Bi-Directional Forwarding Detection

In This Chapter

This section provides information about bi-directional forwarding (BFD) detection.

Topics in this section include:

Applicability

This section is applicable to all of the 7x50 and 7710 series but the timing differs among platforms and these will be indicated. Note that the centralized cpm-np type is only supported by 7750/7450s equipped with SF/CPM 2 or higher. The information contained in this section has been tested with Release 8.0.R4.

Overview

Bi-Directional Forwarding Detection (BFD) is a light-weight protocol which provides rapid path failure detection between two systems. It has been recently published as a series of RFCs (RFC 5880, *Bidirectional Forwarding Detection (BFD)*, to RFC 5884, *Bidirectional Forwarding Detection (BFD) for MPLS Label Switched Paths (LSPs)*.

If a system running BFD stops receiving BFD messages on an interface, it will determine that there has been a failure in the path and notifies other protocols associated with the interface. BFD is useful in situations where two nodes are interconnected through either an optical (DWDM) or Ethernet network. In both cases, the physical network has numerous extra hops which are not part of the Layer 3 network and therefore, the Layer 3 nodes are incapable of detecting failures which occur in the physical network on spans to which the Layer 3 devices are not directly connected.

BFD protocol provides rapid link continuity checking between network devices, and the state of BFD can be propagated to IP routing protocols to drastically reduce convergence time in cases where a physical network error occurs in a transport network.

RFC 5880 define two modes of operation for BFD:

- Asynchronous mode (supported by ALU routers covered in this section) — Uses periodic BFD control messages to test the path between systems
- Demand mode (not supported by ALU router covered in this section)

In addition to the two operational modes, an echo function is defined (ALU routers covered by this section only support response, looping back received BFD messages to the original sender).

The goal of this section is to describe the configuration and troubleshooting for BFD on a link between two peers in the following scenarios:

- BFD for ISIS
- BFD for OSPF
- BFD for PIM
- BFD for Static route
- BFD IES
- BFD for RSVP
- BFD for T-LDP
- BFD support of OSPF CE-PE adjacencies
- BFD over IPSec tunnel
- BFD over VRRP

Figure 265 provides an overview of the possible BFD implementations and shows all protocols that can be bound to a BFD session.

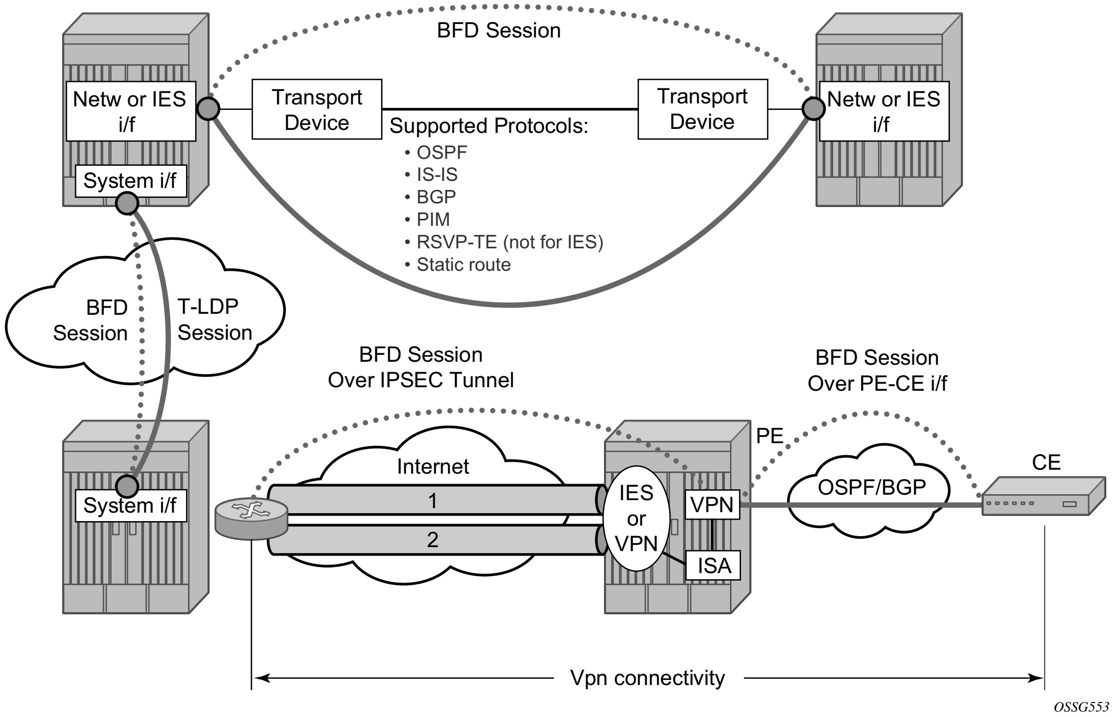

Figure 265: BFD Multi-Scenarios

Configuration

BFD packets are processed both locally (processed on IOM CPU) and centrally (processed on the CPM).

Starting with Release 8, the CPM is able to centrally generate the BFD packets at a sub second interval as low as 10 msec. However it should be noted that the BFD state machine is still implemented in software. It is the BFD packet generation that can be now selectively delegated to CPM hardware as needed. This is applicable where sub second operational requirements for BFD or scaling the number of BFD sessions beyond 250 are required.

Centralized sessions are processed:

- in software by 7x50 SR-1 and ESS-1, 7710 c4 and c12 and 7x50 equipped with SF/CPM 1.
- in hardware by 7x50 equipped with SF/CPM 2 or higher.

Minimum transmitting and receiving Intervals are as follows:

- Centralized sessions:
  - → Minimum 300 ms in 7x50 SR-1 and ESS-1, 7710 c4 and c12
  - → Minimum 100 ms in 7x50 equipped with SF/CPM 1 and in every 7x50 up to Release 7.0
  - → Minimum 10 ms in 7x50 equipped with SF/CPM 2 or higher
- Local sessions:
  - → Minimum 100 ms

The following applications require BFD to run centrally on the SF/CPM and a centralized session will be created independently of the type explicitly declared by the user:

- BFD for IES/VPRN over Spoke SDP
- BFD over LAG and VSM Interfaces
- Protocol associations using loopback and system interfaces (e.g. BFD for T-LDP)
- BFD over IPSec sessions
- BFD sessions associated with multi-hop peering

Figure 266 shows the most relevant scenarios where BFD centralized sessions are used.

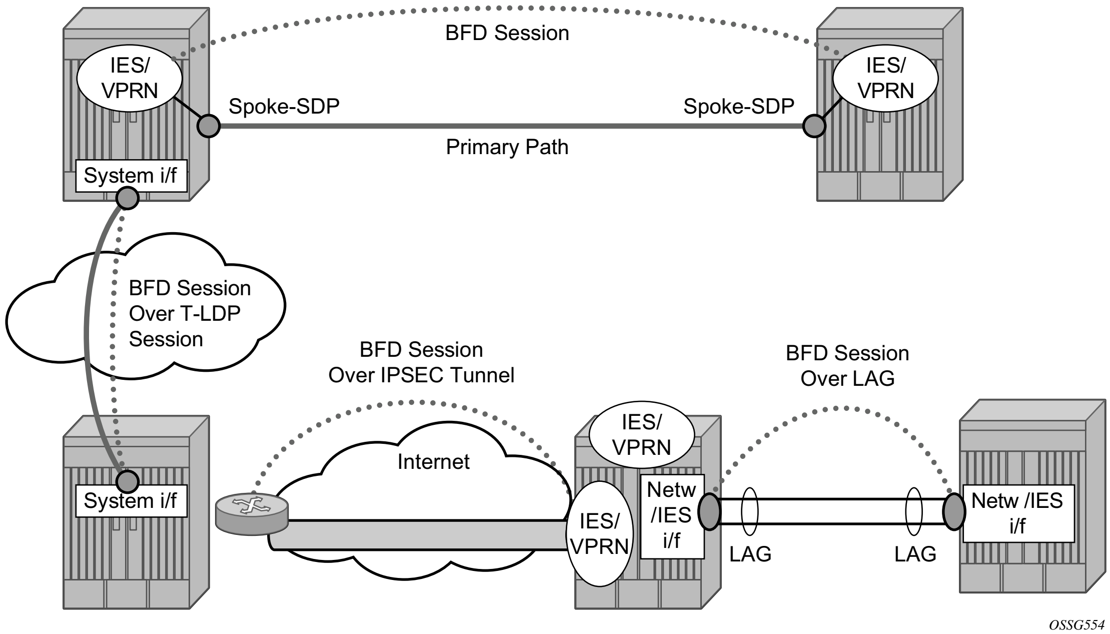

Figure 266: BFD Centralized Sessions

On the other end, when the two peers are directly connected, the BFD session is local by default, but in a 7x50 equipped with SF/CPM 2 or higher, the user can choose which session (local or centralized) to implement.

As general rule, the following steps are required to configure and enable a BFD session when peers are directly connected:

1. Configure BFD parameters on the peering interfaces.
2. Check that the Layer 3 protocol, that is to be bound to BFD, is up and running.
3. Enable BFD under the Layer 3 protocol interface.

