

Troubleshoot Issues on dLFioLL/dLFioATM/dLFioFR/dMLP/dMLFR Interfaces

Document ID: 65334

Introduction

Prerequisites

- Requirements
- Components Used
- Conventions

Background Information

- How to Identify Whether Fragmentation Occurs

Troubleshoot Issues on dLFioATM and dLFioFR Interfaces

Troubleshoot Issues on dLFioLL Interfaces

- Latency Issues with dLFI
- dLFI Scalability Issues

Troubleshoot Issues on dMLP Interfaces

- dMLP Performance Issues
- Incorrect Rate Counters for dMLP Bundle Interfaces

Troubleshoot Issues on dMLFR Interfaces

Troubleshoot Other Error Messages

- FIXBADTXVC Messages
- RSPACCERR Messages

MTU–Related Information

Additional Information You Require on 76xx Series Routers

Case Studies for dDDR

- dCEF Does Not Switch Packets
- Troubleshoot
- Sample Output

Troubleshoot Router/VIP Crashes and Cbus Complex

NetPro Discussion Forums – Featured Conversations

Related Information

Introduction

This document helps you troubleshoot issues on various interfaces when you use these features:

- Distributed Link Fragmentation and Interleaving over ATM (dLFioATM)
- Distributed Link Fragmentation and Interleaving over Frame Relay (dLFioFR)
- Distributed Link Fragmentation and Interleaving over Leased Lines (dLFioLL)
- Distributed Multilink PPP (dMLP)
- Distributed Multilink Frame Relay (dMLFR)

Prerequisites

Requirements

Readers of this document should have knowledge of these topics:

- dLFIoATM
- dLFIoFR
- dLFIoLL
- dMLP
- dMLFR

Components Used

The information in this document is based on these software and hardware versions:

- All Cisco 7500 and 7600 platforms

Note: The information in this document applies to 6500 platforms also.

- Relevant Cisco IOS® Software Releases, which this table lists:

Feature	PAs ¹ Supported	7500 Platforms		7600 Platforms	
		Major Cisco IOS Software Releases	Cisco IOS Releases (Interim)	Major Cisco IOS Software Releases	Cisco IOS Releases (Interim)
dMLP	Chan-PA PA-4T+ PA-8T	<ul style="list-style-type: none"> ◆ 12.0T ◆ 12.0S ◆ 12.1 ◆ 12.1T ◆ 12.2 ◆ 12.2T ◆ 12.3 ◆ 12.3T ◆ 12.2S ◆ 12.1E² 	<ul style="list-style-type: none"> ◆ 12.0(3)T and later ◆ 12.0(9)S and later 	<ul style="list-style-type: none"> ◆ 12.2SX ◆ 12.1E² 	
dLFIoLL	Chan-PA PA-4T+ PA-8T	<ul style="list-style-type: none"> ◆ 12.2T ◆ 12.3 ◆ 12.3T ◆ 12.0S 	<ul style="list-style-type: none"> ◆ 12.2(8)T and later ◆ 12.0(24)S and later 	<ul style="list-style-type: none"> ◆ 12.2SX 	<ul style="list-style-type: none"> ◆ 12.2(17)SXB and later
dLFIoFR	Chan-PA PA-4T+ PA-8T	<ul style="list-style-type: none"> ◆ 12.2T ◆ 12.3 ◆ 12.3T 	<ul style="list-style-type: none"> ◆ 12.2(4)T3 and later 	<ul style="list-style-type: none"> ◆ 12.2SX 	<ul style="list-style-type: none"> ◆ 12.2(17)SXB and later
dLFIoATM	PA-A3 PA-A6	<ul style="list-style-type: none"> ◆ 12.2T ◆ 12.3 ◆ 12.3T 	<ul style="list-style-type: none"> ◆ 12.2(4)T3 and later 	<ul style="list-style-type: none"> ◆ 12.2SX 	<ul style="list-style-type: none"> ◆ 12.2(17)SXB and later
dMLFR	Chan-PA PA-4T+ PA-8T	<ul style="list-style-type: none"> ◆ 12.0S ◆ 12.3T 	<ul style="list-style-type: none"> ◆ 12.0(24)S and later ◆ 12.3(4)T and later 	<ul style="list-style-type: none"> ◆ 12.2SX 	<ul style="list-style-type: none"> ◆ 12.2(17)SXB and later

			later		
QoS on dMLP	Chan-PA PA-4T+ PA-8T	◆ 12.0S ◆ 12.2T ◆ 12.3 ◆ 12.3T	◆ 12.0(24)S and later ◆ 12.2(8)T and later	◆ 12.2SX	◆ 12.2(17)SXB and later
MPLS on dMLP MPLS on dLFIoLL	Chan-PA PA-4T+	◆ 12.2T ◆ 12.3	◆ 12.2(15)T11 and later ◆ 12.3(5a)	◆ 12.2SX	◆ 12.2(17)SXB and later
Distributed DDR	PA-8T-xT1 PA-MC-xE1 PA-MC-xTE1 PA-MCX-xTE1	◆ 12.3T	and later ◆ 12.3(7)T and later	-	-

¹ PA stands for Port Adapter. These PAs support the distributed features:

CT3IP, PA-MC-T3, PA-MC-2T3+, PA-MC-E3, PA-MC-2 E1, PA-MC-2 T1, PA-MC-4T1, PA-MC-8T1, PA-MC-8E1, PA-MC-8TE1+, PA-MC-STM-1.

