to display neighbor BFD session data.

Note: If the remote peer is not reachable over a link group, the router ignores the trunk LAG settings (configured with the *link-group (BFD)* command) for the interface or neighbor.

2 Configuration and Operations Tasks

Note: In this section, the command syntax in the task tables displays only the root command.

To configure BFD, perform the tasks in the sections that follow.

2.1 Configuring a BFD Neighbor

A BFD session is established for each BFD neighbor configured. More than one BFD neighbor can be configured.

To configure a BFD neighbor, perform the tasks described in Table 1. Enter all commands in BFD neighbor configuration mode, unless otherwise noted.

Table 1 Configure a BFD Neighbor

Task	Root Command	Notes
Create a BFD instance and enter BFD router configuration mode.	*router bfd*	Enter this command in context configuration mode.
Create a new BFD neighbor, or select an existing one for modification, and enter BFD neighbor configuration mode.	*neighbor (BFD)*	Enter this command in BFD router configuration mode.

Table 1 Configure a BFD Neighbor

Task	Root Command	Notes
Specify the detection multiplier value.	*detection-multiplier*	The negotiated minimum transmit interval (the minimum desired transmit interval agreed upon by both peers) is multiplied by the detection multiplier value to provide the detection time for the transmitting system in asynchronous mode. The detection time is the time it takes to declare a neighbor as down. For example, if the minimum desired transmit interval was negotiated at 10 ms and the detection multiplier is set to 3, then the detection time is 30 ms. Using the detection multiplier adds robustness to BFD by allowing the system to not bring down a neighbor if only one BFD packet is missed.
Specify the minimum required interval, in milliseconds, between received BFD control packets that the system is capable of supporting.	*minimum receive-interval*	—
Specify the minimum desired transmit interval, in milliseconds, used by the local system when transmitting BFD control packets.	*minimum transmit-interval*	—
Specify whether single-session or multiple-session BFD should be used for a neighbor that is part of a LAG.	*link-group (BFD)*	All deployments of BFD over trunk LAGs are multiple-session by default. Use single-session BFD to interoperate with other vendor equipment and SE with a L2 switch in between.

2.2 Configuring BFD on an Interface

Configuring BFD on an interface establishes a separate BFD session for each neighbor on the interface. Neighbors are learned by the client routing protocol (such as Open Shortest Path First [OSPF]) that has BFD detection enabled.

Note: BFD clients are routing protocols that use BFD to detect communication failures in less than one second. Currently, BFD supports static routes, RSVP links, external Border Gateway Protocol (eBGP), Protocol Independent Multicast (PIM), Intermediate System-to-Intermediate System (IS-IS), and OSPF.

Note: BFD cannot be enabled on bridge, intercontext and loopback interfaces.

To configure BFD on an interface, perform the tasks described in Table 2. Enter all commands in BFD interface configuration mode, unless otherwise noted.

Table 2 Configure BFD on an Interface

Task	Root Command	Notes
Create a BFD instance and enter BFD router configuration mode.	*router bfd*	Enter this command in context configuration mode.
Enable BFD on a named interface and enters BFD interface configuration mode.	*interface*	Enter this command in BFD router configuration mode. The interface must already be configured through the **`interface`** command (in context configuration mode) before BFD can be enabled on it. For more information about the **`interface`** command, see *Configuring Contexts and Interfaces*.
Specify the detection multiplier value.	*detection-multiplier*	The negotiated minimum transmit interval (the minimum desired transmit interval agreed up by both peers) is multiplied by the detection multiplier value to provide the detection time for the transmitting system in asynchronous mode. The detection time is the time it takes to declare a neighbor as down. For example, if the minimum desired transmit interval was negotiated at 10 ms and the detection multiplier is set to 3, then the detection time is 30 ms. Using the detection multiplier adds robustness to BFD by allowing the system to not bring down a neighbor if only one BFD packet is missed.

Table 2 Configure BFD on an Interface

Task	Root Command	Notes
Specify whether single-session or multiple-session BFD should be used for an interface that is part of a LAG.	*link-group (BFD)*	All deployments of BFD over trunk LAGs are multiple-session by default.
Specify the minimum required interval, in milliseconds, between received BFD control packets that the system is capable of supporting.	*minimum receive-interval*	—
Specify the minimum desired transmit interval, in milliseconds, used by the local system when transmitting BFD control packets.	*minimum transmit-interval*	—

2.3 Disabling BFD for an IS-IS Interface

By default, when BFD is enabled for the same interface on which IS-IS has been enabled, BFD is automatically enabled for each neighbor on the interface. To disable BFD for an IS-IS interface, perform the tasks described in Table 3.