Since most of the following procedures share the same first step, it is described only once in the next paragraph and then referred to in the following paragraphs.

BFD Base Parameter Configuration and Troubleshooting

The reference topology for the generic configuration of BFD over two local peers is shown in Figure 267.

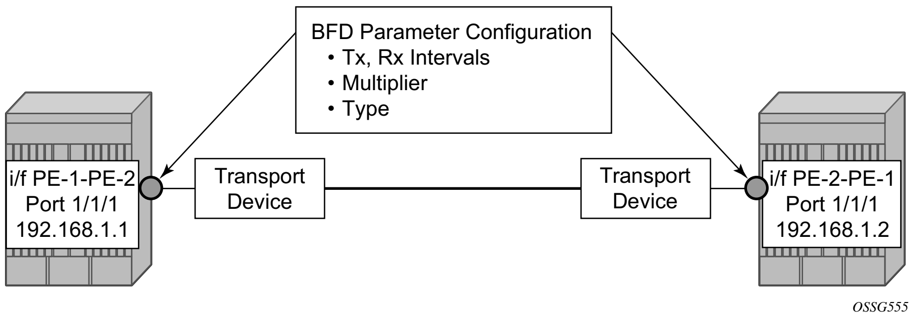

Figure 267: BFD Interface Configuration

To configure BFD between two peers, the user should firstly enable base level BFD on interfaces between PE-1 and PE-2.

On PE1:

```
configure
     router
          interface PE-1-PE-2
               address 192.168.1.1/30
               port 1/1/1
               bfd 100 receive 100 multiplier 3
          exit
     exit
exit
```

On PE2:

```
configure
     router
          interface PE-2-PE-1
               address 192.168.1.2/30
               port 1/1/1
               bfd 100 receive 100 multiplier 3
          exit
     exit
exit
```

The following **show** commands are used to verify the BFD configuration on the router interfaces on PE1 and PE2.

On PE1:

```
A:PE1# show router bfd interface
===============================================================================
BFD Interface
===============================================================================
Interface name                   Tx Interval    Rx Interval    Multiplier
-------------------------------------------------------------------------------
PE-1-PE-2                           100            100            3
-------------------------------------------------------------------------------
No. of BFD Interfaces: 1
===============================================================================
A:PE1#
```

On PE2:

```
A:PE2# show router bfd interface
===============================================================================
BFD Interface
===============================================================================
Interface name                   Tx Interval    Rx Interval    Multiplier
-------------------------------------------------------------------------------
PE-2-PE-1                           100            100            3
-------------------------------------------------------------------------------
No. of BFD Interfaces: 1
===============================================================================
A:PE2#
```

Note that, BFD being an asynchronous protocol, it is possible to configure different tx and rx intervals on the two peers. This is because BFD rx/tx interval values are signaled in the BFD packets while establishing the BFD session.

In 7x50s equipped with SF/CPM 2 or higher, configurable BFD parameters are as follows:

```
bfd <transmit-interval> [receive <receive-interval>] [multiplier <multiplier>] [echo-
receive <echo-interval>] [type <cpm-np>]
no bfd

 <transmit-interval>  : [10..100000] in milliseconds
 <receive-interval>   : [10..100000] in milliseconds
 <multiplier>         : [3..20]
 <echo-interval>      : [100..100000] in milliseconds
 <cpm-np>             : keyword - use CPM network processor
```

Note that it is possible to force the BFD session to be centrally managed by the CPM hardware.

As regards the echo function, it is possible to set the minimum echo receive interval, in milliseconds, for the BFD session. The default value is 100 ms.

If a BFD session is running, it is possible to modify its parameters but to change its type the session must be previously shut down manually. Note that this causes the upper layer protocols bound to it to be brought down as well.

```
configure
   router
      interface PE-2-PE-1
          bfd 10 receive 10 multiplier 3 type cpm-np
      exit
   exit
exit
```

Forcing a centralized session in the case of directly connected peers can be useful when:

- Lower Tx and Rx intervals are requested (up to 10 ms instead of 100 ms supported by local sessions)
- There are no more available local sessions
- Max limit of 500 packet per second per IOM has been reached

The instructions illustrated in following paragraphs are required to complete the configuration and enable BFD.

The BFD session should come up. To verify it, execute a **show router bfd session** command (bound to OSPF in the following example).

```
A:PE1# show router bfd session
===============================================================================
BFD Session
===============================================================================
Interface                     State                 Tx Intvl  Rx Intvl  Multipl
  Remote Address              Protocol              Tx Pkts   Rx Pkts   Type
-------------------------------------------------------------------------------
PE-1-PE-2                     Up (3)                100       100       3
   192.168.1.2                ospf                  165       174       iom
-------------------------------------------------------------------------------
No. of BFD sessions: 1
===============================================================================
A:PE1#
```

If the command gives a negative output, troubleshoot it by firstly checking that the protocol that is bound to it is up: for instance, check the OSPF neighbor adjacency as shown in following example.

```
A:PE-1# show router ospf neighbor
===============================================================================
OSPF Neighbors
===============================================================================
Interface-Name                   Rtr Id          State      Pri  RetxQ   TTL
-------------------------------------------------------------------------------
PE-1-PE-2                        192.0.2.1       Full       1    0       34
-------------------------------------------------------------------------------
...
===============================================================================
A:PE-1#
```

Then check whether a BFD resource limit has been reached (maximum number of local/centralized sessions or maximum number of packet per second per IOM).

If the overloaded limit is the maximum supported number of sessions, the cause is shown by log-id 99. In the reported example, the maximum number of sessions per slot has been reached.

```
A:PE-2# show log log-id 99
===============================================================================
Event Log 99
===============================================================================
Description : Default System Log
Memory Log contents  [size=500   next event=7845  (wrapped)]

7844 2010/10/02 16:43:30.21 UTC MINOR: VRTR #2020 Base 192.168.1.1
BFD Session on node 192.168.1.1 has been deleted.

7843 2010/10/02 16:43:30.21 UTC MAJOR: VRTR #2013 Base Max supported sessions reached
The number of BFD sessions on slot 1 has exceeded 250, constrained by maxSessionsPerSlot”
```

In this case, when one of the running sessions is manually removed or goes down, then the additional configured session will come up. If the limit reached is local (on IOM) it is possible to bring up the session by re-configuring it as centralized, by changing the type.

To check if IOM CPU is able to start more local BFD sessions, execute a **show router BFD session summary** command.

```
A:PE2# show router bfd session summary
=============================
BFD Session Summary
=============================
Termination   Session Count
-----------------------------
central                   0
cpm-np                    1
iom, slot 1             250
iom, slot 2               0
iom, slot 3               0
iom, slot 4               0
iom, slot 5               0
Total                   251
=============================
```

If the **show router bfd session** command reports that the BFD session is down, the check the BFD peer's configuration and state.

The following **log 99** output reports PE-1 logs after a misconfiguration of PE-2 (disabling BFD on the OSPF interface).

As soon as BFD is shutdown on the OSPF interface PE-2-PE-1 of PE-2, the BFD session in PE-1 goes to the down state, then the OSPF adjacency is brought down for approximately 2.8 secs and finally the OSPF state goes back to full, while the BFD session stays in down state.

This state will last until BFD is re-enabled on PE-2 interface.

```
A:PE-1# show log log-id 99
===============================================================================
Event Log 99
===============================================================================
Description : Default System Log
Memory Log contents  [size=500   next event=7  (not wrapped)]

6 2010/10/02 08:47:35.91 UTC WARNING: OSPF #2002 Base VR:  1 OSPFv2 (0)
LCL_RTR_ID 192.0.2.1: Neighbor 192.0.2.2 on PE-1-PE-2 router state changed to full (event
EXC_DONE)

5 2010/10/02 08:47:35.91 UTC MINOR: VRTR #2021 Base 192.168.1.2
BFD: The protocols using BFD session on node 192.168.1.2 have changed.

4 2010/10/02 08:47:33.10 UTC WARNING: OSPF #2002 Base VR:  1 OSPFv2 (0)
LCL_RTR_ID 192.0.2.1: Neighbor 192.0.2.2 on PE-1-PE-2 router state changed to down (event
BFD_DOWN)

3 2010/10/02 08:47:33.10 UTC MINOR: VRTR #2021 Base 192.168.1.2
BFD: The protocols using BFD session on node 192.168.1.2 have changed.
```

```
2 2010/10/02 08:47:33.10 UTC MAJOR: VRTR #2012 Base 192.168.1.2
BFD: Local Discriminator 4009 BFD session to node 192.168.1.2 is down due to noHeartBeat
```

```
A:PE-1# show router bfd session
===============================================================================
BFD Session
===============================================================================
Interface                     State               Tx Intvl  Rx Intvl  Multipl
  Remote Address              Protocols           Tx Pkts   Rx Pkts   Type
-------------------------------------------------------------------------------
PE-1-PE-2                     Down (1)            100       100       3
   192.168.1.2                ospf2               10        0         iom
```

The 2nd column reports the current BFD session state. Possible values are:

- 0 — AdminDown
- 1 — Down
- 2 — Init
- 3 — Up

The **show router bfd session src** <*ip-address*> **detail** command can help in debugging the BFD session.

```
A:PE-1# show router bfd session src 192.168.1.1 detail
===============================================================================
BFD Session
===============================================================================
Remote Address : 192.168.1.2
Admin State    : Up                      Oper State       : Up (3)
Protocols      : ospf2 pim isis static
Rx Interval    : 100                     Tx Interval      : 100
Multiplier     : 3                       Echo Interval    : 0
Recd Msgs      : 24046                   Sent Msgs        : 25723
Up Time        : 0d 00:40:05             Up Transitions   : 1
Down Time      : None                    Down Transitions : 0
                                         Version Mismatch : 0

Forwarding Information

Local Discr    : 4002                    Local State      : Up (3)
Local Diag     : 0 (None)                Local Mode       : Async
Local Min Tx   : 100                     Local Mult       : 3
Last Sent      : 10/08/2010 20:30:27     Local Min Rx     : 100
Type           : iom
Remote Discr   : 4001                    Remote State     : Up (3)
Remote Diag    : 0 (None)                Remote Mode      : Async
Remote Min Tx  : 100                     Remote Mult      : 3
Last Recv      : 10/08/2010 20:30:27     Remote Min Rx    : 100
===============================================================================
```

BFD for IS-IS

The goal of this section is to configure BFD on a network interlink between two 7750 nodes that are IS-IS peers. The topology referred to in this paragraph is shown in Figure 268.

