CHAPTER 3

Troubleshooting Interface Connectivity

This chapter explains how to troubleshoot problems with connectivity between interfaces on the ASR 9000 and interfaces on remote devices. It includes the following topics:

Troubleshooting Ping and ARP Connectivity

Follow the steps in this section to troubleshoot ping andAddress Resolution Protocol (ARP) connectivity problems. The overall approach is to verify that the routing protocol is up, the network topology is properly configured, and neighbors are up and reachable.

This procedure sends ping messages to the remote end and analyzes the resulting responses. For Ethernet interfaces, Address Resolution Protocol (ARP) connectivity is a prerequisite for ping connectivity—ARP must work first before ICMP echo can work. Therefore, check ARP and ensure it is working so that ICMP echo and ping can work. (Optical interfaces do not involve ARP.)

You should trace the path of the ping packets to see if they are getting dropped at any point along the path. Typical steps to locate these drops are:

- Run the command **show inject stats** —See if the packets went from the CPU to the NP.
- Run the command **show interface stats**—Look at the Tx counters to see if the packet was sent.
- Check the remote interface statistics to see if they received the packets.
- Check the remote punt and ARP statistics.
- Run commands in this list (above) on the return patch to check for drops.

Figure 3-1 shows the general approach to troubleshooting ping and ARP connectivity issues. Throubleshooting is required because a ping attempt has failed. In this example, the IGP protocol refers to the protocol currently configured on the network you are troubleshooting—OSPF, EIGRP, IS-IS, or RIP.

Figure 3-1 Example of Troubleshooting Ping and ARP Connectivity Issues

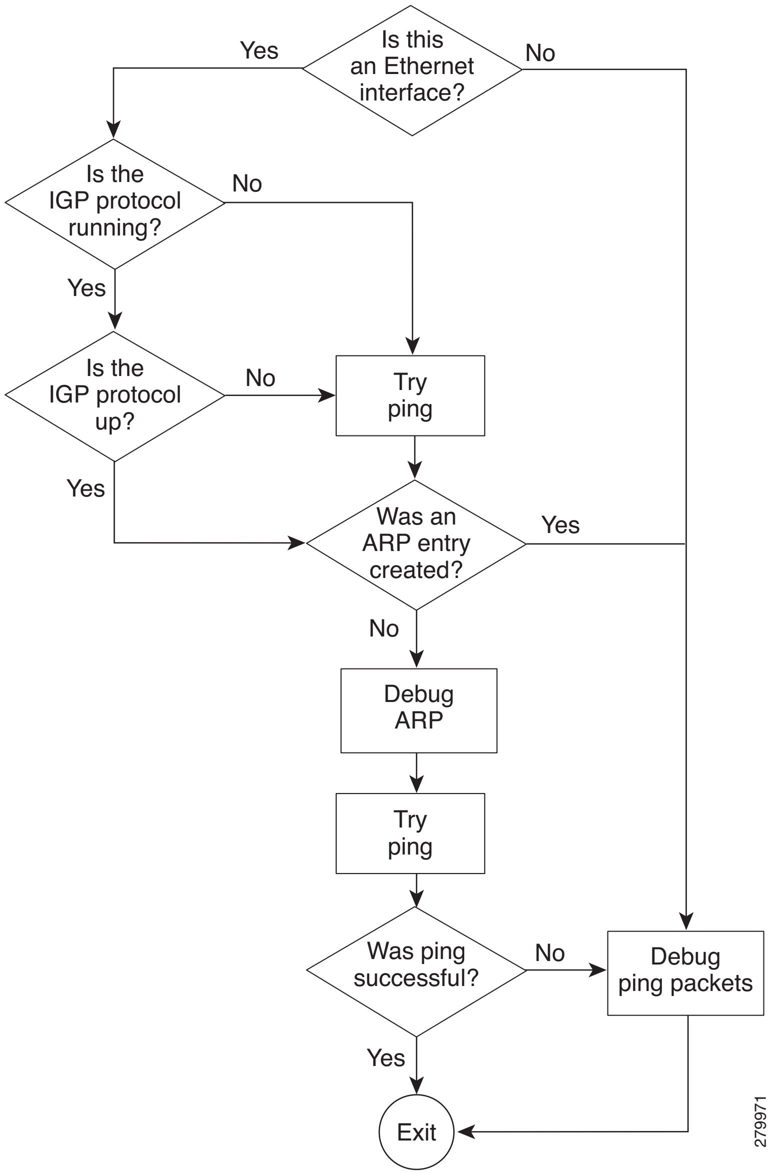

Follow these steps to troubleshoot ping and ARP connectivity issues. See Figure 3-1 to help you locate the steps that apply to your network scenario.

Step 1 Ping the remote end and check for a response. If there is no response, continue with this procedure to determine why the ping was unsuccessful and to connect successfully to the remote end.

Tip Use a systematic process to isolate the location of the failure. Ping the local interface first. If that is successful, ping the directly connected neighbor (single hop). If that is successful, ping the next hop, and so forth.

Step 2 Verify that the interface is configured as Layer 3.

Step 3 Check the routing table to make sure the IP address you are trying to ping has a route in the Routing Information Base (RIB) table.

```
RP/0/RSP0/CPU0:router# show route
```

Step 4 Verify that the IGP protocol is running and the connection to the neighbor is up. The following example assumes OSPF as the IGP.

```
RP/0/RSP0/CPU0:router# show protocols ospf
```

```
RP/0/RSP0/CPU0:router# show ospf neighbor
```

Step 5 Verify that packets are coming in and going out on the Ethernet interface.