² Cisco IOS Software Release 12.1E supports these features on both 7500 and 7600 platforms.

The information in this document was created from the devices in a specific lab environment. All of the devices used in this document started with a cleared (default) configuration. If your network is live, make sure that you understand the potential impact of any command.

Conventions

For more information on document conventions, refer to the Cisco Technical Tips Conventions.

Background Information

This section explains how you can identify whether fragmentation occurs in different distributed features.

If the issue relates to fragmentation counter only, the issue does not affect the functionality. However, if fragmentation does not work, latency issues occur for voice traffic or priority traffic.

How to Identify Whether Fragmentation Occurs

Complete these steps:

1. Collect the output of the **show tech** command from the Route Switch Processor (RSP) and linecard.
2. Collect the output of these commands from the RSP:

◆ **show ppp multilink**

- ◆ **show policy-map [policy name [class name]]**
3. Collect the output of these commands from the Versatile Interface Processor (VIP) console on which you have configured dLFioLL:
- ◆ **show ppp multilink**
 - ◆ **show vip hqf**
 - ◆ **show hqf interface**
 - ◆ **show controllers**
4. Enable the **debug dlfi** and **debug dmlp** commands on the VIP console, if the traffic is below 10 pps. If the traffic is heavy, disable console logs after you enable **debug** commands.

Issue the **show log** command on the linecard console to get these debugs.

Examine the packet debugs to check whether fragmentation occurs. You can infer that fragmentation occurs when the packets go out with these debugs:

```
00:05:53: Se6/3/0/2:0 MLP: I datagramstart FF03003D seq C0000006 size 110
count_fragment: last_frag=110 frag_count=2 total_bytes=122
Se6/3/0/1:0-0 dLFI: O seq 6 size 84
Se6/3/0/1:0 MLP: O datagramstart FF03003D seq C0000006 size 84
Se6/3/0/2:0-0 dLFI: O seq 7 size 38
Se6/3/0/2:0 MLP: O datagramstart FF03003D seq C0000006 size 38
```

Note: For PA–A6, dLFI counters do not work until you apply the fix for Cisco bug ID CSCin66700 (registered customers only) .

Troubleshoot Issues on dLFioATM and dLFioFR Interfaces

This section describes the actions you must perform when issues occur on dLFioATM or dLFioFR interfaces.

In order to troubleshoot issues on dLFioATM or dLFioFR interfaces:

- Collect the output of the **show tech** command from the Route Processor (RP) and linecard.
- Issue the **show interface** command on all interfaces, including the Virtual Access interface.

The creation of the Virtual Access interface depends on the virtual–template associated to the FR–DLCI or ATM–VC. In order to find the Virtual Access interface, issue the **show ip interface brief** command.

- Collect the output of these commands from the VIP console:

- ◆ **show log** if any error messages appear on the ATM or serial interfaces.
- ◆ **show controller**

Note: The dLFI interface does not come up if PPP over ATM or PPP over FR interface is down.

Troubleshoot Issues on dLFioLL Interfaces

This section lists the actions you must perform when you encounter problems on dLFioLL interfaces.

Latency Issues with dLFI

When you use dLFioLL, if you find high latency for voice or any critical traffic (for example, video), collect the debug information listed in this section. The debug information enables you to identify the root cause of the problem, and applies to 7500 and 7600 distributed platforms.

Latency in voice traffic deteriorates voice quality. The peer bundle sees several lost fragments and reordered packets, due to which packet-drops occur.

When you face latency issues with dLFI:

- Collect the output of the **show tech** command from the RSP and linecard.
- Collect the output of these commands from the RSP:

- ◆ **show ppp multilink**
- ◆ **show policy-map [policy name [class name]]**

The output of the **show policy-map** command indicates any drops that occur for the priority class traffic. If drops occur, the QoS priority queue gets filled, and delays the packets. Check whether you have configured sufficient bandwidth for priority traffic.