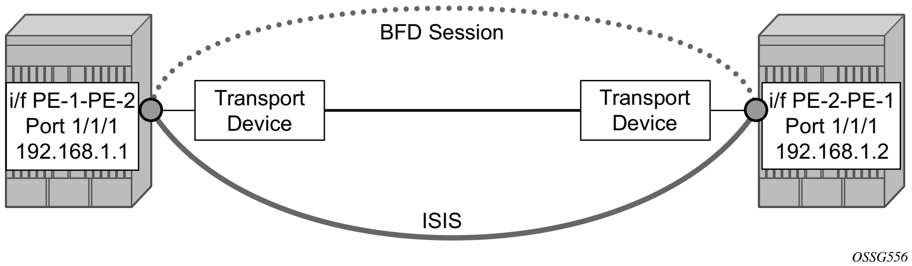

Figure 268: BFD for ISIS

For the base BFD configuration, please refer to BFD Base Parameter Configuration and Troubleshooting on page 1743.

Apply BFD on the ISIS Interfaces.

On PE1:

```
configure
     router
         isis
          interface PE-1-PE-2
              bfd-enable ipv4
          exit
         exit
     exit
exit
```

On PE2:

```
configure
     router
         isis
          interface PE-2-PE-1
              bfd-enable ipv4
          exit
        exit
     exit
exit
```

Finally, verify that the BFD session is operational between PE1 and PE2.

On PE1:

```
A:PE1# show router bfd session
===============================================================================
BFD Session
===============================================================================
Interface                     State                 Tx Intvl  Rx Intvl  Multipl
  Remote Address              Protocol              Tx Pkts   Rx Pkts   Type
-------------------------------------------------------------------------------
PE-1-PE-2                     Up (3)                100       100       3
   192.168.1.2                isis                  165       174       iom
-------------------------------------------------------------------------------
No. of BFD sessions: 1
===============================================================================
A:PE1#
```

On PE2:

```
A:PE2# show router bfd session
===============================================================================
BFD Session
===============================================================================
Interface                     State                 Tx Intvl  Rx Intvl  Multipl
  Remote Address              Protocol              Tx Pkts   Rx Pkts   Type
-------------------------------------------------------------------------------
PE-2-PE-1                     Up (3)                100       100       3
   192.168.1.1                isis                  496       487       iom
-------------------------------------------------------------------------------
No. of BFD sessions: 1
===============================================================================
A:PE2#
```

BFD for OSPF

The goal of this section is to configure BFD on a network interlink between two 7750 nodes that are OSPF peers.

For this scenario, the topology is shown in Figure 269.

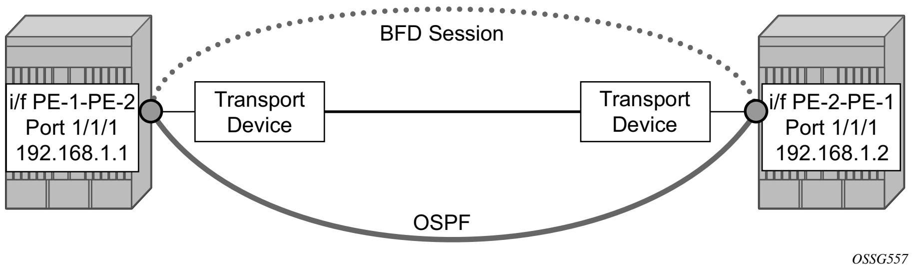

Figure 269: BFD for OSPF

For the base BFD configuration, refer to BFD Base Parameter Configuration and Troubleshooting on page 1743.

Apply BFD on the OSPF Interfaces.

On PE1:

```
configure
     router
         ospf
         interface PE-1-PE-2
             bfd-enable
         exit
         exit
     exit
exit
```

On PE2:

```
configure
     router
         ospf
         interface PE-2-PE-1
             bfd-enable
         exit
         exit
     exit
exit
```

Verify that the BFD session is operational between PE1 and PE2.

On PE1:

```
A:PE1# show router bfd session
===============================================================================
BFD Session
===============================================================================
Interface                     State                 Tx Intvl  Rx Intvl  Multipl
  Remote Address              Protocol              Tx Pkts   Rx Pkts   Type
-------------------------------------------------------------------------------
PE-1-PE-2                     Up (3)                100       100       3
   192.168.1.2                ospf                  170       179       iom
-------------------------------------------------------------------------------
No. of BFD sessions: 1
===============================================================================
A:PE1#
```

On PE2:

```
A:PE2# show router bfd session
===============================================================================
BFD Session
===============================================================================
Interface                     State                 Tx Intvl  Rx Intvl  Multipl
  Remote Address              Protocol              Tx Pkts   Rx Pkts   Type
-------------------------------------------------------------------------------
PE-2-PE-1                     Up (3)                100       100       3
   192.168.1.1                ospf                  501       492       iom
-------------------------------------------------------------------------------
No. of BFD sessions: 1
===============================================================================
A:PE2#
```

BFD for PIM

Since the implementation of PIM uses an Interior Gateway Protocol (IGP) in order to determine its Reverse Path Forwarding (RPF) tree, BFD configuration to support PIM will require BFD configuration of both the IGP protocol and the PIM protocol. Let's assume that IGP protocol is OSPF and that the starting configuration is as described in the previous section.

In this paragraph, configure and enable BFD for PIM on the same interfaces that were previously configured with BFD for OSPF, in reference to the topology shown in Figure 270.

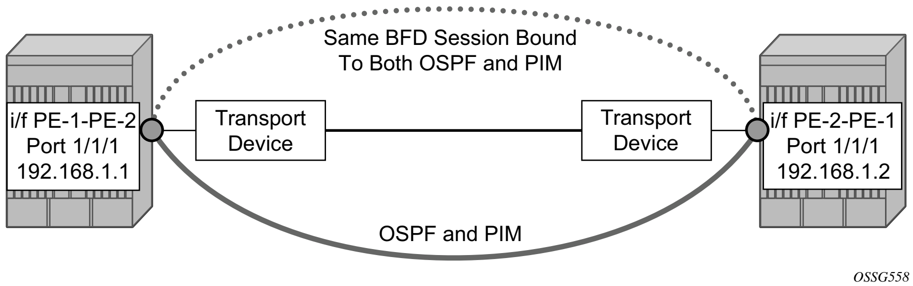

Figure 270: BFD for OSPF and PIM

Since BFD has been already configured on the router interfaces, let's start by applying BFD on the PIM Interface.

On PE1:

```
configure
    router
        pim
         interface PE-1-PE-2
            bfd-enable
         exit
        exit
    exit
exit
```

On PE2:

```
configure
    router
        pim
```

```
          interface PE-2-PE-1
             bfd-enable
          exit
         exit
     exit
exit
```

The final step is to verify whether the BFD Session is operational between PE1 and PE2 for PIM.

On PE1:

```
A:PE1# show router bfd session
===============================================================================
BFD Session
===============================================================================
Interface                     State                 Tx Intvl  Rx Intvl  Multipl
  Remote Address              Protocol              Tx Pkts   Rx Pkts   Type
-------------------------------------------------------------------------------
PE-1-PE-2                     Up (3)                100       100       3
   192.168.1.2                ospf2 pim             3874      3845      iom
-------------------------------------------------------------------------------
No. of BFD sessions: 1
===============================================================================
A:PE1#
```

On PE2:

```
A:PE2# show router bfd session
===============================================================================
BFD Session
===============================================================================
Interface                     State                 Tx Intvl  Rx Intvl  Multipl
  Remote Address              Protocol              Tx Pkts   Rx Pkts   Type
-------------------------------------------------------------------------------
PE-1-PE-2                     Up (3)                100       100       3
   192.168.1.1                ospf2 pim             3137      3145      iom
-------------------------------------------------------------------------------
No. of BFD sessions: 1
===============================================================================
A:PE2#
```

BFD for Static Routes

The following procedures will go through the necessary steps to configure the base level BFD configuration and then apply BFD to the static routes between PE1 and PE2, referring to topology shown in Figure 271.