Be aware of the following ARP behaviors when you are reviewing the display from the **show** commands in this step:

- A normal ping will send an ARP packet out followed by the actual ICMP echo packets. ARP must work before ICMP echo can work. If the system is receiving zero packets back, then there was no ARP reply. Even a single packet back means there was an ARP reply. (The system sends the ARP packet only if there is no ARP entry. Otherwise, it skips the ARP and proceeds with the ICMP echo.)
- By default, the system attempts to ping the remote router five times. If the remote router was recently connected to the network, the first ping will fail because the system needs time to resolve the ARP packet.

```
RP/0/RSP0/CPU0:router# show interfaces location node-id
```

or

```
RP/0/RSP0/CPU0:router# show interfaces type
```

Step 6 Display the ARP information. The IP address that you attempted to ping should be in the output.

```
RP/0/RSP0/CPU0:router# show arp
```

Step 7 If this port was previously attached to another device, or some othe major change has taken place on the remote end, use the **clear arp-cache** command to build a new entry. Verify that the MAC address in the ARP table is correct (see the MAC address in the Hardware Addr column in the example in Step 6).

```
RP/0/RSP0/CPU0:router# clear arp-cache
```

```
RP/0/RSP0/CPU0:router# show arp
```

Step 8 Determine whether an ARP entry exists for the destination IP.

```
RP/0/RSP0/CPU0:router# show arp location node-id
```

a. If an ARP entry does not exist or is incomplete, add a static ARP entry. Ensure that the Tx adjacency points to ‘COMPLETE’.

```
RP/0/RSP0/CPU0:router# show cef {ipv4} prefix hardware egress detail location node-id
```

Caution After you finish using the static ARP entry for troubleshooting purposes, you *must* remove it. If you do not remove the static ARP entry, it will cause traffic to be misdirected.

b. If the ARP entry points to ‘COMPLETE’, it means that the ARP entry is not being updated. Troubleshooting should now focus on why the ARP entry is not getting added (this includes steps such as **show arp, show arp idb**, **show adjacency gig *node-id* detail location *node-id***, and **show arp trace**).

c. If the Tx adjacency still points to ‘PUNT’, it means ARP is adding the entry in the database, but fib_mgr fails to mark the adjacency as ‘COMPLETE’.

d. This could be a fib_mgr, ARP, or AIB problem. Delete and reconfigure the static ARP entry with AIB and CEF debugs on. The debugs show if ARP is adding the entry inside the AIB and if the AIB is informing fib_mgr.

Step 9 Send a burst of traffic to help troubleshoot whether traffic is getting through. This can help in scenarios such as a disconnected cable, intermittent drops at unknown locations, and so forth.

a. Configure a static ARP entry, then send a large number of ping packets (for example, 100 or 1,000 packets) with a zero timeout. This sends out a burst of traffic from the router. When the ping fails, the system displays a dot instead of an exclamation point. If the ping is not possible, the system displays a ‘U’ instead of a dot.

Example:

```
RP/0/RSP0/CPU0:router# ping 192.0.2.55 count 100 timeout 0
Wed Sep 29 15:09:29.809 EDT
Type escape sequence to abort.
Sending 100, 100-byte ICMP Echos to 192.0.2.55, timeout is 0 seconds:
......................................................................
..............................
Success rate is 0 percent (0/100)
```

b. There are parameters to the ping command that can be used to change the time delay for the reply. Try lowering the delay, changing parameters to send 100 ping requests, and so forth.

c. There is a mode that allows you to suppress the ARP request and send out only the ICMP echo packets. See if the pings are failing intermittently or all the time.

Step 10 If ARP connectivity fails, perform the following steps to find out why.

a. Remove the static ARP entry.

b. Local ping—Ping your own interface (the interface your router uses to send out the pings).

c. Determine whether the local ping was successful. If the local ping failed (no response), pings out of that interface will also fail.

d. If you have an ARP entry, verify that there are outgoing/incoming ICMP echo/reply packets.

Note A ping is represented by a dot, exclam point, or capital U. The RSP allows a specific number of seconds for the ping to complete. The ARP is hidden inside this event.

Step 11 If you have to dig deeper into the issue, use the following commands to dump ping packets.

Example

```
RP/0/RSP0/CPU0:router# show route
C 172.21.116.0/24 is directly connected, 2d19h, MgmtEth0/RSP0/CPU0/0
is directly connected, 2d19h, MgmtEth0/RSP1/CPU0/0
L 172.21.116.10/32 is directly connected, 2d20h, MgmtEth0/RSP0/CPU0/0
L 172.21.116.11/32 is directly connected, 2d19h, MgmtEth0/RSP1/CPU0/0
L 172.21.116.12/32 [0/0] via 172.21.116.12, 2d19h, MgmtEth0/RSP0/CPU0/0
O 192.168.12.0/24 [110/2] via 192.168.111.11, 2d19h, GigabitEthernet0/2/0/1
[110/2] via 192.168.121.12, 2d19h, GigabitEthernet0/2/0/2
O 192.168.21.0/24 [110/2] via 192.168.111.11, 2d19h, GigabitEthernet0/2/0/1
[110/2] via 192.168.121.12, 2d19h, GigabitEthernet0/2/0/2
C 192.168.111.0/24 is directly connected, 2d20h, GigabitEthernet0/2/0/1
L 192.168.111.1/32 is directly connected, 2d20h, GigabitEthernet0/2/0/1
```

```
O 192.168.112.0/24 [110/2] via 192.168.111.11, 2d19h, GigabitEthernet0/2/0/1
O 192.168.113.0/24 [110/2] via 192.168.111.11, 2d19h, GigabitEthernet0/2/0/1


RP/0/RSP0/CPU0:router# show protocols ospf

Routing Protocol OSPF 100
  Router Id: 10.144.144.144
  Distance: 110
  Non-Stop Forwarding: Enabled
  Redistribution:
    None
  Area 0
    MPLS/TE enabled
    Loopback0
    GigabitEthernet0/1/0/2
    GigabitEthernet0/1/0/8
    GigabitEthernet0/1/0/18
    GigabitEthernet0/1/0/23
    TenGigE0/4/0/0* Indicates MADJ interface

RP/0/RSP0/CPU0:router# show ospf neighbor
Neighbors for OSPF 100

Neighbor ID     Pri   State           Dead Time   Address         Interface
10.164.164.164  1     FULL/DR         00:00:36    10.147.4.64     GigabitEthernet0/1/0/2
    Neighbor is up for 2d17h
10.166.166.166  1     FULL/DR         00:00:39    10.146.4.66     GigabitEthernet0/1/0/8
    Neighbor is up for 2d17h
10.19.19.19     1     FULL/BDR        00:00:33    10.194.4.19     GigabitEthernet0/1/0/18
    Neighbor is up for 2d18h
10.11.11.11     1     FULL/BDR        00:00:34    10.114.4.11     GigabitEthernet0/1/0/23
    Neighbor is up for 2d18h
10.11.11.11     1     FULL/BDR        00:00:39    10.114.8.11     TenGigE0/4/0/0
    Neighbor is up for 2d18h

Total neighbor count: 5
```