- Collect the output of these commands from the Versatile Interface Processor (VIP) console on which you have configured dLFIoLL:

- ◆ **show ppp multilink**

The output of the **show ppp multilink** command shows the fragment size that relates to each member link. The bundle contains the fragment size, which is minimum among the member links. Based on the fragment size and the member link bandwidth, you can identify whether the serialization delay of the link is too high, and causes latency for voice packets. If the issue is caused by these factors, you must tweak the fragment size. Configure an appropriate value in the **ppp multilink fragment-delay [value]** command such that the fragment size reduces.

- ◆ **show hqf interface**
- ◆ **show vip hqf**
- ◆ **show controller**

The output of the **show controller** command provides information on the tx-queue-limit set on the member links. The tx-queue-limit must be 2 if the link is part of the dLFI bundle. If the tx-queue-limit is not 2, traffic delays occur.

- ◆ **show vip tx-polling-high**

Cisco bug IDs CSCef17891 (registered customers only) , CSCin79544 (registered customers only) and CSCeg28064 (registered customers only) provide the relevant fixes for latency issues with dLFI.

dLFI Scalability Issues

When you configure several dLFIoLL interfaces, you do not get the throughput you require. This section describes the actions you must perform to troubleshoot dLFI scalability issues.

In order to troubleshoot dLFI scalability issues:

- Collect the output of the **show tech** command from the RSP and linecard.
- Collect the output of the **show ppp multilink** command from the RSP.
- Collect the output of these commands from the VIP or Flexwan console:

- ◆ **show proc cpu**
- ◆ **show buffer**

If the CPU utilization reaches 100%, the VIP or Flexwan CPU is unable to handle the traffic of so many dLFI interfaces. Hence, the CPU drops traffic.

Here is a sample **show buffer** command output:

```
show buffer:
:::
Header pools:
Egress Packet Header buffers, 0 bytes (total 2048, permanent 2048):
 0 in free list (2048 min, 2048 max allowed)
 2048 hits, 0 misses
 2048 max cache size, 2047 in cache
 1 hits in cache, 0 misses in cache
MLP Header buffers, 0 bytes (total 200, permanent 200):
 200 in free list (50 min, 200 max allowed)
 0 hits, 0 misses, 0 trims, 0 created

!--- Several misses mean that buffers are running out.

0 failures (0 no memory)
```

Lots of misses and an empty free/cached list imply that buffers are running out. This situation arises when packet reordering occurs. Packet reordering increases if there are more number of member links, and worsens if the link bandwidths are different.

Troubleshoot Issues on dMLP Interfaces

This section describes the issues on dMLP interfaces.

dMLP Performance Issues

When dMLP gives very less throughput on the 7500/VIP:

- Collect the output of the **show tech** command from the RP and linecard.
- Collect the output of these commands from the RP:

- ◆ **show ppp multilink**
- ◆ **show proc cpu**

If the interface is distributed, issue the **show proc cpu** command on the RP. The output enables you to check whether any interface is unable to switch traffic and therefore pushes MLP traffic to the RP, causing high CPU utilization. If you find no such interface, check [1] and [2] in this sample output:

```
Multilink3, bundle name is group3
Bundle up for 02:05:39
Bundle is Distributed
1 lost fragments, 0 reordered, 0 unassigned <----- [1]
14 discarded, 15 lost received, 1/255 load <----- [2]
0x0 received sequence, 0x0 sent sequence
Member links: 1 active, 0 inactive (max not set, min not set)
Se6/3/0/1:0, since 02:05:40, no frags rcvd, 5760 weight, 1496 frag size
```

Many **lost fragments** and **discarded** fragments in the output indicate that the interface reorders a lot of traffic on the multilink interface. The differential delay timer of 30 msec times out, and the interface is unable to buffer too many packets.

- Collect the output of these commands from the VIP on which you have configured the MLP bundle:

- ◆ **show ppp multilink**

The output of the **show ppp multilink** command enables you to identify whether the Multilink interface is distributed (see [1] in the sample output shown here):

```
Multilink3, bundle name is group3
Bundle up for 02:05:39
Bundle is Distributed <----- [1]
1 lost fragments, 0 reordered, 0 unassigned
14 discarded, 15 lost received, 1/255 load
0x0 received sequence, 0x0 sent sequence
Member links: 1 active, 0 inactive (max not set, min not set)
Se6/3/0/1:0, since 02:05:40, no frags rcvd, 5760 weight, 1496 frag size
```

- ◆ **show proc cpu**
- ◆ **show buffer**

The **show buffer** command enables you to check whether the interface has run out of Multilink buffer headers on the VIP.

```
show buffer:
:::::
Header pools:
Egress Packet Header buffers, 0 bytes (total 2048, permanent 2048):
0 in free list (2048 min, 2048 max allowed)
2048 hits, 0 misses
2048 max cache size, 2047 in cache
1 hits in cache, 0 misses in cache
MLP Header buffers, 0 bytes (total 200, permanent 200):
200 in free list (50 min, 200 max allowed)
0 hits, 0 misses, 0 trims, 0 created

!--- Several misses indicate that the interface is
!--- running out of buffers.

0 failures (0 no memory)
```

dMLP load balancing

Some load balancing algorithms that dMLP uses can affect performance if the peer side must work on a specific algorithm. dMLP load balances the outgoing packets based on the bandwidth of each member link. For example, consider three member links of 500Kbps, 1Mbps and 1.5Mbps. You assign a weight of 500 bytes, 1000 bytes and 1500 bytes to these links. The first link can send 500 bytes out. If the last packet is greater than 500 bytes, the packet is still sent out from the same link. Then you move on to next one, and the process continues.