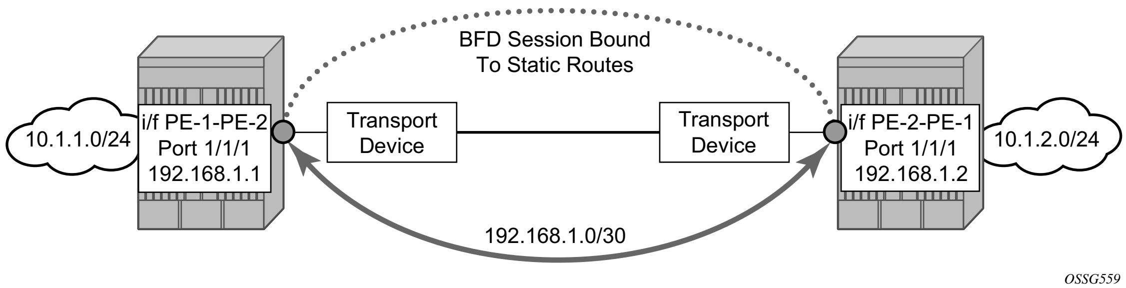

Figure 271: BFD for Static Routes

First, create the static routes for the remote networks both in PE-1 and PE-2.

On PE1:

```
configure
     router
         static-route 10.1.2.0/24 next-hop 192.168.1.2
        exit
exit
```

On PE2:

```
configure
     router
         static-route 10.1.1.0/24 next-hop 192.168.1.1
        exit
exit
```

Next, verify that static routes are populated in the routing table.

On PE1:

```
A:PE1# show router route-table
===============================================================================
Route Table (Router: Base)
===============================================================================
Dest Prefix                                   Type    Proto    Age         Pref
       Next Hop[Interface Name]                                     Metric
-------------------------------------------------------------------------------
10.1.2.0/24                                   Remote  Static   00h20m55s   5
       192.168.1.2                                                  1
```

On PE2:

```
A:PE2# show router route-table
===============================================================================
Route Table (Router: Base)
===============================================================================
Dest Prefix                                   Type    Proto    Age         Pref
       Next Hop[Interface Name]                                     Metric
-------------------------------------------------------------------------------
10.1.1.0/24                                   Remote  Static   00h21m15s   5
       192.168.1.1                                                  1
```

The next step is to configure the base level BFD on PE1 and PE2.

Refer to paragraph BFD Base Parameter Configuration and Troubleshooting on page 1743.

Then apply BFD to the static routing entries using the BFD interfaces as next-hop.

On PE1:

```
configure
     router
         static-route 10.1.2.0/24 next-hop 192.168.1.2 bfd-enable
        exit
exit
```

On PE2:

```
configure
     router
         static-route 10.1.1.0/24 next-hop 192.168.1.1 bfd-enable
        exit
exit
```

Note that BFD cannot be enabled if the next hop is indirect or the **blackhole** keyword is specified.

Finally, show the BFD session status.

On PE1:

```
A:PE1# show router bfd session
===============================================================================
BFD Session
===============================================================================
Interface                      State                 Tx Intvl  Rx Intvl  Multipl
  Remote Address               Protocol              Tx Pkts   Rx Pkts   Type
-------------------------------------------------------------------------------
PE-1-PE-2                      Up (3)                100       100       3
   192.168.1.2                 static                699       661       iom
-------------------------------------------------------------------------------
No. of BFD sessions: 1
===============================================================================
```

On PE2:

```
A:PE2# show router bfd session
===============================================================================
BFD Session
===============================================================================
Interface                      State                 Tx Intvl  Rx Intvl  Multipl
  Remote Address               Protocol              Tx Pkts   Rx Pkts   Type
-------------------------------------------------------------------------------
PE-2-PE-1                      Up (3)                100       100       3
   192.168.1.1                 static                691       729       iom
-------------------------------------------------------------------------------
No. of BFD sessions: 1
===============================================================================
```

BFD for IES

The goal of this section is to configure BFD for one IES service over a spoke SDP.

The IES service is configured in both 7750 nodes, PE1 and PE2, and their interfaces are connected by spoke SDP's. The topology is shown in Figure 272.

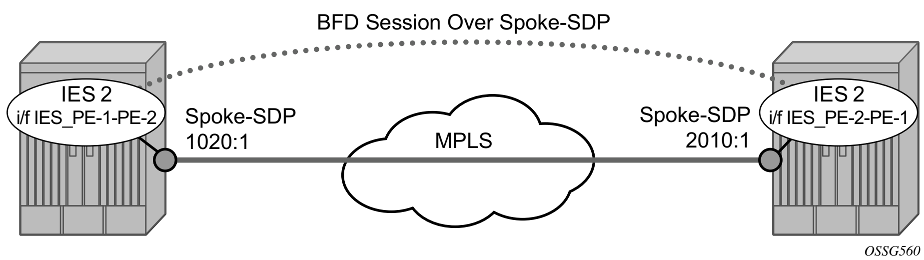

Figure 272: BFD for IES over Spoke SDP

Note that in this scenario BFD is run between the IES interfaces independent of the SDP/LSP paths.

The first step is to configure the IES service on both nodes.

On PE-1:

```
configure
      service
          ies 2 customer 1 create
             interface IES_PE-1-PE-2  create
                address 192.168.3.1/30
                spoke-sdp 1020:1 create
                exit
             exit
             no shutdown
          exit
      exit
exit
```

On PE-2:

```
configure
      service
          ies 2 customer 1 create
            interface IES_PE-2-PE-1  create
```

```
                address 192.168.3.2/30
                spoke-sdp 2010:1 create
                exit
            exit
            no shutdown
          exit
      exit
exit
```

The next step is to add the IES interfaces to the OSPF area domain.

```
On PE-1:
configure
     router
        ospf
            traffic-engineering
            area 0.0.0.0
                interface IES-PE-1-PE-2
                exit
            exit
        exit
     exit
exit
```

On PE-2:

```
configure
     router
        ospf
            traffic-engineering
            area 0.0.0.0
                interface IES-PE-2-PE-1
                exit
            exit
        exit
     exit
exit
```

Then verify that OSPF and the services are up using show commands on both routers.

On PE-1:

```
A:PE-1# show service id 1 base
===============================================================================
Service Basic Information
===============================================================================
Service Id        : 2                   Vpn Id            : 0
Service Type      : IES
Customer Id       : 1
Last Status Change: 09/30/2010 08:09:22
Last Mgmt Change  : 09/30/2010 08:08:31
Admin State       : Up                  Oper State        : Up
SAP Count         : 0
...
===============================================================================
A:PE-1#

A:PE-1# show router ospf neighbor
===============================================================================
OSPF Neighbors
===============================================================================
Interface-Name                   Rtr Id          State      Pri  RetxQ   TTL
-------------------------------------------------------------------------------
IES-PE-1-PE-2                    192.0.2.2       Full       1    0       34
-------------------------------------------------------------------------------
===============================================================================
A:PE-1#
```

On PE-2:

```
A:PE-2# show service id 2 base
===============================================================================
Service Basic Information
===============================================================================
Service Id        : 2                   Vpn Id            : 0
Service Type      : IES
Customer Id       : 1
Last Status Change: 09/30/2010 08:16:50
Last Mgmt Change  : 09/30/2010 08:16:50
Admin State       : Up                  Oper State        : Up
SAP Count         : 0
...
===============================================================================
A:PE-2#

A:PE-2# show router ospf neighbor
===============================================================================
OSPF Neighbors
===============================================================================
Interface-Name                   Rtr Id          State      Pri  RetxQ   TTL
-------------------------------------------------------------------------------
IES-PE-2-PE-1                    192.0.2.1       Full       1    0       33
-------------------------------------------------------------------------------
...
===============================================================================
A:PE-2#
```

Then configure BFD on the IES interfaces.

On PE-1:

```
configure service ies 2
            interface IES-PE-1-PE-2
                bfd 100 receive 100 multiplier 3
            exit
            no shutdown
exit
```

On PE-2:

```
configure service ies 2
            interface IES-PE-2-PE-1
                bfd 100 receive 100 multiplier 3
            exit
            no shutdown
exit
```

Finally, enable BFD on the interfaces under OSPF area 0.

On PE-1:

```
A:PE-1# configure router ospf area 0.0.0.0 interface IES-PE-1-PE-2 bfd-enable
```

On PE-2:

```
A:PE-2# configure router ospf area 0.0.0.0 interface IES-PE-2-PE-1 bfd-enable
```

Note that in case of BFD over spoke SDP, a centralized BFD session is created even if a physical link exists between the two nodes. In fact, the next output shows that BFD session type is cpm-np. This is because the spoke SDP is terminated at the CPM. This is also true for BFD running over LAG bundles.