The following example shows the packet counts for a line card. Note that there are packets being input and output.

```
RP/0/RSP0/CPU0:router# show interfaces location 0/4/CPU0
Wed Sep  1 09:22:03.427 DST
TenGigE0/4/0/0 is up, line protocol is up
  Interface state transitions: 1
  Hardware is TenGigE, address is 001b.53ff.a780 (bia 001b.53ff.a780)
  Layer 1 Transport Mode is LAN
  Description: Connected to P11_CRS-4 10GE 0/2/5/0
  Internet address is 10.114.8.44/24
  MTU 9100 bytes, BW 10000000 Kbit (Max: 10000000 Kbit)
     reliability 255/255, txload 0/255, rxload 0/255
  Encapsulation ARPA,
  Full-duplex, 10000Mb/s, LR, link type is force-up
  output flow control is off, input flow control is off
  loopback not set,
  ARP type ARPA, ARP timeout 04:00:00
  Last input 00:00:00, output 00:00:00
  Last clearing of "show interface" counters never
  5 minute input rate 28000 bits/sec, 39 packets/sec
  5 minute output rate 45000 bits/sec, 39 packets/sec
     2356786692 packets input, 151622450429 bytes, 26 total input drops
     0 drops for unrecognized upper-level protocol
     Received 2 broadcast packets, 2327063140 multicast packets
```

```
              0 runts, 0 giants, 0 throttles, 0 parity
     0 input errors, 0 CRC, 0 frame, 0 overrun, 0 ignored, 0 abort
     30320436 packets output, 4277187228 bytes, 18 total output drops
     Output 1 broadcast packets, 495705 multicast packets
     0 output errors, 0 underruns, 0 applique, 0 resets
     0 output buffer failures, 0 output buffers swapped out
     1 carrier transitions


RP/0/RSP0/CPU0:router# show arp
Wed Sep  1 10:16:22.472 DST

-------------------------------------------------------------------------------
0/4/CPU0
-------------------------------------------------------------------------------
Address         Age        Hardware Addr   State      Type  Interface
10.114.8.11     02:40:44   001b.0c63.67ff  Dynamic    ARPA  TenGigE0/4/0/0
10.114.8.44     -          001b.53ff.a780  Interface  ARPA  TenGigE0/4/0/0

-------------------------------------------------------------------------------
0/1/CPU0
-------------------------------------------------------------------------------
Address         Age        Hardware Addr   State      Type  Interface
10.114.4.11     00:15:22   001b.0c63.67e7  Dynamic    ARPA  GigabitEthernet0/1/0/23
10.114.4.44     -          001b.53ff.87f7  Interface  ARPA  GigabitEthernet0/1/0/23
10.145.4.38     01:43:50   001e.f77d.5219  Dynamic    ARPA  GigabitEthernet0/1/0/27
10.145.4.44     -          001b.53ff.87fb  Interface  ARPA  GigabitEthernet0/1/0/27
10.146.4.44     -          001b.53ff.87e8  Interface  ARPA  GigabitEthernet0/1/0/8
10.146.4.66     02:56:39   0022.0d26.3bc4  Dynamic    ARPA  GigabitEthernet0/1/0/8
10.147.4.44     -          001b.53ff.87e2  Interface  ARPA  GigabitEthernet0/1/0/2
10.147.4.64     00:33:21   0022.0d26.36c4  Dynamic    ARPA  GigabitEthernet0/1/0/2
10.194.4.19     03:16:59   001a.3029.d400  Dynamic    ARPA  GigabitEthernet0/1/0/18
10.194.4.44     -          001b.53ff.87f2  Interface  ARPA  GigabitEthernet0/1/0/18
10.194.8.44     -          001b.53ff.87f0  Interface  ARPA  Bundle-Ether16.162
10.194.12.44    -          001b.53ff.87f0  Interface  ARPA  Bundle-Ether16.163
10.194.16.44    -          001b.53ff.87ec  Interface  ARPA  GigabitEthernet0/1/0/12

-------------------------------------------------------------------------------
0/RSP0/CPU0
-------------------------------------------------------------------------------
Address         Age        Hardware Addr   State      Type  Interface
172.29.52.1     01:51:49   001e.f77d.2a19  Dynamic    ARPA  MgmtEth0/RSP0/CPU0/0
172.29.52.13    03:29:42   0010.79e9.6038  Dynamic    ARPA  MgmtEth0/RSP0/CPU0/0
172.29.52.21    00:50:04   0022.0d5a.a6c4  Dynamic    ARPA  MgmtEth0/RSP0/CPU0/0
172.29.52.22    02:44:58   0001.6443.1678  Dynamic    ARPA  MgmtEth0/RSP0/CPU0/0
172.29.52.27    02:36:46   0012.7fd6.ba08  Dynamic    ARPA  MgmtEth0/RSP0/CPU0/0
172.29.52.28    03:04:37   0012.7fd6.ba09  Dynamic    ARPA  MgmtEth0/RSP0/CPU0/0

RP/0/RSP0/CPU0:router# show cef 192.168.1.1/32 hardware egress detail location 0/4/CPU0
192.168.6.73/32, version 0, internal 0x40040001 (ptr 0x9f613944) [1], 0x0 (0x9eb9bdfc),
0x4500 (0xa0156184)
 Updated Sep 22 20:20:54.369
 remote adjacency to GigabitEthernet0/1/0/23
 Prefix Len 32, traffic index 0, precedence routine (0)
  gateway array (0x9e935bbc) reference count 249, flags 0xd00, source lsd (2),
                [84 type 5 flags 0x101001 (0x9fb06898) ext 0x0 (0x0)]
  LW-LDI[type=5, refc=3, ptr=0x9eb9bdfc, sh-ldi=0x9fb06898]
   via 10.114.4.11, GigabitEthernet0/1/0/23, 10 dependencies, weight 0, class 0 [flags
0x0]
    path-idx 0
    next hop 10.114.4.11
    remote adjacency
     local label 16021      labels imposed {16031}
   via 10.114.8.11, TenGigE0/4/0/0, 13 dependencies, weight 0, class 0 [flags 0x0]
```

```
    path-idx 1
    next hop 10.114.8.11
    local adjacency
     local label 16021      labels imposed {16031}
.
.
.
    TX Adjacency
      Raw data for tx adj struct:
      Raw result1: 0x03000100 0x01000000 0x7e23001b 0x0c6367ff
      Raw result2: 0x95650300 0x00000000 0x00000000 0x00000000
      ---------------------------------------------
         Search Ctrl Flags:
         ------------------
         match         : 1       valid        : 1
         gre_adj       : 0       null_route   : 0
         tx_punt       : 0       tx_drop      : 0
         next_hop_down : 0       adj_complete : 0
         punt_ifib     : 0       nhop_down    : 0
         stop          : 0       match_all_bit: 0
         default_action: 1
         uidb_index          : 0x0001
         l3_mtu              : 9086
         dest mac            : 001b.0c63.67ff
         prefix_adj_cnt_index: 0x95650300

    RX Adjacency
      Raw data for rx adj struct:
      Raw result1: 0x13000100 0x00001300 0x0c000280 0x00000000
      ---------------------------------------------
         Search Ctrl Flags:
         ------------------
         rx_punt       : 0       rx_drop       : 0
         rx_adj_SFP    : 1       rp_destined   : 0
         rp_drop       : 0       match         : 1
         valid         : 1       rx_LAG_adj    : 0
         match_all_bit : 0       pri_adj_down  : 0
         default_action: 1
         rx_adj_field : 0x0013
         egress_ifh   : 0xc000280


    Load distribution: 0 1 (refcount 84)

    Hash  OK  Interface                 Address
    0     Y   GigabitEthernet0/1/0/23   remote
    1     Y   TenGigE0/4/0/0            10.114.8.11
```