Incorrect Rate Counters for dMLP Bundle Interfaces

Sometimes, the rate counters for dMLP are not correct. This section lists the actions you must perform when this issue occurs.

In order to troubleshoot incorrect rate counters for dMLP:

- Collect the **show tech** command from the RP and linecard.
- Collect the output of the **show ppp multilink** command from the RP and from the VIP on which you have configured the bundle.
- Collect the output of these commands from the RP:

- ◆ **show interface multilink [N]**
- ◆ **show interface multilink [N] stats**

Collect the output of these commands from the RP again after 60 seconds.

- Collect the output of the **show interface** and **show interface stats** commands for member links.

The **show interface** command output enables you to compare input and output counters to check whether the rates match or not. Rates could be high if the packet-drops occur at member links, when the traffic pumped through MLP interface exceeds the bandwidth.

Calculate the rates based on the information in the **show interface** command for the multilink or mfr bundle, which you take at a 60-second interval. The formula to calculate rates is $(\text{count1} - \text{count2})/60$.

Similarly, you can calculate the rates shown in the **show interface** command output. In case there is a discrepancy between the calculated throughput of dMLP (via external tgen / set up with ixia..etc), and the throughput that appears on the interface, check whether the Multilink accounts for the drops correctly. Ensure that the packet-drops occur at the member links alone. In the latter case, the rate on the interface is higher than the actual throughput of multilink.

In some error conditions, incorrect or negative output and input counters directly affect the rates. If no updates occur on the output and input counters, the problem lies with the linecard and RP communication. Enable **debug multilink** on the linecard.

Note: Disable console logs before you enable the **debug multilink** command.

Check whether the linecard sends periodic updates to the RP. Cisco recommends you to use a release that integrates bug ID CSCeb31029 (registered customers only), because this bug eliminates output late drops on multilink interface.

Troubleshoot Issues on dMLFR Interfaces

This section explains the actions you must perform when back-to-back **ping** does not work on dMLFR interfaces.

When back-to-back **ping** does not work on dMLFR interfaces:

- Collect the output of these commands from the RP:
 - ◆ **show tech** or **show run**
 - ◆ **show frame-relay multilink**
 - ◆ **show frame-relay map**
 - ◆ **show interface MFR[N]**
 - ◆ **show interface** for member links of MFR bundle
- Collect the output of the **show ppp multilink** command from the VIP console on which MLFR member links lie.
- Check whether any configuration change causes this issue.
- Make sure that MFR interface is Up/Up. If the MFR interface is Up/Up, and **ping** does not work, check the output of the **show frame-relay map** command. The output shows the map:

```
RSP16-2#show frame-relay map | beg MFR
MFR4 (Up): ip 30.0.0.1 dlci 20(0x14,0x440), dynamic,
          broadcast,, status defined, active
```

Build frame-relay map.

The **ping** feature does not work unless you build frame-relay map. If ping still does not work (after you populate the map) you must check where the packet-drop occurs. Check whether the physical interface drops the packets or the CEF drops the packets. Also check whether the linecard reports

receipt or transmission of the ping packet instead of the RP, and vice versa. This information enables you to identify where the packet-drops occur.

- If the MFR interface is Up/Down:
 - ◆ Check whether you have configured the MFR interface as DTE on one side and as DCE on another side.
 - ◆ Check for keepalive failures. Disable keepalive to check whether the MFR interface goes down after three keepalive failures. For example, if the keepalive value is 10 seconds, the MFR interface goes down after 30 seconds.

Troubleshoot Other Error Messages

This section explains how you can troubleshoot the FIXBADTXVC and RSPACCERR messages.

FIXBADTXVC Messages

Sometimes, on a Cisco 7500 router, this error message appears:

```
CBUS-4-FIXBADTXVC: Detected and fixed bad tx vc encap on , bad vc , fixed vc.
```

This message indicates that an interface sends out a packet without the correct VC information in the packet header. The first two bytes of the packet header contain the VC information, which indicates the index of the egress interface. If VC information is absent, the VIP code cannot determine the destination interface of the packet. Therefore, the VIP drops the packet and causes loss of tx-acc for the interface. This error message indicates the detection and correction of the problem. Therefore, this message does not cause any negative impact.