The cpm-np type only exists in 7x50 SR/ESS systems equipped with SF/CPM 2 or higher. When other network elements run centralized BFD sessions like this one, the BFD type is shown as **central**.

```
A:PE-1# show router bfd session
===============================================================================
BFD Session
===============================================================================
Interface                     State                 Tx Intvl  Rx Intvl  Multipl
  Remote Address              Protocols             Tx Pkts   Rx Pkts   Type
-------------------------------------------------------------------------------
IES-PE-1-PE-2                 Up (3)                100       100       3
   192.168.3.2                ospf2                 N/A       N/A       cpm-np
-------------------------------------------------------------------------------
No. of BFD sessions: 1
```

```
A:PE-2# show router bfd session
===============================================================================
BFD Session
===============================================================================
Interface                     State                 Tx Intvl  Rx Intvl  Multipl
  Remote Address              Protocols             Tx Pkts   Rx Pkts   Type
-------------------------------------------------------------------------------
IES-PE-2-PE-1                 Up (3)                100       100       3
   192.168.3.1                ospf2                 N/A       N/A       cpm-np
-------------------------------------------------------------------------------
No. of BFD sessions: 1
```

Note that in the case of centralized BFD sessions, transmitted and received packet counters are not shown.

BFD for RSVP

The goal of this section is to configure BFD between two RSVP interfaces configured in two 7750 nodes.

For this scenario, the topology is shown in Figure 273.

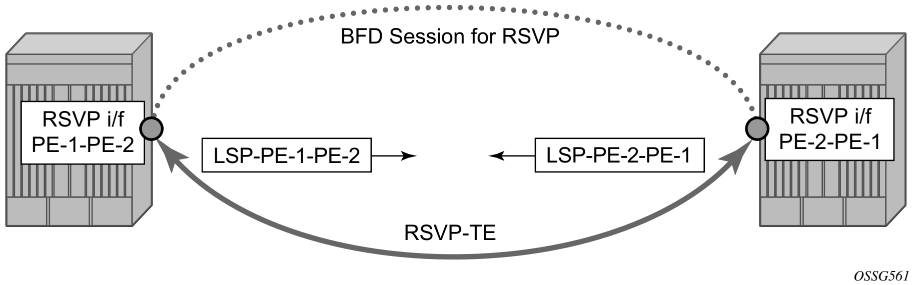

Figure 273: BFD for RSVP

To enable the BFD session between the two RSVP peers, the user should follow these steps:

First, configure BFD on interfaces between PE-1 and PE-2 as described in BFD Base Parameter Configuration and Troubleshooting on page 1743.

Next, configure MPLS, creating the path, the LSP and the interfaces within MPLS (and RSVP).

On PE-1:

```
configure router
        mpls
            interface system
            exit
            interface PE-1-PE-2
            exit
        exit
        rsvp
            interface system
            exit
            interface PE-1-PE-2
            exit
            no shutdown
        exit
        mpls
            path dyn
                no shutdown
            exit
            lsp LSP-PE-1-PE-2
                to 192.0.1.2
                cspf
```

```
                primary dyn
                exit
                no shutdown
            exit
            no shutdown
        exit
exit
```

On PE-2:

```
configure router
        mpls
            interface system
            exit
            interface PE-2-PE-1
            exit
        exit
        rsvp
            interface system
            exit
            interface PE-2-PE-1
            exit
            no shutdown
        exit
        mpls
            path dyn
                no shutdown
            exit
            lsp LSP-PE-2-PE-1
                to 192.0.1.1
                cspf
                primary dyn
                exit
                no shutdown
            exit
            no shutdown
        exit
exit
```

Next, verify that the RSVP session is up.

```
A:PE-1# show router rsvp session
===============================================================================
RSVP Sessions
===============================================================================
From            To              Tunnel LSP   Name                           State
                                ID     ID
-------------------------------------------------------------------------------
192.0.2.2       192.0.2.1       2      516   LSP-PE-2-PE-1::dyn             Up
192.0.2.1       192.0.2.2       1      61446 LSP-PE-1-PE-2::dyn             Up
-------------------------------------------------------------------------------
Sessions : 2
===============================================================================
A:PE-1#
```

Then, apply BFD on the RSVP Interfaces.

On PE1:

```
configure
     router
            rsvp
              interface PE-1-PE-2
                  bfd-enable
              exit
              no shutdown
            exit
     exit
exit
```

On PE2:

```
configure
     router
            rsvp
              interface PE-2-PE-1
                  bfd-enable
              exit
              no shutdown
            exit
     exit
exit
```

Finally, verify that the BFD session is operational between PE1 and PE2.

On PE1:

```
===============================================================================
BFD Session
===============================================================================
Interface                     State                 Tx Intvl  Rx Intvl  Multipl
  Remote Address              Protocols             Tx Pkts   Rx Pkts   Type
-------------------------------------------------------------------------------
PE-1-PE-2                     Up (3)                100       100       3
   192.168.1.2                rsvp                  31515     31506     iom
-------------------------------------------------------------------------------
No. of BFD sessions: 1
===============================================================================
```

On PE2:

```
===============================================================================
BFD Session
===============================================================================
Interface                     State                 Tx Intvl  Rx Intvl  Multipl
  Remote Address              Protocols             Tx Pkts   Rx Pkts   Type
-------------------------------------------------------------------------------
PE-2-PE-1                     Up (3)                100       100       3
   192.168.1.1                rsvp                  31563     31572     iom
-------------------------------------------------------------------------------
No. of BFD sessions: 1
===============================================================================
```

BFD for T-LDP

BFD tracking of an LDP session associated with a T-LDP adjacency allows for faster detection of the liveliness of the session by registering the transport address of an LDP session with a BFD session.

The goal of this paragraph is to configure BFD for T-LDP, referring to the scheme shown in Figure 274.

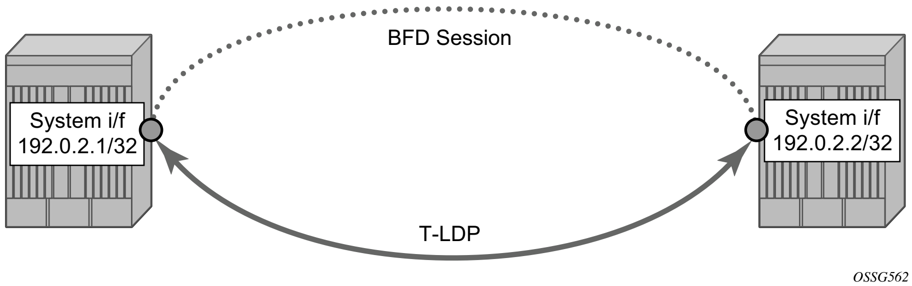

Figure 274: BFD for T-LDP

The parameters used for the BFD session are configured under the loopback interface corresponding to the LSR-ID (by default, the LSR-ID matches the system interface address).

```
configure
    router
       interface system
            address 192.0.2.1/32
            bfd 3000 receive 3000 multiplier 3
       exit
    exit
exit
```

By enabling BFD for a selected targeted session, the state of that session is tied to the state of the underlying BFD session between the two nodes.

When using BFD over other links with the ability to reroute, such as spoke-SDPs, the interval and multiplier values configuring BFD should be set to allow sufficient time for the underlying network to re-converge before the associated BFD session expires. A general rule of thumb should be that the expiration time (interval * multiplier) is three times the convergence time for the IGP network between the two endpoints of the BFD session.

Before enabling BFD, ensure that the T-LDP session is up.

```
On PE-1:
B:PE-1# show router ldp session
==============================================================================
LDP Sessions
==============================================================================
Peer LDP Id        Adj Type   State         Msg Sent  Msg Recv  Up Time
------------------------------------------------------------------------------
192.0.2.2    Targeted   Established   35        41        0d 00:02:50
------------------------------------------------------------------------------
==============================================================================
```

On PE-2:

```
B:PE-2# show router ldp session
==============================================================================
LDP Sessions
==============================================================================
Peer LDP Id        Adj Type   State         Msg Sent  Msg Recv  Up Time
------------------------------------------------------------------------------
192.0.2.1    Targeted   Established   27        23        0d 00:01:32
------------------------------------------------------------------------------
==============================================================================
```

Then, enable the BFD session.

```
configure
     router
         ldp
            targeted-session
                peer 192.0.2.2
                    bfd-enable
                exit
            exit
        exit
     exit
exit
```

Note that the loopback interface can be used to source BFD sessions to many peers in the network.

Finally, check that the BFD session is up.

On PE-1:

```
A:PE-1# show router bfd session
===============================================================================
BFD Session
===============================================================================
Interface                     State                 Tx Intvl  Rx Intvl  Multipl
  Remote Address              Protocols             Tx Pkts   Rx Pkts   Type
-------------------------------------------------------------------------------
system                        Up (3)                100       100       3
```

```
   192.0.2.2                  ldp                   N/A       N/A       cpm-np
```

On PE-2:

```
A:PE-1# show router bfd session
===============================================================================
BFD Session
===============================================================================
Interface                     State                 Tx Intvl  Rx Intvl  Multipl
  Remote Address              Protocols             Tx Pkts   Rx Pkts   Type
-------------------------------------------------------------------------------
system                        Up (3)                100       100       3
   192.0.2.1                  ldp                   N/A       N/A       cpm-np
```

When the T-LDP session comes up, a centralized BFD session is always created even if the local interface has a direct link to the peer.