Table 3 Disable BFD for an IS-IS Interface

Task	Root Command	Notes
Access IS-IS router configuration mode.	*router isis*	Enter this command in context configuration mode.
Access IS-IS interface configuration mode.	*interface (IS-IS)*	Enter this command in IS-IS router configuration mode. Only one IS-IS instance can be running on an interface.
Disable BFD for an IS-IS interface.	*disable-bfd (IS-IS)*	Enter this command in IS-IS interface configuration mode. For more information about IS-IS, see *Configuring IS-IS*.

2.4 Disabling BFD for an OSPF Interface

By default, when BFD is enabled for the same interface on which OSPF has been enabled, BFD is automatically enabled for each neighbor on the interface. To disable BFD for an OSPF interface, perform the tasks described in Table 4.

Table 4 Disable BFD for an OSPF Interface

Task	Root Command	Notes
Access OSPF router configuration mode.	*router ospf*	Enter this command in context configuration mode.
Access OSPF area configuration mode.	*area*	—
Enter OSPF interface configuration mode.	*interface (OSPF)*	—
Disable BFD for an OSPF interface.	*disable-bfd (OSPF)*	Enter this command in OSPF interface configuration mode. For more information about OSPF, see *Configuring OSPF*.

2.5 Disabling BFD for a PIM Interface

By default, BFD is enabled on PIM interfaces and for each neighbor on the interface. To disable BFD for a PIM interface, perform the tasks described in Table 5.

Table 5 Disable BFD for a PIM Interface

Task	Root Command	Notes
Access interface configuration mode for the PIM interface on which you want to disable BFD.	*interface (BFD)*	Enter this command in context configuration mode.
Disable BFD for the PIM interface.	*no pim bfd*	For more information about PIM, see *Configuring IP Multicast*.

2.6 Enabling BFD for a Static Route

By default, BFD is disabled for all static routes, but you can enable BFD for a particular static route. When BFD detects a communication failure to the next hop specified for a static route that has BFD enabled, the static route is withdrawn.

To enable BFD for a static route, perform the task described in Table 6.

Table 6 Enable BFD for a Static Route

Task	Root Command	Notes
Enable BFD for a static route.	*ip route* *ipv6 route*	Enter this command in context configuration mode. Use the **bfd** keyword to enable BFD for a static route. For more information about static routes, see *Configuring Basic IP Routing*.

2.7 Enabling FRR Triggered by BFD to the IP Next Hop

By default, BFD detection is disabled on RSVP links. You must enable BFD at both the router RSVP level and for any individual neighbor that requires FRR.

Note: FRR for RSVP supports single-hop IP BFD only. The SmartEdge router does not support Multiprotocol Label Switching (MPLS) BFD.

Consider the following rules and restrictions when enabling FRR triggered by BFD to the IP next hop:

- BFD is supported on one next hop for each interface only and runs outside the RSVP tunnel.
- Label Distribution Protocol (LDP)-over-RSVP Label Switched Paths (LSPs) do not support BFD.
- If a port-down event occurs simultaneously with a BFD event, the port-down event takes precedence over the BFD event.
- If protected LSPs traverse a LAG for which BFD is enabled and the number of remaining active links falls below the minimum number of links required for the LAG to be active, BFD triggers FRR.
- If a protected link in a LAG fails and the number of remaining active links is greater than or equal to the minimum number of links required for the LAG to be active, traffic that was traversing the failed link is redistributed to the remaining links in less than 50 milliseconds after the failed link is detected.

Note: Traffic redistribution is supported only when an LSP is protected.

To enable FRR triggered by BFD to the IP next hop, perform the task described in Table 7.