However, you must perform further code analysis to determine whether a process did not add the *board encap*, or whether the incorrect VC information was added. Sometimes, this error occurs when the interface adds the VC information at the wrong location in the packet.

Collect the output of these commands when this error occurs:

- **show running config**
- **show log**

Decode the traceback to identify the process that generates packets without adding the *board encap*. Based on the process, a code analysis reveals the path that calls *datagram_out*, or enqueues the packet on the hold queue without adding the *board encap*.

Also, you can use the config and traceback to determine whether the interface adds wrong VC information, or adds the VC information at the wrong location in the packet. The VC information forms the first two bytes of the packet (except when you add QoS classification information, where the QoS classification information forms the first 8 bytes, after which the VC information appears).

These bugs fix similar issues:

- CSCdu40952 (registered customers only)
- CSCdy46586 (registered customers only)
- CSCdu68154 (registered customers only)
- CSCds58693 (registered customers only)

RSPACCERR Messages

Cisco allocates an accumulator for each physical interface on a Cisco 7500 router. You can compute the limit of the accumulator during an MEMD carve. Alternatively, a dynamic adjustment occurs through the CLI when you issue the **tx-queue-limit <new limit>** command, or due to multilink configuration.

The RSPACCERR message appears when the value of the tx-acc is greater than the limit, and a reset occurs. This error indicates some sort of accumulator credit mismanagement. A wrong interface probably receives wrong credit for the acc of another interface.

Collect the output of these commands when the RSPACCERR error occurs:

- **show controller cbus**

The output of the **show controller cbus** command enables you to determine whether any interface has lost tx-acc credits. Hence, you can infer that the wrong interface received the credits.

- **show log**

The logs enable you to co-relate the network events that trigger this issue.

- **show run**

In addition, try to find out whether any change in router configuration triggered this error originally.

Note: After the first occurrence of the RSPACCERR error, every interface reset triggers this error message.

This issue has no real fixes.

Sometimes, faulty VIPs corrupt the packet contents. In channelized interfaces, the first two bytes of the packet denote the index of the egress interface on a plughole. If the information is incorrect, the interface loses credit, or a wrong interface receives credit. If a wrong interface receives credits, the tx-acc value of that interface becomes greater than the permissible limit. Occasionally, this issue occurs when the primary link of the dMLP/dMFR bundle changes, wherein the information on the tx-acc and the VC index are inconsistent for a brief period on different VIPs in the system.

Here are the related Cisco bugs:

- CSCeb57127 (registered customers only)
- CSCdv46696 (registered customers only)

MTU-Related Information

In Cisco IOS Software Releases earlier than 12.3(7)T and 12.0(28)S, the advertised Maximum Receive Reconstructed Unit (MRRU) value is 1524 by default. Therefore the Peer Bundle restricts the MTU to 1524. A router that runs IOS MLP drops any packet of size greater than 1524.

However, a 7500 router that runs dMLP has this caveat:

- The router does not drop non-fragmented MLP Packets greater than 1524. The router drops fragmented packets. So, when the MTU size is more than 1524 in the network, an IOS MLP router can drop packets greater than 1524, and the 7500 router can drop fragmented MLP packets.

In Cisco IOS Software 12.3(7)T, 12.0(28)S, and later, Cisco IOS Software Release 12.0(28)S3 includes an MRRU negotiation feature. After the introduction of this feature, Cisco fixed bug ID CSCin73658 (registered

customers only) to address the caveat.

Additional Information You Require on 76xx Series Routers

You must collect some additional information on 76xx series routers in case the traffic does not pass through the Multilink /MFR interfaces, or when these interfaces fail to come up. This issue can occur because the internal VLAN allocation for these interfaces is incorrect. This issue can also occur if packets to and from the linecard do not reach the correct VLAN.

Collect this information when such problems occur on the 76xx series routers:

From the RP:

```
7600-1#sh cwan vlan | inc MFR
 1232 231 MFR1
 1233 233 MFR2
 1234 239 MFR5
7600-1#sh cwan vlan | inc Multilink
 1235 242 Multilink1
 1236 246 Multilink2

From LC on slot 6
FlexWan-ENH6/0#sh cwan vlan | inc MFR
 1232 (4D0 ) 000a.000b.000c MFR1
 1233 (4D1 ) 000a.000b.000c MFR2
 1234 (4D2 ) 000a.000b.000c MFR5
FlexWan-ENH6/0#sh cwan vlan | inc Multilink
 1235 242 Multilink1
 1236 246 Multilink2
```

The VLAN IDs (in the first column) must match for all the Multilink/MFR interfaces. If VLAN IDs do not exist, or are incorrect, traffic issues can occur.