BFD Support of OSPF PE-CE Adjacencies

This feature, introduced with Release 8.0, extends BFD support to OSPF within a VPRN context when OSPF is used as the PE-CE protocol. In this section, the topology shown in Figure 275.

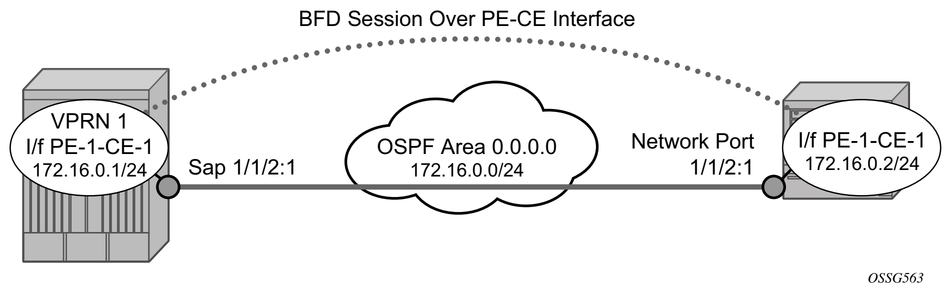

Figure 275: BFD for OSPF PE-CE I/F

First, configure the VPRN service interface PE-1-CE-1 on PE-1 with BFD parameters.

```
config
      service
           vprn 1 customer 1 create
              route-distinguisher 1:1
              vrf-target target:1:1
              interface PE-1-CE-1  create
                  address 172.16.0.1/24
                  bfd 100 receive 100 multiplier 3
                  sap 1/1/1:1 create
                  exit
              exit
              ospf
                  area 0.0.0.0
                      interface PE-1-CE-1
                      exit
                  exit
              exit
              no shutdown
           exit
      exit
exit
```

Next, configure the router interface on CE-1 and add it to the OSPF area 0 domain.

```
configure
     router
        interface CE-1-PE-1
            address 172.16.0.2/24
            port 1/1/1:1
            bfd 100 receive 100 multiplier 3
        exit
        ospf
            area 0.0.0.0
                interface CE-1-PE-1
                exit
            exit
        exit
     exit
exit
```

Then, ensure that OSPF adjacency is up.

On PE-1:

```
A:PE-1>config>service>vprn# show router 1 ospf neighbor
===============================================================================
OSPF Neighbors
===============================================================================
Interface-Name                   Rtr Id          State      Pri  RetxQ   TTL
-------------------------------------------------------------------------------
PE-1-CE-1                        192.0.2.5       Full       1    2       33
-------------------------------------------------------------------------------
No. of Neighbors: 1
===============================================================================
```

On CE-1:

```
A:CE-1# show router ospf neighbor
===============================================================================
OSPF Neighbors
===============================================================================
Interface-Name                   Rtr Id          State      Pri  RetxQ   TTL
-------------------------------------------------------------------------------
CE-1-PE-1                        192.0.2.1       Full       1    0       31
-------------------------------------------------------------------------------
No. of Neighbors: 1
===============================================================================
```

Then, enable BFD on the PE-1-CE-1 interface on PE-1.

```
configure service vprn 1 ospf area 0.0.0.0 interface PE-1-CE-1 bfd-enable
```

Enable BFD on the CE-1-PE-1 interface on CE-1.

```
configure router ospf area 0.0.0.0 interface CE-1-PE-1 bfd-enable
```

Finally, check that the BFD sessions are up in both PE-1 and CE-1.

```
A:PE-1# show router 1 bfd session
===============================================================================
BFD Session
===============================================================================
Interface                     State                 Tx Intvl  Rx Intvl  Multipl
  Remote Address              Protocols             Tx Pkts   Rx Pkts   Type
-------------------------------------------------------------------------------
PE-1-CE-1                     Up (3)                100       100       3
   172.16.0.2                 ospf2                 6331      6340      iom
-------------------------------------------------------------------------------
No. of BFD sessions: 1


A:CE-1# show router bfd session
===============================================================================
BFD Session
===============================================================================
Interface                     State                 Tx Intvl  Rx Intvl  Multipl
  Remote Address              Protocols             Tx Pkts   Rx Pkts   Type
-------------------------------------------------------------------------------
CE-1-PE-1                     Up (3)                100       100       3
   172.16.0.1                 ospf2                 6691      6682      iom
-------------------------------------------------------------------------------
No. of BFD sessions: 1
```

BFD within IPSec Tunnels

The ability to assign a BFD session to a given static LAN-to-LAN IPSec tunnel that provides heart-beat mechanism for fast failure detection has been introduced in Release.8.0.

IPSec needs a Multi-service Integrated Service Adapter (MS-ISA) installed, so this scenario is only applicable to 7750 SR-7/12 equipped with IOM-2 or 3.

In this section, the topology is shown in Figure 276.s

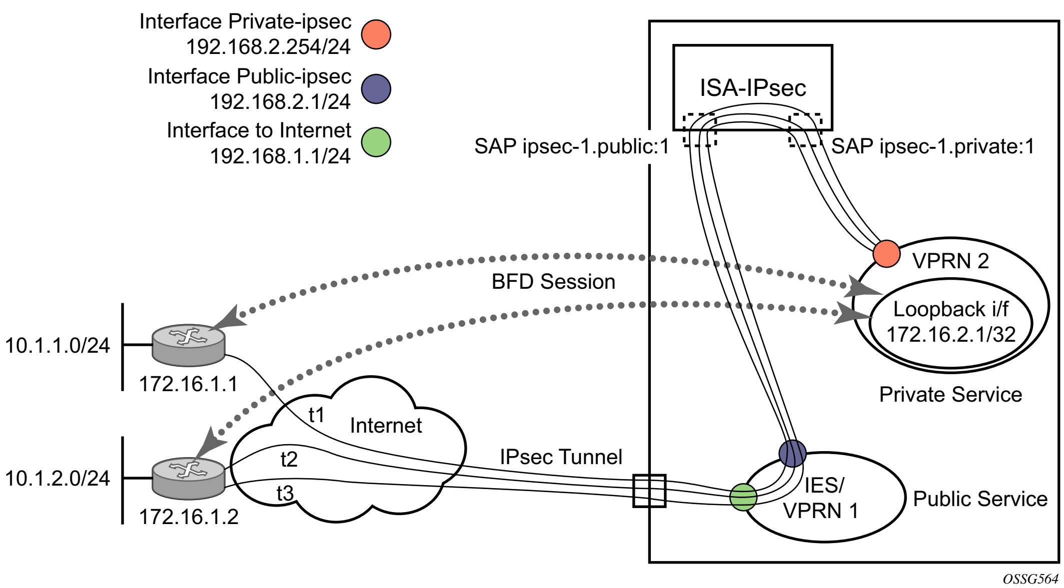

Figure 276: BFD Sessions within IPSec Tunnels

The first step is to configure MS-ISA card as **type isa-tunnel**.

```
configure
    card 1
        card-type iom3-xp
        mda 1
            mda-type isa-tunnel
        exit
        mda 2
            mda-type m10-1gb-sfp-b
        exit
    exit
exit
```

Next, instantiate the tunnels t1, t2 and t3 from the private service (in this example, VPRN 2) to the peers passing through the public service (in this example VPRN 1, but it could be instead an IES).

Since the configuration of IPSec tunnels is out of the scope of this section, only relevant command lines are reported to configure the interfaces shown in Figure 276.

```
configure service
        vprn 1 customer 1 create
            route-distinguisher 1:1
            interface toInternet  create
                address 192.168.1.1/24
                sap 1/2/1 create
                exit
            exit
            interface public-ipsec  create
                address 192.168.2.1/24
                sap tunnel-1.public:1 create
                exit
            exit
            no shutdown
        exit
        vprn 2 customer 1 create
            ipsec
                security-policy 1 create
                    entry 10 create
                        local-ip 192.168.3.1/32
                        remote-ip any
                    exit
                exit
            exit
            route-distinguisher 1:2
            interface private-ipsec tunnel create
                sap tunnel-1.private:1 create
                    ipsec-tunnel t1 create
                     local-gateway-address 192.168.2.254 peer 172.16.1.1 delivery-service 1
                        exit
                    exit
                    ipsec-tunnel t2 create
                     local-gateway-address 192.168.2.254 peer 172.16.1.2 delivery-service 1
                        exit
                    exit
                    ipsec-tunnel t3 create
                     local-gateway-address 192.168.2.254 peer 172.16.1.2 delivery-service 1
                        exit
                    exit
                 exit
            exit
            interface loop  create
                address 172.16.2.1/32
                loopback
            exit
            static-route 10.1.1.0/24 ipsec-tunnel t1
            static-route 10.1.2.0/24 ipsec-tunnel t2 metric 1
            static-route 10.1.2.0/24 ipsec-tunnel t3 metric 5
            no shutdown
```

Then configure the BFD parameters within loopback interface loop (refer to BFD Base Parameter Configuration and Troubleshooting on page 1743).

```
configure service vprn 2
            interface loop
                bfd 100 receive 100 multiplier 3
            exit
exit

And finally enable BFD within the tunnels.

configure service
         vprn 2
            interface private-ipsec tunnel
                sap tunnel-1.private:1
                    ipsec-tunnel t1
                        bfd-enable service 2 interface loop dst-ip 172.16.1.1
                    exit
                    ipsec-tunnel t2
                        bfd-enable service 2 interface loop dst-ip 172.16.1.2
                        bfd-designate
                    exit
                    ipsec-tunnel t3
                        bfd-enable service 2 interface loop dst-ip 172.16.1.2
exit all
```

The BFD-enable parameters are as follows:

- **service** <*service-id*> — Specifies the service-id where the BFD session resides.
- **interface** <*interface-name*> — Specifies the name of the interface used by the BFD session.
- **dst-ip** <*ip-address*> — Specifies the destination address to be used for the BFD session.