Troubleshooting Bidirectional Forwarding Detection

Bidirectional Forwarding Detection (BFD) is a detection protocol designed to provide fast forwarding path failure detection times for all media types, encapsulations, topologies, and routing protocols. This section contains the following subsections:

- Using show and debug Commands, page 3-82
- BFD Sessions in Down State, page 3-83
- BFD Sessions Flap, page 3-83

- BFD Sessions Down on Neighboring Router, page 3-85
- BFD Sessions Are Not Created on the LC, page 3-85

Using show and debug Commands

SUMMARY STEPS

1. **show bfd [ipv4 | all] [location** *node-id*]
2. **show bfd client [detail]**
3. **show bfd [ipv4 | all] session** [**detail** | [**interface** *ifname]* | [**location** *node-id*]] [**detail**]
4. **show bfd counters packet [interface** *ifname*] **location** *node-id*
5. **show bfd trace** {**adjacency | error | fsm | packet**} [**interface** *ifname*] [**location** *node-id*]
6. **show tech-support routing bfd {file | location | rack}**

DETAILED STEPS

	Command or Action	Purpose

Step 1	`show bfd [ipv4 | all] [location node-id]`
	Example:
	`RP/0/RSP0/CPU0:router# show bfd location 0/4/CPU0`
		View general BFD information on the Route Switch Processor (RSP), such as the number of sessions. Use the **location** keyword to display information for a specific LC. If not specified, information for all locations displays.

Step 2	`show bfd client [detail]`
	Example:
	`RP/0/RSP0/CPU0:router# show bfd client detail`
		View BFD clients. Use the **detail** keyword to display more information.

Step 3	`show bfd [ipv4 | all] session [detail | [interface ifname] | [location node-id] ] [detail]`
	Example:
	`RP/0/RSP0/CPU0:router# show bfd session interface Gig2/1/0/0 detail`
		View BFD session information. Filter results using the following parameters and keywords:
		• **location**—BFD sessions hosted on this location.
		• **interface**—BFD sessions on the specified interface (no wildcards).
		• **detail**—Detailed session information: statistics, number of state transitions.

Step 4	`show bfd counters packet [interface ifname] location node-id`
	Example:
	`RP/0/RSP0/CPU0:router# show bfd counters packet interface POS 0/3/0/0 location 0/3/cpu0`
		View packet counters information. Filter results using the following parameters and keywords:
		• **location**—Packet counters for BFD sessions hosted on this location.
		• **interface**—Packet counters for BFD sessions on the specified interface (no wildcards).
		• **invalid**—Invalid packet counter information.

	Command or Action	Purpose
Step 5	**show bfd trace** {**adjacency** | **error** | **fsm** | **packet**} [**interface** *ifname*] [**location** *node-id*] **Example:** `RP/0/RSP0/CPU0:router# show bfd trace fsm location 0/4/CPU0`	View tracing information from the RSP. Filter results using the following parameters and keywords: • **adjacency**—Traces generated when BFD receives an adjacency update from the Adjacency Information Base (AIB) Finite State Machine (FSM) display. • **error**—Traces generated when an error is detected. • **fsm**—Traces generated when there is a state change in a session. • **packet**—Traces generated when there is a change in a Tx or Rx packet. • **location**—Traces for BFD traces on the specified interface. **Note** Log the trace to a file to save the results.

Step 6	sh**ow tech-support routing bfd {file** | **location** | **rack}**
	Example:
	`RP/0/RSP0/CPU0:router# show tech-support routing bfd location 0/1/CPU0`
		View BFD debugging information.

BFD Sessions in Down State

To troubleshoot BFD sessions in the down state, perform the following steps.