If the 7600 router operates in truncated or compact mode (fabric switching) you must also verify the FPOE program for these interfaces. For example, consider MFR1 from the example. The VLAN ID is 1232 (0x4d0 in hex). Check the flood entry for this interface in the supervisor. Use these commands:

```
7600-1# show fabric
show fabric active:
Active fabric card in slot 5
Backup fabric card in slot 6
show fabric mode:
Fabric module is not required for system to operate
Modules are allowed to operate in bus mode
Truncated mode allowed, due to presence of N/A
Module Slot      Switching Mode
 3                Crossbar
 5                Crossbar
 6                Bus
```

This table illustrates how to decide the operating mode of 7600:

Slot(s) Switching mode	7600 operating mode
All Slots in Crossbar Mode	Compact Mode
All Slots in Bus Mode	Bus Mode
	Truncated Mode

Some slots in Crossbar Mode and Some in Bus Mode	
--	--

In order to verify that the FPOE entries are correct, issue these commands:

```
FlexWan-ENH2/0#sh cwan vlan | i Multilink51
1444 (5A4 ) 000d.662d.e880 Multilink51
1444 (5A4 ) Multilink51 Multilink51
```

In this example, you have configured Multilink 51 in slot 2.

On the RP, issue this command:

```
A#show fabric fpoe interface mul 51
fpoe for Multilink51 is 9
```

You also need to verify if the same FPOE program is present for interface Multilink 51 on the SP also. Issue this command on the SP:

```
A-sp#show platform software table fpoe slot 8 start 0x85A4 end 0x85A4
Index   Block0   Block1   Block2   Block3   Block4   Block5   Block6   Block7
-----+-----+-----+-----+-----+-----+-----+-----+-----
0x85A4  0x3FDFF  0x3FDFF  0x3FDFF  0x3FDFF  0x3FDFF  0x3FDFF  0x3FDFF  0x3FDFF
```

Note: Obtain the index from ORing 0x8000 with the hex value of vlan_id.

This entry gives a value of 0x3FDFF. Start from the right–most bit, which is bit 1. This value has a '0' in bit 10, which is the expected value for the interface Multilink 51. Assume that the output of the RP command **show fabric fpoe interface** is N (in this case, 9). You must set the bit in the (N+1)th position (in this case bit 10) to 0 in the output of **show platform software table fpoe slot 8 start 0x85A4 end 0x85A4** from the SP. If this entry is incorrect, packet–switching in truncated mode is also incorrect.

Here is a list of possible error messages that can appear when you configure Multilink bundles on SIP1 linecard and switch over from hardware implementation to software implementation:

- **Member link spread across bays:** This error message indicates that member links are from different bays. Freedm is unable to perform multilink functionality, and the bundle switches to Software Mode.
- **link_count exceeds max links in the bundle:** This error message indicates that no links in the bundle became greater than 12. Freedm is unable to perform multilink functionality, and the bundle switches to Software Mode.
- **Interleaving enabled with multiple links:** This error occurs when you have enabled interleaving, and there is more than one link. Freedm is unable to perform multilink functionality, and the bundle switches to Software Mode.
- **Member links are not full T1/E1:** This error occurs when you add fractional T1/E1 to the bundle, and the bundle switches to Software Mode. Freedm is unable to perform multilink functionality. When the link fails to be added to the bundle, these error messages appear:
 - ◆ **Links spread across bays or dCEF disabled:** This error occurs when you try to add links from a different linecard.
 - ◆ **Link incompatible with existing links in the bundle:** This error occurs when you add links of a different type to the bundle. For example, you add one link from CTE1 and another one from CT3.

Case Studies for dDDR

Distributed Dial on Demand Routing (dDDR) allows the use of DCEF on dialer interfaces. dDDR supports these Dialer interface types:

- Legacy Dialer interfaces
- Dialer Profile interfaces
- Dialer Profile interfaces with MLPPP
- Legacy Dialer interfaces with MLPPP

Note: dDDR does not support QoS currently. dDDR currently supports only the channelized PAs, namely, PA-MC-xT1, PA-MC-xE1 (x = 4,8) and PA-MCX-TE1+.

dCEF Does Not Switch Packets

Use the **show interface xxx stat** command to check whether CEF switches packets or dCEF switches packets.

Troubleshoot

Complete these steps:

1. Issue the **show adjacency detail** command on the VIP and RP to check for Adjacency. If Adjacency does not exist, proceed to step 2.

If you find Adjacency, issue the **show cef interface** on the VIP to identify the populated vectors on the incoming interfaces. Then, collect the output of the **show cef interface** command on the RP. Collect **show log** command outputs on both RP and VIP, to identify the state changes of the links. Also check with the customer if there were any link flaps before the occurrence of the problem. If you use a Dialer profile configuration, collect the output of the **show ppp multilink** command on the RP and VIP. Check for any failed ipc messages in the output. If you find failed ipc messages, perform Online Insertion and Removal (OIR) on the VIP. Otherwise, proceed to step 4.