The following statements are to be taken into consideration to correctly configure BFD in this environment:

- BFD over IPSec sessions are centralized, managed by the hardware on the CPM.
- Only BFD over static lan-to-lan tunnel is supported in Release 8.0 (not dynamic).
- Only one BFD session is allowed between a given source/destination address pair.
- Each tunnel can be associated to only one BFD session but multiple tunnels can be associated to the same BFD session.
- In case of multiple tunnels sharing the same BFD session, one IPSec tunnel carries BFD traffic: the BFD-DESIGNATED tunnel.

Referring to Figure 276 and to the above configuration, the tunnels t2 and t3 share the same BFD-session. Tunnel t2 is the bfd-designated tunnel, the BFD session runs within it and the other tunnel t3 shares its BFD session. If the BFD session goes down, the system will bring down both the designated tunnel t2 and the associated tunnel t3.

The state machine in Figure 277 shows the decision process in case of shared BFD sessions.

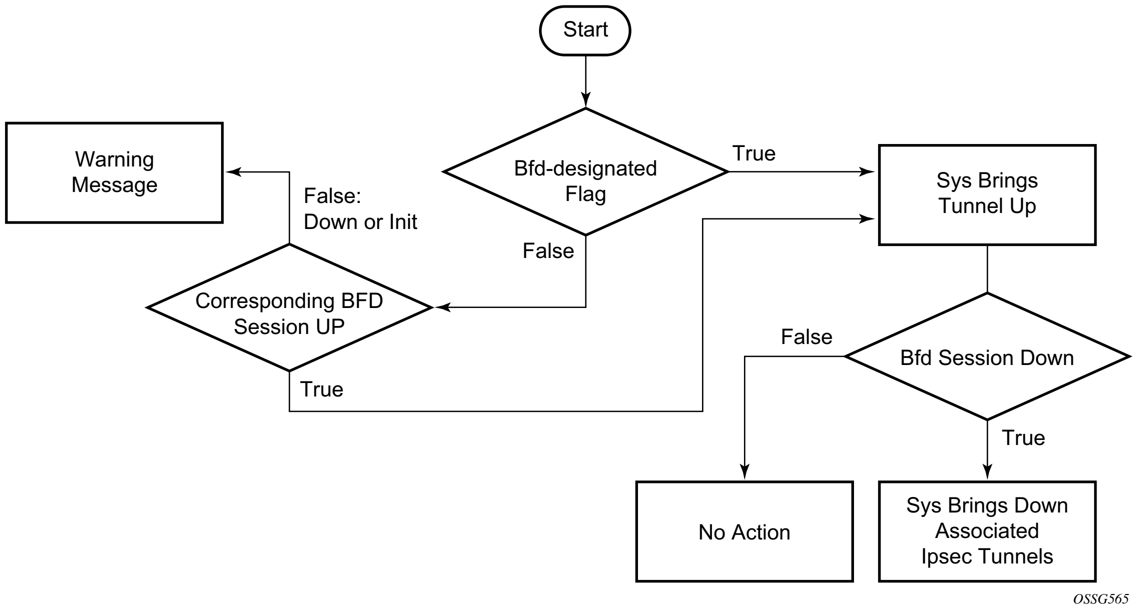

Figure 277: Logic for Shared BFD Sessions

BFD for VRRP

This feature assigns a BFD session to provide a heart-beat mechanism for the given VRRP/SRRP instance. It should be noted that there can be only one BFD session assigned to any given VRRP/SRRP instance, but there can be multiple SRRP/VRRP sessions using the same BFD session.

In this section, the topology is shown in Figure 278.

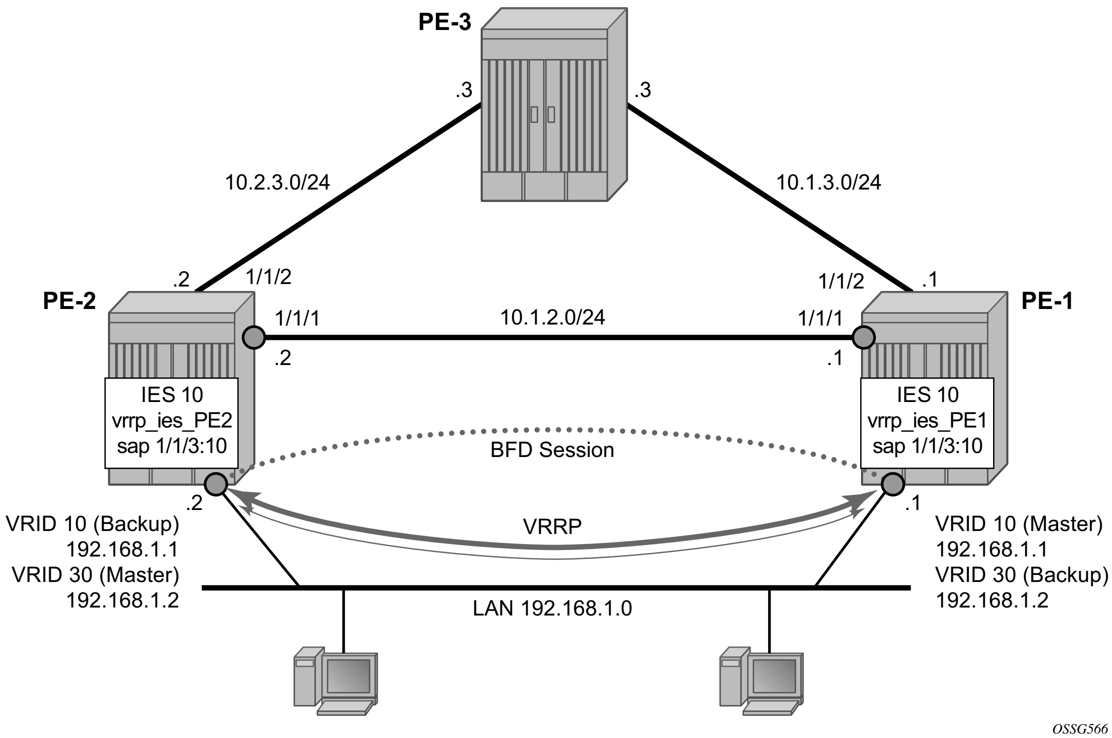

Figure 278: BFD for VRRP

First, create the LAN subnet. Two PE routers are connected by IES or VPRN services (in following examples IES 10 is created in both routers).

On PE-1:

```
configure service ies 10 customer 1 create
            interface vrrp_ies_PE1 create
                address 192.168.1.1/24
                sap 1/1/3:10 create
                exit
            exit
            no shutdown
```

```
        exit
```

On PE-2:

```
configure service ies 10 customer 1 create
            interface vrrp_ies_PE2 create
                address 192.168.1.2/24
                sap 1/1/3:10 create
                exit
            exit
            no shutdown
        exit
```

Verify that the IES services are operational (**show service service-using**) and verify that you can ping the remote interface IP address.

Next, configure the VRRP parameters for both PE-1 and PE-2, enable VRRP on the IES interface that connects to the 192.168.1.0/24 subnet.

In this section, the configurations are shown for the VRRP owner mode for master but any other scenario for VRRP can be configured (non owner mode for master).

In the following examples two VRRP instances are created on the 192.168.1.0/24 subnet:

VRID = 10 Master (owner) = PE-1
Backup = PE-2
VRRP IP = 192.168.1.1
VRID = 30 Master (owner) = PE-2
Backup = PE-1
VRRP IP = 192.168.1.2
Host 1 is configured with default gateway = 192.168.1.1
Host 2 is configured with default gateway = 192.168.1.2

On PE-1:

```
configure service ies 10 interface vrrp_ies_PE1
                vrrp 10 owner
                    backup 192.168.1.1
                exit
                vrrp 30
                    backup 192.168.1.2
                    ping-reply
                    telnet-reply
                    ssh-reply
                exit
```

On PE-2:

```
configure service ies 10 interface vrrp_ies_PE2
                vrrp 10
                    backup 192.168.1.1
                    ping-reply
                    telnet-reply
                    ssh-reply
                exit
                vrrp 30 owner
                    backup 192.168.1.2
                exit
```

To bind the VRRP instances with a BFD session, add the following command under any VRRP instance: **bfd-enable** *service-id* **interface** *interface-name* **dst-ip** *ip-address*.

Note that the IES service-id must be declared where the interface is configured.