Step 1 Verify IP connectivity. Verify there is no IP packet loss.

```
RP/0/RSP0/CPU0:router# ping local-remote-address
```

Step 2 Ensure that the router and remote device are configured with the following parameters:

- Number of BFD sessions they can support
- Timers to support the police rates

BFD Sessions Flap

To check various BFD parameters, perform the following steps. Also see:

- BFD Sessions Down on Neighboring Router, page 3-85
- BFD Sessions Are Not Created on the LC, page 3-85

Step 1 Verify IP connectivity.

```
RP/0/RSP0/CPU0:router# ping local-IP-address
```

Step 2 View input and output counters.

```
RP/0/RSP0/CPU0:router# show interface
```

Step 3 View session detail information.

```
RP/0/RSP0/CPU0:router# show bfd session detail
```

Step 4 View session packet counters.

```
RP/0/RSP0/CPU0:router# show bfd counter
```

Step 5 View SPP counters.

```
RP/0/RSP0/CPU0:router# show spp node
```

Note SPP means software packet processing, but is more commonly referred to as vector path processing (VPP).

Step 6 View resource usage.

```
RP/0/RSP0/CPU0:router# monitor process
```

Step 7 View IP connectivity. Verify there is no IP packet loss.

```
RP/0/RSP0/CPU0:router# ping local remote address
```

If the following message appears, the BFD flap is a result of the application flap.

```
bfd_agent[104]: %BFD-6-SESSION_REMOVED : BFD session to neighbor 192.168.1.1 on interface
Gi0/5/0/0 has been remove
```

Step 8 Verify that the SPP is not losing packets.

```
RP/0/RSP0/CPU0:router# show spp node location
```

Step 9 Check LC CPU and memory usage.

```
RP/0/RSP0/CPU0:router# monitor processes location
```

Step 10 Check the local interface counters.

```
RP/0/RSP0/CPU0:router# show interfaces type interface-name
```

Step 11 Check any QoS policies applied to the interface.

```
RP/0/RSP0/CPU0:router# show policy-map interface
```

Step 12 Repeat Step 1 through Step 11on the remote end.

BFD Sessions Flap Because of Local Echo Failure

BFD sessions flap may be locally triggered because the router detects echo failure.

Examine LC CPU utilization:

```
RP/0/RSP0/CPU0:router# monitor process location
```

Examine the SPP process on the LC CPU to determine the delay encountered by BFD echo packets:

```
RP/0/RSP0/CPU0:router# show bfd trace performance reverse location
```

Rule out BFD echo packet loss: **`show bfd counters packet location`**

BFD Sessions Flap Because of SPP Process Restart

If BFD failure detection is configured to be within 1 second, the BFD session would flap if the SPP process is restarted on the LC.

BFD Sessions Down on Neighboring Router

The neighbor router sends this message to indicate its BFD is going down:

```
LC/0/6/CPU0:Aug 8 16:42:56.821: bfd_agent[104]: %L2-BFD-6-SESSION_STATE_DOWN: BFD session
to neighbor 192.168.1.1 on interface Gi0/5/0/0 has gone down. Reason: Nbor signalled down
```

BFD Sessions Are Not Created on the LC

Up to 1024 BFD sessions are allowed per LC. Configuring more than 1024 BFD sessions may result in random BFD sessions not being created.

Troubleshooting Ethernet CFM

Ethernet Connectivity Fault Management (CFM) monitors, detects, and diagnoses remote network faults end-to-end across the network. It does this using keepalives and MAC-based ping and traceroutes. Unlike most other Ethernet protocols which are restricted to a single physical link, CFM frames can transmit across the entire end-to-end Ethernet network.

This section describes how to troubleshoot problems with CFM on the local ASR 9000 router. For more information on how to use CFM to troubleshoot problems across the network, see the “Ethernet CFM” section in *Cisco ASR 9000 Series Aggregation Services Router Interface and Hardware Component Configuration Guide*.

Figure 3-2 shows an example of maintenance domains across a network.

Figure 3-2 CFM Maintenance Domains Across a Network

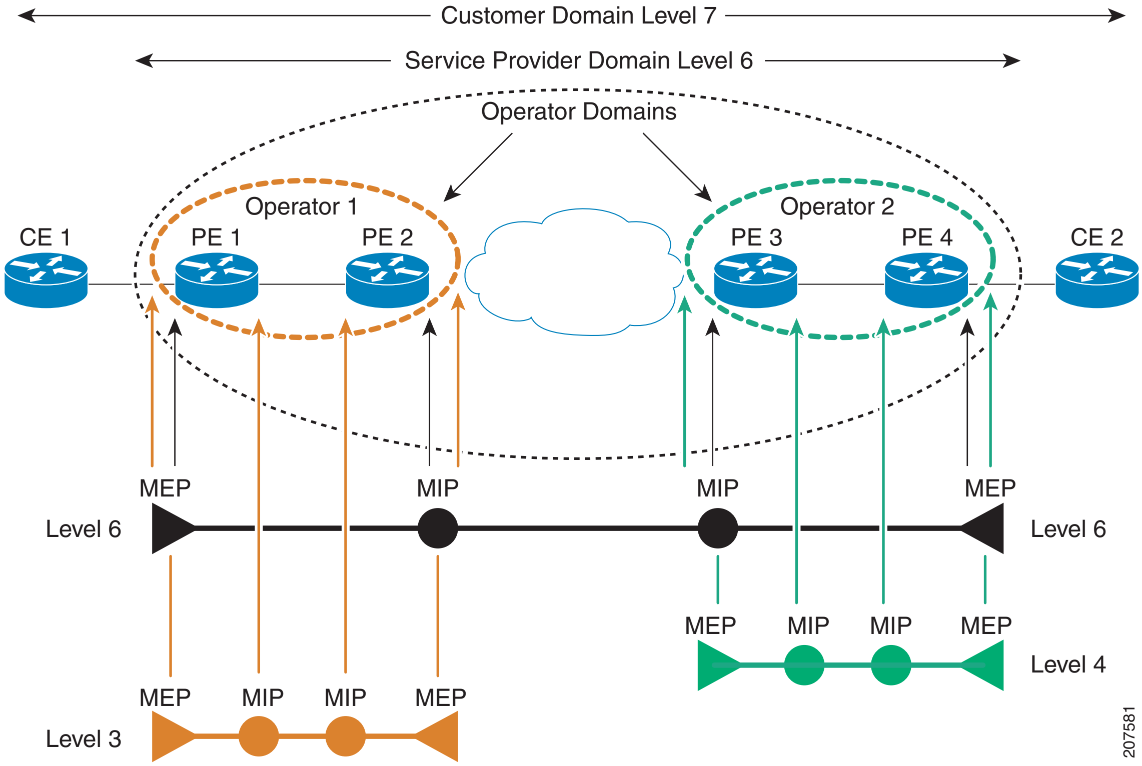

This section contains the following topics:

- Using show and debug Commands, page 3-87
- MEPs Are Not Created, page 3-88
- MIPs Are Not Created, page 3-88
- No CCMs are Received at the MEP or Peer MEPs Are Not Seen, page 3-89
- Peer MEP Defects and Mismatches Are Seen, page 3-90
- Remote Defect Indication Received, page 3-91
- Peer MEP Times Out But No Alarm Or Action Occurs, page 3-92
- No Debugs or Counters for Higher-Level Packets at a MEP or MIP, page 3-92
- Dropped CFM PDUs, page 3-92
- CFM ping Or traceroute Returns a “not found” Error, page 3-93
- AIS Messages Are Not Sent, page 3-93

Tip For an extensive discussion of CFM usage and CFM command examples, see *Cisco ASR 9000 Series Aggregation Services Router Interface and Hardware Component Configuration Guide.*

Using show and debug Commands

The show and debug commands in this section are useful for troubleshooting CFM. Further details on these commands can be found in *Cisco ASR 9000 Series Aggregation Services Router Interface and Hardware Component Command Reference*.

These commands are useful for checking the validity of the configuration commands:

- show run ethernet cfm—View all CFM global configuration.
- show ethernet cfm configuration-errors—Displays any problems that have been detected in the CFM configuration.

These commands are useful for verifying CFM functions:

- show ethernet cfm local maintenance-points—Displays a summary of the MEPs and MIPs that have been created.
- show ethernet cfm local meps—Displays information about local MEPs, including continuity check messages (CCMs), details about the types of packets being sent and received, and counters for each packet type.
- show ethernet cfm peer meps—Displays information about peer MEPs, including details about any peer MEP defects that have been detected.
- show ethernet cfm traceroute-cache—Displays the contents of the traceroute cache, that is, the result of recent traceroute operations.
- show ethernet cfm interfaces ais—Displays a summary of interfaces where AIS messages are being sent or received.
- show ethernet cfm interfaces statistics—Displays counters for CFM PDUs that are dropped per interface.
- show ethernet cfm ccm-learning-database—Displays the contents of the CCM learning database, which the system uses when it responds to received traceroute (linktrace) messages.
- debug ethernet cfm packets—Enables debugging of sent and received CFM PDUs.
- debug ethernet cfm protocol-state—Enables debugging of major CFM protocol state-machine operations.

If you need to collect information to provide to Cisco, the following commands are also useful, in addition to those listed above. Note that many of these commands require the cisco-support task ID.

- show tech-support ethernet
- show ethernet cfm trace
- show ethernet cfm interfaces status
- show ethernet cfm services
- debug ethernet cfm platform—Displays platform-specific debugging information for CFM.
- debug ethernet oam platform—Displays platform-specific debugging information for OAM.
- show spp node—Displays SPP counters.
- show spp sid stats—Check the SPP stream ID (SID) statistics to see that CFM traffic is injected and punted.
- show spp client—Displays information from the RSP about traffic on the SPP. Check to see if there are any SPP drops.

Note To clear the spp counters, run the command **clear spp {client | interface | node-counters} location** *node-id*. This command clears client statistics, interface statistics, and per-node counters, depending on the keyword you use.

MEPs Are Not Created

If MEPs have been configured but have not been created, follow the troubleshooting steps in this section.

Step 1 Display information about errors that might be preventing configured CFM operations from becoming active, as well as any warnings that have occurred.

```
show ethernet cfm configuration-errors
```

Step 2 Display a list of local maintenance points that have been created. Verify that the list contains the expected nodes.

```
show ethernet cfm local maintenance-points
```

Step 3 Display operational states of local MEPs. Verify that the states are as expected.

```
show ethernet cfm local meps
```

MIPs Are Not Created

This section explains what to do if MIP creation has been configured but MIPs have not been created as expected. Understand the factors that can impact MIP creation, then troubleshoot the specific MIP creation issues.

Understanding the Factors that Impact MIP Creation

Configuring MIP creation for a service does not guarantee that MIPs will actually be created for all interfaces in that service. MIPs are only created on interfaces that are correctly configured for Layer 2 switching, that is, interfaces that:

- Are configured as Layer 2 interfaces
- Have an appropriate encapsulation configured
- Have been added to a bridge domain or point-to-point xconnect.

The CFM standard (IEEE 802.1ag-2007) specifies an algorithm that is used to determine whether a MIP should be created, and at what level. For details of this implementation, see the **mip auto-create** command in *Cisco ASR 9000 Series Aggregation Services Router Interface and Hardware Component Command Reference*.

Troubleshooting MIP Creation Issues

Follow these steps to troubleshoot MIP creation issues.

Step 1 Display information about errors that might be preventing configured CFM operations from becoming active, as well as any warnings that have occurred.

```
show ethernet cfm configuration-errors
```

Step 2 Display a list of local maintenance points that have been created. Verify that the list contains the expected nodes. Check for MEPs configured on the interface, and for MIPs enabled on a service at a lower level.

```
show ethernet cfm local maintenance-points
```

Step 3 If MIP creation is not functioning, verify that the bridge domain or xconnect is configured correctly. To verify or troubleshoot these bridge domain and xconnect configurations, see Chapter 9, "Troubleshooting L2VPN and Ethernet Services."

No CCMs are Received at the MEP or Peer MEPs Are Not Seen

This section explains how to troubleshoot the following conditions:

- Continuity check messages (CCMs) are not seen at one or more maintenance end points (MEPs).
- Peer MEPs are not seen.