2. If Adjacency is absent on the LC, proceed to step 3. Absence of Adjacency at the RP indicates a problem with the CEF code. Issue the **show dialer** command to check whether the dialer protocol is up.

Issue the **show isdn status** command in order to verify dDDR:

```
BOX2002#show isdn status
Global ISDN Switchtype = primary-net5
ISDN Serial3/1/0:23 interface
***** Network side configuration *****
dsl 0, interface ISDN Switchtype = primary-net5
Layer 1 Status:
ACTIVE
Layer 2 Status:
TEI = 0, Ces = 1, SAPI = 0, State = MULTIPLE_FRAME_ESTABLISHED
Layer 3 Status:
0 Active Layer 3 Call(s)
Active dsl 0 CCBS = 0
The Free Channel Mask: 0x807FFFFF
Number of L2 Discards = 0, L2 Session ID = 6
```

Ensure that the ISDN status is MULTIPLE_FRAME_ESTABLISHED. This status indicates that the physical layer is ready for ISDN connectivity.

Next, issue the **show dialer** command:

```
EDGE#show dialer
  Serial6/0:0 - dialer type = ISDN
  Idle timer (120 secs), Fast idle timer (20 secs)
  Wait for carrier (30 secs), Re-enable (15 secs)
  Dialer state is data link layer up
  Time until disconnect 119 secs
  Current call connected never
  Connected to 54321
    Serial6/0:1 - dialer ISDN =
  Idle timer (120 secs), Fast idle timer (20 secs)
  Wait for carrier (30 secs), Re-enable (15 secs)
  Dialer state is idle
```

From the output, you can infer that the type of dialer is ISDN. ISDN implies legacy dialer configuration and PROFILE implies dialer profile configuration. The output also displays the Dialer State, which is the current state of the dialer. The state of an unconnected dialer interface is idle. The Idle timer resets whenever interesting traffic arrives. If this timer ever expires, the interface immediately disconnects. You can configure the Idle timer parameter.

If the dialer protocol is down, check your configuration to identify the reason. If you do not see any problem with the configuration, file a bug with the dialer team. Provide the output of the **show running config** and **show log** commands.

If dialer protocol is up, collect the output of the **show cef interface** command. This output indicates why CEF Adjacency is absent for the dialer interface. Provide the output of the **show running config** and **show log** commands from the RP. Also, collect **show interface** command output.

3. If Adjacency is absent on the LC, but present on the RP, there can be some problem with the distributed feature. Collect the output of the **show cef interface** command on the RP and VIP. On the RP, you can examine this output to find out why the dialer is not distributed. For example PPP not up . Then check the configuration.

If you use Dialer profile configuration, collect the output of the **show ppp multilink** command on the LC, and check for failed ipc configuration messages.

If you see any such messages, perform OIR of the linecard. Otherwise, proceed to step 4. Ensure that you collect the outputs of the **show log** and **show tech** commands from the RP and VIP.

4. If you have reached this step, the Adjacencies are fine. A possible reason for traffic to be CEF-switched, even though the Adjacencies are fine, is loss of tx-acc. To troubleshoot this issue, collect the output of the **show controller cbus** command. Here is a sample output:

```
Serial0/0/0/3:0, txq E8001A00, txacc E8000082 (value 5), txlimit 5
```

Issue the **show controller cbus** command repeatedly, and verify the output to see whether the 'value' parameter appears to be stuck at a low value. If so, loss of tx-acc is the issue. Complete these steps:

- a. Collect the output of the **show controller serial** command on the VIP five times.
- b. Collect the outputs of the **show tech** and **show tech cef** commands on the LC and RP.

Note: Remember these points:

- ◆ dCEF switching does not occur on the interface if:

- ◇ The tx-acc value falls below 4.
- ◇ The acc value falls below 2.

- ◆ If the acc value falls to 0, the interface no longer forwards traffic. The interface does not even

forward control packets, and leads to the 'Not transmitting' and 'output frozen' status.

- ◆ If you observe that the tx-acc value of some interfaces is more than the limit while others have a very small value, and if you have configured dMLP or dMLFR on these links, the problem can be due to these bugs:

- ◇ CSCdx67602 (registered customers only)
- ◇ CSCin47712 (registered customers only)
- ◇ CSCed29590 (registered customers only)
- ◇ CSCin62978 (registered customers only)

On the 7500 platform, examine the output of the **show interface** command. If you find that the number of output buffer failures is on an increase, you can infer that the system is running low on MEMD buffers. This situation results in a false indication of 'Not transmitting'.