Table 7 Enable FRR Triggered by BFD to the IP Next Hop

Task	Root Command
Enable BFD on the RSVP client. BFD does not work if it is not first enabled at the RSVP level:	

Table 7 Enable FRR Triggered by BFD to the IP Next Hop

Task	Root Command
Enter context configuration mode.	`context`
Enter RSVP router configuration mode.	`router rsvp`
Enable BFD on the RSVP client.	`bfd`
Enable BFD for a specific next-hop neighbor:	
Enter context configuration mode.	`context`
Create a BFD instance and enter BFD router configuration mode.	`router bfd`
Enable BFD for the next-hop neighbor.	`interface`
Verify the BFD configuration for the RSVP interface.	`show rsvp interface`

2.8 Enabling BFD for an eBGP Neighbor

By default, BFD is disabled for all eBGP neighbors, but you can enable BFD for a particular eBGP neighbor. When BFD detects a communication failure to the eBGP neighbor, the neighbor is reset.

To enable BFD for an eBGP neighbor, perform the task described in Table 8.

Table 8 Enable BFD for an eBGP Neighbor

Task	Root Command	Notes
Enable BFD for an eBGP neighbor.	*bfd (BGP neighbor)*	Enter this command in BGP neighbor configuration mode. BFD can be enabled only for eBGP neighbors; enabling BFD for an internal Border Gateway Protocol (iBGP) neighbor generates an error message. For more information about BGP, see *Configuring BGP*.

2.9 BFD Operations

To manage BFD functions, perform the appropriate tasks described in Table 9. Enter the **`show`** command in any mode. Enter the **`debug`** commands in exec mode.

Table 9 BFD Operations Tasks

Task	**Root Command**
Enable the generation of BFD debug messages.	*debug bfd*
Enable the generation of BFD debug messages for all OSPF instances.	*debug ospf bfd*
Display active BFD session information for neighbors in the current context.	*show bfd session*

3 Configuration Examples

The sections that follow provide BFD configuration examples.

3.1 BFD Neighbor

A BFD session is established for each BFD neighbor configured. More than one BFD neighbor can be configured. The following example configures the BFD neighbor, **192.168.0.24**, sets the minimum desired transmit interval to **30** ms, sets the minimum receive interval to **30** ms, and the sets detection multiplier to **4:**

```
[local]Redback#configure

[local]Redback(config)#context local

[local]Redback(config-ctx)#router bfd

[local]Redback(config-bfd)#neighbor 192.168.0.24

[local]Redback(config-bfd-nbr)#minimum receive-interval 30

[local]Redback(config-bfd-nbr)#minimum transmit-interval 30

[local]Redback(config-bfd-nbr)#detection-multiplier 4

[local]Redback(config-bfd-nbr)#end
```

3.2 BFD Interface

Configuring BFD on an interface establishes a separate BFD session for each neighbor on the interface. Neighbors are learned by the client routing protocol (such as OSPF) that has BFD detection enabled. The following example configures BFD on the interface, **foo**, sets the minimum desired transmit interval to **25** ms, sets the minimum receive interval to **40** ms, and the sets detection multiplier to **2:**

Note: BFD can be configured on a dual stack (IPv4 and IPv6) interface.

```
[local]Redback#configure
[local]Redback(config)#context local
[local]Redback(config-ctx)#router bfd
[local]Redback(config-bfd)#interface foo
[local]Redback(config-bfd-if)#minimum receive-interval 25
[local]Redback(config-bfd-if)#minimum transmit-interval 40
[local]Redback(config-bfd-if)#detection-multiplier 2
[local]Redback(config-bfd-if)#link-group single-session
[local]Redback(config-bfd-if)#end
```

```
[local]Redback(config-ctx)#router bfd
[local]Redback(config-bfd)#neighbor 3.3.3.2
[local]Redback(config-bfd)#minimum transmit-interval 10
[local]Redback(config-bfd)#minimum receive-interval 10
[local]Redback(config-bfd)#neighbor 5001:1:2::2
[local]Redback(config-bfd)#minimum transmit-interval 10
[local]Redback(config-bfd)#minimum receive-interval 10
[local]Redback(config-bfd)#neighbor 6001:1:1:1::2
[local]Redback(config-bfd)#minimum transmit-interval 10
[local]Redback(config-bfd)#minimum receive-interval 10
!
[local]Redback(config-ctx)#ipv6 route 3ffe::/64 5001:1:2::2 bfd
[local]Redback(config-ctx)#ipv6 route 3ffe::/64 6001:1:1:1::2 bfd
!

[local]Redback(config-ctx)#link-group to-dev3 dot1q
[local]Redback(config-ctx)#dot1q pvc 101
!
[local]Redback(config-ctx)#link-group to-dev3-via-xc dot1q
 [local]Redback(config-ctx)#dot1q pvc 300
```