CFM MEPs exchange CCMs periodically according to parameters configured on the system. These CCMs are multicast to all other MEPs in the service at the same level. When the local MEP receives a CCM, it creates an entry in the peer MEP table. If CCMs are not being exchanged correctly, perform the following steps.

Step 1 Verify that CCM is enabled and there is a supported encapsulation on the interface.

```
RP/0/RSP0/CPU0:router# show running-config
```

Step 2 View configured MEPs and maintenance intermediate points (MIPs).

```
RP/0/RSP0/CPU0:router# show ethernet cfm local maintenance-points
```

Step 3 Display operational states of local and peer MEPs. Verify that CCM is enabled and the states are as expected.

```
RP/0/RSP0/CPU0:router# show ethernet cfm local meps
RP/0/RSP0/CPU0:router# show ethernet cfm local meps verbose
RP/0/RSP0/CPU0:router# show ethernet cfm peer meps
RP/0/RSP0/CPU0:router# show ethernet cfm peer meps detail
```

Step 4 View packets seen by the CFM PI. Enable all of the options. The output shows if packets are dropped, forwarded, or processed.

```
RP/0/RSP0/CPU0:router# debug ethernet cfm packets packet-type ccm
```

Step 5 View remote MEPs shown by the specific LC CFM instance. If CCMs are not received, the peer does not display.

Step 6 View CFM SID statistics seen by the SPP. This displays any CFM traffic that is injected and punted.

```
RP/0/RSP0/CPU0:router# show spp sid stats
```

Step 7 View SPP drops.

```
RP/0/RSP0/CPU0:router# show spp client location location
```

Step 8 Show bi-state alarms, to check for invalid encapsulation.

Step 9 Check for dropped PDUs as described in the "Dropped CFM PDUs" section on page 3-92.

Peer MEP Defects and Mismatches Are Seen

This section explains what to do when MEPs are exchanging CCMs, but there are defects or mismatches in the received CCMs. Perform the following steps to locate the specific problem. Then take appropriate corrective action.

Step 1 Run the following commands to obtain the output you will need for detailed troubleshooting.

```
RP/0/RSP0/CPU0:router# show ethernet cfm local meps
RP/0/RSP0/CPU0:router# show ethernet cfm local meps verbose
RP/0/RSP0/CPU0:router# show ethernet cfm peer meps
RP/0/RSP0/CPU0:router# show ethernet cfm peer meps detail
```

Step 2 Wrong level—Check the output of the commands in Step 1 to see if CCMs are being received at a lower level than the level of the local MEP. This indicates a misconfiguration, for example:

- The domain level is configured incorrectly on the local device or the peer device.
- An intended MEP at the lower level has not been configured, and as a result the CCMs it would consume are reaching the local MEP.
- The forwarding path within the network has been misconfigured, such that CCM packets are being received from an unintended source.

Step 3 Cross-connect (wrong MAID)—Check the output of the commands in Step 1 to see if CCMs are being received with an maintenance association identifier (MAID) that does not match the MAID configured locally for the service. The MAID is formed from the maintenance domain identifier (MDID) and the short maintenance association name (SMAN). By default, the MDID is set to the name of the domain and the SMAN is set to the name of the service. A crossconnect error indicates a misconfiguration, for example:

- The domain name or ID is configured incorrectly on the local device or on the peer device.
- The service name or ID is configured incorrectly on the local device or on the peer device.
- The forwarding path within the network has been misconfigured, such that CCM packets are being received from an unintended source.

Step 4 Wrong interval—Check the output of the commands in Step 1 to see if CCMs are being received with a CCM interval that does not match the locally configured CCM interval. This indicates that the interval is configured incorrectly on either the local device or the peer device. For a given service, the same CCM interval must be configured on all devices.

Step 5 Loop (local MAC address received)—Check the output of the commands in Step 1 to see if CCMs are being received with the source MAC equal to the MAC address of the interface for the local MEP. This indicates that there is a loop in the network such that the local device is receiving its own packets, or that two devices in the network are configured with the same MAC address.

Step 6 Configuration (local MEP ID received)—Check the output of the commands in Step 1 to see if CCMs are being received from a peer MEP with the same MEP ID as the local MEP. This defect indicates that two MEPs are configured with the same MEP ID. Across the entire network, each MEP in the service must be configured with a different MEP ID.

Step 7 Peer interface down—Check the output of the commands in Step 1 to see if CCMs are being received that indicate the interface on the peer MEP is down, or that the interface on every peer MEP is STP blocked. This indicates a problem with the operational state of the network.

Step 8 Missing (crosscheck)—Check the output of the commands in Step 1. If crosscheck is configured specifying this peer MEP, but no CCMs are being received, the peer MEP is missing. This might indicate a failure in the network.

Step 9 Unexpected (crosscheck)—Check the output of the commands in Step 1. If crosscheck is configured and CCMs are being received from a peer MEP that is not specified, these CCMs are unexpected. This may indicate a misconfiguration or that CCMs are being received from an unintended source.

Step 10 Remote defect received—Check the output of the commands in Step 1. If received CCMs indicate that the peer MEP has detected a defect, take the action recommended in the "Remote Defect Indication Received" section on page 3-91.

Remote Defect Indication Received

This section explains what to do if the local router receives a remote defect indication (RDI) from a peer router.

Understanding How RDIs are Exchanged

When a MEP detects a defect (as described in the "Peer MEP Defects and Mismatches Are Seen" section on page 3-90), it sets the remote defect indication (RDI) in the CCMs it is sending. When another MEP receives the RDI, it recognizes that the peer MEP sending the CCMs has detected a defect.

- In a point-to-point service, most defects will be detected by both MEPs; therefore both MEPs will send the RDI and both will receive the RDI. However, a unidirectional failure in the network could cause one of the MEPs to detect a crosscheck missing defect, while the other MEP does not detect any defect. In this case, the RDI sent by the first MEP serves to notify the second MEP of the problem.
- In a multipoint service, if there is a defect on any MEP or pair of MEPs, all other MEPs in the service receive the RDIs from the MEP or MEPS that detected the defect.