On PE-1:

```
configure service ies 10 interface vrrp_ies_PE1
                vrrp 10 owner
                bfd-enable 10 interface vrrp_ies_PE1 dst-ip 192.168.1.2
                exit
                vrrp 30
                bfd-enable 10 interface vrrp_ies_PE1 dst-ip 192.168.1.2
                exit
```

On PE-2:

```
configure service ies 10 interface vrrp_ies_PE2
                vrrp 10 owner
                bfd-enable 10 interface vrrp_ies_PE2 dst-ip 192.168.1.1
                exit
                vrrp 30
                bfd-enable 10 interface vrrp_ies_PE2 dst-ip 192.168.1.1
                exit
```

The parameters used for the BFD are set by the BFD command under the IP interface.

Note that unlike the previous scenarios, the user can enter the commands above, enabling the BFD session, even if the specified interface (vrrp_ies_PE1) has not been configured with BFD parameters.

If it has not been configured yet, the BFD session will be initiated only after the following configuration.

On PE-1:

```
configure service ies 10 interface vrrp_ies_PE1
                bfd 1000 receive 1000 multiplier 3
```

On PE-2:

```
configure service ies 10 interface vrrp_ies_PE2
                bfd 1000 receive 1000 multiplier 3
```

Finally, verify that the BFD session is up (for instance on PE-1):

```
A:PE1>show router bfd session src 192.168.1.1 detail
===============================================================================
BFD Session
===============================================================================
Remote Address : 192.168.1.2
Admin State    : Up                         Oper State       : Up (3)
Protocols      : vrrp
Rx Interval    : 100                        Tx Interval      : 100
Multiplier     : 3                          Echo Interval    : 0
Recd Msgs      : 7404                       Sent Msgs        : 7412
Up Time        : 0d 00:04:26                Up Transitions   : 2
Down Time      : None                       Down Transitions : 1
                                            Version Mismatch : 0

Forwarding Information

Local Discr    : 4006                       Local State      : Up (3)
Local Diag     : 1 (Detect time expired)    Local Mode       : Async
Local Min Tx   : 100                        Local Mult       : 3
Last Sent      : 12/14/2010 17:44:34        Local Min Rx     : 100
Type           : iom
Remote Discr   : 4003                       Remote State     : Up (3)
Remote Diag    : 1 (Detect time expired)    Remote Mode      : Async
Remote Min Tx  : 100                        Remote Mult      : 3
Last Recv      : 12/14/2010 17:44:34        Remote Min Rx    : 100
===============================================================================
```

This session is shared by all the VRRP instances configured between the specified interfaces.

When BFD is configured in a VRRP instance, the following command gives details of BFD related to every instance:

```
show router vrrp instance interface vrrp_ies_PE1
===============================================================================
VRRP Instances for interface vrrp_ies_PE1
===============================================================================
-------------------------------------------------------------------------------
VRID 10
```

```
-------------------------------------------------------------------------------
Owner               : Yes                  VRRP State        : Master
Primary IP of Master: 192.168.1.1 (Self)
Primary IP          : 192.168.1.1          Standby-Forwarding: Disabled
VRRP Backup Addr    : 192.168.1.1
Admin State         : Up                   Oper State        : Up
Up Time             : 12/14/2010 16:47:47 Virt MAC Addr     : 00:00:5e:00:01:0a
Auth Type           : None
Config Mesg Intvl   : 1                    In-Use Mesg Intvl : 1
Base Priority       : 255                  In-Use Priority   : 255
Init Delay          : 0                    Init Timer Expires: 0.000 sec
Creation State      : Active
-------------------------------------------------------------------------------
BFD Interface
-------------------------------------------------------------------------------
Service ID          : 10
Interface Name      : vrrp_ies_PE1
Src IP              : 192.168.1.1
Dst IP              : 192.168.1.2
Session Oper State  : connected
-------------------------------------------------------------------------------
Master Information
-------------------------------------------------------------------------------
Primary IP of Master: 192.168.1.1 (Self)
Addr List Mismatch  : No                   Master Priority   : 255
Master Since        : 12/14/2010 16:47:47
-------------------------------------------------------------------------------
Masters Seen (Last 32)
-------------------------------------------------------------------------------
Primary IP of Master   Last Seen             Addr List Mismatch     Msg Count
-------------------------------------------------------------------------------
192.168.1.1            12/14/2010 16:47:47   No                             0
192.168.1.2            12/14/2010 17:39:57   No                             5
-------------------------------------------------------------------------------
Statistics
-------------------------------------------------------------------------------
Become Master       : 7                    Master Changes    : 7
Adv Sent            : 347577               Adv Received      : 5
Pri Zero Pkts Sent  : 6                    Pri Zero Pkts Rcvd: 0
Preempt Events      : 0                    Preempted Events  : 0
Mesg Intvl Discards : 0                    Mesg Intvl Errors : 0
Addr List Discards  : 0                    Addr List Errors  : 0
Auth Type Mismatch  : 0                    Auth Failures     : 0
Invalid Auth Type   : 0                    Invalid Pkt Type  : 0
IP TTL Errors       : 0                    Pkt Length Errors : 0
Total Discards      : 0

-------------------------------------------------------------------------------
VRID 30
-------------------------------------------------------------------------------
Owner               : No                   VRRP State        : Backup
Primary IP of Master: 192.168.1.2 (Other)
Primary IP          : 192.168.1.1          Standby-Forwarding: Disabled
VRRP Backup Addr    : 192.168.1.2
Admin State         : Up                   Oper State        : Up
Up Time             : 12/14/2010 17:39:49 Virt MAC Addr     : 00:00:5e:00:01:1e
Auth Type           : None
Config Mesg Intvl   : 1                    In-Use Mesg Intvl : 1
Master Inherit Intvl: No
```

```
Base Priority       : 100               In-Use Priority   : 100
Policy ID           : n/a               Preempt Mode      : Yes
Ping Reply          : Yes               Telnet Reply      : Yes
SSH Reply           : Yes               Traceroute Reply  : No
Init Delay          : 0                 Init Timer Expires: 0.000 sec
Creation State      : Active
-------------------------------------------------------------------------------
BFD Interface
-------------------------------------------------------------------------------
Service ID          : 10
Interface Name      : vrrp_ies_PE1
Src IP              : 192.168.1.1
Dst IP              : 192.168.1.2
Session Oper State  : connected
-------------------------------------------------------------------------------
Master Information
-------------------------------------------------------------------------------
Primary IP of Master: 192.168.1.2 (Other)
Addr List Mismatch  : No                Master Priority   : 255
Master Since        : 12/14/2010 17:39:57
Master Down Interval: 3.609 sec (Expires in 3.000 sec)
-------------------------------------------------------------------------------
Masters Seen (Last 32)
-------------------------------------------------------------------------------
Primary IP of Master   Last Seen             Addr List Mismatch     Msg Count
-------------------------------------------------------------------------------
192.168.1.1            12/14/2010 17:39:57   No                             0
192.168.1.2            12/14/2010 17:54:03   No                        342583
-------------------------------------------------------------------------------
Statistics
-------------------------------------------------------------------------------
Become Master       : 6                 Master Changes    : 11
Adv Sent            : 4441              Adv Received      : 342583
Pri Zero Pkts Sent  : 1                 Pri Zero Pkts Rcvd: 0
Preempt Events      : 0                 Preempted Events  : 5
Mesg Intvl Discards : 0                 Mesg Intvl Errors : 0
Addr List Discards  : 0                 Addr List Errors  : 338989
Auth Type Mismatch  : 0                 Auth Failures     : 0
Invalid Auth Type   : 0                 Invalid Pkt Type  : 0
IP TTL Errors       : 0                 Pkt Length Errors : 0
Total Discards      : 0
===============================================================================
```

Finally, for troubleshooting: it could be that the BFD session between the two IP interfaces is up but (in one or both peers) the command **show router vrrp instance interface** *interface-name* gives the following output regarding BFD for one or more VRID's.

```
--------------------------------------------------------------------------------
BFD Interface
-------------------------------------------------------------------------------
Service ID          : None
Interface Name      : vrrp_ies_PE2
Src IP              : 0.0.0.0
Dst IP              : 192.168.1.1
Session Oper State  : notConfigured
---------------------------------------------------------------------------------
```

To fix this, check that BFD has been correctly configured for the VRRP istances.

For instance, in the following example, the cause of the misconfiguration is that the IES service-id is not declared in the bfd-enable command:

```
configure service ies 10 interface vrrp_ies_PE2
                vrrp 10 owner
                bfd-enable interface vrrp_ies_PE2 dst-ip 192.168.1.1
                exit
```

Conclusion

BFD is a light-weight protocol which provides rapid path failure detection between two systems and it is useful in situations where the physical network has numerous intervening hops which are not part of the Layer 3 network.

BFD is linked to a protocol state. For BFD session to be established, the prerequisite condition is that the protocol to which the BFD is linked must be operationally active. Once the BFD session is established, the state of the protocol to which BFD is tied to is then determined based on the BFD session's state. This means that if the BFD session goes down, the corresponding protocol will be brought down.

In this section every scenario where BFD could be implemented has been described, including the configuration, show output and troubleshooting hints.