Analyze the output of the **show controllers** command on the linecard of a distributed system or on the router of a non-distributed system. Check for the 'tx-count' information in the output for the relevant interface.

If the tx-count remains a non-zero entity continuously even when there is no traffic, the driver hardware and software are probably out of synch. Follow the specific driver model to troubleshoot further. Some of the known fixes in this area include:

- CSCec07487 (registered customers only)
- CSCdw00005 (registered customers only)
- CSCdw18116 (registered customers only)
- CSCdt48893 (registered customers only)
- CSCin62978 (registered customers only)

In case of channelized drivers on a 7500 platform, check whether the output of the **show controllers** command from the VIP indicates any 'bad vc' errors. Bad VC errors indicate that the VIP received packets for which the VIP was unable to determine the outgoing interface. This issue can occur either because the packets were received even before you configured the interface in the VIP, or because the packets were corrupt, and hence the first two bytes of the packets that were to indicate the outgoing interface were incorrect.

Sample Output

```
Router#show controller interface Serial1/4:0
```

```
!--- output Suppressed
```

```
.  
. .  
rom rev 0x1012B f/w rev 1.2.4, h/w rev 131, PMC freedm rev 1 idb = 0x62042920  
ds = 0x60D29EC0, plx_devbase=0x54804000, pmc_devbase=0x54800000  
Enabled=TRUE, DSX1 linestate=0x4,  
fastsend = 0x6044F090, post_compress = 0x60299498, lovenote = 0x0  
RSR: 0x5000, PLPRCR: 0x5000  
pak_to_host=0x6044F5D0, vip_memd_ifcntl=0x60A1CCC0,  
tx acc=0xE80000CA Tx bad vc=5 Null Tx acc inc=0 stuck reset 0
```

Here are the bugs that relate to problems with OIR:

- CSCec07487 (registered customers only)
- CSCdw00005 (registered customers only)
- CSCdw18116 (registered customers only)

- CSCdt48893 (registered customers only)

However, an occurrence of this issue in normal operation without OIR events indicates hardware problems. The **shut** and **no shut** commands on the dialer interface solve the problem.

Troubleshoot Router/VIP Crashes and Cbus Complex

When the router or VIP crashes, ensure that you collect the crashinfo file. Without the crashinfo file, Cisco can provide very little analysis. For information on Cbus Complex refer to What Causes %RSP-3-RESTART: interface [xxx], output stuck/frozen/not transmitting Messages?. Here are a few reasons for router crashes and Cbus Complex:

- **Configuration Changes:** Sometimes a change of configuration on the near end or far end can cause a router crash. This problem can occur when you turn dCEF on or off on the router. If MLPPP is unconfigured on the fly on near or far end routers.
- **Watchdog timeout:** If a process refuses to relinquish control of the CPU for a long time, the box crashes. This issue occurs when lines flap a number of times. To verify whether line-flapping is the cause of the crash, collect the output of the **show tech** and **show log** commands.
- **Output stuck:** When an interface no longer transmits on a 7500 system, each of the VIPs reload. Interfaces stop transmission when they lose tx-acc. If the problem occurs frequently, issue the **debug cbus output-stuck** command. Refer to What Causes %RSP-3-RESTART: interface [xxx], output stuck/frozen/not transmitting Messages? for more information.

Please gather the complete output of the **show tech** command from the RSP. Also check whether you made any configuration changes that triggered the problem.

Note: The router usually comes back online after the crash.

NetPro Discussion Forums – Featured Conversations

Networking Professionals Connection is a forum for networking professionals to share questions, suggestions, and information about networking solutions, products, and technologies. The featured links are some of the most recent conversations available in this technology.

NetPro Discussion Forums – Featured Conversations for Router and IOS Architecture
Network Infrastructure: LAN Routing and Switching
Network Infrastructure: WAN Routing and Switching

Related Information

- **Distributed Multilink Point-to-Point Protocol for Cisco 7500 Series Routers**
- **Distributed Multilink Frame Relay**
- **Distributed Link Fragmentation and Interleaving over Leased Lines**
- **VoIP over PPP Links with Quality of Service – LFI**
- **What Causes %RSP-3-RESTART: interface [xxx], output stuck/frozen/not transmitting Messages**
- **What Causes a "%RSP-3-RESTART: cbus complex"?**
- **Troubleshooting TechNotes – Cisco 7500 Series Router**
- **Router Product Support Page**
- **Technical Support & Documentation – Cisco Systems**

[Contacts & Feedback](#) | [Help](#) | [Site Map](#)

© 2007 – 2008 Cisco Systems, Inc. All rights reserved. [Terms & Conditions](#) | [Privacy Statement](#) | [Cookie Policy](#) | [Trademarks of Cisco Systems, Inc.](#)

Updated: Jun 24, 2008

Document ID: 65334