Locating the Source of RDIs and Resolving Defects

Step 1 Run the following commands to obtain the output you will need for troubleshooting RDIs.

```
RP/0/RSP0/CPU0:router# show ethernet cfm peer meps
RP/0/RSP0/CPU0:router# show ethernet cfm peer meps detail
```

Step 2 Determine the peer MEP from which the RDI is being received.

Step 3 Log into the peer device and follow the steps in the "Peer MEP Defects and Mismatches Are Seen" section on page 3-90.

Peer MEP Times Out But No Alarm Or Action Occurs

This section explains what to do if a peer MEP times out without generating an alarm or automatic corrective action. You need to verify that crosscheck is configured and functioning correctly.

If the local MEP stops receiving CCMs from a peer MEP, it times out by default after 3.5 times the CCM interval. This is refered to as a loss of continuity. By default, a loss of continuity does not trigger any other actions, such as log messages, SNMP traps, AIS, or Ethernet fault detection (EFD). These actions are only triggered if crosscheck is configured for the service and the peer MEP is specified. In this case, the loss of continuity causes a crosscheck missing defect, and this in turn triggers the other actions.

Step 1 Check that MEP crosscheck is configured for the service.

```
RP/0/RSP0/CPU0:router# show running-config
```

Step 2 Check for the crosscheck missing defect.

```
RP/0/RSP0/CPU0:router# show ethernet cfm peer meps
```

No Debugs or Counters for Higher-Level Packets at a MEP or MIP

CFM packets at a higher level than the highest MEP or MIP configured on the interface are forwarded by the network processor and are not handled by the software. Therefore, it is not possible to display debugs or counters for these packets. In addition, certain packets at the same level as a MIP are forwarded by the network processor. Again it is not possible to display debugs or counters for these.

Dropped CFM PDUs

If CFM PDUs are not reaching the expected destination or are not being processed as expected, it is possible that they are being dropped. The PDU drops could be caused by any of these reasons:

- Dropped by the network processor
- Dropped when being passed to software due to exceeding the supported CFM packet rate of 16,000 packets per second per line card

Note Note that the CFM packet rate limit (16,000 CFM packets per second per line card) includes all CFM packet types, including linktrace (traceroute) and loopback (ping) packets, as well as CCMs and Ethernet SLA probes. Normally, the number of linktrace or loopback packets is low; however, the use of "continuity-check auto-traceroute" can cause a high number of linktrace packets to be sent, if a number of peer MEPs time out in quick succession.

- Dropped because the interface or the forwarding node is down
- Dropped because the PDU is invalid or not formed properly
- Dropped because a higher level MEP was reached
- Dropped due to an unknown PDU type
- Dropped because the configured maximum MEPs limit (default 100) has been reached for the service

To display information for troubleshooting dropped CFM PDUs, perform the following steps. Take corrective actions based on the outputs of the commands in these steps.

Step 1 Enable packet debugging to determine whether forwarded packets are being received at the MIP.

```
RP/0/RSP0/CPU0:router# debug ethernet cfm packets [received dropped interface
gigabitEthernet node-id]
```

Step 2 Display the statistics of the CFM PDUs per interface. Look for any drops ted to packets that are improperly formed, invalid, wrong level, or unknown type.

```
RP/0/RSP0/CPU0:router# show ethernet cfm interfaces statistics
```

Step 3 Display the local MEPs and look for discarded CCMs. Discarded CCMs might indicate the the configured maximum MEPs limit (default 100 MEPs per service) is reached.

```
RP/0/RSP0/CPU0:router# show ethernet cfm local meps verbose
```

Step 4 View peer MEPs seen by every local MEP.

```
RP/0/RSP0/CPU0:router# show ethernet cfm peer meps
```

Step 5 Check the STP status on the interfaces with MEPs or MIPs. CFM PDUs originating at MEPs on a STP block port get forwarded, however, PDUs forwarded on a MIP are subject to the STP port state. This means that if MIP is on a port that is STP blocked, then CFM PDUs will be dropped at the MIP.

Step 6 View STP state and CFM peer MEP status.

```
RP/0/RSP0/CPU0:router# show spanning-tree mst mstp
```

CFM ping Or traceroute Returns a “not found” Error

This section explains what to do if you perform a CFM ping or traceroute and receive a “not found” error.

For the ping or traceroute commands, the target is specified by means of a MAC address or a MEP ID. If the target is specified as a MAC address, the MAC address is copied directly into the message. However, if a MEP ID is specified, the system looks in the peer MEP table to find the MAC address for the corresponding peer MEP. If there is no peer MEP for the service with the specified MEP ID, or if there is more than one peer MEP for the service with the specified MEP ID, this lookup fails and the system returns a “not found” error.

View the peer MEPs and check that there is an entry for the MEP ID that was being used as the target MEP ID in the ping or traceroute command.

```
RP/0/RSP0/CPU0:router# show ethernet cfm peer meps
```

AIS Messages Are Not Sent

Alarm indication signal (AIS) can be enabled in configuration, either for MEPs or explicitly on an interface. The system sends AIS messages when it detects a peer MEP defect, when it receives AIS or LCK messages, or when the interface is down. AIS messages are sent in one of two ways:

- If there is another MEP on the interface at a higher level, and in the same direction, the AIS messages are sent internally from the lower level MEP to the next highest level MEP. In this case, no actual PDUs are transmitted.
- Otherwise, if there is a MIP on the interface then AIS PDUs are transmitted at the level of the MIP. If there is no MIP on the interface, no AIS messages are transmitted.

Use the following steps for troubleshooting.

Step 1 Verify that AIS is enabled in the configuration.

```
RP/0/RSP0/CPU0:router# show running-config
```

Step 2 Verify that there is a MIP.

```
RP/0/RSP0/CPU0:router# show ethernet cfm local maintenance-points
```

Step 3 Display the information published in the interface AIS table, including a record of the AIS transmissions. Determine whether AIS messages are actually being sent.

```
RP/0/RSP0/CPU0:router# show ethernet cfm interfaces ais
```

Step 4 Determine whether the system should be sending AIS messages.

```
RP/0/RSP0/CPU0:router# show ethernet cfm local meps detail
```