3.3 BFD for IS-IS

By default, when BFD is enabled for the same interface on which IS-IS has been enabled, BFD is automatically enabled for each neighbor on the interface; however, you can disable BFD for the interface. The following example disables BFD for the IS-IS interface **isis-foo:**

```
[local]Redback(config)#context local

[local]Redback(config-ctx)#router isis ip-backbone

[local]Redback(config-isis)#interface isis-foo

[local]Redback(config-isis-if)#disable-bfd

[local]Redback(config-isis-if)#end
```

The following example enables BFD on an IS-IS interface.

```
[local]Redback(config-ctx)#router bfd interface r1_switch
minimum transmit-interval 1000
minimum receive-interval 1000
detection-multiplier 3
link-group single-session

neighbor 1.1.1.2
minimum transmit-interval 1000
minimum receive-interval 1000
detection-multiplier 3
link-group single-session
        !
[local]Redback(config-ctx)# router isis Router1
net 47.0001.1111.2222.1111.00
address-family ipv4 unicast
interface r1_switch
address-family ipv4 unicast
```

3.4 BFD for OSPF

By default, when BFD is enabled for the same interface on which OSPF has been enabled, BFD is automatically enabled for each neighbor on the interface; however, you can disable BFD for the interface. The following example disables BFD for the OSPF interface **ospf-foo:**

```
[local]Redback(config)#context local

[local]Redback(config-ctx)#router ospf 15

[local]Redback(config-ospf)#area 0

[local]Redback(config-ospf-area)#interface ospf-foo

[local]Redback(config-ospf-if)#disable-bfd

[local]Redback(config-ospf-if)#end
```

3.5 BFD for Static Routes

By default, BFD is disabled for all static routes, but you can enable BFD for a particular static route. The following examples enable BFD for the static route to **1.1.1.1/24** with the next hop, **2.2.2.2 and IPv6 static route to 2001::5:**

```
[local]Redback(config)#context local

[local]Redback(config-ctx)#ip route 1.1.1.1/24 2.2.2.2 bfd

[local]Redback(config-ctx)#end
```

```
[local]Redback(config)#context local

[local]Redback(config-ctx)#ipv6 route 3000::1/64 2001::5 bfd

[local]Redback(config-ctx)#end
```

3.6 BFD for PIM

By default, BFD is enabled for all PIM interfaces. If BFD is enabled on a PIM interface, BFD is automatically enabled for each neighbor on the interface; however, you can disable BFD for the interface. The following example disables BFD for the PIM interface **pim1**:

```
[local]Redback(config)#context local

[local]Redback(config-ctx)#interface pim1

[local]Redback(config-if)#no pim bfd

[local]Redback(config-if)#end
```

3.7 FRR Triggered by BFD to the IP Next Hop

By default, BFD is disabled for all RSVP LSPs, but you can enable BFD at both the router RSVP level and for any individual neighbor that requires FRR. The following example shows how to enable FRR triggered by BFD to the IP next hop. In this example, BFD is enabled on the RSVP client and for the next hop on the RSVP interface called `rsvp1`:

First, enable BFD on the RSVP client. All IP interfaces configured on that RSVP client have BFD enabled:

```
[local]Redback#configure

[local]Redback(config)#context local

[local]Redback(config-ctx)#router rsvp

[local]Redback(config-rsvp)#bfd
```

Next, enable BFD for a specific next hop neighbor:

```
[local]Redback(config)#context local

[local]Redback(config-ctx)#router bfd

[local]Redback(config-bfd)#interface rsvp1
```

3.8 BFD for eBGP

By default, BFD is disabled for all eBGP neighbors, but you can enable BFD for a particular eBGP neighbor. The following example enables BFD for the eBGP neighbor, 8.8.8.2:

```
[local]Redback(config)#context local

[local]Redback(config-ctx)#router bgp 100

[local]Redback(config-bgp)#neighbor 8.8.8.2 external

[local]Redback(config-ospf-bgp-neighbor)#bfd

[local]Redback(config-ospf-bgp-neighbor)#end
```

```
[local]Redback(config)#context local
[local]Redback(config-ctx)#router bgp <100>
 neighbor 2000::2 external
bfd
 neighbor 10.1.1.2 external
bfd
```